

AI & ML at QCon San Francisco 2026

Sessions

Skills, Memory, or Fine-Tuning? The Engineering Loop Behind Self-Improving Agents

As agents become mainstream, everyone wants to improve theirs either by making fewer mistakes on existing tasks or by taking on harder ones. This usually happens once an agent is already deployed in production.

Abhinav Sinha

CEO @Lucidic AI, Previously @Stanford AI Lab, @Citadel and Susquehanna International Group, and @Apple

Why Most Platform Teams Fail: The Adoption Problem Nobody Wants to Own

We have all seen the moment: the platform goes live, the launch deck looks sharp, the portal is polished, the golden paths are documented, and yet teams quietly continue doing things the old way. Not always because the platform is bad, but because adoption was assumed, not owned.

Shweta Vohra

Architecture Leader @Booking.com, Author of "Decoding Platform Engineering Patterns" & "Dear Software and AI Architect", 24+ Years Experience Building Cloud, Platform, and AI Systems

AI Inside and Out: New Guardrails for When AI Is Both Your Tool and Your User

AI is changing our systems from two directions simultaneously - and most teams are only watching one.

Adapt or Drift: Resilience Engineering When AI Moves the Operating Point

Resilient systems do not stay resilient by standing still. They survive by adapting. But adaptation comes with risk: under sustained pressure, organizations and systems can slowly drift toward failure while still

appearing to operate normally.

Andrew Hatch

Engineering Leader and SRE Manager @Cisco ThousandEyes, With 25+ Years Building Software, Operations, SRE, and Platform Teams Across Australia, India, and the United States

Building a Data Platform with AI

Data platform sits at the intersection of infrastructure, governance, operations, and user experience.

Adaptive Systems in Production: What Recommendation Systems Can Teach Us About Agents

As organizations race to build AI agents, many teams are encountering challenges that feel new: evaluation uncertainty, feedback loops, behavioral drift, exploration versus exploitation, and maintaining user trust in systems that continuously adapt.

But these challenges are not new.

Mallika Rao

Senior Engineering Manager @Zocdoc, Previously @Netflix, @Twitter and @Walmart

The Revenge of the Data Scientist: Why Reliable AI Needs Evals, Traces, and Metrics

Most teams can now ship an AI prototype by calling a foundation-model API. The hard part is knowing whether that system works when real users, messy data, and business consequences arrive.

Hamel Husain

Machine Learning Engineer, 20+ Years in Applied AI, Machine Learning, and Data Science

Progressive Failure Modes of Modern AI Serving Systems

Inference platforms fail in layers. Most organizations focus on model quality while underestimating the systems engineering required to operate production AI workloads safely and reliably at scale.

Abi Aryan

AI Infrastructure Engineer and Educator

Trainings

Building AI Agents: A Practical Framework for Choosing the Right SDK

Thursday Nov 19

The AI agent ecosystem has exploded with options—Google's Agent Developing Kit, Claude SDK, OpenAI Agents, LangGraph, Pydantic AI, and more emerging monthly. But which one should you actually use?

Hien Luu

AI/ML Leader, Advisor, Speaker, and Author

AI-Assisted Coding Master Class: Build Your Mission Control

Thursday Nov 19

Most engineers aren't spending enough time setting up their environment to get the most out of their coding agents.

Sepehr Khosravi

Machine Learning Platform Engineer @Coinbase, Award Winning Instructor @UC Berkeley - Gen-AI Bootcamp, Founder @AI Squads
