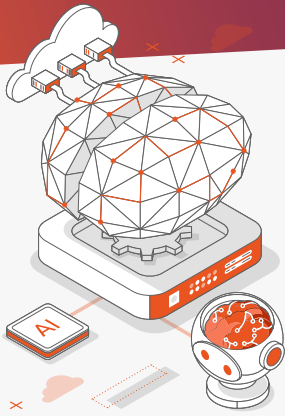


Bring an optimized LLM to market – and secure it

5-day workshop to accelerate time-to-market, elevate in-house expertise, and bring your AI-powered system to production.



Enterprises often face challenges when implementing GenAI projects in production, such as scalability, security, and integration issues. This retrieval augmented generation (RAG) workshop is designed to help your team prepare for the launch of a production-ready, optimized LLM.

By the end of the workshop, you will have a comprehensive understanding of RAG and LLM theory, architecture, and best practices. Together, we will build and deploy solutions tailored to your needs.

Retrieval Augmented Generation (RAG)

RAG can enhance the accuracy and relevance of your LLM outputs by enabling the model to reference additional sources outside of its training data. RAG benefits include:


Up-to-date information

Enhanced Accuracy

Broad knowledge base

Context-specific response

5-Day workshop content




1st Day

Data & AI, MLOps introduction and architecture. Learning about best practices ensures reliability and success in shipping models to production.



2nd Day

Deep-dive into LLM and RAG architecture with security and data and AI confidentiality considered in the design.



3rd - 5th Day

By the end of the workshop, you will learn and gain the following:

- An open-source MLOps architecture customized for your specific use case
- High-level and low-level designs of the AI infrastructure
- Hands-on experience using the latest open-source tools for your use cases
- Written content for learning and upskilling about RAG.

Benefits of joining the workshop

Build the right RAG architecture

Build the architecture for an AI-powered search system tailored to your use case in 5 days.

Learn and use best-in-class RAG tooling

Easily unlock the benefits of RAG with open source tooling for your entire data and machine learning lifecycle.




5 days
with Canonical experts on-site

Become an open source RAG expert

Elevate the internal expertise of your team. Learn how to make the most of cutting-edge open source technologies.

Run production-grade RAG on any hardware and cloud

Run RAG-enabled LLMs on any hardware and cloud in production and at any scale.



10 Participants
per workshop

Ship fast and achieve value

Accelerate time-to-market and product delivery to stay ahead of the competition.

Secure your stack with confidence

Secure your GenAI projects and learn best practices for maintaining your software stack. Learn to protect your code, data and machine learning models in production with confidential AI.



\$ 30,000
Fixed cost

[Contact us](#) to book your MLOps workshop. Learn more about MLOps in our [guide](#). Listen to our [podcast](#) for conversations about open source, machine learning and the latest trends in the industry.