

# Platform Engineering at QCon San Francisco 2025

## Sessions

### Designing AI Platforms for Reliability: Tools for Certainty, Agents for Discovery

**Monday Nov 17**

Modern AI platforms don't have to choose between deterministic precision and probabilistic exploration—they need both.

**Aaron Erickson**

Senior Manager and Founder of the DGX Cloud Applied AI Lab @NVIDIA, Previously Engineer @ThoughtWorks, VP of Engineering @New Relic, CEO and Co-Founder @Orgspace

---

### Continuous Delivery for Foundational Platforms

**Monday Nov 17**

Platform teams frequently inherit systems that were never architected for their current scale, yet are so foundational that downtime can halt the business.

**Ian Nowland**

CEO @Junction Labs, Author of O'Reilly's Platform Engineering, Previously SVP Core Engineering at Datadog and Leader of AWS Nitro

---

### Confidently Automating Changes Across a Diverse Fleet

**Monday Nov 17**

Maintaining up-to-date and secure software across a polyglot fleet is a challenge for any engineering organization. Manual migrations and urgent updates disrupt productivity and require coordination across many teams.

**Casey Bleifer**

Senior Software Engineer @Netflix

## Beyond Line Charts: Why Some Diversity in Telemetry Visualization Is Long Overdue

**Monday Nov 17**

For decades, visualization of service metrics overwhelmingly converges to line charts. The time-centric nature of real-time telemetry further cemented this phenomenon via storage layouts and domain-specific query languages.

**Yao Yue**

Founder & Chief Executive Officer @IOP Systems, Platform Engineer, Distributed System Aficionado, Cache Expert

---

## From Staff Platform Engineer to a 16z Founder: What I Wish I'd Known

**Monday Nov 17**

**The Technical Leadership Challenge:** Every successful engineering team faces the same inflection point: transforming beloved internal tools into a legitimate platform business.

**Gonzalo (Glo) Maldonado**

CTO & Cofounder of @Mocksi.ai, Staff Platform Engineer Turned Startup Founder, Turns Developer Platforms into Products, Temporal & MCP Fanatic

---

## Microservices Platforms: When Team Topologies Meets Microservices Patterns

**Monday Nov 17**

When many teams work on a large, complex application, the microservice architecture potentially enables them to work independently and deliver a continuous stream of changes.

**Chris Richardson**

Creator of microservices.io, Java Champion, & Core Microservices Thoughtleader

---

## Building Resilient Platforms: Insights from 20+ Years in Mission-Critical Infrastructure

**Monday Nov 17**

In this talk, Matthew will describe lessons learned from over 20+ years of building scalable, secure and stable infrastructure platforms for software in financial services (electronic trading, credit card processing etc.), the talk is relevant to anyone building platforms for mission-critical

**Matthew Liste**

Head of Infrastructure @American Express, Previously @JPMorgan Chase and @Goldman Sachs

---

## Keeping the Mainline Green Across Diverse Language Monorepos

**Monday Nov 17**

At Uber's scale, ensuring an always-green mainline while processing hundreds of changes per hour is a massive challenge— especially when those changes span multiple language monorepos supporting dozens of business-critical apps.

**Dhruva Juloori**

Senior Software Engineer @Uber, Core Contributor to SubmitQueue (Uber's CI System at Scale), Expert in Machine Learning, Distributed Systems, and Developer Productivity

---

## Architecting a Centralized Platform for Data Deletion at Netflix

**Monday Nov 17**

What does it take to safely delete data at Netflix scale? In large-scale systems, data deletion cuts across infrastructure, reliability, and performance complexities.

**Vidhya Arvind**

Tech Lead & a Founding Architect for the Data Abstraction Platform @Netflix, Previously @Box and @Verizon

**Shawn Liu**

Senior Software Engineer @Netflix, Building Reliable and Extensible Systems for Consumer Data Lifecycle at Scale

---

# Platform Engineering: Lessons from the Rise and Fall of eBay Velocity

**Monday Nov 17**

Once a stock market darling and a pioneering hyperscaler in the 1990s and early 2000s, eBay has been in steady decline since the 2010s. A household name with a flat business, eBay has been unable to make substantive strides in its market reach or its engineering outcomes in the last 15 years.

**Randy Shoup**

SVP Engineering @Thrive Market, Previously @eBay, @Google, @Stitch Fix

---

# Directing a Swarm of Agents for Fun and Profit

**Monday Nov 17**

Coding agents are a new tool, which many of us are trying to figure out how to use effectively.

**Adrian Cockcroft**

Technology Advisor and Consultant @OrionX.net, Previously VP Open Source and Sustainability @Amazon, Cloud Architect @Netflix, Distinguished Engineer @eBay

---

# From VR to Flat Screens: Bridging the Input and Immersion Gap

**Tuesday Nov 18**

Expanding the market of your VR game to other platforms may seem like an easy way to increase your revenues but the journey is all but simple. When what makes your game unique is the presence and immersion it provides, reproducing the experience on a flat screen is a daunting task.

**Dany Lepage**

VP of Technology @Lucky VR - Building Virtual Vegas, xUbisoft, xFramestore, xNvidia

---

# Producing the World's Cheapest Tokens: A How-to Guide

**Wednesday Nov 19**

AI inference is expensive, but it doesn't have to be. In this talk, we'll break down how to systematically drive down the cost per token across different types of AI workloads.

**Meryem Arik**

Co-Founder and CEO @Doubleword (Previously TitanML), Recognized as a Technology Leader in Forbes 30 Under 30, Recovering Physicist

---

## **The Rust High Performance Talk You Did Not Expect**

**Wednesday Nov 19**

Rust runs faster, but it slows down engineers, right? This was our team's assumption when we decided to rewrite our code from Kotlin into Rust. But we were wrong in completely unexpected ways.

**Ruth Linehan**

Software Engineer @Momento, Previously APIs/Webhooks @GitHub and @Puppet

---

## **Instrumentation at Scale: Having Your Performance Cake and Eating It Too**

**Wednesday Nov 19**

In high-performance code, a single misplaced counter increment can cost more than the operation it's measuring. That creates a paradox: instrument too much and you slow the system down; instrument too little and you miss the insights you need to continuously deliver.

**Brian Martin**

Co-founder and Software Engineer @IOP Systems, Focused on High-Performance Software and Systems, Previously @Twitter

---

## **Accelerating LLM-Driven Developer Productivity at Zoox**

**Wednesday Nov 19**

Over the past year, Zoox has invested in integrating Large Language Models (LLMs) into the entire developer lifecycle through a companywide initiative called Zoox Intelligence (ZI).

**Amit Navindgi**

Staff Software Engineer, Developer Experience @ Zoox, Leading Applied AI Initiatives

---

## How Netflix Shapes our Fleet for Efficiency and Reliability

**Wednesday Nov 19**

Netflix runs on a complex multi-layer cloud architecture made up of thousands of services, caches, and databases. As hardware options, workload patterns, cost dynamics and the Netflix products evolve, the cost-optimal hardware and configuration for running our services is constantly changing.

### **Joseph Lynch**

Principal Software Engineer @Netflix Building Highly-Reliable and High-Leverage Infrastructure Across Stateless and Stateful Services

### **Argha C**

Staff Software Engineer @Netflix - Leading Netflix's Cloud Scalability Efforts for Live

---

## Powering the Future: Building Your GenAI Infrastructure Stack

**Wednesday Nov 19**

Behind every productivity leap is a rock-solid platform. Go under the hood with Intuit's GenOS team to see how vector stores, prompt management, RAG pipelines, and agent orchestration come together to serve ~100 million users.

### **Merrin Kurian**

Distinguished Engineer @Intuit

---

## Realtime and Batch Processing of GPU Workloads

**Wednesday Nov 19**

SS&C Technologies runs 47 trillion dollars of assets on our global private cloud. We have the primitives for infrastructure as well as platforms as a service like Kubernetes, Kafka, NiFi, Databases, etc.

### **Joseph Stein**

Principal Architect of Research & Development @SS&C Technologies, Previous Apache Kafka Committer and PMC Member

---

## From ms to $\mu$ s: OSS Valkey Architecture Patterns for Modern AI

**Wednesday Nov 19**

As AI applications demand faster and more intelligent data access, traditional caching strategies are hitting performance and reliability limits.

**Dumanshu Goyal**

Uber Technical Lead @Airbnb Powering \$11B Transactions, Formerly @Google and @AWS

---

## The Time it Wasn't DNS

**Wednesday Nov 19**

In January of 2023, the Microsoft Azure Wide Area Network experienced a global outage. If you were a Microsoft customer at the time, you were impacted by this outage.

**Sean Klein**

Principal Technical Program Manager - Modern Incident Analysis @Microsoft Azure

---

## Write-Ahead Intent Log: A Foundation for Efficient CDC at Scale

**Wednesday Nov 19**

As companies grow, so does the complexity of keeping distributed systems in sync. At DoorDash, we tackled this challenge while building a high-throughput, domain-oriented data platform for capturing changes across hundreds of services.

**Vinay Chella**

Engineering Leader @DoorDash - Specializing in Distributed Systems, Streaming & Storage Platforms, Apache Cassandra Committer, Previously Engineering Leader @Netflix

**Akshat Goel**

Staff Software Engineer, Core Infra at @DoorDash, Previously Senior Software Engineer @Amazon

---