QCon SAN FRANCISCO by InfoQ

Observability & SRE at QCon San Francisco 2025

Sessions

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

---

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

---

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

---

Enhancing Reliability Using Service-Level Prioritized Load Shedding at Netflix

Monday Nov 17

How does Netflix maintain a seamless viewing experience for millions of users, especially during traffic spikes or when backend datastores are overloaded? Autoscaling can help during traffic spikes, but it costs money, takes a few minutes to kick in, and capacity may not always be available.

Anirudh Mendiratta

Staff Software Engineer, Playback Lifecycle @Netflix, Previously @Amazon Prime Video and @fuboTV

Benjamin Fedorka

Staff Software Engineer, Productivity Engineering @Netflix

Monolith Down: Cleaning Up After the Great Identity Migration Disaster

Tuesday Nov 18

One does not simply migrate a monolith. Imagine a team working on a monolith-to-microservices migration of a healthcare portal. A foundational first step - migrating to a commercial identity provider - takes 9 months, only to bring the entire portal crashing down on release day.

Sonya Natanzon

VP of Engineering @Heartflow, Decomplexifier, Software Architect, Healthcare and Life Sciences Specialist, and International speaker

---

Modernizing Relevance at Scale: LinkedIn's Migration Journey to Serve Billions of Users

Tuesday Nov 18

How do you deliver relevant and personalized recommendations to nearly a billion professionals—instantly, reliably, and at scale? At LinkedIn, the answer has been a multi-year journey of

architectural reinvention.

Nishant Lakshmikanth

Engineering Manager @LinkedIn, Leading Infrastructure for "People You May Know" and "People Follows", Previously @AWS and @Cisco

---

The Human Toll of Incidents & Ways To Mitigate It

Wednesday Nov 19

Have you ever wondered what it's like to respond to a significant incident? Walk through an hour by hour reconstruction of an incident response or two, focusing on what it was like to be "in the room" and the human response to the incidents.

Kyle Lexmond

Production Engineer @Meta, Previously @AWS and @Twitter

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 Incidents Refuse to End

Wednesday Nov 19

As engineers, we’re used to managing failure, but long-running outages hit differently. They stretch teams, systems, and assumptions about how incidents “should” play out.

Vanessa Huerta Granda

Resiliency Manager @Enova, Co-Author of the Howie Guide on Post Incident Analysis

QCon
SAN FRANCISCO by InfoQ

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

Week-Long Outage: Lifelong Lessons

Wednesday Nov 19

Routine database upgrades should be straightforward, especially with familiar, well-established technology. We were confident heading into our Elasticsearch upgrade, equipped with a solid plan and excited to see performance gains like we had seen from past upgrades.

Molly Struve

Staff Site Reliability Engineer @Netflix

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