

LOGICAL CLOCKS

ENTERPRISE HOPSWORKS

Hopsworks is a full-stack platform for Data Science built on Hops, the world's fastest and most scalable Hadoop distribution. Hopsworks has unique support for project-based multi-tenancy, scale-out ML pipelines and managed GPUs-as-a-resource.

MULTI-TENANCY

Hopsworks provides Projects as a privacy-by-design sandbox for data, including sensitive data, and for managing collaborating teams (like GitHub). Datasets can be shared between projects (like Dropbox). Each Project has its own Conda environment, enabling Python dependencies to be managed by the developers themselves.

HOPS

Hops is a next-generation distribution of Apache Hadoop that supports distributed metadata. Our improved HDFS, called HopsFS, supports both 16X higher throughput than HDFS* and with NVMe disks an order of magnitude lower latency for reading/writing small files**.

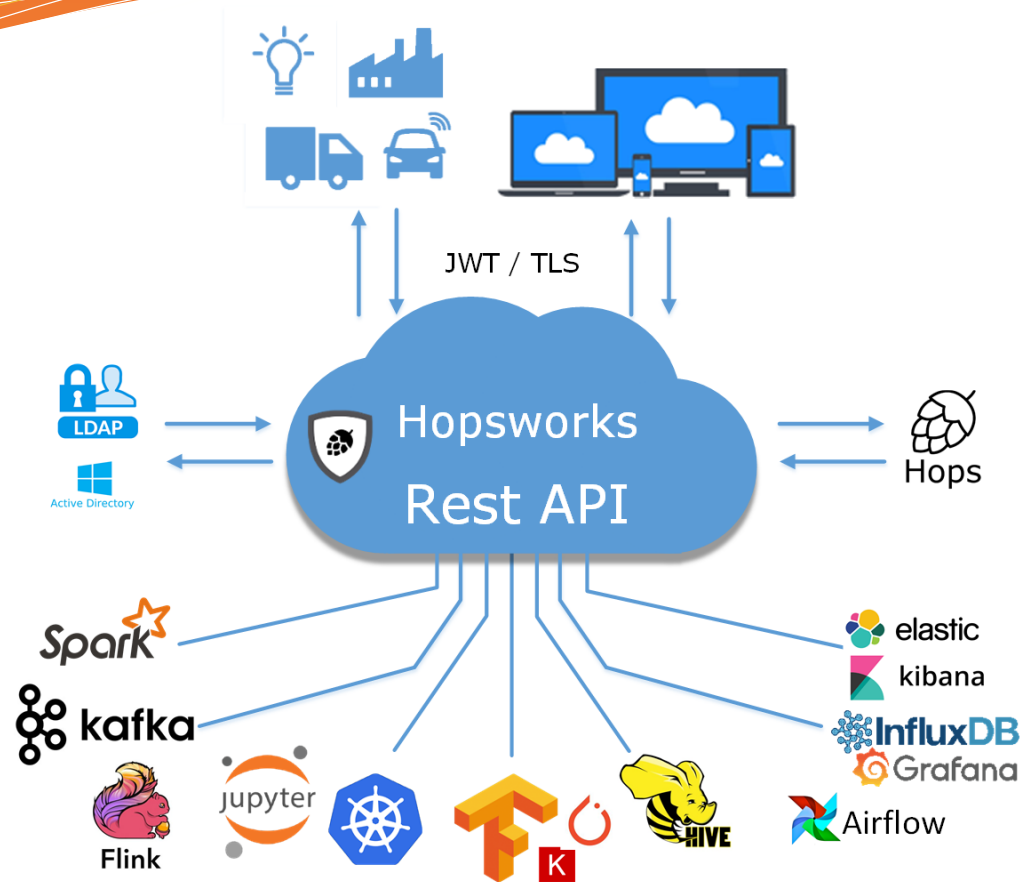
HOPSMML

HopsML is our framework for writing end-to-end machine learning workflows in Python. We support Airflow to orchestrate workflows with: ETL in PySpark or TensorFlow, a Feature Store, hyperparameter optimization parallelized over many GPUs in Keras/TensorFlow/PyTorch, and distributed training on many GPUs. Jupyter notebooks can be used to write parts of the pipeline, and TensorBoard used to visualize experiment runs. Models can be deployed in Kubernetes (built-in or external) and monitored in production using Kafka/Spark-Streaming.

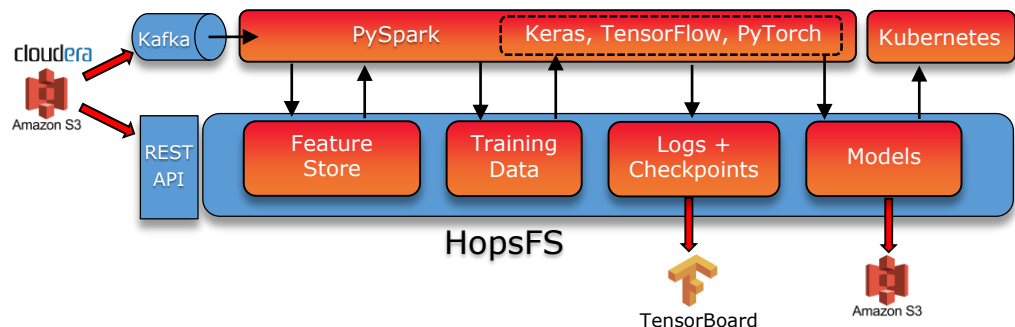
TLS SECURITY

Uniquely in Hadoop, Hops supports TLS certificates for authentication: users, services, jobs and for in-flight encryption. At-rest encryption is also supported using ZFS-on-Linux.

Hopsworks®



Scale-out ML Workflows



Silicon Valley Office
470 Ramona St
Palo Alto
California,
USA

Stockholm Office
Box 1263,
Isafjordsgatan 22
Kista,
Sweden

UK Office
IDEALondon,
69 Wilson St
London, EC2A2BB,
U.K

For more information, visit us on the Web at:

www.logicalclocks.com

*<https://www.usenix.org/conference/fast17/technical-sessions/presentation/niazi>
**<http://2018.middleware-conference.org/index.php/accepted-papers/>