

# Apache Spark and Its Role in the Enterprise Data Hub

Mike Olson, Chief Strategy Officer, Cloudera  
[mike.olson@cloudera.com](mailto:mike.olson@cloudera.com), [@mikeolson](https://twitter.com/mikeolson)



# Spark Unifies and Simplifies Hadoop

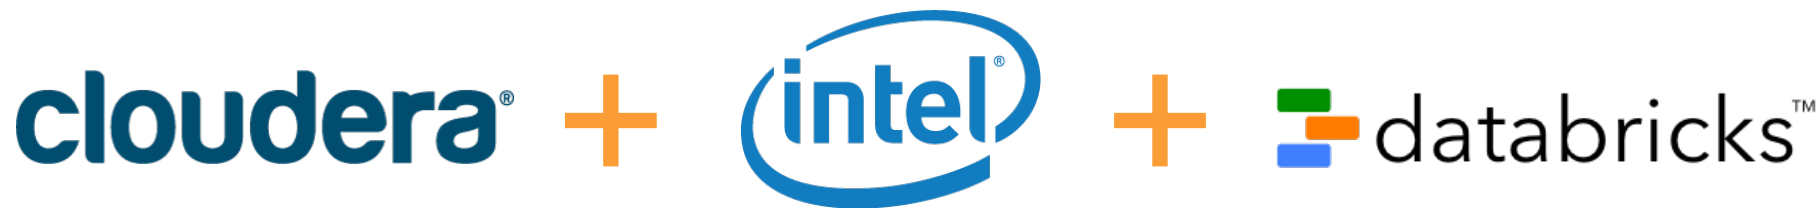
[illegible]

Batch Processing

Stream Processing

Machine Learning

The Hadoop logo, featuring the word "hadoop" in a stylized, lowercase, blue font with a black outline, set against a yellow background.

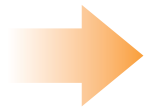


Developing and supporting Spark together  
to ensure customer success

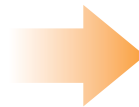
# Spark at Cloudera

---

October  
2013



February  
2014



July  
2014

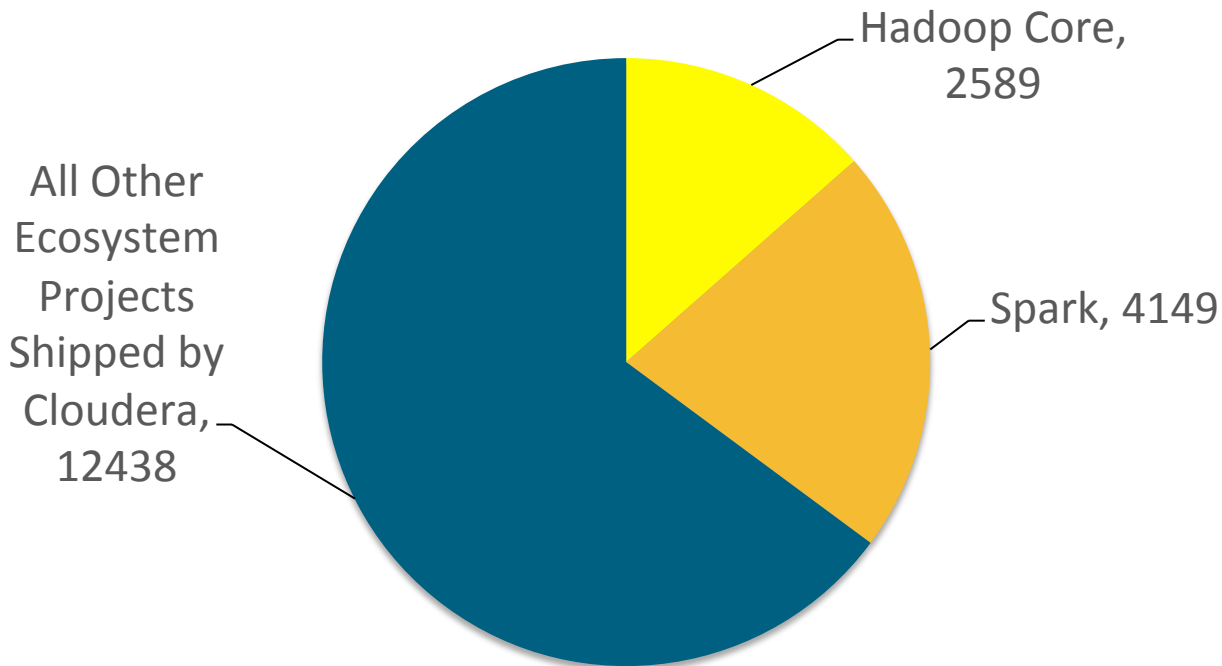
Databricks and  
Cloudera  
partner

Spark support  
added to CDH

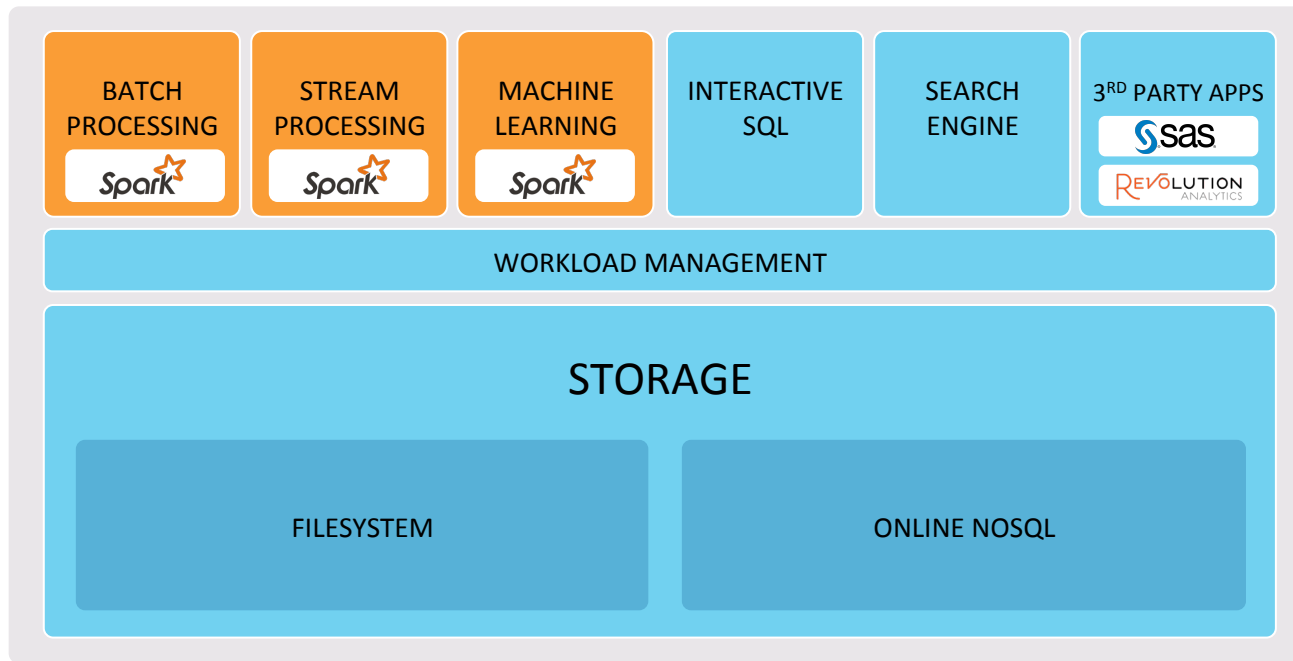
Continuing  
support &  
innovation

# Spark is a Core Component of Hadoop

## Commit Activity Past 12 Months



# Fully Integrated into CDH



- Integrated and supported part of our platform
- Diverse use cases in production
- Well-trained support and external trainings

# Customer Adoption



CONCUR



RelayHealth



Search personalization  
through machine  
learning investigations



Fast processing of  
millions of stock  
positions and future  
scenarios



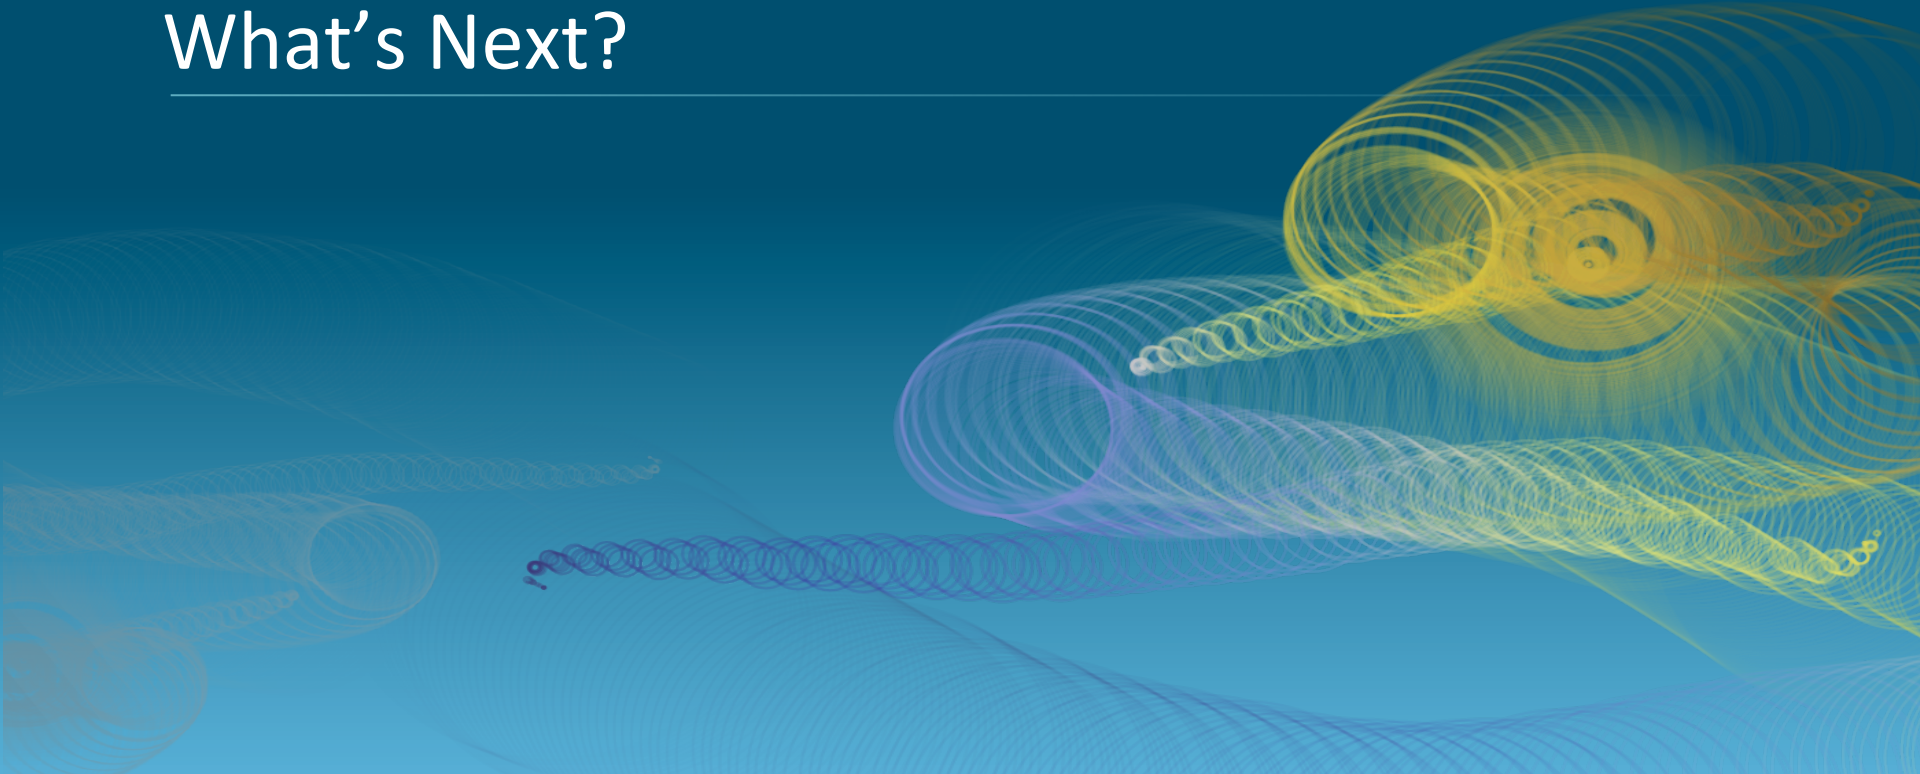
Genomics research  
using Spark pipelines



Predictive modeling  
of disease conditions

# What's Next?

---



# Cloudera Developer Training for Apache Spark

---

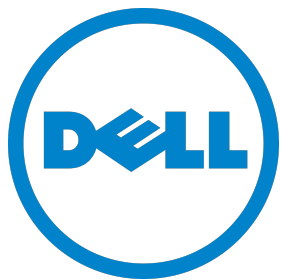
The only hands-on deep dive  
into building unified  
applications with Spark



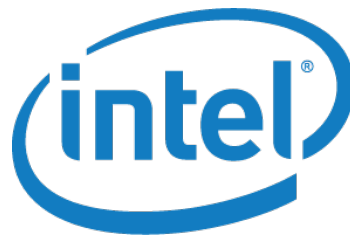
Public GA: Aug 5, Redwood City

# Dell In-Memory Appliances for Cloudera Enterprise

- Simplifies and speeds up complex cluster deployments
- Includes Cloudera Enterprise and ScaleMP's Versatile SMP (vSMP) architecture
- Built on the Intel(R) Xeon(R) processor-based Dell R920 hardware
- Optimized for Spark



**cloudera®**



# Spark as the Standard Processing Engine



Spark



Spark



Spark



Spark

# Bringing the Communities Together

The Hive and Spark communities are coming together to drive consolidation in the Hadoop ecosystem



# Hive on Spark

cloudera®

IBM

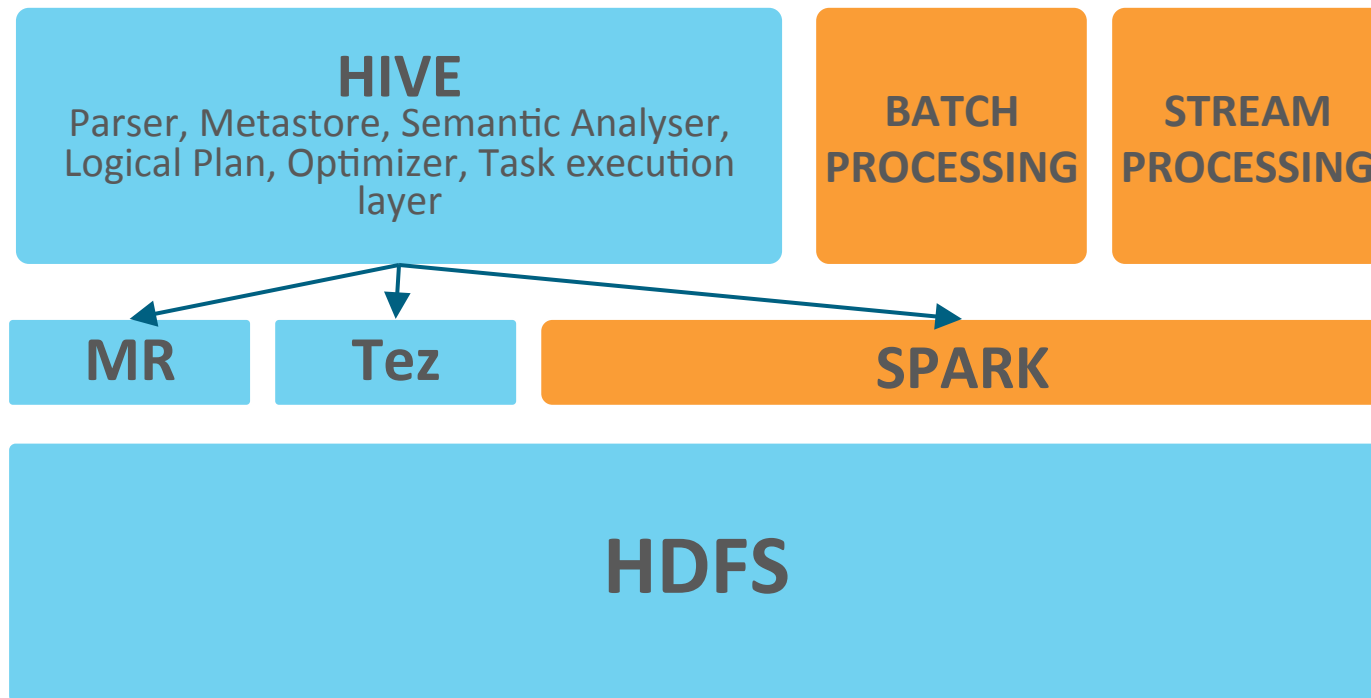


 databricks™

MAPR®

 intel®

# Architecture



# Our SQL on Hadoop Vision



BI and SQL  
Analytics



Batch  
Processing



Mixed Spark and  
SQL Applications



**cloudera®**

Ask Bigger Questions

Thank you!

---

Mike Olson

[mike.olson@cloudera.com](mailto:mike.olson@cloudera.com)

[@mikeolson](#)