QCon SAN FRANCISCO by InfoQ

Modern Languages & Infrastructure at QCon San Francisco 2025

Sessions

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

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

---

Rust at the Core - Accelerating Polyglot SDK Development

Monday Nov 17

Developing SDKs for your users in multiple languages can come at a high cost - especially if you need to implement complex logic client side, but traditionally options for sharing logic across those languages have been quite limited.

Spencer Judge

Engineering Manager @Temporal Technologies, previously Senior Software Engineer @Transparent

Systems, Senior Software Engineer @ Tableau Software

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

---

Engineering at AI Speed: Lessons from the First Agentically Accelerated Software Project

Tuesday Nov 18

Claude Code is the first developer tool built specifically to maximize AI development velocity.

Adam Wolff
Engineer and Individual Contributor to Claude Code @Anthropic, Previously @Robinhood, @Facebook

---

Why Fetch When You Can Sync? Building Local-First Apps on a Sync Engine Architecture

Tuesday Nov 18

Front-end has long been about reactivity frameworks and client-side state management. However, the alpha in these is receding.

James Arthur
Co-founder and CEO @ElectricSQL, Previously Co-Founder and CTO @Hazy and @Opendesk

---

Automating the Web With MCP: Infra That Doesn't Break

Tuesday Nov 18

AI agents are only as strong as the infrastructure beneath them. In this talk, we'll walk through the

architecture behind Browserbase's model context protocol (MCP), built to support stateful browser automation at scale.

Paul Klein

Founder @Browserbase, previously Director of Self-Service & Engineering Manager @Mux, Co-Founder & CTO @Stream Club, Technical Lead @Twilio Inc.

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

---

When Every Bit Counts: How Valkey Rebuilt Its Hashtable for Modern Hardware

Wednesday Nov 19

Ever wondered what happens when a bunch of performance-obsessed developers decide their blazing-fast database isn't quite blazing-fast enough?

Madelyn Olson

Principal Engineer @AWS, Maintainer of the Open-Source Valkey Project

---

Accelerating Performance by Incrementally Integrating Rust Into Existing Codebase

Wednesday Nov 19

In order to improve the performance of existing applications and services, we can identify the most performance-critical pieces and reimplement them in Rust as opposed to completely rewriting the applications from scratch.

Lily Mara

Staff Engineer @Discord, Author of "Refactoring to Rust", Previously Engineering Manager @OneSignal

---

From ms to µs: 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

Python, Numba, and Algorithm Design: Building Efficient Models in Financial Services

Wednesday Nov 19

The popularity of Python means insurance and financial services companies have a growing body of actuaries, quantitative developers, and software engineers capable of building innovative and customized solutions for both data management and modeling.

Chad Schuster

Principal @Milliman Focusing on Risk Management, Modeling, and Technology Consulting Services

---

QCon
SAN FRANCISCO by InfoQ