

GLOBALLY RECOGNIZED · VENDOR-NEUTRAL · ANSI-ALIGNED

Certified Context Engineer

— (CCE)

Become a Globally Recognized Context Engineering Professional

Earn the CCE credential — gain the skills to design, build, and govern production-grade context systems for LLM applications, RAG pipelines, and AI agents, and command a **20–60% higher salary**.







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

Enroll at gsdcouncil.org/context-engineering-certification







Why This Certification, Why Now

AI teams are racing to move beyond clever prompts into production-grade context systems — while managing the cost, security, and reliability risks that come with them. This is the skillset that bridges both sides.

- **Prompting alone doesn't scale**
Production AI needs instructions, memory, retrieved knowledge, and tools managed as one coherent system — not one-off prompt tricks.
- **RAG and agents are now the default**
Enterprises are shipping copilots, RAG systems, and AI agents that live or die on the quality of their context.
- **Context is a new attack surface**
Prompt injection, data leakage, and tool misuse create risks that traditional AI training doesn't cover.
- **Leaders need governed AI, not just demos**
Businesses want AI systems that are grounded, auditable, and cost-controlled — not a black box.

FROM THEORY TO PRACTICE

What You'll Be Able to Do

-  Explain the difference between prompt engineering and context engineering in enterprise AI systems
-  Design layered context architectures using instructions, memory, retrieved knowledge, tools, and metadata
-  Build RAG-based context pipelines with chunking, embeddings, retrieval, ranking, and source attribution
-  Create memory-enabled AI workflows while applying privacy, governance, and retention controls
-  Implement tool calling, function calling, and MCP-based integrations for agentic AI applications
-  Evaluate context quality using groundedness, relevance, safety, cost, and observability metrics

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 working AI engineers, RAG-systems practitioners, and agentic-AI experts designing production context systems.

RECOMMENDED BY

Forbes · Indeed · TechTarget · CareerSidekick · LeanIX · People Managing People · Authentic



ABICB Accredited



ANSI Member



Vendor-Neutral

See our [Wall of Fame & verified reviews](#) → · trustpilot.com/review/gsdouncil.org · reviews.io

Our Global Clients

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

Oracle

Amazon

Microsoft

Cisco

Veritas

HP

Bank of America

Deloitte

Volkswagen

IS THIS FOR YOU?

Who Is This Certification For?

- AI engineers, prompt engineers, and GenAI developers
- Software developers building LLM applications and agents
- Data scientists and ML engineers working with RAG systems
- Solution architects and enterprise AI leaders
- Business automation teams using AI copilots and workflow agents
- Prompt engineers moving into context engineering and production AI design
- Anyone looking to formalize their skills in effective context engineering for AI agents

PRE-REQUISITES

No mandatory prerequisites. The program is designed for professionals with a basic understanding of generative AI, LLMs, and prompt engineering. Familiarity with APIs, JSON, and basic software workflows helps; beginner-level Python or JavaScript and awareness of cloud, databases, or enterprise architecture make the content easier to follow but are not required to get started.

Learn from Experts

Learn from experienced practitioners and industry leaders who design, update, and teach the program — bringing real-world context engineering and AI systems 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

36+ hours across the full context engineering lifecycle, with hands-on labs and real frameworks built into every module. Each module pairs videos and e-books with templates and Learn-By-Doing activities.

1

Foundations of Context Engineering

From isolated prompt writing to complete context system design

2

LLM Context Anatomy, Tokens & Attention Behavior

How models consume, prioritize, and degrade information in the context window

3

Prompt-to-Context Design Patterns

Repeatable context patterns that support reliable AI behavior

4

Context Architecture & Information Hierarchy

Architecting context as layered information rather than a flat prompt

5

Structured Context & Output Control

Predictable inputs and outputs for systems that require consistency

6

Retrieval-Augmented Context & RAG Fundamentals

Dynamically adding external knowledge into the model context

7

Advanced Retrieval, Ranking & Context Assembly

The right evidence, at the right time, in the right order

8

Memory Systems for AI Applications

Short-term and long-term memory for personalized, stateful AI

What You'll Learn

9

Context Compression & Token Optimization

Reduce token cost and noise while preserving critical information

10

Tool Use, Function Calling & Model Context Protocol

Connect models to tools, APIs, and enterprise systems safely

11

Agentic Workflows & Multi-Agent Context Orchestration

Context for autonomous agents performing multi-step tasks

12

Security, Guardrails & Context Governance

Protect AI systems from unsafe instructions, leakage, and misuse

13

Context Evaluation, Observability & Quality Metrics

Measure grounded, useful, and safe outputs with real metrics

14

Capstone — Production Context Engineering Blueprint

Design an enterprise-grade context engineering solution end-to-end

Module-by-Module

Module 1 — Foundations of Context Engineering

- What context engineering means in modern AI systems
- Prompt engineering vs context engineering vs AI application architecture
- The LLM application stack: model, prompt, context, tools, data, and workflow
- The six context inputs: instructions, user intent, history, retrieved knowledge, tools, and metadata
- The LLM-as-operating-system analogy: model as reasoning engine, context window as working memory
- Typical business use cases: assistants, copilots, RAG systems, agents, and enterprise automation

Module 2 — LLM Context Anatomy, Tokens, and Attention Behavior

- Tokens, context windows, and model input limits
- Instruction hierarchy: system, developer, user, retrieved, and tool messages
- Attention, recency bias, and position sensitivity
- Context overload, distraction, and context rot
- Latency, cost, and quality trade-offs in long-context applications
- Designing context for clarity, sufficiency, relevance, and economy

Module 3 — Prompt-to-Context Design Patterns

- Role, task, context, format, and constraint patterns
- Few-shot, example-driven, and counter-example prompting
- Chain-of-thought alternatives: reasoning scaffolds, rubrics, and structured deliberation
- Step-back, self-check, and critique-revision patterns
- Reusable prompt templates and prompt versioning
- Designing context blocks for different personas, domains, and tasks

Module 4 — Context Architecture and Information Hierarchy

- Context layers: policy, task, domain knowledge, user state, tools, and output rules
- Instruction priority and conflict resolution
- Context mapping: what the model must know, may know, and must not access
- Context isolation and quarantine for untrusted content
- Metadata-driven context assembly
- Context portability across models, teams, and workflows

Module 5 — Structured Context and Output Control

- Structured context blocks with headings, XML-like tags, and markdown sections
- JSON schema, function schemas, and typed outputs
- Format enforcement, validation, and retry loops
- Separating instructions, evidence, examples, and user data
- Designing context for APIs, workflows, and downstream systems
- Handling ambiguity, missing information, and escalation behavior

Module-by-Module

Module 6 — Retrieval-Augmented Context and RAG Fundamentals

- RAG as context construction rather than simple document search
- Document ingestion, cleaning, and normalization
- Chunking strategies: fixed-size, semantic, hierarchical, and sliding-window chunks
- Embeddings, vector search, and semantic similarity
- Vector databases and indexes: Pinecone, Weaviate, pgvector, FAISS, and similar tools
- Grounded response generation with citations and source traceability

Module 7 — Advanced Retrieval, Ranking, and Context Assembly

- Hybrid search: keyword, semantic, and metadata filtering
- Query transformation, expansion, and decomposition
- Re-ranking, relevance scoring, and evidence selection
- Context assembly pipelines and retrieval orchestration
- Handling noisy, duplicated, and contradictory sources
- Knowledge freshness, provenance, and confidence scoring

Module 8 — Memory Systems for AI Applications

- Conversation memory and session state
- Short-term vs long-term memory design
- Episodic, semantic, and procedural memory
- User preference stores and enterprise knowledge stores
- Memory write, update, retrieval, and deletion rules
- Privacy, consent, and safe handling of remembered information

Module 9 — Context Compression and Token Optimization

- Token budgeting and context cost estimation
- Sliding windows, rolling summaries, and state compaction
- Selective inclusion, truncation, and prioritization strategies
- Summarization, extractive compression, and learned compression
- Preventing loss of critical facts during summarization
- Measuring quality impact after compression

Module 10 — Tool Use, Function Calling, and Model Context Protocol

- Function calling and tool schemas
- Tool descriptions, argument design, and result handling
- Model Context Protocol concepts and use cases
- Building and consuming MCP servers for data, files, APIs, and developer tools
- Tool selection, tool routing, and tool execution constraints
- Securing tool access with authentication, authorization, and audit trails

Module-by-Module

Module 11 — Agentic Workflows and Multi-Agent Context Orchestration

- Agent loops: plan, act, observe, reflect, and revise
- Context routing for task-specific agents and sub-agents
- Multi-agent collaboration, delegation, and handoff patterns
- Workflow orchestration using frameworks such as LangGraph, AutoGen, and similar tools
- State management, checkpoints, and recovery in agent workflows
- Human-in-the-loop review for high-risk actions

Module 12 — Security, Guardrails, and Context Governance

- Prompt injection, jailbreaks, and indirect injection attacks
- Data leakage risks in RAG, memory, and tool-enabled systems
- Context isolation, input sanitization, and policy enforcement
- Guardrails for content, actions, tools, and data access
- Enterprise governance: ownership, approvals, change control, and auditability
- Responsible AI considerations for transparency, fairness, and compliance

Module 13 — Context Evaluation, Observability, and Quality Metrics

- Evaluation dimensions: relevance, sufficiency, faithfulness, provenance, safety, and usability
- Hallucination detection and groundedness testing
- Retrieval evaluation: precision, recall, MRR, and hit rate
- LLM-as-judge, rubric-based testing, and human review
- A/B testing for prompt and context versions
- Observability tools such as LangSmith, tracing, logs, and feedback loops

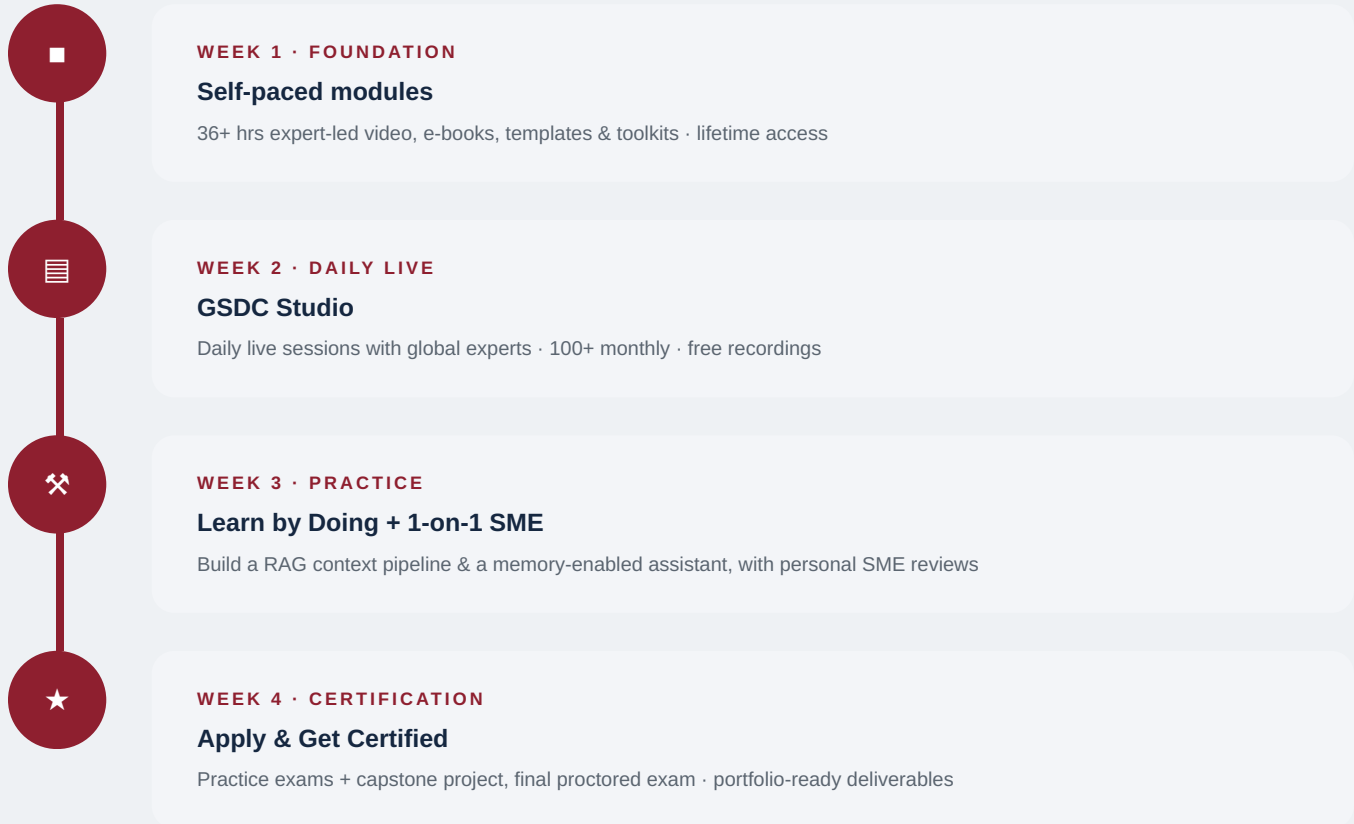


Capstone Project · Production Context Engineering Blueprint

— Define a business problem, design the context architecture, RAG pipeline, memory policy & tool interfaces, build security, governance & evaluation controls, and present a deployment-ready blueprint with cost, latency, and risk trade-offs — SME-reviewed, portfolio-ready. Plus 1-on-1 and daily expert sessions.

A Guided 4-Week Path to Certification







Self-paced study, daily live sessions, personal mentoring, and hands-on practice — finish in as little as one week.



Move at your own pace. This is a suggested roadmap — driven learners can finish in as little as one week.

ALL IN ONE ENROLLMENT

Everything Included

- | | |
|--|--|
|  <p>Expert-Led Videos
36+ hours, self-paced</p> |  <p>GSDC Studio
100+ live monthly sessions</p> |
|  <p>SME Connect
1-on-1 with industry experts</p> |  <p>Learn By Doing
Hands-on demos: RAG context pipeline & memory-enabled assistant</p> |
|  <p>Capstone Project
Design a production context engineering blueprint</p> |  <p>Job Support
Career help + free GSDC membership</p> |

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.

IS THIS FOR YOU?

Who Is This Certification For?


YOUR ROLE	HOW CCE HELPS YOU	CAREER OUTCOME
AI / Prompt Engineer	Move from one-off prompt tricks to layered, production-grade context systems.	→ Context Engineer / AI Engineer
Software Developer	Build LLM applications with RAG, memory, structured outputs, and tool calling.	→ LLM Application Developer
Data Scientist / ML Engineer	Design, ground, and evaluate retrieval pipelines with real quality metrics.	→ RAG / AI Platform Engineer
Solution Architect	Architect governed multi-agent context systems with proper guardrails.	→ AI Solution Architect
Automation / Copilot Team	Orchestrate workflow agents with reliable, cost-controlled context.	→ AI Workflow Lead
Risk & Compliance Professional	Apply guardrails, governance, and auditability to AI-driven systems.	→ AI Governance Lead


A professional-level program — built for practitioners already working with generative AI who want to formalize production-grade context design.



7-Day Money-Back Guarantee

Enroll with zero risk — full 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

1

CERTIFICATION

- ✓ Self-paced expert-led videos + GSDC Studio
- ✓ 3 SME Connect (1-on-1) sessions
- ✓ Capstone Project + Job Support + membership
- ✓ Exam + 1 free retake & practice

Bundle Access

MOST POPULAR

3

CERTIFICATIONS

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

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

Enroll Now → gsdcouncil.org/context-engineering-certification

WhatsApp [+1 251-333-1717](https://wa.me/9742513331717) · Global Skill Development Council · US · Singapore · Switzerland

LIMITED-TIME OFFER — EXTRA 10% OFF

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