






FOSDEM'17 - HPC, Big Data & Data Science devroom

Event	Speakers	Start	End
Saturday			
Opening	Vasia Kalavri	10:30	10:35
Portability of containers across diverse HPC resources with Singularity	Michael Bauer, César Gómez-Martín	10:35	11:00
The birth of HPC Cuba <i>How supercomputing is being made available to all Cuban researchers using FOSS</i>	Dieter Roefs, Hector Cruz Enriquez	11:00	11:25
Optimized and reproducible HPC Software deployment <i>with free software and GNU Guix</i>	Ludovic Courtès, Pjotr Prins	11:30	11:55
Reproducible HPC Software Installation on Cray Systems with EasyBuild	Guilherme Peretti-Pezzi	12:00	12:25
Putting Your Jobs Under the Microscope using OGRT	Georg Rath	12:30	12:55
Dask - extending Python data tools for parallel and distributed computing	Joris Van den Bossche	13:00	13:25
Purely Functional GPU Programming with Futhark	Troels Henriksen	13:30	13:55
 The Marriage of Cloud, HPC and Containers	Adam Huffman	14:00	14:10
 Quickstart Big Data	Olaf Flebbe	14:10	14:20
 Extending Spark Machine Learning Pipelines <i>Going beyond wordcount with Spark ML</i>	Holden Karau	14:25	14:35
 Using BigBench to compare Hive and Spark versions and features <i>BigBench in Hive and Spark</i>	Nicolas Poggi, Alejandro Montero	14:35	14:45
 Making Wiki Gardening Tasks Easier Using Big Data and NLP	Bee Padalkar	14:45	14:55
A field guide to the machine learning zoo	Theodore Vasiloudis	15:00	15:25
Intelligently Collecting Data at the Edge <i>Intro to Apache NiFi and MiNiFi</i>	Andy LoPresto	15:30	15:55
Postgres MPP Data Warehousing joins Hadoop ecosystem <i>Making two elephants dance</i>	Roman Shaposhnik	16:00	16:25
BigPetStore on Spark and Flink <i>Implementing use cases on unified Big Data engines</i>	Marton Balassi	16:30	16:55
Democratizing Deep Learning with Tensorflow on Hops Hadoop	Jim Dowling, Gautier Berthou	17:00	17:25
Kafka Streams and Protobuf <i>stream processing at trivago</i>	Clemens Valiente	17:30	17:55
Not less, Not more. Exactly Once Large-Scale Stream Processing in Action.	Paris Carbone	18:00	18:25
Why you should care about SQL for big data and how Apache Calcite can help	Christian Tzolov	18:30	18:55
Closing	Vasia Kalavri	18:55	19:00

Supported by:

