QCon LONDON by InfoQ

Architecture at QCon London 2026

Sessions

Unconference: Modern Data Engineering

Tuesday Mar 17

Unconference: Architectures You've Always Wondered About

Monday Mar 16

Stitching Together Traces in a World of Async Callbacks

Tuesday Mar 17

Details coming soon.

Evolving Wise Architecture to Power a Global Account

Monday Mar 16

Details coming soon.

Andrei Tognolo
Staff Engineer @Wise, 19+ Years in Software Engineering, Previously Senior Consultant @ThoughtWorks

---

Evolution of Booking.com's Ranking Platform

Monday Mar 16

Details coming soon.

QCon LONDON by InfoQ

Navigating the Realities of Global Multi-Cloud Resilience

Monday Mar 16

Details coming soon.

Building an AI Ready Global Scale Data Platform

Wednesday Mar 18

As organizations move from single-cloud setups to hybrid and multi-cloud strategies, they are under pressure to build data platforms that are both globally available and AI-ready.

George Peter Hantzaras

Engineering Director, Core Platforms @MongoDB, Open Source Ambassador, Published Author

Explicit Semantics for AI Applications: Ontologies in Practice

Wednesday Mar 18

Modern AI applications struggle not because of a lack of models, but because meaning is implicit, fragmented, and brittle. In this talk, we'll explore how making semantics explicit (using ontologies and knowledge graphs) changes how we design, build, and operate AI systems.

Jesus Barrasa

Field CTO for AI @Neo4j

Ontology Driven Observability: Building the E2E Knowledge Graph at Netflix Scale

Tuesday Mar 17

As Netflix scales hundreds of client platforms, microservices, and infrastructure components, correlating user experience with system performance has become a hard data problem, not just an observability one.

QCon LONDON by InfoQ

Prasanna Vijayanathan
Engineer @Netflix
Renzo Sanchez-Silva
Engineer @Netflix

From DVDs to Global Streaming: How Netflix's Commerce Architecture Actually Evolved

Monday Mar 16

Netflix didn't start as a global streaming platform. It began as a US-centric DVD-by-mail business, with a commerce system designed for one country, one currency, and relatively simple payment flows.

As Netflix expanded internationally, those early assumptions began to break.

Kasia Trapszo
Principal Engineer @Netflix, Leading Architecture for the Commerce Platform

---

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

From S3 to GPU in One Copy: Rethinking Data Loading for ML Training

Tuesday Mar 17

ML training pipelines treat data as static. Teams spend weeks preprocessing datasets into WebDataset or TFRecords, and when they want to experiment with curriculum learning or data mixing, they reprocess everything from scratch.

Onur Satici

Staff Engineer @SpiralDB & Core Maintainer of Vortex (LF AI & Data), Previously Building Distributed Systems @Palantir

---

The Influence Toolkit

Monday Mar 16

Senior IC roles are often said to be about influence – but what do we mean by influence? And what are the tools that a senior IC can employ to influence individuals and teams within a large organisation?

Henry Wilson

Senior Architect @BBC

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

Introducing Tansu.io -- Rethinking Kafka for Lean Operations

Tuesday Mar 17

What if Kafka brokers were ephemeral, stateless and leaderless with durability delegated to a pluggable storage layer?

QCon LONDON by InfoQ

Peter Morgan
Founder @tansu.io

The Right 300 Tokens Beat 100k Noisy Ones: The Architecture of Context Engineering

Wednesday Mar 18

Your agent has 100k tokens of context. It still forgets what you told it two messages ago.

Patrick Debois
AI Product Engineer @Tessl, Co-Author of the "DevOps Handbook", Content Curator at AI Native Developer Community

---

How to Find Resilience Bugs in Systems that Don't Exist

Wednesday Mar 18

Building correct distributed systems takes thinking outside the box, and the fastest way to do that is to think inside a different box. One different box is "formal methods", the discipline of mathematically verifying software and systems.

Hillel Wayne
Author of "Logic for Programmers" and "Learn TLA+", Thought Leader in the Space of Empirical Software Engineering

---

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.

QCon LONDON by InfoQ

Saranya Vedagiri
Senior Staff Engineer @eBay

Beyond Context Windows: Building Cognitive Memory for AI Agents

Tuesday Mar 17

AI agents are rapidly changing how users interact with software, yet most agentic systems today operate with little to no intelligent memory, relying instead on brittle context-window heuristics or short-term state.

Karthik Ramgopal
Distinguished Engineer & Tech Lead of the Product Engineering Team @LinkedIn, 15+ Years of Experience in Full-Stack Software Development

---

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
