



Pathfinder Research and
Training Foundation


POWERED
BY



Claude 101

Building the Modern AI Stack

Five working systems. One certificate that says you built them.



DATES

**Starting
24th August**

DURATION

25+ hours

VENUE

**PRTF
GREATER
NOIDA**

TOTAL SEATS

50 only

05

SYSTEMS YOU BUILD AND
KEEP

01

CO-BRANDED
CERTIFICATE, PRTF x
CLOUDTHAT


50

SEATS IN THE FOUNDING
BATCH

25+

HOURS OF LIVE
INSTRUCTION

Taught on campus, on your own laptop, with your hands on the keyboard for most of the twenty-five hours. Every session is a build session; you leave each day with a system that ran on your own machine.



PRTF CAMPUS, GREATER NOIDA

EVERYONE IN YOUR BATCH HAS A GENERATIVE AI CERTIFICATE.

Almost none of them have shipped an agent.

There is a difference between having used an AI tool and having built one. The first is now assumed. The second is still rare enough to be worth interviewing for. Over five days you will write the retrieval, the tool calls, the evaluation harness and the agent loop yourself, and leave with the code.



DAY 03, TOOL USE • PRTF CAMPUS

~40%

OF ENTERPRISE LLM
SPEND

~54%

OF THE AI CODING
MARKET

70%

OF THE FORTUNE 100
ARE CUSTOMERS


You are not learning a chatbot. You are learning the system that enterprises are actually buying.

LEVEL I

A multi-tool assistant

You define tool schemas, handle tool-call message blocks, return results to the model and run a multi-turn tool conversation. By the end of the session the assistant decides for itself which tool to reach for.

- SPECIMEN I • function calling, tool schemas, multi-turn tool loops




LEVEL II

A retrieval pipeline over a real document corpus

You chunk a PDF corpus, embed it, run vector search against BM25 lexical search, then combine both into hybrid retrieval and measure which answers improve. This is the single most requested capability in Indian AI job postings and most candidates have only read about it.

- SPECIMEN II • chunking, embeddings, BM25, hybrid retrieval, contextual retrieval




LEVEL III

An evaluation harness

You generate a test set, run responses against it, and grade them with a second model call and with code. Anyone can produce an output. Being able to prove an output is good is what separates a demo from a system.

- SPECIMEN III • test set generation, model-graded and code-graded evals




LEVEL IV

An MCP server and its client

You build a Model Context Protocol server that exposes a filesystem tool and a data resource, inspect it, then write the Python client that calls it. This is the protocol the industry standardised on for connecting models to real systems.

- SPECIMEN IV • MCP server, resources, prompts, inspector, Python client




LEVEL V

An end-to-end agentic system

The capstone. Retrieval, tools and MCP combined into one agent that researches, decides and reports, then presented to the room for review.

- SPECIMEN V • chaining, routing, parallelisation, sub-agents, capstone review



THE COMPLETE CURRICULUM

The full module list, taught hands-on across **25 + hours**. Every day closes with a workshop, and every lab is written on your own machine —twenty-two of them across the week.

FOUNDATIONS

■ HANDS-ON LAB

Drive Claude Code from VS Code and the terminal

5 MODULES · 5H

Claude and Anthropic

Safety principles, the Haiku, Sonnet and Opus family, Constitutional AI, and where each model earns its place.

How Claude works

Tokens, context windows, system prompts, temperature, and the difference between streaming and batch.

Claude Code 101

An agent rather than a chat window: terminal, VS Code and JetBrains, the agentic loop, permissions, and Plan Mode.

The AI fluency framework

Discover, discern, develop, deploy — fluency against literacy, and what ethical use looks like in practice.

Review and workshop

Choosing model and mode, then first prompts written in both Claude.ai and Claude Code.

THE API AND PROMPTING

■ HANDS-ON LAB

Build a multi-turn chat loop, then an evaluated Q-and-A assistant

5 MODULES · 5H

Accessing the API

The console, key hygiene, curl and Python, and the three parameters every call needs.

System prompts and control

Personas, temperature, top-p and top-k, streaming responses, and forcing clean JSON output.

Prompt engineering

XML structure, few-shot and zero-shot, chain of thought, constraints, and the anti-patterns to unlearn.

Evaluation workflows

Test set generation, automated runs, and grading by model or by code.

Workshop

Prompt, craft and evals combined into a single working build.

The second half moves from single calls to systems: retrieval, the Model Context Protocol, agent architecture, governance, and the capstone you present on the final afternoon.

TOOL USE AND RETRIEVAL

■ HANDS-ON LAB

Run retrieval over a real PDF corpus with live tools attached

5 MODULES · 5H

MODEL CONTEXT PROTOCOL AND AGENTS

■ HANDS-ON LAB

Ship an MCP server, then the client that drives it

5 MODULES · 5H

GOVERNANCE AND CAPSTONE

■ HANDS-ON LAB

Design and present a full agentic system for review

5 MODULES · 5H

Tool use

Schemas, tool-call blocks, returning results, multi-turn loops, and the built-in tools.

Retrieval

Chunking, embeddings and vector search, BM25, hybrid indexes, and contextual retrieval.

Advanced features

Extended thinking, vision, PDF processing, citations, and the Files API.

Prompt caching

When to cache, the constraints, what it does to cost and latency, and how to read hit rates.

Workshop

Retrieval and live tools together in one agentic assistant.

Introduction to MCP

Why the protocol exists, how servers, clients and hosts fit together, and when it beats direct tool use.

Building a server

Setup, tool schemas, the inspector, resources and prompts.

Building a client

A Python client that connects, invokes tools and reads resources.

Agents and workflows

Agent against chatbot: parallelisation, chaining, routing, and when each is the right shape.

Workshop

An agent designed end to end, reviewed by the group.

Claude Code workflows

Explore, plan, code, commit — plus /compact, /context, CLAUDE.md and sub-agents.

Skills, hooks and MCP

Attaching servers, skills as bundles, hooks for control, and Computer Use.

AI fluency, applied

Course and workflow design, academic integrity, and institutional policy.

Security and compliance

Secrets management, rate limits and cost, data privacy, and audit logging.

Capstone and wrap-up

Final assessment: API, MCP and Claude Code combined, presented by teams.

Twenty-five modules, twenty-two labs, one capstone.

THE STACK

THIRTY-FIVE TECHNIQUES • TWENTY-FIVE HOURS

FOUNDATIONS tokens • context windows • system prompts • temperature • top-p and top-k • streaming vs batch

PROMPT CRAFT XML-structured prompts • few-shot and zero-shot • chain of thought • negative examples and constraints • structured JSON output • anti-patterns

RETRIEVAL chunking strategies • embeddings and vector search • BM25 • hybrid multi-index retrieval • contextual retrieval

AGENTS AND MCP tool schemas • multi-turn tool conversations • MCP servers • MCP clients • the inspector • chaining • routing • parallelisation • sub-agents • hooks • CLAUDE.md • Computer Use

PRODUCTION eval harnesses • model-graded and code-graded scoring • prompt caching • cost and rate-limit management • API key hygiene • audit logging

Every item on this page is taught with your hands on the keyboard, not on a slide.

Four things, and you are ready.

☐	PYTHON	You can write a Python function and read an error message without panicking.
☐	A USER, SO FAR	You have used ChatGPT or Claude, but you have never called an API.
☐	A LAPTOP	You have a laptop you are allowed to install things on.
☐	PROOF	You want proof of what you can do, not vocabulary to describe it.

NOT REQUIRED: machine learning theory · a GPU · prior cloud experience · a computer science degree

WHO TEACHES IT, AND WHAT YOU LEAVE WITH

A certified Claude instructor from CloudThat. Delivered by CloudThat, a corporate training organisation whose certifications are globally recognised as well as across the Indian technology industry.

THE CERTIFICATE

A co-branded certificate of completion issued jointly by PRTF School of Management, an AICTE-approved institution, and CloudThat. It records the program, the twenty-five plus hours, and the systems you built.

The part after the certificate.

Building the five systems is the point. Being seen for them is the next step — and that does not end on day five. PRTF's placement desk carries your work forward to the people who hire.

01 Your five builds, packaged

A reviewed repository and a one-page project sheet an interviewer can read in two minutes — written with you on the final afternoon.

02 Mock technical interviews

Three sessions on the exact material you wrote — retrieval, tool use, evaluation and MCP — with written feedback after each one.

03 Internship and job referrals

Your profile goes to PRTF's hiring network and to CloudThat's corporate clients as AI engineering and internship roles open.

04 On-campus interview drives

Shortlisted participants are put in front of visiting recruiters at the drives PRTF runs through the year.

MOCK INTERVIEWS	REFERRAL NETWORK	SUPPORT WINDOW
3 included	PRTF × CLOUDTHAT	12 months

SAID PLAINLY

This is placement assistance, not a placement guarantee. No program can promise an offer. What this one does is make sure the systems you built are in front of the people who hire for them.



THE INFRASTRUCTURE

Where you will build.

The sessions run on PRTF's own campus in Greater Noida, and for the five days all of it is open to you: the lab, the classrooms, the library and the campus around them.



■ ACCESS TO ALL



Participants have access to all of it for the duration of the program.

THE PRICE

independent
benchmarks for
equivalent authorized
training

~~₹25,000~~

Almost all in-person costs
about the same and gives you
**no instructor, no labs, no
room and no certificate.**

FOUNDING BATCH FEE

₹11,999

Near 30 hours in person on
campus, all five build sessions,
the complete codebase you
write, the co-branded
certificate, lunch and
refreshments, printed
materials, and API access for
the labs.

HOW TO REGISTER

01

Reserve a seat at
pathfinderfoundation.co.in,
or call the office.

02

Pay the fee to
confirm.

03

Setup instructions
and the reading list
arrive a week before
day one.

REGISTER

pathfinderfoundation.co.in

CONTACT

+91
8860005458

CAMPUS

PRTF,
GREATER
NOIDA

Straight answers.

I am in my second year. Is it too early?

No, if you can already write Python, used LLMs before and can bring a laptop then the year does not matter.

Will this get me a job?

No program can promise that. What it can do is give you five systems to talk about when everyone else is describing a course they watched.

What if I miss a day?

Each day builds on the last, so attendance matters. If you miss a day we share that day's code and lab notes, and you can sit the same day with the next batch at no cost.

What is the refund policy?

A full refund on cancellation up to ten days before day one, a half refund up to five days before, and none after that.

Will the API key be provided?

Yes, all participants will be provided the claude API key, this is included in the program fees itself.



Pathfinder Research and Training Foundation

30/7, KP-III, GREATER NOIDA, NCR

+91 88600 05458 • PATHFINDERFOUNDATION.CO.IN

DELIVERED WITH CLOUDTHAT • CLOUDTHAT.COM

SCAN TO ENROL



pathfinderfoundation.co.in