INTAI

DOWNLOADABLE GUIDE

RISK

AI Risk Management Guide

A practical overview of how organisations can identify, assess, manage and monitor the risks created by AI systems and AI-assisted work.

INTAI LIMITED
AI Governance & Assurance Consultancy

www.intai.co.uk
2026 Edition

Contents

1 Understanding the Shape of AI Risk

2 Identifying AI Risk Systematically

3 Assessing Likelihood & Impact

4 Treating & Mitigating Risk

5 Monitoring Risk Over Time

6 Embedding Risk Management into Governance

AI Risk Management Guide

A practical overview of how organisations can identify, assess, manage and monitor the risks created by AI systems and AI-assisted work.

> AI risk is often discussed as though it were a single category, when in practice it spans data, security, fairness, safety, operational, regulatory and reputational dimensions — each requiring somewhat different identification and treatment approaches. This guide provides a practical overview of how organisations can identify, assess, manage and monitor the risks created by AI systems and AI-assisted work, in a way that scales proportionately with organisational size and risk exposure.

CHAPTER 1 · SHAPE

Understanding the Shape of AI Risk

AI risk is not one thing. Treating it as a single, undifferentiated category tends to produce risk management that is either too shallow to catch specific problems, or so broad it becomes impossible to prioritise. A more useful approach separates AI risk into distinct categories, each with its own typical causes and mitigations.

RISK CATEGORY	TYPICAL CONCERN
Data & Privacy	Unlawful, excessive or poorly protected use of data
Security	Vulnerability to attack, manipulation or unauthorised access
Bias & Fairness	Unjustified discriminatory outcomes across groups
Safety & Reliability	Unpredictable or harmful behaviour under real conditions
Regulatory	Non-compliance with applicable law and standards
Reputational	Public or stakeholder perception of AI failures or misuse

> *A risk register that treats every AI risk the same way is a register that is not actually managing risk.*

KEY TAKEAWAYS

- AI risk spans several distinct categories, each with different typical causes and mitigations.
- Treating AI risk as undifferentiated tends to produce either shallow or unmanageably broad risk management.
- Effective risk management separates these categories explicitly rather than blending them together.

CHAPTER 2 · IDENTIFICATION

Identifying AI Risk Systematically

Risk identification should be a structured exercise applied consistently to every AI system, not an informal judgement made once at a system's inception. A consistent structure ensures risks are not missed simply because a particular reviewer did not happen to think of them.

A Practical Starting Structure

Working through each risk category systematically — even briefly — for every system catches risks that an unstructured review would likely miss. This need not be time-consuming for lower-risk systems, but should be applied consistently rather than selectively.

1. Has this system been assessed against each major risk category, not only the one most obviously relevant?
2. Would a different reviewer, using the same structure, likely identify the same risks?

KEY TAKEAWAYS

- Risk identification should be structured and consistent, not dependent on a single reviewer's judgement.
- Working through defined categories systematically catches risks an unstructured review would likely miss.
- Consistency matters more than depth for lower-risk systems, provided the structure is genuinely applied.

CHAPTER 3 · ASSESSMENT

Assessing Likelihood & Impact

Once identified, risks should be rated consistently for likelihood and impact, producing a comparable rating across different systems and risk categories. This comparability is what allows an organisation to prioritise its limited risk management attention sensibly.

RATING	LIKELIHOOD GUIDE	IMPACT GUIDE
Low	Rare or unlikely to occur	Minimal effect if it did occur
Moderate	Plausible within a reasonable timeframe	Noticeable effect requiring active response
High	Likely or already occurring	Significant harm, cost, or exposure

Ratings should be assessed twice: once for inherent risk, before considering existing controls, and again for residual risk, after accounting for them. The gap between the two is a useful indicator of how much a given control is actually reducing risk in practice.

KEY TAKEAWAYS

- Consistent likelihood and impact ratings allow risk to be compared and prioritised across systems.
- Risk should be assessed both before controls (inherent) and after them (residual).
- The gap between inherent and residual risk indicates how effective existing controls actually are.

CHAPTER 4 · TREATMENT

Treating & Mitigating Risk

Not every identified risk needs to be eliminated — some can be reasonably accepted, some transferred, and some avoided entirely by not proceeding with a given use case. What matters is that the treatment decision is deliberate and recorded, not left implicit.

Four Treatment Options

Accept: proceed as-is, having judged the residual risk acceptable. Mitigate: introduce further controls to reduce likelihood or impact. Transfer: shift some risk to a third party, such as through insurance or contractual allocation. Avoid: do not proceed with the use case in its current form.

> *An unrecorded risk decision is indistinguishable, later, from a risk nobody noticed.*

KEY TAKEAWAYS

- Risk treatment includes acceptance, mitigation, transfer and avoidance — not only mitigation.
- The treatment decision should be deliberate and explicitly recorded, not left implicit.
- A documented decision protects the organisation later, distinguishing a considered acceptance from an oversight.

CHAPTER 5 · MONITORING

Monitoring Risk Over Time

AI risk is not static. A system assessed as low-risk at deployment can become higher-risk over time as its use expands, its underlying data shifts, or the regulatory environment around it changes. Risk management that stops at initial assessment misses this entirely.

1. Is there a defined trigger for reassessing a system's risk, beyond a fixed annual schedule?
2. Would a material change in how a system is used be noticed and flow back into a risk reassessment?

KEY TAKEAWAYS

- AI risk changes over time as systems, data and regulation evolve, not only at initial deployment.
- Risk management should include defined triggers for reassessment, not rely solely on a fixed schedule.
- Changes in how a system is actually used should feed back into risk assessment.

CHAPTER 6 · EMBEDDING

Embedding Risk Management into Governance

Risk management delivers the most value when it is genuinely integrated into an organisation's wider governance activity, rather than existing as a parallel process that governance bodies reference only occasionally.

This means risk ratings inform governance approval decisions directly, the risk register is a standing agenda item for the AI governance body, and material risks are escalated through the same reporting lines as other organisational risk, rather than a separate, disconnected channel.

> **WHERE TO GO NEXT**
>
> INTAI's AI Risk Assessment Template and AI Risk Register Template provide the practical documents to apply this approach system by system, and the AI Risk Checklist offers a quick way to scan for common risk areas.

KEY TAKEAWAYS

- Risk management delivers most value when integrated directly into wider governance activity.
- Risk ratings should inform governance approval decisions, not sit in a separate, disconnected process.
- Material AI risk should be escalated through the same reporting lines as other organisational risk.

INTAI

Want support applying this guide?

INTAI works with organisations across Europe — from SMEs to public sector bodies and larger enterprises — to turn AI governance guidance into a practical, proportionate programme. If you would like support applying the ideas in this guide to your own organisation, our team is glad to help.

www.intai.co.uk | hello@intai.co.uk

© 2026 INTAI Limited. This guide is provided for general guidance purposes only and does not constitute legal, regulatory or professional advice. Organisations should adapt it to their own context and seek independent advice where appropriate.