

# Monday, June 12, 2017

## Kino

09:30—09:45 Welcome Speech

## Kino

09:45—10:30

## Keynote

Free and Open Source Software Today—Karen Sandler

### Kesselhaus

### Maschinenhaus

### Palais Atelier

### Moon Lounge

10:30—11:00

Coffee Break at Palais

Coffee Break at Palais

Coffee Break at Palais

Coffee Break at Palais

11:00—11:40

**scale** » Intermediate  
**Apache Spark? If only it worked**—Marcin Szymaniuk

**scale** » Intermediate  
**Distributed and Native Hybrid optimizations for Machine Learning Workloads**—Suneel Marthi

**search** » Intermediate  
**Apache Solr Learning to Rank For The Win!**—Diego Ceccarelli, Michael Nilsson

**stream** » Intermediate  
**Using Apache Beam to create a unified benchmarking framework...**—Ismael Melja, Jean-Baptiste Onofré

11:50—12:10

**search** » Intermediate  
**Apache Lucene 7 – What's coming next?**—Uwe Schindler

**scale** » Intermediate  
**Driving digital innovation through a modern approach to data analytics**—Gopi Krishnamurthy

**stream** » Beginner  
**How to build a recommendation system overnight**—Raam Rosh-Hai

**stream** » Intermediate  
**The History and Future of the Blurring of Stream Processing & OLTP**—John Hugg

12:20—13:00

**scale** » Intermediate  
**Integration Patterns for Big Data Applications**—Michael Häussler

**store** » Beginner  
**Challenges of Monitoring Distributed Systems**—Nenad Bozic

**search** » Intermediate  
**How does a Lucene Query actually work?**—Alan Woodward

**search** » Intermediate  
**Embracing diversity: searching over multiple Languages**—Tommaso Teofili, Suneel Marthi

13:00—14:00

Lunch Break at Palais

Lunch Break at Palais

Lunch Break at Palais

Lunch Break at Palais

14:00—14:40

**search** » Intermediate  
**BM25 is so yesterday**—Grant Ingersoll

**scale** » Intermediate  
**365 days of Spark!**—Mathieu Bastien

**stream** » Intermediate  
**Containerizing Distributed Pipes**—Hagen Tönnies

**scale** » Intermediate  
**Feeding Word2vec with tens of Billions of items...**—Simon Dollé

14:50—15:10

**scale** » Advanced  
**Aggregating, reporting, and visualizing 10s of billions of records...**—Frank Conrad

**scale** » Intermediate  
**Design Patterns for Calculating User Profiles in Real Time**—Igor Mazor

**scale** » Beginner  
**10 Ways to Ally for Women in Technology**—Heather VanCura

**stream** » Intermediate  
**Anti-Spam and Machine Learning at Lovoo**—Juan de Dios Santos Rivera

15:20—16:00

**stream** » Intermediate  
**Experiences running Flink at Very Large Scale**—Stephan Ewen

**store** » Intermediate  
**Building a fence around your Hadoop environments**—Lars Francke

**scale** » Beginner  
**Rolling out Enterprise Kubernetes Clouds at SAP**—Thomas Fricke, Stefan Vetter

**search** » Intermediate  
**Migrating a Real-time News Classification Engine to Lucene**—Marvin Justice

16:00—16:30

Coffee Break at Palais

Coffee Break at Palais

Coffee Break at Palais

Coffee Break at Palais

16:30—17:10

**scale** » Intermediate  
**Metaphors we Compute By**—Alvaro Videla

**stream** » Beginner  
**Queryable State or How to Build a Billing System...**—Maximilian Bode, Konstantin Knauf

**search** » Intermediate  
**Running slow queries with Lucene**—Adrien Grand

**store** » Intermediate  
**One Database to Rule 'em All**—Stefanie Janine Stölting

17:20—18:00

**search** » Intermediate  
**Building a Vibrant Search Ecosystem at Bloomberg**—Ken LaPorte

**stream** » Beginner  
**A big data streaming recipe**—Konstantin Gregor

**scale** » Advanced  
**Optimization of Public Cloud for Efficiency and Scale**—Erik Onnen

**stream** » Intermediate  
**Balancing Heroes and Pokemon in Real Time**—Fokko Driesprong/ Vincent Warmerdam

From 18:00

Get-together at Palais