CORSEN AI Academy

Mistral Deployment: Cloud and Self-Hosting

Course syllabus

Version 2026-11

Overview

Deploy Mistral everywhere: AI Studio, Azure AI, Google Vertex, AWS Bedrock, Snowflake, and self-hosting with vLLM, TensorRT-LLM, and TGI.

Learning objectives

- Compare AI Studio, Azure AI, Vertex, Bedrock, and Snowflake for hosting Mistral models
- Structure organizations, workspaces, and pricing tiers within Mistral AI Studio
- Instantiate the MistralAzure client in Python and deploy to production
- Deploy a Mistral model via self-hosting with vLLM, TensorRT-LLM, or TGI
- Size GPU infrastructure and arbitrate between managed API and internal hosting

Audience and prerequisites

Developers, engineers, and integrators looking to leverage Mistral AI via API in their applications.

Confirmed experience in application development and deployment to production.

Duration and format

Estimated duration : 1 h 45 min. 100% online self-paced course, accessible 24/7 from a browser. Progress saved, resume anytime. Estimated duration: 1 h 45 min.

Content

- Mistral AI Studio (3 lessons)
- Azure AI (3 lessons)
- Google Vertex, Bedrock, and Others (3 lessons)
- Self-Hosting (4 lessons)
- Conclusion (1 lessons)

Assessment

Self-assessment based on the Key takeaways concluding each lesson, and direct implementation of examples on your own examples. A final quiz of 5 corrected questions at the end of the last lesson validates acquired knowledge.

Accessibility

Self-paced course accessible to people with disabilities: text content compatible with screen readers, keyboard navigation, and compliant contrasts. For any adaptation needs, email hello@corsen.ai.

Price and recognition

Free

Certificate of completion issued by Corsen AI.

Access time

Immediate after account creation and email confirmation.

Corsen AI — Paris, France — SIRET 10297040700015
academy@corsen.ai — https://academy.corsen.ai