SQL-on-Everything with Apache Drill

Julien Le Dem, Principal Architect at Dremio
VP Apache Parquet
Apache Pig PMC
julien@dremio.com | @J_

Big Data Apps meetup
January 27, 2016

© 2015. All Rights Reserved.

From Proprietary to Open Source…

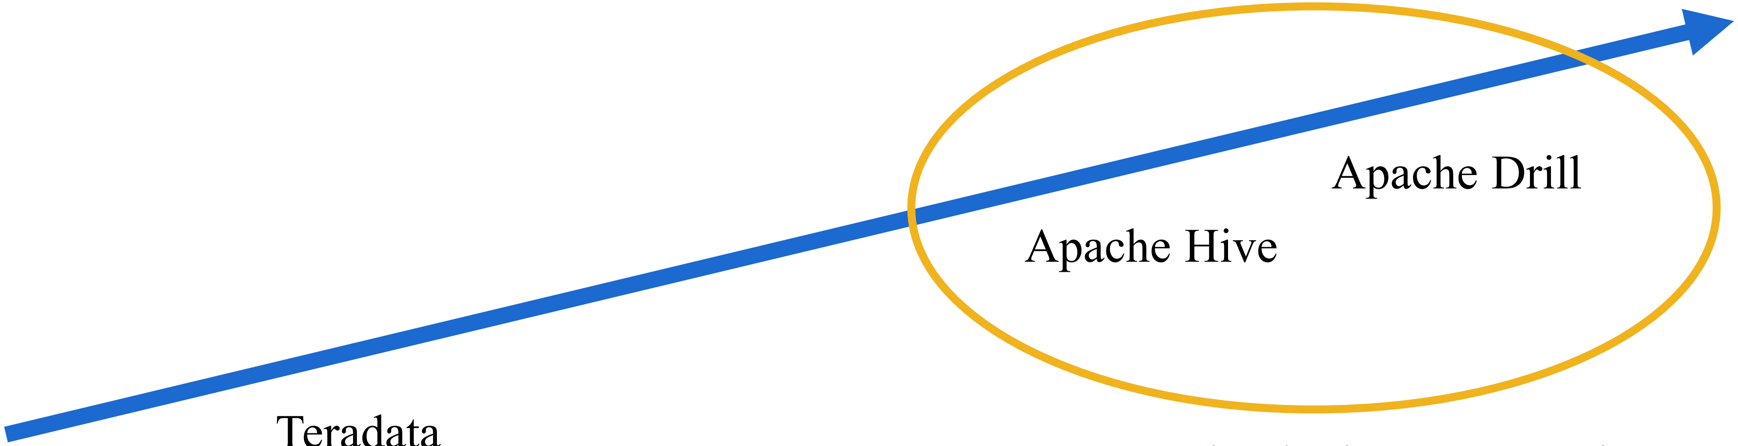

Apache Drill

Apache Hive

Teradata
Oracle
SQL Server
…

Open source technologies are emerging as the preferred approach for large-scale SQL-based analytics

© 2015. All Rights Reserved.

Apache Drill: Open Source Schema-Free SQL Engine

CLI

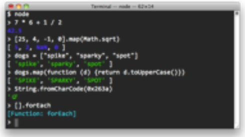

```
$ node
> 7 * 6 + 1 / 2
42.5
> [25, 4, -1, 0].map(Math.sqrt)
[ 5, 2, NaN, 0 ]
> dogs = ["spike", "sparky", "spot"]
[ 'spike', 'sparky', 'spot' ]
> dogs.map(function (d) {return d.toUpperCase()})
[ 'SPIKE', 'SPARKY', 'SPOT' ]
> String.fromCharCode(0x263a)
'☺'
> [].forEach
[Function: forEach]
>
```

Tableau, Excel, Qlik, …

Custom Applications

node js Microsoft .NET iOS python

Apache Drill

Kudu · MongoDB · HBase · Elasticsearch* · MapR · RDBMS · Phoenix

HDFS · NAS · Local Files · Amazon S3 · Azure Blob Storage

* Currently being developed

© 2015. All Rights Reserved.

Apache Drill: Open Source Schema-Free SQL Engine

Open Source Apache Project

- Contributors from many companies including Dremio, MapR and Hortonworks
- 3-year engineering effort, 200K+ lines of code

Innovative Schema-free Engine

- Point-and-query vs. schema-first
- No data loading, schemas or ETL
- Handles complex (eg, JSON) data natively

Extreme Scale & Performance

- Scales from one laptop to 1000s of servers
- High performance via columnar execution & dynamic query compilation

Extensible Architecture

- Pluggable high-speed datastore connectors (eg, MongoDB, Amazon S3)
- Custom operators and UDFs

© 2015. All Rights Reserved.

JSON Model, Columnar Speed

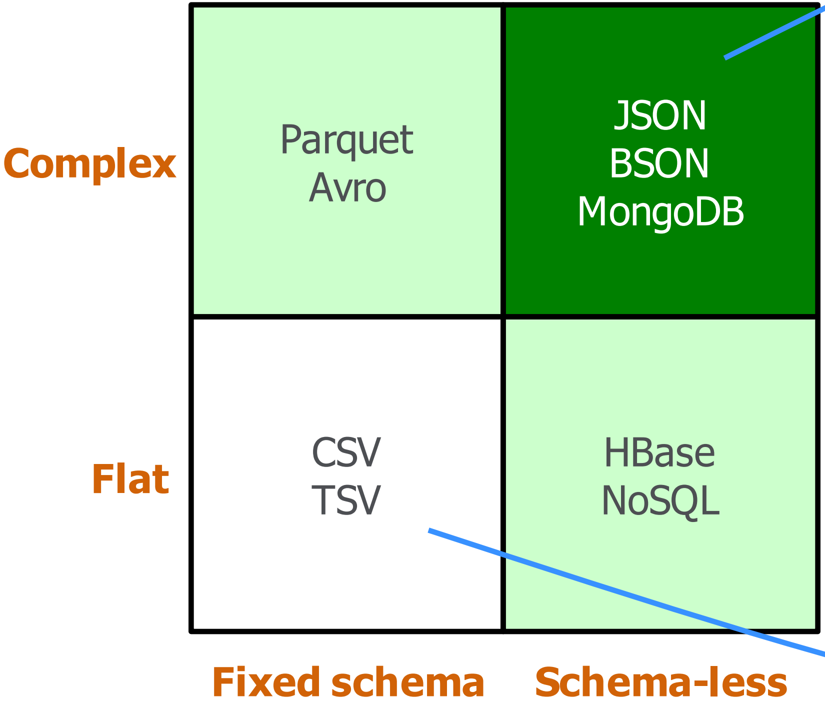

	Fixed schema	Schema-less
Complex	Parquet Avro	JSON BSON MongoDB
Flat	CSV TSV	HBase NoSQL

Apache Drill table

```
{
  name: {
    first: Michael,
    last: Smith
  },
  hobbies: [ski, soccer],
  district: Los Altos
}
{
  name: {
    first: Jennifer,
    last: Gates
  },
  hobbies: [sing],
  preschool: CCLC
}
```

RDBMS/SQL-on-Hadoop table

Name	Gender	Age
Michael	M	6
Jennifer	F	3

© 2015. All Rights Reserved.

Drill Supports *Schema Discovery On-The-Fly*

Schema Declared In Advance

- Fixed schema
- Leverage schema in centralized repository (Hive Metastore)

Schema Discovered On-The-Fly

- Fixed schema, evolving schema or schema-free
- Leverage schema in centralized repository or self-describing data

SCHEMA ON WRITE

ORACLE

Microsoft SQL Server

IBM DB2

SCHEMA BEFORE READ

HIVE

SCHEMA ON THE FLY

APACHE DRILL

© 2015. All Rights Reserved.

Apache Drill is Not Just SQL-on-Hadoop

	Drill	SQL-on-Hadoop (Hive, Impala, etc.)
Use case	Self-service, in-situ, SQL-based analytics	Teradata offload
Deployment model	Standalone or co-located with NoSQL/Hadoop	Hadoop service
User experience	Point-and-query	Ingest data → define schemas → query
Data model	Schema-free JSON (like MongoDB)	Relational (like Postgres)
Data sources	NoSQL, Cloud Storage, Hadoop, SaaS, local files (including multiple instances)	A single Hadoop cluster
Data management	Logical, by IT or end-users (self-service)	Physical, by IT only
1.0 availability	Q2 2015	Q2 2013 or earlier

© 2015. All Rights Reserved.

Omni-SQL (“SQL-on-Everything”)

> ***Drill: Omni-SQL***
>
> *Whereas the other engines we're discussing here create a relational database environment on top of Hadoop, Drill instead enables a SQL language interface to data in numerous formats, without requiring a formal schema to be declared. This enables plug-and-play discovery over a huge universe of data without prerequisites and preparation. So while Drill uses SQL, and can connect to Hadoop, calling it SQL-on-Hadoop kind of misses the point. A better name might be SQL-on-Everything, with very low setup requirements.*

Andrew Brust, ZDNet

© 2015. All Rights Reserved.

ARCHITECTURE

© 2015. All Rights Reserved.

Everything Starts With a Drillbit…

- High performance query executor
- In-memory columnar execution
- Directly interacts with data, acquiring knowledge as it reads
- Built to leverage large amounts of memory
- Networked or not
- Exposes ODBC, JDBC, REST
- Built-in Web UI and CLI
- Extensible

Single process (daemon or CLI)

© 2015. All Rights Reserved.

Data Lake, More Like Data Maelstrom

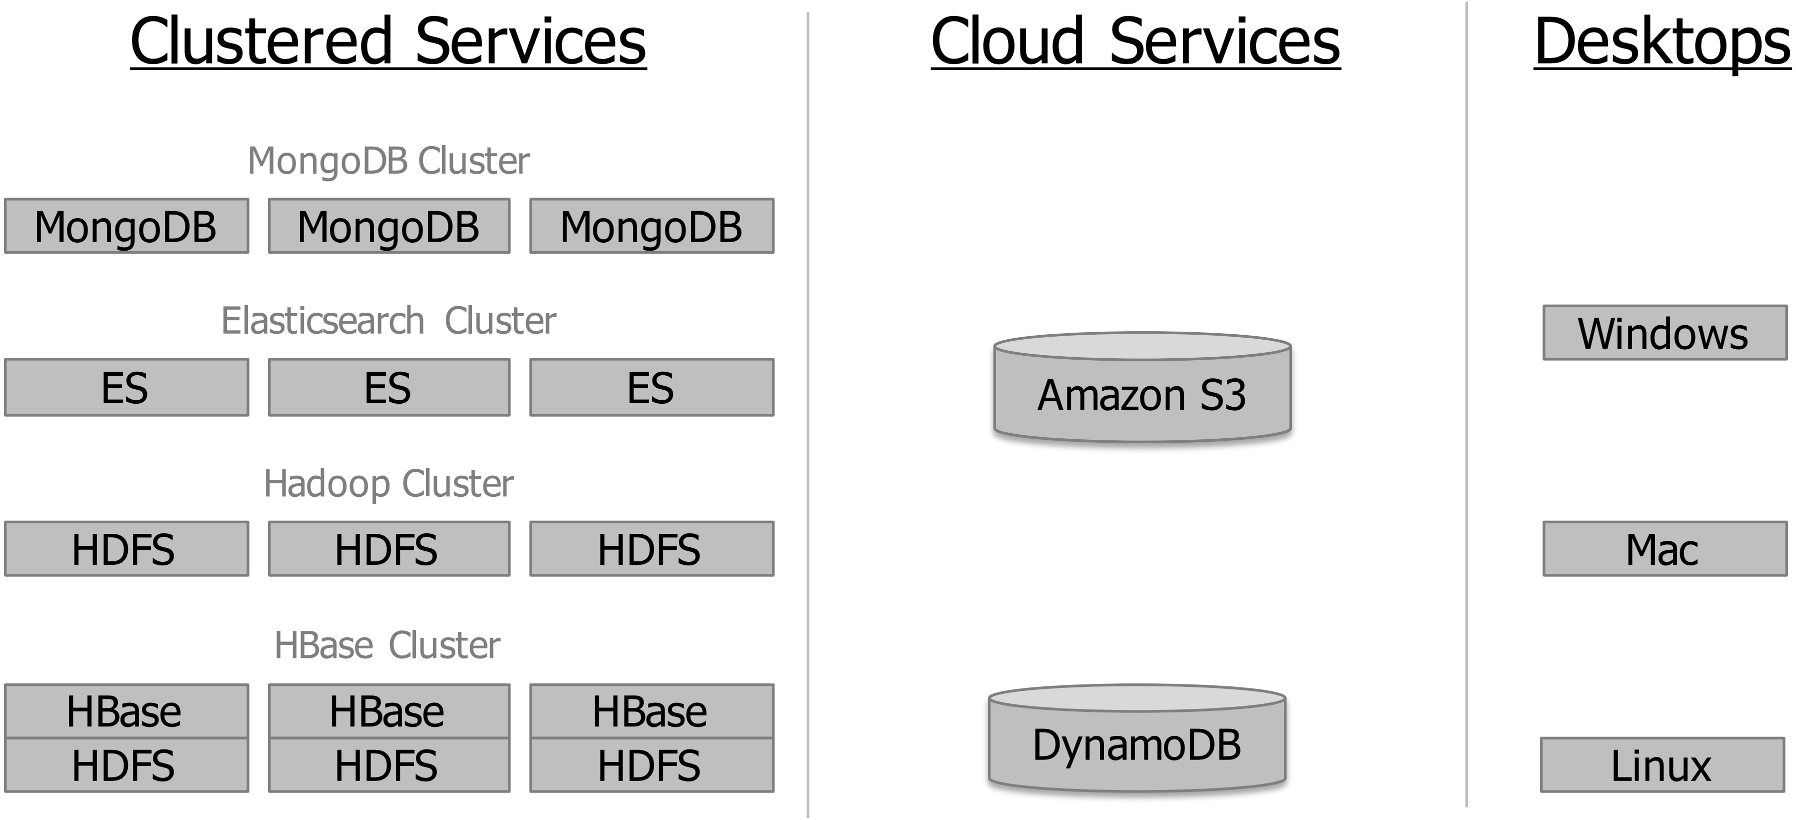

Clustered Services

MongoDB Cluster

MongoDB | MongoDB | MongoDB

Elasticsearch Cluster

ES | ES | ES

Hadoop Cluster

HDFS | HDFS | HDFS

HBase Cluster

HBase / HDFS | HBase / HDFS | HBase / HDFS

Cloud Services

Amazon S3

DynamoDB

Desktops

Windows

Mac

Linux

© 2015. All Rights Reserved.

Run Drillbits Wherever; Whatever Your Data

Clustered Services

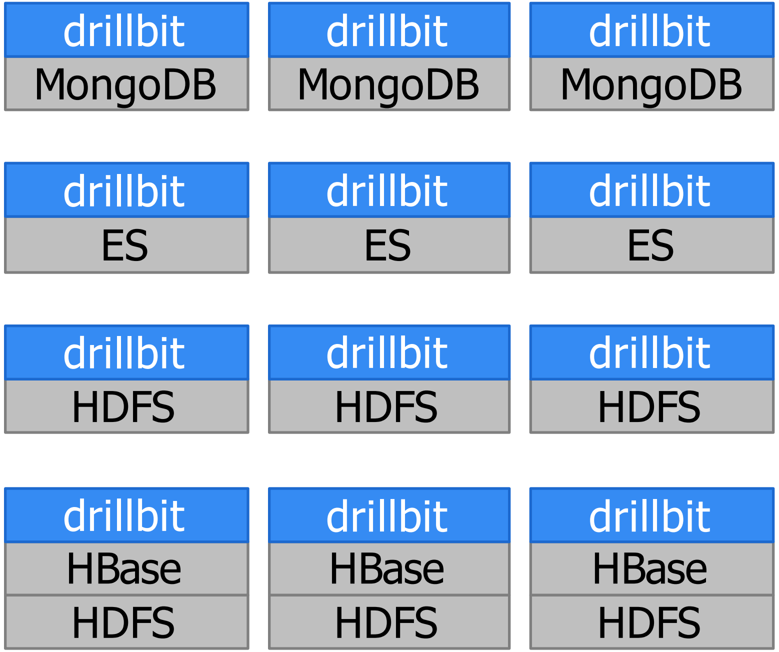

drillbit drillbit drillbit
MongoDB MongoDB MongoDB

drillbit drillbit drillbit
ES ES ES

drillbit drillbit drillbit
HDFS HDFS HDFS

drillbit drillbit drillbit
HBase HBase HBase
HDFS HDFS HDFS

Cloud Services

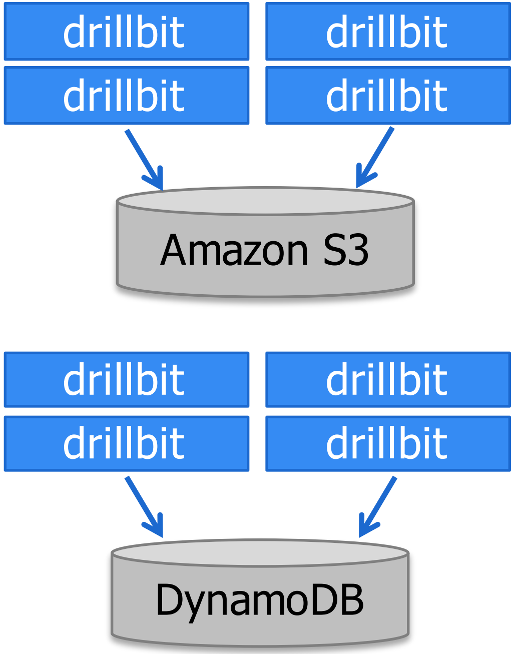

drillbit drillbit
drillbit drillbit

Amazon S3

drillbit drillbit
drillbit drillbit

DynamoDB

Desktops

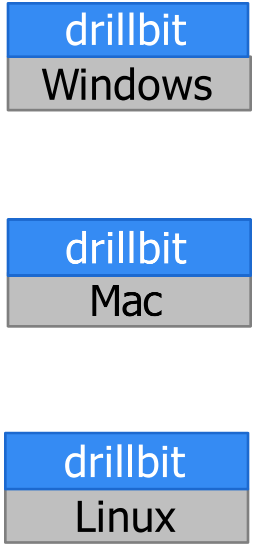

drillbit
Windows

drillbit
Mac

drillbit
Linux

© 2015. All Rights Reserved.

Deployment Modes

Embedded Mode

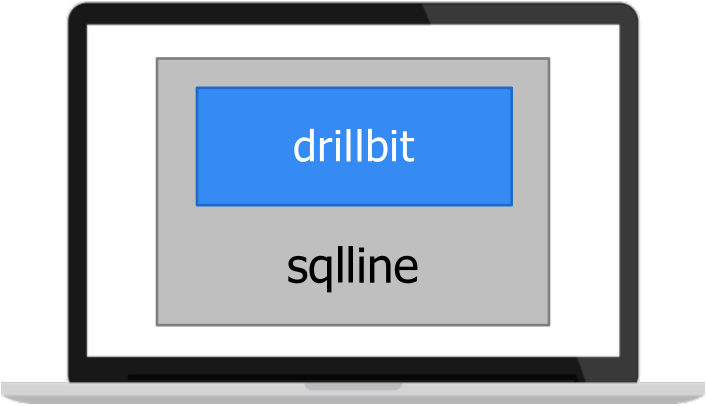

Distributed Mode (aka Drill Cluster)

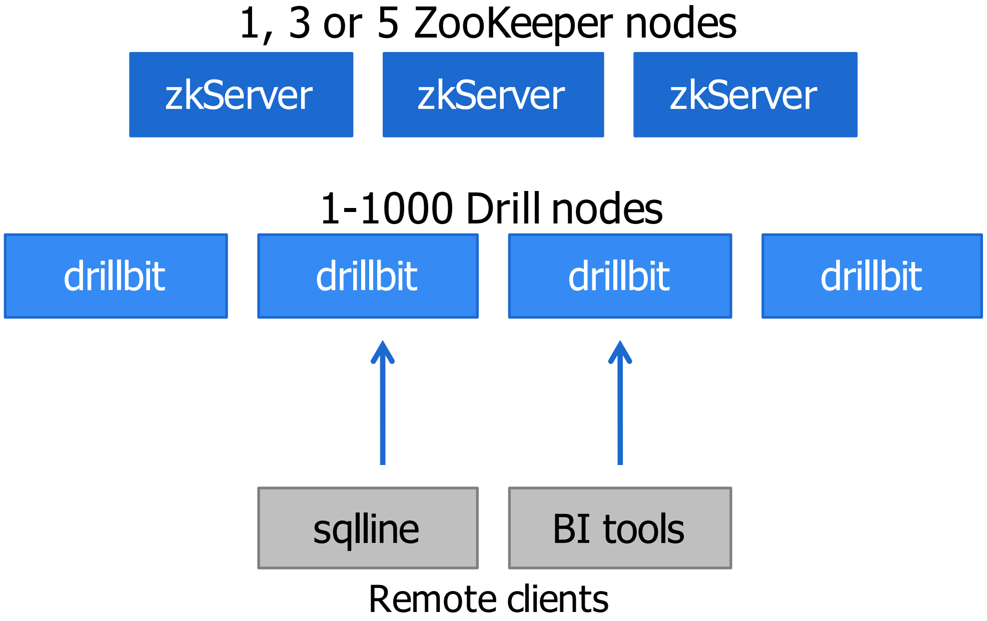

© 2015. All Rights Reserved.

APACHE DRILL

Extensible Datastore Architecture

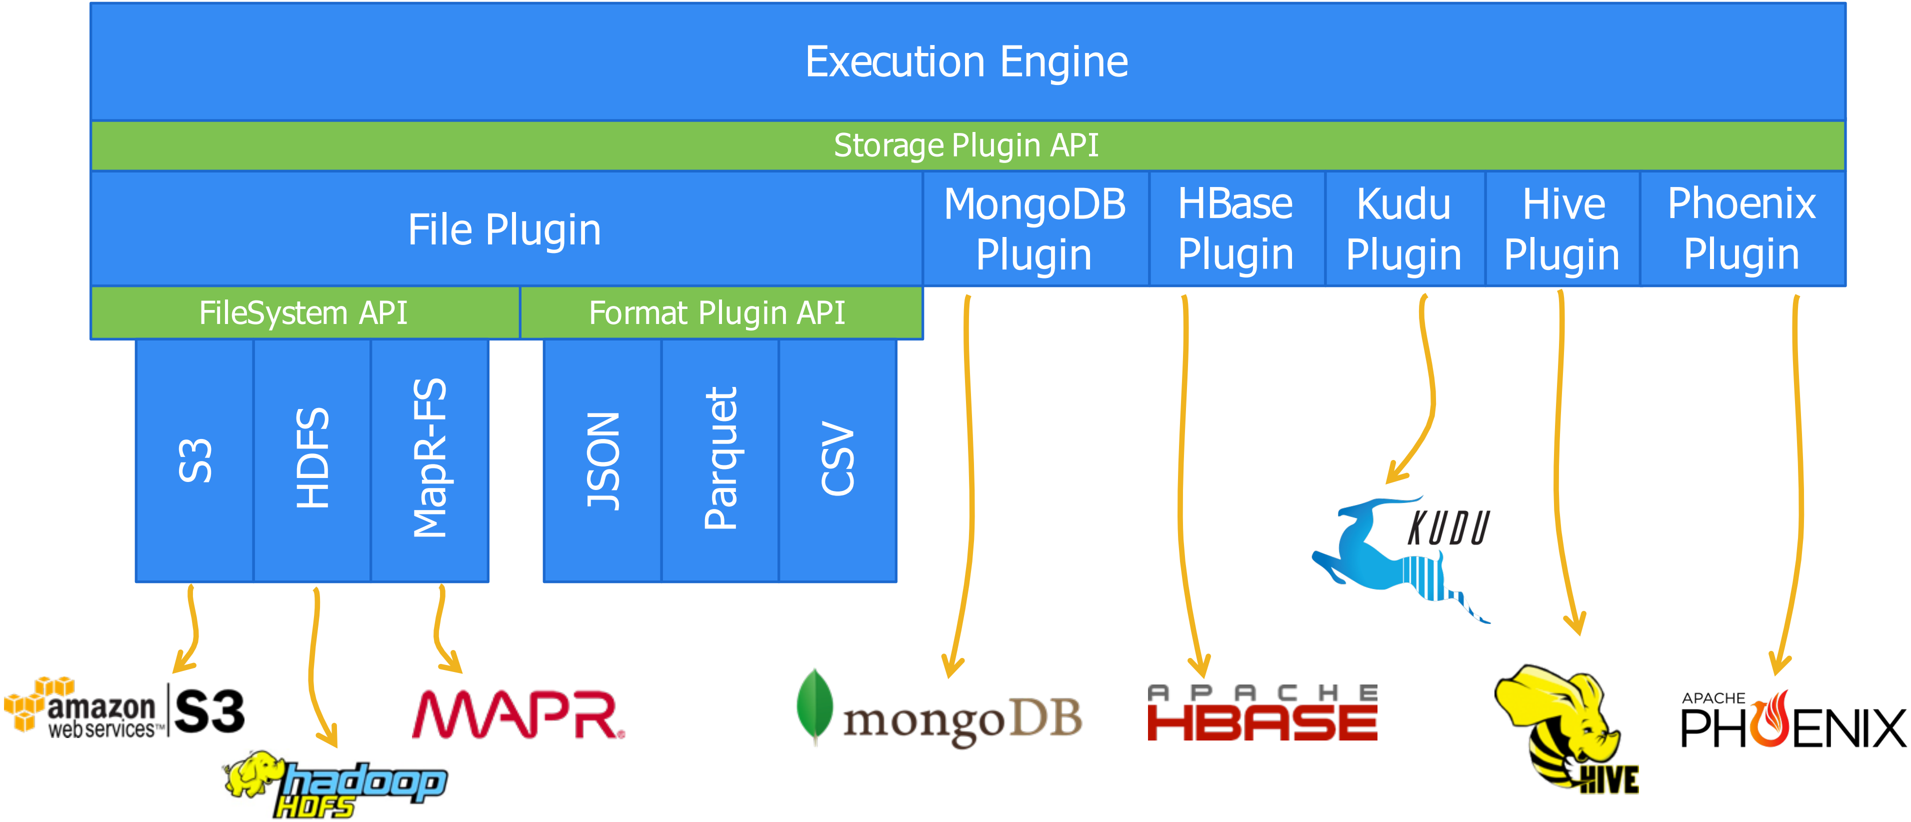

© 2015. All Rights Reserved.

GETTING STARTED

© 2015. All Rights Reserved.

Install Drill

From the browser: http://drill.apache.org/download

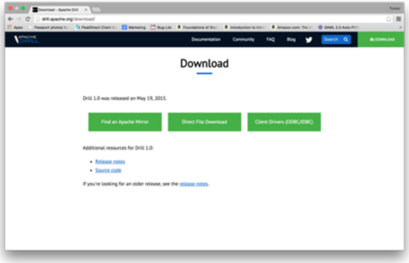

Download

Drill 1.0 was released on May 19, 2015.

Additional resources for Drill 1.0:

- Release notes
- Source code

If you're looking for an older release, see the release notes.

From the command line:

```
$ curl -L http://www.dremio.com/drill-latest.tgz | tar xz
```

(Also make sure you have Java 1.7+ installed…)

© 2015. All Rights Reserved.

Install a Local MongoDB Instance

Install and run MongoDB: http://docs.mongodb.org/manual/installation/

Import the dataset of Yelp businesses:

```
$ mongoimport --db yelp --collection business --drop –file
yelp/business.json
$ mongo
> use yelp;
> db.business.findOne().pretty();
{
    "_id" : ObjectId("55921ddfc6c0a4a2d8ef700c"),
    …
}
```

© 2015. All Rights Reserved.

Test Your Drill Setup

Start Drill shell with an embedded drillbit daemon:

```
$ apache-drill-x.y.z/bin/drill-embedded
```

(If you are using the tutorial EC2 instance, type `drill` instead of …`/drill-embedded`)

Access the Drill web UI at localhost:8047:

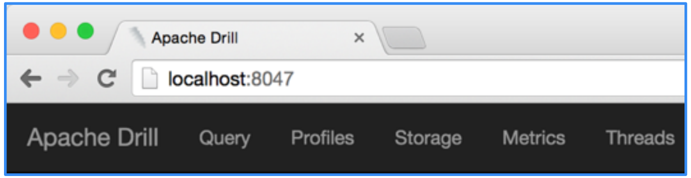

Enable the MongoDB plugin by clicking Enable in the Storage tab:

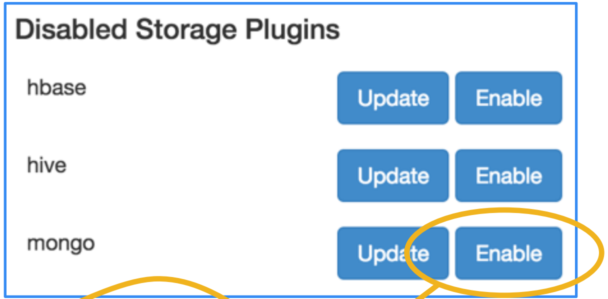
Disabled Storage Plugins

© 2015. All Rights Reserved.

Run Your First Query

```
> SELECT name FROM mongo.yelp.business LIMIT 1;
+--------------------+
|        name        |
+--------------------+
| Eric Goldberg, MD  |
+--------------------+

> SELECT name FROM dfs.root.`/opt/tutorial/yelp/business.json`
LIMIT 1;
+--------------------+
|        name        |
+--------------------+
| Eric Goldberg, MD  |
+--------------------+
```

© 2015. All Rights Reserved.

Referencing a Table

```
SELECT * FROM production.website.users;
```

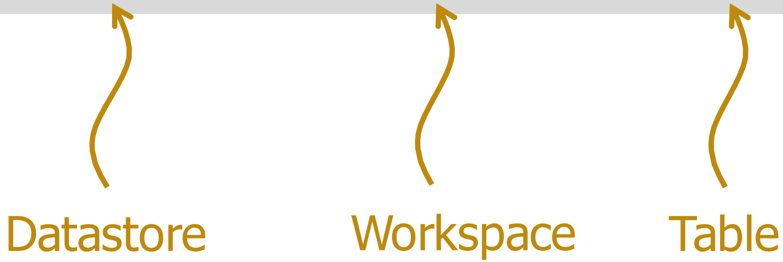

Datastore

Workspace

Table

© 2015. All Rights Reserved.

Namespaces & Tables

Storage Plugin Type	Workspace	Table
mongo	Database	Collection
hive	Database	Table
hbase	Namespace	Table
file (HDFS cluster, S3, …)	Directory	File or directory
…	…	…

User defines these in the datastore configuration

© 2015. All Rights Reserved.

Joining Across Datastores is Easy!

Alias to a specific file system (S3, HDFS, local, NAS)

```
> SELECT *
  FROM dfs.root.`/opt/tutorial/yelp/review.json` r,
       mongo.yelp.business b
  WHERE r.business_id = b.business_id;
```

Alias to a specific MongoDB cluster

© 2015. All Rights Reserved.

Joining Across Datastores

Data

- Local file: yelp/review.json (`dfs.yelp.`review.json``)
- MongoDB collection: yelp.business (`mongo.yelp.business`)

Question

- What’s the name of the business with the most reviews on Yelp?

© 2015. All Rights Reserved.

```
> SELECT b.name AS name, COUNT(*) AS reviews
  FROM dfs.tutorial.`yelp/review.json` r,
       mongo.yelp.business b
  WHERE r.business_id = b.business_id
  GROUP BY b.business_id, b.name
  ORDER BY reviews DESC
  LIMIT 3;
+--------------------+-----------+
|        name        |  reviews  |
+--------------------+-----------+
| Mon Ami Gabi       | 3695      |
| Earl of Sandwich   | 3263      |
| Wicked Spoon       | 3011      |
+--------------------+-----------+
```

© 2015. All Rights Reserved.

Accessing Array Elements

```
> SELECT categories FROM business LIMIT 2;
+-------------------------------------------+
|                categories                 |
+-------------------------------------------+
| ["American (Traditional)","Restaurants"]  |
| ["Chinese","Restaurants"]                 |
+-------------------------------------------+

> SELECT categories[0] FROM business LIMIT 2;
+-------------------------+
|         EXPR$0          |
+-------------------------+
| American (Traditional)  |
| Chinese                 |
+-------------------------+
```

© 2015. All Rights Reserved.

Accessing Data in Maps

- Use the dot (.) notation to access nested fields
- Must specify the table name/alias when accessing nested fields
  - `<table>.<field>.<nested field 1>.<nested field 2>`
  - When Drill sees one or more dots, it assumes the first is a table name/alias

```
> SELECT attributes FROM business;
> SELECT business.attributes FROM business;
> SELECT b.attributes FROM business b;

> SELECT attributes.Parking FROM business;
Error: PARSE ERROR: Table 'attributes' not found
> SELECT business.attributes.Parking FROM business;
> SELECT b.attributes.Parking FROM business b;
```

© 2015. All Rights Reserved.

FLATTEN

- FLATTEN converts single record with array field into multiple records
  - One output record for each array element
- Non FLATTENed fields are repeated in each of the output records

```
> SELECT categories
  FROM business LIMIT 2;
+-----------------------------------------+
|                categories                |
+-----------------------------------------+
| ["American (Traditional)","Restaurants"] |
| ["Chinese","Restaurants"]                |
+-----------------------------------------+
```

→

```
> SELECT FLATTEN(categories)
  FROM business LIMIT 4;
+-------------------------+
|         EXPR$0          |
+-------------------------+
| American (Traditional)  |
| Restaurants             |
| Chinese                 |
| Restaurants             |
+-------------------------+
```

© 2015. All Rights Reserved.

Non-FLATTENed Fields are Repeated

```
> SELECT name, categories FROM business LIMIT 2;
+-------------------------------+--------------------------------------------+
|             name              |                 categories                 |
+-------------------------------+--------------------------------------------+
| Deforest Family Restaurant    | ["American (Traditional)","Restaurants"]   |
| Chang Jiang Chinese Kitchen   | ["Chinese","Restaurants"]                  |
+-------------------------------+--------------------------------------------+

> SELECT name, FLATTEN(categories) FROM business LIMIT 4;
+-------------------------------+--------------------------+
|             name              |          EXPR$1          |
+-------------------------------+--------------------------+
| Deforest Family Restaurant    | American (Traditional)   |
| Deforest Family Restaurant    | Restaurants              |
| Chang Jiang Chinese Kitchen   | Chinese                  |
| Chang Jiang Chinese Kitchen   | Restaurants              |
+-------------------------------+--------------------------+
```

© 2015. All Rights Reserved.

BI

© 2015. All Rights Reserved.

ODBC and JDBC

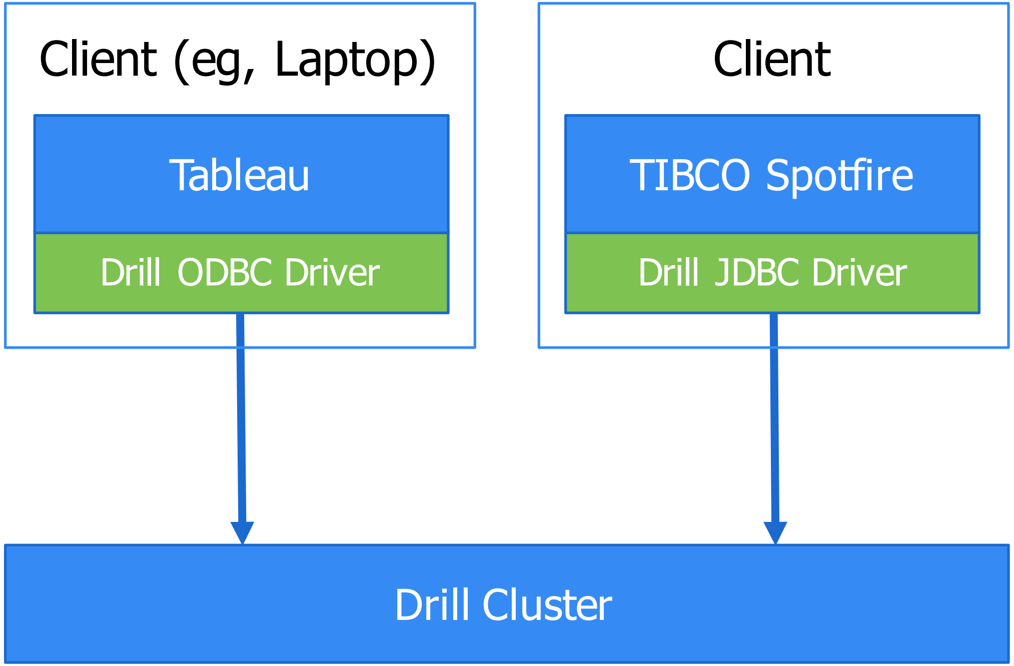

- Drill includes standard ODBC/JDBC drivers
  - ODBC for native apps
  - JDBC for Java apps
- User installs the driver on the client
  - The same machine as the BI tool
- Driver communicates with Drill cluster(s)
- Make sure driver and cluster are compatible versions

© 2015. All Rights Reserved.

Install the ODBC Driver

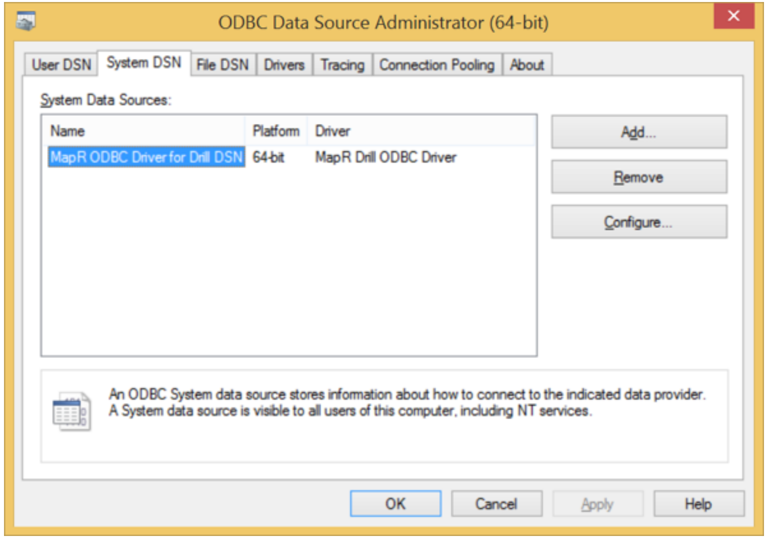

Name	Platform	Driver
MapR ODBC Driver for Drill DSN	64-bit	MapR Drill ODBC Driver

An ODBC System data source stores information about how to connect to the indicated data provider. A System data source is visible to all users of this computer, including NT services.

- Download & install the Drill ODBC driver from the Drill website (drill.apache.org)
- If you will use Tableau, install Tableau before the ODBC driver or get the TDC file separately
- Open the system's ODBC Data Source Administrator

© 2015. All Rights Reserved.

Configure the ODBC Driver

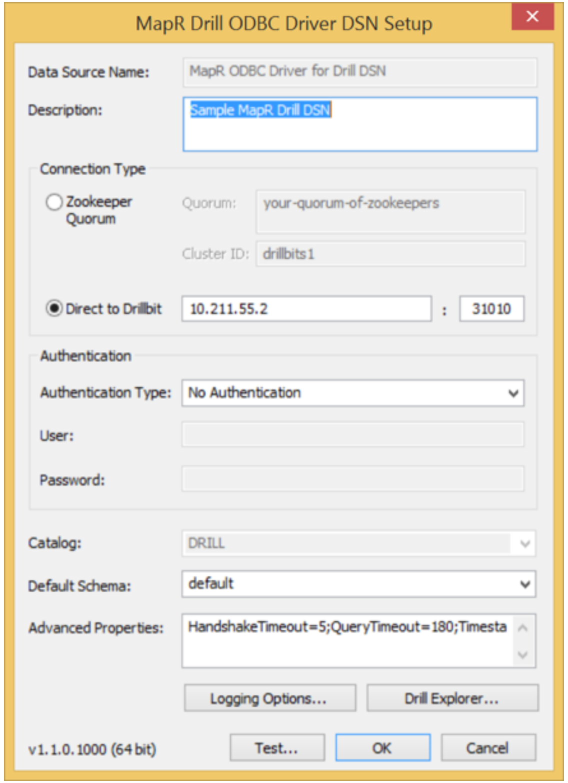

MapR Drill ODBC Driver DSN Setup

Data Source Name: MapR ODBC Driver for Drill DSN

Description: Sample MapR Drill DSN

Connection Type

○ Zookeeper Quorum — Quorum: your-quorum-of-zookeepers — Cluster ID: drillbits1

● Direct to Drillbit — 10.211.55.2 : 31010

Authentication

Authentication Type: No Authentication

User:

Password:

Catalog: DRILL

Default Schema: default

Advanced Properties: HandshakeTimeout=5;QueryTimeout=180;Timesta

v1.1.0.1000 (64 bit)

- Two connection options:
  - Connect to a specific node (drillbit) in the Drill cluster
  - Connect to any node (drillbit) in the Drill cluster
- Click “Test...” to make sure the ODBC driver can connect to Drill

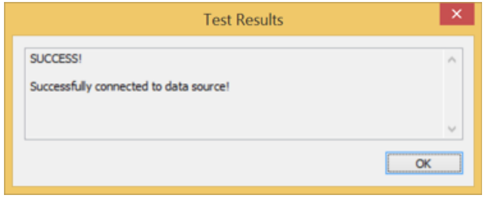

Test Results

SUCCESS!

Successfully connected to data source!

© 2015. All Rights Reserved.

Connect with Tableau

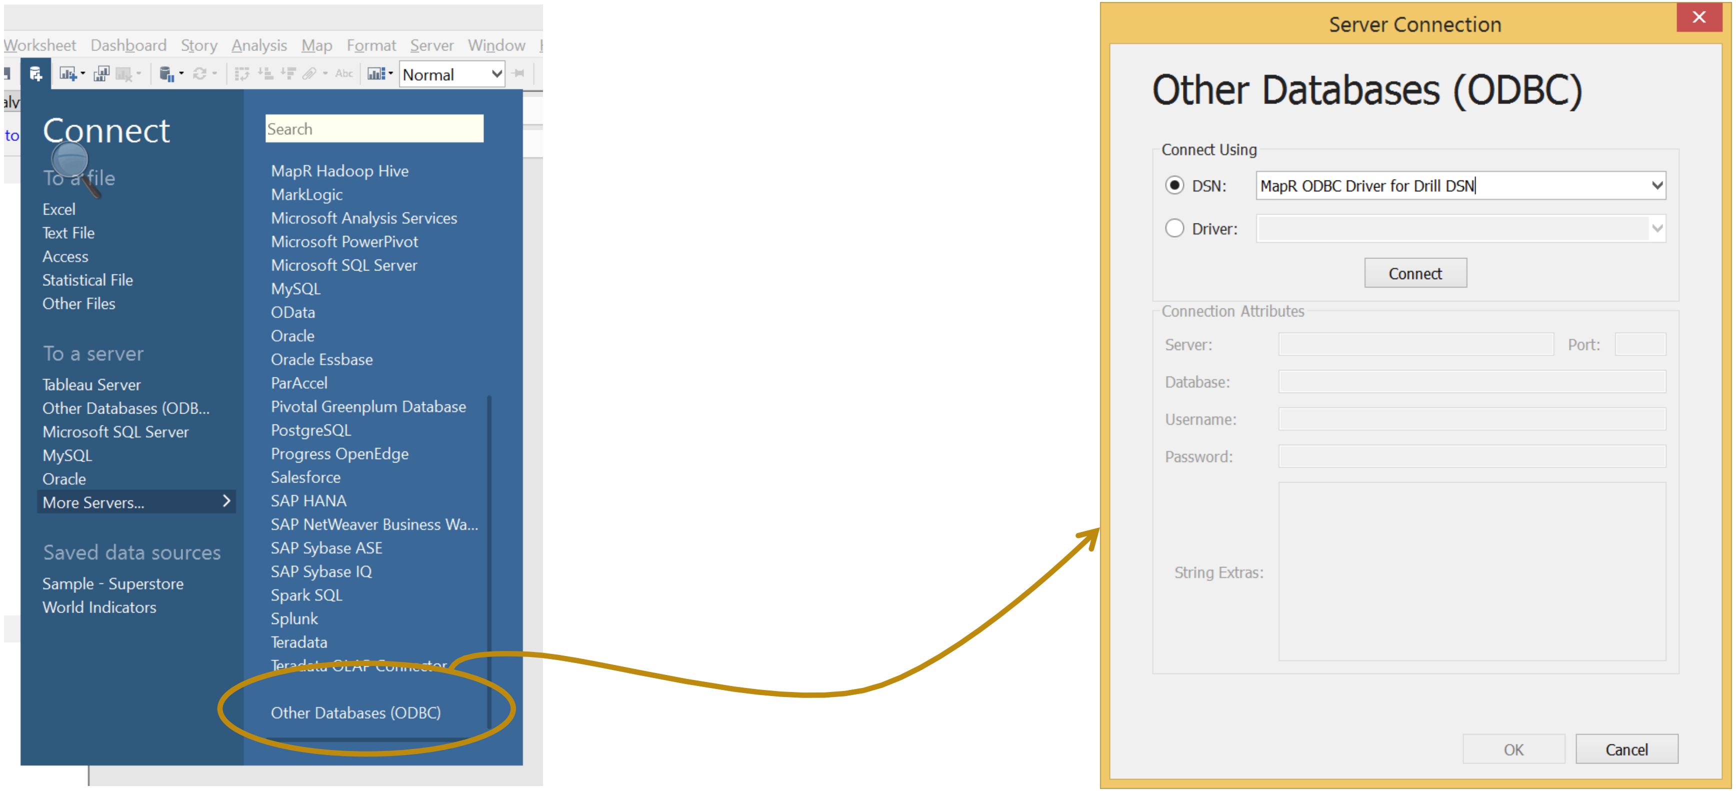

● DSN: MapR ODBC Driver for Drill DSN

○ Driver:

© 2015. All Rights Reserved.

Choose ODBC DSN

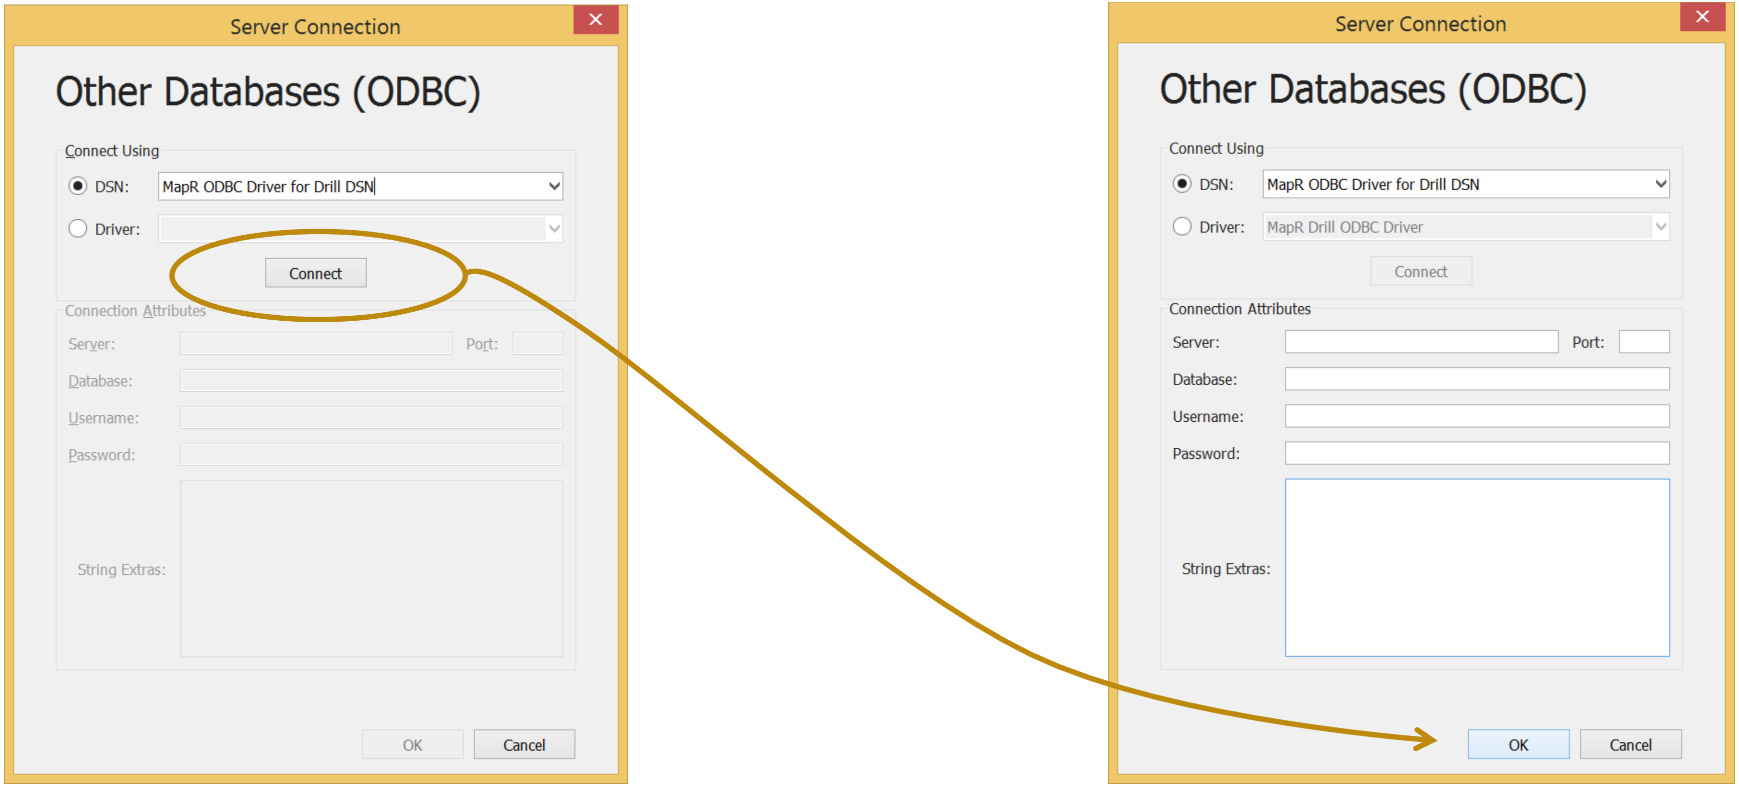

Server Connection

Other Databases (ODBC)

Connect Using

● DSN: MapR ODBC Driver for Drill DSN

○ Driver:

Server Connection

Other Databases (ODBC)

Connect Using

● DSN: MapR ODBC Driver for Drill DSN

○ Driver: MapR Drill ODBC Driver

© 2015. All Rights Reserved.

Choose Drill Workspace & Table

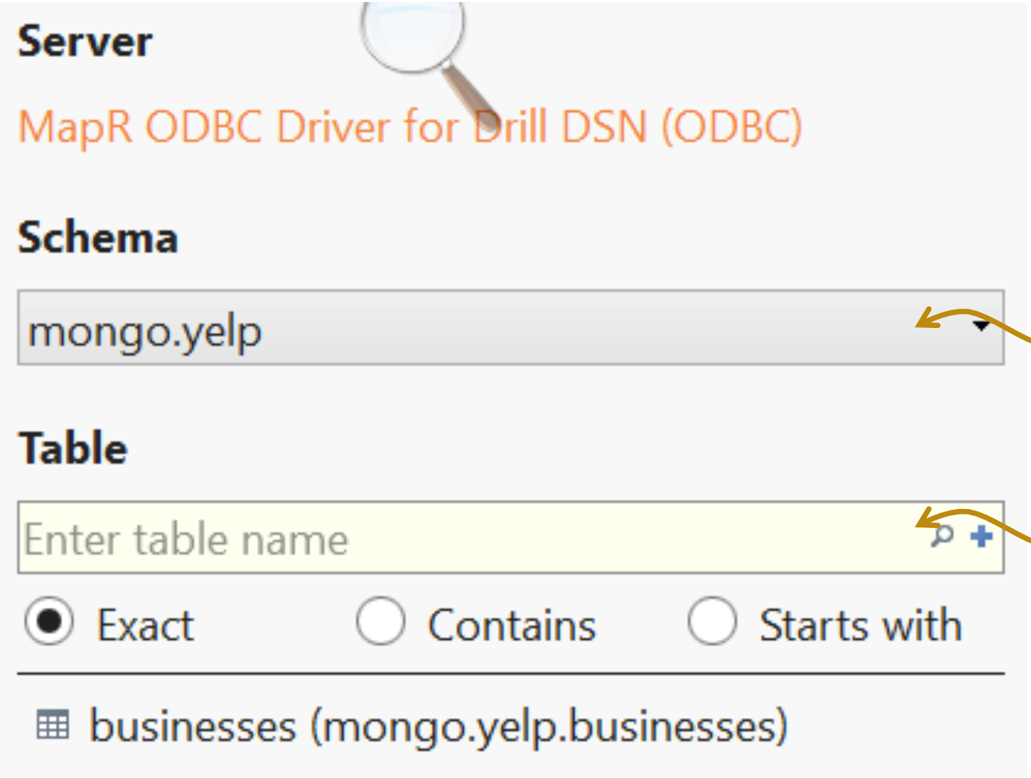

Server

MapR ODBC Driver for Drill DSN (ODBC)

Schema

mongo.yelp

Table

Enter table name

● Exact ○ Contains ○ Starts with

businesses (mongo.yelp.businesses)

All <datastore>.<workspace> combinations will be shown here

In a file system datastore, currently only views will be shown here (an upcoming feature will enable selection of any file/directory representing a dataset)

© 2015. All Rights Reserved.

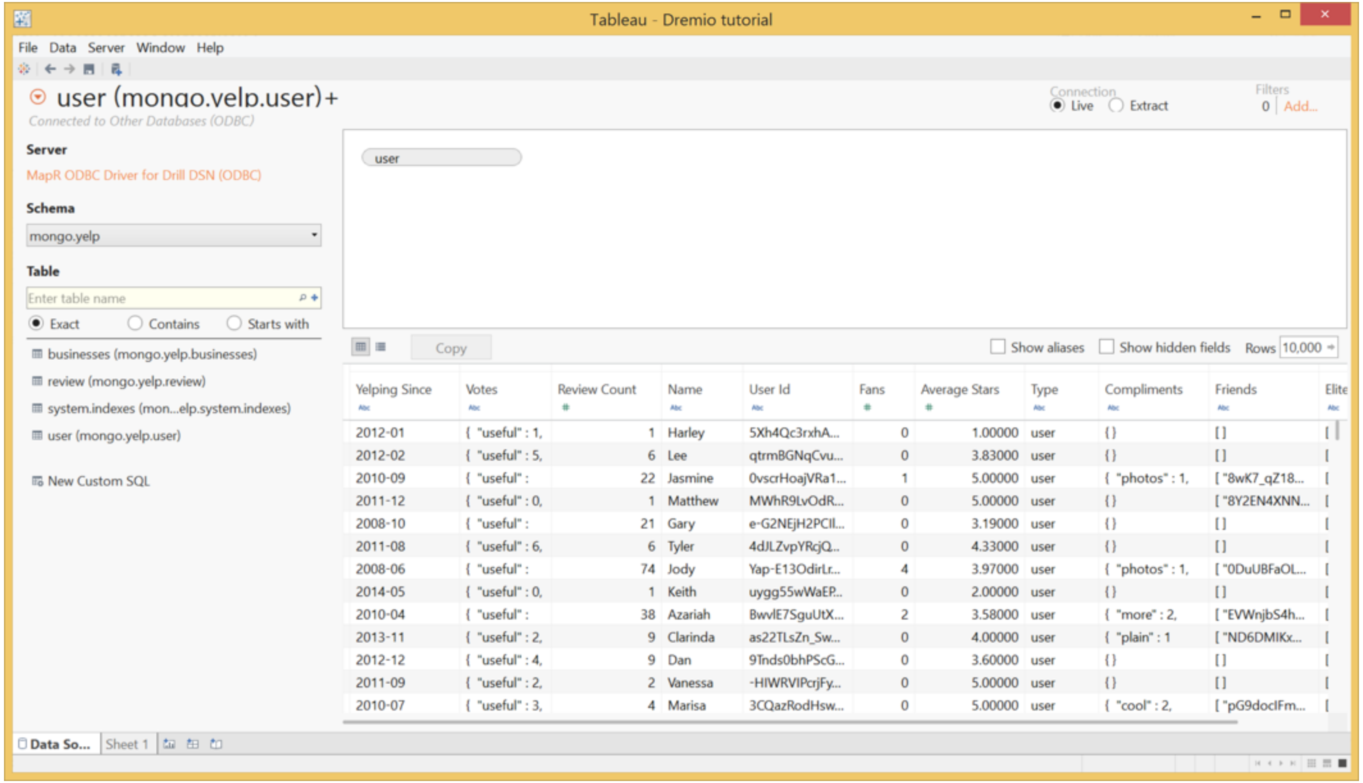

user (mongo.yelp.user)+

Connected to Other Databases (ODBC)

Connection ● Live ○ Extract

Filters 0 Add...

Server

MapR ODBC Driver for Drill DSN (ODBC)

Schema

mongo.yelp

Table

Enter table name

● Exact ○ Contains ○ Starts with

- businesses (mongo.yelp.businesses)
- review (mongo.yelp.review)
- system.indexes (mon...elp.system.indexes)
- user (mongo.yelp.user)
- New Custom SQL

user

□ Show aliases □ Show hidden fields Rows 10,000

Yelping Since	Votes	Review Count	Name	User Id	Fans	Average Stars	Type	Compliments	Friends	Elite
2012-01	{ "useful" : 1,	1	Harley	5Xh4Qc3rxhA...	0	1.00000	user	{ }	[ ]	[
2012-02	{ "useful" : 5,	6	Lee	qtrmBGNqCvu...	0	3.83000	user	{ }	[ ]	[
2010-09	{ "useful" :	22	Jasmine	0vscrHoajVRa1...	1	5.00000	user	{ "photos" : 1,	["8wK7_qZ18...	[
2011-12	{ "useful" : 0,	1	Matthew	MWhR9LvOdR...	0	5.00000	user	{ }	["8Y2EN4XNN...	[
2008-10	{ "useful" :	21	Gary	e-G2NEjH2PCIl...	0	3.19000	user	{ }	[ ]	[
2011-08	{ "useful" : 6,	6	Tyler	4dJLZvpYRcjQ...	0	4.33000	user	{ }	[ ]	[
2008-06	{ "useful" :	74	Jody	Yap-E13OdirLr...	4	3.97000	user	{ "photos" : 1,	["0DuUBFaOL...	[
2014-05	{ "useful" : 0,	1	Keith	uygg55wWaEP...	0	2.00000	user	{ }	[ ]	[
2010-04	{ "useful" :	38	Azariah	BwvlE7SguUtX...	2	3.58000	user	{ "more" : 2,	["EVWnjbS4h...	[
2013-11	{ "useful" : 2,	9	Clarinda	as22TLsZn_Sw...	0	4.00000	user	{ "plain" : 1	["ND6DMIKx...	[
2012-12	{ "useful" : 4,	9	Dan	9Tnds0bhPScG...	0	3.60000	user	{ }	[ ]	[
2011-09	{ "useful" : 2,	2	Vanessa	-HIWRVIPcrjFy...	0	5.00000	user	{ }	[ ]	[
2010-07	{ "useful" : 3,	4	Marisa	3CQazRodHsw...	0	5.00000	user	{ "cool" : 2,	["pG9docIFm...	[

© 2015. All Rights Reserved.

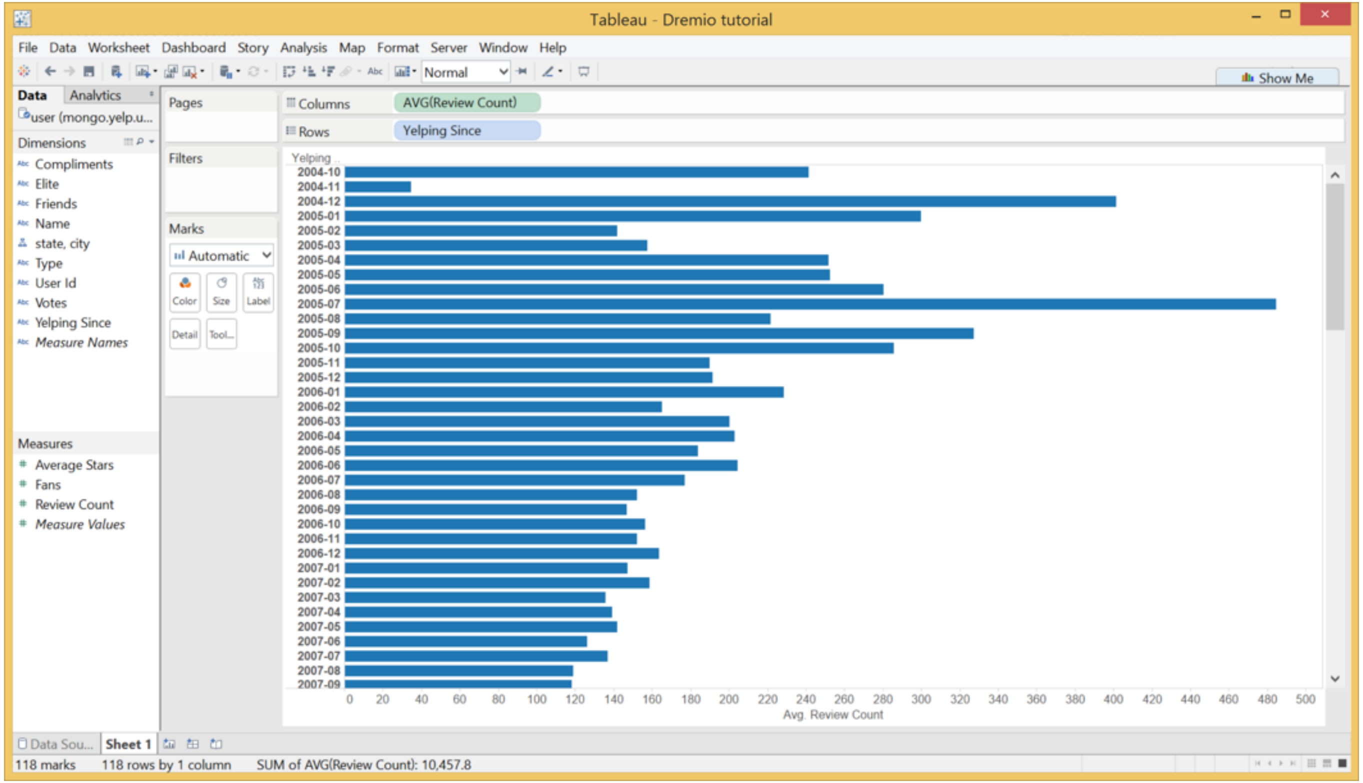

Tableau - Dremio tutorial

Columns: AVG(Review Count)

Rows: Yelping Since

Marks: Automatic

118 marks 118 rows by 1 column SUM of AVG(Review Count): 10,457.8

© 2015. All Rights Reserved.

REST

© 2015. All Rights Reserved.

Using the REST API

Client application sends a properly structured HTTP request to any drillbit in the cluster

Part	Required Content
Method	POST
Headers	`Content-Type: application/json`
Payload	JSON

```
curl \
  --header "Content-Type: application/json" \
  --request POST \
  --data '{"queryType": "SQL", "query": "SELECT * FROM …"}' \
  http://localhost:8047/query.json
```

© 2015. All Rights Reserved.

Run a Query via REST

```
$ curl --header "Content-Type: application/json" \
       --request POST \
       --data '{"queryType": "SQL", "query": "SELECT name FROM
dfs.tutorial.`yelp/business.json` LIMIT 3"}' \
       http://localhost:8047/query.json | python -m json.tool

{
  "columns": ["name"],
  "rows": [
    {"name": "Eric Goldberg, MD" },
    {"name": "Pine Cone Restaurant"},
    {"name": "Deforest Family Restaurant"}
  ]
}
```

© 2015. All Rights Reserved.

Thank You!

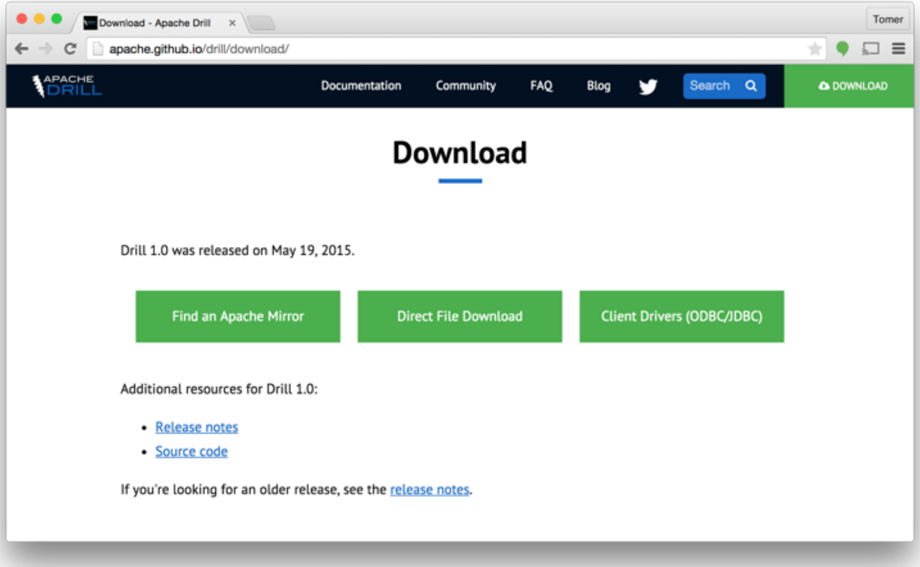

Download

Drill 1.0 was released on May 19, 2015.

Additional resources for Drill 1.0:

- Release notes
- Source code

If you're looking for an older release, see the release notes.

- Download at drill.apache.org
- Ask questions:
  - user@drill.apache.org
- Tweet
  - @ApacheDrill

© 2015. All Rights Reserved.