educage.

SCOPE	GDPR · UK GDPR · EU AI Act
FORMAT	Approval checklist, 3 pages

GOVERNANCE CHECKLIST

AI Agent Approval Checklist

25 governance questions every DPO should ask before approving an AI agent under GDPR, UK GDPR and the EU AI Act.

A practical governance checklist for organisations deploying AI agents.

WHY THIS MATTERS

Modern AI agents don't simply answer questions.
They access systems, act, and decide.

They can access business systems, process personal data, execute workflows, call APIs, browse the internet and make autonomous decisions. Governance decisions cannot rely on technical testing alone - before deployment, ask whether the agent is appropriately governed, monitored and controlled.

- Business systems
- Personal data
- Workflows & APIs
- Autonomous decisions

Eight Governance Domains

Work through every domain before sign-off

1 PURPOSE

- Clearly defined business purpose?
- Is AI genuinely necessary?
- Purpose documented?

2 PERSONAL DATA

- What personal data is processed?
- Any special category data?
- Is data minimisation applied?

3 ACCESS

- Which systems can it access?
- Can it write or modify data?
- Can it execute code?
- Can it browse the internet?
- Can it call third-party APIs?

4 SECURITY

- Could it operate outside scope?
- Least-privilege permissions?
- Emergency kill switch?
- Independently logged activity?

5 HUMAN OVERSIGHT

- Meaningful oversight defined?
- Can critical actions be reviewed?
- Can autonomous acts be stopped?

6 RISK MANAGEMENT

- Has the DPIA been reviewed?
- Foreseeable misuse documented?
- AI-specific risks assessed?

7 INCIDENT RESPONSE

- How is unauthorised behaviour detected?
- Would a breach be identified quickly?
- Is the IR plan AI-ready?

8 ACCOUNTABILITY

- Who owns the AI agent?
- Who approves deployment?
- Who reviews the logs?
- Who is responsible if it exceeds scope?

Trust is not an AI governance control.
Independent verification is.

educage.

SCOPE	GDPR · UK GDPR · EU AI Act
FORMAT	Self-assessment, 25 items

EXECUTIVE SELF-ASSESSMENT

Executive Self-Assessment

Twenty-five questions to evaluate AI governance readiness - not technical implementation. Answer at board or leadership level before approving deployment.

01	Is there a clearly defined, documented business purpose?	14	Are its activities independently logged?
02	Has it been confirmed that AI is genuinely necessary?	15	Is meaningful human oversight defined for this agent?
03	Has the personal data it will process been identified?	16	Can critical actions be reviewed before or after the fact?
04	Has special category data involvement been confirmed?	17	Can autonomous actions be paused or stopped by a human?
05	Has data minimisation been considered in its design?	18	Has the DPIA been reviewed or updated for this agent?
06	Have the systems it can access been fully mapped?	19	Have foreseeable misuse scenarios been documented?
07	Has its ability to write or modify data been assessed?	20	Have AI-specific risks been formally assessed?
08	Has its ability to execute code been assessed?	21	Can unauthorised or unexpected behaviour be detected?
09	Has its ability to browse the internet been assessed?	22	Would an AI-related data breach be identified quickly?
10	Has its ability to call third-party APIs been assessed?	23	Is the incident response plan ready for AI incidents?
11	Has the risk of it exceeding its intended scope been assessed?	24	Is there a named owner accountable for this agent?
12	Are its permissions based on least privilege?	25	Is it clear who approves deployment and reviews logs?
13	Is there an emergency kill switch or stop mechanism?		

Y N

AI Governance Readiness - score one point for every "Yes"

21–25

✓ Governance appears mature

Purpose, data, access, oversight and accountability are documented and controlled. Maintain review cycles.

15–20

⚠ Review recommended

Review recommended before deployment. Close the specific gaps identified above and re-assess.

0–14

Significant gaps identified

Significant governance gaps. Do not approve deployment until foundational controls are in place.

BEFORE APPROVAL

If more than five answers are "No", the organisation should review its AI governance controls before approving deployment.

educage.

SCOPE	GDPR · UK GDPR · EU AI Act
FORMAT	One-page summary

VISUAL SUMMARY

AI Agent Governance at a Glance

A one-page reference for the full lifecycle of AI agent governance - from policy through to continuous improvement.

GOVERN

- Policies
- Ownership
- AI strategy
- Accountability

PROTECT

- Access controls
- Least privilege
- Independent logging
- Human oversight
- Kill switch

MONITOR

- Continuous monitoring
- Incident detection
- Audit trails
- Risk reviews

IMPROVE

- DPIA updates
- Supplier reviews
- Lessons learned
- Governance reviews
- Training

Good Practice

AI governance should be reviewed throughout the system's lifecycle - not only before deployment.

⚠ Key Risk

An AI agent that exceeds its intended scope may create governance, security and privacy risks - even when its original task was legitimate.

✓ Practical Tip

Treat AI agents like privileged users, not software features.

Keep this checklist on file. Use it alongside the AI Agent Risk Assessment Canvas when approving new AI agents or reviewing existing deployments.