QCon SAN FRANCISCO by InfoQ

Architecture 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

---

How to Build an Exchange

Monday Nov 17

These days it is possible to achieve fairly good performance on cloud provisioned systems. We discuss the design of a high performance, strongly consistent system which maintains constant service in the face of regular updates to core logic.

Frank Yu

Director of Engineering @Coinbase, Previously Principal Engineer and Director @FairX

Compiling Workflows into Databases: The Architecture That Shouldn't Work (But Does)

Monday Nov 17

What if everything you know about building distributed systems is backwards?

Jeremy Edberg

CEO of DBOS, Creator of Chaos Engineering, Tech Editor for 'AWS for Dummies'; Previously Founding Reliability Engineer @Netflix, and First Engineer @Reddit

Qian Li

Co-founder, Architect @DBOS, Stanford CS Ph.D., Co-organizer of South Bay Systems

Parting the Clouds: The Rise of Disaggregated Systems

Monday Nov 17

Cloud systems are undergoing an architectural shift. Traditional shared-nothing designs struggle to deliver the elasticity, availability, and operational simplicity that the cloud demands.

Murat Demirbas

Principal Research Scientist @MongoDB Research, Previously Principal Applied Scientist @AWS and a Professor of Computer Science at the University at Buffalo (SUNY)

Maximizing Success with Limited Time, Resources, and Energy: Lessons from Startup Engineering

Monday Nov 17

Startups are the harshest environments for engineers. Limited time, resources, and energy force teams to make decisions under pressure — and those decisions can determine whether a company survives or stalls.

David Gudeman

Co-Founder and CTO @Velocity AI

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-critic

Matthew Liste

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

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

---

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

---

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

---

Stripe's Docdb: How Zero-Downtime Data Movement Powers Trillion-Dollar Payment Processing

Tuesday Nov 18

Stripe processes over $1 trillion in payments annually with industry-leading reliability, powered by its custom-built document database, DocDB, built on top of open source MongoDB. Stripe's DocDB serves over five million queries per second from Stripe's product applications.

Jimmy Morzaria

Staff Software Engineer @Stripe, Previously Software Engineer on Amazon QLDB and Amazon Managed Streaming for Kafka

---

QCon SAN FRANCISCO by InfoQ

Migrating Uber Eats Feeds to Webview

Tuesday Nov 18

Uber Eats has many surfaces developed using native-first design. Historically these were built on the Android and iOS stacks. To accelerate development and enable rapid iteration and experimentation, while preserving the native-first design, a webview-powered stack was developed.

Nick DiStefano

Sr Staff Engineer @Uber, Previously iOS Lead @Tumblr

Accelerating Netflix Data: A Cross-Team Journey from Offline to Online

Tuesday Nov 18

At Netflix, certain use cases demand the rapid transfer of massive datasets—such as 50 TB—from offline to online systems. Doing this efficiently, without disrupting applications interacting with our online systems, presents a significant challenge.

Rajasekhar Ummadisetty

Software Engineer @Netflix - Driving Scalable Data Abstractions, Leader in Distributed Systems and Data Management, Previously @Amazon and @Facebook

Ken Kurzweil

Software Engineer @Netflix - Leading a Data Movement Team Focused on Data Infrastructure Innovation, Previously @Amazon, @Shutterfly, and @Gannett Media

---

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

---

QCon SAN FRANCISCO by InfoQ

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

---

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

---

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

---

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 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

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

---