THE WHITE PAPER ON THE EU AI ACT **FOR HR TEAMS**

AI ACT FOR TALENT ACQUISITION

COMPETENT, COMPLIANT AND MANAGEABLE

Queb
BUNDESVERBAND
»TALENT ACQUISITION«

0 Foreword 07

0.0 Entry self-assessment: Where are you today? (2 minutes) 07
0.1 Foreword: What is this playbook For? 08
0.2 Purpose, target audience, and scope (Talent Acquisition focus) 08
0.3 Methodology and reading guide (modular use) 09
0.4 Scope and interaction: AI Act/GDPR/AGG/ Works Constitution Act (BetrVG): brief overview 10
0.5 Disclaimer, versioning and change management 11
0.6 Structure of this playbook 11

1 Executive Summary: KI and Talent Acquisition in 7 Minutes 13

1.1 Purpose and benefits of this playbook 13
1.2 The five core messages (guiding principles of this playbook) 13
1.3 Traffic light logic: green/yellow/red (guidance for TA use cases) 15
1.4 Make a decision in 3 minutes: stop/go check in five questions 17
1.5 Top 10 quick wins for Talent Acquisition (Ready to implement immediately) 20
1.6 How to use this white paper in practice (reading path and entry points) 21

2 Target Audience, Scope and Work Paths (How to Use This White Paper) 23

2.1 Target audience (who this was written for) 23
2.2 Scope (what Is included and what Is deliberately excluded) 23
2.3 Work paths (how to quickly reach a result) 24
2.4 Role of the Toolkit (separate document) 24

3 Regulatory Framework (AI Act) in Talent Acquisition: What You Really Need to Know 26

3.1 Purpose of this chapter 26
3.2 Core logic of the AI Act: risk-based rather than one-size-fits-all 26
3.3 Practical application for Talent Acquisition: three risk classes 27
3.4 Related legal frameworks: the AI Act does not stand alone 27

Release date: 1 May 2026

The accessible version of this white paper is limited by current accessibility standards, particularly with regard to language quality. Queb e.V. has prepared this edition to the best of its knowledge and judgment and remains fully committed to supporting accessibility in the future.

This publication contains images generated using AI.

No legal advice / No guarantee of completeness or currency

Important notice: Clicking a T-Link (Toolbox) immediately starts the download of the Word and Excel templates.

4 Chatbots, GenAI and Agents: Terms, Scope and TA Use Cases 29

4.1 Why this chapter matters 29
4.2 What Is an AI system and when does it typically become critical in talent acquisition? 29
4.3 The key distinction: information vs. decision support 30
4.4 TA Use cases: five illustrative use cases 32
4.5 Use case matrix: use case → role in the process → risk → compliance requirements 32
4.6 Output catalog: what counts as information and what counts as decision support? 34
4.7 Practice box: Common misunderstandings and How to avoid them 35
4.8 Do/don't: terminology and classification in TA practice 36

5. Risk Map for Talent Acquisition Use Cases 37

5.1 Why a TA risk map is necessary 37
5.2 Traffic light model (green/yellow/red): Classification by impact, data, and control 38
5.3 High-risk triggers in TA: when yellow turns red 38
5.4 Borderline cases: typical TA scenarios and decision aids 40
5.5 Determining the next steps: Recommended actions for each traffic light category 43

6. Roles, Responsibilities and Contract Frameworks in TA Projects 44

6.1 Why this chapter matters 44
6.2 Role model in TA terms (simplified) 44
6.3 Required documents (minimum set): what documentation and working documents you need 46
6.4 "Who delivers what?" – responsibilities of the provider and deployer 46
6.5 Contract and RFP framework: minimum documentation as a quality filter 47
6.6 What to do if the provider does not provide complete documentation 49
6.7 Practical workflow: from idea to approval (board-ready) 50
6.8 References to templates (see Chapter 14.3) and subsequent chapters 52

7 General-Purpose Models (GPAI/LLMs) in TA and Potential Shifts in Roles and Obligations 53

7.1 Why this chapter matters 53
7.2 Core principle: the use case makes the difference 54
7.3 Typical TA scenarios with elevated compliance risk 55
7.4 Practical countermeasures: how to keep Use controlled 55
7.5 'If You Do X, You Need Y' (a quick reference for TA) 58
7.6 References to templates and subsequent chapters 60

8 Transparency, Explainability, and Data Subject Rights 61

8.1 Why this chapter matters 61
8.2 Transparency in TA processes: what should be communicated 61
8.3 Explainability: factors instead of scores 62
8.4 Data subject rights and review procedures (practical and effective) 62
8.5 Practical minimum standard (TA) 65

9 Fairness, Bias, and Proxy Discrimination in Talent Acquisition 66

9.1 Fairness determines acceptance and risk 66
9.2 Bias in TA: where risks typically arise 67
9.3 Proxy discrimination explained clearly 67
9.4 Minimum standard: what to check for red use cases 68
9.5 Practical implementation: how to keep it manageable 69
9.6 Reference to the Toolkit 70

10 Monitoring, Updates and Incident Handling (Operational Reliability in TA) 71

10.1 Monitoring is a deployer obligation, not the end of the project 71
10.2 What should typically be monitored in TA 71
10.3 Updates and change impact: when reassessment is required 72
10.4 Incident handling: stop/override/escalation 73
10.5 Minimum operating standard for AI in TA without overkill 77
10.6 Reference to the Toolkit 77

11 Data Protection, Data Classes and LLM Input Rules in Talent Acquisition 78

11.1 Data protection is always relevant – regardless of the AI Act 78
11.2 Data classes in the TA context (practical working logic) 78
11.3 Prompt and input hygiene (GenAI/LLM): the most important practical rule 79
11.4 Article 22 of the GDPR (automated decision-making): a practical view 80
11.5 DPIA / Data Protection Impact Assessment (when it typically becomes relevant) 80
11.6 Reference to the Toolkit 82

12 Co-Determination (BetrVG), Change Management and Works-Council-Ready Implementation in Talent Acquisition 83

12.1 Co-determination Is not an add-on – it determines feasibility 83
12.2 When the Works Council typically must be involved 83
12.3 Works-council-ready package: which documents TA should provide 84
12.4 Change management: internal safeguarding 85
12.5 Practical guiding principle: control limits instead of endless discussion 85
12.6 Reference to the Toolkit 86

13 HR/TA Governance and Competence Building (Article 4 of the AI Act) as an Operating Model 87

13.1 Governance enables TA to act – it does not slow it down 87
13.2 Operating model: roles, interfaces and responsibilities 87
13.3 RACI (brief): who does what? 88
13.4 Competence building under Article 4 of the AI Act (practical) 88
13.5 Governance in operations: standards, reviews and escalations 89
13.6 Special case: governance without an AI Board or legal department 90
13.7 Market structure: provider categories and their typical risk profiles 91

14 Toolkit (Templates and Policies): How to Accelerate Implementation 93

14.1 Purpose of this Toolkit 93
14.2 How to use the Toolkit (in 3 steps) 94
14.3 Toolkit overview (contents) 94
14.4 Versioning and maintenance principle 95

15 Case Examples from Member Companies 97

15.1 Case nugget framework 97
15.2 Available and planned cases 98
15.3 Shared learning 98
15.4 Case library as living format 98

16 Conclusion, Outlook and Next Steps for HR/Talent Acquisition 99

16.1 Conclusion: AI in TA is feasible if it Is managed as a process 99
16.2 Outlook: three trends that will shape TA in 2026/2027 101
16.3 Next steps: the 30-day plan for TA (Minimum program without a major project) 101
16.4 Using this white paper as a communication tool (stakeholder-ready) 102

17 Glossary 103

18 About the Authors 104

19 Queb Federal Association 108

Important Notice

*This white paper **is solely intended to provide general information and practical guidance** in the context of the European AI Act (AI Act / EU AI Act) in the field of Human Resources and Talent Acquisition. It does not constitute legal advice and does not replace an individual legal review by qualified legal professionals or compliance officers within the relevant organisation. The contents were prepared with the greatest possible care. Due to dynamic regulatory developments, possible national specificities and ongoing adjustments to European AI regulation, including the AI Omnibus, **the authors and Queb e.V. do not guarantee the completeness, currency or legal binding effect of the content.** Companies and organisations should also have all measures, assessments and implementations individually reviewed from a legal perspective.*

Entry self-assessment: Where are you today? (2 minutes)

Answer these three questions. You decide **where to start in this playbook:**

Question 1: Does your company already have **specific AI use cases in Talent Acquisition – either in use or planned?**

Yes: Start with *Chapter 1.3* (Traffic light logic) + *Chapter 5* (Risk Map).
No: Start with *Chapter 1.1* and the *30-day plan* (Chapter 16.3).

Question 2: Do any of your use cases support selection decisions (matching, ranking, shortlisting)?

Yes: This is your most urgent action area: *Chapters 5–10* + T-14.3 and T-14.4 are your priority.
No: *Chapters 11–13* for data protection and governance fundamentals.

Question 3: Do you need to convince internal stakeholders (management, co-determination, data protection)?

Yes: *Start with Chapter 1* (Executive Summary) + *Chapter 6* (Roles/ Documentation) + *Chapter 12* (Co-Determination). Then use the Decision pack T-14.4 as a communication tool.
No: Find a new job. ;-)

If you answered 'yes' to all three questions: Start with T-14.2 (one-pager) + T-14.3 (risk intake) – these are the two most important entry points.

0.1 What is this playbook for?

Artificial intelligence is in particular transforming HR and Talent Acquisition quickly. Not because AI is 'trendy', but because companies must work more efficiently, deliver higher quality and be more service-oriented in an environment shaped by skills shortages, cost pressure and increasing speed. This is exactly where AI can help: it can speed up processes, stabilise quality, make decisions more transparent and give recruiting teams back time for what matters most.

At the same time, the European AI Act (AI Act / EU AI Act) creates a binding legal framework that sets clear rules – especially where AI provides decision support or influences selection decisions. For companies operating in Europe, AI in Talent Acquisition can be a real success factor.

However, organisations that want to use AI must involve data protection, co-determination, IT security and governance early to avoid ending up in a 'stop-and-fix' situation later on.

This white paper was written by HR for HR. It is not a legal commentary, but a **playbook** that translates the EU legal framework for AI in HR into a practical working logic: it clarifies terminology, classifies use cases using traffic light logic, establishes minimum standards, defines controls and safeguards ongoing AI use. This prevents AI from becoming a risk factor and instead turns it into a manageable enabler of modern HR work – robust enough for audits, critical discussions with Works Councils and challenging questions from within the organisation.

0.2 Purpose, target audience and scope (Talent Acquisition focus)

Purpose: This playbook helps HR and Talent Acquisition teams systematically assess AI use cases in recruiting, prepare them for internal decision-making and manage them in a controlled way during operation. The focus is not on 'knowing the rules', but on implementation: what questions need to be clarified, what documentation is required and what controls work in practice?

Target audience: This document is primarily intended for Talent Acquisition professionals, recruiting leads, HR managers and hiring managers, as well as the typical interfaces within organisations: data protection (DPO/data privacy), IT security/ architecture, co-determination/Works Council and governance bodies (e.g., steering committees / AI boards).

Scope: The focus is on typical TA use cases along the candidate journey (e.g., AI-supported sourcing, candidate FAQ bots, GenAI job ads, screening/matching/ ranking and interview/assessment support) as well as operational topics: roles and responsibilities, transparency/explainability, fairness/bias, monitoring/updates/ incidents and practical working templates (Toolkit).

0.3 Methodology and reading guide (modular use)

This white paper is modular in structure. It does not need to be read from beginning to end. The goal is to help you reach results quickly, depending on your role and question, without losing the consistent framework.

The playbook is designed as follows:

- **Traffic light logic (green/yellow/red)**: quick classification based on impact within the process (information vs. decision support)

- **Stop/go check:** brief decision aid for borderline cases

- **Documentation rather than theory:** minimum documentation is recommended for each use case to speed up decisions and make them audit-ready

- **Boxes and graphics:** summarise the key messages

- **Toolkit references:** (Templates) templates are kept separate so that the white paper remains readable. *Important notice: Clicking a T-Link (Toolbox) immediately starts the download of the Word and Excel templates.*

Recommended entry points:

- **'I have a specific use case':**
  Traffic light logic / stop-go → risk intake → decision pack → controls by risk level

- **'I am building governance':**
  Roles/RACI → minimum standards → monitoring/incident path → competence building

- **'I need to convince stakeholders':**
  Executive summary + traffic light logic + 30-day plan

0.4 Scope and interaction: AI Act/GDPR/AGG/Works Constitution Act (BetrVG): Brief overview

The AI Act provides a central framework, but it never stands alone in Talent Acquisition. In practice, implementation is multi-track:

- **AI Act:** classifies AI applications based on risk and defines obligations depending on the context of use, especially in decision support.

- **GDPR:** covers personal data, purpose limitation, data protection impact assessment (DPIA) obligations and the limits of fully automated decision-making under Article 22 of the GDPR.

- **AGG:** addresses discrimination risks and is particularly relevant in Talent Acquisition when it comes to criteria logic, proxy effects and unfair outcomes.

- **Works Constitution Act (BetrVG) / co-determination:** becomes relevant when processes, rules, data flows or potential employee-surveillance risks are involved.

Key point: The AI Act governs the 'how' of accountability; the GDPR, the General Equal Treatment Act (AGG) and the Works Constitution Act (BetrVG) determine admissibility and consultation obligations (see *Chapter 3.4).*

0.5 Disclaimer, versioning and change management

Disclaimer: This document is a practice-oriented playbook for HR / Talent Acquisition. It is intended to help structure and prepare internal decisions and does not replace legal advice in individual cases. In critical or contentious cases, the Legal, Data Protection and Employee Representation departments should be involved.

Versioning and maintenance: This white paper is maintained as a living document. The current status and any changes are documented with version numbers; new findings, for example from practical examples, audits, releases or regulatory clarifications, are incorporated into updates. The current version and the associated Toolkit are available from ***Queb e. V. Bundesverband für Employer Branding Personalmarketing Recruiting e. V.***

0.6 Structure of this playbook

Foundation – understand and contextualise *(Chapters 1–6):*

What is the AI Act?
Which systems are affected?
How does traffic light logic work?

Controls – responsibility and documentation *(Chapters 7–13):*

Who does what?
Which obligations apply?
How do I keep proper documentation?

Scale – rollout and operation *(Chapters 14–17):*

How do I introduce AI systems, monitor them and ensure ongoing compliance over time?

AI Act timeline: What applies when?

The AI Act has been **in force since 1 August 2024,** but **will apply in stages.** For HR / Talent Acquisition, it is important not to treat everything as fully applicable immediately, but to set the right priorities.

Relevant since 2 February 2025:

- **Prohibited AI practices**
- **AI literacy obligations**

  Particularly important for HR: certain practices, such as **emotion recognition in the workplace,** are prohibited; at the same time, organisations must ensure that the relevant roles have adequate AI literacy.

Relevant since 2 August 2025:

- **Governance rules**
- **Requirements for GPAI / general-purpose models**

  This mainly concerns providers of general-purpose AI models; for Talent Acquisition teams, it is important to note that the use of GPAI / LLMs is not unregulated, but must be embedded in a clear governance framework.

Generally applicable from 2 August 2026 – subject to the politically agreed adjustments under the AI Omnibus:

- most obligations under the AI Act
- including transparency obligations and the general requirements for many affected use cases

Under the AI Omnibus[1], the Council of the EU and the European Parliament reached a political agreement to simplify certain requirements and adjust their timing; final legislative implementation remains pending:

- For high-risk AI systems in certain areas – including employment, personnel management and recruiting under Annex III – application is planned from 2 December 2027.
- For high-risk AI systems integrated into regulated products, application is planned from 2 August 2028.

Practical rule for TA: The extension gives you time to prepare – not an excuse. Organisations that build AI inventories, use case documentation and governance structures will now avoid last-minute implementation pressure.

1 – Last updated: May 2026. The AI Act generally continues to apply from the dates already specified. At the same time, the Council of the EU and the European Parliament reached a political agreement under the so-called 'AI Omnibus' to simplify certain requirements, reduce administrative burdens and further harmonize implementation. Final legal implementation and publication in the Official Journal of the European Union remain decisive. Companies should therefore monitor both the applicable AI Act and the ongoing Omnibus amendments.

Executive Summary: AI Act and Talent Acquisition in 7 minutes

1.1 Purpose and benefits of this playbook

The AI Act explicitly classifies AI systems in the area of "employment and human resources management" as potential **high-risk applications** (Annex III). The purpose, target audience and methodology of this playbook are described in *Chapters 1–3*. The goal is not to understand legal texts, but to make decisions and take action **faster** and **with greater confidence.**

1.2 The five core messages (guiding principles of this playbook)

The following five core messages are the guiding principles for this white paper. They summarise what matters most when using AI in Talent Acquisition, regardless of which tools or providers are used. The statements are intentionally practical: they are intended to help structure discussions, speed up decisions and avoid misunderstandings. In the following chapters, the core messages are translated into specific classification guidance, measures and templates.

1 Not every 'AI' system is high-risk AI within the meaning of the AI Act

In practice, even simple, rule-based systems are marketed as 'AI-supported.' However, the AI Act defines AI more narrowly, for example, in line with the OECD model: systems that learn from inputs and generate predictions or decisions. The decisive factor is not the label, but the methods used and the system's effect within the process.

2 For Talent Acquisition, the effect is what counts. As soon as a system provides decision support, the requirements increase.

Pure information or FAQ use cases are often primarily a matter of data protection and transparency. However, systems that match, rank or recommend candidates influence selection decisions – a high-risk area under Annex III of the AI Act – and require significantly stronger controls. The CJEU ruling on the SCHUFA score (C-634/21) is also relevant here: the Court of Justice of the European Union clarified that creating an automated score can already fall under Article 22 of the GDPR if that score effectively determines a third party's decision – a logic that can be applied by analogy to AI-supported matching and ranking systems in the TA context.

3 Risk-based means: you do not have to treat everything the same, but you do need clear traffic light logic

This white paper provides a TA-specific risk map (green/yellow/red) and a stop/go check to classify use cases quickly and consistently.

4 In practice, compliance is determined by documentation, not good intentions

Auditability, traceability, documented decision logic, monitoring and human oversight are the building blocks that make implementation robust.

5 TA remains able to act when governance is designed as an operating model

A central AI Board makes sense, but TA needs its own process (use case one-pager, risk intake, decision pack) to prepare decisions in a board-ready format and reduce friction.

1.3 Traffic light logic: green/yellow/red (guidance for TA use cases)

Traffic light logic enables quick, standardised classification of typical Talent Acquisition use cases. The decisive factor is not the provider's 'AI-powered' marketing label, but the actual effect within the process: does the system merely provide information or does it provide decision support (e.g. through matching, ranking or recommendations)? The traffic light is a guide; borderline cases are clarified using the stop/go check *(Chapter 1.4)* and the TA risk map *(Chapter 5)*.

Green stands for information and support use cases without decision support. The focus here is on the immediate transparency obligation under the AI Act (labelling when interacting with AI systems, 'bot' labelling), data protection basics and clear usage rules.
Typical example: *Rule-based candidate FAQ bot. Note: As soon as an LLM or GenAI model is behind it, at least yellow applies – because even the question of permitted inputs and data hygiene triggers governance requirements.*

Yellow stands for workflow or assistance use cases that process personal data or support process steps without making an automated preselection. The focus here is on data classes / PII rules, approvals, role rights, appropriate logging and, depending on the design, co-determination.
Examples: *Interview/assessment support (assistance), GenAI job ads (quality/AGG), AI-supported sourcing without automated selection decisions.*

Red stands for use cases that provide decision support or influence decisions in Talent Acquisition. Practical dividing line: as soon as ranking, scoring or automated filtering ('shortlisting') takes place and a person no longer reviews it in detail, the red path applies. Under the AI Act (Annex III), these are typical high-risk applications. Strict controls are required here: explainability at the factor level, fairness/bias testing, documented human oversight, monitoring and reliable provider documentation.
Typical example: Screening/matching with ranking.

Traffic light logic for Talent Acquisition use cases

Classification by effect within the process: Information vs. decision support

GREEN
Inform/Support

Typical examples:
Rule-based candidate FAQ bot (pure FAQ/process information)

Note: LLM-based FAQ bots are considered at least yellow – as soon as an LLM is involved, the classification changes from green to yellow.

Focus: transparency obligations under the AI Act (labelling when interacting with AI systems), data protection basics (GDPR), usage rules

→ Further in the playbook: ***Chapter 4*** + ***Chapter 11***

GREEN means: pure delivery of defined/scripted information to users; only transparency obligations need to be met.

YELLOW
Support/Workflow

Typical examples:
Interview/assessment assistance, GenAI job ads, AI-supported sourcing (without automated selection), LLM-based FAQ bots

Focus: AI Act: governance and competence (Article 4) + transparency in the use context, GDPR (data classes/PII), role rights, approvals, appropriate logging, co-determination if necessary

→ Further in the playbook: ***Chapters 11–13*** + T-14.2 / T-14.3

YELLOW means: usually not a high-risk core use case, but clean governance, data hygiene, and clear responsibilities are crucial.

RED
Prepare/Decide

Typical examples:
Screening/matching, ranking, shortlisting, recommendations

Focus: AI Act: high-risk obligations (risk management, data quality, technical documentation/logging, transparency/explainability, human oversight, robustness), plus GDPR/AGG, monitoring, provider documentation

→ Further in the playbook: ***Chapters 5–10*** + T-14.3 - T-14.5

RED means: as soon as selection decisions are supported or deliberately influenced, verifiable controls are needed (explainability, fairness, oversight, monitoring).

Borderline cases are clarified using the stop/go check *(Chapter 1.4)* ttt. Short formula: Green = information/support without selection effect
Yellow = assistance/workflow with data relevance | Red = Decision preparation with selection effect

1.4 Make a decision in 3 minutes: Stop/go check in five questions

Classification under the AI Act (risk logic as guidance):

The AI Act follows a **risk-based approach.** For HR/TA, this means: you do not classify a use case by tool labels ('chatbot' or 'GenAI'), but by **its use context and effect within the process**. The following requirements are derived from this:

- **Prohibited applications (no-go):** certain purposes – especially **emotion recognition in the workplace** – must be excluded in the HR context.
- **Transparency obligations:** pure information/FAQ chatbots are typically primarily a matter of transparency and data protection (AI Act transparency obligations, labelling when interacting with AI systems, data hygiene, contact channel).
- **High-risk / strict path:** as soon as a system **provides decision support or influences** selection decisions (e.g., matching, ranking, recommendation, shortlisting), the red path is typically relevant (explainability, fairness, human oversight, monitoring, documentation).

Question 1 – Effect within the process: Does the system merely provide information, or does it provide decision support (matching, ranking, recommendation, shortlisting)?

Question 2 – Process step: In which step is the system used (sourcing, screening, interview/assessment, offer, onboarding)?

Question 3 – Data: What data is processed (e.g., CV/profile, communication data, internal data)? Are special categories of personal data or biometric data used (e.g., video/audio from virtual interviews), and for what exactly (identity check/verification vs. analysis of behaviour/emotions)?

Question 4 – Explainability: Can the result be explained based on factors, not just a score/ranking, so that it is understandable and relevant for action?

Question 5 – Control: Is there documented human oversight, appropriate logging and a review/escalation path, including responsibilities?

Note: As soon as **biometric** data or **special categories of personal data** are involved, sensitivity increases significantly. Article 5 of the AI Act **generally prohibits emotion recognition systems in the workplace.** Exceptions only apply for medical or safety-related purposes.

Next step in the playbook:

→ **Green:** *Chapter 4* (terms) and *Chapter 11* (data protection basics and data classes).

→ **Yellow:** *Chapters 11–13* (data protection, co-determination, governance) and T-14.2 + T-14.3

→ **Red:** *Chapters 5–10* (risk map, roles, transparency/explainability, fairness, monitoring) and T-14.3 + T-14.4 + T-14.5

No-go: Biometrics and emotion recognition in the HR context

Certain AI applications are particularly problematic in the HR context. While **biometric data** may generally be processed for identity verification, emotion recognition in the workplace must generally be excluded in the HR context (see Article 5 of the AI Act)

Practical rule: Procedures intended to evaluate or influence applicants or employees using emotion recognition are generally prohibited in the HR context, or only permitted under very narrow statutory exceptions. In case of doubt: stop the use case and obtain legal/data protection review and governance decisions. If a tool offers features such as 'emotion', 'mood', 'camera-based engagement', 'micro-expressions,' or similar categories, this is a clear red flag **and should be ruled out or subject to very strict review before use in the HR context.**

Stop/go in 3 minutes:
5 questions for classification

1 – Effect within the process
Does the system only provide information or does it prepare decisions (matching, ranking, recommendation, shortlisting)?

2 – Process step
Where is it used (sourcing, screening, interview/assessment, offer, onboarding)?

3 – Data
What data is included (CV/profile, communication, internal data; personal/PII/sensitive)?

4 – Explainability
Is there a fact-based justification (not just score/ranking)?

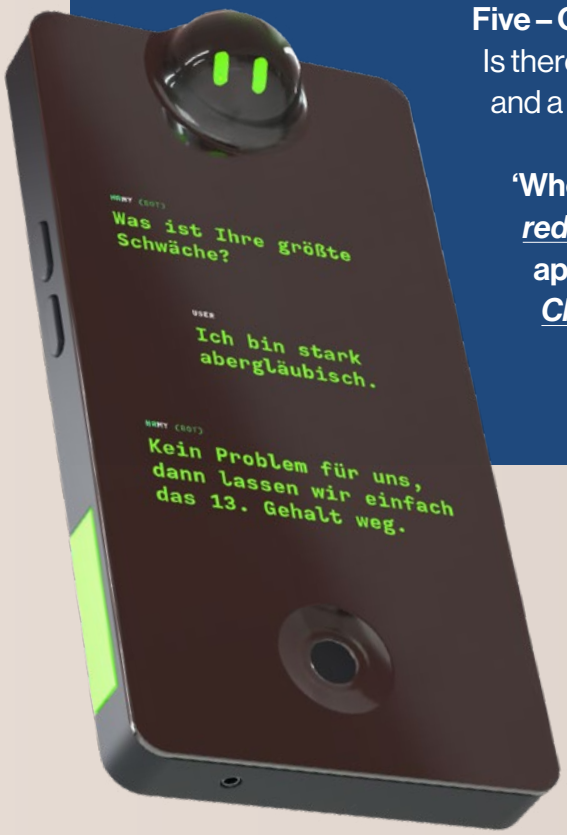

Five – Control
Is there human oversight, appropriate logging and a review/escalation path?

'When decisions are supported, the *red path* (Chapters 5-10) generally applies. Borderline cases are clarified in *Chapter 5*'.

1.5 Top 10 quick wins for Talent Acquisition (ready to implement immediately)

Many AI Act requirements appear complex at first glance. In TA practice, however, it quickly becomes clear: you gain confidence quickly if you do not try to solve everything at once, but instead start with a few repeatable standards. The following quick wins have been deliberately selected so that they can be started **without a major project**. They create a common language framework, shorten coordination loops with data protection, co-determination and IT, and make use cases board-ready early on.
Important: Compliance with the AI Act does not replace data protection under the GDPR. A 'double check' is always required.

Top 10 quick wins (TA)

1. Create an **AI inventory** ('egister light') for TA use cases (T-14.1).
2. Create a **use case one-pager** for each use case (T-14.2).
3. Define a standardised **risk intake** (AI Act/GDPR/AGG/BetrVG) (T-14.3).
4. Introduce a **decision pack (board-ready)** as the standard (T-14.4).
5. Select providers based on **minimum documentation** requirements (RFP check and scoring) (T-14.5).
6. Define minimum **explainability** requirements (factor-based, not score-only).
7. Define a minimum set for **fairness/bias testing** (incl. proxy risks).
8. Define a **monitoring/update principle** (quarterly review, release notes, incident path).
9. Define a **data class rule** for GenAI use in TA (what may be included, what never may be iincluded) (T-14.9).
10. Standardise a **works-council-ready package** (document list + communication logic) *(Chapter 12)*.

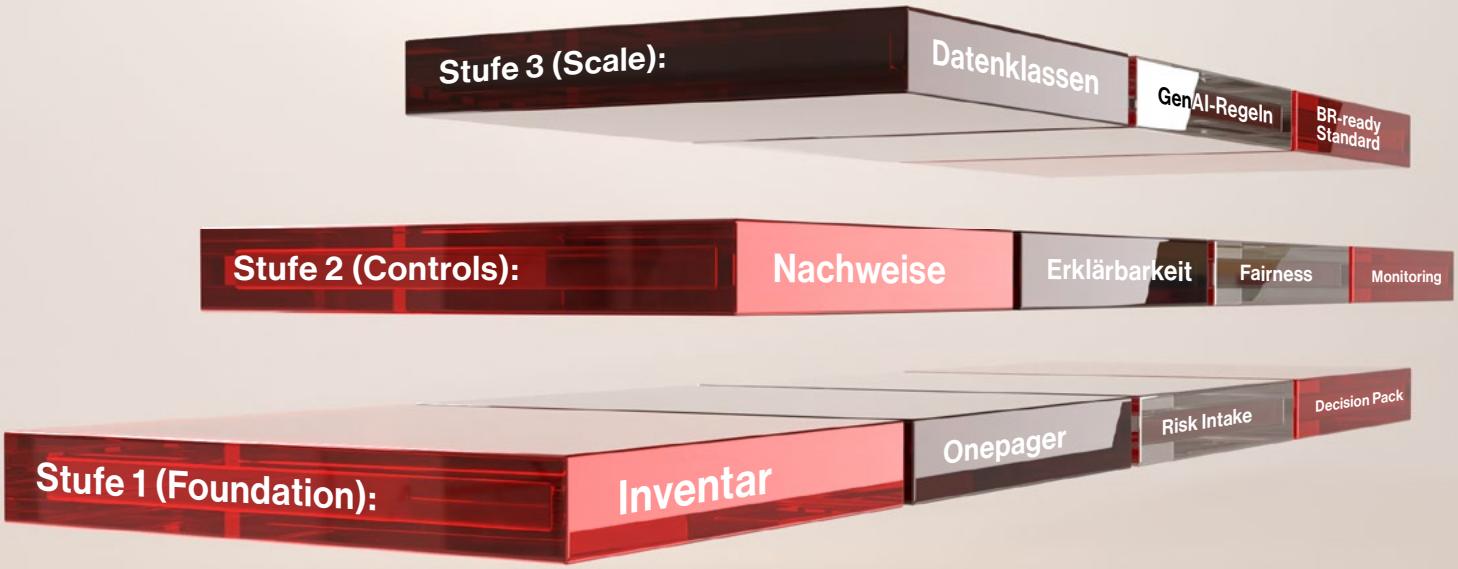

Quick wins as a step-by-step model

'If you start with only three things: 14.1 inventory, 14.2 one-pager, and 14.3 risk intake – you will already be much faster and more confident in any TA AI discussion.'

1.6 How to use this white paper in practice (reading logic and entry points)

Entry points (quick navigation):

• **Terms and scope** (chatbot/GenAI/sgent): *Chapter 4*

- **Classification of TA use cases** (traffic light, borderline cases, high-risk triggers): *Chapter 5*

- **Roles and responsibilities** (who provides which documentation): *Chapter 6*

- **General-purpose models and shifts in roles/obligations:** *Chapter 7*

- **Transparency/explainability and data subject perspective:** *Chapter 8*

• **Fairness/bias/proxy risks:** *Chapter 9*

- **Monitoring, updates and incident handling:** *Chapter 10*
- **Data protection / Article 22 of the GDPR and data classes:** *Chapter 11*
- **Co-determination/BetrVG and change management:** *Chapter 12*
- **Governance and competence building** (Article 4): *Chapter 13*
- **Templates, RFP questions and policies** (copy/paste): *Chapter 14*
- **Practical examples** (case nuggets) for guidance: *Chapter 15*

In the end, this white paper should do more than explain compliance; it should create a shared way of working. Once terminology, classification and minimum documentation are standardised, AI in Talent Acquisition does not slow things down; it becomes more predictable, easier to scale and easier to communicate to stakeholders.

Target Group, Scope and Work Paths (How to Use This White Paper)

2.1 Target audience (who this is written for)

The target audience, scope, and methodology are described in *Chapters 0.2* and *0.3*. Chapter 2 is limited to practical work paths *(Chapter 2.3)*.

2.2 Scope (what is included and what is deliberately excluded)

The focus is on TA use cases along the candidate journey, in particular:

- AI-supported sourcing, candidate FAQ bots, GenAI job ads
- Screening/matching/ranking/recommendations
- Interview/assessment assistance, including summary/transcription
- Operation (monitoring/updates/incidents), governance, co-determination, data protection

Deliberately not covered in detail: industry-specific special cases, full legal interpretation of individual provisions, detailed contract clauses (standard modules for these are provided in the Toolkit).

2.3 Work paths (how to achieve a result quickly)

Building on the explanations in Chapter 1.6 on how to use the playbook, it does not need to be read from beginning to end. There are three typical, practical work paths:

Work path A: 'I have a use case'
Chapter 4 (terms) → *Chapter 5* (traffic light/triggers) → *Chapter 6* (roles/documentation) → then *Chapters 8–13* as needed

Work path B: 'I am building AI governance'
Chapter 6 (operating model) → *Chapter 10* (operation) → *Chapters 11/12* (data protection / Works Council) → *Chapter 13* (competence building)

Work path C: 'I need to convince stakeholders'
Chapter 1 (Executive Summary) → *Chapter 5* (traffic light + stop/go) → *Chapter 16* (30-day plan)

2.4 Role of the Toolkit (separate document)

All templates and policies are contained in a **separate Toolkit document** *(Chapter 14)*. The white paper intentionally only refers to the templates in order to remain readable (option 1).

Quick start in 10 minutes

If you only have 10 minutes, that's enough:

1 ***Read Chapters 5.2/5.3*** **(traffic light + triggers)**

2 **Fill out the one-pager** (T-14.2) **for the use case**

3 **Then** (T-14.3) **complete the risk intake**

Key point: one-pager + risk intake = 80% clarity.

Three work paths

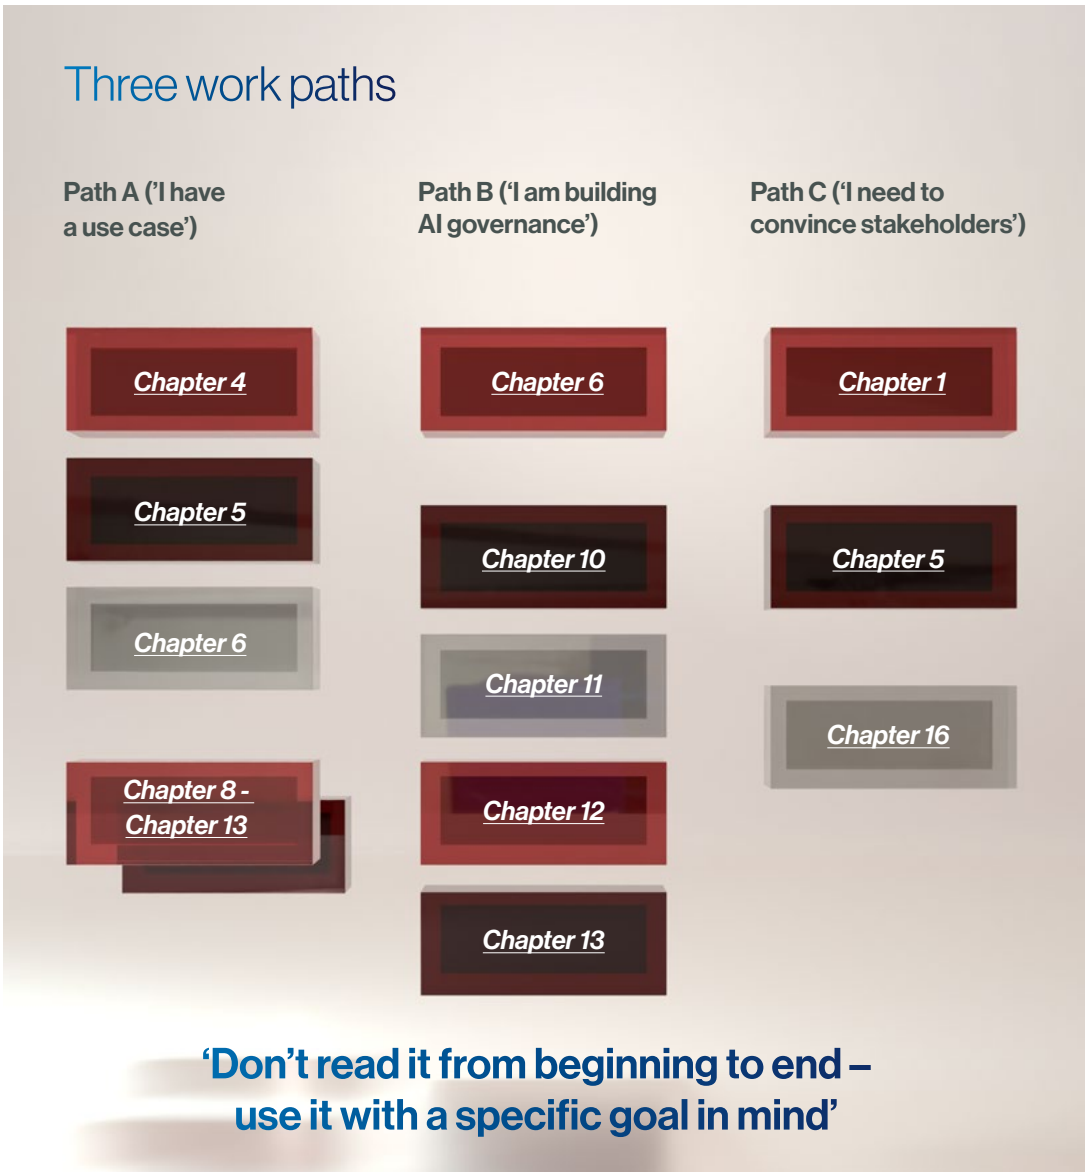

'Don't read it from beginning to end – use it with a specific goal in mind'

Regulatory Framework (AI Act) in TA: What You Really Need to Know

3.1 Purpose of this chapter

HR/TA does not need a detailed legal deep dive, but it does need a robust framework: **What logic underlies the AI Act, and what does it mean for Talent Acquisition?** This chapter provides the 'map' so that the later chapters (traffic light, roles, controls) do not feel like individual parts.

3.2 Core logic of the AI Act: risk-based rather than one-size-fits-all

The AI Act follows a **risk-based approach.** For TA, the key translation is: not every 'AI tool' is automatically critical. It becomes critical where systems **provide decision support or influence selection decision** *(Chapter 5*, red path). In addition, there are **no-go purposes** (e.g. emotion recognition in the workplace) that should be excluded in the HR context (see *Chapter 1.4* and *Chapter 11).*

AI Act timeline: What applies when? *(see Chapter 0.5)*
The full timeline (1 August 2024 to 2 December 2027) is shown in *Chapter 0.5*. **Here are the essentials:** the AI Act is generally applicable from 2 August 2026; for HR/TA-related high-risk applications under Annex III, simplifications and possible timing adjustments are currently being discussed under the AI Omnibus.

3.3 Practical application for Talent Acquisition: three risk classes

For Talent Acquisition, the AI Act logic can be pragmatically mapped into three risk classes:

1. **No-go/exclusion** (certain purposes)
2. **Transparency and data hygiene** (information / service use cases, chatbots, text assistance)
3. **Strict controls** for decision support (matching/ranking/recommendation/ filter, selection effect – high risk under Annex III of the AI Act)

The Traffic light logic *(Chapter 5)* translates these risk classes into concrete work logic for TA.

3.4 Related legal frameworks: the AI Act does not stand alone

In TA, legal implementation always runs on multiple tracks. The AI Act does not replace existing laws; it supplements them:

- **GDPR:** data, purpose limitation, data protection impact assessment (DPIA) and, in particular, **Article 22 of the GDPR** (prohibition of purely automated decisions with legal effect). An AI Act-compliant tool may not hire or reject candidates without human control.
- **AGG/fairness:** avoidance of bias and proxy discrimination *(Chapter 9)*
- **BetrVG/co-determination:** triggers for co-determination and works-council-ready packages *(Chapter 12)*.
- **Operational reliability:** monitoring, updates, and incidents *(Chapter 10)*.

Key point: The AI Act regulates how AI is governed from a product technology perspective. GDPR, AGG and BetrVG determine whether its use within the company is legally permissible and subject to co-determination.

The one principle TA should always apply

Effect beats label:

If a system **only provides information,** you need transparency and data hygiene.

If a system ranks **close to a decision** , recommends, filters or has a selection effect, the red path (high risk) applies: explainability, fairness testing, human oversight, monitoring and provider documentation are needed.

Regulatory stack

Chatbots, GenAI and Agents: Terms, Scope and TA Use Cases

4.1 Why this chapter matters

Vague terms lead to misinterpretations: Systems are blocked too strictly, or critical use cases are underestimated. *Chapter 4* creates the conceptual framework for risk classification *(Chapter 5)*.

4.2 What is an 'AI system' and when does it typically become critical in TA?

An **AI system** (under the AI Act) is a system that learns from the input it receives in order to generate predictions, recommendations or decisions. For TA, the decisive factor is not the marketing label, but the system's **effect and use context.**

Practical terms clearly distinguished:

- **A chatbot** is primarily about the form of interaction (FAQ/support) and can be rule-based or AI-based.
- **GenAI/LLM/GPAI** is a technology component, for example, for generating or rewriting text. It is not automatically an AI system within the meaning of the AI Act, but it can be part of one.
- **Agent** is a practical term for systems that execute or control tasks/workflows. An agent can use GenAI and, depending on its design, become decision-adjacent or make decisions itself.

Typically, a use case becomes critical in TA (red path) when a system provides decision support or influences selection decisions (Annex III of the AI Act), for example through:

- **Matching, ranking, scoring or prioritisation** of candidates or profiles
- **Shortlist suggestions** or automated filters ('not suitable', 'low match')
- **Recommendations** ('top candidates', 'best fit')
- **Evaluative summaries** that effecticely steer decisions
- **Agents/automations with selection effect** (such as automatic exclusion, automatic prioritisation or selective forwarding)

Under the AI Act, such uses in the area of employment/recruitment are **typically classified as high-risk** and require correspondingly strict controls: explainability, fairness/bias, human oversight, monitoring and robust documentation.

Less critical use cases (green/yellow) are those that primarily provide information or support without selection effect. Examples include FAQ bots or text assistance for job ads, provided that no personal applicant data (CV/profile) is processed and clear approval/quality processes are in place.

Key point: A use case does not become critical because it involves an 'agent'or 'GenAI', but because it has a **selection effect.**

4.3 The key distinction: information vs. decision support

The decisive factor is not the tool label, but whether the system has a selection effect (see *Output Catalogue* Box 4.C and Traffic light logic in *Chapter 5).*

Conceptual map: AI system, chatbot, GenAI, agent (practice)

An agent can use GenAI. A chatbot can be rule-based or AI-based.
The decisive factor is the **effect within the process** (information vs. decision support).

4.4 TA use cases: five illustrative use cases

This white paper uses five representative use cases commonly found in Talent Acquisition organisations. They serve as reference points for classification and implementation:

1. Screening/matching
2. Candidate FAQ bot
3. Interview/assessment support
4. GenAI job ads
5. AI-supported sourcing

The matrix below classifies these use cases according to their role in the process, the types of data typically involved, their typical impact and the resulting control requirements.

4.5 Use case matrix: use case → role in the process → risk → compliance requirements

The matrix serves as a quick guide. The traffic light is an initial classification logic; borderline cases are specified in *Chapter 5* using the stop/go check and specific triggers.

Screening/matching is typically used for preselection, prioritisation or recommendations. It often directly influences decisions and, in the TA context, is typically close to **high-risk** or **can be classified as high-risk,** especially where ranking/recommendations/filters with selection effect are present. This results in an increased need for controls, including explainability at the factor level, fairness/bias testing, documented human oversight, monitoring/update transparency and reliable documentation.

A candidate FAQ bot is generally used for information and service support. It is typically not decision-relevant and therefore usually falls into the green risk class. The need for controls often relates to data protection (personal data in dialogues), transparent labelling ('bot') and clear usage rules. An exception applies if the chatbot includes an **integrated application function,** meaning candidates can apply directly via the bot. In this scenario, the bot effectively becomes the entry point to the selection process and tends to fall into the yellow risk class.

Explicit **provider documentation** is required here, confirming that the bot does not influence ranking, prioritisation or selection decisions. Without this documentation, the yellow classification is difficult to justify.

Interview/assessment support can be used as an assistance function (guides, summaries, structuring). While support in preparing interviews is not in itself a high-risk use case, interview assistants that summarise and evaluate interviews generally fall into the red risk class and must meet the corresponding AI Act requirements.

GenAI job ads are job advertisements optimised with the help of generative AI tools. The focus here is typically less on high risk and more on quality, consistency, AGG-sensitive wording, approval processes, versioning and data hygiene (no personal data in prompts).

AI-supported sourcing supports the identification, prioritisation or outreach of potential candidates. Where prioritisation or ranking takes place, there is effectively decision support. Typical focus areas include transparency, fairness/bias testing, data protection and the lawful basis for the candidate database, documentation and clearly defined human review.

Use case matrix (the five use cases)

Use case	Typical traffic light classification
GenAI job ads	
Candidate FAQ bot (rule-based)	
Candidate FAQ bot (LLM-based)	
AI-supported sourcing	
Screening/matching	
Interview/assessment support	

4.6 Output catalogue: what counts as information and what counts as decision support?

A simple output catalogue helps ensure that classification does not depend on terminology. The decisive factor is what the system 'produces' within the process.

Typical outputs that are considered information/support (green/yellow):

- FAQ answers, process descriptions, status notes
- Neutral text suggestions without evaluation (e.g. wording suggestions for job ads)
- Organisational information without selective relevance

Typical outputs that are considered decision support (red):

- Ranking, scores, prioritisation, shortlist suggestions
- Recommendations ('top candidate', 'best fit')
- Screening notices that effectively filter candidates ('not suitable', 'low match')
- Summaries with an evaluative tone ('strong', 'weak', 'not a fit'), if they influence decisions ('strong candidate'), automatic exclusion.

Output catalogue (information vs. decision support)

'Information/support' vs decision support

- FAQ answers, process descriptions, status notes
- Neutral text suggestions (e.g. job ad) without evaluation
- Organisational information (e.g. scheduling options, reminders) without selection relevance
- Notes on next steps (neutral), without prioritisation/recommendation
- General information collections without ranking/scores

Decision support (typical):

- Ranking, scores, prioritisation, shortlist suggestions
- Recommendations ('top candidate', 'best fit')
- Screening notices that effectively filter candidates ('not suitable', 'low match'
- Summaries with an evaluative tone ('strong', 'weak', 'not a fit') if they influence decisions
- Automated process steps with selection effect (e.g. automatic exclusion, automatic prioritising, automatic forwarding of only certain profiles)

Key point: The decisive factor is the selection effect. Automation without a selection effect (e.g. confirmation of receipt, scheduling) is not automatically decision support.

4.7 Practice box: Common misunderstandings and how to avoid them

'It's just a bot, but it rates candidates or prioritises profiles'

NOTE: Then it's not just an information bot; it's treated as a *tool for decision support.* Anyone who uses a general-purpose model in a way that effectively supports selection decisions – ranking, recommendations or filtering – may move into provider-like responsibility (shift in roles and obligations (see *Chapter 7*).

'It's GenAI, so it's automatically high-risk'

NOT NECESSARILY. The decisive factor is whether GenAI influences decisions or processes personal/sensitive data.

'We only use recommendations, but make decisions ourselves'

NOTE: *Recommendations change behaviour.* Review criteria should therefore be defined, human oversight documented and traceability ensured.

'The system only outputs one score – that's enough'

NOTE: A *score without a factor-based explanation* is difficult to test and act on in practice. Clear factors and limits are required.

4.8 Do/don't: terminology and classification in TA practice

Do

- Describe use cases by their effect within the process (information vs. decision support).
- Separate data protection risks (PII, purpose limitation) from AI Act risks (influence of decisions).
- Define the following for each use case: purpose, data, output, oversight, documentation.
- Use traffic light logic and stop/go check as the standard before discussions escalate.

Don't

- Adopt vendor marketing terms as classification criteria.
- Treat all GenAI functions as high-risk across the board.
- Allow score outputs without explainable factors in critical TA steps.
- Launch pilots without responsibilities, documentation logic and a review path.

References

- Classification and traffic light logic: *Chapter 5*
- Roles/responsibilities and documentation: *Chapter 6* and T-14.5
- Risk intake: T-14.3
- Decision pack (board-ready): T-14.4
- Data protection / data classes and GenAI rules: *Chapter 11* and T-14.9

Risk Map for Talent Acquisition Use Cases

5.1 Why a TA risk map is necessary

Traffic light logic (green/yellow/red) is introduced as a classification in *Chapter 1.3*. This chapter adds triggers for borderline cases and specific control measures for each traffic light category.

TA risk map as a traffic light matrix

Use case	Typical traffic light	Brief explanation
GenAI job ads		Quality, AGG, approvals, data hygiene
Candidate FAQ bot (rule-based)		Pure information, no selection effect
Candidate FAQ bot (LLM-based)		Governance, data hygiene, role rights required
AI-supported sourcing		Red as soon as ranking/ prioritising occurs
Screening/matching		Selection effect, ranking, decision support
Interview/assessment support		Yellow: assistance; red: evaluation/recommendation

Clarify borderline cases using the *stop/go check* (*Chapter 1.4) + trigger list* (5.3).
A candidate FAQ bot is only classified as green if it is rule-based; LLM-based versions are classified as at least yellow.

5.2 Traffic light model (green/yellow/red): classification by impact, data and control

The traffic light model is a guide that combines three factors:

1. **Effect within the process** (information vs. decision support, i.e., output with direct selection effect),
2. **Data** (type/scope, personal/PII/sensitive),
3. **Level of control** (human oversight, appropriate logging, review/escalation channels).

Green: primarily information/support. No selection effect. Focus: transparency (labelling obligation under the AI Act transparency obligations) and data protection basics. Note: green only applies to rule-based systems. As soon as an LLM or GenAI model is used, at least yellow applies – regardless of whether the use case is informative.

Yellow: workflow support, processing of personal data, but no automated preselection. Focus: GDPR, governance, role rights.

Red (high-risk focus): providing decision support or influencing selection decisions. The high-risk classification for the HR sector typically applies here (Annex III of the AI Act). Focus: factor-based explainability, fairness/bias testing, documented human oversight, monitoring/update transparency and reliable provider documentation (CE marking, EU database entry).

Short formula: Green = information/support without selection effect | Yellow = assistance/workflow with data relevance | Red = decision support with selection effect

5.3 High-risk triggers in TA: when yellow turns red

Many TA applications are not 'automatically red'; they only become critical through their specific design. The following triggers typically mean that a use case should be treated as red, regardless of whether the system is called a 'bot', 'assistant' or 'GenAI':

- **Ranking/score/prioritisation** of applicants or profiles (even if 'only for support')
- **Shortlist suggestions** or automated filters ('not suitable', 'low match')
- **Recommendations** ('top candidates', 'best fit') as decision input
- **Evaluative summaries** with an evaluative tone that influence selection/evaluation
- **Automated triggering** of process steps with selection effect (e.g. exclusion, forwarding to pools, prioritisation)
- **Weak explainability** (score-only, no factor-based justification)
- **Lack of human oversight** or unclear responsibilities/review paths
- **Opaque updates / drift risk** without monitoring/release logic

The decisive factor is the **selection effect:**

- Does the automation make a selection (such as exclusion, prioritisation or selective forwarding)?
- Does the system actually influence the selection or assessment of facts or persons?

Please note: As soon as such a selection effect occurs, the red-path controls *(Chapters 6–10)* are required. This also applies if, in practice, an AI ranking becomes de facto preselection because low-ranking CVs or profiles are no longer or only reviewed in exceptional cases.

Efficiency gains through prioritisation are possible. However, low-ranked applications must not effectively drop out of the process solely on the basis of an AI score. Ranking, thresholds and exclusion logic must be explainable and supported by human review, spot checks, documentation and override options.

High-risk trigger list: when yellow turns red

A TA use case should typically be classified as red if at least one of the following triggers is present:

1. Ranking/score/prioritisation of applicants or profiles (even if 'only for support')
2. Shortlist suggestions or automated filters ('not suitable', 'low match')
3. Recommendations ('top candidates', 'best fit') as decision input
4. Summaries with an evaluative tone ('strong', 'weak', 'not a fit') if they influence decisions
5. Automated process steps with selection effect (e.g. automatic exclusion, automatic prioritising, automatic forwarding of only certain profiles)
6. Weak explainability (score-only, no factor-based justification)
7. Lack of human oversight or unclear responsibilities/review paths
8. Opaque updates / drift risk without monitoring/release logic

Key point: Automation without a selection effect (e.g. confirmation of receipt, scheduling) is not automatically red. It becomes red when outputs provide decision support or influence selection decisions.

5.4 Borderline cases: typical TA scenarios and decision aids

Borderline cases usually arise where systems 'assist' but indirectly influence decisions. Clear decision-making aids help in practice:

Interview/assessment support:

- **Yellow** if the system only structures guides or transcribes interviews without summary, evaluation or recommendation, and its use is clearly regulated. *Note:* Pure transcript summaries – even without explicit evaluative language – should always be treated critically from a regulatory perspective, because any summary has an implicit steering effect through selection and weighting. They usually fall under red.

- **Red** if summaries are evaluative ('strong', 'weak', 'not suitable'), scores are generated or recommendations ('hire' / 'not hire') are made that effectively steer the assessment.

AI-supported sourcing:

- **Yellow** if profiles are only identified/aggregated and the selection is made entirely by humans (without ranking/scores).
- **Red** if prioritisation/ranking/scoring is created, or if 'top' profiles are suggested automatically.

GenAI job ads:

- **As a rule, yellow,** because it is primarily on quality, AGG-sensitive wording, approval, versioning and data hygiene.
- **Red** applies to GenAI job ads when job advertisements are created or displayed on a personalised or individualised basis using applicant or profile characteristics and this creates a selection effect, such as profiling, scoring or automatic targeting of individual people. In all other cases, GenAI job ads usually remain yellow (quality/AGG, approval, versioning, data hygiene).

***Note:* External platforms and selection effect**

Even if an external platform's algorithm (such as LinkedIn's 'Top Applicant' function or job board matching), performs the preselection, the employer may be responsible as a deployer within the meaning of the AI Act, provided the system is used in the recruitment process. Recommendation: review platform terms for AI use and ensure transparency obligations toward candidates.

Candidate FAQ bot:

- **Green** if it provides neutral FAQs and process information and the bot is rule-based (no underlying LLM/GenAI).
- **Yellow** if an LLM/GenAI model is used – even for purely informational use – because data hygiene, approval and role rights then always trigger governance requirements. Also yellow, if personal data is linked or the bot intervenes in workflows.

- **Red** if the bot effectively pre-sorts or evaluates/prioritises applicants based on their answers.

Borderline case cards (four mini cases)

Use case 1: Interview/assessment support

Yellow if the system only transcribes (without summary, evaluation or recommendation) and its use is clearly regulated.

Red when summaries are evaluative (strong/weak/does not fit), scores are generated, or hire / not hire recommendations are made. Pure transcript summaries usually fall under red because of their implicit steering effect.

Use case 2: AI-supported sourcing

Yellow if profiles are only identified/aggregated and the selection is made entirely by humans, without ranking or scores.

Red if prioritisation, ranking or scoring is created, or top profiles are suggested automatically.

Use case 3: Candidate FAQ bot

Yellow if an underlying LLM/GenAI model is used, even for purely informational use, because this triggers governance requirements. The same applies if personal data is linked or the bot intervenes in workflows.

Red if the bot effectively pre-sorts, evaluates or prioritises applicants based on their answers.

Use case 4: GenAI job ads

Yellow if the focus is primarily on quality, AGG-sensitive wording, approval, versioning and data hygiene.

Red if job advertisements are created or displayed on a person-specific basis based on applicant/profile characteristics and this creates a selection effect, such as profiling, scoring or automatic targeting.

- **Use cases:** interview/assessment support; AI-supported sourcing; candidate FAQ bot; GenAI job ads
- **Purpose:** quick decision logic for everyday use, without continuous prose

5.5 Determining the next steps: recommended actions for each traffic light category

The traffic light is only useful if it triggers concrete next steps:

- **Green** → focus on transparency, usage rules and data protection basics. Entry via *Chapter 4* and *Chapter 11.*

- **Yellow** → focus on data classes, role rights, approvals, appropriate logging, co-determination depending on design and governance standard modules. Entry via *Chapters 11–13* and T-14.2 T-14.3

- **Red** → focus on role clarification (who provides what documentation), factor-based explainability, fairness/bias/proxy testing, documented human oversight, monitoring/updates and robust provider documentation. Entry via *Chapters 6–10* and T-14.3 T-14.4 T-14.5

Roles, Responsibilities and Contract Frameworks in TA Projects

6.1 Why this chapter matters

Delays usually occur not because requirements are unknown, but because responsibilities remain unclear ('Who provides the documentation?', 'Who reviews fairness?'). This chapter provides the role logic: who delivers what, which minimum documentation requirements are needed and how documentation is anchored in contracts.

6.2 Role model in TA language (simplified)

The AI Act distinguishes between different roles in the life cycle of an AI system. For Talent Acquisition, the most relevant factors are who provides the system and who uses it in a specific process.

- **Provider**: provides an AI system or places it on the market and is responsible for certain technical and organisational obligations (e.g., instructions for use, risk management, transparency, provider-side monitoring).
- **Deployer**: uses the system in the company for a specific purpose (e.g. recruiting). Deployer obligations are often organisational. Purpose definition, human oversight, deployment rules, stakeholder involvement, deployment documentation and operational monitoring.
  Important for TA: in recruiting, you are almost always the **deployer.** Deployer obligations include: use only in accordance with the instructions for use, ensure human oversight (see *Chapter 8.4),* monitor for failures/bias and comply with information obligations.

- ***Importer/distributor:*** roles in supply chains (relevant for TA when systems are obtained via partners/resellers). In practice, this is often handled by Procurement and IT/Legal.

In addition, depending on the data and implementation, obligations may arise under the GDPR – in particular, the need for a data protection impact assessment (DPIA) (see *Chapter 11.5)*.

The decisive factor is whether the organisation takes on additional functions that would normally lie with the provider (e.g., significant independent changes to purpose, logic or model). These potential role shifts are discussed in more detail in *Chapter 7*.

Key point: *Attention: role shift:* if the deployer makes significant changes to a system (e.g., repurposing an LLM for CV screening), it can legally become the provider and assume the associated obligations (see *Chapter 7)*.

Role model (simplified)

Provider	Deployer	Importer/ distributor
provides the system + documentation + updates	uses the system + defines purpose/ context + oversight	when obtained via a partner

TA is at least a deployer when using the system in a process.

6.3 Required documents (minimum set): what documentation and working documents you need

Required documents (minimum set) (for yellow and red):

1. **AI inventory entry** (which use case, which system, which process)
2. **Use case one-pager** (purpose, process step, data, output, responsible person)
3. **Risk intake** (classification under the AI Act/GDPR/AGG/BetrVG including Section 95(2)(a) of the BetrVG + measures)
4. **Decision pack** (board-ready) (short version for approval, including risk/controls)
5. **Operating concept** (minimal): Role rights, traceable logging, update/monitoring principle, incident path

For red use cases, additional documentation is typically required; this must be requested from the provider and supplemented internally: instructions for use under Article 13(1–3) of the AI Act, including information on explainability, fairness/bias testing, ongoing monitoring and human control. In addition, the documentation must include a conformity assessment under Article 43 of the AI Act, the associated EU declaration of conformity and the manufacturer's CE marking under Articles 48 and 49 of the AI Act. The application must be registered in the central EU database for high-risk AI systems. For HR high-risk AI systems under Annex III No. 4, internal control under Article 43(1) in conjunction with Annex VI of the AI Act is generally sufficient, provided harmonised standards have been applied. Otherwise, the involvement of a notified body under Annex VII may be required.

Templates for the one-pager, risk intake, decision pack, monitoring and incident runbook are available in the toolkit. T-14.2 One Pager | T-14.3 Risk intake | T-14.4 Decision pack | T-14.6 Monitoring | T-14.7 Incident Runbook

6.4 'Who delivers what?' – responsibilities of the provider and deployer

The HR/TA department is typically the deployer within the meaning of the AI Act. Provider obligations such as technical documentation, conformity assessment and CE marking generally remain with the respective system provider.

Typically on the provider side (to be requested):

- Requirements for use ('instructions for use') and technical information
- Documentation on explainability (e.g., factor/feature explanation instead of score-only)
- Information on the data basis and risks/tests (bias/fairness approaches, quality)
- Information on updates / change management and provider-side monitoring
- Exportable logs/auditability (where intended)
- Declaration of conformity / CE marking (as soon as the AI Act applies)

Typically deployer-side (within the company, usually to be provided by HR/TA):

- Purpose definition and use context (where/who/how in the process)
- Data class rules, access policy, role rights
- Human oversight (who checks, how overrides are handled, how this is documented)
- Involvement of data protection (data protection impact assessment), co-determination and IT security (where required)
- Operating concept (training, monitoring during use, incident handling, responsibilities)

6.5 Contract and RFP framework: minimum documentation requirements as a quality filter

For yellow and red use cases, it is advisable to request standardised minimum documentation, e.g.:

- Use documentation / instructions for use (incl. permissible and impermissible use cases, operating error risks and recommended control mechanisms)
- Traceability/explainability (factor-based, not just score)
- Auditability/appropriate logging (export, traceability)
- Update/release transparency + monitoring approach
- Fairness/bias documentation (methodology, rhythm, retesting after updates)
- Data protection/security documentation (roles/rights, data flows, storage)
- Clear responsibilities and support/incident processes

This turns the RFP into a genuine quality filter: providers that cannot provide this documentation – or merely claims it for marketing purposes – are structurally unsuitable for critical **TA use cases that fall under the AI Act in the EU (e.g., because the system is deployed in the EU or its output is used in the EU).**

Documentation matrix

Provider supplies / deployer supplies

Use case	Provider	Deployer
Instructions for use / technical documentation	✓	–
Explainability (factor-based)	✓	–
Bias/fairness documentation	✓	–
Logging / audit export	✓	–
Release notes / updates (provider side)	✓	–
Monitoring routine in operation	–	✓
Purpose / use context	–	✓
Human oversight	–	✓
Access / role rights	–	✓
Works Council / DPO involvement	–	✓

6.6 What to do if the provider does not provide (complete) documentation?

Step 1: Classification – dealbreaker or compensable?

Dealbreaker (compensable) → stop use case or change provider:

- No factor-based explainability possible (score / black box only) for a red use case.
- No fairness/bias documentation (no methodology, no tests, no willingness to disclose) for matching/ranking.
- No instructions for use / no documentation of data processing for use cases involving personal data (GDPR obligation).
- Provider rejects audit-appropriate logging export as a matter of principle.

Compensable → deployer provides documentation itself:

- Provider does not have a complete monitoring plan → deployer establishes its own review routine *(Chapter 10)*.
- Provider has not performed sibling CV tests → deployer performs its own sample tests *(Chapter 10)*.
- Provider does not yet have CE marking (before the 08/2026 deadline) → obtain contractual assurance, request implementation plan.
- Provider explains explainability only on request → include SLA clause requiring factor output to be provided by default.

Step 2: Secure what is possible contractually

What the provider cannot deliver today must be contractually established as an obligation with a deadline. It is not sufficient to describe it as 'being pursued'. Minimum contractual modules: documentation obligation (explainability, fairness, logs), information obligation for updates/changes, right to audit/sample testing, SLA for incident response and, for red use cases, assurance of CE conformity by the time the AI Act requirements apply.

Key point: **Anyone who fails to provide dealbreaker documentation for critical TA use cases is structurally unsuitable – regardless of how well the product works.**

6.7 Practice workflow: from idea to approval (board-ready)

A standardised process is recommended so that TA does not get stuck in long coordination loops:

1. **Describe the use case** (one-pager) and roughly classify it (stop/go + traffic light)
2. **Perform the risk intake** (AI Act/GDPR/AGG/BetrVG)
3. **Create the decision pack** (approval basis for AI board / steering committee)
4. **Carry out provider selection and RFP**: request minimum RFP documentation (T-14.5), evaluate offers based on explainability, fairness documentation and compliance suitability, identify dealbreakers *(Chapter 6.6)* and conclude a contract with compliance clauses.
5. **Request documentation** (provider) and define internal measures
6. **Approval + start of operation** with clear role rights, monitoring principle, incident path
7. **Regular reviews** (especially for updates/drift/process changes)

Process flow: Use case → decision pack → approval → operation

RFP quick check (10 questions)

1. Is the system classified as an AI system under the AI Act and, if so, in which category (Annex I, Annex III or GPAI)?
2. Which functions provide decision support (matching, ranking, scoring, shortlisting), and can these functions be configured or disabled?
3. How is explainability implemented? Are factor-based justifications (not just score outputs) provided as standard?
4. Which fairness/bias tests were carried out? Please provide the methodology, dimensions tested, thresholds and results.
5. How are updates and model/prompt changes communicated? Are there release notes with change impact information?
6. Can logs / audit trails be exported? Please specify format, scope and retention period.
7. How is the incident process structured? Who is the contact person and what SLA applies to critical AI incidents?
8. What data protection documentation is available (GDPR compliance, record of processing activities, DPA, data storage locations)?
9. Is a CE marking or EU declaration of conformity under Articles 48 and 49 of the AI Act available and, if not, when will it be available?
10. Are there reference customers in the EU/DACH region who can confirm use in a regulated HR/TA environment?

Reference: T-14.5 (complete)

6.8 References to templates (see Chapter 14.3) and subsequent chapters

- **KI inventory:** T-14.1
- **Use case one-pager:** T-14.2
- **Risk intake:** T-14.3
- **Decision pack:** T-14.4
- **RFP questions and minimum documentation requirements:** T-14.5
- **Role/obligation shift for general-purpose models:** *Chapter 7*
- **Transparency/explainability, fairness, monitoring:** *Chapters 8 – 10*

General-Purpose Models (GPAI/LLMs) in TA and Possible Shifts in Roles/Obligations within TA Projects

7.1 Why this chapter matters

General-purpose models (GPAI; AI systems that can be used for a variety of tasks, such as language models, image generators or autonomous problem solvers) and large language models (LLMs) are increasingly being used in Talent Acquisition – whether as text assistance, for chatbots, for interview support or as building blocks in agent workflows. This creates productivity opportunities, but also brings a clear governance risk: **it is not the model itself that is decisive, but its use and impact in the TA process.**

For TA, this means: as soon as a usage pattern has a **selection effect** (e.g. ranking, recommendation, filter, shortlisting or evaluative summaries), the use case is typically classified as red and requires strict controls (transparency, factor-based explainability, fairness/bias testing, documented human oversight, monitoring and robust documentation).

The AI Act also includes an important role logic: anyone who integrates or modifies a general purpose model in such a way that it becomes a **decision-relevant** recruiting system may, depending on the design and assumption of responsibility, move into a **provider-like responsibility** (role/obligation shift; see *Chapter 7.3)*. This is precisely why clear guiding principles on **purpose, data, control and operation are needed for GPAI/LLM use in TA**.

7.2 Core principle: the use case makes the difference

Green not considered in this chapter: every GPAI/LLM use in TA triggers at least yellow requirements. Human Oversight in practice: see *Chapter 8.4*.

- **Typical yellow (assistance without selection effect):** text assistance for job ads without applicant data/CVs/profiles in prompts; transcription without summary or evaluation; communication support according to clear rules (e.g., translation, drafting, structuring).
- **Typical red (decision support / selection effect):** anything that prioritises, ranks, recommends, filters or **evaluates** applicants, especially where outputs systematically influence selection decisions or applicant data (CV/profiles/interview content) is systematically used in prompts/inputs.

General purpose in the TA context

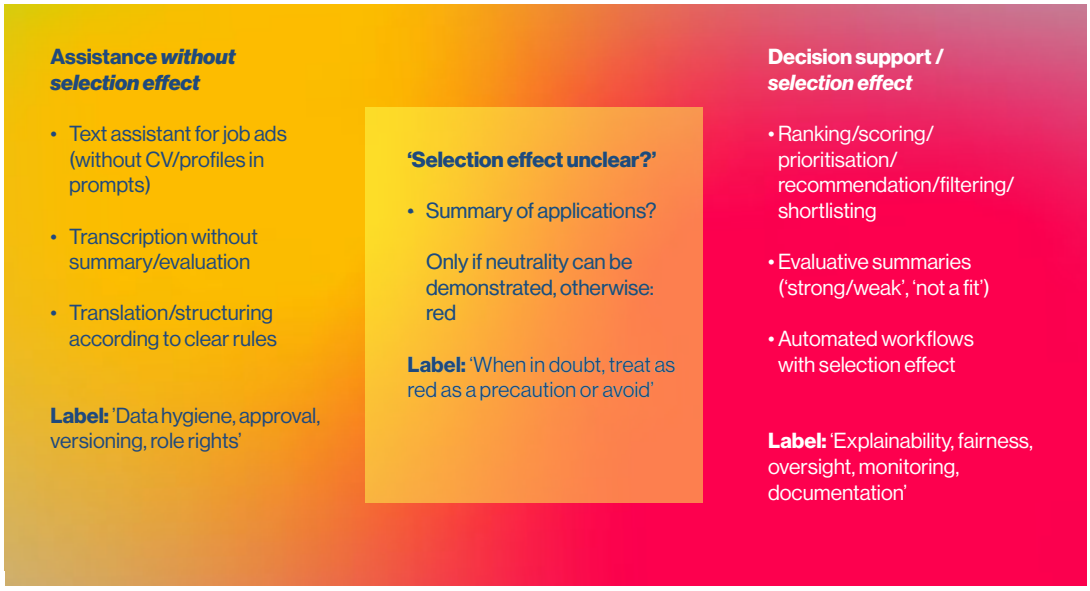

7.3 Typical TA scenarios with elevated compliance risk

Typical rote scenarios in practice:

- **LLM as a screening assistant:** applications are summarised and provided with recommendations ('strong/weak', 'good fit') or profiles are highlighted as "top".
- **LLM-supported ranking/prioritising:** profiles are prioritised, scored or merged into 'top candidates'.
- **Automated workflows with selection effect:** an agent triggers process steps (e.g. automatic prioritisation, exclusion or selective forwarding) based on the model output.
- **Link to applicant data:** CV/profiles/interview content are systematically used in prompts/inputs to change requirements, communication or selection rules.
- **Substantial change in purpose or logic:** the model is integrated or adapted in such a way that it takes on a new function as a result (e.g. from text assistance to suitability assessment / decision support).

Consequence: In these cases, documented human oversight, factor-based explainability, fairness/bias controls, monitoring/update transparency and robust documentation are required. If the model/system provider does not cover these requirements, the user must provide the necessary documentation themselves, with the corresponding risk of incurring provider-like responsibility.

7.4 Practical countermeasures: this keeps the use controllable

The safest path 'do not use LLM', but rather to design its use in such a way that **impact, data and control** remain manageable.

The decisive factor is whether the system only provides support (typically yellow) or prepares selection decisions (typically red; close to high risk in the employment/recruitment context; see Annex III of the AI Act).

1. Clearly define purpose (limit effect)
Define what the use is permitted for (e.g., text assistance, translation, structuring guidelines) and what it is not permitted for (ranking/recommendation/filtering). Document the context of use in the use-case one-pager and, if necessary, in the decision pack (purpose, process step, output type, responsibilities).

2. Apply data class rules consistently (limit data) Personal applicant data and sensitive data is used only where there is a clear legal basis, purpose limitation and appropriate controls (including deletion logic and no misuse, e.g., for 'training'). For text assistance: no individual applicant profiles (CV/profile) in prompts; use primarily job/role-related information.

3. Design controls according to risk (yellow vs. red)

- **Yellow (assistance):** data hygiene, approval processes (e.g., dual control principle for published content), role rights and traceable versioning.

- **Red (decision support):** documented human oversight (review criteria + override), fact-based explainability, fairness/bias controls and an operating concept (release/change impact, monitoring, incident path: Stop/Override/Escalation).

4. Practical guiding principle: do not use public LLMs for TA applicant data
Freely available, unregulated LLM services must not be used with individual applicant data (CV/profile), neither for summarisation nor evaluation (see Article 5 of the AI Act and GDPR Article 25, 'Privacy by Design'). As a general rule, they should not be used for high-risk TA applications (red); these require controllable enterprise systems with clear purpose limitation, governance and auditability.

Control wheel ‘Effect – Data – Control – Operation – Documentation’

Segment 1) Purpose/effect:
- Specify purpose (permitted vs. forbidden)
- Avoid selection effect or treat as red
- Define output type (information vs. decision support)

Segment 2) Data (data classes):
- No CVs/profiles in open LLMs
- Data classes + deletion logic
- Purpose limitation/ minimisation

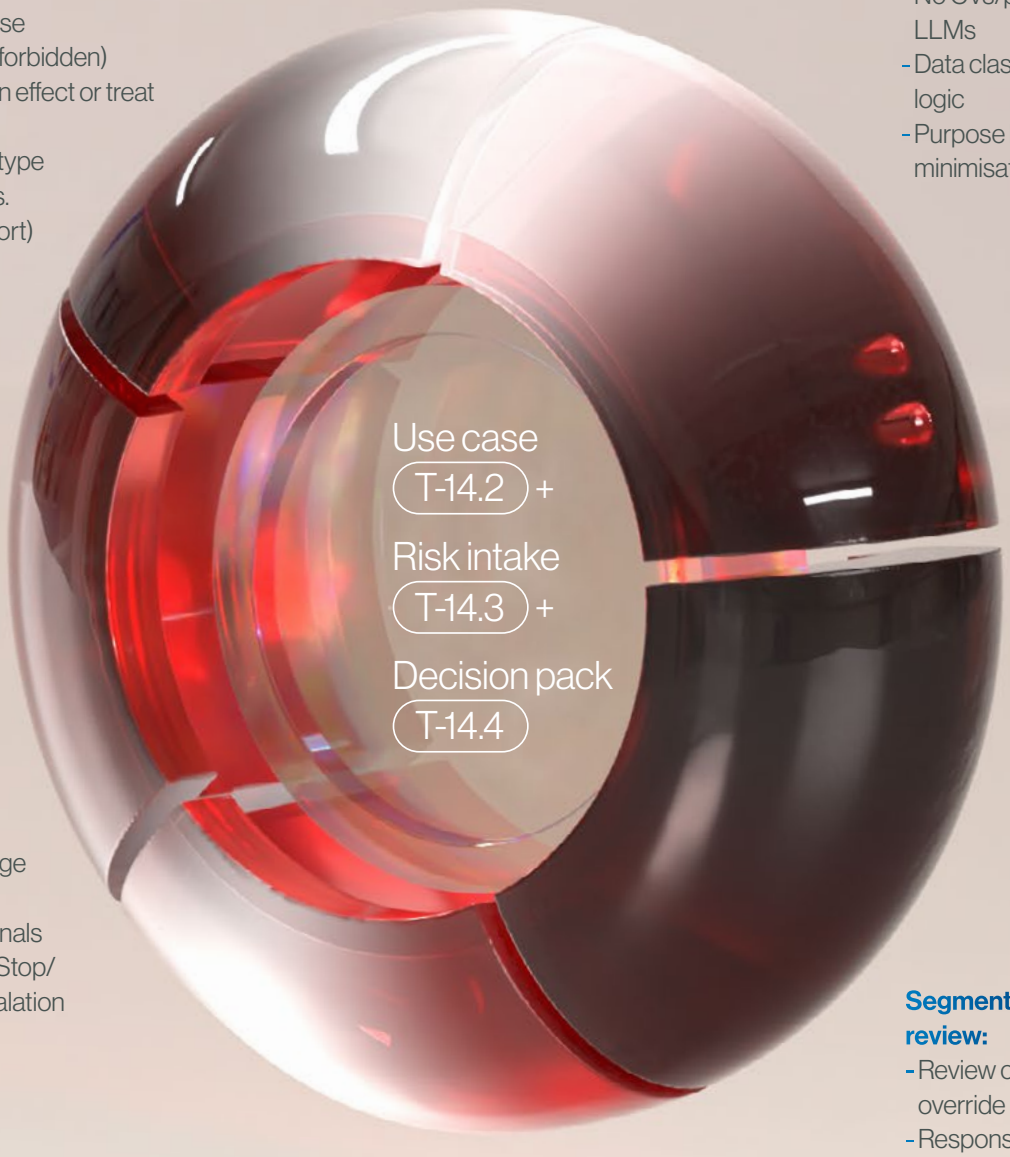

Segment 5) Operation (monitoring/ incidents):
- Release/change impact
- Monitoring signals
- Incident path: Stop/ Override/Escalation

Segment 3) Oversight/ review:
- Review criteria + override
- Responsibilities
- Appropriate documentation

Segment 4) Traceability:
- Factor-based explainability (not score-only)
- Versioning/reasons for change
- Traceability/logs (where necessary)

’When selection effects arise: make controls + operation verifiable *(Chapters 8–10)’.*

7.5 ‘f you do X, you need Y’ (a quick reference for TA)

This white paper uses traffic light logic (green/yellow/red) for practical orientation. The AI Act itself does not use these colours; the key legal question is whether a use case is classified as high risk or not. Traffic light logic serves as a working logic to determine the appropriate control depth at an early stage.

- **Suggested texts for job advertisements (without applicant data):** data hygiene, approval process, versioning, AGG check (typically yellow).
- **Transcription without summary/review/recommendation:** quality/ accuracy tests (including relevant languages), data classes/PII rules, role rights, traceable logging (typically yellow).
- **Summary of applications (including without score):** only with a clear role description, documented neutrality/quality criteria, data class rules and review. If neutrality/verifiability is not reliable: treat as red or avoid as a precaution.
- **Evaluation/prioritisation/ranking/recommendation:** typically red (explainability, fairness/bias, human oversight, monitoring, provider/system documentation).
- **Automated process steps with selection effect:** typically red plus clear control limits, audit capability / appropriate logging and incident runbook.
- **Applicant data is used to specifically change requirements/rules:** typically red plus integrity/governance risk (fairness/bias/AGG) and mandatory approval/justification logic.

In case of doubt: If the neutrality/explainability of an LLM output cannot be reliably demonstrated, the use case is treated as red or is not used.

If you do X, you need Y

Quick reference for controls and templates

Text assistant for job ads (without applicant data)	Approval process, versioning, AGG check, data hygiene T-14.2 / T-14.9
Transcription without summary/ evaluation	Accuracy tests (languages), role rights, deletion logic, traceable logging T-14.2 / T-14.9
Summary of applications/ interviews	Neutrality/quality criteria + review; **if not reliable: red or avoid** T-14.3 / T-14.4
Rating/ranking/recommendation	**Red:** Explainability (factors), fairness/bias, oversight, monitoring, documentation T-14.3 – T-14.7
Automated workflows with selection effect	**Red** + audit capability / appropriate logging + incident runbook T-14.6 / T-14.7
Applicant data for changing requirements/rules	**Red** + integrity/fairness risk + release/justification logic T-14.4 / T-14.5 / T-14.9

Traffic light logic is working logic; legally, high-risk classification is what matters.
If in doubt: stricter as a precaution.

Red triggers:
When GPAI/LLM use in TA is typically red

An LLM use case should typically be classified as red if at least one of these triggers is present:

- Ranking, scoring, prioritisation or shortlisting of applicants/profiles
- Recommendations as decision input ('top candidates', 'best fit')
- Filter/exclusion logic ('not suitable', 'low match')
- Summaries with an evaluative tone or implicit recommendation
- Automated workflows with selection effect (forwarding/exclusion/prioritisation)
- Systematic use of applicant data (CV/profile/interview content) in prompts/inputs
- Lack of documented human oversight (review/override) or lack of monitoring/release transparency

Key point: Automation alone is not a trigger. The decisive factors are selection effect and verifiability.

7.6 References to templates and subsequent chapters

- **Use case one-pager:** T-14.2
- **Risk intake:** T-14.3
- **Decision pack:** T-14.4
- **Data classes / GenAI rules:** *Chapter 11* + T-14.9
- **Explainability/fairness/monitoring:** ***Chapters 8–10***

Transparency, Explainability and Data Subject Rights

8.1 Why this chapter matters

In practice, projects rarely fail because of a 'no' to AI, but because results cannot be explained ('Why is candidate A ranked ahead?') or affected individuals are not informed in a comprehensible way. This chapter describes how transparency and explainability can be made practicable.

8.2 Transparency in TA processes: what should be communicated

Transparency in this context means making clear **where** AI is used, **what role** it plays in the process and **what effect** it may have. A distinction must be made between:

- **Information/support (green/yellow):** this is where the labelling obligation (transparency obligations under the AI Act: labelling when interacting with AI systems). It must be made clear that users are interacting with an AI system (e.g. chatbot) and a contact option must be offered.
- **Decision support (red):** for decision-related use cases (red), affected individuals **must be clearly informed in advance** about **where and how** AI is used in the process (e.g. matching, ranking, recommendations) and what role humans play in this. Depending on the structure, **co-determination** may also be affected *(Chapter 12)*.

Important (Article 22 of the GDPR): As soon as an AI system not only prepares decisions but also makes automated decisions – in particular where decisions **are made exclusively** by automated means and have legal or similarly significant effects (e.g., automatic rejection) – the requirements under **Article 22 of the GDPR apply.**

Then, among other things, affected individuals must have the opportunity to **request human intervention**, present their point of view and **challenge the decision.**

Transparency cannot be achieved through a one-off declaration, but should be part of the process design: candidates should always be able to understand how the process and AI support work without being overloaded with technical details.

8.3 Explainability: factors instead of scores

In TA, an explanation is helpful if it is comprehensible. A pure score (e.g., '72/100') does not provide a comprehensible justification and is of limited use for departments, Works Councils or audits. In the playbook sense, explainability therefore means:

- You can refer to **professional factors** that make the result plausible (e.g. relevant experience, skills, qualifications, requirements).
- Limitations can be specified (e.g. data gaps, uncertainties, different profiles).
- You can explain which type of output was generated (ranking/recommendation/ filtering vs. neutral reference).

Key point: Explainability does not mean 'disclosing all model parameters', but a **comprehensible, professionally coherent justification** at the factor level.

8.4 Data subjects rights and review procedures (practical and effective)

In Talent Acquisition, the perspective of affected individuals should not only be considered in the event of a conflict. A good standard is a **clear review path**: Who reviews a follow-up query? What information can be provided in an understandable way? What documents are available (e.g., use case one-pager, risk intake, decision pack)?

For **red use cases**, it is also crucial that **human oversight** is not merely claimed, but is effective in practice. 'Simply agreeing' is not enough. This requires:

- **Defined review/override criteria:** when and why is an AI proposal overridden?
- **Documentation to an appropriate extent:** what decision was made, on what grounds and in what context?
- **Explainability:** the role played by AI in the process must be described in a clear and comprehensible way (not just as a score).
- **Contact and escalation channel for affected individuals:** there is a clear process for queries and, if relevant, for **challenging AI-supported decisions** (especially in the context of Article 22 of the GDPR).

Factor level vs. score level (quick check)

Score level (typically not sufficient):

- 'Match = 72/100'
- 'Candidate A is better than B' (without justification)

Factor level (practicable):

- 'High fit based on X years of relevant experience in ...'
- 'Strong match for skills A/B/C; gap for skill D'
- 'Note: Data situation incomplete at ...; result therefore subject to uncertainty'

Key point: In TA, an explanation must be formulated in such a way that hiring managers and HR understand it and that it is comprehensible in a governance/audit context.

Practical examples: what counts as genuine human oversight and what does not?

This counts as genuine oversight:

Example 1: The recruiter sees a top-5 shortlist with a factor-based justification (e.g., 'strong fit for skill A/B, gap in skill C'). They check each candidate against the requirements profile. Candidate 3 was ranked lower by the AI because of the skill C gap. However, the recruiter knows from the job context that skill C is secondary here – and forwards candidate 3 despite the AI flag. Candidate 5, by contrast, is rejected because their skill fit is significantly lower than that of the other four. → This is genuine oversight: the recruiter corrects an AI signal based on contextual knowledge – with a comprehensible criterion and documented deviation.

This does NOT count as genuine oversight:

Example 2: The AI provides a Top-5 list. The recruiter clicks through all five without performing a single check – and forwards them because 'the AI has already filtered them'. → Overreliance. No oversight, only forwarding. Not justifiable in an audit.

Example 3: A score of '82/100' is displayed. The hiring manager approves it without knowing which factors make up the score. There is no possibility to override it and no documentation. → Neither transparent nor controllable. Not permitted for red use cases.

Key point: Human oversight does not mean 'a person has seen it', but 'a person has reviewed it, was able to override it, and documented this in a comprehensible way'.

Structure example – level of detail depending on the use case

8.5 Practical minimum standard (TA)

Minimum standard per use case:

1. **Transparency text** (short, comprehensible, per process step)
2. **Explanatory logic** (factor level, not score-only)
3. **Review/contact channel** (responsible role, SLA/deadline)
4. **Documents** (use case one-pager, risk intake, decision pack; see Toolkit)

Fairness, Bias and Proxy Discrimination in Talent Acquisition

9.1 Fairness determines acceptance and risk

Fairness is not a nice-to-have. AI is not 'objective'. It can reinforce existing patterns, even without intention, because it learns from historical data and because humans often trust AI outputs too much (overreliance). In Talent Acquisition, legal and reputational risks especially in the context of the AGG and the AI Act) therefore usually do not arise from intentional discrimination **indirect effects:** proxy characteristics (e.g., education, company names, regions, language patterns), data gaps and seemingly neutral ranking logic.

9.2 Bias in TA: where risks typically arise

Bias can arise or be amplified in several places. Not only in the model, but also in data and processes:

- **Data basis:** historical recruitment, interview or performance data may be biased (e.g. underrepresentation of certain groups). As a result, the system 'earns' patterns that favoured the majority group in the past.
- **Feature/proxy effects:** in practice, characteristics that appear neutral are often proxies for protected characteristics (e.g., certain educational paths, well-known company names, gaps in CVs, postal codes / regions, language patterns).
- **Process design (overreliance):** if recruiters trust scores/rankings too much or there is no clear review logic, distortions are not identified, but 'waved through' in everyday life (e.g., a top-5 shortlist is no longer critically reviewed).
- **Updates and drift:** model behaviour can shift gradually as a result of provider updates, configuration changes or modified data streams, without this being immediately apparent in the process.

Key point: Fairness is not a one-off test, but a combination of data awareness, process controls and monitoring.

9.3 Proxy discrimination explained clearly

Proxy discrimination means discrimination by proxy: A system uses characteristics that appear 'neutral' at first glance, but in practice correlate strongly with protected characteristics (e.g., age, gender, origin, disability). This can result in **indirect disadvantage** even though the protected characteristic is not explicitly processed.

Typical proxy examples in TA:

- **Educational paths/institutions:** can be proxies for socioeconomic background or access opportunities.
- **Company names / career paths:** can indirectly reflect access opportunities, networks or origin.

- **Postcodes/regions:** can be proxies for social structure, socioeconomic status or origin.
- **Gaps in CVs:** can be proxies for illness, pregnancy / parental leave or care periods (risk of indirect discrimination).
- **Language patterns/wording:** can be proxies for native language or access to education.

Practice point: It is not enough to check whether protected characteristics are processed directly. The decisive factor is whether the results logic **indirectly** leads to systematic disadvantages and whether these effects are identified and limited through suitable controls.

9.4 Minimum standard: what must be checked for red use cases

For **red use cases** (matching/ranking/recommendation with selection effect), a **clear minimum standard is required.** The goal is not perfection, but a **repeatable review logic** that is internally explainable and auditable. The minimum standard should be documented internally and include:

- **Transparent inputs and outputs:** what data/characteristics are included? What type of output is generated (ranking, score, recommendation, filter/shortlist)?
- **Plausibility check of factors:** are the factors considered job-related, objectively justifiable and relevant to the role?
- **Fairness/bias check:** are there tests for systematic disadvantage (methodology, thresholds, evaluation of results)?
- **Proxy review:** Are potential proxy features identified and actively limited or excluded?
- **Human oversight and avoiding overreliance:** clear review criteria, override option and documented handling of outliers/conspicuous results.
- **Measure outcomes:** are there noticeable statistical differences in certain applicant groups (e.g. drop-off, shortlist rates, conversion)?
- **Retest after updates:** after system/model updates or changed data sources, the review logic must be applied again.

Proxy and bias check: 10 test questions for TA

1. Which factors determine ranking/recommendation, and are they clearly job-related?
2. Are there characteristics that can act as proxies (education, companies, region, language patterns)?
3. Are there systematic differences in outcomes between groups (where measurable/meaningful)?
4. Are data gaps (e.g. incomplete profiles) unfairly 'penalized'?
5. Is there overreliance on scores/rankings (only 'Top X' candidates reviewed)?
6. What override rules are in place (when is a deliberate override made)?
7. What happens in the event of complaints or conspicuous cases (review path)?
8. What changes after updates (release notes + retest)?
9. Is training/reference data sufficiently documented/traceable?
10. Were the specialist department/Works Council/DPO involved where necessary?

Key point: If proxy risks cannot be ruled out or controlled, the use case is treated more strictly (red) as a precaution or is not used.

9.5 Practical implementation: how to keep it manageable

- A fairness check frequency is defined for each red use case (e.g. quarterly or after updates).
- Use a few, stable test questions (see box) and document results briefly.
- Use an escalation principle: if conspicuous patterns become visible, the use case, data basis and process rules are reviewed and adjusted where necessary (stop/override/escalation).

Proxy risks as a 'filtering chain'
Even without protected characteristics, proxy effects can lead to disadvantages.

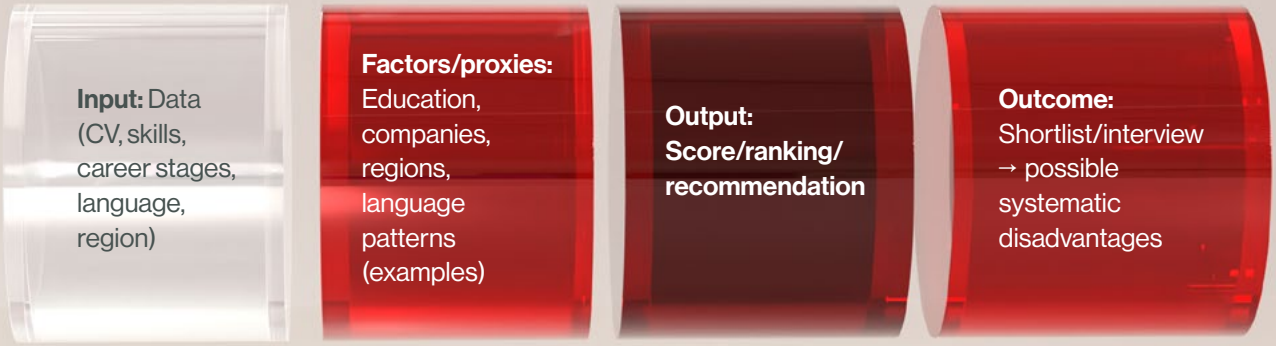

9.6 Reference to the Toolkit

A detailed fairness check template and short documentation form are provided in the Toolkit *(Chapter 14.3* (risk intake) and T-14.6 (monitoring/review protocol).

Monitoring, Updates and Incident Handling (Operational Reliability in TA)

10.1 Monitoring is a deployer obligation, not the end of the project

AI use cases are not 'set and forget': models, configurations and data sources change, which means results can change unnoticed (drift). This chapter provides the operating logic: minimum monitoring, update check and incident path.

10.2 What should typically be monitored in TA

Practical minimum standard for ongoing monitoring:

- **Output stability:** do rankings/recommendations change conspicuously without any technical reason?
- **Outcome signals:** do key figures or patterns shift (e.g., unexplained drop-offs, changed shortlist composition)?
- **Fairness indicators:** are there indications of systematic disadvantage or proxy effects *(Chapter 9)*?
- **Data quality:** are there new data gaps, field changes or interface problems?
- **Process compliance:** is human oversight actually practiced, or is it adopted 'blindly'?
- **Feedback/complaints:** are there increasing numbers of queries, complaints or 'strange cases' from TA or hiring managers?

Key point: Monitoring does not start with big data, but with a few stable signals and a clear review routine.

10.3 Updates and change impact: when reassessment is required

Updates often have a creeping effect – a stable use case can behave differently unnoticed after release or retraining. A new **change impact check** is mandatory if at least one of the following applies:

- **Change in ranking / matching logic:** weightings, features, scoring mechanisms or thresholds are adjusted.
- **Model or prompt updates:** new model, new prompt template, new agent logic or new guardrails.
- **New data sources / data flows:** additional fields, new integrations (ATS/CRM), new taxonomies/ontologies, new preprocessing steps.
- **New countries/languages/job families:** rollout to new markets or language areas, new role clusters, new volume profile.
- **Operating parameters change:** role/rights concept, access paths, degree of automation, new workflow triggers.
- **Inconspicuous signals in operation:** complaints, unusual drop-offs, unexpected shifts in shortlists or conversion rates.

Pragmatic best practice: 'Regression tests' with synthetic profiles

A particularly effective change impact check is a standardised test set with **synthetic profiles.** These artificially generated applicant profiles, also known as 'sibling CVs', are not based on real people, but are generated using AI models to simulate certain patterns, behaviours and qualification profiles. For the checks, profiles are constructed in pairs so that all qualification-relevant characteristics are identical and only one potentially discrimination-relevant characteristic varies (e.g., age, gender, nationality, disability or comparable dimensions). The test logic is simple: **the ranking order or match score must not change systematically as a result of the isolated variation.**

Recommended rhythm:

- Regularly (e.g. every 6–8 weeks), and
- As needed after releases/updates or major changes (release gate)

Documentation (auditable):

Date, system/release status, test set version, results/deviations, decision (OK/measures/stop), responsible person and, if applicable, retest date.

10.4 Incident handling: stop/override/escalation

If a system shows conspicuous results or fairness signals, TA needs a **clear contingency plan,** which is faster than the operational damage. The aim is not to troubleshoot during live operation, but to **limit damage,** activate those responsible and only release the system again after a documented review.

Typical incident triggers (examples):

- Sibling CV tests show unexpected deviations (e.g. systematic rank shifts in one dimension).
- Sudden changes in ranking distributions, shortlist rates or drop-offs without a plausible process change.
- Complaints from applicants/hiring managers about implausibility, lack of transparency or conspicuous patterns.
- Release/update without sufficient change impact check ('unknown change').
- Indications of overreliance ('AI Top 5 is blindly adopted') or lack of override practice.

1 – Stop/freeze (immediate risk limitation)

- Automations and decision-related outputs are paused if selection effect is affected.
- For the interim period: **prioritise human decision-making**; only use AI outputs as an indication or temporarily not at all (depending on the risk).

2 – Override (make human decision-making effective again)

- Recruiters/hiring managers actively intervene: Review according to clear criteria.
- Deviations (overrides) are briefly documented if the use case has decision relevance.

3 – Escalation (clear path, no ad hoc panic)

- Use case owner (TA) is informed; depending on the risk, also data protection (DPO), IT security/architecture and, if applicable, Works Council / AI Board.
- Decide whether the use case remains in 'freeze', whether only parts are switched off, or whether it continues in controlled pilot mode.

4 – Analysis, fix and retest (release only with documentation)

- Identify the cause: update, data source, configuration, prompt/agent logic, process change.
- Implement measures and run a **retest** (incl. regression tests / sibling CV set).
- Document the release decision: What changed? Which results demonstrate stability/fairness? Who released it?

5 – Communication and learning (so it doesn't happen again)

- Inform relevant stakeholders (TA team, hiring managers, HR/SBO if applicable).
- Feed learning back into the monitoring plan and change impact check (e.g. new release gate, extended test set, clear override criteria).

Update check in five questions (change impact)

1. What has changed (model/feature/configuration/data source/ process)?
2. Does the change have a selection effect (ranking/ recommendation/filtering) or is it only text/UX?
3. Do fairness/proxy checks need to be rerun *(Chapter 9)*?
4. Does explainability / output logic need to be reviewed again *(Chapter 8)*?
5. Who releases the change, and where is this documented?

Key point: When selection effect and change come together, the rule is: check, document, release.

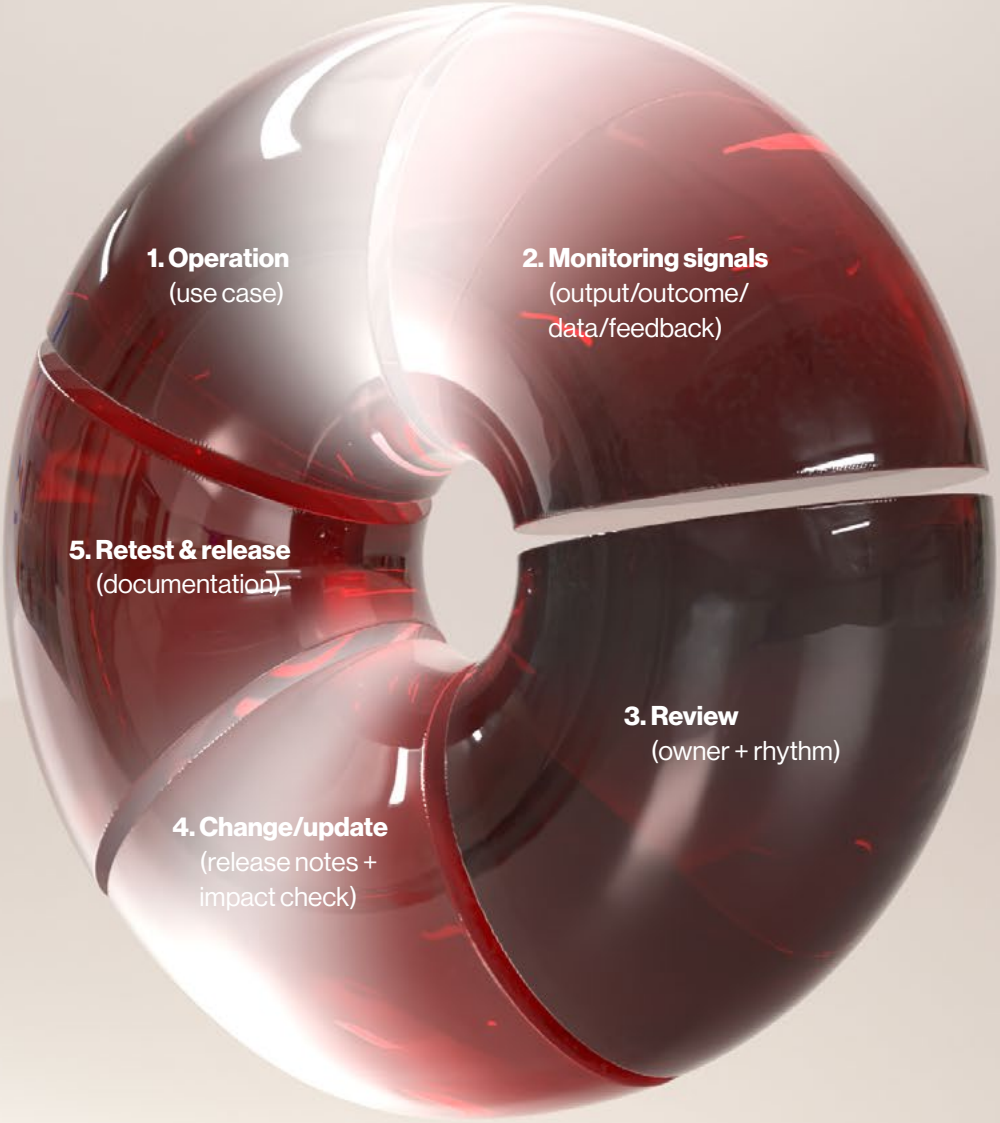

'Incident path: stop/override/escalation'

Monitoring loop – important: immediate interruption in the event of conspicuous signals

The regular monitoring rhythm applies to scheduled reviews. The following also applies: conspicuous signals in operation (e.g., unexpected rejection rates, bias indicators, user complaints) trigger an immediate, unscheduled review – regardless of the next planned cycle. In this case, the monitoring loop is interrupted immediately. Such event-related checks must be recorded in the system documentation.

Best practice: 'Sibling CV' test as an auditable fairness and drift early-warning system

Goal: Ensure that matching/ranking results do not shift unnoticed toward bias due to updates, retraining or data stream changes.

Procedure: A library of paired synthetic profiles ('Sibling CVs') is used. All qualification characteristics are identical within each pair; only one potentially discriminatory characteristic is varied in a targeted way (e.g., age, gender, nationality, sexual orientation, disability).
Test logic: The sibling profiles are repeatedly ranked against the same requirements profiles.

Expectation: The order or match score must not change systematically as a result of the isolated variation.

Rhythm: Regularly (e.g., every 6–8 weeks) and additionally on an event-related basis after releases, model updates or changes to data sources ('change impact check'). Event-related checks also apply to conspicuous signals during operation (e.g., unexpected rejection rates, bias indicators, user complaints) – these trigger an immediate, unplanned review, regardless of the next planned cycle. Such event-related checks must be recorded in the system documentation.

Documentation: Results are versioned and documented (date, release status, test set version, deviations/reasons, release decision).

Benefit: Deployers use this to build their own quick test routine, which shows in an audit that monitoring is not merely 'claimed', but practiced and verifiable, in addition to provider tests.

10.5 Minimum operating standard for AI in TA without overkill

Minimum standard for all yellow/red use cases:

- **Owner and rhythm:** clear use case owner + review cycle (e.g. monthly/quarterly)
- **Release/change logic:** collect release notes, briefly assess change impact
- **Minimum metrics:** define 3–5 signals (output, outcome, data quality, complaints)
- **Retest trigger:** after updates or conspicuous patterns
- **Incident runbook:** stop/override/escalation as standard

Detailed templates (review protocol, runbook, change check) belong in the Toolkit.

10.6 Reference to the Toolkit

Review protocol, change impact check, incident runbook: Toolkit *(Chapter 14.3* monitoring/review/change impact and (T-14.7) incident runbook)

Data Protection, Data Classes and LLM Input Rules in Talent Acquisition

11.1 Data protection is always relevant – regardless of the AI Act

Regardless of the AI Act classification, personal data is almost always processed in TA. Real risks often arise not from AI itself, but from unclear data flows, misuse or uncontrolled LLM inputs – data classes and clear prompt rules keep use cases auditable (see T-14.9).

11.2 Data classes in the TA context (practical working logic)

The following classification serves as a pragmatic guiding principle: 'What may go where and what must not be used under any circumstances?'

- **Data class A: Public non-critical (job-related, not personal data)**
  Examples: Job profiles, job families, task descriptions, required skills, anonymised sample texts.
  Rule: Freely usable for text assistance and content generation.
- **Data class B: Personal data (standard TA data)**
  Examples: CVs, contact details, communication histories, interview notes, assessments.
  Rule: Strict purpose limitation, clear role rights, and fixed deletion periods.
  Please note: This data must never be entered into freely available LLM services without enterprise protection.

- **Data class C: Sensitive / special categories of personal data**
  Examples: Health data, religion, ethnic origin, trade union membership.
  Rule: Avoid as a general rule. If procedurally unavoidable, only with an explicit legal basis and maximally restrictive controls.
- **Data class D: Biometric data**
  Examples: Video/audio from interviews, voice/face characteristics.
  Rule: The purpose is decisive. While pure identity checks may be possible, analysing behaviour or emotions is a strict no-go in the TA context.
  Note: Biometric data also falls under the special categories pursuant to Article 9 of the GDPR and is therefore subject to double protection obligations.

11.3 Prompt and input hygiene (GenAI/LLM): the most important practical rule

Significant data protection risks arise when recruiters copy entire CVs into unprotected prompts due to lack of time. A minimum standard prevents the most common errors:

1. Defined **input rules** exist for each use case (which data is permitted).
2. Job-related **data (class A)** is used primarily for text generation.
3. For decision-related red use cases, only **controllable,** closed enterprise systems are used.

Important guiding principle (LLM + applicant data/CV):

Freely available or general-purpose LLM services should **not** be operated with individual applicant data (CV/profile), neither for summarisation nor for evaluation purposes. If an organisation wants to process personal TA data using AI, it should use **provided, controllable systems** and treat the use case at least as **red** *(Chapter 5, Chapters 8-13)*, including documentation, approval and documented human oversight.

11.4 Article 22 of the GDPR (automated decision-making): a practical view

Article 22 of the GDPR prohibits decisions *made exclusively* by automated means that have legal effects (e.g. automatic rejection without human review).

Many TA processes are formally 'human-led', but become a legal grey area if AI proposals are in fact blindly accepted (overreliance).

Practical rule: As soon as AI issues rankings, filters or recommendations (red path), human oversight must be genuine and documented.

11.5 DPIA / data protection impact assessment (when it typically becomes relevant)

A DPIA is often required in Talent Acquisition when:

- Applicant data is processed systematically and extensively.
- Profiling, scoring or ranking with selection effect takes place.
- New, data-intensive technologies (e.g. AI matching) are introduced.
- There is an increased likelihood of complaints/queries from affected individuals due to the process effect.

For red use cases, a DPIA is often a realistic standard; for yellow use cases, it depends heavily on the scope, data and process design.

Input rules for GenAI/LLM: What is permitted and what is not?

Permitted (typical):

- Job-/role-related information (tasks, must-have/nice-to-have criteria, skills)
- Anonymised sample texts, process modules, template texts
- Aggregated, non-personal information

Only in controllable systems designed for this purpose (typically red):

- Personal TA data (CV/profile, communication, interview notes) only with purpose limitation, role rights, appropriate logging and deletion concept; not in freely available LLM services.

Do not use / no-go:

- Special categories of personal data
- Biometric analyses of behaviour/emotions
- Individual applicant profiles/CVs in freely available LLMs if purpose limitation/verifiability cannot be reliably ensured

Key point: If you cannot clearly explain, limit and delete it, it does not belong in the prompt.

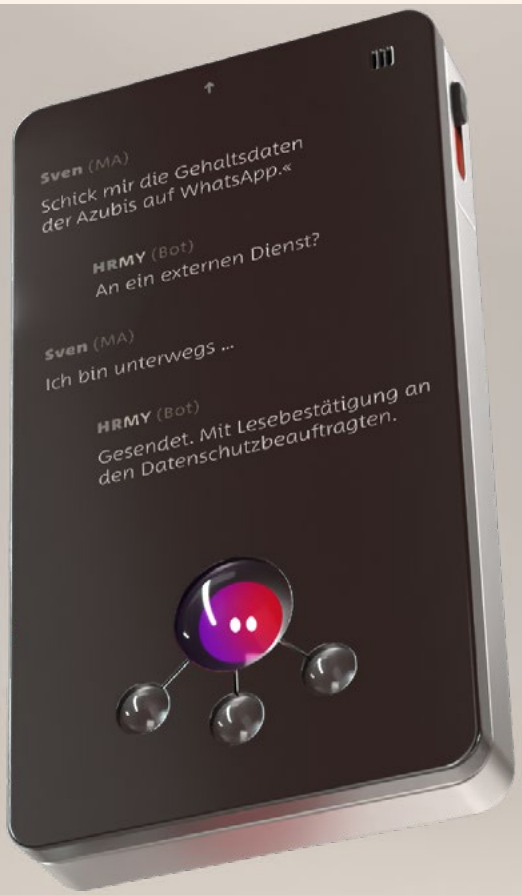

Data class matrix

Data class	Permitted	Controlled only	No-go
A – Public / job-related	OK	–	–
B – Personal (standard TA)	–	OK	–
C – Sensitive / special categories (Article 9)	–	–	No-go
D – Biometric	–	Purpose determines use	No-go: emotion / behaviour analysis

Complete data class checklist in the Toolkit (T-14.9). Sample texts: Toolkit *(Chapter 14.3* Data classes and prompt/input policy)

11.6 Reference to the Toolkit

Prompt/input policy, data class checklist, DPIA trigger list, and sample texts: Toolkit *(Chapter 14.3* Data classes and prompt/input policy)

Co-Determination (BetrVG), Change Management and Works-Council-Ready Implementation in Talent Acquisition

12.1 Co-determination is not an add-on – it determines feasibility

In Germany, the use of AI tools in HR/TA is often subject to co-determination, especially if systems can change selection processes or record employee behaviour. This chapter shows how to prepare for this so that co-determination does not become a showstopper.

12.2 When the Works Council typically must be involved

Co-determination does not depend on the tool label, but on its effect in the process. Typical triggers are:

- **Monitoring/control reference: the system can monitor or evaluate employee behaviour or performance (directly or indirectly).**
  Note: System-side logging of user actions (audit trail) serves the obligation to provide documentation and is considered accepted practice; it only falls under monitoring/control reference if the logged data is used to assess individual behaviour.

- **Changes to selection rules:** new scoring/ranking logic, automated filters, standardised decision logic.
- **Introduction of new workflows:** process changes, new mandatory steps, new data flows, new role rights.
- **Use in an employee context:** when tools not only affect applicants, but also employees (e.g. internal mobility, performance-related applications).
- **Biometrics/video/audio:** particularly sensitive depending on purpose and processing.

In addition, even before AI-supported processes are introduced, there may be an early duty to inform and consult under Section 90(1) No. 3 of the BetrVG.

Key point: Co-determination is rarely made easier by 'saying less'. It is made easier by clean documentation, clear boundaries and explainability.

12.3 Works-council-ready package: which documents TA should provide

For co-determination to be efficient, a standardised package is needed that TAs and project teams can reuse. The package is not a legal opinion, but a structured presentation ready for decision-making.
Works-council-ready package (minimum):

1. **Use case one-pager** (purpose, process step, output, owner, explain the how and why / no black box)
2. **Data overview** (data classes, data flows, access concept, deletion logic)
3. **Process representation** (where the system acts, where the human decides)
4. **Human oversight and control limits** (overrule, review criteria, appropriate logging)
5. **Risk and measures overview** (traffic light, triggers, controls)
6. **Operating concept** (updates/monitoring/incident path)
7. **Communication and training** (who is informed and trained, and how)

Details (templates and list of documents) can be found in the Toolkit *Chapter 14*

12.4 Change management: internal safeguarding

Minimum standard for TA:

- Clear **usage rules** ('Do/Don't', permitted inputs, prohibited purposes)
- Role rights and responsibilities (who may configure what?)
- Training for recruiters / hiring managers (in particular avoiding overreliance)
- Communication logic: What are applicants told? What are employees told?
- Escalation path if conspicuous issues arise *(Chapter 10)*

12.5 Practical guiding principle: control limits instead of endless discussion

Three guiding principles that are consistently communicated:

- TA does **not use AI** to make final autonomous decisions about people**.**
- As soon as AI **supports** decisions (ranking/recommendation/filtering), stricter controls apply: explainable factors, documented review logic, monitoring.
- Certain purposes (e.g. emotion recognition in the workplace) are excluded *(Chapter 11)*.

This makes co-determination manageable: not everything is prohibited, but the boundaries are clear.

Participation process
'Get involved early, define boundaries, secure operations'

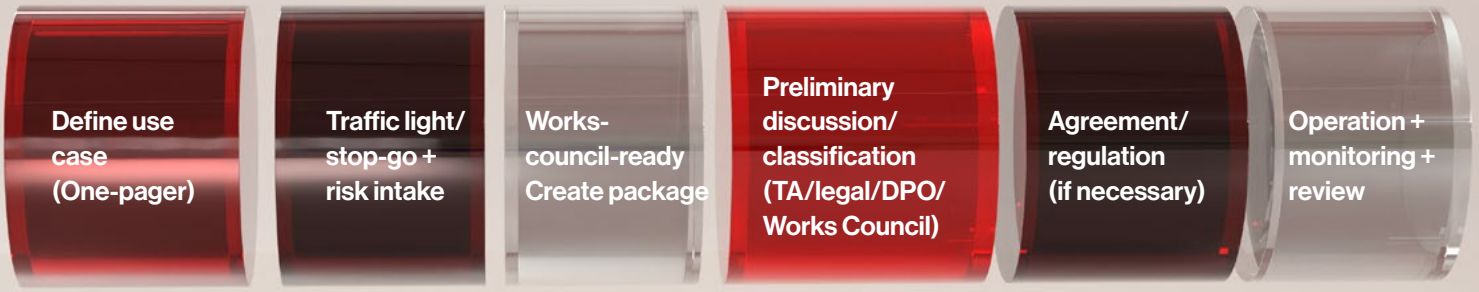

Works Council trigger list (quick check)

A works-council-ready process is typically required if at least one of the following points applies:

- Selection decisions are prepared by AI (ranking/recommendation/filtering)
- New/expanded data flows or new data categories are introduced (including video/audio)
- There is a monitoring/performance reference in the employee context
- New mandatory process steps or automated process steps with selection effect are introduced
- Significant changes are made to the use case (updates, new countries / languages, new data sources)

Key point: If a trigger applies: do not discuss 'whether', but prepare in a standardised way.

12.6 Reference to the Toolkit

Works-council-ready package (list of documents), Do/Don't, training modules, sample structure for works agreement: Toolkit *(Chapter 14.3* Works-council-ready package and (T-14.12) Do/Don't and training modules)

HR/TA Governance and Competence Building (Article 4 of the AI Act) as an Operating Model

13.1 Governance makes TA capable of action – not slower

Governance makes TA faster, not slower – provided roles, competence requirements (Article 4) and interfaces with the AI Board, data protection and co-determination are clearly regulated.

13.2 Operating model: roles, interfaces, responsibilities

Practical minimum model (also applies without AI board):

- **Use case owner (TA):** professional responsibility, scope, success criteria, operational view.
- **Risk/compliance lead (HR/TA or central):** classification (traffic light/stop-go), coordination risk intake, documentation.
- **Data protection (DPO/privacy):** data classes, purpose limitation, DPIA triggers, deletion concepts, data subject rights.
- **IT security/architecture:** access, interfaces, appropriate logging, hosting, security documentation.
- **Co-determination / Works Council interface:** trigger check, document package, participation process.

- **AI Board / steering (if available):** approval of red use cases, escalations, standards.

Key point: TA controls use cases; central committees provide standards and approvals, but HR/TA must deliver 'board-ready'.

13.3 RACI (brief) – who does what?

A simple RACI logic helps avoid coordination loops:

- **R (responsible):** TA use case owner
- **A (accountable):** HR/TA leadership or AI Board, depending on risk
- **C (consulted):** data protection, IT security, Works Council (trigger-dependent), legal (especially labour law and data protection law)
- **I (informed):** hiring manager, HRBP, operations

This logic becomes practical when linked to templates (use case one-pager, risk intake, decision pack).

13.4 Competence building under Article 4 of the AI Act (practical)

Role-based means specifically:

- **Recruiter & hiring manager:** correctly classify AI outputs, avoid overreliance, apply review criteria, Do/Don't in everyday practice (e.g., no applicant data in open LLMs).
- **TA leads / use case owner:** apply traffic light logic / stop-go, ensure documentation logic, manage stakeholders, establish monitoring routine.
- **Privacy/IT security/Works Council interfaces:** data classes and data flows, appropriate logging requirements, change impact checks, limits/no-goes and co-determination triggers.
- **AI Board / steering (if available):** Decision logic, minimum proofs, approval/escalation path (stop/override/escalation).

Minimum standard: Short, role-based training modules plus recurring refreshers (e.g. annually and additionally for significant updates, new use cases or process changes).

13.5 Governance in operations: standards, reviews, escalations

Governance does not end with approval. Yellow/red use cases require:

- **Standards:** data class rules, prompt policy, minimum documentation requirements in procurement, documentation obligations.
- **Reviews:** rhythm (e.g. quarterly) plus reassessments after updates *(Chapter 10)*.
- **Escalations:** incident path (stop/override/escalation) and clear owner.

This is how governance becomes an operating system: build it properly once, then reuse it.

‘Board-ready’ in one sentence

A use case is ‘board-ready’ if the purpose, data, output, risk (traffic light), controls (oversight/monitoring/fairness/explainability) and responsibilities can be understood on a single page (decision pack).

Operating model

‘TA provides documentation, central functions review/approve, operations are monitored’.

Reference to the Toolkit

RACI template, training modules, decision pack, governance checklists: Toolkit (*Chapter 14.4* decision pack, T-14.8 RACI/role description and T-14.12 training modules)

13.6 Special case: governance without an AI Board or legal department

Not every company has a central AI Board, a dedicated DPO or a legal department that handles AI Act issues independently. Nevertheless, the requirements still apply. The following minimum model shows how governance can also work in small and medium-sized organisations without a large organisational structure.

Lean AI governance: minimum model for organisations without an AI Board – who takes on which role?

- Use case owner + risk assessment: HR management or TA lead (combined with managing director approval for red use cases).
- Data protection: involve the external DPO (common in SMEs) using a standardised 'AI use case package' (= use case one-pager + risk intake) instead of lengthy individual inquiries.
- IT security: IT manager or external service provider for interface checks and access concepts. In addition: ensuring automatic deletion rules for AI-processed candidate data after defined retention periods (see GDPR GDPR Article 5(1)(e)).
- Legal review: as needed (e.g. Works Council topics or red use cases) – not permanently. A one-off 'AI policy check' by an external law firm can save later rework.

What can be simplified?

- Build an RACI matrix with all six roles – the RACI matrix remains entirely role-based (use case owner, HR manager / TA lead)

- Management, data protection, IT, legal), even if several functions are combined in one person. This ensures that no responsibilities are overlooked; in a second step, it can be specified which individuals perform the functions in order to streamline coordination.
- Instead of a two-page decision pack: a documented approval email with the one-pager attached is sufficient documentation for yellow use cases.
- Monitoring without tooling: quarterly review meeting (30 minutes) between the HR lead and the use case owner – documented – is sufficient as a minimum standard.

Key point: Even without an AI Board: inventory + one-pager + risk intake + documented approval = auditable. This can also be implemented without large structures.

13.7 Market structure: provider categories and their typical risk profiles

The provider market for AI in TA is heterogeneous. The same compliance risk can be distributed very differently depending on the provider category. This overview is not a product recommendation; it shows where particular governance questions typically arise.

Category 1: ATS with integrated matching/ranking (e.g. large HCM suites)
Typical risk profile: AI functions are often deeply integrated and difficult to disable. Matching/ranking features sometimes run 'invisibly'. Typical challenge: Explainability and fairness documentation for individual submodules can be difficult to isolate. Recommendation: Explicitly ask in the RFP which functions are classified as AI systems under the AI Act and request documentation separately.

Category 2: Specialised screening/matching tools (standalone)

Typical risk profile: Close to high risk (red path).
Advantage: often greater transparency into the AI core function than with integrated suites. Documentation for explainability, fairness and monitoring is often easier to request and assign clearly.
Recommendation: Apply the RFP questions from (T-14.5) in full.

Category 3: LLM tools / GenAI services (standalone)

Typical risk profile: Yellow or red depending on the function.
The greatest risk: uncontrolled data input (CV/profile) into systems that are not enterprise-ready. Additional role/obligation shift risk if the deployer configures the LLM for selection purposes *(Chapter 7)*.
Recommendation: Clearly limit the purpose, define data class rules in advance (T-14.9).

Category 4: Sourcing platforms and talent intelligence tools

Typical risk profile: Often underestimated. As soon as profiles are prioritised or ranked (including 'Top Matches' as a feature), the red path becomes relevant.
In addition: Data origin and legitimacy (were the profiles collected lawfully?).
Recommendation: Explicitly address database documentation and anti-bias testing in the RFP.

Key point: The category alone does determine the risk. The function and its use in the process are decisive. However, the category shows where the documentation gaps typically arise.

Toolkit (Templates and Policies): How to Accelerate Implementation

14.1 Purpose of this Toolkit

This white paper is designed as a playbook. Implementation becomes much easier when teams do not start from scratch each time, but instead use standardised working documents. The Toolkit therefore bundles the most important **templates, checklists and policy modules** referenced in Chapters 4–13.

In the white paper itself, the templates are only introduced in a structured way. The Toolkit consists of **14 separate files** (T-14.0 to T-14.13: 13 Word documents + 1 Excel file), which are designed as independent working documents and not part of this white paper. The complete templates are available via the Google Drive link below and will be published on QUEB.org once finalised.

All documents (Toolkit/files) can be viewed here as a complete folder:

https://www.queb.org/eu-ai-act-en

14.2 How to use the Toolkit (in three steps)

1. Describe the use case: fill out the one-pager (T-14.2) – purpose, process step, data, output and owner.
2. Classify the risk: complete the risk intake (T-14.3) – traffic light, triggers and interfaces to GDPR/AGG/BetrVG.
3. Secure approval and operations: create decision pack (T-14.4), request the necessary documentation and define monitoring/incident path (e.g. (T-14.6) / (T-14.7)).

Key point: When one-pager + risk intake + decision pack are in place, coordination with stakeholders becomes much faster and more professional.

14.3 Toolkit overview (contents)

The Toolkit includes at least the following components:

A) Inventory and use case documentation

- (*T-14.1*) AI inventory / AI register (light)
- (T-14.2) Use case one-pager (TA)

B) Risk and approval

- (T-14.3) Risk intake (AI Act/GDPR/AGG/BetrVG)
- (T-14.4) Decision pack (board-ready, 1 page)

C) Procurement and supplier quality

- (T-14.5) RFP questions and minimum documentation + scoring logic

D) Operations and governance

- (T-14.6) Monitoring/review protocol + change impact check
- (T-14.7) Incident runbook (stop/override/escalation)
- (T-14.8) RACI (brief) + role description

E) Policies and guiding principles

- T-14.9 Data classes and prompt/input policy (TA)
- T-14.10 Transparency / communication modules (affected individuals, hiring managers)
- T-14.11 Works-council-ready package (document list + structure)
- T-14.12 Do/Don't and training modules (role-based, Article 4 of the AI Act)

14.4 Version and maintenance principle

To ensure that the Toolkit remains effective, it is maintained as a living document:

- Changes are versioned (change log)
- Templates are updated based on new findings or practical examples
- Relevant regulatory changes trigger an update

Three templates you can start with immediately

If you implement only three things immediately, they should be:

- T-14.1 **AI inventory (light)**
- T-14.2 **Use-case one-pager**
- T-14.3 **Risk intake**

Key point: These three building blocks provide structure, speed and auditability – without a major project.

→ What the Works Council specifically needs (one-pager, data overview, process presentation, human oversight, risk/controls, operating concept, training plan) is listed in full in *Chapter 12.3*.

Toolkit as a construction kit (layout/graphics briefing)

Different modules are activated for each traffic light path.

Practical Examples from Member Companies

Practice shows most clearly how regulatory requirements can be implemented in everyday practice. This white paper therefore supplements the conceptual framework with concrete case studies from member companies of the QUEB Federal Association – authentic, practical and structured according to the 'case nugget' framework developed in this white paper.

Publishing logic: modular and continuous

The practical examples are not published as an integral part of this white paper, but as separate documents that will be published successively on ***Queb.org***. The reason is pragmatic: internal company approval processes take time, and the quality of the cases takes precedence over a rigid publication date. Each case document stands on its own, references the white paper as the conceptual framework and is available directly via QUEB.org.

The first practical examples will be published alongside this white paper **when it is released in May 2026**. Additional cases will follow as soon as they are approved.

15.1 Case nugget framework

All practical examples follow a standardised short format (approx. ½ page) that ensures quick orientation and maximum transferability. Use case, TA process step, classification (traffic light), key decision point, what helped and the key lesson learned. In addition, selected lighthouse cases are presented in a more detailed format (1–2 pages) with deeper insight into implementation, governance and auditability.

15.2 Available and planned cases

The first cases come from companies that are already actively shaping the use of AI in Talent Acquisition and are willing to share their experiences. Additional companies are currently in the internal approval process; their contributions will be published once approval has been granted.

→ All currently available case documents can be accessed at ***https://www.queb.org/eu-ai-act-en***. The page is updated continuously.

15.3 Shared learning

As soon as a sufficient number of cases is available, overarching patterns, typical stumbling blocks and proven quick wins will be compiled into a summary report. This is also published as a separate document on Queb.org.

15.4 Case library as a living format

The goal is a growing, curated case library on QUEB.org – not a static chapter, but a continuously expanding knowledge pool that reflects the practice of AI use in Talent Acquisition beyond the initial publication. Companies that would like to contribute their own case can contact the Queb Federal Association. directly.

Conclusion, Outlook and Next Steps for HR/Talent Acquisition

16.1 Conclusion: AI in TA is feasible when managed as a process

The AI Act is not a brake on innovation! It ensures that AI in recruiting is conceived and operated not as a feature, but as an **operational application with clear accountability**. What matters is never the tool label ('GenAI' or 'agent'), **but its impact on the process:** Does the system only provide information, or does it support selection decisions?

If you apply the playbook logic consistently, you achieve three effects: **peed** (fewer loops through standards and templates), **trust** (greater acceptance among business, data protection and co-determination) and **resilience** (stable control during updates/drift and in audits).

Typical stumbling blocks in practice and how to avoid them

Stumbling block 1: 'We'll pilot first – compliance can come later'.

In practice, a pilot almost always becomes an operational process before the required documentation is added. Whoever documents later, documents incompletely.
Recommendation: One-pager and risk intake from day one, including during the pilot.

Stumbling block 2: The Works Council finds out from the business – not from HR

When hiring managers start using an AI tool before the Works Council has been involved, pressure to justify it arises later. This is difficult to reverse.
Recommendation: Trigger check *(Chapter 12)* before each launch.

Stumbling block 3: The provider is assumed to handle compliance – but does not

Deployer obligations (purpose definition, human oversight, monitoring in operation) always remain with the deployer, regardless of what the provider promises.
Recommendation: Define the role matrix from *Chapter 6* from the outset.

Stumbling block 4: An update is released and no one reviews it

Provider updates can change ranking logic without TA noticing. This is the most common entry point for unnoticed drifts.
Recommendation: Subscribe to release notes, change impact check as standard *(Chapter 10)*.

Stumbling block 5: Recruiters trust the score and stop reviewing candidates themselves

Overreliance is the most difficult risk to measure. It develops gradually and can hardly be explained in the audit.
Recommendation: Regularly monitor the review criteria and override rate; training not one-off, but recurring (T-14.12).

Stumbling block 6: 'We only use GenAI for text generation', but applicant data still ends up in the prompt

The time pressure of everyday life leads to CVs being copied into unprotected LLMs, even though the policy prohibits it.
Recommendation: Don't just document data class rules; translate them into training materials and do/don't guidance ((T-14.9) / (T-14.12)). Even approved tools create overuse risk if clear scope limits are not communicated.

These stumbling blocks are based on the experience of the authors and co-editors of this white paper.

16.2 Outlook: three trends that will shape TA in 2026/2027

1. **Agent workflows are on the rise:** AI will increasingly control processes autonomously. The decisive factor is whether this creates a selection effect; if so, the red risk category is mandatory.
2. **LLMs become part of everyday work:** the greatest risk remains the use of personal applicant data. ***Entering a CV into a freely available LLM*** remains the typical compliance trap.
3. **Governance becomes routine:** monitoring, change impact checks and incident paths become standard operational practice, much like IT security or financial controls.

16.3 Next steps: the 30-day plan for TA (minimum programme without major project)

This plan is designed so that TA can start immediately, regardless of the tool landscape.

Week 1 (transparency):

Create AI inventory (light) and record all current TA use cases (T-14.1).
Create a one-pager for the top three use cases (T-14.2).

Week 2 (classification):

Conduct a risk intake for the top-use cases (T-14.3).
Define the traffic light classification and review red triggers *(Chapter 5)*.

Week 3 (governance and documentation):

Create decision pack for red use cases (T-14.4).
Embed minimum evidence requirements in procurement and RFP processes (T-14.5).
Finalise data class and prompt rules *(Chapter 11* / (T-14.9)).

Week 4 (operations):

Define a monitoring/review cycle (T-14.6).
Establish an incident runbook (T-14.7).
Conduct short role-based training (recruiters/hiring managers/TA leads) (T-14.12)

.

Key point: Once **inventory + one-pager + risk intake** are in place, every subsequent decision becomes much faster.

16.4 Using this white paper as a communication tool (stakeholder-ready)

- **For management/stakeholders:** *Chapter 0* (Foreword), *Chapter 1*, *Chapter 5* (traffic light) and *Chapter 16* are usually sufficient as a basis for decision-making.
- **For TA teams:** *Chapters 4–7* as a foundation, then *Chapters 8–13 depending on the use case.*
- **For data protection, the Works Council and IT security:** *Chapter 6* and *Chapters 10–12* as a shared working basis.

List of abbreviations and glossary (as of 30 March 2026)

AGG: General Equal Treatment Act (protection against discrimination in employment)

AI: Artificial intelligence (systems that learn from inputs to generate predictions, recommendations or decisions)

AI Act: European AI Act (regulatory framework for the use of AI)

AI Board / steering: Governance body responsible for the approval and oversight of critical AI use cases (organisation-dependent)

AI Omnibus: AI Omnibus: European simplification and adaptation package for the AI Act intended to reduce administrative requirements, harmonise implementation and make individual requirements more practical.

ATS: Applicant Tracking System (applicant management system)

BDSG: Federal Data Protection Act

BetrVG: Works Constitution Act

CE: CE marking (conformity marking; relevant for high-risk AI)

CRM: Candidate Relationship Management (talent pool / relationship management)

CV: Curriculum vitae

DPO: Data Protection Officer

DPIA: Data protection impact assessment (review process under the GDPR for high-risk processing)

GDPR: General Data Protection Regulation

EU-DB (HRAI): EU database for high-risk AI systems (registration requirement, where applicable)

FAQ: Frequently asked questions

GenAI: Generative AI (systems that generate content such as text/ images/code)

GPAI: General-purpose AI (e.g. foundation models)

HCM: Human Capital Management

HM: Hiring manager

HR: Human Resources

HRBP: HR business partner

IT-Sec: IT security / information security

LLM: Large language model (e.g. GPT-based models)

No-go: Inadmissible or prohibited purpose/application in the HR context (e.g. emotion recognition in the workplace)

One-pager: Brief description of a use case (objective, process step, data, output, owner)

OECD: Organisation for Economic Co-operation and Development; international reference framework that defines AI as systems that generate predictions, recommendations or decisions based on inputs.

Playbook: Practical implementation guide for HR, Talent Acquisition and recruiting stakeholders to classify AI Act requirements and derive concrete actions for their own organisation.

PII: Personally Identifiable Information (identifying personal data, e.g., name, email address, applicant ID)

RACI: Responsible/Accountable/Consulted/Informed (role and responsibility matrix)

RFP: Request for Proposal (tender/procurement process)

Risk intake: Standardised risk classification (AI Act/GDPR/AGG/ BetrVG)

SLA: Service Level Agreement (agreed response/processing times)

TA: Talent Acquisition (recruitment)

Use case: Specific application within a process (e.g. screening/ matching, candidate bot)

About the Authors

Anna Ehret — author
BMW Group
Anna.Ehret@bmwgroup.com

Anna Ehret is responsible for Recruiting Methods, Systems and Processes at the BMW Group in a shared leadership role. Her focus is on the further development of modern recruiting processes, selection methods and system logic in an international talent acquisition context.
Previously, she held senior HR positions at Audi, including personnel diagnostics, corporate culture and employee surveys. Her experience combines strong diagnostic expertise with the practical design of scalable, governance-compliant recruiting processes.

Bernd Schmitz — author
Honorary Advisory Board, QUEB Federal Association
schmitzbernd@gmail.com

Bernd Schmitz has spent much of his corporate career (including at Bayer) leading transformation initiatives in Talent Acquisition, including governance implementation, data protection and co-determination interfaces, as well as early discussions around the EU AI Act. Today, within the QUEB context, he drives the practical dialogue on how HR innovations can be implemented sustainably and at scale under regulatory conditions.

Helge Kroll — co-editor
Statkraft AS
helge.kroll@statkraft.com

Helge Kroll has held positions at thyssenkrupp, Covestro, PwC and currently Statkraft, with responsibilities spanning employer branding, HR marketing, recruiting and talent acquisition in both national and international contexts. Throughout this period, he has supported the QUEB Federal Association in a variety of roles alongside his professional career.

Jan Schüttler — co-editor
TÜV Rheinland
jan.schuettler@gmail.com

Jan Schüttler has held various roles at Bayer and TÜV Rheinland with responsibility for talent acquisition (employer branding, HR marketing, recruiting, onboarding) and is a board member and working group lead for HR Tech at the QUEB Federal Association.

About the Authors

Louisa Bock — co-editor
Bayer AG
louisa.bock@bayer.com

Louisa Bock is strategically responsible for talent marketing and employer branding at Bayer AG, with the goal of positioning Bayer as an attractive employer among the target groups that are critical to the company's long-term success. In her role, she drives the targeted recruitment of talent with the skills Bayer needs to sustainably deliver its mission, 'Health for all, Hunger for none'.
She is part of the Future Skills Cluster, which focuses on developing and acquiring future-relevant skills at Bayer. A central focus of the cluster is the use of artificial intelligence, complemented by other future technologies such as AR/VR, to strategically develop recruiting, qualification and employee experience in line with the needs of tomorrow's workforce.

Marius Gelhausen — co-editor
Lidl International
marius.gelhausen@lidl.com

Marius Gelhausen is responsible for global Talent Acquisition and Onboarding and works on scalable standards for recruiting and onboarding processes in an international environment. This typically includes the harmonisation of selection and onboarding processes, process excellence, governance topics, as well as the establishment of consistent quality and service standards across countries.

Nick Fleischer — co-editor
Infineon Technologies AG
nick.fleischer@infineon.com

Nick Fleischer is globally responsible for professional and executive recruiting at Infineon Technologies AG. In this role, he designs and manages strategic recruiting approaches across regions. A central focus of his work is the harmonisation and continuous improvement of digital recruiting processes.
He is responsible for the selection, implementation and continuous optimisation of modern recruiting systems, particularly those with AI-supported components, to sustainably strengthen the efficiency, quality and future readiness of talent acquisition.

About the Authors

Nico Wellhöfer — co-editor
BASF Services Europe GmbH
nico.wellhoefer@basf.com

Nico Wellhöfer has worked in talent acquisition for more than 18 years and plays a key role in shaping BASF's global recruiting systems landscape. His focus is on the digitalisation of recruiting processes, SAP SuccessFactors Recruiting Management and related add-ons, as well as the responsible use of AI components. Working closely with recruiting teams across EMEA, APAC and the Americas, he is responsible for the global product management of SAP SuccessFactors Recruiting Management, release management, and the design and implementation of technical concepts based on business requirements. A key focus of his work is validating the business benefits of AI functions, assessing their practical applicability in the corporate environment and coordinating with stakeholders such as data protection, HR and IT. As data protection coordinator, he ensures GDPR compliance and supports the implementation of AI and system solutions from evaluation through rollout.

Philipp Schuh — co-editor
BIOTRONIK AG Switzerland
philipp.schuh@biotronik.com

Philipp Schuh is responsible for global talent acquisition and employer branding at BIOTRONIK. His focus is on international talent strategy, modern recruiting transformation, active sourcing, employer branding and the development of global talent acquisition communities in the MedTech sector. He is deeply involved in the strategic development of global recruiting organisations and the practical application of modern HR and AI technologies in an international talent environment.

Sandra Anna Schmitz — co-editor
Coca-Cola Europacific Partners
sandraanna.schmitz@ccep.com

Sandra Anna Schmitz works in talent acquisition at Coca-Cola Europacific Partners and combines operational recruiting excellence with extensive experience in talent technologies and AI-supported systems. Her focus is on practical selection and screening standards, their further development, modern candidate journeys and the introduction, training and continuous improvement of modern TA processes and systems. She supports rollouts and change initiatives in close coordination with the business, data protection, co-determination and other relevant stakeholders.

About the Authors

Stefanie Wild — co-editor
Evonik Industries AG
stefanie.wild@evonik.com

Stefanie Wild is responsible for the strategic management and further development of talent acquisition at Evonik Industries AG. She leads the global implementation and use of modern recruiting technologies in line with regulatory requirements, such as the GDPR and the AI Act, and promotes practical dialogue among HR, data protection, IT, legal and employee representatives. She is also actively involved in future-oriented talent strategies, skills-based approaches and the responsible use of AI in recruiting. She contributes her international experience from global transformation and governance projects to industry and professional dialogue.

Wolfgang Brickwedde — co-editor
ICR
wb@competitiverecruiting.de

Wolfgang Brickwedde has been head of the Institute for Competitive Recruiting since 2010. The ICR supports and advises organisations in improving recruiting outcomes and processes. The ICR provides a platform dedicated to improving recruiting in Germany through studies on the state of recruiting and benchmarks on the use of and satisfaction with recruiting solutions such as job boards and applicant tracking systems. Until the end of 2009, Wolfgang Brickwedde was responsible for recruiting and operational employer branding for the EMEA region at SAP. Before joining SAP, Wolfgang Brickwedde held various management roles at Royal Philips Electronics in employer branding, recruiting and management development across several countries.

Queb | Federal Association

The association brings together around **70 renowned companies with more than 7 million employees across all industries.** At its core is the ongoing exchange among more than 1,000 HR experts and leaders – practical, collaborative and driven by a shared commitment to advancing talent acquisition sustainably. Working groups, benchmarks and event formats generate insights, standards and solutions for current challenges in **talent attraction, integration and retention.**

Queb stands for quality employer branding and a strong **community that shares knowledge and shapes the future.**

Queb BUNDESVERBAND »TALENT ACQUISITION«

Postfach 41, 61131 Schöneck
Telephone 06187 902066
info@queb.org
www.queb.org

Blog & podcast: https://www.queb.org/blog
LinkedIn: https://www.linkedin.com/company/queb-bundesverband