

GLOBALLY RECOGNIZED · VENDOR-NEUTRAL · ANSI-ALIGNED

Certified LLM Ops Professional

— (CLLMOP)

Become a Globally Recognized LLM Ops Professional

Earn the CLLMOP credential — operate large-language-model applications in production end to end: model selection and serving, prompt & context engineering, RAG, fine-tuning, evaluation, observability, security and scaling. **Today's offer: \$1200 → \$600.**



4.8/5 Reviews.io (96%) · 4.4/5 Trustpilot · 2,50,000+ certified

Enroll at gsdcouncil.org/llm-ops-certification

Why This Certification, Why Now

Everyone can call an LLM API. Far fewer can run one reliably, safely and affordably in production. That operational layer — serving, retrieval, evaluation, observability and governance — is where certified LLMOps professionals earn their seat.



LLMs behave unlike traditional software

Non-determinism, prompt sensitivity and token economics break the assumptions classic MLOps was built on.



Prototypes are easy; production is hard

Latency, cost, hallucinations and drift decide whether an LLM app survives contact with real users.



RAG and fine-tuning need real engineering

Retrieval quality, evaluation harnesses and guardrails separate a demo from a dependable product.



Cost, safety and governance are board-level

Token spend, prompt injection, PII and compliance (GDPR, SOC 2) can't be bolted on after launch.

FROM THEORY TO PRACTICE

What You'll Be Able to Do



Select and serve LLMs with the right inference architecture, quantization and GPU sizing



Engineer production-grade prompts, structured outputs, function calling and context management



Design and optimize RAG pipelines with vector databases, hybrid search and re-ranking



Choose between prompting, RAG and fine-tuning (SFT / PEFT / LoRA) for each use case



Evaluate quality with golden datasets, LLM-as-a-Judge and automated production quality gates



Run observability, cost tracking, security guardrails and scalable multi-model systems

Why GSDC

**4.4 / 5**

Trustpilot · Excellent

**4.8 / 5**

Reviews.io · 96% recommend

2,50,000+

professionals certified

500+

Fortune clients

190+

countries served

Since 2016

US · SG · Switzerland

- ✓ Independent, vendor-neutral certification body established in 2016 — offices in the US, Singapore, and Switzerland.
- ✓ Accredited by ABICB (Accreditation Board for International Certification Bodies) and a proud member of ANSI.
- ✓ Alumni work at 500+ Fortune companies across 190+ countries; 100+ Authorized Training Partners worldwide.
- ✓ Certification built and reviewed with LLM, MLOps and platform-engineering practitioners running generative-AI systems in production.

RECOMMENDED BY**Forbes · Indeed · TechTarget · CareerSidekick · LeanIX · People Managing People · Authentic****ABICB Accredited****ANSI Member****Vendor-Neutral***Read verified reviews: trustpilot.com/review/gsdccouncil.org · reviews.io*

TRUSTED BY 2,50,000+ PROFESSIONALS



Our Global Clients

Professionals across technology, cloud, finance, and enterprise transformation hold GSDC certifications.

Oracle

Amazon

Microsoft

Cisco

Veritas

HP

Bank of America

Deloitte

Volkswagen

IS THIS FOR YOU?

Who Is This Certification For?

- ML / MLOps and platform engineers
- Data scientists and AI/ML engineers moving to LLMs
- Backend and software engineers building LLM applications
- AI product managers and engineering leads
- DevOps / SRE and cloud engineers supporting AI workloads
- Consultants and architects designing production LLM systems

PRE-REQUISITES

Basic familiarity with Python, APIs and cloud concepts is recommended. Experience in software engineering, data science, machine learning or DevOps is beneficial. No advanced research background is required, and no prior LLM or MLOps certification is required.

Learn from Experts

Learn from experienced practitioners and industry leaders who design, update, and review the program — bringing real-world AI, machine-learning and platform-engineering expertise to every module.



Maxim Salnikov
MICROSOFT
Digital and App Innovation
Business Lead, Western
Europe



Trevor Wiseman
THE CIRCUIT
VP of Technology & AI
Governance



Swathi Adimulam
ORACLE
OCI Cloud & AI Architect



Hayk Hakobyan
INSIGHTGENIE
Co-Founder



Baris Dirim
CULTUREASY
Human Systems Architect

What You'll Learn

Ten modules across the full LLMOps lifecycle — from foundations and serving to RAG, fine-tuning, evaluation, observability, security and scaling — with real-world scenarios and Learn-By-Doing activities in every module.

1

Foundations of LLMOps

LLMOps vs MLOps, LLM system characteristics, lifecycle, roles, and build-vs-buy.

2

Model Selection & Serving Architecture

Proprietary vs open-weight models, inference serving (vLLM/TGI/Triton), GPU sizing and quantization.

3

Prompt & Context Engineering

Production prompts, templates & versioning, structured outputs, function calling, context management.

4

Retrieval-Augmented Generation (RAG)

RAG architecture, embeddings, vector DBs (pgvector/Pinecone/Weaviate/Qdrant), hybrid search, re-ranking.

5

Fine-Tuning & Model Customization

Prompting vs RAG vs fine-tuning; SFT, PEFT, LoRA/QLoRA, dataset curation, RLHF/DPO.

6

Evaluation & Quality Assurance

Evaluation methods, golden datasets, LLM-as-a-Judge, hallucination/toxicity, A/B tests, quality gates.

7

Production Infrastructure & MLOps Integration

Production architecture, CI/CD for prompts/models/indexes, Docker, Kubernetes, registries, deployment.

8

Observability, Monitoring & Cost Management

Observability, tracing (OpenTelemetry), latency/token monitoring, drift detection and cost management.

9

Security, Safety & Governance

LLM security, prompt injection, PII redaction, GDPR/SOC 2, guardrails, red-teaming and governance.

10

Scaling, Multi-Model Systems & Capstone

Scaling, multi-model routing, fallback/resilience, agentic systems, HA/DR, and the capstone build.

**Capstone Project · Production-Ready LLM Application**

Design and build an end-to-end production LLM application — serving, a prompt/RAG pipeline, evaluation, observability, security guardrails and scaling — SME-reviewed and portfolio-ready. Plus 1-on-1 and daily expert sessions.

Module-by-Module

Module 1 — Foundations of LLMOps

- Introduction to LLMOps and its evolution from MLOps
- Characteristics of LLM systems: non-determinism, prompt sensitivity, token economics
- LLMOps lifecycle and operational workflows
- Key roles and responsibilities in LLMOps
- Build vs. buy decisions for LLM solutions
- Challenges in operationalizing LLM applications

Module 2 — Model Selection & Serving Architecture

- Proprietary vs. open-weight LLMs and licensing considerations
- LLM model selection criteria
- Inference serving architectures and frameworks
- vLLM, TGI, Triton, and Ray Serve
- GPU sizing and resource optimization
- Model quantization: INT8, INT4, GPTQ, and AWQ
- Latency, throughput, batching, and inference optimization

Module 3 — Prompt & Context Engineering

- Production-grade prompt engineering
- Prompt templates, versioning, and management
- System prompt design
- Structured outputs and JSON mode
- Function calling and tool use
- Context window management and optimization
- Chunking and context compression strategies
- Prompt testing and regression management

Module 4 — Retrieval-Augmented Generation (RAG)

- RAG architecture and core concepts
- Embedding models and vector representations
- Vector database selection and architecture
- pgvector, Pinecone, Weaviate, and Qdrant
- Document ingestion, chunking, and indexing
- Semantic, keyword, and hybrid search
- Re-ranking and retrieval optimization
- Data freshness and access control
- Common RAG challenges and failure modes

Module 5 — Fine-Tuning & Model Customization

- Prompting vs. RAG vs. fine-tuning
- Fine-tuning strategies and use cases
- Supervised Fine-Tuning (SFT)
- Parameter-Efficient Fine-Tuning (PEFT)
- LoRA and QLoRA
- Dataset curation and preparation
- Synthetic data generation
- RLHF and DPO fundamentals
- Fine-tuning evaluation and production considerations

Module-by-Module

Module 6 — Evaluation & Quality Assurance

- LLM evaluation principles and methodologies
- Golden datasets and evaluation harnesses
- Rubric-based evaluation
- LLM-as-a-Judge
- Measuring accuracy, relevance, and task success
- Hallucination and toxicity evaluation
- Offline and online evaluation
- A/B testing and shadow deployments
- Regression testing and continuous evaluation
- Automated quality gates for production deployments

Module 7 — Production Infrastructure & MLOps Integration

- Production architecture for LLM applications
- CI/CD for prompts, models, and RAG indexes
- Docker and containerization
- Kubernetes for LLM workloads
- KServe and Seldon
- Model registries and version management
- Deployment strategies for LLM applications
- Caching and performance optimization
- Integration of LLMOps with MLOps workflows

Module 8 — Observability, Monitoring & Cost Management

- LLM application observability fundamentals
- Logging and distributed tracing
- OpenTelemetry for LLM systems
- Monitoring latency, errors, throughput, and token usage
- Input and output quality monitoring
- Model and data drift detection
- Cost tracking and token-level cost analysis
- Model routing and cost optimization
- Alerting and incident management
- LLM performance optimization

Module 9 — Security, Safety & Governance

- LLM security fundamentals and threat landscape
- Prompt injection and jailbreak attacks
- Data leakage and insecure tool use
- PII detection and redaction
- Privacy and compliance considerations
- GDPR and SOC 2 requirements
- LLM guardrails and safety controls
- Access control and authentication
- Audit logging and secure context handling
- LLM red-teaming and security testing
- AI governance and risk management

Module 10 — Scaling, Multi-Model Systems & Capstone

- Scaling LLM applications for production
- Multi-model architecture and model routing
- Model fallback and resilience strategies
- Cost-aware model selection
- Agentic systems and production considerations
- Tool-calling reliability and orchestration
- Failure recovery and retry mechanisms
- Disaster recovery and high availability
- Platform team patterns for LLMOps
- End-to-end production readiness
- Capstone: design and build a production-ready LLM application

A Guided Path to Certification

Self-paced study, daily live sessions, personal mentoring, and hands-on practice — finish at your own pace.



STEP 1 • FOUNDATION

Self-paced modules

36+ hrs expert-led video, e-books, templates, cheat sheets & toolkits · lifetime access



STEP 2 • DAILY LIVE

GSDC Studio

Daily live sessions with global experts · 100+ monthly · free recordings



STEP 3 • PRACTICE

Learn by Doing + 1-on-1 SME

Build and deploy an LLM application for a real scenario, with personal SME reviews



STEP 4 • CERTIFICATION

Apply & Get Certified

Practice exams + capstone project, final proctored exam · portfolio-ready deliverables

Move at your own pace. Study self-paced with lifetime access to all materials, live sessions, and recordings.

ALL IN ONE ENROLLMENT

Everything Included



Expert-Led Videos

36+ hours, self-paced



SME Connect

3 personalized 1-on-1 sessions



Capstone Project

Build a production-ready LLM application



GSDC Studio

100+ live monthly sessions



Learn By Doing

Hands-on LLMops production scenarios



Job Support

LinkedIn Enhancer, Resume Builder + free membership

Exam & What to Expect

This is a certification, not a course. Study self-paced with the materials provided (plus live support), then sit a proctored online exam. Lifetime access to materials, plus 2 practice exams.

40

EXAM QUESTIONS

MCQ

EXAM FORMAT

65%

PASSING SCORE

90 min

DURATION

English

LANGUAGE

No

OPEN BOOK

5 Years

VALIDITY

Yes

FREE RETAKE

IS THIS FOR YOU?

Where CLLMOP Takes You

YOUR ROLE	HOW CLLMOP HELPS YOU	CAREER OUTCOME
ML / MLOps Engineer	Extend MLOps skills to LLM serving, evaluation and cost control.	→ LLMOps Engineer
Data Scientist / ML Engineer	Ship LLM features with RAG, fine-tuning and evaluation.	→ Applied LLM Engineer
Software / Backend Engineer	Build reliable, tool-using LLM applications in production.	→ LLM Application Engineer
DevOps / SRE / Cloud	Run GPU workloads, observability and guardrails at scale.	→ AI Platform / LLM Infra Engineer
AI Product / Eng Lead	Own quality, cost and safety across LLM products.	→ Head of LLM Engineering
Architect / Consultant	Design multi-model, governed, production LLM systems.	→ LLM Solutions Architect



7-Day Money-Back Guarantee

Enroll with zero risk — full 100% refund within 7 days if you feel you haven't gained the skills you're looking for. Includes a complimentary exam retake and lifetime access to your materials.

[Chat on WhatsApp](#)[Book a 10-min advisor call](#)

Choose Your Plan

Single Access

~~\$800~~ **\$400**

1

CERTIFICATION

- ✓ Self-paced expert-led videos + GSDC Studio
- ✓ 3 SME Connect (1-on-1) sessions
- ✓ Book of Knowledge + Capstone + Job Support
- ✓ Exam + 1 free retake & practice · membership (\$109 free)

Enroll Now

Bundle Access

MOST POPULAR

~~\$1200~~ **\$600**

3

CERTIFICATIONS

- ✓ Everything in Single, across 3 certifications — just \$200 each
- ✓ Unlimited SME Connect (1-on-1)
- ✓ Exam + 2 free retakes & practice
- ✓ Most-chosen by serious learners

Enroll Now

Region-specific pricing shown at enrollment · Teams? Ask about GSDC for Business.

Enroll Now → gsdcouncil.org/llm-ops-certification

WhatsApp +1 251-333-1717 · Global Skill Development Council · US · Singapore · Switzerland

Limited-Time Offer — Extra 10% Off

Use code **UPSKILL10** at checkout — valid for 72 hours only