DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Prompt Engineering Advanced Training

Design agentic systems, optimise and operate prompts in production, and govern LLM use with security, privacy and responsible AI controls.

LIVE ONLINE **Max 8-10 hours**	CORPORATE (DAY-WISE) **5 days · Days 1-5**	SELF-PACED VIDEO **1-year LMS access**
BATCH SIZE **Max 10 learners**	NEXT BATCHES **1st of every month · weekends: first Saturday**	CERTIFICATE **Verifiable certificate of completion**

Foundation	Essential	Intermediate	Advanced
1 day · Included	2 days · Included	3 days · Included	5 days · This curriculum

INSIDE THIS CURRICULUM

What you will achieve	Outcomes and level focus
Who should attend	Audience and prerequisites
Agenda	53 topics · Foundation to Advanced
Hands-on labs	4 labs
Formats and fees	Live, 1:1, self-paced, corporate
Trainer and reviews	Rajesh Kumar
Enrolment	Batches, payment, contacts

WHY DEVOPSSCHOOL

- Rated 4.6 on Google
- 1,500+ verifiable certificates issued in the public registry
- Trainer Rajesh Kumar: 20 years in industry
- Small batches: max 10 learners, every session recorded

Course page: www.devopsschool.com/courses/prompt-engineering/prompt-engineering-advance-training.html
Talk to us: WhatsApp +91 99057 40781 · +1 (469) 756-6329 · contact@DevOpsSchool.com · reply within 1 hour

OUTCOMES

What you will achieve

By the end of the Prompt Engineering course path, you will be able to:

1

Write reliable, effective prompts

Use structure, examples, reasoning and grounding to get accurate, consistent LLM outputs.

2

Build LLM-powered applications

Integrate prompts with APIs, tool calling, RAG and MCP, and evaluate them systematically.

3

Operate AI safely at scale

Run agents and prompts in production with security, governance and continuous evaluation.

WHAT THE ADVANCED LEVEL ADDS

Agents, optimisation, PromptOps and governance

Hands-on: Ship a governed AI agent. Build a tool-using agent for a team workflow, add evals, guardrails and tracing, and document its governance and rollback plan.

WHO SHOULD ATTEND

- Experienced Prompt Engineering practitioners responsible for quality, security and scale.
- Senior engineers, architects and leads who design and troubleshoot Prompt Engineering solutions.
- Teams making Prompt Engineering an organisation-wide standard.

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Integrate prompts into applications with LLM APIs, tool calling and RAG, then evaluate, secure and troubleshoot them systematically.
- Basic computer skills and a laptop with a stable internet connection for the live sessions and labs.
- This curriculum includes every earlier level, so it is complete on its own.

AGENDA · FOUNDATION TO ADVANCED

Prompt Engineering Advanced Training agenda

53 topics. Each level builds on the one before it, so this curriculum lists every earlier level in full. Days are for corporate (day-wise) delivery; live online batches cover the same agenda in max 8-10 hours.

DAY 1 — Foundation: Prompt engineering principles, tools and team practice

Included from Foundation

1. Introduction to prompt engineering: overview of principles and practices and how LLMs turn prompts into outputs
2. Why prompts matter: real-world examples and case studies showing the impact of good and poor prompts
3. How LLMs work: tokens, context windows, temperature, top_p and why outputs vary between runs
4. Key concepts and techniques: clear instructions, context, constraints, examples and an explicit output format
5. Understanding constraints and challenges: hallucination, knowledge cut-off, context limits, cost and latency
6. Problem-solving with prompts: decomposing tasks, step-by-step instructions and explicit decision criteria
7. Faster iteration: optimising prompt iteration cycles in playgrounds and minimising rework
8. Collaboration and communication: teamwork on prompts, shared prompt libraries, reviews and team conventions
9. Tools and technologies: ChatGPT, Claude, Gemini and Copilot, API playgrounds and automation tools
10. Version control for prompts: storing prompts in Git, naming, change history and reuse
11. Agile methodologies and prompt engineering: Agile principles, sprints and managing timelines and deliverables
12. Case studies and best practices: analysing successful prompt projects, common challenges and effective solutions
13. Continuous improvement and learning: feedback loops, evaluating prompt performance and ongoing skill development
14. Risk management: identifying risks, mitigation and contingency plans, balancing speed and quality

DAY 2 — Essential: Core prompting techniques and output control

Included from Essential

15. Prompt anatomy: role, task, context, constraints, examples and output format sections
16. System prompts and personas: setting behaviour, tone and boundaries for assistants
17. Zero-shot and few-shot prompting: choosing and formatting examples and avoiding example bias
18. Chain-of-thought and reasoning models: when to ask for step-by-step reasoning and when the model reasons itself
19. Delimiters and structure: XML tags, Markdown headings and quoting source material clearly
20. Structured outputs: JSON mode, JSON Schema-constrained responses, tables and data extraction
21. Prompt chaining: splitting work into sequential prompts and passing outputs between steps
22. Task patterns: summarisation, classification and extraction prompts with quality checks
23. Writing and rewriting: tone control, audience adaptation, style guides and length control
24. Code prompts: generating, explaining, reviewing and refactoring code and writing tests from specifications
25. Multimodal prompts: images, screenshots, PDFs and charts as input, and prompts for image generation
26. Reducing hallucination: grounding in provided sources, asking for citations and allowing 'I don't know'
27. Reusable prompt patterns: parameterised prompts, variables and a shared prompt library

DAY 3 — Intermediate: APIs, tools, RAG, evaluation and prompt security

Included from Intermediate

28. LLM APIs: calling OpenAI, Anthropic and Gemini APIs with messages, roles and generation parameters
29. Tool and function calling: defining tools with JSON Schema and handling tool results in the conversation
30. Retrieval-augmented generation: chunking, embeddings, vector search and grounding prompts in retrieved context

31. Context engineering: selecting, ordering and compressing context, long-context strategies and prompt caching
32. Model Context Protocol: connecting assistants to tools and data through MCP servers
33. Frameworks and prompt management: LangChain, LlamaIndex, LangSmith or Langfuse
34. Evaluation: golden datasets, rubric scoring, LLM-as-judge and regression tests with promptfoo
35. Prompt injection and jailbreaks: direct and indirect injection and the OWASP Top 10 for LLM Applications
36. Guardrails: input and output filtering, PII redaction, refusal handling and moderation APIs
37. Self-consistency and verification: generating several answers, voting, self-critique and checker prompts
38. Cost and latency: token budgeting, model selection, batching, caching and streaming
39. Open-weight models: prompting Llama, Mistral or Qwen locally with Ollama and model chat formats
40. Troubleshooting prompts: ambiguity, format drift, overlong context and inconsistent outputs

DAYS 4-5 **Advanced: Agents, optimisation, PromptOps and governance** **New at this level**

41. Agentic workflows: ReAct, planner-executor and multi-agent patterns with tool use
42. Agent frameworks: Claude Agent SDK, OpenAI Agents SDK and LangGraph, with human-in-the-loop checkpoints
43. Automated prompt optimisation: DSPy signatures and optimisers, meta-prompting
44. Fine-tuning vs prompting: when to fine-tune, distillation and prompt-only alternatives
45. Prompt operations: versioning, A/B testing prompts, staged rollout and rollback, tracing and observability
46. Evaluation at scale: offline and online evals, human review queues and drift monitoring
47. Security hardening: least-privilege tools, sandboxing, output validation and LLM red teaming
48. Governance and responsible AI: usage policies, bias testing, transparency, EU AI Act and NIST AI RMF awareness
49. Data privacy: handling confidential data, enterprise data controls and retention settings
50. Domain applications: incident summaries and runbooks for DevOps, customer support, analytics and documentation
51. Model migration: porting prompts across models and versions with regression suites
52. Enterprise enablement: prompt standards, a Centre of Excellence and reusable prompt catalogues
53. Capstone: design, evaluate and govern an AI agent that automates a real team workflow

PRACTICE

Hands-on labs

You build these yourself during the course, guided by the trainer.

FOUNDATION **Improve prompts for real tasks**
In hands-on exercises and group discussion, rewrite weak prompts for your own work scenarios and compare the outputs before and after.

ESSENTIAL **Build a team prompt library**
Create tested, parameterised prompts for summarising, extracting to JSON, drafting and code review, with guidance on when to use each.

INTERMEDIATE **Build and evaluate a RAG assistant**
Build a document Q&A assistant with retrieval and tool calling, then measure it with a golden dataset and harden it against prompt injection.

ADVANCED

Ship a governed AI agent

Build a tool-using agent for a team workflow, add evals, guardrails and tracing, and document its governance and rollback plan.

HOW YOU CAN LEARN

Formats and fees

The same agenda in four formats. Fees are in Indian rupees as of November 2026; the course page shows them in your currency.

FORMAT	DURATION	WHAT YOU GET	FEE
Live & interactive (online)	Max 8-10 hours	Instructor-led public batch of max 10 learners with demos, labs, Q&A and assignment review. Weekday batches start on the 1st of every month, weekend batches on the first Saturday of every month.	**INR 34,999**
1:1 live training (online)	Max 8-10 hours, one-to-one	The same agenda, taught one-to-one by the trainer. Starts immediately, without waiting for the next batch; sessions planned around your working hours, at your pace.	**INR 99,999**
Self-paced video	1-year LMS access	All four Prompt Engineering levels, plus 20+ more courses and 50+ tools, at your own pace.	**INR 833/month** billed yearly at INR 9,996
Corporate (day-wise)	5 days for this level	A private batch for your team, on demand. Foundation 1 day, Essential 2, Intermediate 3, Advanced 5 (cumulative).	**Quote on request**

ASSESSMENT AND CERTIFICATE

Your progress is checked through hands-on labs and assignments at every level. On completion you receive a DevOpsSchool certificate of completion with a unique credential ID that anyone can verify at devopsschool.com/certificates.

INCLUDED WITH EVERY LIVE BATCH

- Every live session is recorded and stays in your LMS for a year.
- Can't attend your batch? Move to next month's batch.
- Free demo class before you pay.
- Free 10-15 minute call with an advisor or the trainer to check the course fits you.

YOUR TRAINER

Learn from Rajesh Kumar

Principal Architect, trainer and mentor in DevOps, SRE, DevSecOps and platform engineering, with 20 years across 8+ software MNCs, including Staff DevOps/SRE Engineer for a leading Japanese mobile-payment platform and Principal DevOps Architect and Manager in Europe.

Profile: www.rajeshkumar.xyz · Videos: youtube.com/@TheDevOpsSchool

20 years in industry	**200+** corporate trainings
10,000+ engineers mentored	**70+** organisations coached

REVIEWS

What learners say

Rated 4.6 on Google · more reviews at devopsschool.com/testimonials and in our video testimonials.

> “Rajesh is very good trainer. Rajesh was able to resolve our queries and question effectively. We really liked the hands-on examples covered during this training program.”
>
> Indrayani, India

> “The training was very useful and interactive. Rajesh helped develop the confidence of all.”
>
> Abhinav Gupta, Pune

START YOUR PROMPT ENGINEERING TRAINING

Join the next batch

Weekday batches start on the 1st of every month and weekend batches on the first Saturday of every month, with max 10 learners each. 1:1 training can start immediately; corporate batches run on demand. We reply within 1 hour.

Pay by UPI, credit and debit cards, net banking or PayPal. Before you pay, read our refund policy. All levels of this course: www.devopsschool.com/courses/prompt-engineering/

ENROL OR RESERVE A SEAT
Prompt Engineering Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com