Strata+Hadoop WORLD

PRESENTED BY

O'REILLY cloudera

strataconf.com
#StrataHadoop

SQL in Hadoop: To Boldly Go Where No Data Warehouse has Gone Before

Emma McGrattan

SVP Engineering, Actian Corp

Who is Actian?

$140M Revenues + Profitable

10,000+ Customers

Global Presence: 8 world-wide offices, 7x 24 multinational support model

"Actian is now very powerfully positioned in the big data and analytics markets." Bloor

"Fast becoming a big data powerhouse to challenge the market." Forrester

☑

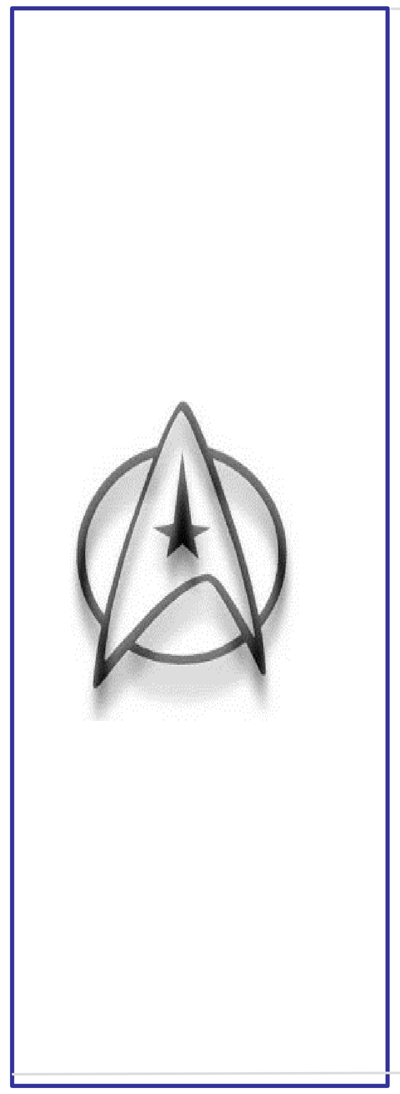

amazon webservices™

Goldman Sachs

Bank of America

Department of Defense United States of America

future group Udna hai. Aaj. Abhi.

YAHOO!

OfficeMax

salesforce.com

inferenda access today's most powerful data

The Home Depot

Sheetz Fresh Food Made to Order

is cool ENTERTAINMENT

autometrics

Groupe adeo

OPERA SOLUTIONS

EVERNOTE

ebay enterprise

GE

intuit

badoo

Actian Vortex

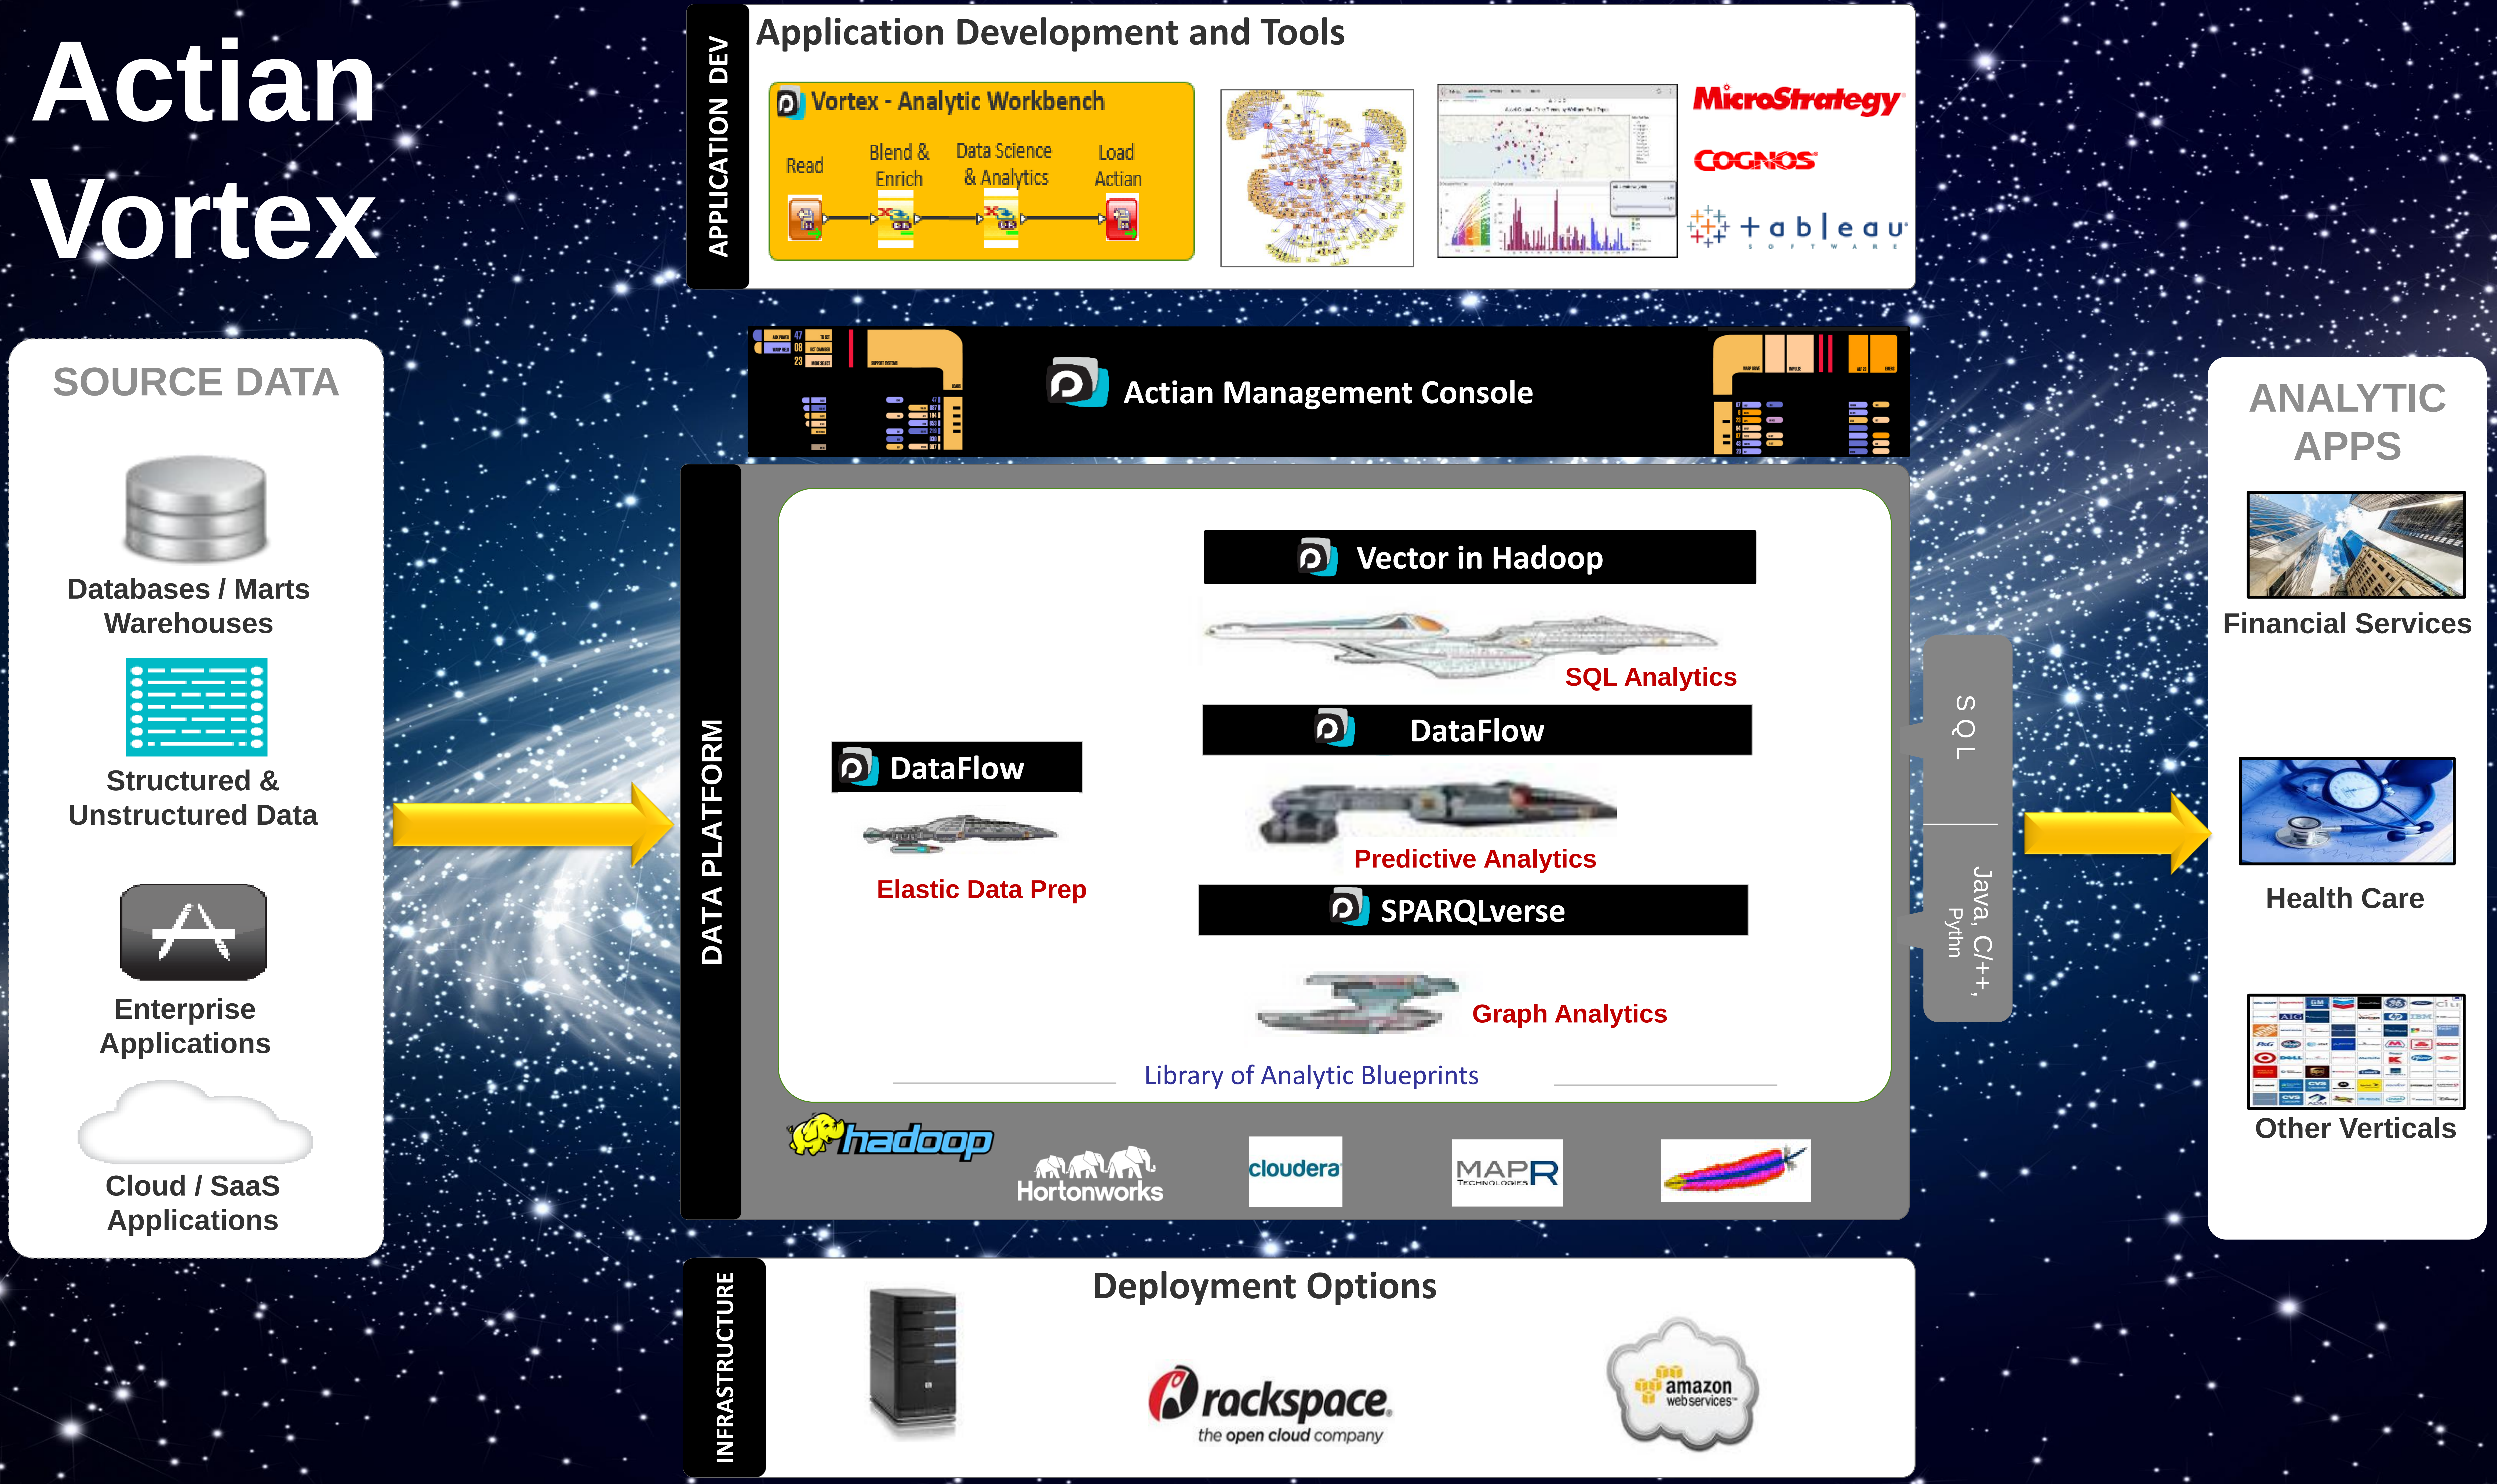

SOURCE DATA

Databases / Marts Warehouses

Structured & Unstructured Data

Enterprise Applications

Cloud / SaaS Applications

APPLICATION DEV

Application Development and Tools

Vortex - Analytic Workbench

Read → Blend & Enrich → Data Science & Analytics → Load Actian

Actian Management Console

DATA PLATFORM

Vector in Hadoop

SQL Analytics

DataFlow

Predictive Analytics

SPARQLverse

Graph Analytics

DataFlow

Elastic Data Prep

Library of Analytic Blueprints

SQL

Java, C/++, Pythn

ANALYTIC APPS

Financial Services

Health Care

Other Verticals

INFRASTRUCTURE

Deployment Options

Actian Vector in Hadoop Architecture

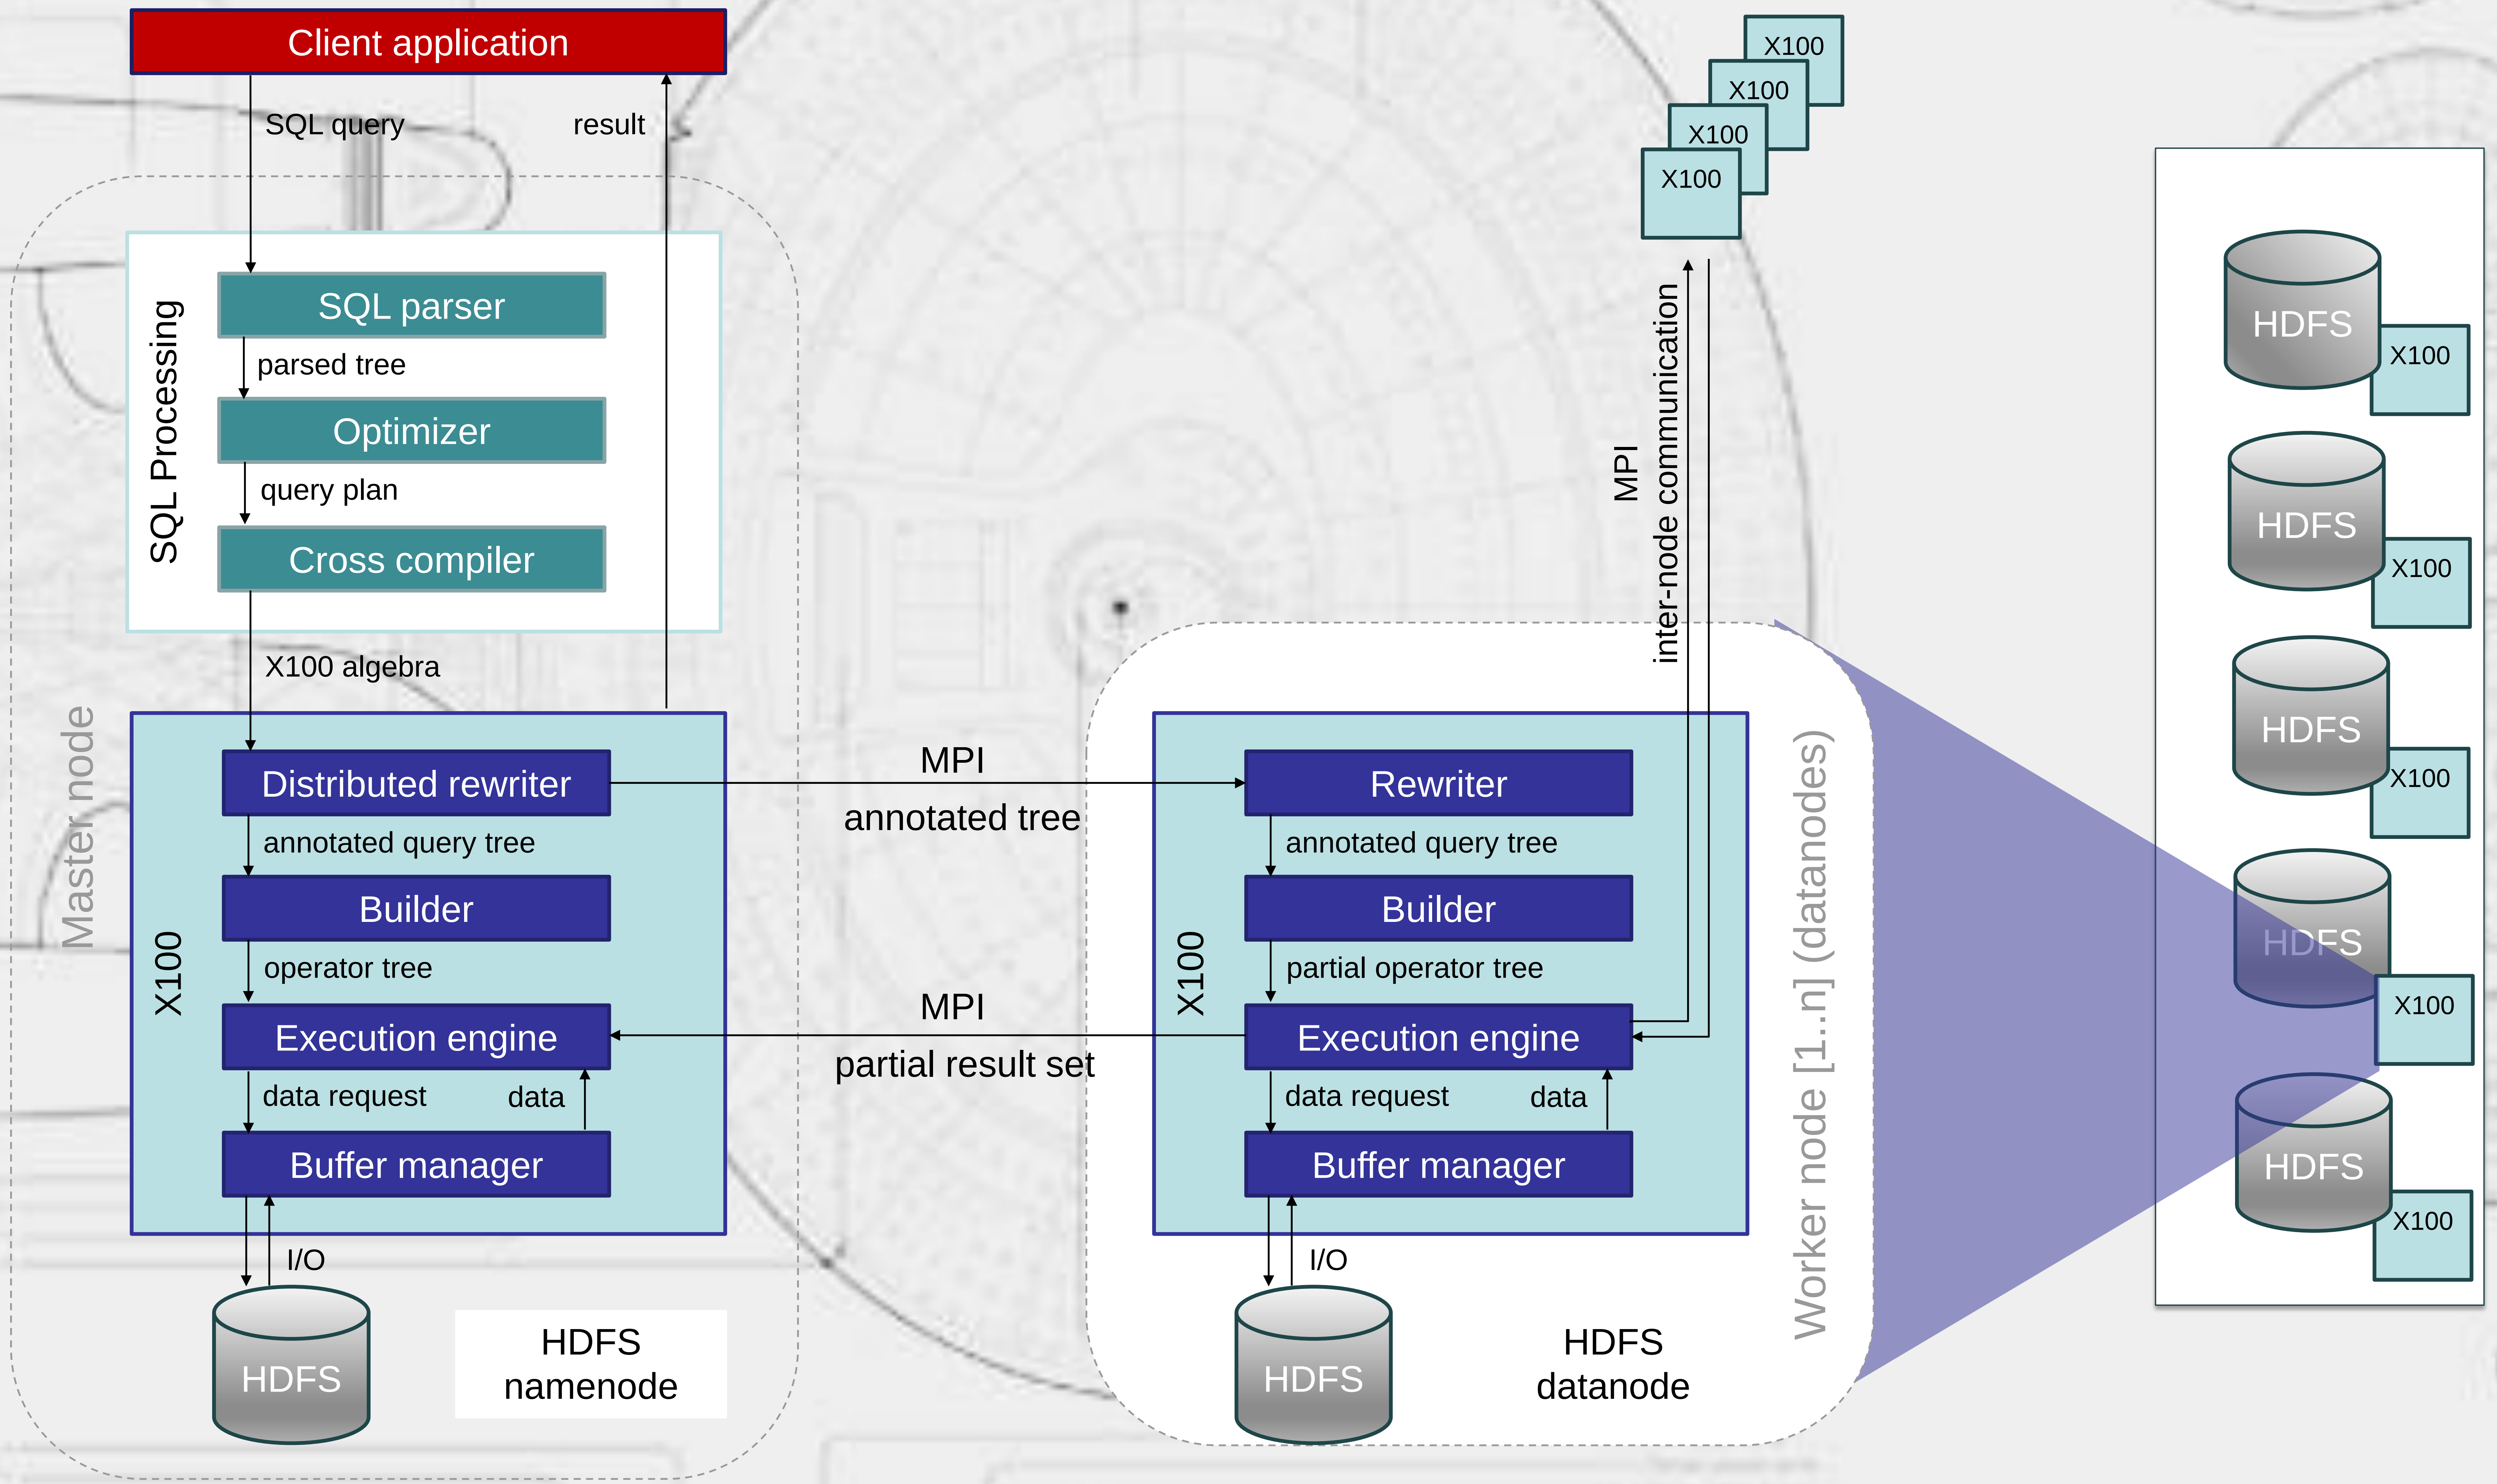

Vector: Built for Warp Speed

1 Vectorized End-to-End

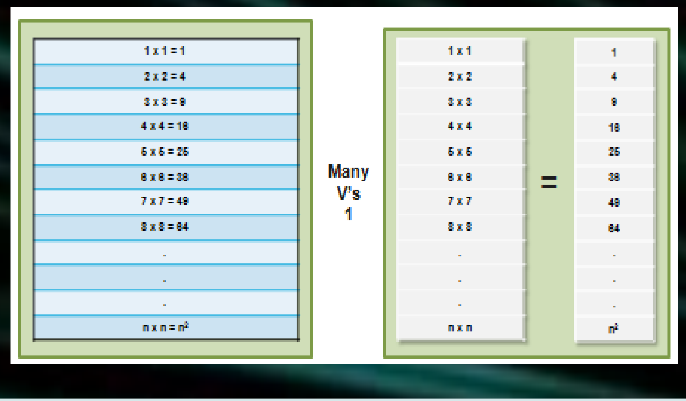

Single
Instruction
Multiple
Data

2 Exploiting Chip Cache

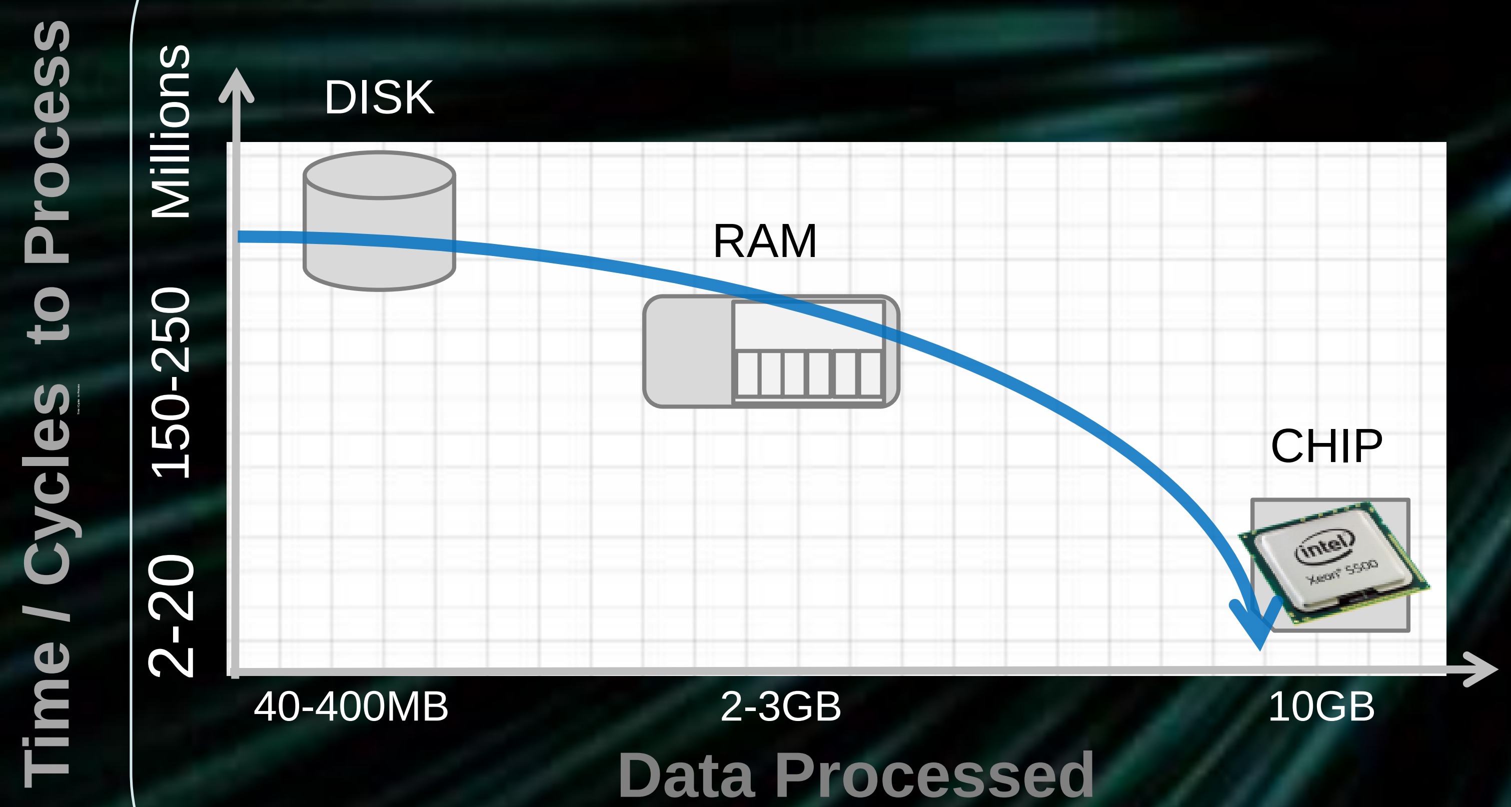

Process data on chip – not in RAM

3 Update Capability

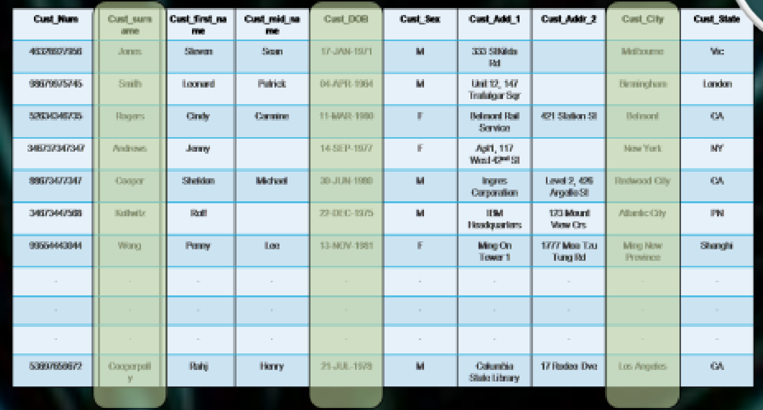

Cust_Num	Cust_surname	Cust_first_name	Cust_mid_name	Cust_DOB	Cust_Sex	Cust_Add_1	Cust_Addr_2	Cust_City	Cust_State
45378557958	Jones	Steven	Sean	17-JAN-1971	M	333 StKilda Rd		Melbourne	Vic
98879975746	Smith	Leonard	Patrick	04-APR-1984	M	Unit 12, 147 Trafalgar Sqr		Birmingham	London
52634048735	Rogers	Cindy	Carmine	11-MAR-1980	F	Belmont Rail Service	421 Station St	Belmont	CA
348737347347	Andrews	Jenny		14-SEP-1977	F	Apt1, 117 West 42nd St		New York	NY
89673477347	Cooper	Sheldon	Michael	30-JUN-1980	M	Ingres Corporation	Level 2, 456 Argyle St	Redwood City	CA
34873447588	Katheiz	Ralf		22-DEC-1975	M	IBM Headquarters	123 Mount View Crs	Atlantic City	PN
90654443044	Wong	Penny	Lee	13-NOV-1981	F	Ming On Tower 1	1777 Mao Tzu Tung Rd	Ming New Province	Shanghi
.	.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.	.
53897858672	Cooperpall y	Rahj	Henry	21-JUL-1978	M	Columbia State Library	17 Rodeo Dve	Los Angeles	CA

Limit I/O
Efficient real time updates
Update & Delete individual records

4 Smart Compression

Maximize throughput
Vectorized decompression

5 Storage Indexes

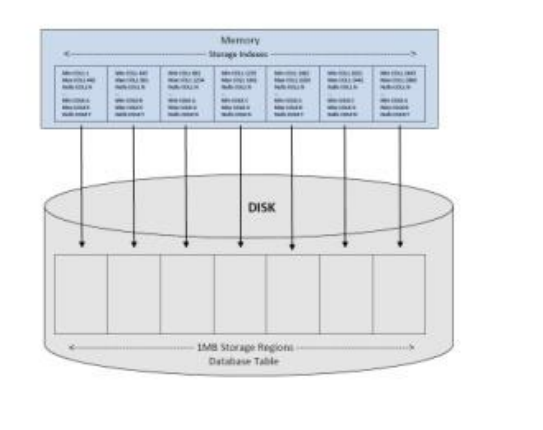

Quickly identify candidate data blocks
Minimize I/O

6 YARN Integration

Intelligent Block Placement
Dynamic Resource Management

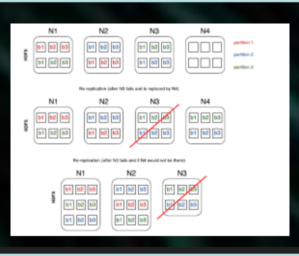

Vector: To Boldly Go…

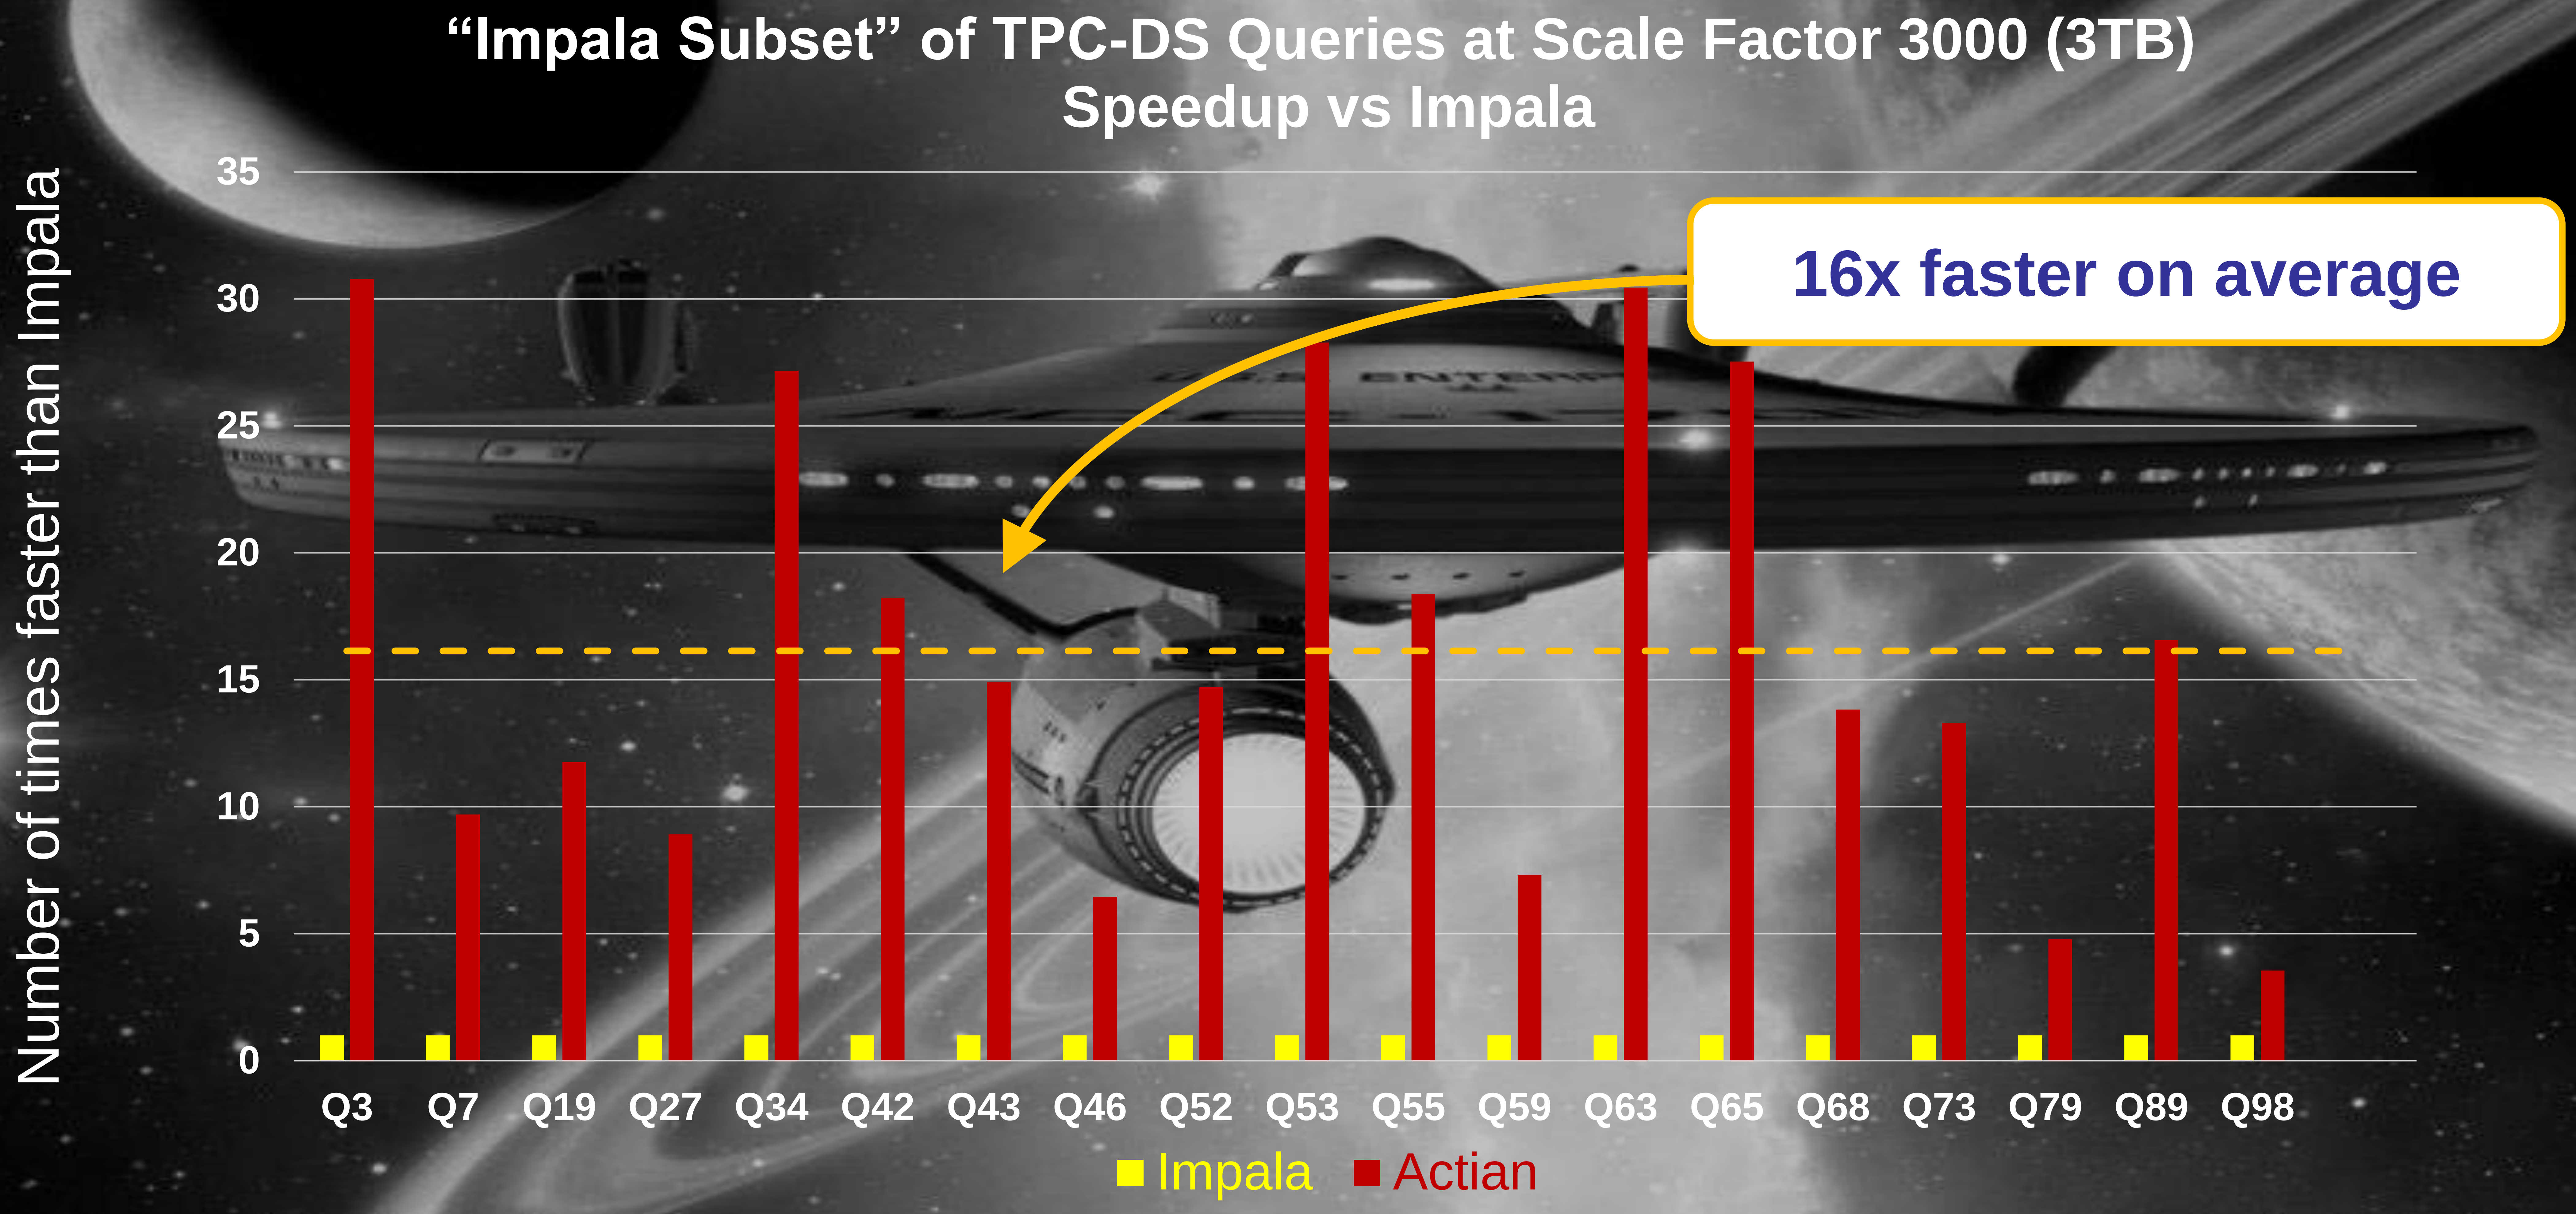

16x faster on average

Both Executed on the same hardware and software environment:
5 Node Cluster with 64GB of RAM per node and 24 x 1TB Hard Disks.

The SQL Behind the Actian Numbers

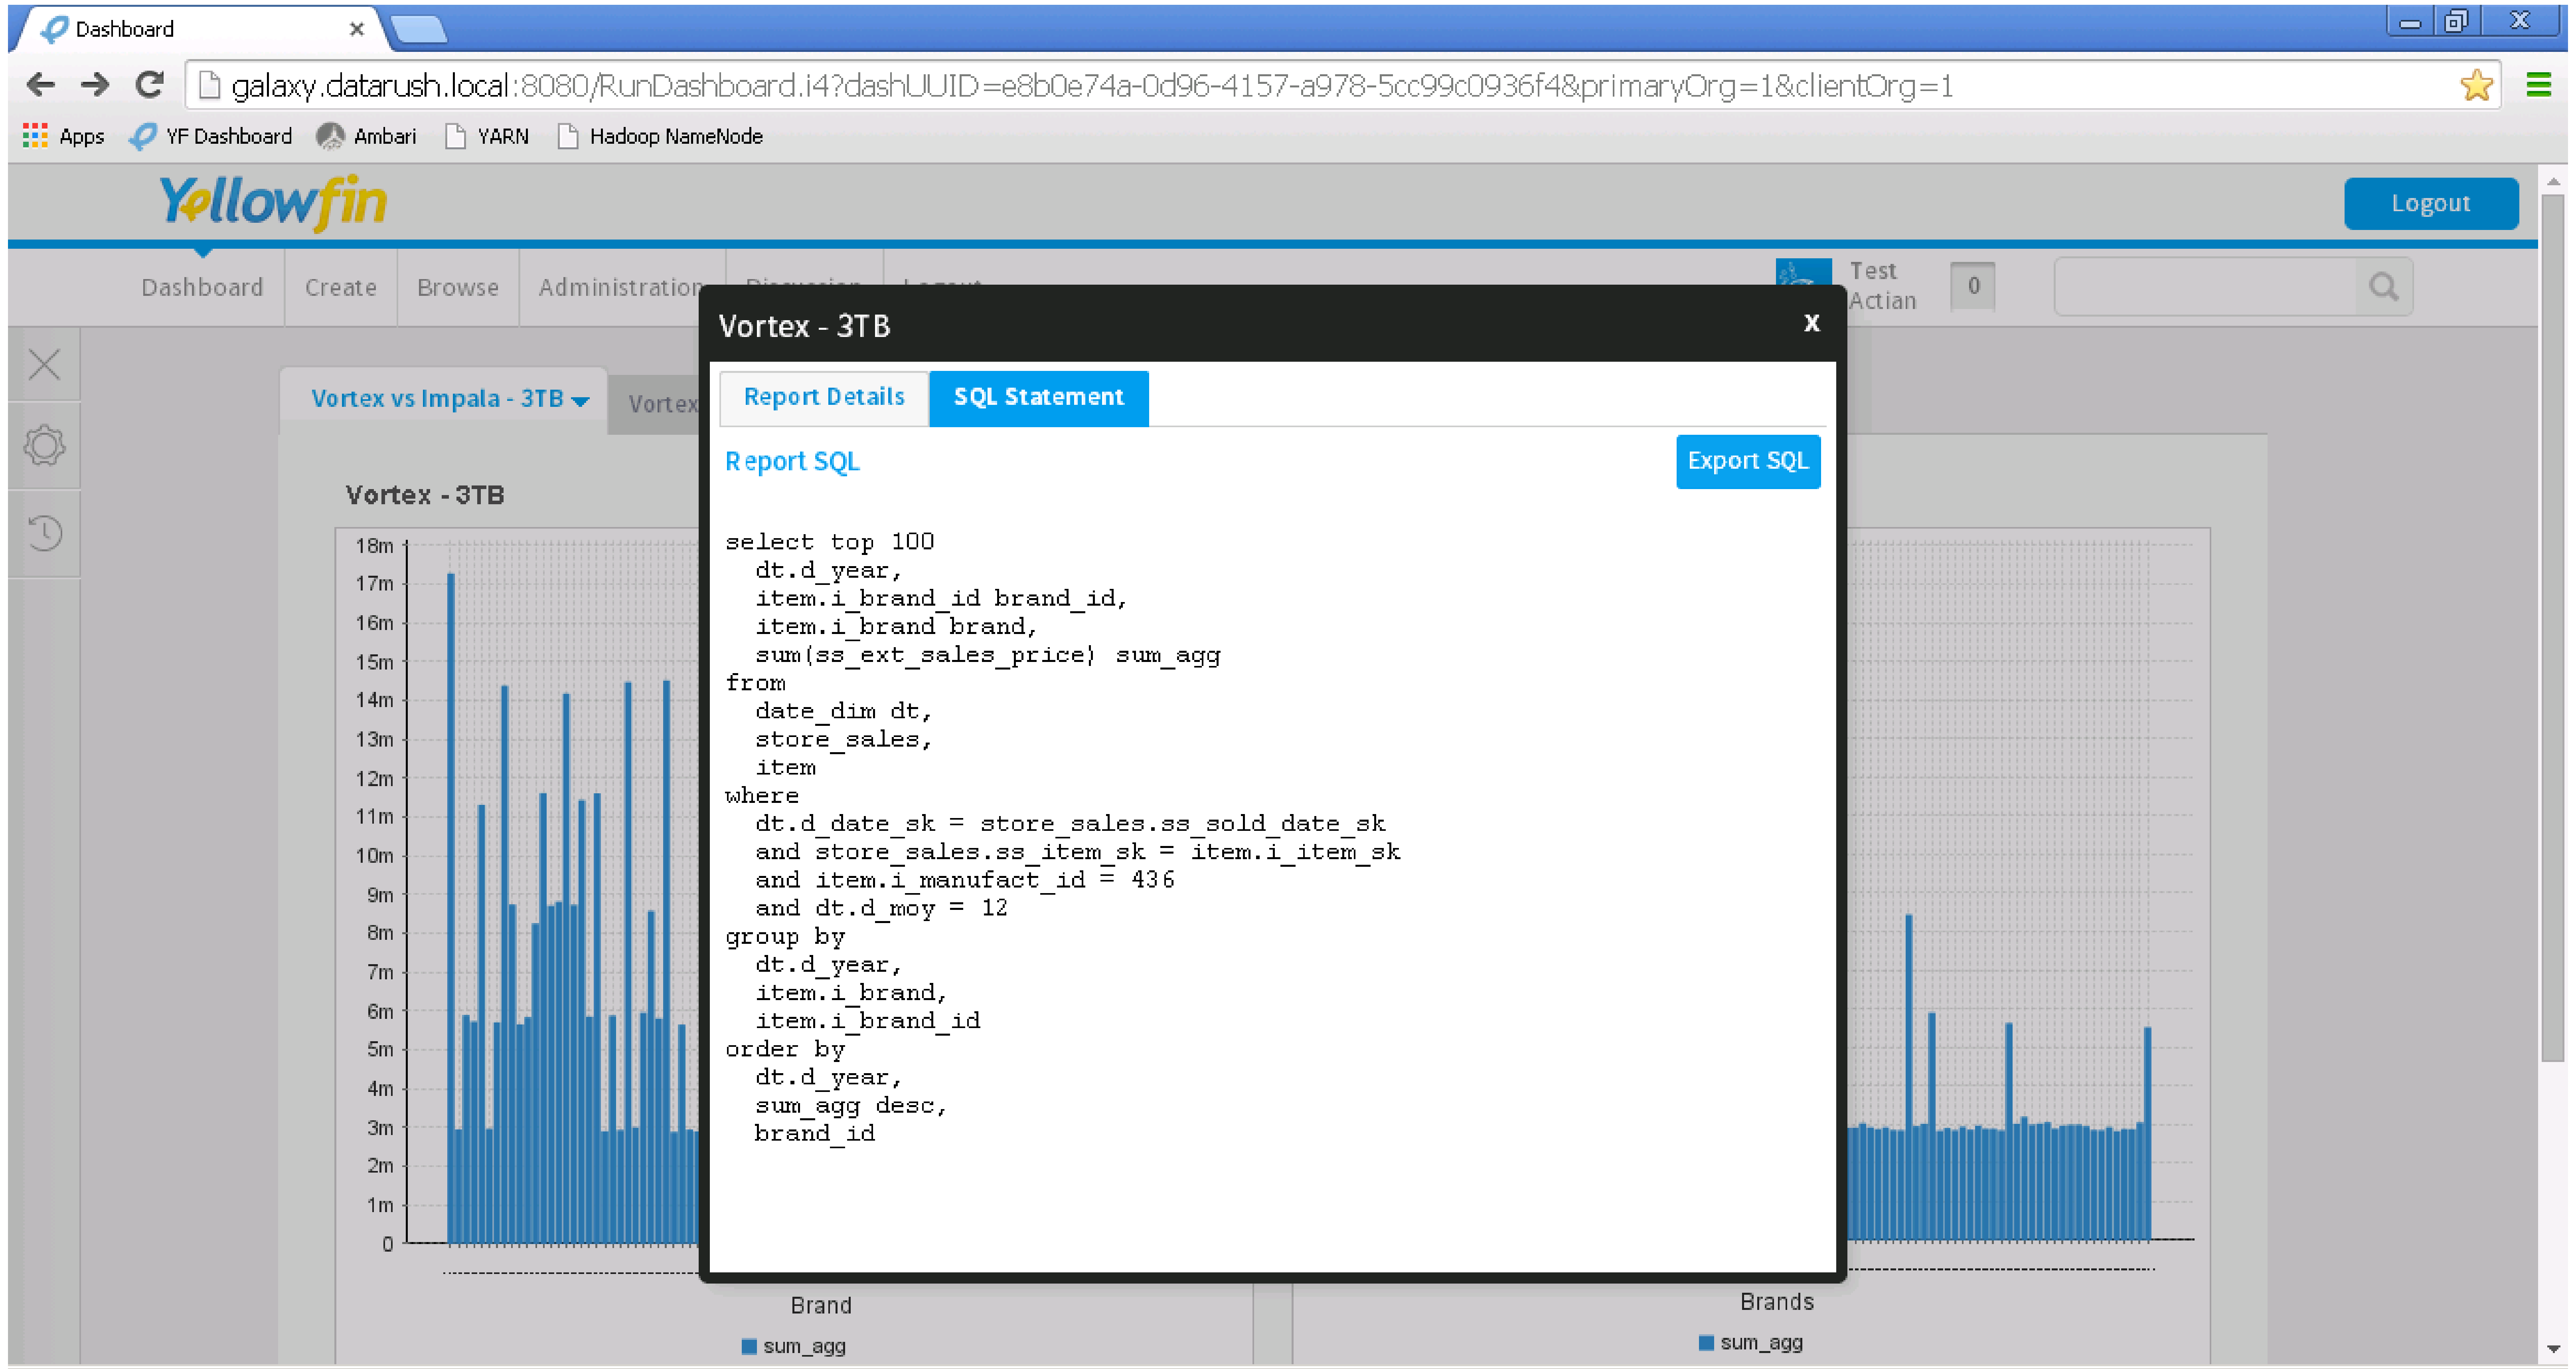

```
select top 100
  dt.d_year,
  item.i_brand_id brand_id,
  item.i_brand brand,
  sum(ss_ext_sales_price) sum_agg
from
  date_dim dt,
  store_sales,
  item
where
  dt.d_date_sk = store_sales.ss_sold_date_sk
  and store_sales.ss_item_sk = item.i_item_sk
  and item.i_manufact_id = 436
  and dt.d_moy = 12
group by
  dt.d_year,
  item.i_brand,
  item.i_brand_id
order by
  dt.d_year,
  sum_agg desc,
  brand_id
```

The Impala Equivalent Uses “Hints”

Note the use of partition keys

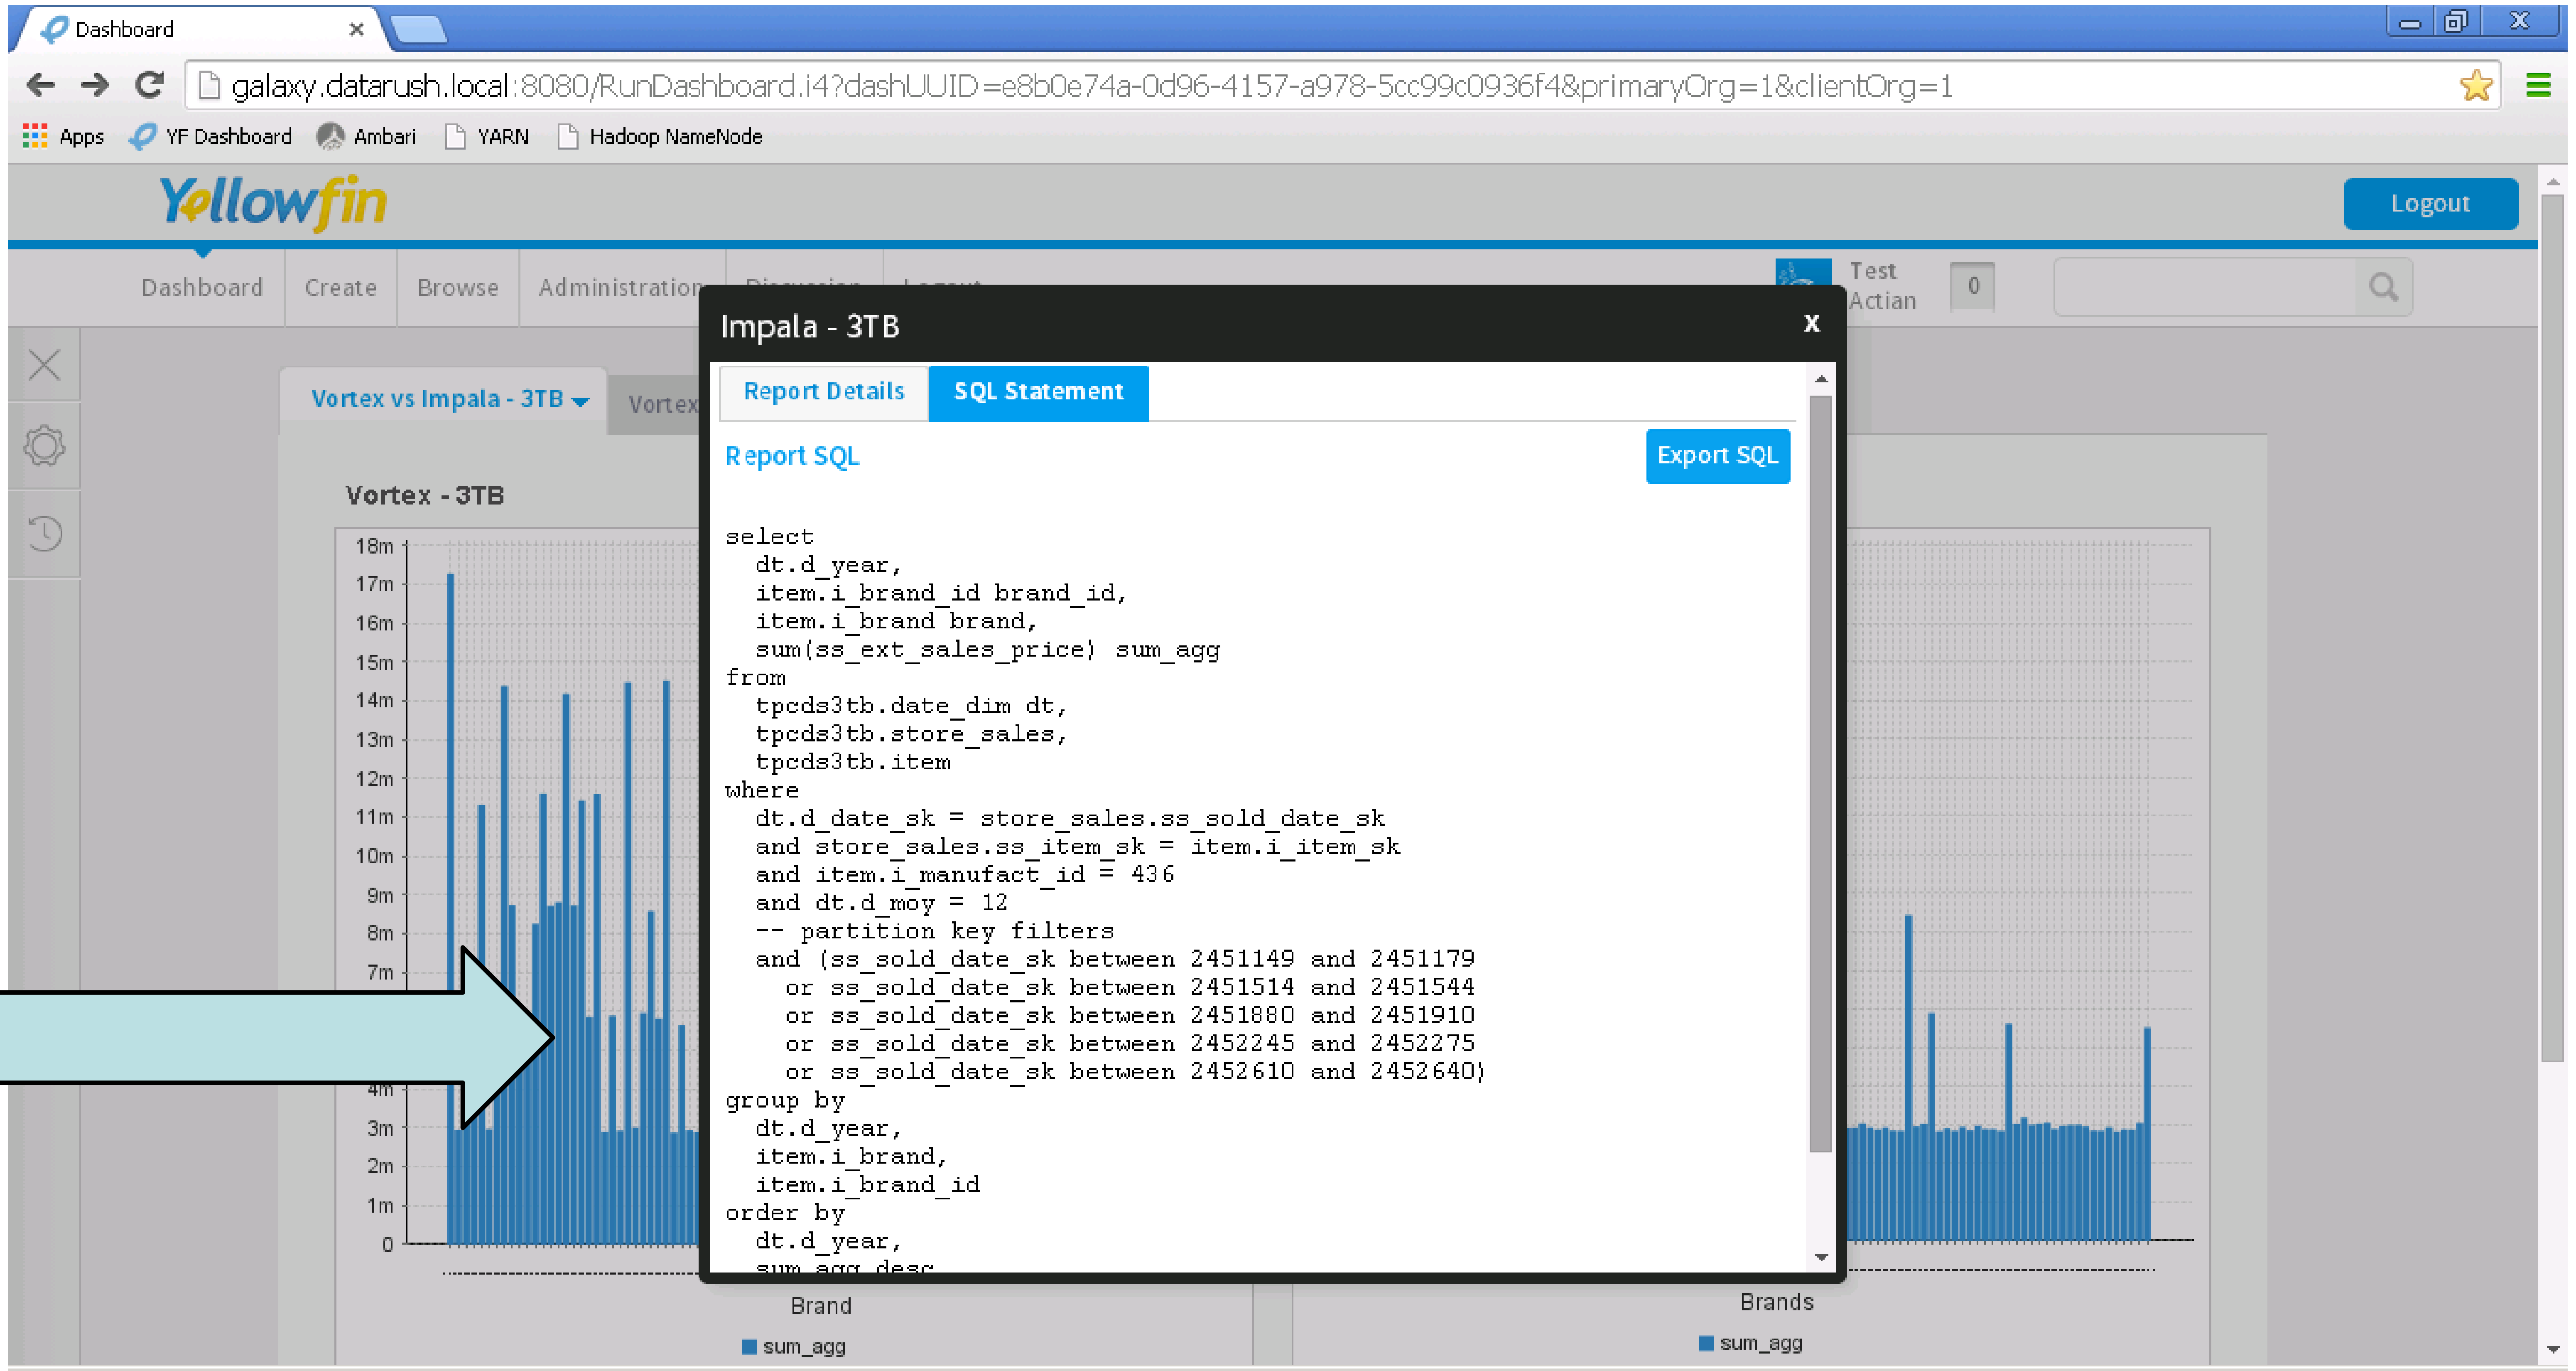

```
select
  dt.d_year,
  item.i_brand_id brand_id,
  item.i_brand brand,
  sum(ss_ext_sales_price) sum_agg
from
  tpcds3tb.date_dim dt,
  tpcds3tb.store_sales,
  tpcds3tb.item
where
  dt.d_date_sk = store_sales.ss_sold_date_sk
  and store_sales.ss_item_sk = item.i_item_sk
  and item.i_manufact_id = 436
  and dt.d_moy = 12
  -- partition key filters
  and (ss_sold_date_sk between 2451149 and 2451179
    or ss_sold_date_sk between 2451514 and 2451544
    or ss_sold_date_sk between 2451880 and 2451910
    or ss_sold_date_sk between 2452245 and 2452275
    or ss_sold_date_sk between 2452610 and 2452640)
group by
  dt.d_year,
  item.i_brand,
  item.i_brand_id
order by
  dt.d_year,
  sum_agg desc
```

CAPTAIN JAMES T. KIRK

I'M SORRY, I CAN'T HEAR YOU OVER THE
SOUND OF HOW AWESOME I AM.

Trickle Update Support

- The Kobayashi Maru of Hadoop
  - The design paradigm for HDFS is for data to be written once and read ever after.
  - Appending updated records to the end of a column/table or rewriting the entire table significantly impacts system performance
- The Solution – Positional Delta Trees
  - Enable on-line updates, without impacting read performance
  - Keep track of the tuple position of Inserts/Modifies/Deletes
  - Designed to make merging in of these updates fast by providing the tuple positions where differences have to be applied at update time.

Positional Delta Trees

Stable Table

SID	VAL	RID
0	A	0
1	B	1
2	C	2
3	D	3
4	E	4
5	F	5
6	G	6
7	H	7
8	I	8
9	J	9
10	K	10
11	L	11

Updates

INSERT a BEFORE A
INSERT a BEFORE a
MODIFY C INTO b
DELETE D
MODIFY F INTO C
INSERT d BEFORE G
INSERT d BEFORE G
INSERT e BEFORE G
DELETE H
DELETE I
INSERT i BEFORE J
INSERT h BEFORE J
INSERT g BEFORE J
MODIFY J INTO m
MODIFY K INTO h
DELETE L

Update Table

SID	VAL
0	A
1	B
2	C
3	D
4	E
5	F
6	G
7	H
8	I
9	J
10	K
11	L

SID	VAL	RID
0	a	0
0	a	1
0	A	2
1	B	3
2	b	4
4	E	5
5	c	6
6	d	7
6	d	8
6	e	9
6	G	10
9	I	11
9	H	12
9	G	13
9	M	14
10	h	15

SID	9	-
DELTA	3	1
CHILD		

SID	5	7
DELTA	1	3
CHILD		

SID	9	-
DELTA	2	-1
CHILD		

SID	0	0	2	3
Type	I	I	M	D
Value	a	a	b	-
RID	0	1	4	5

SID	5	6	6	6
Type	M	I	I	I
Value	c	d	d	e
	6	7	8	9

SID	7	8	9	-
Type	D	D	I	-
Value	-	-	i	-
	11	11	11	11

SID	9	9	9	-
Type	I	I	M	-
Value	h	g	m	-
	12	13	14	-

SID	10	11	-	-
Type	M	D	-	-
Value	h	-	-	-
	15	16	-	-

Data Security

- Access Control
- Role Separation
  - System Administrator & Database Administrator should not have access to all data
- Security Auditing
  - Ability to audit who accessed, or attempted to access, what and when
- Encryption
  - Data at rest – minimize performance impact by enabling at column level
  - Data in motion
  - File system

It’s SQL in Hadoop, Jim, but not as we know it

Highest Performing and Fully Industrialized SQL in Hadoop

Full ANSI SQL 92 support – enables use of ALL standard BI tools and apps

Fully ACID compliant – brings transactional integrity to Hadoop to prevent inaccurate results

Hadoop distribution agnostic - avoids vendor lock-in and provides customer flexibility

Update Capability – provides ability to update without impacting read performance

Highest Concurrency – allows your customers to have simultaneous users and tasks run without long wait times

Highly Performant – up to 30x faster than our closest competitor, Impala

Native DBMS Security - authentication, user and role-based security, data protection, and encryption

Mature, proven planner and fastest optimizer ensures customers can maximize number of nodes, CPU, memory and cache

Native in-Hadoop YARN – manage Hadoop resources automatically to prevent inefficiencies

Collaborative architecture - query native Hadoop file formats (like Parquet) without ingestion

Open APIs - allow read access to our block format

Beam it Down, Scotty!

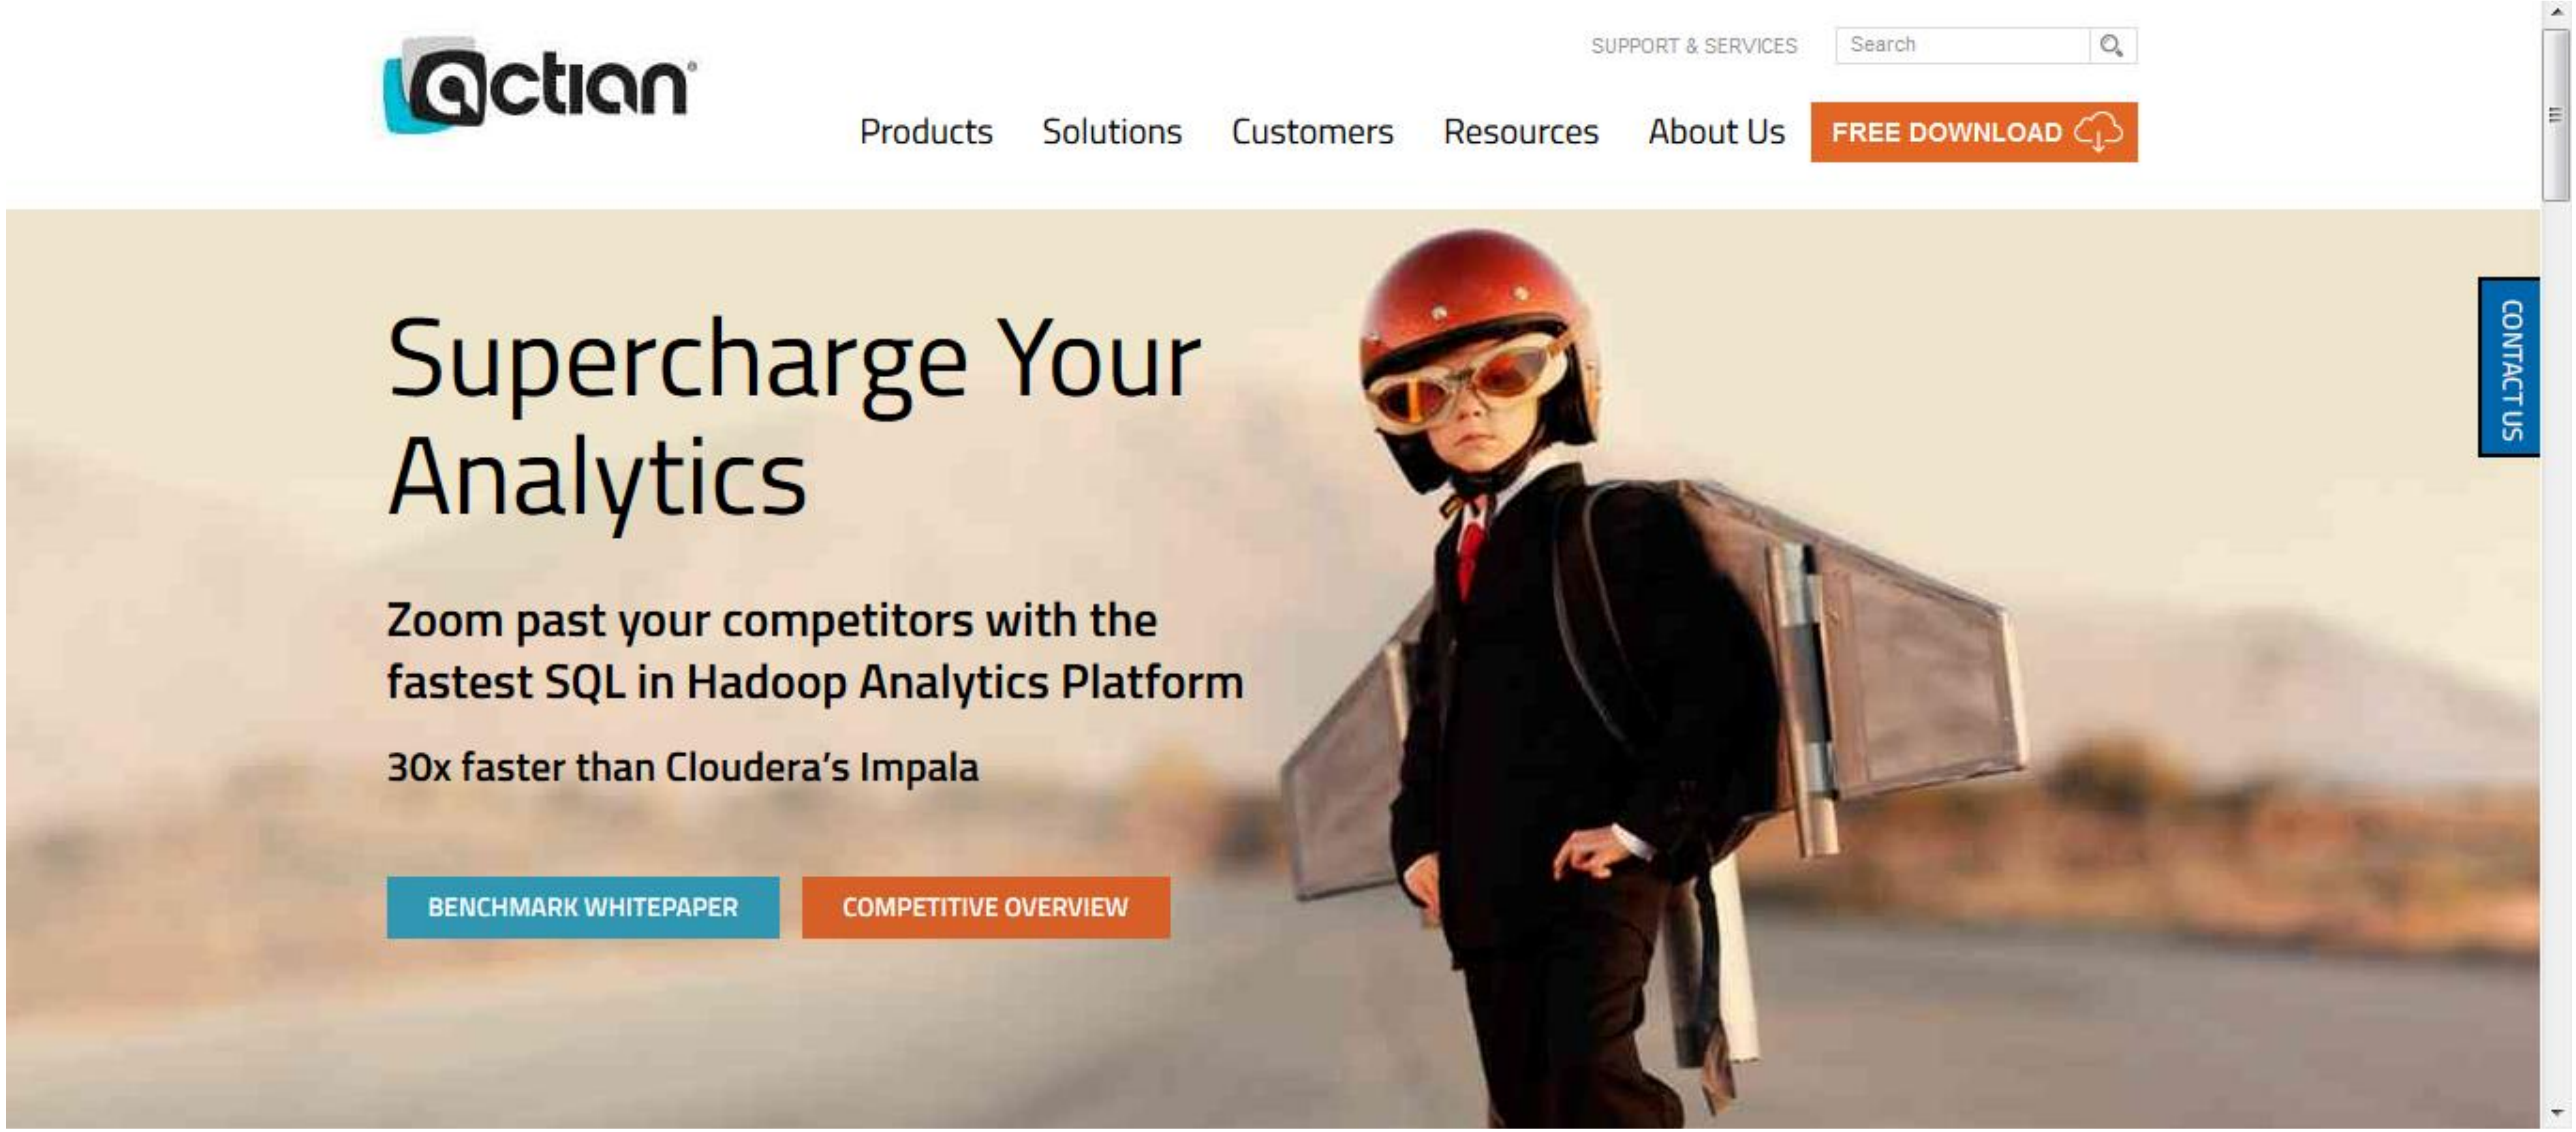

Supercharge Your Analytics

Zoom past your competitors with the fastest SQL in Hadoop Analytics Platform

30x faster than Cloudera's Impala

Strata+Hadoop WORLD

PRESENTED BY

O'REILLY cloudera

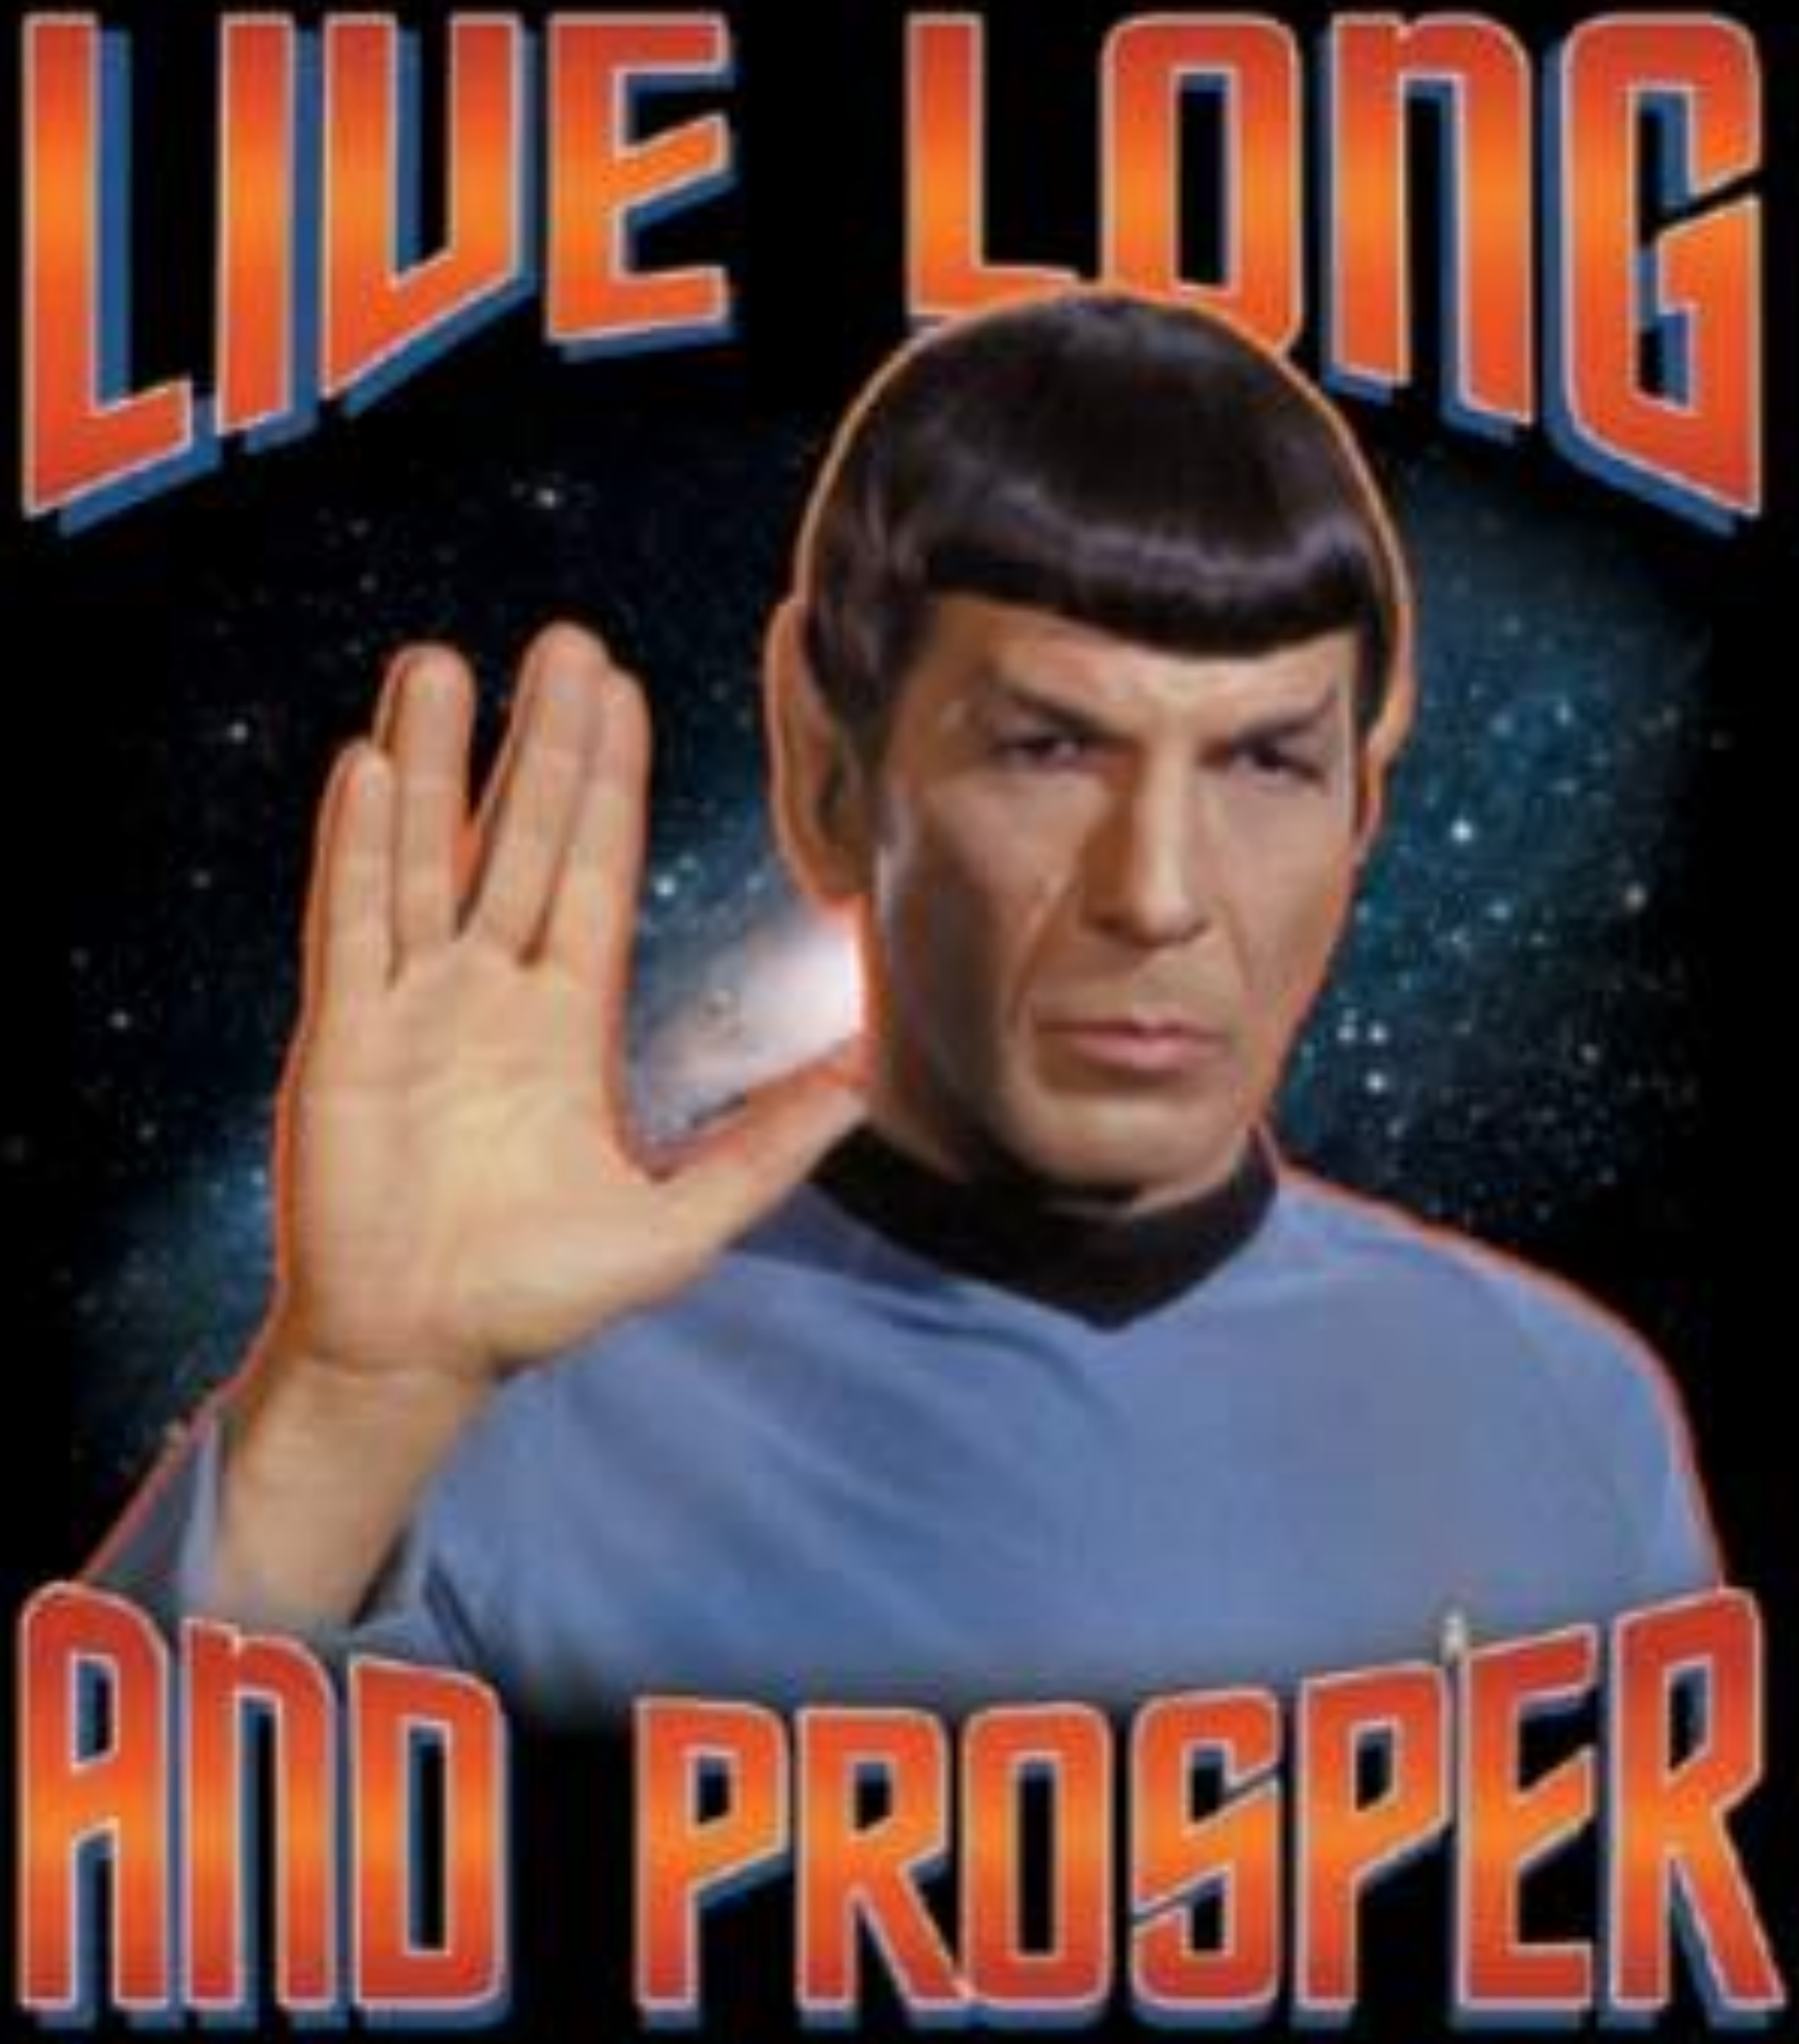

strataconf.com

#StrataHadoop