

Platform Engineering at QCon London 2026

Sessions

Unconference: Building Engineering Teams

Tuesday Mar 17

Connecting the Dots: Modern Data Engineering & Architectures (Limited Space - Registration Required)

Tuesday Mar 17

Unconference: Tech of Finance Industry

Monday Mar 16

Unconference: Connecting Systems

Monday Mar 16

Wrangling Telemetry at Scale: A Guide to Self-Hosted Observability

Tuesday Mar 17

Observability is supposed to help you tame complexity, but your Observability stack can quickly become just as complex as the systems it's meant to watch. For most teams, the answer is to pay someone else to deal with it.

Colin Douch

Site Reliability Engineer @DuckDuckGo

How Eve Online Leverages Head Based Sampling to Observe "Fun"

Tuesday Mar 17

A unique pattern in video game software is real-time interactions to express the personality of users.

Here we will talk about how we instrument the universe of New Eden to identify the traffic that matters, even the "fun" parts!

Nicholas Herring

Technical Director, Eve Online @CCP Games, Refiner of Internet Spaceships and Explorer of Feral Gordian Knots of Python

Orienting, Understanding, Playing, Thriving: Debugging your Organisation

Tuesday Mar 17

Debugging is both an art and a science. But more than that, it's an activity undertaken with deep intention: to understand and improve your systems. In the purely technical realm, we have an extraordinary range of tooling and techniques that can help us tackle this problem.

Hazel Weakly

Fellow @Nivenly Foundation; Director, Haskell Foundation; Experienced Leader Focusing on Organizational Change, Developer Experience, and Resilience Engineering

Modernising Retail at Scale: Architecting a Cloud to Edge Platform for a Global Enterprise

Monday Mar 16

Modern retail depends on a technology platform that operates consistently across digital and physical channels - from customer experiences to colleague tools and the supply chain that supports them.

Jayesh Bhayani

Sr Director Technology @Tesco PLC, Board Member at Tesco Technology and Services Europe, Board Audit Committee Member of Tesco Mobile

Behind Booking.com's AI Evolution: The Unpolished Story

Monday Mar 16

It's easy to look at a mature AI platform and imagine a grand blueprint. Ours began with none. What started as a few data scientists hacking on Perl scripts and Mysql queries has grown into an AI platform that impacts millions of travel decisions every day.

Jabez Eliezer Manuel

Principal Engineer @Booking.com - Building Next-Gen AI Platform

Enchant Your AI and APIs with eBPF Magic

Monday Mar 16

It is a common occurrence to see applications thrown over the fence, landing somewhere in production without a second thought about their lifecycle or how they may need maintaining in the future to connect to more efficient API endpoints.

Dan Finneran

Principal Community Advocate at Isovalent @Cisco

Managing Asynchronous APIs at Scale

Monday Mar 16

When event-driven architectures are small, teams can reason about events through word-of-mouth. They know who publishes what, who consumes it, and how messages flow through the system. Teams manage their own infrastructure or raise tickets to request changes.

Ian Cooper

Senior Principal Engineer @Just Eat Takeaway

Shielding the Core: Architecting Resilience with Multi-Layer Defenses

Wednesday Mar 18

High-demand events can cause sudden traffic spikes that overwhelm even well-designed systems. In ticketing platforms, millions of users — alongside increasingly sophisticated automated agents — may arrive simultaneously, placing extreme pressure on backend services.

Anderson Parra

Staff Software Engineer @SeatGeek

Blurring the Lines: Engineering & Data Teams in the Age of AI

Tuesday Mar 17

Every senior engineer knows the feeling: a model makes a bad decision, a customer complains, and suddenly you're debugging a system that spans three teams, two pipelines, and a machine learning model nobody fully owns. Where do you even start?

Lada Indra

Head of Data Platform @Pleo, Previously Head of Data @Legend and Director API Platform BI & Data @Vonage

Ethical AI Is an Engineering Problem

Wednesday Mar 18

The systems we design encode the values we choose

Clara Higuera Cabañes

Responsible AI Program Lead @BBVA - Driving Innovation & Social Impact in AI, PhD Thesis Advisor, Member of OdiselA, Previously Data Scientist @BBC

AI is an Amplifier: Scale High Performance, Not Dysfunction

Wednesday Mar 18

AI adoption in software development is nearly universal, yet the outcomes for teams are highly variable. Why do some organizations see massive productivity gains while others see their delivery stability crash? DORA's latest research provides a key insight: AI acts as an amplifier.

Nathen Harvey

Lead of DORA and Product Manager @Google Cloud

Move Fast, Don't Break Trust: Shipping Constantly with Humans and Beyond

Monday Mar 16

Moving fast only works if your systems are engineered with trust and verification in mind. In this talk, we showcase how teams can scale to the point where shipping hundreds of changes a day becomes boring, routine, and safe.

Suhail Patel

Senior Staff Engineer @Monzo Leading the Platform and Data Functions, Previously @Citymapper

The Realities of Building an AI Native Fintech Startup

Monday Mar 16

For fifty years, \$400 trillion in professionally managed assets sat untouched by external technology. The most analytically demanding investment workflows required a combination of PhD-level quant finance and advanced engineering that no outside vendor could sustain commercially.

David Lin

Founder and CEO @Linvest21, Previously CTO @JPMorgan

Your Multicloud Strategy Is a Product Problem - Treat It Like One

Monday Mar 16

You can't really "opt out" of multicloud anymore. Between cloud concentration risk, SaaS sprawl, and increasing regulatory expectations, most enterprises end up operating across multiple clouds whether they planned to or not. The hard part isn't having two or three providers.

Luis Henrique Albinati Junior

Executive Director, Product Strategy @JP Morgan Chase, Previously @Microsoft, @Oracle, and @Red Hat

Surabhi Mahajan

Executive Director @JP Morgan Chase

Machine Learning at the Edge of Scale and Speed: Nanosecond Inference at the CERN Large Hadron Collider

Wednesday Mar 18

The CERN Large Hadron Collider (LHC) produces $O(10,000)$ exabytes of raw data annually from high-energy proton collisions. Handling this volume under strict compute and storage limits requires real-time event filtering capable of processing millions of collisions per second.

Thea Klaeboe Arrestad

Particle Physics and Real-Time ML @CERN @ETH Zürich

The Rise of the Streamhouse: Idea, Trade-Offs, and Evolution

Tuesday Mar 17

Over the last decade, streaming architectures have largely been built around topic-centric primitives—logs, streams, and event pipelines—then stitched together with databases, caches, OLAP engines, and (increasingly) new serving systems.

Giannis Polyzos

Principal Streaming Architect @Ververica

Anton Borisov

Principal Data Architect @Fresha

APIs for Agents: Rethinking API Programs in the MCP Era

Monday Mar 16

As API programs mature, a familiar gap emerges: some teams operate with strong standards, reusable platforms, and clear governance, while others rely on informal guidance and best-effort consistency.

Jim Gough

Distinguished Engineer, API Platform Lead Architect @Morgan Stanley, Co-Author of Optimizing Java

Andreea Niculcea

Vice President @Morgan Stanley

The Ladder Is Missing Rungs: Engineering Progression When AI Ate the Middle

Tuesday Mar 17

Career progression in engineering has traditionally followed a predictable path: junior tasks teach fundamentals, mid-level work builds judgment, senior roles require synthesis across systems.

Alasdair Allan

Scientist, Author, Hacker, Maker, Journalist, CTO @Negrone Venture Studios, Interim CTO @Evaro

Building an AI Gateway Without Frameworks: One Platform, Many Agents

Tuesday Mar 17

Early AI integrations often start small: wrap an inference API, add a prompt, ship a feature. At Zoox, that approach grew into Cortex, a production AI gateway supporting multiple model providers, multiple modalities, and agentic workflows with dozens of tools, serving over 100 internal clients.

Amit Navindgi

Staff Software Engineer @Zoox

Jatin Aneja

Leading Developer Experience @Zoox, Previously Director of Site Reliability Engineering @AppLovin

Rewriting All of Spotify's Code Base, All the Time

Tuesday Mar 17

We don't need LLMs to write new code. We need them to clean up the mess we already made.

In mature organizations, we have to maintain and migrate the existing codebase. Engineers are constantly balancing new feature development with endless software upkeep.

Jo Kelly-Fenton

Engineer @Spotify

Aleksandar Mitic

Software Engineer @Spotify

From Fan-Out to Fast: Sub-100ms API Design in Distributed Systems

Monday Mar 16

A "simple" API request rarely stays simple. In distributed systems, one call quickly turns into fan-out across gateways, services, caches, and databases — and your p99 becomes the sum of every hop and every flaky dependency.

Saranya Vedagiri

Senior Staff Engineer @eBay

Building on Bedrock: A Security Philosophy from Bootloader to Runtime

Tuesday Mar 17

In Minecraft, every world is built from blocks. At the very bottom lies bedrock: an unbreakable foundation that everything else rests on. Above it sit layers of stone, dirt, sand, and other materials.

Alex Zenla

Founder & CTO @Edera

Why Governance Matters: The Key to Reducing Risk Without Slowing Down

Tuesday Mar 17

When you hear “governance,” you might think of red tape, bureaucracy, or someone telling you what you can’t do. But real governance is about alignment and reducing technical risk. And that matters more than ever.

Sarah Wells

Independent Consultant and Author
