

SRE CON_®



VIRTUAL EVENT • October 12–14, 2021

Sponsored by USENIX, the Advanced Computing Systems Association

Important Dates

- Proposals due: ~~Wednesday, June 30, 2021, 11:59 pm UTC~~ **Tuesday, July 6, 2021 12:01 am UTC (Extended)**
- Notifications to presenters: **Wednesday, August 4, 2021**
- Confirmation of acceptances and deadline for program materials: **Tuesday, August 10, 2021**
- Pre-recorded videos due: **Monday, September 13, 2021**
 - If changes are needed, you will be notified as soon as possible and will need to submit your changes within one week from that date.
- Conference start: **Tuesday, October 12, 2021**

Overview

For this SREcon, we want to hear how emerging technologies are influencing, and shaping the world of SRE. We would like to learn about innovations around remote working, scaling, observability, resilient architectures, and security—all of which have been essential to us as an industry. This includes new processes, tools, or patterns that you have been using or developing, as well as things SREs have discovered and learnt about systems and data management through their application. We want to hear how your year of scaling up and out has gone, what is going to make next year easier, and how might we move forward as we continue to grow as a community.

We are interested in how SREs work along many dimensions but would like to bring extra attention to emerging technologies, facilitation of incident management, machine learning, and new scaling requirements and lessons learned. Talks might consider such questions as:

- Have you seen innovative applications of technology to enhance SRE tooling, processes, and management of complexity within our systems?
- As fields of Artificial Intelligence, Machine Learning, Robotics, and Cybersecurity continue to grow, what does practicing SRE mean across these areas?
- As we embraced remote working through the pandemic, how has this influenced how we build, practice, and scale SRE? What were your experiences on working as a distributed infrastructure team?
- With 34% annual growth of roles (2020 LinkedIn US Emerging Jobs Report) and the SRE job title becoming more coveted, how do we ensure our community can sustain this rate of growth, stay abreast of developments in the industry, and keep the culture from being diluted?

Please join us in creating an excellent program for SREcon21. Since 2016, participants have 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 2021 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, including people that are under-represented in, or excluded from, technology, including but not limited to: people of all colors, women, LGBTQ people, people with disabilities, neurodiverse participants, students, veterans, and others with unique characteristics. Similarly, we welcome participants from diverse professional roles: QA testers, security teams, DBAs, network administrators, compliance experts, UX designers, government employees, scientists. 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 SREcon21.

Proposals

We are looking for proposals to fit into a global virtual conference, with talks spanning different time zones so we can be as inclusive of all attendees and locations as possible. All talks will be pre-recorded and scheduled like a normal conference for attendees to watch. During presentations speakers will be able to interact with the audience in real time! More details will follow upon acceptance of your proposal in the submission system.

Due to the nature of a virtual conference we will not have workshops or lightning talks this year and instead will be supporting panels and video conversations.

Below are some ideas of what you or your team could submit.

Talks

- 15 or 30-minute talk (with corresponding Q&A)

Suggested Topics

- Tell us about applying SRE principles to production management of emerging technologies.
 - Do all principles we know and love apply?

- Are there gaps in our understanding and practices?
- What can SRE uniquely bring to developing these areas?
- Tell us about building SRE teams and practices for brand new systems or organizational functions.
 - How did your organization decide how and when SRE is required?
 - What challenges or nuances did you discover in doing this?
 - Did you use any tried and tested frameworks to help manage scaling of this nature?
 - How did you measurably demonstrate the value of SRE?
- Tell us about your wild debugging stories that took you into the internals of complex systems.
 - Did your investigation cross team, product, or international boundaries?
 - Was it an emergency investigation or the product of months or years?
 - Was it a “wicked problem,” something where the fix created other problems?
- Tell us about the deep technical or people skills you have developed expertise in as an SRE.
 - How do you disseminate this knowledge and coach others in your organization?
- What are new and unexpected types of incidents you are encountering?
 - Has this changed how you facilitate incident management?
 - How are you learning from and adapting to them?
- Most organizations had to adapt their operations rapidly at the start of the pandemic. Tell us what has continued to work and evolve since, and what has already been abandoned.
 - For anyone able to return to physical offices, what have you decided to continue from working remotely?
 - For organizations staying remote, what are the best practices you are taking from a year of cross-industry experimentation?
- Tell us about ways to stay ahead of developments in the technology industry.
 - How do we ensure SRE continues to stay relevant?
 - What is practical today, versus aspirational for the future?

Video Panels

- Have you worked at an organization that dealt with unexpected scale during the pandemic (e.g., organizations that focused on food delivery, streaming, medical services, media, COVID-19 work)?
- Compare and contrast how SRE is practiced across different types of systems or technical domains. What are the nuances you have discovered or experienced?
- Tell us about a challenging incident that you have experienced and resolved. Bring some of your team members onto this panel, and tell us the story about what happened and how the team coordinated, reacted to, and solved the problem.

Video-Conversations

- Do you have a topic you are wanting to be interviewed on? A topic you are wanting to interview someone else on? We would love to hear it.

Principles Track

With the popularity of the Core Principles track at previous SREcon conferences, we’ve decided to continue this component at SREcon20 Americas. Talks in this track will focus on providing an in-depth understanding of how the underlying technologies used by many SREs function, and why it’s important to know these details when supporting and scaling your infrastructure.

For this track, we are looking for a number of topics, such as:

- **Performance** (e.g., CPU affinity, bottlenecks)
- **Databases** (e.g., how is data stored on disk in MySQL, PostgreSQL, etc.?)
- **Observability** (e.g., monitoring overview, events vs. metrics, whitebox vs. blackbox, visualizations, debugging, etc.)
- **Distributed Systems** (e.g., consistency and consensus, Hadoop, MapReduce, Jupyter Notebooks, Containers)
- **Network** (e.g., SD WAN, HTTP routing and load balancing, DNS)

Operationalizing Machine Learning (OpML) Track

Recognizing the affinity between the challenges of operationalizing machine learning at scale with other SRE realms of practice, we are inviting the community from the OpML (or MLOps) space to participate in a dedicated track at SREcon. We will have specialist program committee members evaluate proposals for the OpML track. Submissions will be reviewed based on their applicability to the problems of Operational ML, originality, technical merit, topical relevance, and the likelihood of leading to insightful discussions that will influence practices of ML and its variants in production and benefit attendees.

For this track, we are looking for a number of topics, such as:

- Challenges of production training and re-training including scale, transitioning from experimental models to production models, etc.
- Systems for orchestrating, diagnosing, monitoring, and managing ML in production.
- Applying existing DevOps and SDLC tools and practices to the ML operational lifecycle.
- Diagnostics of ML algorithms operating on live data (e.g., drift detection).
- New model introduction into production (e.g., staging, A/B test).
- Governance of ML models and deployment processes, model risk management, adaptations for new regulations.
- Systems, Tools, Approaches for Model/Dataset version control, and lineage.
- Data governance approaches as they relate to production ML.
- Technical advances for addressing ML regulatory requirements (e.g., GDPR, Schrems II).

- Privacy and security challenges in production ML.
- Bringing research on Explainable ML into production use.
- Experiences with bringing ML techniques to production and scaling ML in production.
- Industry-specific best practices for ML production (e.g., edge computing/IoT, healthcare).
- Containerized ML workflows (e.g., Kubernetes for scale in ML).
- Use of Cloud APIs and Cloud based ML services in Production, experience and best practices.

Speaker Information

To see the details of what we want to know about your proposal to speak, we encourage you to view the talks submission system.

If you are a new presenter or would just like some extra help, please reach out. We can provide support via practice sessions, review of your pre-recorded videos, and a helpful guide on how to record your presentation to get the best results possible.

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.

If you have questions about this Call for Participation, feel free to drop us a message at srecon21chairs@usenix.org.

Background (Overarching goals of the world-wide 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 (<https://www.usenix.org/srecon#past>).

Conference Organizers

Program Co-Chairs

Daria Barteneva, Microsoft
 Frances Rees, Google
 Heidi Waterhouse, LaunchDarkly
 Vanessa Yiu, Goldman Sachs

Program Committee

Maria Acosta, Fermi National Accelerator Laboratory
 Blake Bisset, Microsoft
 Matt Brown, Google
 Jonah Horowitz, Apple
 Matthew Huxtable, Ziglu
 Adam Iwanicki, Google
 Jordan Li, Goldman Sachs
 Ana Margarita Medina, Gremlin
 Effie Mouzeli, Wikimedia Foundation
 Niall Murphy, relyAblit.ie
 Laura Nolan, Slack
 Herve Quiroz, Google
 Gráinne Sheerin, Google
 Harpreet Singh, DBS
 Stig Sorensen, Bloomberg L.P.
 Mohit Suley, Microsoft
 Nisha Talagala, Pyxeda AI
 Javier Turegano, Slack
 Avleen Vig, Facebook Inc.
 Jesse Ward, Salesforce
 Jamie Wilkinson, Google
 Jen Wohlner, DataDog
 Joel Young, LinkedIn
 Denise Yu, GitHub

Steering Committee

Kurt Andersen, Blameless
 Narayan Desai, Google
 Sabrina Farmer, Google
 Andrew Fong, Dropbox
 Liz Fong-Jones, Honeycomb
 Casey Henderson, USENIX Association
 Xiao Li, LinkedIn
 John Looney, Facebook Inc
 Laura Nolan, Slack
 Emil Stolarsky
 Avleen Vig, Facebook Inc
 Fernanda Weiden