

# SRE CON — EUROPE MIDDLE EAST AFRICA

DUBLIN, IRELAND • 7–9 October, 2025



Rev. 26/03/2025

*Sponsored by USENIX, the Advanced Computing Systems Association*

## Important Dates

- Talk proposals due: **Wednesday, 21 May, 2025, 23:59 UTC**
- Notification to talk presenters: **Wednesday, 18 June, 2025**
- Confirmation of acceptances and deadline for program materials: **Wednesday, 9 July, 2025**

## Overview

AI and Machine Learning are fundamentally reshaping how production systems are designed, deployed, and operated. Unlike traditional software, where reliability primarily concerns infrastructure and code, AI-driven systems depend on large-scale data pipelines, evolving datasets, and complex model behaviors. Ensuring reliability in this new landscape is no longer just about uptime and error budgets—it's about data integrity, model stability, and delivering predictable model performance over time. As companies increasingly integrate AI into their platforms, managing data quality and operating ML models is becoming a critical reliability concern.

At SREcon25 Europe/Middle East/Africa, we will dive deep into the new **reliability challenges posed by Data and AI-driven systems**. How do we define SLOs for models whose outputs don't have a clear error signal? What observability tools do we need to detect silent failures, data drift, and performance degradation? How do we debug AI incidents, safely roll back models, and ensure that automated decisions remain trustworthy over time? And how do we mitigate security risks such as prompt injection, adversarial attacks, and model poisoning? How do we empower SREs to select the right security models for safeguarding data? We invite SREs, data engineers, and AI practitioners to share best practices, real-world lessons, and innovative approaches to building and operating reliable AI systems at scale.

In 2025, SREcon Europe/Middle East/Africa also introduces a new format: the **InFocus track**. Each day, this track will center on a specific theme. This year's themes are "Data and AI Reliability," "Platform Engineering," and "Reliability in Finance." The main conference tracks will continue to feature a broad range of SRE topics that complement and extend these themes, including Full-Stack Observability, AI and Automation, SRE and Culture, and Systems Engineering.

## Themes

### InFocus: Data and AI Reliability

This track explores how AI-driven systems introduce new reliability challenges, requiring stronger data management, observability, and operational strategies to ensure model stability and trustworthiness.

- **Ensuring Data Quality and Integrity:** Operationalizing data validation, contracts, and governance.
- **Observability for AI and Data Pipelines:** Monitoring model performance, data drift, and silent failures. Detecting and mitigating failures that don't trigger traditional alerts. Defining SLOs for models with non-deterministic outputs.
- **Scaling ML and Data Infrastructure:** Managing latency, cost, and reliability in AI-powered systems. Architecting training and inference tooling and platforms.
- **Integrations to AI Systems:** Unique challenges of AI-powered integrations, including Chat, RAG, function calling, agents, and code completion.
- **ML Model Deployment and Rollbacks:** Safe model versioning, rollback strategies, and progressive rollouts.
- **Debugging AI Systems in Production:** Troubleshooting ML failures across models, pipelines, and data sources.
- **Security Risks in AI Systems:** Mitigating prompt injection, adversarial attacks, and model poisoning.

### InFocus: Platform Engineering

This track explores how modern SRE are building platforms that provide self-serviceable, automated, and scalable reliability solutions.

- **SRE and Platform Engineering: Collaboration or Convergence?** Defining ownership and responsibilities in a platform-driven world. Building best practices into developer workflows and tooling.
- **Platform as a Product:** Steering and funding of platform initiatives depends on measuring the business value provided by platform and reliability.
- **Reducing Toil Through Self-Service:** Designing platforms that empower teams to manage reliability at scale.

- **Observability in Platform Engineering:** Building monitoring, tracing, and diagnostics directly into platforms. Defining SLOs and Error budgets following best principles.

### InFocus: Reliability in Finance

Since the first SREcon companies in the finance sector have strived to adopt SRE with varying success. This track will share the secrets of how financials have adopted SRE culture. The finance sector encompasses broader institutions such as exchanges, trading systems, high-frequency low-latency companies, hedge funds, and insurance. Each institution creates unique challenges for SRE adoption.

- **Observability and Data Pipelines:** Monitoring at a scale is incredibly important, especially for companies at the heart of the economy. Defining SLOs and error budgets to help stability in those institutions is not straightforward.
- **High Availability in Financial Systems:** The reliability of systems for banking, trading, and payments are of critical importance to the global economy. How do these requirements impact software and data architecture? What are the implications to incident and failure response?
- **Navigating Regulatory Constraints:** How financial institutions align SRE practices with compliance requirements (e.g., SOX, Basel III, DORA) while supporting developers and business needs.
- **How Are Decisions Made?** Many financial companies have a long process and series of meetings to reach a decision. How do you reduce this toil through governance, standard frameworks, or self-service platforms, and be compliant with regulatory concerns?

Whatever the journey, this track is your opportunity to help share the stories around SRE adoption within the financial companies and shape where the next steps the Financial industry could take in their journey.

### AI and Automation in SRE

Reducing toil with automation has always been a central pillar of SRE activities. This track focuses on how automation has been evolving over the past year, and how AI is starting to be leveraged as part of those automations.

- **Auto-Remediation: Blessing or Curse?** Can AI automate incident response, or does it impact the ability of human operators to deal with incidents that the auto-remediation tools cannot handle?
- **AI-Assisted Incident Response:** How can AI support incident detection, triage, and mitigation?
- **AI Postmortem Tooling:** Automating incident documentation and refining postmortems with AI-powered tools.

### Systems Engineering and Principles

This track explores core technologies that power large-scale distributed systems, focusing on their architecture, strengths, limitations, and recent advancements. Talks should provide deep technical insights that help

engineers design, support, and scale infrastructure effectively. We especially welcome topics that bridge applied theory (e.g., queueing theory, load balancing, consistency models) with real-world challenges. We seek proposals on many topics, such as:

- **Performance and Scalability:** OS tuning, kernel optimizations, hardware bottlenecks, and performance engineering at scale.
- **Networking and Load Balancing:** SD-WAN, DNS, proxies, IP protocols, BGP, and traffic engineering strategies.
- **Storage and Databases:** Data persistence, consistency models, performance tuning in MySQL/PostgreSQL, and distributed storage.
- **Security and Low-Level Systems:** TPMs, HSMs, firmware security, transport encryption, and filesystem integrity.

Additionally, we always welcome submissions on all topics relevant to both the human side and the technical aspects of software operations, availability, and reliability.

### Full-Stack Observability

Observability is evolving beyond logs, metrics, and traces to provide deeper, actionable insights spanning from mobile/client side over backend systems to data-pipeline observability. This track will explore observability techniques, best practices, and tools that enable engineers to proactively ensure system reliability.

- **Mobile Observability:** Ensuring reliability and performance across mobile networks, devices, and operating systems.
- **Data Observability:** Detecting schema changes, ensuring data quality, and monitoring large-scale data infrastructure.
- **Observability for AI and Data Pipelines:** Monitoring model performance, data drift, and silent failures in machine learning systems.
- **Observability-Driven Development:** Embedding observability into engineering workflows from day one.

### Culture and SRE Maturity

SRE is more than just a set of technical practices—it's a culture that shapes how organizations approach reliability, collaboration, and growth. This track will explore how companies evolve their SRE maturity, scale teams, and create sustainable on-call practices.

- **Establishing SRE in Small and Medium-Sized Companies:** How to get started with SRE? What is different from hyperscaler environments?
- **Measuring SRE Impact:** How do organizations quantify the value of reliability initiatives and platform products?
- **SRE Maturity Models:** How do organizations progress through different stages of SRE maturity? What frameworks exist to assess reliability capabilities, and how can teams benchmark their progress?

## Talks

- Talks
  - 15-minute talks followed by 5 minutes for Q&A
  - 35-minute talks followed by 10 minutes for Q&A

## Participant Information

SREcon participants come from a wide variety of backgrounds: small startups, tech giants with tens of thousands of employees, finance, and enterprise sector companies adopting or expanding SRE in their organizations, and academia. New speakers are encouraged to submit talks; many of our best talks have come from people with new perspectives to share, and the previous year most certainly has given us all new experiences and stories we can share and learn from.

We welcome and encourage participation from all individuals in any country. We also welcome participants from diverse professional roles: QA testers, customer experience/support, security teams, DBAs, network administrators, compliance experts, UX designers, health care professionals, scientists, and economists. Regardless of who you are or the job title you hold, if you are a technologist who faces unique challenges and shares our areas of interest, we encourage you to be a part of SREcon25 Europe/Middle East/Africa.

## Speaker Information

To see the details of what we want to know about your proposal to speak, we encourage you to click through the talks submission system, linked from the Call for Participation web page.

If you are a new presenter or would just like some extra help, please reach out. We can provide support via practice sessions.

Both presenters and organizers may withdraw or decline proposals for any reason, even after initial acceptance. Speakers must submit their own proposals; third-party submissions, even if authorized, will be rejected.

SREcon25 Europe/Middle East/Africa is scheduled to be held in Dublin. All presentations will require in-person participation.

If you have questions about this Call for Participation, feel free to drop us a message at [srecon25emea\\_chairs@usenix.org](mailto:srecon25emea_chairs@usenix.org).

## Background (Overarching goals of the worldwide SREcon conferences)

SREcon is a gathering of engineers who care deeply about site reliability, systems engineering, and working with complex distributed systems at scale. Our purpose is to be inclusive as we bring together ideas representative of our diverse community, whether its members are focusing on a global scale, launching new products and ideas for a small business, or pivoting their approach to unite software and systems engineering. SREcon challenges both those new to the profession as well as those who have been involved in SRE or related endeavors for years. The conference culture is built upon respectful collaboration amongst all participants in the community through critical thought, deep technical insights, continuous improvement, and innovation.

For more information on the themes and programs of past conferences, see the list of past conferences available at [www.usenix.org/srecon#past](http://www.usenix.org/srecon#past).

## Conference Organizers

### Program Co-Chairs

Heinrich Hartmann, Zalando  
Stephane Dudzinski, Reddit

### Program Committee

TBA

### Steering Committee

Kurt Andersen, Clari  
Daria Barteneva, Microsoft  
Liz Fong-Jones, Honeycomb  
Casey Henderson-Ross, USENIX Association  
John Looney, Reddit  
Laura Nolan  
Mike Rembetsy, Bloomberg  
Emil Stolarsky  
Avleen Vig, Twilio