

# AI296

## Develop, Test, and Run Granite Family LLMs with Red Hat Enterprise Linux AI

---

An introduction to fine-tuning, serving, and using generative AI models with Red Hat Enterprise Linux AI.

Develop, Test, and Run Granite Family LLMs with Red Hat Enterprise Linux AI (AI296) provides students with the fundamental knowledge about how to fine-tune, serve, and use generative AI models tailored to specific business needs. This course helps students build core skills to train models in a secure and private environment.

This course is based on Red Hat Enterprise Linux AI 1.5.

### Course content summary

- Generative AI Fundamentals: Capabilities, Challenges, Models, and Techniques
- Granite Models For Enterprise Generative AI
- Training Large Language Models with Red Hat Enterprise Linux AI
- Deploying Trained Models with Red Hat Enterprise Linux AI

### Audience for this course

- Data Scientists and AI specialists who want to learn more about Granite models and the process of training and aligning GenAI models for specific business requirements.
- Developers with basic AI/ML knowledge or machine learning engineers who are interested in learning how to develop, test, and run AI-based enterprise applications in an enterprise and secure environment.
- System administrators who are interested in learning how to manage the company's hybrid cloud infrastructure to boost the adoption of generative AI and LLMs, while enforcing security policies.
- Other professionals, such as SMEs and domain experts, who are interested in learning how to contribute to the training of secure models tailored to their business needs.

### Prerequisites for this course

- A basic understanding of AI & ML is recommended but not required.
- Familiarity with the Linux command line.

## Outline for this course

- **Generative AI Fundamentals: Capabilities, Challenges, Models, and Techniques**  
Define Generative AI, its importance, benefits, and challenges. Learn to select suitable GenAI and LLM models and techniques, considering their capabilities, use cases, and limitations.
- **Granite Models For Enterprise Generative AI**  
Describe and understand Granite models, and their capabilities to solve enterprise use cases, compared to larger, closed models.
- **Training Large Language Models with Red Hat Enterprise Linux AI**  
Enable technical and non-technical stakeholders to collaborate in the creation, training, and deployment of large language models (LLMs) optimized for specific business needs by using Red Hat Enterprise Linux AI and IBM Granite.
- **Deploying Trained Models with Red Hat Enterprise Linux AI**  
Deploy, serve, and operate generative AI models trained with Red Hat Enterprise Linux AI.

## Impact on the individual

As a result of attending this course, you will comprehend the fundamentals of generative AI. You will be able to understand the capabilities of the Granite models to solve enterprise use cases. Finally, you will be able to train, deploy, and serve large language models with Red Hat Enterprise Linux AI.