



399-HOUR EXCELLENCE CURRICULUM

2027 COHORT

AS300 — AI Platform Engineer

Intensive 399-hour curriculum to master Red Hat OpenShift AI (RHOAI), LLM/ML model deployment, MLOps/LLMOps pipelines, and GenAI applications in production.



DURATION

399 h

57 days (12 weeks)

CERTIFICATIONS

EX267 AI/ML

Exams + Retakes

PEDAGOGY MODEL

75% / 25%

Red Hat + Capstone

FORMAT

100% Live Remote

Sync + Labs 24/7

TARGET PROFILES & CAREER OPPORTUNITIES

Cloud Engineers, Senior DevOps, Backend Developers, and Architects mastering Enterprise AI Platforms and MLOps in production.

2027 TUITION

€7,490 incl. VAT

(20% VAT included)

✓Eligible for OPCO, Corporate training & Public funding



MANIFESTO & VISION

Aperta Scientia & Our Core Values



Aperta Scientia — 'open knowledge' in Latin — is a higher education engineering institution dedicated to AI, Hybrid Cloud, and Quantum. In an era of mass automation, we cultivate the deep desire to understand, training critical-thinking, sovereign builders capable of architecting and operating resilient production platforms.



The Desire to Learn

Learning is far more than consuming generated answers: we awaken intellectual curiosity and the joy of mastering system internals.



Thinking Minds & Builders

We train engineers who direct technology with architectural rigor, capable of designing and debugging without digital crutches.



Open Source Culture

Open source is our cornerstone: Red Hat technologies, Linux RHEL, Kubernetes, and active contribution to an enduring global ecosystem.



Equal Opportunity

Admissions rest purely on motivation and intellectual potential. Personalized accommodations for learners with disabilities (PSH).



Aperta Scientia Tech Radar — Ongoing Research & Intelligence

Leading-edge benchmarks on distributed LLM serving (vLLM, KServe, RHOAI), InstructLab fine-tuning, GitOps, and Quantum (Qiskit).

[Explore Radar →](#)



Why this curriculum? Industry Challenges

Industrializing AI models requires robust, scalable, and sovereign infrastructure. Enterprises no longer just need data scientists; they urgently require engineers capable of deploying, securing, and operating GPU clusters, low-latency LLM serving pipelines (vLLM, KServe), and end-to-end MLOps platforms.



Red Hat

A Historic Technology Partnership

Over 20 years of continuous Red Hat ecosystem expertise

RHLS Enterprise (365 Days)

Unlimited 24/7 access to all official Red Hat courses, downloadable e-books, and chaptered video modules with transcripts.

400+ Hours Dedicated Cloud Labs

Real bare-metal and multi-node clusters hosted on Red Hat's global infrastructure for hands-on practice in production conditions.

Expert Certified Instructors

100% live synchronous virtual classrooms led by certified Red Hat instructors at the peak industry level.

Exam Vouchers Included

Official remote proctored certification exam sittings, including free retake sessions.



The 75 / 25 Pedagogical Model

75% of the time is devoted to mastering official Red Hat courses (theory + guided labs).
The remaining 25% is dedicated to the industrial capstone project and official certification exams.

Red Hat Partner →

57–Day Pedagogical Curriculum (399 Hours)

A structured 4–stage progression combining container foundations, advanced operations, specialization, and capstone engineering.

STAGE 1 · 13 DAYS (91H)

Container Foundations

Podman, Linux CLI, Skopeo, Buildah, rootless containers & OpenShift.



STAGE 2 · 17 DAYS (119H)

OpenShift Admin & Dev

Cluster administration, GitOps (ArgoCD), Tekton CI/CD pipelines.



STAGE 3 · 8 TO 12 DAYS

AI & MLOps

Red Hat OpenShift AI (AI267), LLM Serving, InstructLab.



STAGE 4 · 19 DAYS (133H)

Certifications & Capstone

Intensive exam prep bootcamp, certification sittings, and capstone.



Code	Module / Official Course	Duration	Key Competency
D0188	Podman & Containers Architecture	3d / 21h	Rootless containerization, Podman, Buildah
D0180	OpenShift I: Containers & Kubernetes	5d / 35h	Kubernetes & OpenShift Platform Fundamentals
D0280	OpenShift Administration II	5d / 35h	Production cluster administration, security, RBAC
D0380	OpenShift Administration III: Scaling	5d / 35h	GitOps with ArgoCD, scaling and observability
AI267	Red Hat OpenShift AI (RHOAI) & MLOps	5d / 35h	Model Serving, Data Science pipelines, Jupyter
AI-LAB	Distributed LLM Serving (vLLM, KServe, Triton)	3d / 21h	High-performance serving & GPU orchestration
AI-FT	Sovereign Fine-Tuning (InstructLab & Granite)	3d / 21h	Model alignment and Enterprise RAG
B00T	Red Hat EX267 AI/ML Certification Bootcamp	5d / 35h	Intensive exam drills & official certification
PR0J	Sovereign AI Capstone Project & Jury Defense	15d / 105h	Full AI platform deployed on production GPUs



ARCHITECTURE, GITOPS & SPECIALIZATION

Enterprise AI Platforms & MLOps



An in-depth focus on the four key pillars of enterprise sovereign AI: Red Hat OpenShift AI (AI267), distributed LLM serving (vLLM/KServe), InstructLab fine-tuning, and enterprise RAG.

AI267 — Red Hat OpenShift AI & Plateforme MLOps (5 Jours / 35h)

Préparation EX267

Déploiement et exploitation d'environnements interactifs et pipelines de données pour équipes Data Science.

- ✓ Workbenches Jupyter, PyTorch et Ray sur OpenShift
- ✓ Pipelines MLOps reproductibles avec KubeFlow & Elyra
- ✓ Registre de modèles & suivi des expérimentations
- ✓ Partage et allocation dynamique des ressources GPU

AI-LAB — Inférence LLM Distribuée (vLLM & KServe) (4 Jours / 28h)

High-Throughput

Serving haute performance de grands modèles de langage ouverts (Granite, Mistral, LLaMA).

- ✓ Moteur vLLM & optimisation mémoire PagedAttention
- ✓ Serveurs d'inférence KServe & ModelMesh serverless
- ✓ Découpage et mutualisation GPU (MIG / time-slicing)
- ✓ Optimisation des temps de latence TTFT & débit tokens/s

AI-FT — Alignement & Fine-Tuning InstructLab (4 Jours / 28h)

Souveraineté IA

Entraînement incrémental et spécialisation métier de modèles LLM sans fuite de données.

- ✓ Génération de données synthétiques avec InstructLab
- ✓ Fine-Tuning efficace (LoRA / QLoRA) sur clusters GPU
- ✓ Alignement de compétences et connaissances métiers
- ✓ Évaluation et benchmarks comparatifs des modèles entraînés

AI-RAG — Architectures RAG & Gouvernance TrustyAI (4 Jours / 28h)

Guardrails & Prod

Génération augmentée par récupération et sécurisation des applications d'IA en production.

- ✓ Base vectorielle Milvus & indexation sémantique
- ✓ Pipelines LangChain & LlamaIndex d'entreprise
- ✓ Détection de dérive et explicabilité avec TrustyAI
- ✓ Garde-fous de sécurité, conformité RGPD & AI Act

ENTERPRISE STANDARDS & RED HAT ALIGNMENT



The 12 Core Technologies of the Track

Every technology is directly anchored into official curriculum courses to guarantee operational excellence in production.

 Red Hat OpenShift AI Red Hat OpenShift AI v2.16 AI267 Jupyter/Ray workbenches & KubeFlow.	 kubernetes Kubernetes & GPU DRA v1.31 D0280 • AI267 Dynamic DRA allocation & MIG GPU slicing.	 InstructLab InstructLab & RHEL AI v0.20+ AI-FT • AI267 Private sovereign LoRA GPU fine-tuning.
 vLLM vLLM Serving Engine v0.6+ AI-LAB • AI267 PagedAttention & ultra-low latency.	 KServe KServe & ModelMesh v0.14 AI267 Serverless serving & dynamic auto-scaling.	 IBM IBM Granite Open Models v3.0 AI-FT • AI-LAB Open sovereign foundation LLMs code & logic.
 PyTorch PyTorch Framework v2.5 AI267 Tensor compute & distributed CUDA GPU.	 Ray & KubeRay Cluster v2.35 AI267 Distributed training & AI scaling.	 JupyterLab Workbenches v4.2 AI267 Interactive workbenches with GPU hooks.
 TrustyAI & Guardrails v1.4 AI-RAG Model explainability, drift & AI Act.	 Milvus Milvus Vector Database v2.4 AI-RAG Vector database & hybrid semantic search.	 LangChain LangChain & LlamaIndex v0.3 AI-RAG Autonomous agents & enterprise RAG.



CAREER ROADMAP & MARKET INSIGHTS

Target Roles, Job Market & Salaries



TARGET JOB POSITIONS

- ✓ Ingénieur Plateforme IA (AI Platform Engineer)
- ✓ Architecte Infrastructure IA & GPU

- ✓ Ingénieur MLOps / LLMOps
- ✓ DataOps / AI Systems Engineer

MARKET SALARY BENCHMARKS

Sources: Hays IT · Michael Page · Robert Half

JUNIOR / RECONVERSION

45k€ – 55k€

Starting annual gross salary

CONFIRMED (3–5 YRS)

60k€ – 85k€

High recruiter demand

FREELANCE DAILY RATE

550€ – 900€

Average daily rate

EMPLOYER PARTNERS & HIRING NETWORK

**devoteam**

European leader in digital transformation and cloud consulting, hiring top certified platform engineers.

STARTX

Part of Devoteam

Over 20 years of premier expertise on the Red Hat, OpenShift and Kubernetes ecosystem in Europe.



PERFORMANCE-BASED EXAMS & BOOTCAMP

Red Hat Certifications & Intensive Bootcamp



Red Hat certifications are 100% hands-on performance-based evaluations on real live clusters. Our 6-day intensive bootcamp ensures complete readiness.

EXAMEN EX267

Hands-on · 3h

Red Hat Certified Specialist in OpenShift AI & ML

Validates deploying enterprise RHOAI, GPU orchestration, LLM serving, and MLOps.

TITRE RNCP NIVEAU 7

EQF 7

AI / DevOps Platform Engineer

Formal qualification registered with France Compétences (EQF 7).

+15% à +25% Average Salary Premium for Certified Engineers**Source: Red Hat Survey →**

According to the *Global Knowledge IT Skills & Salary Report*, Red Hat certified engineers earn on average 15% to 25% higher salaries and advance twice as fast into lead and architect roles.

Timed Drills & Expert Coaching

5 days of timed mock evaluations with in-depth coaching by certified Red Hat instructors.

Vouchers & Free Retakes Included

Every official exam voucher includes a free retake attempt. Secure remote 24/7 delivery.

PRODUCTION SIMULATION • 15 DAYS (105H)

Industrial Capstone & Jury Defense



During the final 15 days, students collaborate in small engineering teams (2–3 members) on architecting, deploying, and operating a production-grade enterprise platform.

SPRINT 1 (J1–J4)

Core & Topology

Multi-node setup, topology & RBAC.

SPRINT 2 (J5–J8)

GitOps & CI/CD

Tekton pipelines & declarative ArgoCD.

SPRINT 3 (J9–J12)

Workloads

vLLM serving, KServe & RAG.

SPRINT 4 (J13–J15)

Runbook & Jury

Production runbook & oral defense.



45-min Defense before CTO & Tech Directors Jury

Live architecture presentation, demonstration on real clusters, disaster recovery drill, and technical Q&A with enterprise leaders.

OFFICIAL JURY





2027 Tuition, Funding & Admissions

Certified career tracks eligible for comprehensive corporate and public funding support.

2027 TUITION

€7,490 incl. VAT

(20% VAT included)

Eligible funding: OPCO direct settlement (Atlas, Akto...), France Travail POEC, AIF, corporate training plans, phased self-funding.

ACCESSIBILITY & QUALITY

100% accessible to individuals with disabilities, featuring tailored workstation and exam accommodations.

Disability Officer Contact: care@apertascientia.io

4-STEP ADMISSION PROCESS

01. Candidature

Official online form.

02. Entretien

30-min meeting.

03. Financement

Agreement & OPCO file.

04. Intégration

RHLS & Discord setup.

CONTACT US & OFFICIAL NETWORKS

APERTA SCIENTIA



Phone

+33185 39 02 26



Email

contact@apertascientia.io



LinkedIn

[Aperta Scientia](#)



X (Twitter)

[@ApertaScientia](#)



Discord

[ApertaScientia](#)

[Apply →](#)