



[https://fosdem.org/2018/schedule/track/hpc,\\_big\\_data,\\_and\\_data\\_science](https://fosdem.org/2018/schedule/track/hpc,_big_data,_and_data_science)

Installing software for scientists on a multi-user HPC system <i>A comparison between conda, EasyBuild, Guix, Nix &amp; Spack</i>	Kenneth Hoste	09:00	09:25
Binary packaging for HPC with Spack	Todd Gamblin	09:30	09:55
Tying software deployment to scientific workflows <i>Using Guix to make software deployment a first-class citizen</i>	Ludovic Courtès	10:00	10:25
Combining CVMFS, Nix, Lmod, and EasyBuild at Compute Canada	Bart Oldeman	10:30	10:55
Behind the scenes of a FOSS-powered HPC cluster at UCLouvain <i>Ansible or Salt? Ansible AND Salt!</i>	Damien François	11:00	11:25
How DeepLearning can help to improve geospatial DataQuality , an OSM use case.	Olivier Courtin	11:30	11:55
Modules v4 <i>Pushing forward user environment management</i>	Xavier Delaruelle	12:00	12:10
Scale Out and Conquer: Architectural Decisions Behind Distributed In-Memory Systems	Akmal Chaudhri	12:10	12:20
The Magnificent Modular Mahout <i>An extensible library for distributed math and HPC</i>	Trevor Grant	12:20	12:30
Tools for large-scale collection and analysis of source code repositories <i>Open source Git repository collection pipeline</i>	Alexander Bezzubov	12:30	12:40
Slurm in Action: Batch Processing for the 21st Century	Georg Rath	12:40	12:50
The Julia programming language <i>A look at arrays and Trilinos integration</i>	Bart Janssens	13:00	13:25
Does data security rule out high performance? <i>Is it possible to maintain the anonymity of people contributing their data while not hindering ambitious science?</i>	Adam Huffman	13:30	13:55
CrateDB: A Search Engine or a Database? Both! <i>How We Built a SQL database on top of Elasticsearch and Lucene</i>	Maximilian Michels	14:00	14:25
Scaling Deep Learning to hundreds of GPUs on HopsHadoop	Fabio Buso	14:30	14:55
AI on Microcontrollers <i>uTensor brings Deep-Learning to MCUs</i>	Neil Tan	15:00	15:25
Productionizing Spark ML Pipelines with the Portable Format for Analytics	Nick Pentreath	15:30	15:55
Accelerating Big Data Outside of the JVM <i>Adventures with Apache Arrow, Spark, BEAM, Flink, and friends</i>	Holden Karau	16:00	16:25
Nexmark A unified benchmarking suite for data-intensive systems with Apache Beam	Ismaël Mejía	16:30	16:55

lightning talks

Supported by:

