Technical AI Safety Specialist

Welcome to Group Technology, where we pride ourselves on engineering solutions that directly impact Nordea's 2030 goals to modernize data technology and accelerate AI.

We’re seeking a Technical AI Safety Specialist to build deep, hands-on technical expertise in the engineering-heavy risk areas that come with scaling AI and agentic systems across topics like AI gateway architecture, MCP servers and agentic tooling, red-teaming methodology, evaluation frameworks, and the cybersecurity implications of new AI capabilities. You’ll partner closely with engineering, architecture, and information security teams to turn that technical understanding into pragmatic governance guidance, technical risk assessments, and internal white papers. This role sits within the AI Safety team, which brings together architectural governance, AI risk operations and advisory capabilities to enable safe, trusted and compliant AI across Nordea. As AI and agentic systems become more complex, the team needs stronger technical expertise to understand how these systems are designed, secured, evaluated and operated, and to translate that understanding into practical guardrails, risk assessments and guidance.

Nordea is a place where traditions meet tomorrow. We're not just a bank; we're a tech employer on a mission to evolve finance securely and responsibly. Together, we impact millions of people's daily lives by ensuring they can access our solutions anytime, anywhere, while safeguarding their personal data and wealth. Join us in making an impact on the banking industry.

This position is a full-time, permanent role with hybrid working location based in Helsinki, Finland, Copenhagen, Denmark, or Stockholm, Sweden.

About our team

Meet the AI Safety team, part of AI Portfolio Delivery within Group AI, responsible for enabling safe, trusted and compliant AI across Nordea through efficient, scalable and holistic risk governance. AI and agentic systems are moving fast, and so is our remit. We’re building deeper technical capability so that our governance keeps pace with new

architectures, tooling, and emerging risks, helping to streamline AI deployment and maintenance while ensuring effective risk management and regulatory adherence.

Our culture emphasizes collaboration, continuous learning, and innovation in an open environment where diverse ideas thrive. We support each other's growth while pushing technology boundaries, prioritizing work-life balance, and open communication.

What you will be doing:

- Build and maintain deep technical knowledge of engineering-heavy AI risk areas, including AI gateways (e.g. Kong), MCP servers and agentic tooling, evaluation frameworks, and red-teaming methodology
- Partner with engineering, platform, and cybersecurity teams to understand how AI and agentic systems are built and operated, and translate that understanding into governance guidance
- Track emerging AI developments and assess the technical and security implications of new capabilities, protocols, and attack surfaces
- Author technical white papers, position papers, and risk assessments on AI infrastructure and emergent risk topics, written for both technical and non-technical audiences
- Contribute technical input to governance frameworks, evaluation criteria, and technical controls for AI systems, without owning their day-to-day implementation
- Engage with senior technical stakeholders across the organisation to pressure-test governance approaches against real engineering constraints
- Act as a technical peer to engineering and cybersecurity teams in cross-functional discussions on AI security protocols and technical controls
- Help shape what technical AI safety expertise should look like as the team’s scope and remit continue to evolve
- Translate deep technical understanding into governance positions, RACI, and decision processes that first- and second-line stakeholders can act on

Who you are

You have a technical/engineering background and enjoy going deep on hard technical problems. You may not write production code day to day anymore, but you can hold your own in a conversation with a platform engineer about gateway architecture, or with a security engineer about red-teaming an LLM. You’re comfortable operating in ambiguity – this role and team will keep evolving as the AI landscape does – and you’re good at translating technical depth into guidance that non-technical stakeholders can act on.

We're looking for someone with most of the following experience and skills:

- BSc or MSc in Computer Science, Software Engineering, Cybersecurity, or a related technical discipline
- 5+ years of experience in software engineering, platform engineering, security engineering, or a similar hands-on technical role
- 5+ years of experience working with AI/ML system architectures, including how models, data pipelines, and infrastructure fit together
- Experience with API gateways or similar infrastructure (e.g. Kong) and/or agentic protocols such as MCP
- Understanding of cybersecurity principles as they apply to AI and agentic systems, including common attack surfaces
- Familiarity with red-teaming or adversarial testing approaches for AI systems
- Hands-on experience with cloud platforms (e.g. AWS, GCP) in an AI context
- Ability to communicate technical depth clearly to both technical and non-technical stakeholders

It would be ideal if you also have:

- Experience with AI evaluation frameworks or model/system testing methodologies
- Experience writing technical white papers, research summaries, or position papers
- Exposure to AI governance, risk, or compliance work, and an interest in bridging it with engineering
- Understanding of regulatory frameworks for AI in financial services or other regulated industries
- Experience with vendor or third-party technical risk assessments for AI platforms
- Familiarity with data governance and privacy considerations for AI systems

We encourage applications from candidates who are passionate about the technical side of AI safety and meet most of these criteria. We value diverse perspectives and welcome applicants from all backgrounds.

If this sounds like you, get in touch!

Next steps

We kindly ask you to submit your application as soon as possible, but no later than 16/08/2026. Any applications or CVs sent by email, direct messages, or any other channel than our application forms, will not be accepted or considered.

At Nordea, we know that an inclusive workplace is a sustainable workplace. We deeply believe that our diverse backgrounds, experiences, characteristics, and traits make us better at serving our customers and communities.

If you have any questions about our recruitment process, please reach out to our tech recruiter and main point of contact anna.dahlstrom@nordea.com.

Nordea	Ansøgningsfrist	Kontaktperson
Taastrup, 2630 Taastrup www.nordea.dk	16. august 2026 Technical AI Safety Specialist	anna.dahlstrom@nordea.com