



JOB OPPORTUNITY

Reference Number
MND/PO/J/26/024

Date
18 August 2026

Maldives Innovation
and Digital Company

+960 3335777
info@mindco.mv

G. Niusha, 5th Floor
Lily Magu, Malé 20130
Republic of Maldives

Post	Senior AI Engineer
No of Positions	01
Employment Type	Contractual Full-Time
Terms of Employment	This is a one-year assignment from the commencement of work. Contract renewal eligibility requires both the successful completion of probation and a satisfactory annual performance review
Location	Head Office / Male'
Remuneration	Based on qualifications and experience

BACKGROUND

MINDCo is seeking an experienced and delivery-focused **Senior AI Engineer** to design, build, and operate production-grade AI/ML systems for its digital government ecosystem. Reliable AI systems depend on high-quality data, the role will also provide technical ownership of the data engineering capabilities required to support AI, machine learning, and analytics workloads.

The successful candidate will be highly proficient in modern AI and LLM development and capable of working across the lifecycle from operational data and curated datasets to evaluated models and deployed AI services, while collaborating with platform, product, and engineering teams on broader infrastructure where required.

OBJECTIVE OF THE ASSIGNMENT

To recruit a Senior AI Engineer who can design and deliver reliable, secure, and well-evaluated AI capabilities, including adapted and fine-tuned models, retrieval-augmented generation, agentic applications, and the data foundations required to support them, while establishing appropriate standards for AI quality, security, governance, and responsible use across MINDCo's digital products.

KEY RESPONSIBILITIES

The role will provide technical leadership and direct implementation for critical AI and AI-related data components while collaborating with existing engineering teams rather than acting as the sole implementer of all organizational data infrastructure:

AI Engineering

- Define and maintain MINDCo's AI technical architecture and engineering standards, including model selection, build-versus-buy decisions, hosted versus self-hosted models, inference architecture, retrieval infrastructure, agent architecture, and integration patterns for AI-enabled services.
- Fine-tune, adapt, and evaluate LLMs and other ML models for MINDCo use cases, including low-resource Dhivehi language capabilities such as text processing, embeddings, OCR, speech recognition, text-to-speech, translation, transliteration, and information retrieval.
- Build retrieval-augmented generation (RAG) pipelines, agentic workflows, and tool/API integrations that connect models to operational systems.
- Establish benchmarks, automated evaluation suites, and feedback loops; implement guardrails and human oversight mechanisms for accuracy, safety, security, and responsible AI use.
- Own model deployment, versioning, inference optimization, and monitoring - from experiment tracking through production serving, quality monitoring, cost optimization, and drift detection.

Data Engineering

- Collect, clean, label, version, and document datasets prepared specifically for model training and fine-tuning.
- Build and maintain the batch and real-time pipelines, ETL/ELT processes, and data infrastructure required to support AI, analytics, and the central data platform, in collaboration with relevant engineering teams.
- Support BI tools and curated datasets that enable non-technical stakeholders to access self-serve analytics.
- Ensure compliance with data security and privacy standards, and monitor data quality, integrity, and observability with automated alerting.

Collaboration

- Work closely with R&D, product, and engineering teams to translate business requirements into AI and data solutions.
- Participate in design and architecture reviews and incorporate stakeholder feedback.
- Prepare technical documentation and support handover, knowledge sharing, and internal capability building.

REQUIRED QUALIFICATION & EXPERIENCE

Education

- Bachelor's degree in Computer Science, Artificial Intelligence, Data Science, Engineering, or a related field.
- A Master's degree in a relevant field is an advantage.
- Equivalent professional experience may be considered.

Mandatory Technical Requirements

- At least 4 years of experience in AI/ML engineering, data engineering, or a similar role, with demonstrable experience taking ML systems to production.
- Hands-on experience fine-tuning and deploying LLMs and other models, including parameter-efficient fine-tuning, embeddings, vector databases, and RAG architectures.
- Experience building LLM applications: prompt engineering, agent frameworks, structured outputs/function calling, and inference serving.
- Strong proficiency in Python for ML and pipeline development, and SQL for complex querying.
- Strong hands-on experience with cloud infrastructure for data and AI workloads, including GPU/accelerated compute for training and inference.
- Experience with cloud data warehouses, data modeling, ETL/ELT frameworks, workflow orchestration, and BI tooling.
- Solid understanding of data governance, security protocols (encryption, masking), and access control mechanisms (RBAC).
- Familiarity with MLOps practices such as experiment tracking, model registries, and CI/CD for models.
- Understanding of privacy and responsible AI considerations specific to AI systems.

Preferred / Optional Skills

- Experience working in Agile/Scrum environments.

Added Advantage

- Experience with low-resource language NLP (ideally Dhivehi or related scripts) and OCR systems.
- Contributions to open-source model ecosystems.
- Experience deploying and operating self-hosted or on-premises model inference.

COMPETENCIES

- Excellent problem-solving and analytical skills.
- Strong sense of ownership over the AI and data stack.
- Ability to balance research and experimentation with delivery deadlines.
- Strong communication and collaboration skills across technical and non-technical teams.
- Ability to work independently and manage multiple workstreams simultaneously.

DELIVERABLES

- Production-ready fine-tuned models and LLM-powered features integrated with core services.
- Evaluation framework with benchmarks and continuous monitoring for model quality, safety, and drift.
- Curated, cleaned, and versioned AI-ready datasets for training and fine-tuning, with documented lineage.
- MLOps infrastructure covering model registry, deployment workflows, and lifecycle monitoring.
- Documented data infrastructure with production-ready batch and real-time pipelines and support for priority self-serve BI and analytics requirements.
- Governance framework covering data and model security, access controls, and quality monitoring.

APPLICATION SUBMISSION

Interested candidates are invited to submit their completed job application form together with an up-to-date CV to talent@mindco.mv no later than 16:00 hrs on 30th August 2026.

ADDITIONAL DOCUMENTS FOR SHORTLISTED CANDIDATES

Please be informed that applicants who are shortlisted will be contacted for an interview and will be required to provide the following documents via email to talent@mindco.mv upon notification:

- A recent passport-size photograph (digital copy)
- A scanned copy of the National ID card
- Scanned copies of educational certificates
- Employment reference letters from previous employers
- Valid police report

OUR COMMITMENT TO DIVERSITY & INCLUSION:

At MINDCo, we champion an inclusive and respectful workplace where everyone has equal opportunities to thrive. We welcome applications from all qualified individuals, and we strongly encourage women to apply.