QCon SAN FRANCISCO by InfoQ

Platform Engineering at QCon San Francisco 2026

Sessions

Stop Fighting the Remodel: API Design for Models That Outgrew Their Boundaries

Monday Nov 16

Every domain model has a breaking point. At Netflix, ours arrived all at once—Ads, Games, Live events, and entirely new kinds of content. Each pulled the core entities in different, sometimes contradictory directions.

Aarti Gupta

Staff Engineer @Netflix

Saturation: How Your Software Will Fail at Scale

Monday Nov 16

Even in the ethereal world of software, everything has a limit. And, once that limit is reached, very bad things can happen. In this talk, we will explore the problem of saturation: when a software system runs into one of these limits.

Lorin Hochstein

Staff Software Engineer @Airbnb, Writes @surfingcomplexity.blog, Previously @Netflix and Member of the Resilience in Software Foundation

---

What Is Incident Severity, but a Lie Agreed Upon?

Monday Nov 16

Is your incident severity a complex word problem? Do you arbitrate which severity an incident should be during the incident itself? Do you know what the difference between your severities are? Do you even have incident severity in your org?

Em Ruppe

Technical Incident Commander @Twilio, Previously @SendGrid and @Twilio, and Product and Training

QCon SAN FRANCISCO by InfoQ

@Jeli.io

Inside Asimov: BlackRock's Platform Architecture for Applied AI at Scale

Tuesday Nov 17

BlackRock's Portfolio Management Group is transforming how investment teams conduct research by combining deep investment expertise with shared AI capabilities.

Dan Wolf

Head of Platform Engineering for PMG Tech @BlackRock

Sri Annapurna Bhupatiraju

Director and Principal Engineer @BlackRock, One of the Founding Members of BlackRock AI Labs

---

Building a Data Platform with AI

Tuesday Nov 17

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

Mouli Mukherjee

Engineering Manager @OpenAI

From Datasets to Data Products: Powering Trusted AI for Social Impact

Tuesday Nov 17

When an AI agent acts on a nonprofit's behalf, it's only as trustworthy as the data underneath it. At Blackbaud, we build software for the social impact sector, and nonprofits have relied on our data for decades.

Shannon McKenzie

Sr Principal Engineer @Blackbaud - Building the Data Platform for Social Impact

---

QCon
SAN FRANCISCO by InfoQ

Failing in Layers: Lessons from Rebuilding Netflix's ML Serving Platform

Tuesday Nov 17

Everyone talks about model quality. Fewer people talk about what happens when that model has to serve real traffic without falling over, or why the system ended up shaped the way it did in the first place.

Rajat Shah

Staff Software Engineer @Netflix

Wenbing Bai

Senior Software Engineer, AI Platform Team @Netflix

How Netflix Retrofitted Its SDLC for AI Coding Agents

Wednesday Nov 18

Agentic tools amplify what's already in the codebase. The agent works through your code, your architecture decisions, and the documentation around them. It follows what it can read and guesses at the rest.

Mark Khuzam

Senior Software Engineer @Netflix, Previously Led the Consumer Web Platform Team @OpenTable

---

Platform Engineering's Second Act: From Vending Machine to Passport Control

Wednesday Nov 18

Three years ago on the QCon SF stage, I made the case for "Acceleration, Autonomy, and Accountability" as the pillars of a successful platform. Those pillars haven't moved. AI has just rewritten what each one requires, and the platform team's job along with it.

Smruti Patel

SVP of Engineering @apollographql

Alex Mann

Senior Engineering Manager @apollographql

QCon SAN FRANCISCO by InfoQ

No More Small Systems, How AI May Bring Large System Problems to Everyone

Wednesday Nov 18

Once you figure out how to use AI to make your organization productive, you will soon find you have a whole new set of problems. While some of those problems will be novel thanks to LLMs, many of them are just the result of there being "more stuff".

Adam Bender

Technical Lead for Enterprise Engineering @Google

How LinkedIn Runs Compute: Self-Managed Kubernetes on 500,000+ Bare-Metal Servers

Wednesday Nov 18

LinkedIn's Kubernetes-based compute platform spans more than 500k bare-metal machines, 5M pods, with thousands of developers doing more than 110k deploys a week. We don't expose raw Kubernetes to developers and instead, provide a curated set of platform offerings.

Ronak Nathani

Principal Staff Software Engineer, Compute Infra @LinkedIn, Podcast Host @Software Misadventures

---

Observability: The Missing Control Plane for Agentic Engineering

Wednesday Nov 18

As AI agents move from generating code to testing, deploying, and operating it, the important question is no longer "How autonomous is the agent?" It's "What can this agent safely decide without a human?"

Rachael Wonnacott

Global Head of Machine Learning and Applied AI Engineering @Capital Group, OpenUK Ambassador & Founding Member of FastFlowConf, Previously Head of Cloud Engineering @Fidelity International

---

Why Platform Teams Fail: From Developer Platforms to Agentic AI

QCon SAN FRANCISCO by InfoQ

Wednesday Nov 18

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.

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

---

Building a Migration Platform: Moving 100+ Netflix RDBMS Workloads to Aurora PostgreSQL

Wednesday Nov 18

In late 2024, Netflix made a bet: consolidate the vast majority of our relational database use cases onto a single engine: Amazon Aurora PostgreSQL.

Ammar Khaku

Distributed Systems Engineer @Netflix with an Affinity for Online Data Stores

Kshitij Gupta

Software Engineer @Netflix, Building Databases and Distributed Systems

---