

LUMINARY LIVE SESSION

Build And Launch A *Customer-Facing* MCP

Thursday 20 August 2026

What We'll Cover

- 01 What An MCP Actually Is
- 02 How Do You Build An MCP
- 03 Live Review
- 04 Q&A



Your Hosts



Akhil Raj Gupta

Founder at Luminary, ex Airbnb



Bharat Batra

AI Product Lead at Superhuman

01



What An MCP Actually Is

USB-C for AI



Model

The agent doing the work — Claude, ChatGPT, Cursor.



Context

Everything the model can see.



Protocol

An open standard for putting your product in there.

Open-sourced by Anthropic, 25 Nov 2024.

Its standing today — a universal way for an external product to plug into any AI agent.

The Evolution Of MCPs



Nov 2024

Open sourced by Anthropic

You ran it on your own machine, with your own keys in a config file.



Mar 2025

Remote servers arrive, with OAuth

Until now it was local and ungoverned. Now a company could host one, and customers could sign in.



May 2025

Official Claude Integrations launch

Ten shipped on day one — Atlassian, Linear, Stripe, PayPal, Asana, Sentry.



Today

A growth lever for SaaS

~10,000 published in the official registry. Almost every major software company has one.



Under The Hood

THE CLIENT SIDE 

The model's context

The menu of tools arrives once, on connect — the model reads it like documentation.

The model

Picks one tool, fills in the arguments, sends a request — then reads what comes back and continues.


At startup, the app opens a connection



The menu travels into the context



The model picks a tool and calls it



The result returns; the agent continues



 THE SERVER SIDE

Written instructions

In the server's code — one entry per tool.

```
name: "get_campaign_performance"  
description: "Retrieves performance metrics  
for campaigns... The primary tool for daily  
campaign monitoring and optimization  
decisions."
```

Product copy, written for a robot.



The machinery

The MCP server, holding credentials — wired to your existing REST API.

02



How Do You Build An MCP



What Actually Gets Built

A list of tools

Names, descriptions, and what each one takes. This is where the thinking goes.

A thin server

Wraps the API you already have.

**Not one tool per endpoint —
one tool per job.**

Sign-in

OAuth, so the agent acts as that customer — with their permissions, not yours.

A listing

A registry entry and a one-click install, so agents can find you.

Three of these are plumbing. **The tools and their descriptions are the product work.**

Scaffolding Tools — What To Keep In Mind

Running example: an MCP for an online store.

PRINCIPLE	 AVOID	 AIM FOR
Naming Agents use ToolSearch to load tools into context, so names matter a lot more than we think	get · fetch · lookup	search_orders · search_customers
Descriptions Common pitfalls include lack of consistency or clarity. E.g - Same parameter defined differently in different tools	search_orders → “the ref” get_order → “order number”	both → “order ID, e.g. ORD-4821”
Actionable responses Find the most common tool paths and make sure each one returns outputs and errors an agent can actually act on.	400 Bad Request	refund \$42 exceeds order total \$31.50
Tool composition Combine tools to cut latency and context consumption — huge lift in the customer experience	get_customer → list_orders → get_order	search_orders(customer_email)

Existing REST API vs Agent-Specific APIs

Wrap the API you have

One tool for every endpoint you already expose. Fastest path to shipping.

POST /v2/calendar/events

The agent inherits your API's shape — not the customer's job.

Build agent-specific APIs

One tool for every job a customer actually asks for.

book_a_meeting

Atomic

Markdown

Async when needed

Agent ergonomics

03

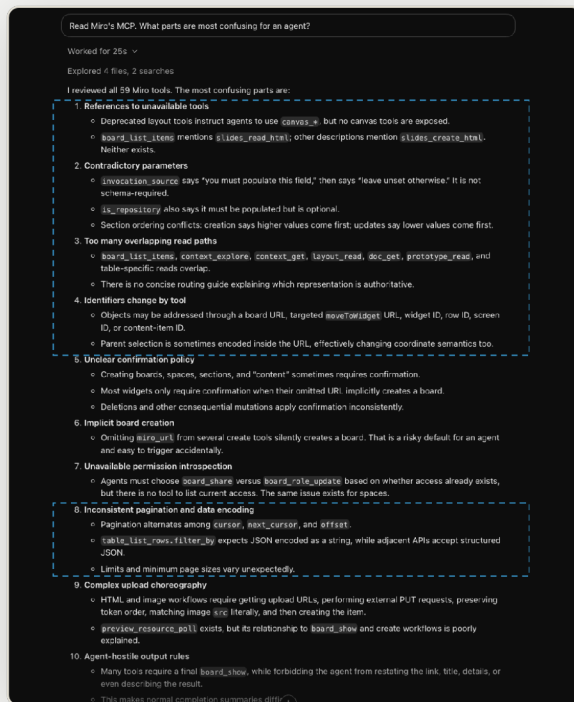


Live Review

Let's Open Up Some MCPs In Claude Code



Getting Feedback On Your MCP From The End User





A Quick Word From Our *Sponsors*

It's us. We're the sponsors. 😊

AI PM: Ship Your First Customer-Facing Agent

01 Make It Work

Take a real customer-facing agent apart — the loop, tool calls, context, memory, the system prompt. Leave week one with your own v0 running.

02 Make It Yours

Turn it into a harness that beats a generic assistant in your environment — proprietary context, UX, and tools a horizontal agent can't match.


03 Make It Reliable

Build an eval set that catches failures the way your hardest user would, and a playbook for choosing the right fix.

04 Make It Safe

Guardrails, observability, and a rollout that contains the blast radius — plus your day-zero launch checklist.

Plus a Demo Day capstone, a bonus module on background agents, a guest lecture on solving agent quality — and a **private 60-minute 1:1**.



Starts 11 Sep 4 weeks ~6 hrs / week \$750 (₹70,000)

Register at theluminary.ai or on Maven — use code

LIVESESSION10

for an extra 10% off the founding price.

04



Q&A