

Claude & Agentic AI Masterclass

The Claude API, MCP servers and agents you can actually put in production

Claude API

MCP

Agent SDK

Tool use

Bedrock

Trainer Venu Katragadda · Live online · Recordings for life

₹30,000

COURSE SYLLABUS

Claude & Agentic AI Masterclass

A full engineering course built around one model family — the API, the protocol, the agent patterns and the production economics.

100 hours LIVE INSTRUCTION	10 modules STRUCTURED PATH	3 projects PORTFOLIO READY	₹30,000 EMI AVAILABLE
-------------------------------	-------------------------------	-------------------------------	---------------------------------

Level	Intermediate to Advanced
Mode	Live, instructor-led online sessions. Every session recorded; lifetime access.
Trainer	Venu Katragadda — 14+ years in industry (L&T Infotech, Capgemini, Brillio, Manthan Systems). 1,200+ professionals trained.
Batch timings	Morning 6:00–7:30 AM IST · Morning 7:30–9:00 AM IST · Evening 7:30–9:00 PM IST · Weekend (Sat & Sun)
Certification	Course completion certificate · portfolio of 3 shipped agents
Enquiries	+91-8500002025 / +91-9247159150 (WhatsApp) · info@databrickstraining.com

What you will be able to do

Every outcome below is demonstrated in a lab or a graded project — not just discussed in a slide.

- Use the Claude Messages API fluently: system prompts, tool use, streaming, structured output, vision, files
- Cut cost 60–90% with prompt caching, the Batch API and model selection between Haiku/Sonnet/Opus tiers
- Design and build MCP servers that expose your own tools, resources and prompts to any MCP client
- Build multi-step agents with the Agent SDK: subagents, memory, context compaction, permissions
- Use Claude Code as an engineering tool: hooks, slash commands, plugins, CI automation
- Deploy on the Anthropic API, Amazon Bedrock and Google Vertex AI with proper security and observability

Who this masterclass is for

- Engineers building on Claude
- Teams standardising on MCP for internal tooling
- GenAI engineers who want agent depth
- Developer-experience and platform teams

Prerequisites

- Solid Python (TypeScript examples also provided)
- Basic REST/API experience
- Anthropic API key or Bedrock/Vertex access

Tools and technologies

Claude API · Claude Code · Claude Agent SDK · Model Context Protocol (MCP) · Python SDK · TypeScript SDK · Amazon Bedrock · Google Vertex AI · FastAPI · LangGraph · OpenTelemetry · Docker · GitHub Actions

Curriculum — 100 hours across 10 modules

Sessions run live in small batches. Roughly 60% of each session is hands-on lab work.

MODULE 1 Claude fundamentals and the Messages API

10 hrs

- Model family and tiers: choosing between fast/balanced/most-capable for a given task
- Messages API: roles, system prompts, multi-turn state, stop reasons, token counting
- Streaming, retries, idempotency, rate limits and error taxonomy
- Vision and document input; the Files API; PDF handling
- SDKs in Python and TypeScript; async patterns
- Lab: production-ready API client with retry, timeout and cost logging

MODULE 2 Prompt design for Claude

10 hrs

- System prompt architecture: role, constraints, output contract, examples
- XML tags, long-context organisation, document-first prompting
- Extended thinking: budgets, when it pays for itself, interaction with tools
- Structured output: JSON schemas, tool-based extraction, validation loops
- Prefill, stop sequences, and controlling verbosity and format
- Lab: prompt regression suite with 40 golden cases

MODULE 3 Tool use and function calling

10 hrs

- Tool schemas, tool_choice, parallel tool calls, tool results and errors
- Designing tools an LLM can actually use: naming, granularity, idempotency
- Server-side tools: web search, code execution, computer use — capabilities and limits
- Client-side orchestration loops; the tool-runner pattern
- Failure handling: timeouts, partial results, retries, human escalation
- Lab: a 6-tool assistant with deterministic fallbacks

MODULE 4 Cost, latency and reliability engineering

10 hrs

- Prompt caching: cache breakpoints, TTLs, hit-rate measurement, cache-aware prompt layout
- Batch API for offline workloads; throughput planning
- Model cascades and routing; when a cheap model plus verification beats one expensive call
- Token budgets, context compaction, summarisation strategies for long sessions
- Load testing, SLOs, graceful degradation
- Lab: take a workload from \$0.31 to \$0.04 per request with measurements

MODULE 5 Model Context Protocol (MCP) in depth

12 hrs

- Why MCP exists; architecture: hosts, clients, servers, transports (stdio, HTTP/SSE)
- Server primitives: tools, resources, prompts, sampling, roots, elicitation
- Building a server in Python and TypeScript; schema design and validation
- Auth patterns: OAuth, tokens, per-user scoping, multi-tenant servers
- Security: tool poisoning, confused deputy, prompt injection through resources, allowlists
- Testing MCP servers; the inspector; publishing and versioning
- Lab: Project 2 — ship a real MCP server with an auth story

MODULE 6 Agents with the Claude Agent SDK

12 hrs

- Agent loop anatomy: gather context, act, verify, repeat
- Subagents and delegation; when parallel agents help and when they thrash
- Memory: session state, files-as-memory, compaction, retrieval-backed memory
- Permission systems, sandboxing, allow/deny rules, dry-run modes
- Hooks and lifecycle events for policy enforcement and logging
- Long-running and scheduled agents; resumability
- Lab: Project 3 — autonomous workflow with approval gates

MODULE 7 Claude Code as an engineering platform

10 hrs

- CLI workflows: planning mode, diffs, checkpoints, worktrees
- CLAUDE.md, project context, skills, slash commands, subagent definitions
- Hooks for tests, linting, secret scanning and policy checks
- Plugins and marketplaces; distributing tooling to a team
- Headless mode in CI: automated review, migration and refactor jobs
- Lab: build a team plugin that encodes your engineering standards

MODULE 8 RAG and data integration with Claude

10 hrs

- Long context vs retrieval — measured trade-offs at 200K+ tokens
- Contextual retrieval; citation-first answering; grounding checks
- Connecting Claude to warehouses: Databricks, Snowflake, BigQuery via MCP or tools
- Permission-aware retrieval and row-level security in agent contexts
- Lab: Project 1 knowledge agent

MODULE 9 Deployment, security and governance

8 hrs

- Anthropic API vs Amazon Bedrock vs Google Vertex AI: features, regions, pricing, IAM
- Secrets, key rotation, per-tenant isolation, egress control
- OWASP LLM Top 10 applied to agentic systems; blast-radius design
- Observability: OpenTelemetry traces, token dashboards, eval-in-production
- Usage policies, data retention, and what not to send to any model
- Lab: deploy the knowledge agent on Bedrock with private networking

MODULE 10 Evaluation, capstone and career

8 hrs

- Evaluating agents: task success rate, trajectory analysis, cost per resolved task
- LLM-as-judge rubrics for agent output; regression suites for tool behaviour
- Red-teaming agents: injection through tool results, over-permissioned tools, runaway loops
- Capstone presentation and architecture defence
- Agentic-AI interview questions, portfolio framing and resume rewrite
- Keeping up: reading changelogs, deprecation notices and the MCP spec

Hands-on projects

You leave with work you can demo in an interview. Each project is code-reviewed and returned with written feedback.

PROJECT 1 Internal knowledge agent

A Claude-powered assistant over your company's documents with citations, permission-aware retrieval, prompt caching and a full trace/eval dashboard.

PROJECT 2 Custom MCP server

Build, test and publish an MCP server exposing a real system (database, ticketing, internal API) with auth, rate limits and a tool-safety review.

PROJECT 3 Autonomous engineering workflow

A Claude Code + Agent SDK pipeline that triages issues, proposes patches, runs tests and opens PRs — with human approval gates and blast-radius limits.

Frequently asked questions

Is this an Anthropic-official course?

No. It is an independent training programme built on Anthropic's public documentation and SDKs. We are not affiliated with or endorsed by Anthropic, and Claude is a trademark of Anthropic.

Why a whole course on one model family?

Because agentic engineering — MCP, tool design, permissioning, context management — is deep enough to deserve it, and the patterns transfer. If you want breadth across providers, take the Generative AI course; many students do both.

Do the APIs change during the course?

Constantly. Content is refreshed every batch and there is a standing 'what changed' segment. Deprecated parameters and models are flagged explicitly.

What does the API usage cost?

Roughly \$25–\$50 across the course. Prompt caching and the small model tier are taught in Module 4 specifically so your own bills stay low.

How to enrol

- Book a free demo. Sit in on a live session with the current batch — no card, no commitment. WhatsApp +91-9247159150 or call +91-8500002025.
- Pick your batch. Weekday morning, weekday evening or weekend. All times IST.
- Pay the fee (₹30,000, instalments and EMI available) and get your workspace, material and calendar invites.
- Start building. Labs from session one; recordings posted the same day.

What is included

- 100 hours of live instruction by Venu Katragadda
- Lifetime access to all session recordings
- 3 code-reviewed portfolio projects
- Preparation for: Course completion certificate · portfolio of 3 shipped agents
- Resume review, LinkedIn optimisation and one mock interview with written feedback
- Free re-attend of any future batch of the same masterclass
- Doubt-clearing sessions and a private student group

Refund policy in brief

When you cancel	Refund
Before the batch starts	Up to 90%
Within 5 days of start, max 2 classes attended	Up to 70%
After 7 days, or more than 5 classes attended	No refund — free transfer to a later batch

Full policy: databrickstraining.com/refund-policy/

Talk to us

Call	+91-8500002025 · +91-9247159150
WhatsApp	+91-9247159150 — usually answered within the hour
Email	info@databrickstraining.com
Website	databrickstraining.com/courses/claude-agentic-ai-masterclass/
Foundation tier	databrickstraining.in — shorter 70–80 hour courses if this one is too advanced
Institute	Sreyobhilashi IT · Udyam UDYAM-TS-02-0334643 · 112, Samshraya Homes, opp. PF Office, Balanagar, Kukatpally, Hyderabad, Telangana 500037

Sreyobhilashi IT is an independent training provider and is not affiliated with, endorsed by or sponsored by Databricks, Amazon, Microsoft, Google, Snowflake, dbt Labs or Anthropic. Placement assistance is a best-effort service, not a guarantee of employment. Cloud platforms change frequently; confirm current pricing and service behaviour with the vendor's documentation.