

# Rohit D. Tavare

rohit.d.tavare@gmail.com  
github.com/rohittavare  
/in/rohittavare  
rtavare.com  
(408) 547-7779

## Experience

### Production Engineer Meta Menlo Park, CA

Sep 2021 - Present

Python, Claude, OpenClaw, Thrift, SQL, Presto, Tupperware, C/C++

#### AI-Related Contributions

- ▶ Automated code revision & iteration - intelligently respond & integrate code review & CI signals. Direct authoring & test plans from the code review window.
- ▶ Automated alert investigator
  - Aggregate signals across multiple datasets, services & teams
  - Propose root cause & next steps for the human investigator
  - Self-learning through human feedback & documented manual investigations
- ▶ Plan, prioritize and execute entire projects & code migrations E2E using AI agents
  - navigate around custom client code, prioritize impact & assess risk across a scope too large of a human
  - Build & execute complex multi-step development, iteration, testing & rollout pipelines to automate entire projects

#### Capacity Turn Ups Aug 2024 - Present

- ▶ Building and scaling frameworks to automate the fulfillment of 80% Resources in the Meta Capacity Planning (MCP) flywheel - Data-driven workflow engine, automated validation & guardrails framework, E2E system telemetry, explainability & escalation tooling for blocked turn-ups.
- ▶ Work with multiple stakeholders across several product groups: Data, FB Blue, Serverless, Core Systems, Storage
- ▶ Lead the efforts to align & automate Meta's core infrastructure & serverless stack according to Meta Capacity Planning.
- ▶ Owning roadmap OKRs for improving testing & deployment reliability.
  - Defining, prioritizing and executing on the most critical projects to unblock MCP business outcomes.
  - Conducting retrospectives to outline the highest value adds.
  - Lead team effort to bring codebase test coverage up from 30% to 80%

#### Capacity Orchestration/ICSP Aug 2023 - Aug 2024

- ▶ Enable hardware fungibility to 80% of statically sized fleet through scaling the resize & onboarding API of the RRU autosizing abstraction
- ▶ Lead the team effort to establish production readiness of the next generation of service control plane. E2E testing, SLAs, DR runbooks.
- ▶ Implement E2E test framework for Meta's in-house autoscaling solution
- ▶ Guiding Junior engineer to design & implement updated SLAs for the autoscaling system

#### Capacity Modeling Sep 2021 - Aug 2023

- ▶ Scaling modeling & prediction systems to efficiently size more of Meta's fleet
  - enable \$19.5M yearly savings by saving on average 27k servers from the Instagram and FB Blue.
  - 900k predictions per hour (3x throughput, 5x better peak latency)
- ▶ Telemetry & SLA measurement to better detect, classify and prioritize reliability & performance regressions.
  - 32% drop in error rate
  - 80% drop in module latency
  - stability improvements for DR modeling algorithms.
- ▶ Systematic testing & CI/CD improvements to bring service deployment reliability up from 8% to 95%. Better explainability, coverage & accessibility of testing.

## Education

### UCLA

Grad. 2021

Coursework: AI/ML, Management, Entertainment Design, Real/Complex Analysis, Graph Theory, Quantum Computing

## Other Experience

### Daily Bruin 2017 - 2021

Bruinwalk Senior Staff  
Bruinwalk Editor  
Bruinwalk Web Developer

### Facebook 2020

Production Engineering Intern

### MemVerge 2019

Internship

### Cape Analytics 2018

Internship - Platform Team

## Skills

### Languages

Python, Thrift, Shell, SQL, SQL (Presto), C/C++, JavaScript

### Tools

Proficient with git, Linux, gdb, Tupperware, Docker, NodeJS, PostgreSQL and vim.

### Foreign Languages

French (*Advanced*)  
Marathi (*Conversational*)  
Hindi (*Beginner*)

## Interests

Astronomy • Music • Consumer Tech • Automotive • Aerospace