

AI & ML at QCon London 2026

Sessions

Evolution of Booking.com's Ranking Platform

Monday Mar 16

Details coming soon.

From Copilots to Orchestrators: A Three-Month Playbook for Training AI-Native Engineering Teams

Tuesday Mar 17

Most engineering teams are stuck treating AI as autocomplete. Engineers have GitHub Copilot installed (or Claude or Cursor or whatever), they're generating snippets faster, but leaders can't connect usage to business outcomes—and developers are shipping code they don't fully understand.

Krystal Flores

Staff Software Engineer @Crunchyroll, Previously @Carta, @Lob, @Simple Habit, and @Nordstromrack.com|HauteLook

Team Topologies as the 'Infrastructure for Agency' with AI

Tuesday Mar 17

Details coming soon.

Matthew Skelton

CEO & Principal @Conflux, Co-Author of "Team Topologies", Leader in Modern Organizational Dynamics for Fast Flow

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

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

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 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, and Head of Documentation, CTO @Negroni 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

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

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

Refreshing Stale Code Intelligence

Tuesday Mar 17

Coding models are helping software developers move even faster than ever before, but weirdly, they're not keeping up with our fast progress. The models that power code generation are often based on months to years old snapshots of open source code.

Jeff Smith

CEO & Co-Founder @ 2nd Set AI, AI Engineer, Researcher, Author, Ex-Meta/FAIR

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

Reliable Retrieval for Production AI Systems

Tuesday Mar 17

Search is central to many AI systems. Everyone is building RAG and agents right now, but few are building reliable retrieval systems.

Lan Chu

AI Tech Lead and Senior Data Scientist
