A look ahead at Spark's development

Reynold Xin @rxin

Spark Summit EU, Amsterdam

Oct 29th, 2015

databricks™

Spark stack diagram

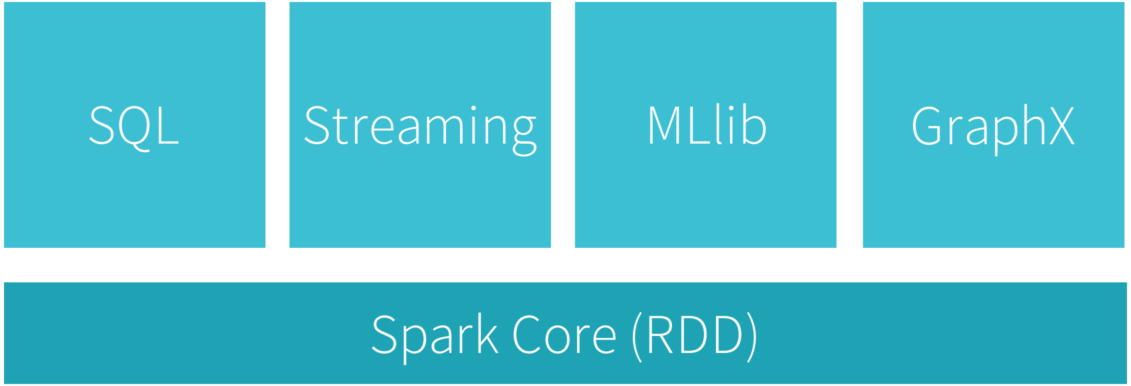

Spark stack diagram (a different take)

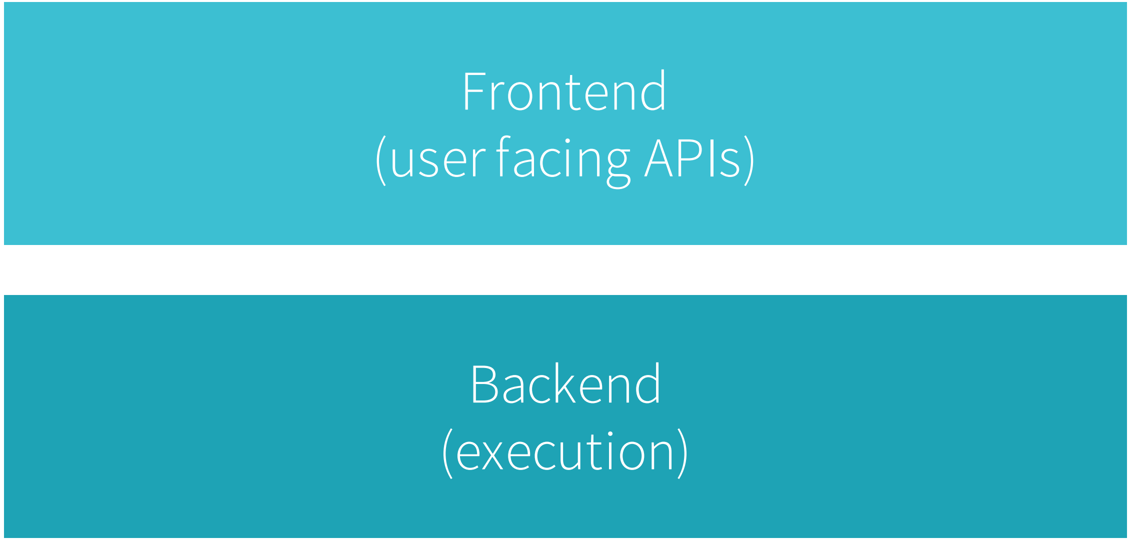

Spark stack diagram
(a different take)

Frontend
(RDD, DataFrame, ML pipelines, …)

Backend
(scheduler, shuffle, operators, …)

Last 12 months of Spark evolution

Frontend

DataFrames

Data sources

R

Machine learning pipelines

. . .

Backend

Project Tungsten

Sort-based shuffle

Netty-based network

. . .

Last 12 months of Spark evolution

Frontend

DataFrames

Data sources

R

Machine learning pipelines

. . .

Backend

Project Tungsten

Sort-based shuffle

Netty-based network

. . .

DataFrame:
A Frontend Perspective

Spark DataFrame

Scalable data frame for Java, Python, R, Scala

Similar APIs as single-node tools (Pandas, dplyr), i.e. easy to learn

```
> head(filter(df, df$waiting < 50))  # an example in R
##  eruptions waiting
##1     1.750      47
##2     1.750      47
##3     1.867      48
```

Spark RDD Execution

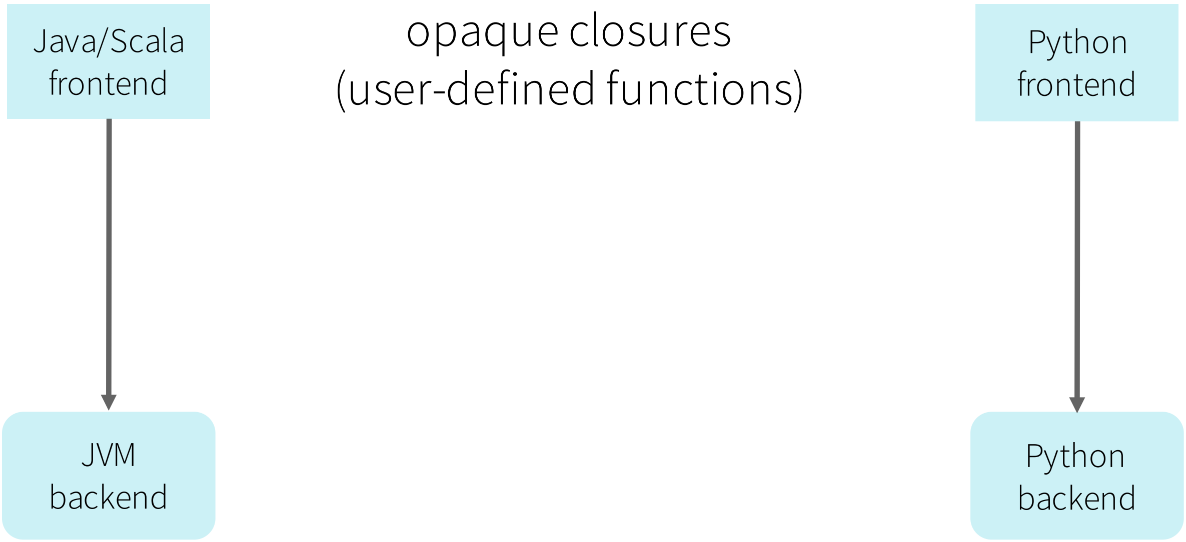

Spark DataFrame Execution

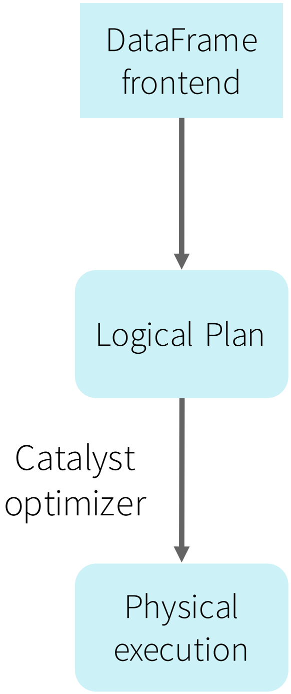

Intermediate representation for computation

Spark DataFrame Execution

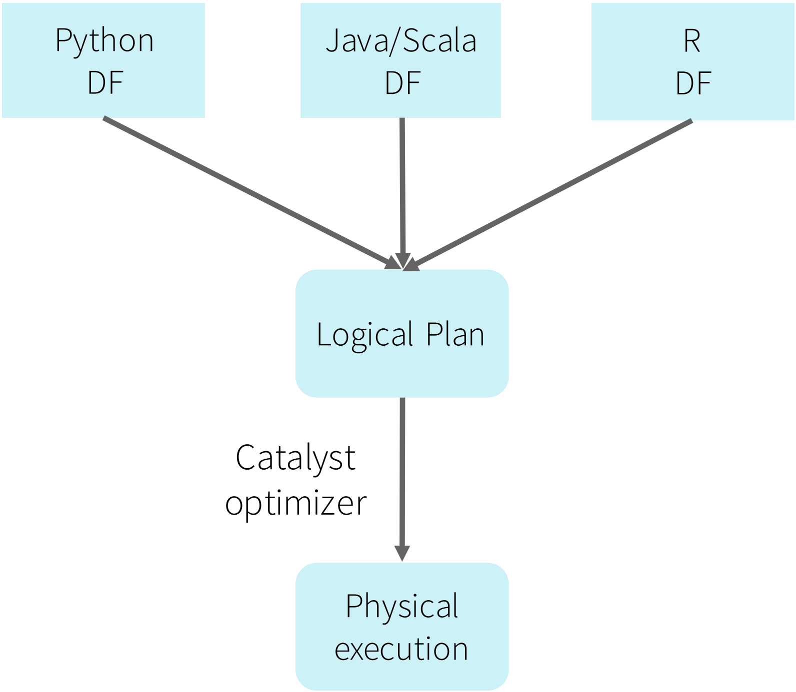

Simple wrappers to create logical plan

Intermediate representation for computation

Benefit of Logical Plan: Simpler Frontend

Python : ~2000 line of code (built over a weekend)

R : ~1000 line of code

i.e. much easier to add new language bindings (Julia, Clojure, …)

Performance

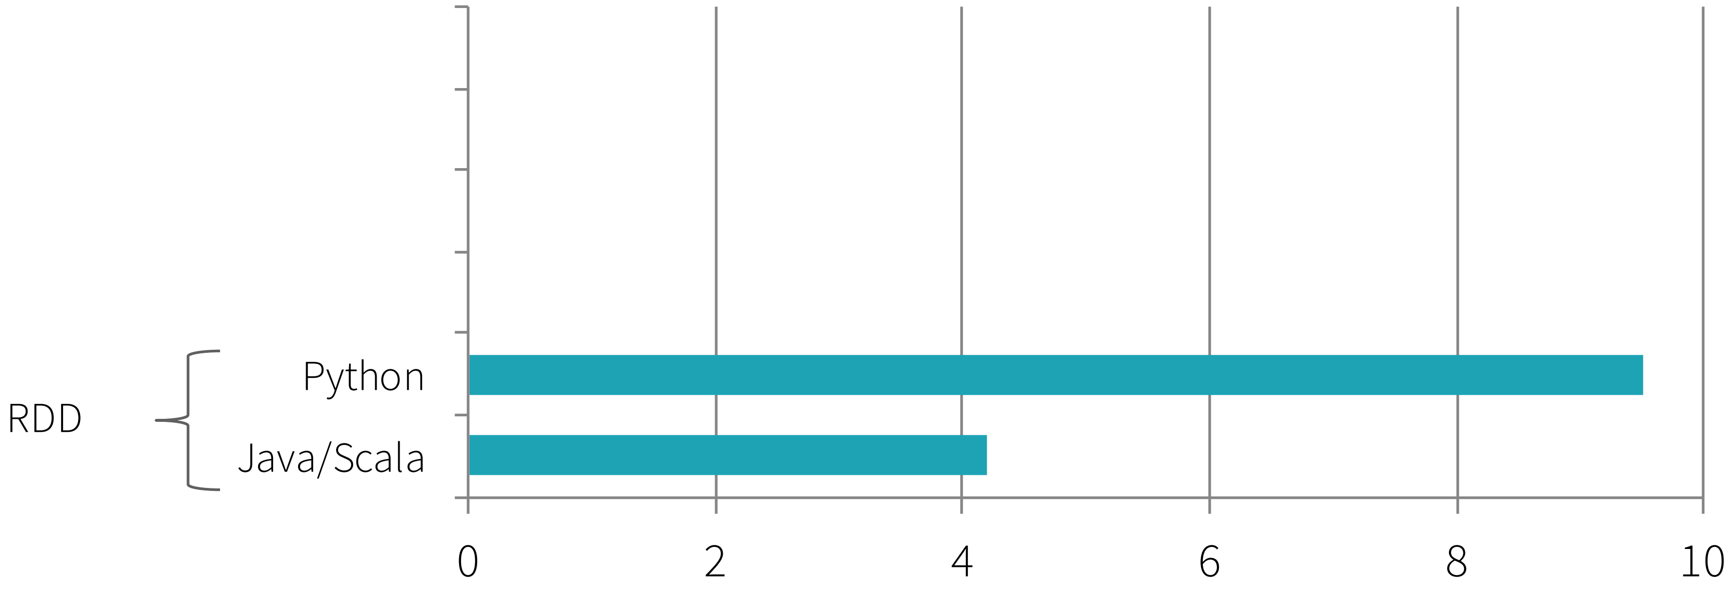

Runtime for an example aggregation workload

Benefit of Logical Plan: Performance Parity Across Languages

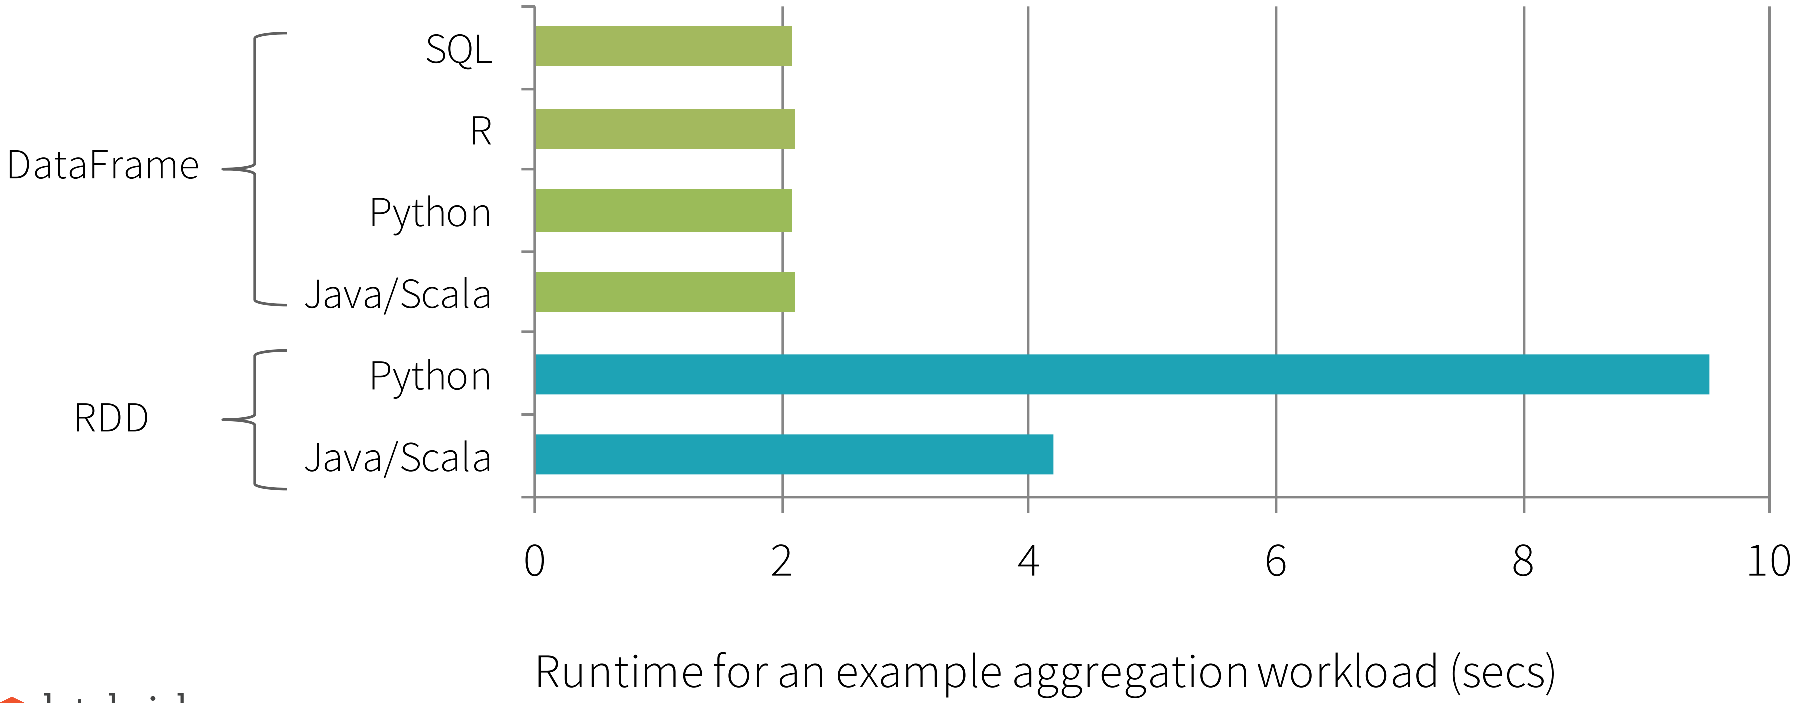

databricks

Tungsten:
A Backend Perspective

Hardware Trends

Storage

Network

CPU

Hardware Trends

	2010
Storage	50+MB/s (HDD)
Network	1Gbps
CPU	~3GHz

Hardware Trends

	2010	2015
Storage	50+MB/s (HDD)	500+MB/s (SSD)
Network	1Gbps	10Gbps
CPU	~3GHz	~3GHz

Hardware Trends

	2010	2015	
Storage	50+MB/s (HDD)	500+MB/s (SSD)	10X
Network	1Gbps	10Gbps	10X
CPU	~3GHz	~3GHz	☹

Project Tungsten

Substantially speed up execution by optimizing CPU efficiency, via:

(1) Runtime code generation
(2) Exploiting cache locality
(3) Off-heap memory management

From DataFrame to Tungsten

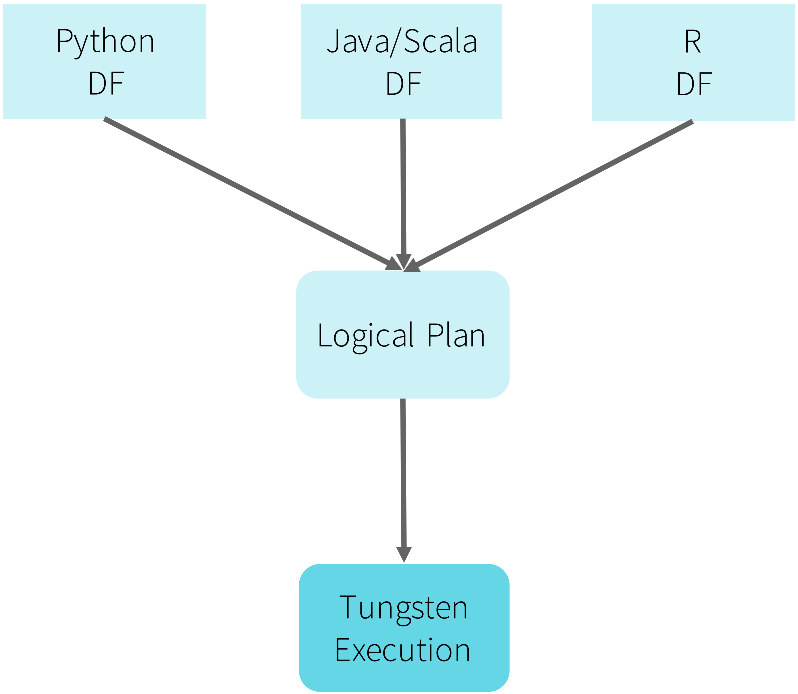

Initial phase in Spark 1.5

More work coming in 2016

3 Things to Look Forward To

Dataset API in Spark 1.6

Typed interface over DataFrames / Tungsten

```
case class Person(name: String, age: Int)

val dataframe = read.json(“people.json”)
val ds: Dataset[Person] = dataframe.as[Person]

ds.filter(p => p.name.startsWith(“M”))
      .groupBy(“name”)
      .avg(“age”)
```

Dataset

"Encoder" to specify type information
so Spark can translate it into DataFrame
and generate optimized memory layouts

Checkout SPARK-9999

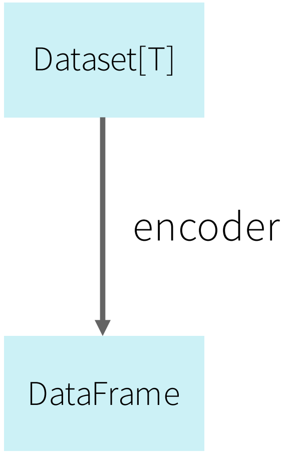

Streaming DataFrames

Easier-to-use APIs (batch, streaming, and interactive)

```
val stream = read.kafka("...")
stream.window(5 mins, 10 secs)
  .agg(sum("sales"))
  .write.jdbc("mysql://...")
```

And optimizations:

- Tungsten backends
- native support for out-of-order data
- data sources and sinks

Microsoft Azure

SALES 1-800-867-1389 MY ACCOUNT

Why Azure Products Documentation Pricing Partners Blog Resources Support

BLOG > ANNOUNCEMENTS , VIRTUAL MACHINES

Largest VM in the Cloud

THURSDAY, JANUARY 8, 2015

DREW MCDANIEL
Principal Program Manager, Azure

G-Series Size Details

VM Size	Cores	RAM		
Standard_G1	2	2		
Standard_G2	4	5		
Standard_G3	8	11		
Standard_G4	16	22		
Standard_G5	32	448 GiB	6596 GB	64

AWS Announces X1 Instances For EC2 With 2TB Of Memory, Launching Next Year

Posted Oct 8, 2015 by Frederic Lardinois (@fredericl)

926 SHARES

Amazon today announced a massive new instance type for its AWS EC2 compute service. The

CrunchBase

Amazon

FOUNDED
1994

OVERVIEW
Amazon is an e-commerce retailer forme
to provide consumers with products in tw
It offers users with merchandise and cont
purchased for resale from vendors and th
by third-party sellers. Operating in North
and international markets, Amazon provi
services through websites such as amazo
amazon.ca. It also enables authors, music
filmmakers, ...

LOCATION
Seattle, WA

CATEGORIES
E-Commerce, Crowdsourcing, Groceries, Co
Goods, Delivery, Software, Retail, Internet

FOUNDERS

3D XPoint

- DRAM latency
- SSD capacity
- Byte addressible

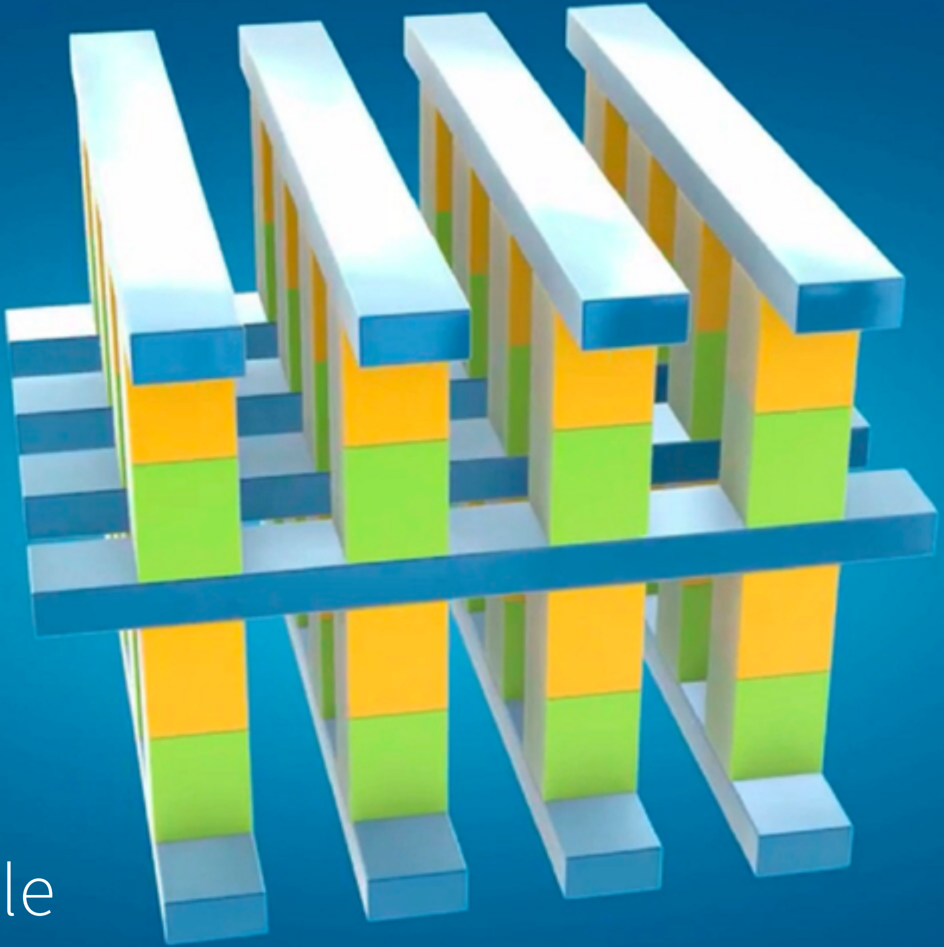

Unified API, One Engine, Automatically Optimized

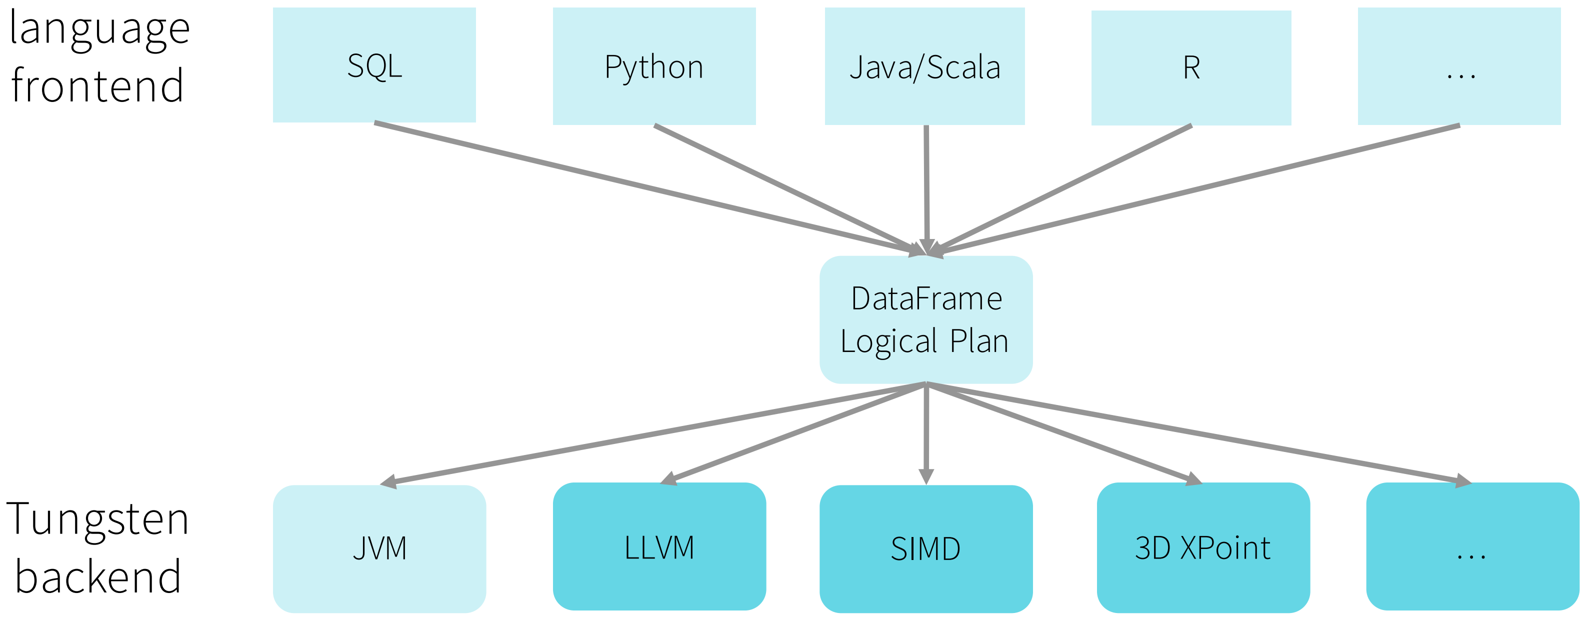

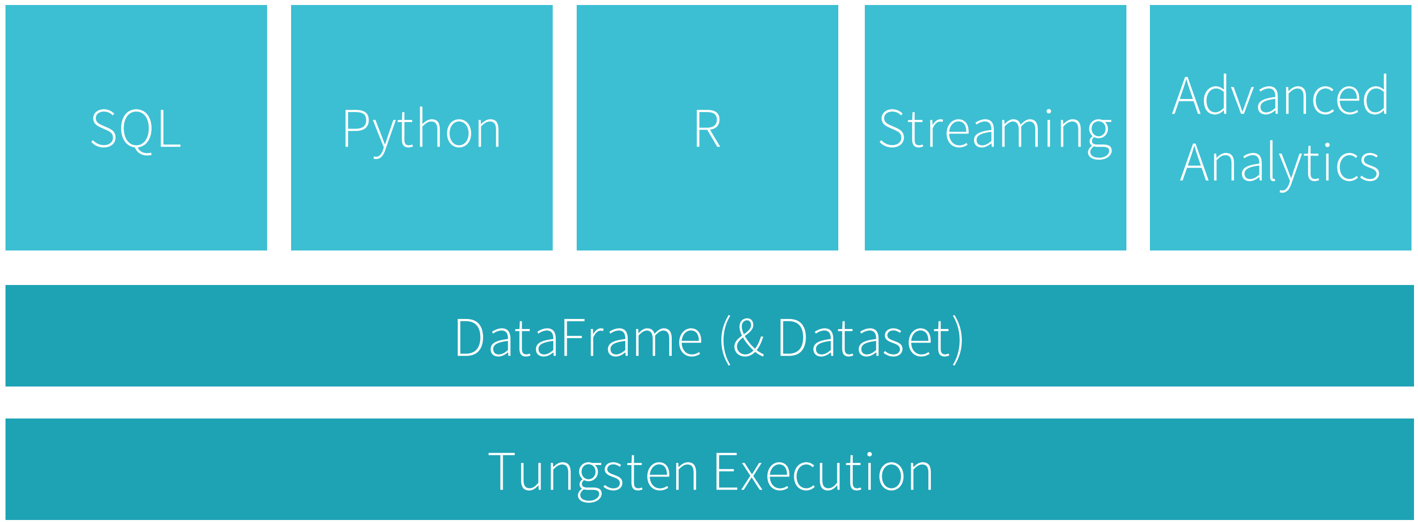
SQL
Python
R
Streaming
Advanced Analytics
DataFrame (& Dataset)
Tungsten Execution

Office Hours Today @ Databricks booth

	Topic Area
10:30 – 11:30	Spark general (Reynold)
13:00 – 14:00	R and data science (Hossein)
13:30 – 14:30	machine learning (Joseph)
14:00 – 15:00	Spark, YARN, etc (Andrew)