

EXECUTIVE BRIEF

What Is an AI Risk Assessment? A Practical Guide to Identifying and Managing AI Risk

An AI risk assessment is a structured review of an AI system to identify risks and decide how to manage them before deployment. What it is, when to run one, the nine areas it should review, and how to set risk levels.

EXECUTIVE SUMMARY

An AI risk assessment is a structured review of an AI system to identify risks and decide how to manage them before deployment. What it is, when to run one, the nine areas it should review, and how to set risk levels.

KEY TAKEAWAYS

- An AI risk assessment is a structured review of an AI system or use case to identify risks and decide how to manage them, before deployment.
- Risk depends on context: the same capability can be low risk in one use and high risk in another.
- Review nine areas: purpose, data, impact on people, accuracy, bias and fairness, security and privacy, human oversight, vendor risk, and legal alignment.
- Score risk by likelihood and impact, and match the depth of review to the risk level.
- It is not a one-time stamp: reassess when the system, vendor, data, or use case changes.
- The assessment identifies risk; the AI risk register tracks and manages it over time.

Read the full guide at ai-governance-jobs.com/insights/how-to-conduct-an-ai-risk-assessment/
Executive Brief · AI Governance Essentials · Source: GRC-Careers.org