


goto;

conference



Join the conversation #gotocph





Click 'engage'
to rate sessions
and ask questions

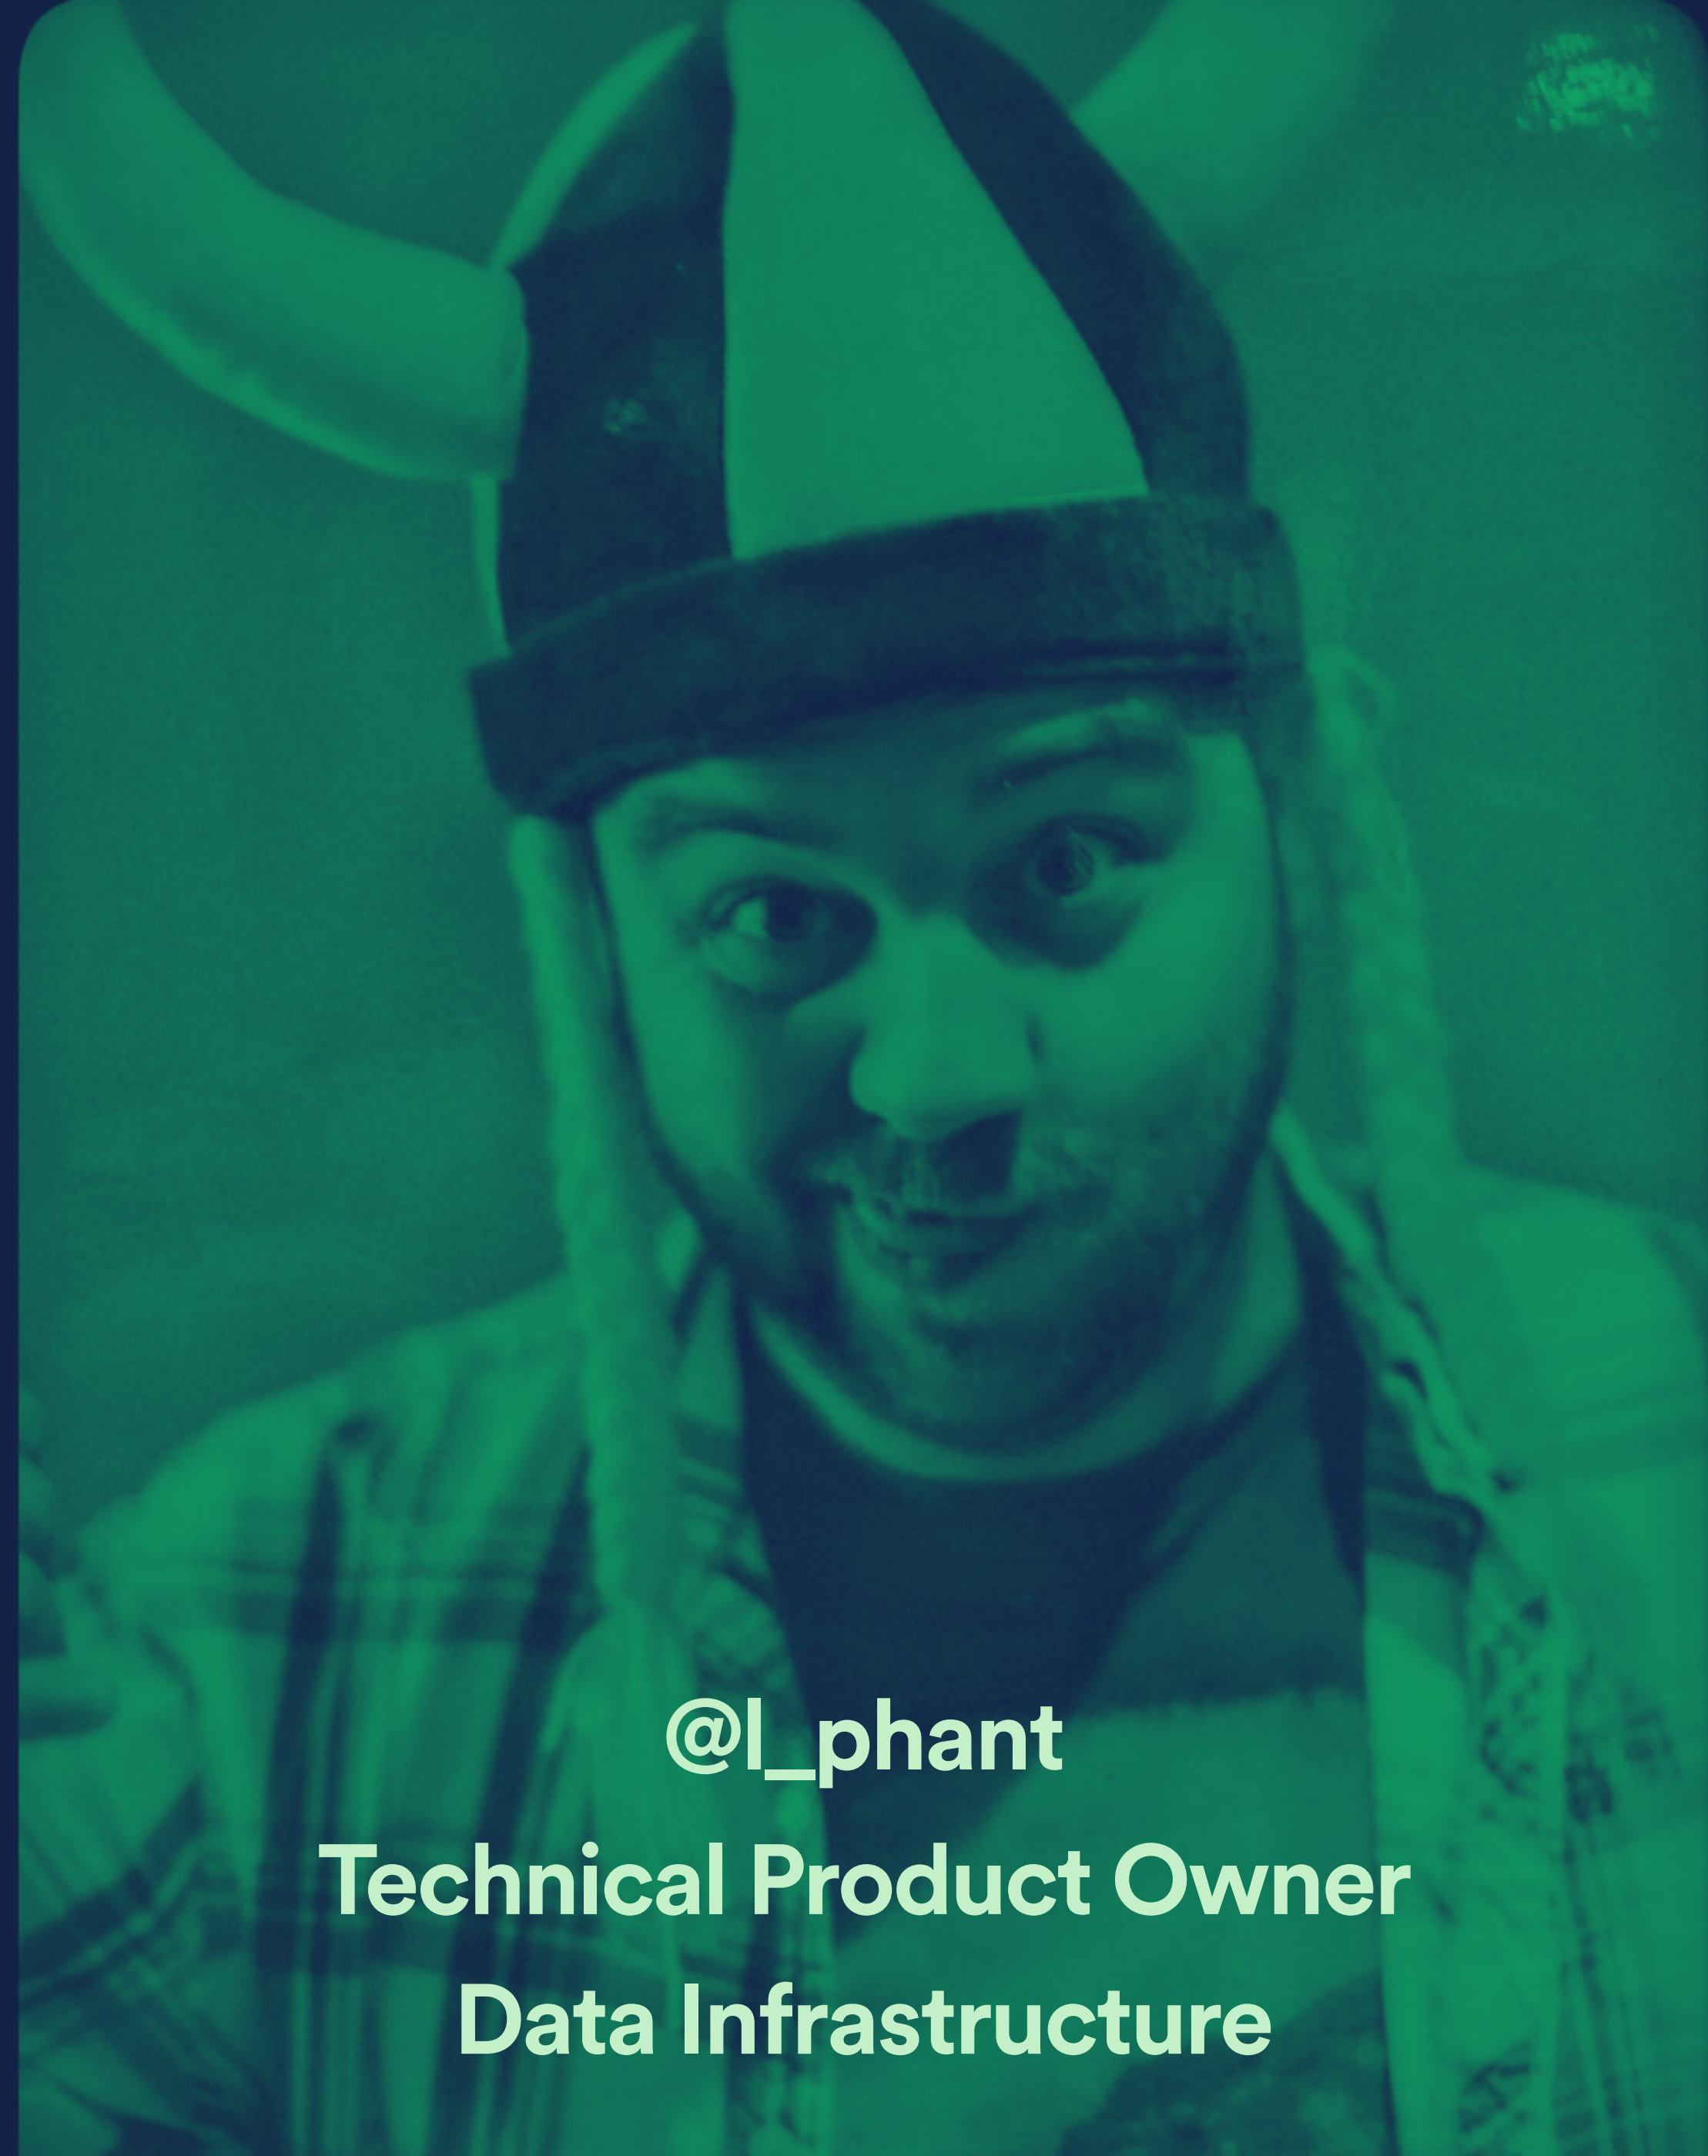


The Evolution of Hadoop at Spotify

Rafal Wojdyla (rav@spotify.com)

Josh Baer (jbx@spotify.com)





@l_phant

**Technical Product Owner
Data Infrastructure**



@ravwojdyla

**Data Engineer
Hadoop Team**

Overview

- Growing Pains
- Gaining Focus
- The Future

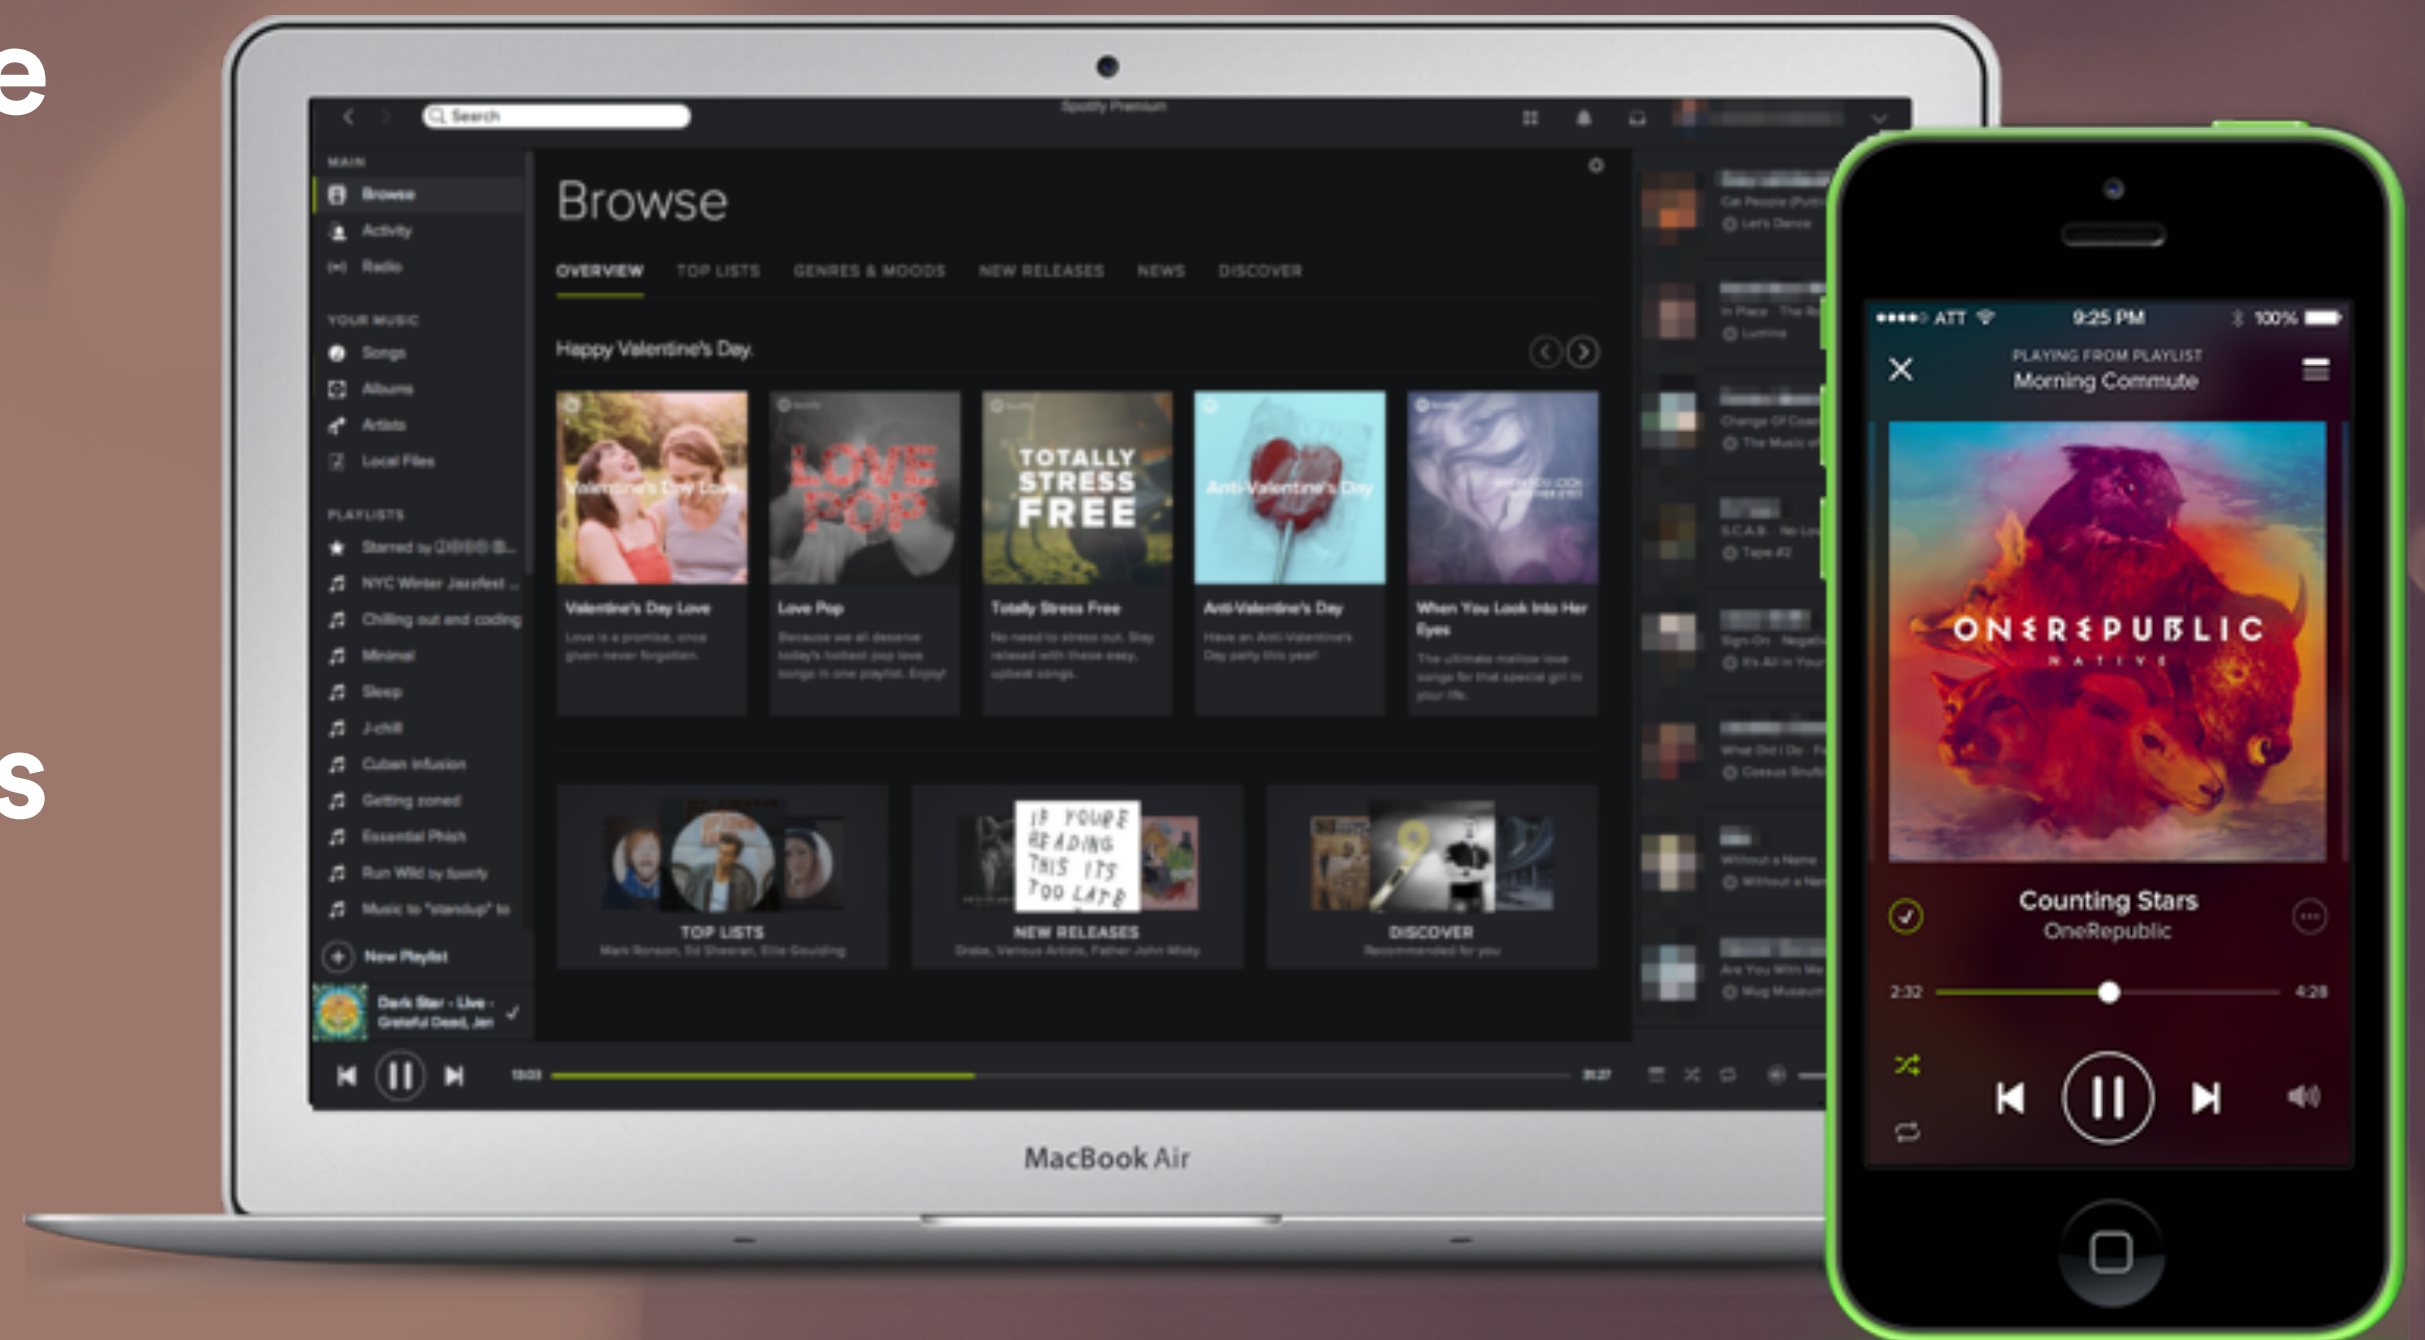


Growing Pains



What is Spotify?

- Music Streaming Service
- Launched in 2008
- Free and Premium Tiers
- Available in 58 Countries



A person is riding a skateboard on a paved surface. The skateboard has a graphic design on its deck. The rider's legs and feet are visible, wearing dark clothing and sneakers. The background is a blurred beach scene with several palm trees and a clear sky. The entire image has a warm, sepia-toned filter.

75+ Million Active Users

A sepia-toned photograph of two young women sitting on a concrete ledge overlooking a city street. The woman on the left is wearing sunglasses and a light-colored top, while the woman on the right is wearing a dark top and has her hair in a bun. They are both looking down at a smartphone held by the woman on the right. The background shows a busy city street with cars and buildings under a hazy sky. The text '30+ Million Songs' is overlaid in the center in a large, white, sans-serif font.

30+ Million Songs

A sepia-toned photograph of a crowd at a festival. In the foreground, a woman with long, wavy hair and sunglasses is smiling, wearing a patterned scarf. To her right, another woman with long hair and sunglasses is also smiling. In the background, several people have their arms raised in the air. A speed limit sign is visible on the left side of the image. The overall atmosphere is festive and joyful.

1+ Billion Plays/Day

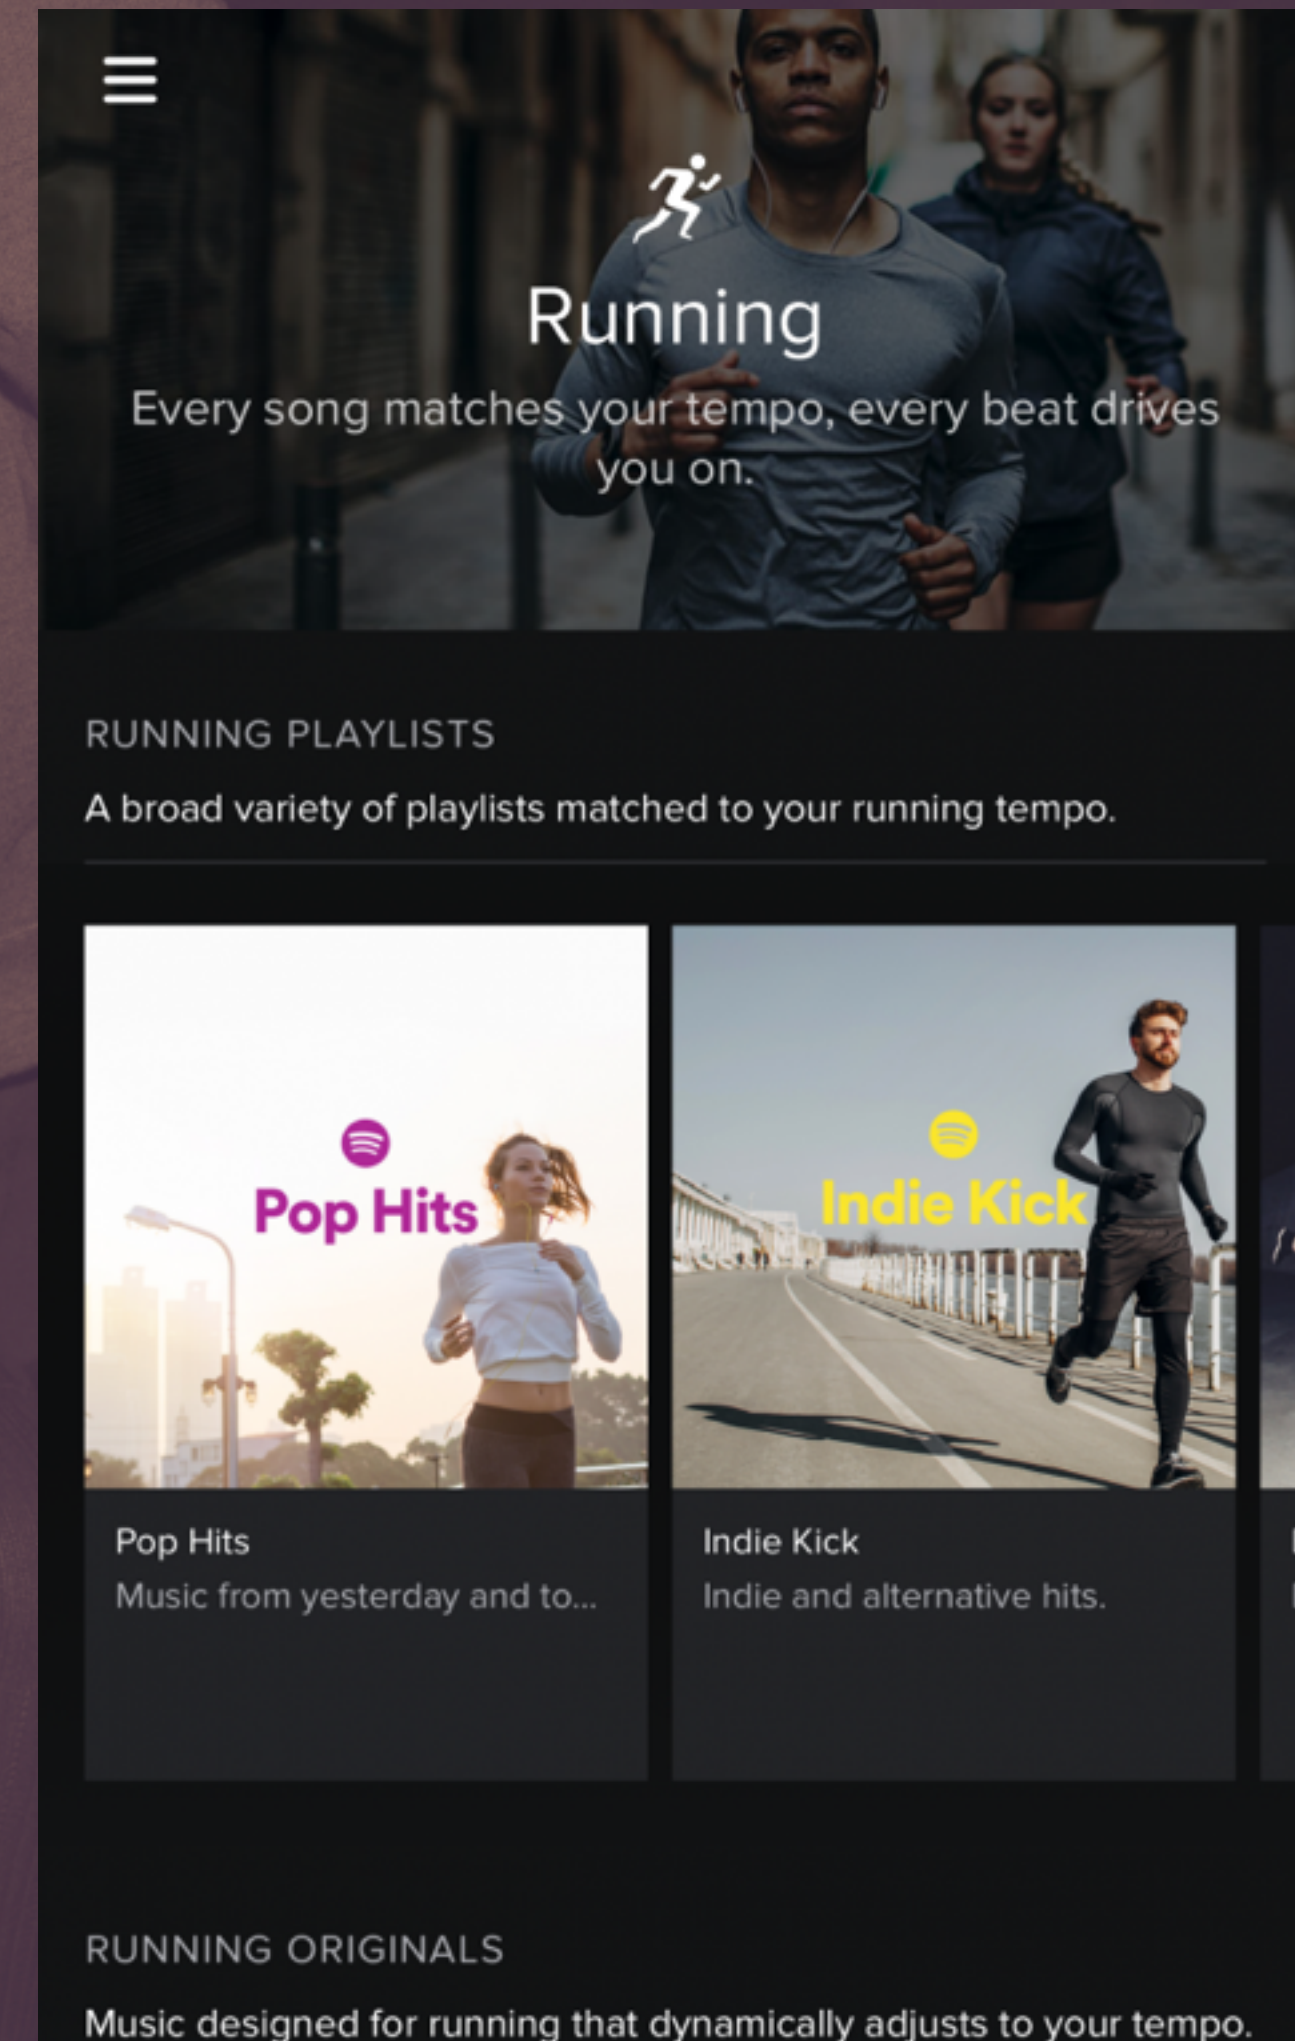
What is Spotify?

- Data Infrastructure
 - 1700 Hadoop Nodes
 - 62 PB Storage
 - 30 TB/day from user logs
 - 400 TB/day generated by Hadoop



Powered by Data

- Running App
 - Matches music to running tempo
 - Personalized running playlists in multiple tempos



Powered by Data

MAIN

Browse

Activity

Radio

YOUR MUSIC

Songs

Albums

Artists

Local Files

PLAYLISTS

Discover Weekly by S...

Modern Psychedelia ...

Grateful Dead and Fr...

The Most Beautif...

Loading...

New Playlist

Orchids

Monster Rally

+

Discover Weekly

PLAYLIST

Discover Weekly

Your weekly mixtape of fresh music. Enjoy new discoveries and deep cuts chosen just for you. Updated every Monday, so save your favourites!

Created by: Spotify · 30 songs, 2 hr 14 min

PAUSE

FOLLOWING

FOLLOWERS
1


Available Offline

	SONG	ARTIST	ALBUM		
+	The Unfallen Kingdom	Gramatik	Water 4 The Soul	6 days ago	3:48
+	Planet Caravan	Black Sabbath	Paranoid	6 days ago	4:29
<div></div> +	Orchids	Monster Rally	Return to Paradise	6 days ago	2:57
+	Their Pie	Mark Orton	Nebraska	6 days ago	2:06
+	Sound & Color	Alabama Shakes	Sound & Color	6 days ago	3:02
+	The First Daffodils	Bibio	Silver Wilkinson	6 days ago	2:57
+	Metronomic Underground - LP- Wagon Christ Mix	Stereolab	Emperor Tomato Ketchup	6 days ago	7:55
+	For 12	Other Lives	Tamer Animals	6 days ago	4:11
+	Knee-Deep in the North Sea	Portico Quartet	Knee-Deep In The North Sea	6 days ago	4:47
+	Om Nashi Me	Edward Sharpe & The Magnetic Zeros	Edward Sharpe & The Magnetic Zeros	6 days ago	6:16
+	The Lemon Of Pink	The Books	The Lemon Of Pink	6 days ago	4:40

0:18

2:56

LYRICS



CHART

Denmark Top 50

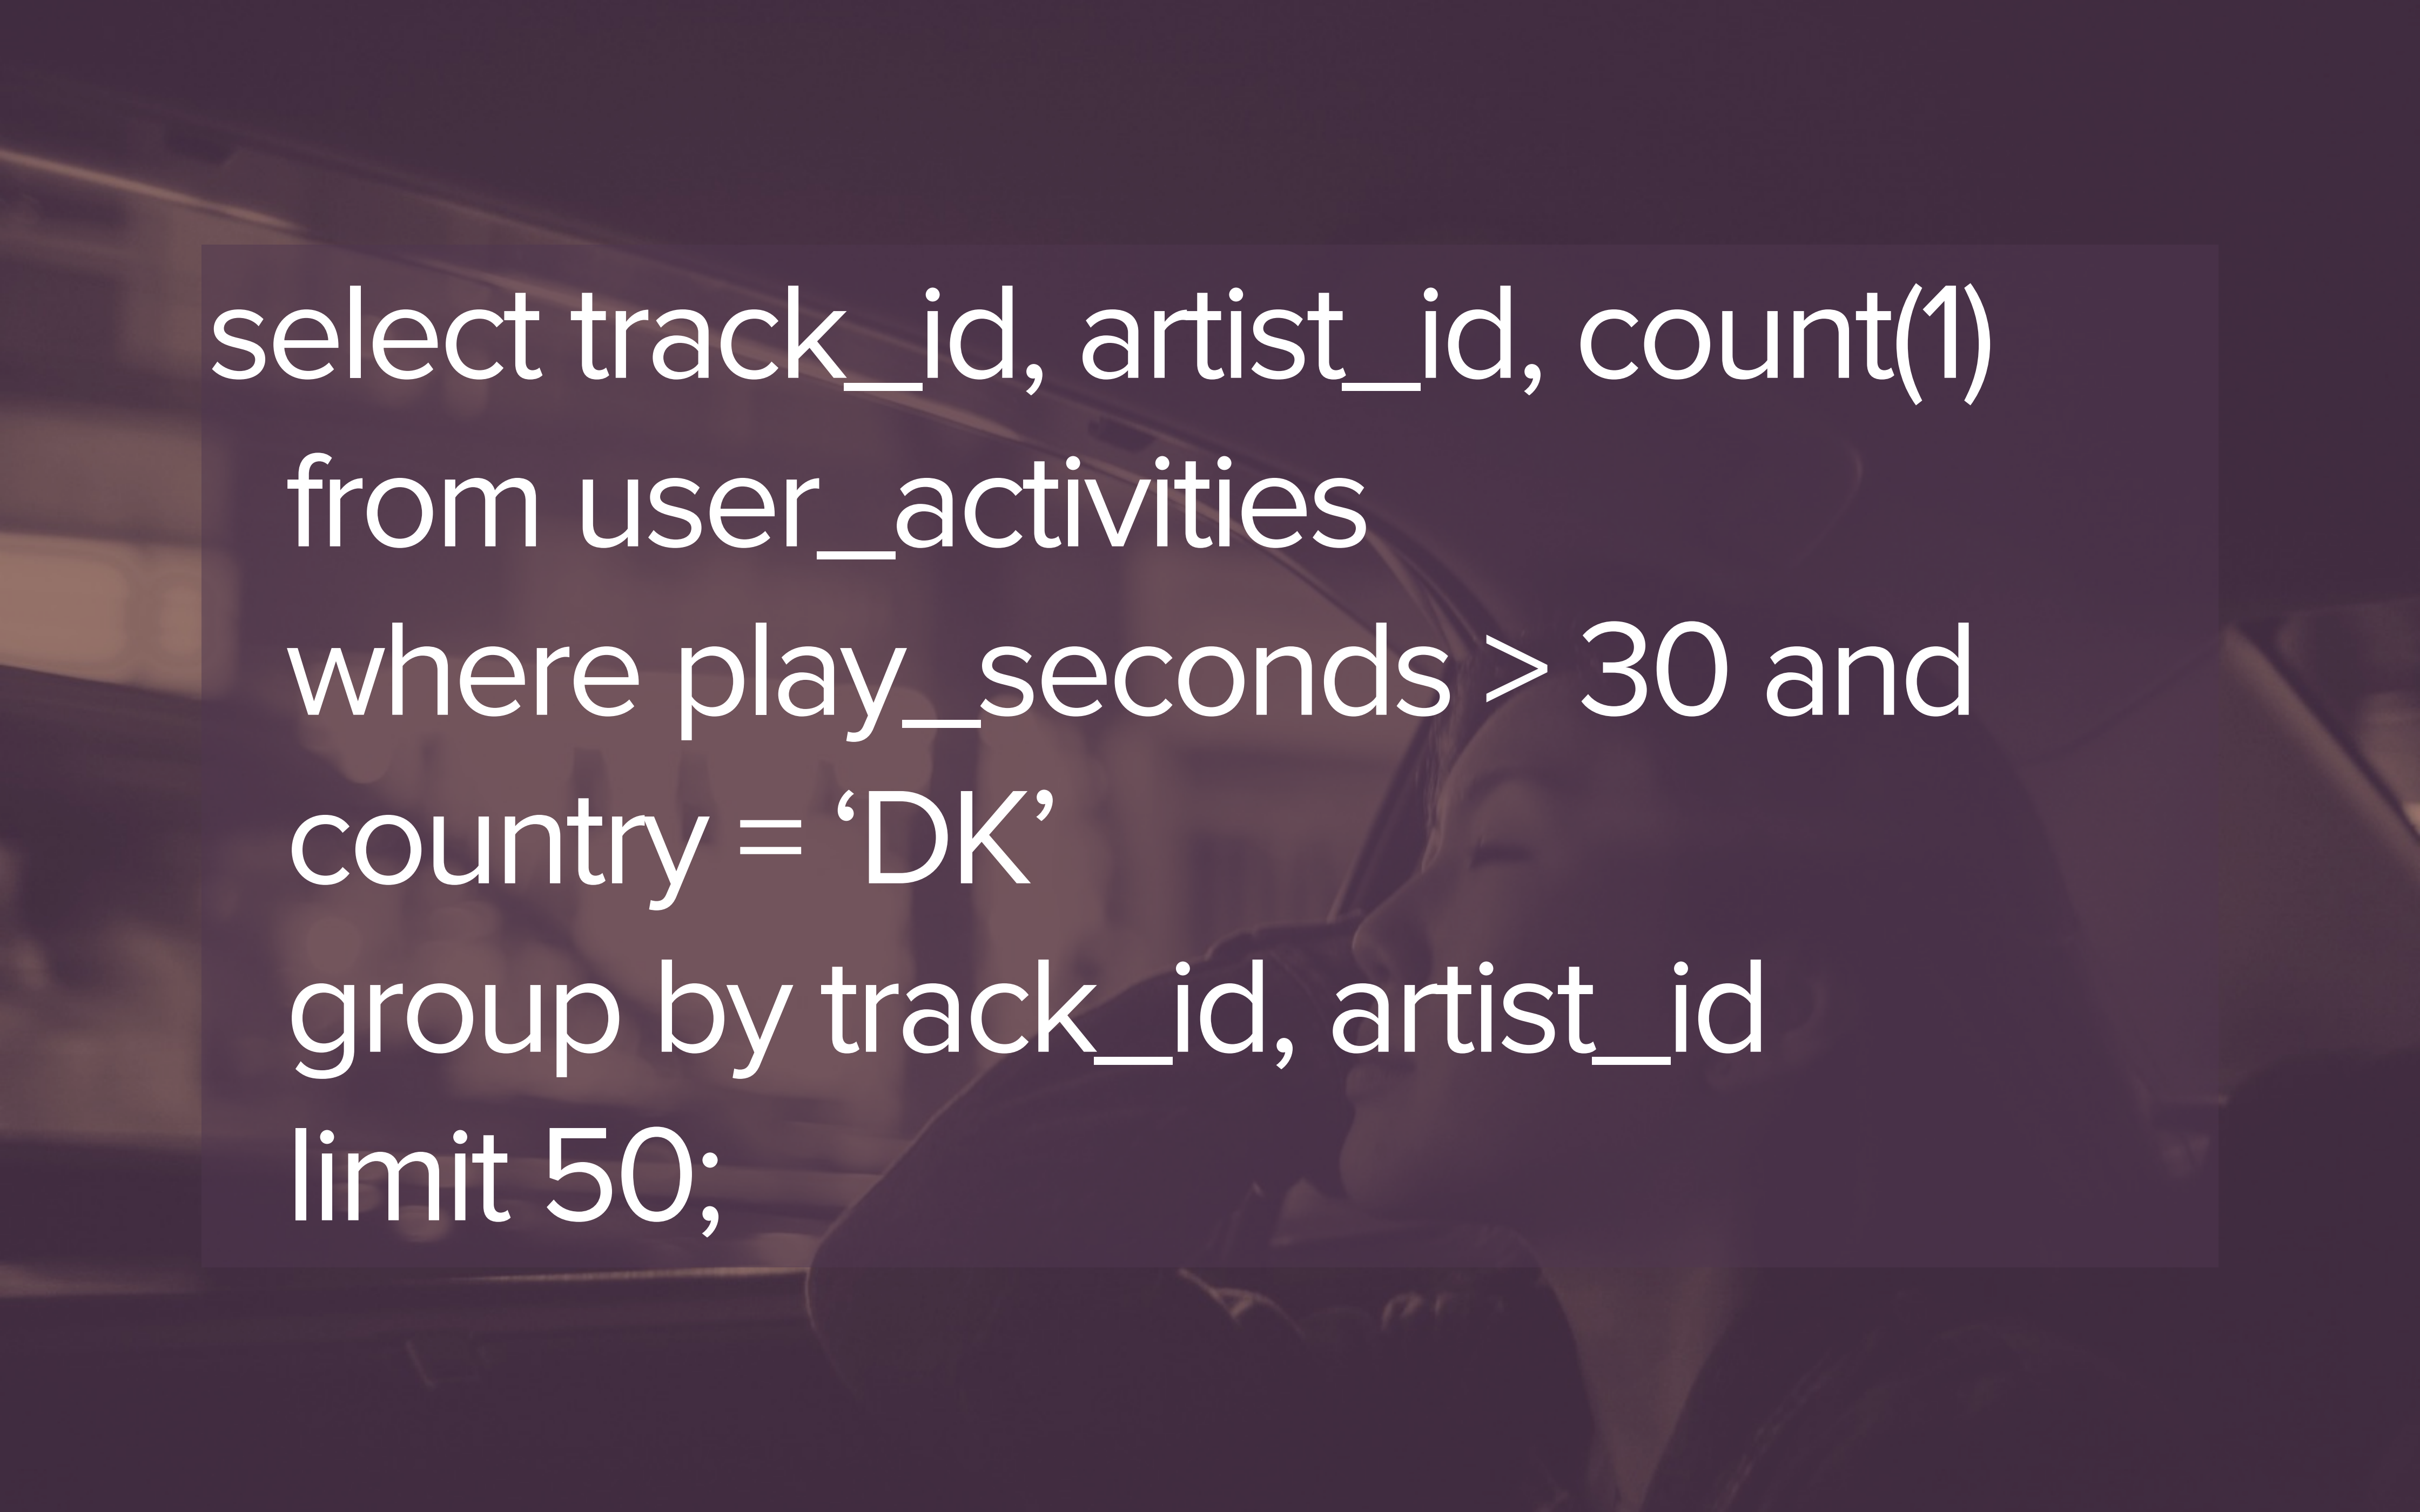
Your daily update of the most played tracks in Denmark right now.

● 1 New Entry • Last Updated: 10 hours ago


▶ PLAY

#		SONG	ARTIST	DAILY PLAYS
1	– +	What Do You Mean?	Justin Bieber	80,239
2	▲ +	Never Forget You	Zara Larsson, MNEK	61,690
3	▼ +	C'est La Vie	Gilli, MellemFingaMuzik	61,024
4	– +	Locked Away	R. City, Adam Levine	56,797
5	– +	Tyveri	hasan shah, Gilli	53,225
6	– +	Uanset	Rasmus Seebach	51,677
7	▲ +	Hula Hoop	OMI	48,481
8	▲ +	Stitches	Shawn Mendes	47,888
9	– +	Can't Feel My Face	The Weeknd	45,220
10	▲ +	Fucking Smuk	Page Four	44,900

EXPLICIT

A background image of a person singing into a microphone, overlaid with a semi-transparent dark purple rectangle containing SQL code.

```
select track_id, artist_id, count(1)
from user_activities
where play_seconds > 30 and
country = 'DK'
group by track_id, artist_id
limit 50;
```

Moving Data to Hadoop

- Raw data is complicated
 - Often dirty
 - Evolving structure
 - Duplication all over
- Getting data to a central processing point is HARD

```
10.123.133.333 - - [Mon, 3 June 2015 11:31:33 GMT] "GET /api/admin/job/
aggregator/status HTTP/1.1" 200 1847 "https://my.analytics.app/admin"
"Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_4) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/43.0.2357.81 Safari/537.36"
```

```
10.123.133.222 - - [Mon, 3 June 2015 11:31:43 GMT] "GET /api/admin/job/
aggregator/status HTTP/1.1" 200 1984 "https://my.analytics.app/admin"
"Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_4) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/43.0.2357.81 Safari/537.36"
```

```
10.123.133.222 - - [Mon, 3 June 2015 11:33:02 GMT] "GET /dashboard/
courses/1291726 HTTP/1.1" 304 - "https://my.analytics.app/admin"
"Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_4) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/43.0.2357.81 Safari/537.36"
```

```
10.321.145.111 - - [Mon, 3 June 2015 11:33:03 GMT] "GET /api/loggedInUser
HTTP/1.1" 304 - "https://my.analytics.app/dashboard/courses/1291726"
"Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_4) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/43.0.2357.81 Safari/537.36"
```

```
10.112.322.111 - - [Mon, 3 June 2015 11:33:03 GMT] "POST /api/
instrumentation/events/new HTTP/1.1" 200 2 "https://my.analytics.app/
dashboard/courses/1291726" "Mozilla/5.0 (Macintosh; Intel Mac OS X
10_10_4) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.81
Safari/537.36"
```

```
10.123.133.222 - - [Mon, 3 June 2015 11:33:02 GMT] "GET /dashboard/
courses/1291726 HTTP/1.1" 304 - "https://my.analytics.app/admin"
"Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_4) AppleWebKit/537.36
(KHTML, like Gecko) Chrome/43.0.2357.81 Safari/537.36"
```




LogArchiver

- Original method to transport logs from APs to HDFS
- Lasted from 2009 - 2013
- Relies on rsync/scp and cron to move files around



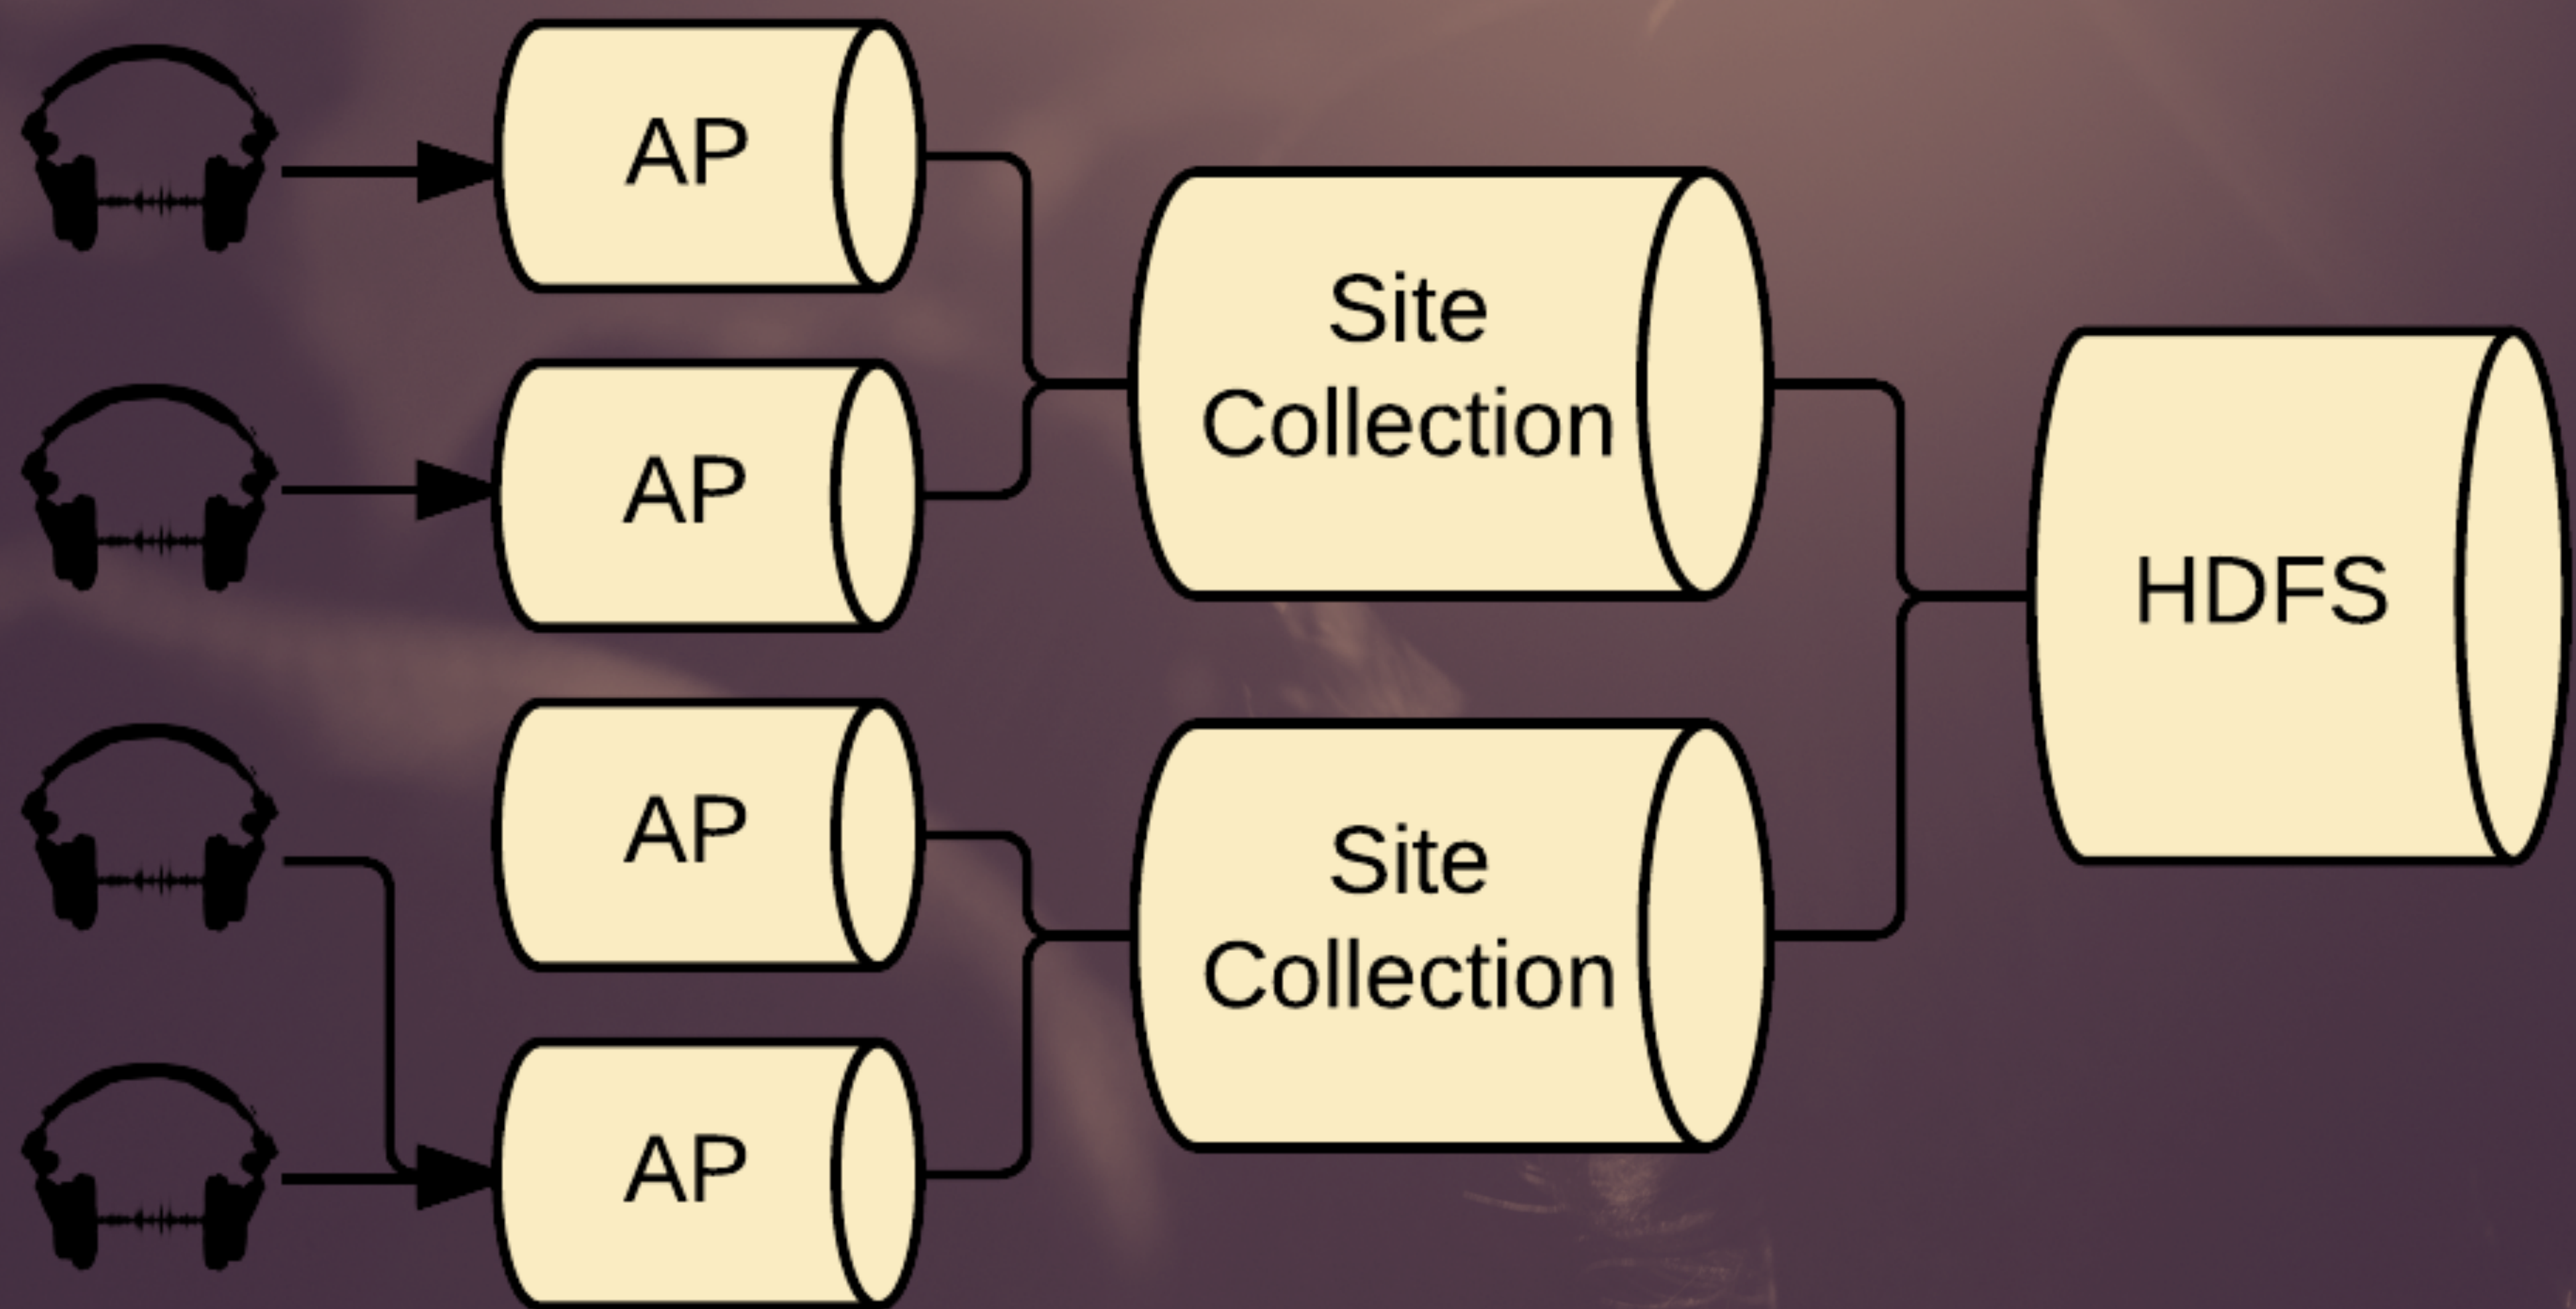
FAIL



ERR, LESSON?



Log -> HDFS latency reduced from hours to seconds!



Workflow Management Fail!

5	*	*	*	*	spotify-core	hadoop jar merge_hourly_logs.jar
15	*	*	*	*	spotify-core	hadoop jar aggregate_song_plays.jar
30	*	*	*	*	spotify-analytics	hadoop jar merge_song_metadata.jar
0	1	*	*	*	spotify-core	hadoop jar daily_aggregate.jar
0	2	*	*	*	spotify-core	hadoop jar calculate_toplist.jar



Luigi



<https://github.com/spotify/luigi>

The background of the slide is a photograph of three people riding bicycles on a sandy beach at sunset. The sky is a deep orange-red, and the ocean is visible in the distance. The people are silhouetted against the bright light of the setting sun.

```
[data-sci@sj-edge-a1 ~] $ hdfs dfs -ls /data
```

```
Found 3 items
```

```
drwxr-xr-x    - hdfs hadoop      0 2015-01-01 12:00 lake
drwxr-xr-x    - hdfs hadoop      0 2015-01-01 12:00 pond
drwxr-xr-x    - hdfs hadoop      0 2015-01-01 12:00 ocean
```

```
[data-sci@sj-edge-a1 ~] $ hdfs dfs -ls /data/lake
```

```
Found 1 items
```

```
drwxr-xr-x    - hdfs hadoop 1321451 2015-01-01 12:00 boats.txt
```

```
[data-sci@sj-edge-a1 ~] $ hdfs dfs -cat /data/lake/boats.txt
```

```
...
```




snakebite

<https://github.com/spotify/snakebite>


```
$ time for i in {1..100}; do hdfs dfs -ls / > /dev/null; done  
real 3m32.014s  
user 6m15.891s  
sys 0m18.821s
```

```
$ time for i in {1..100}; do snakebite ls / > /dev/null; done  
real 0m34.760s  
user 0m29.962s  
sys 0m4.512s
```



Gaining Focus




Hadoop Availability

- In 2013:
 - Hadoop expanded to 200 nodes
 - Critical but not very reliable
- Created a 'squad' with two missions:
 - Migrate to a new distribution with Yarn
 - Make Hadoop reliable

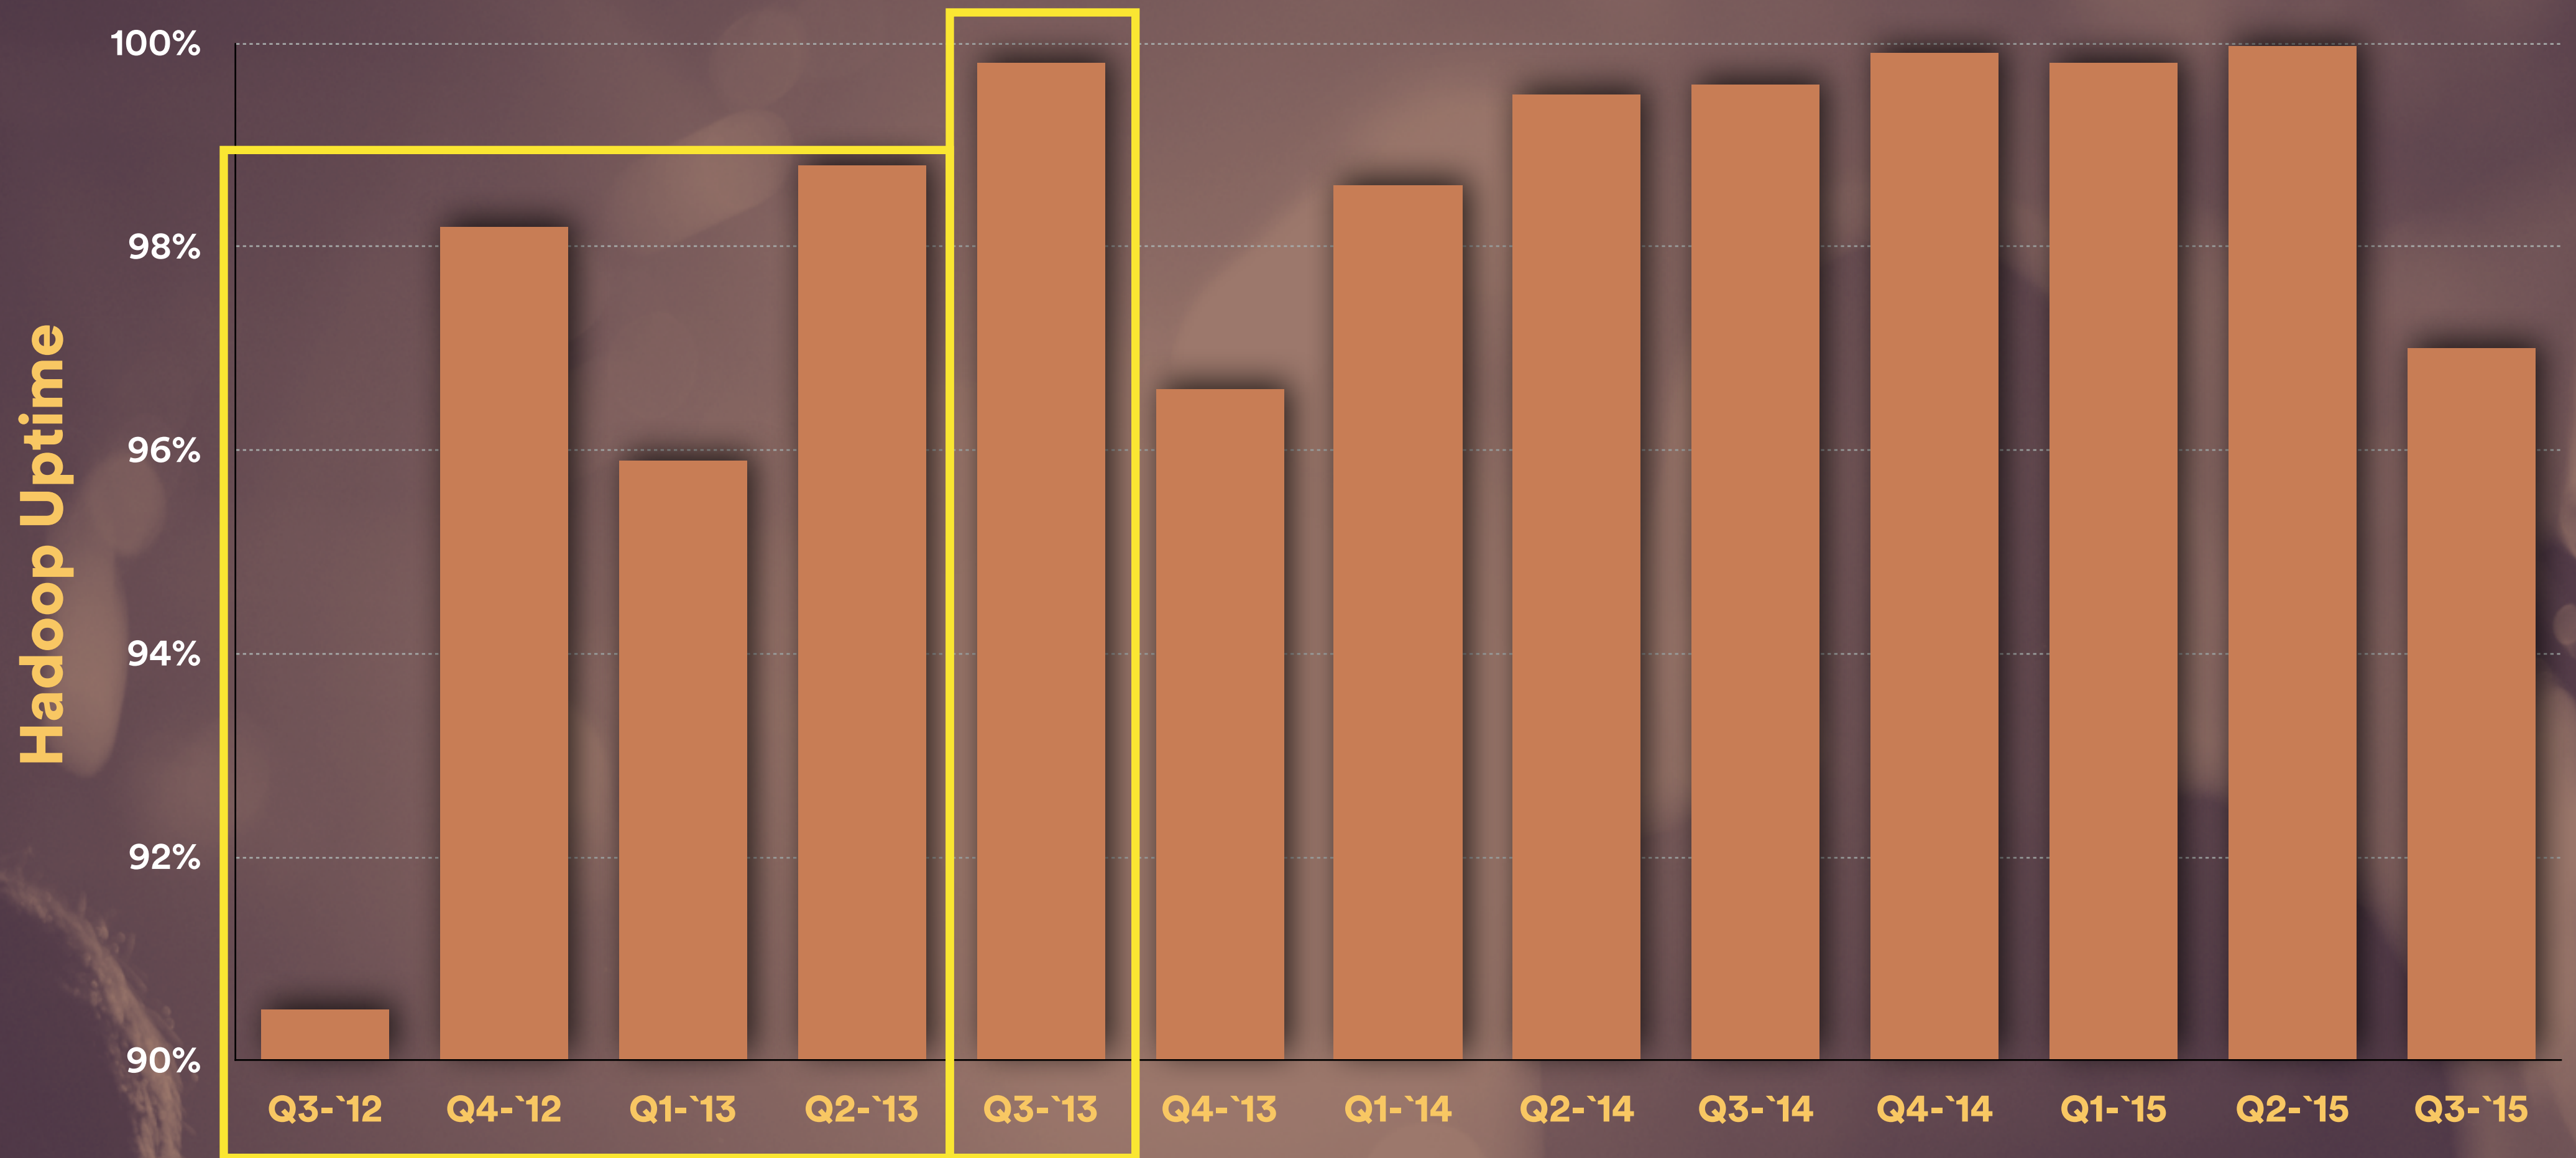
How did we do?




How did we do?




How did we do?




How did we do?



How did we do?



How did we do?



What happened in the last quarter?

- Expanded our cluster from ~1200 nodes to ~1700 nodes
- When you scale Hadoop, the bugs in the code scale with it
 - HDFS-5790
 - HDFS-6425

A classical sculpture of a man, likely a Greek or Roman figure, shown in a state of distress. He is leaning forward, with his right hand covering his face, suggesting grief or shame. The sculpture is made of a light-colored material, possibly marble or plaster, and is set against a dark, textured background. The lighting highlights the contours of his body and the texture of his hair and clothing.

**Uhh ohh.... I think I
made a mistake**



```
[2014.03.12 16:48:02 | data-sci@edge-1 | /home/data-sci/development] $ snakebite rm -R /team/disco/ test-10/
```


A clay sculpture of a person, possibly a woman, with their head bowed and hands covering their face in a gesture of distress or grief. The sculpture is light-colored and textured, contrasting with the dark, solid background. The lighting highlights the contours of the body and the texture of the clay.

```
$ snakebite rm -R /team/disco/ test-10/
```


A clay sculpture of a person, possibly a woman, with their head buried in their hands. The sculpture is made of a light-colored material, possibly clay or plaster, and is set against a dark, solid background. The person's head is tilted back, and their hands are pressed against their face, covering their eyes and nose. The sculpture is overlaid with a semi-transparent brown filter, which makes the text in the foreground stand out.

disco/ test-10





MOTHER OF GOD



\$ snakebite rm -R /team/disco/ test-10/

OK: Deleted /team/disco


Goodbye Data!
(1PB)

Lessons Learned

- “Sit on your hands before you type” - Wouter de Bie
- Users will always want to retain data!
- Remove superusers from ‘edgenodes’
- Moving to trash = client-side implementation

The Wild Wild West

Pre-Production



Going from Python to JVM

- Most of our jobs were Hadoop (python) streaming
- Lots of failures, slow performance
- Had to find a better way




Moving from Python to Crunch

- Investigated several frameworks*
- Selected Crunch:
 - Real types - compile time error detection, better testability
 - Higher level API - let the framework optimize for you
 - Better performance #JVM_FTW

*thewit.ch/scalding_crunchy_pig

Crunch vs Hadoop Streaming Benchmark



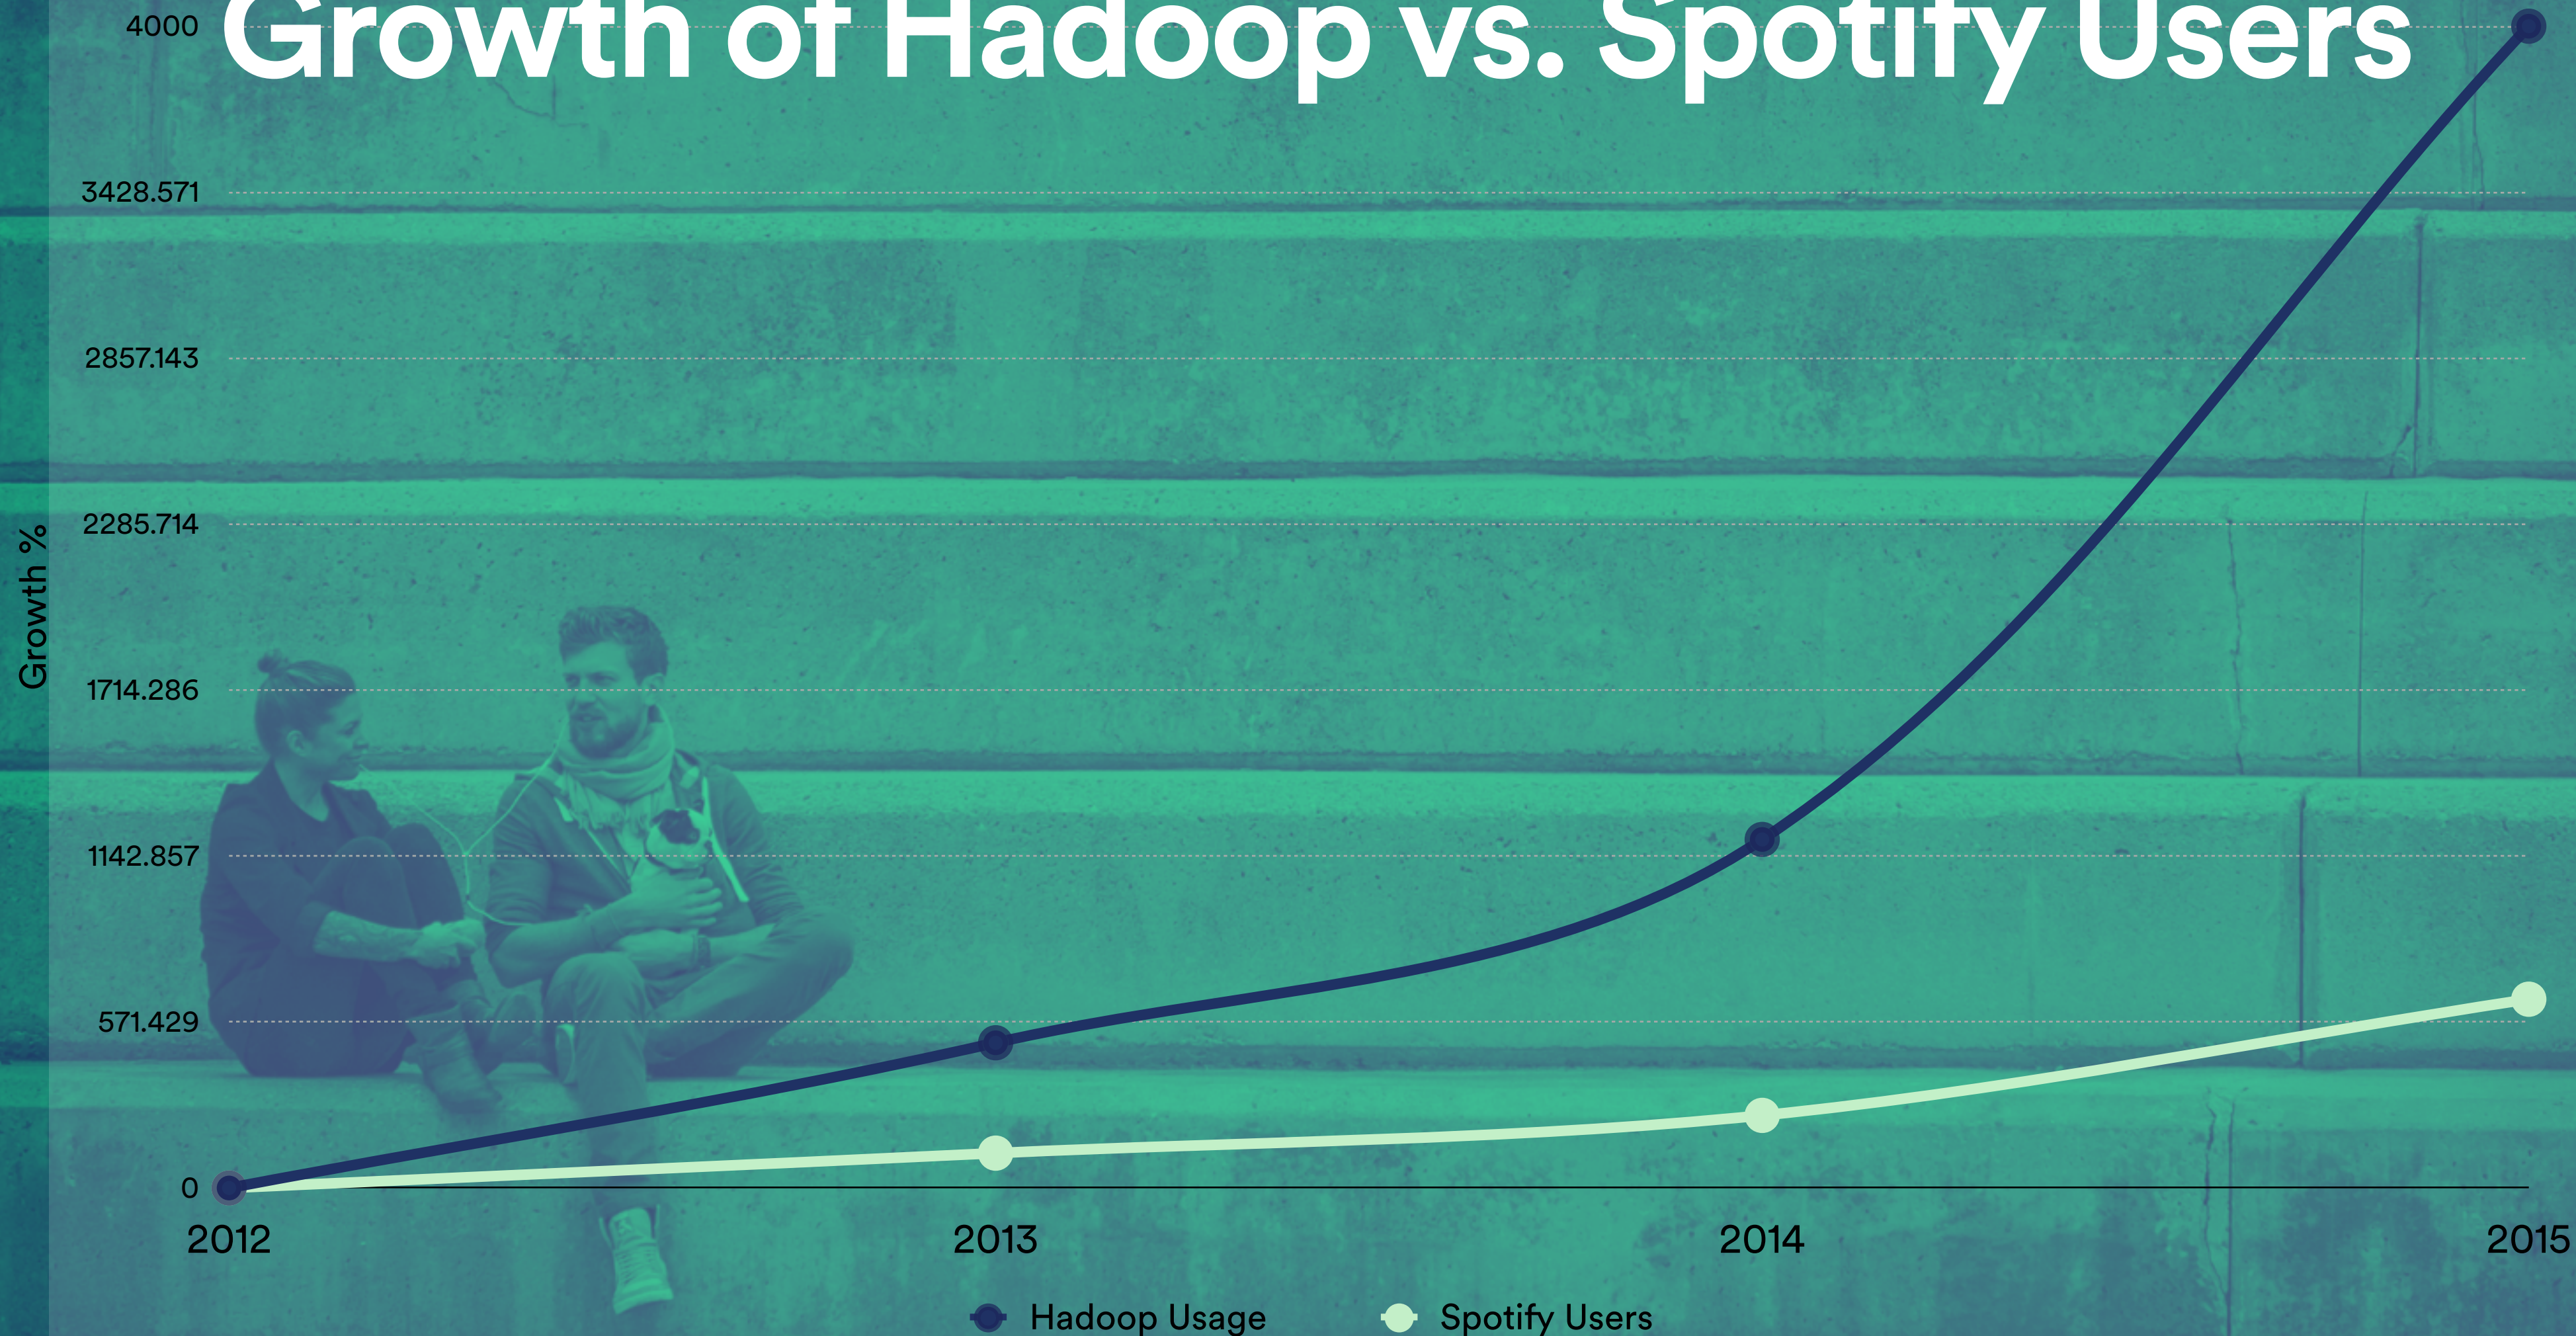
Let's Review

- Getting data into Hadoop
- Deploying data pipelines
- Increasing availability and reliability of infrastructure
- Killing it with performance

The Future



Growth of Hadoop vs. Spotify Users



A man and a woman are running together, captured in a dynamic pose. The image is overlaid with a teal gradient, which serves as a background for the text. The woman is in the foreground, leaning forward, and the man is slightly behind her, also running. Both are smiling, conveying a sense of energy and movement.

Explosive Growth

- Increased Spotify Users
- Increased Use Cases
- Increased Engineers

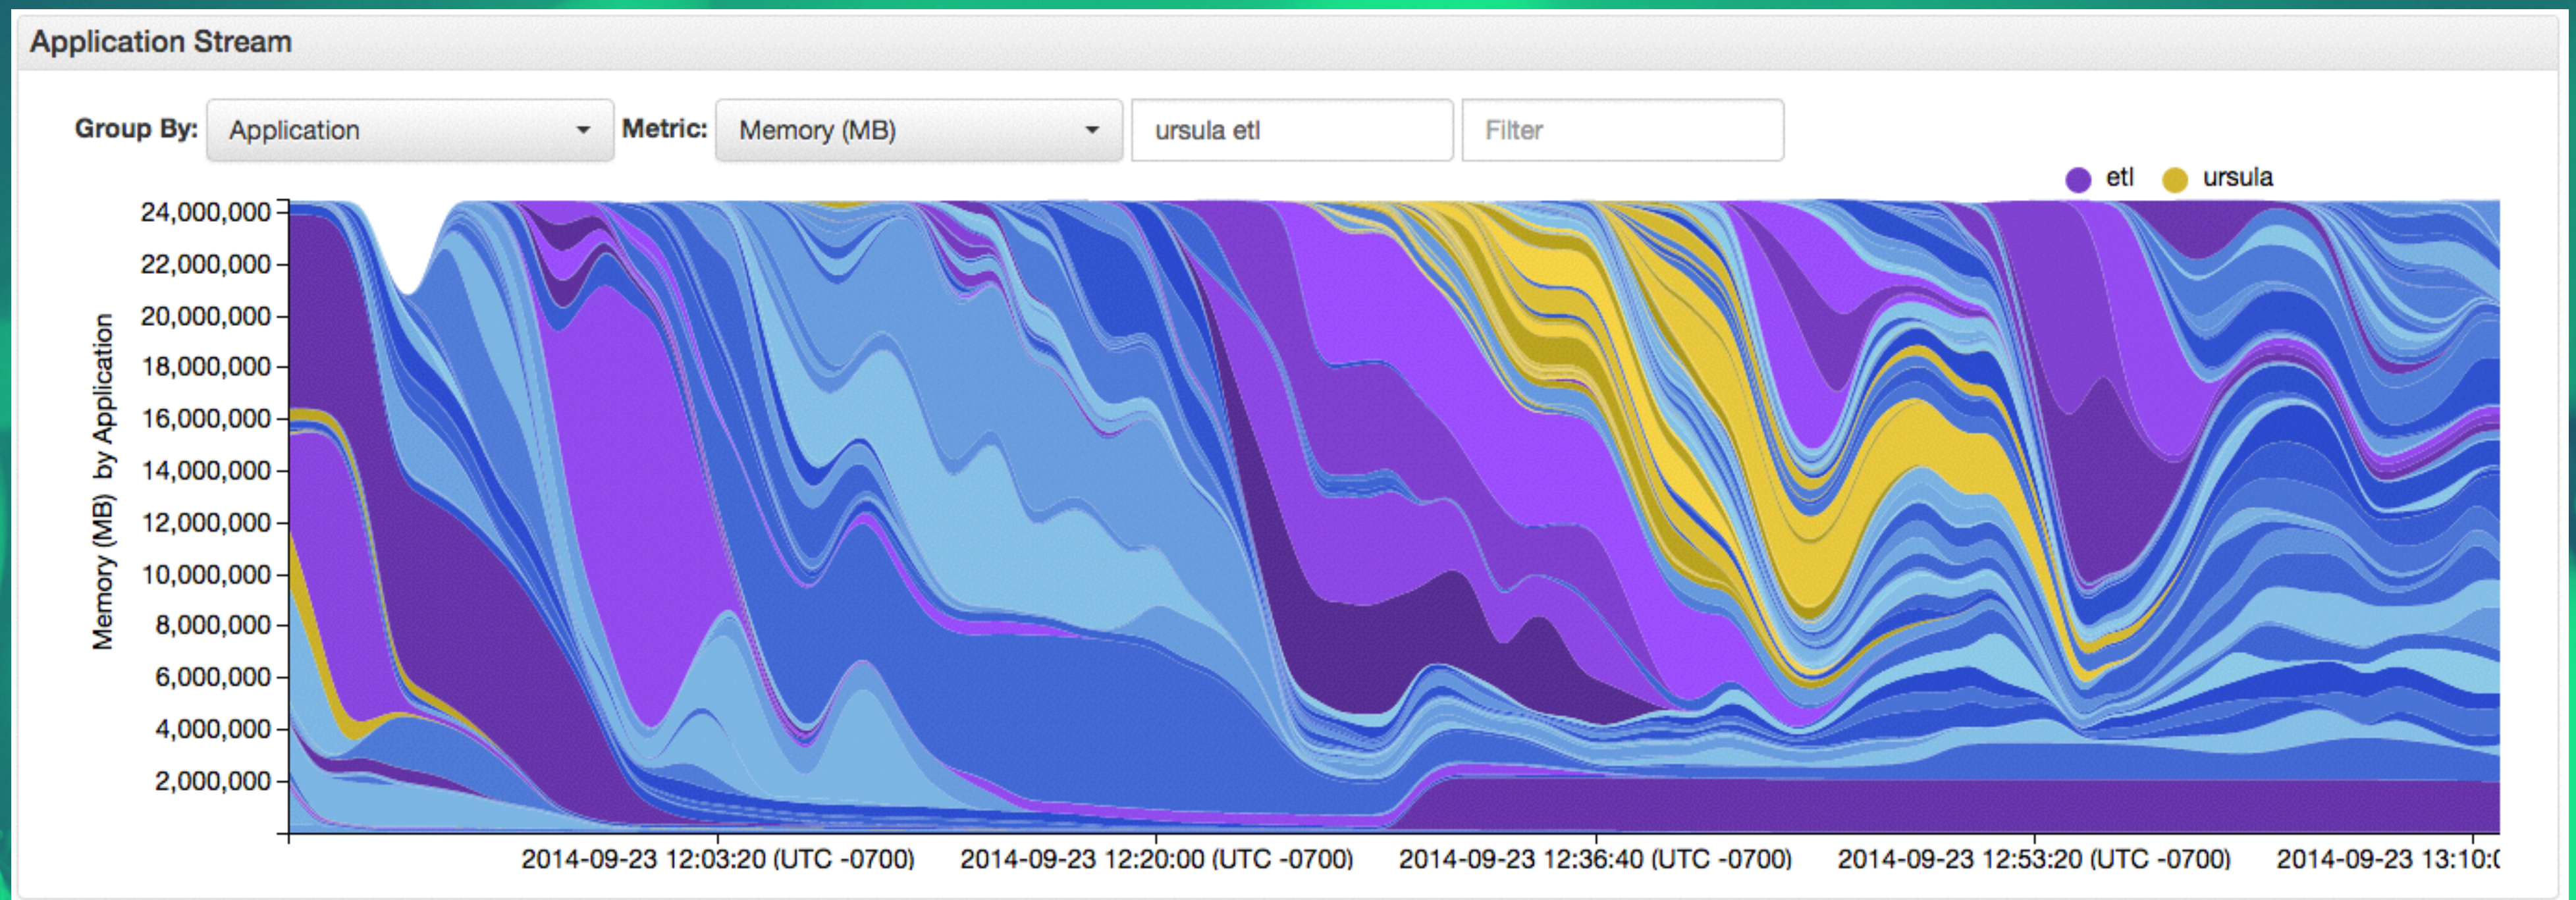


Scaling Machines: Easier
Scaling People: Harder

User Feedback: Automate it!



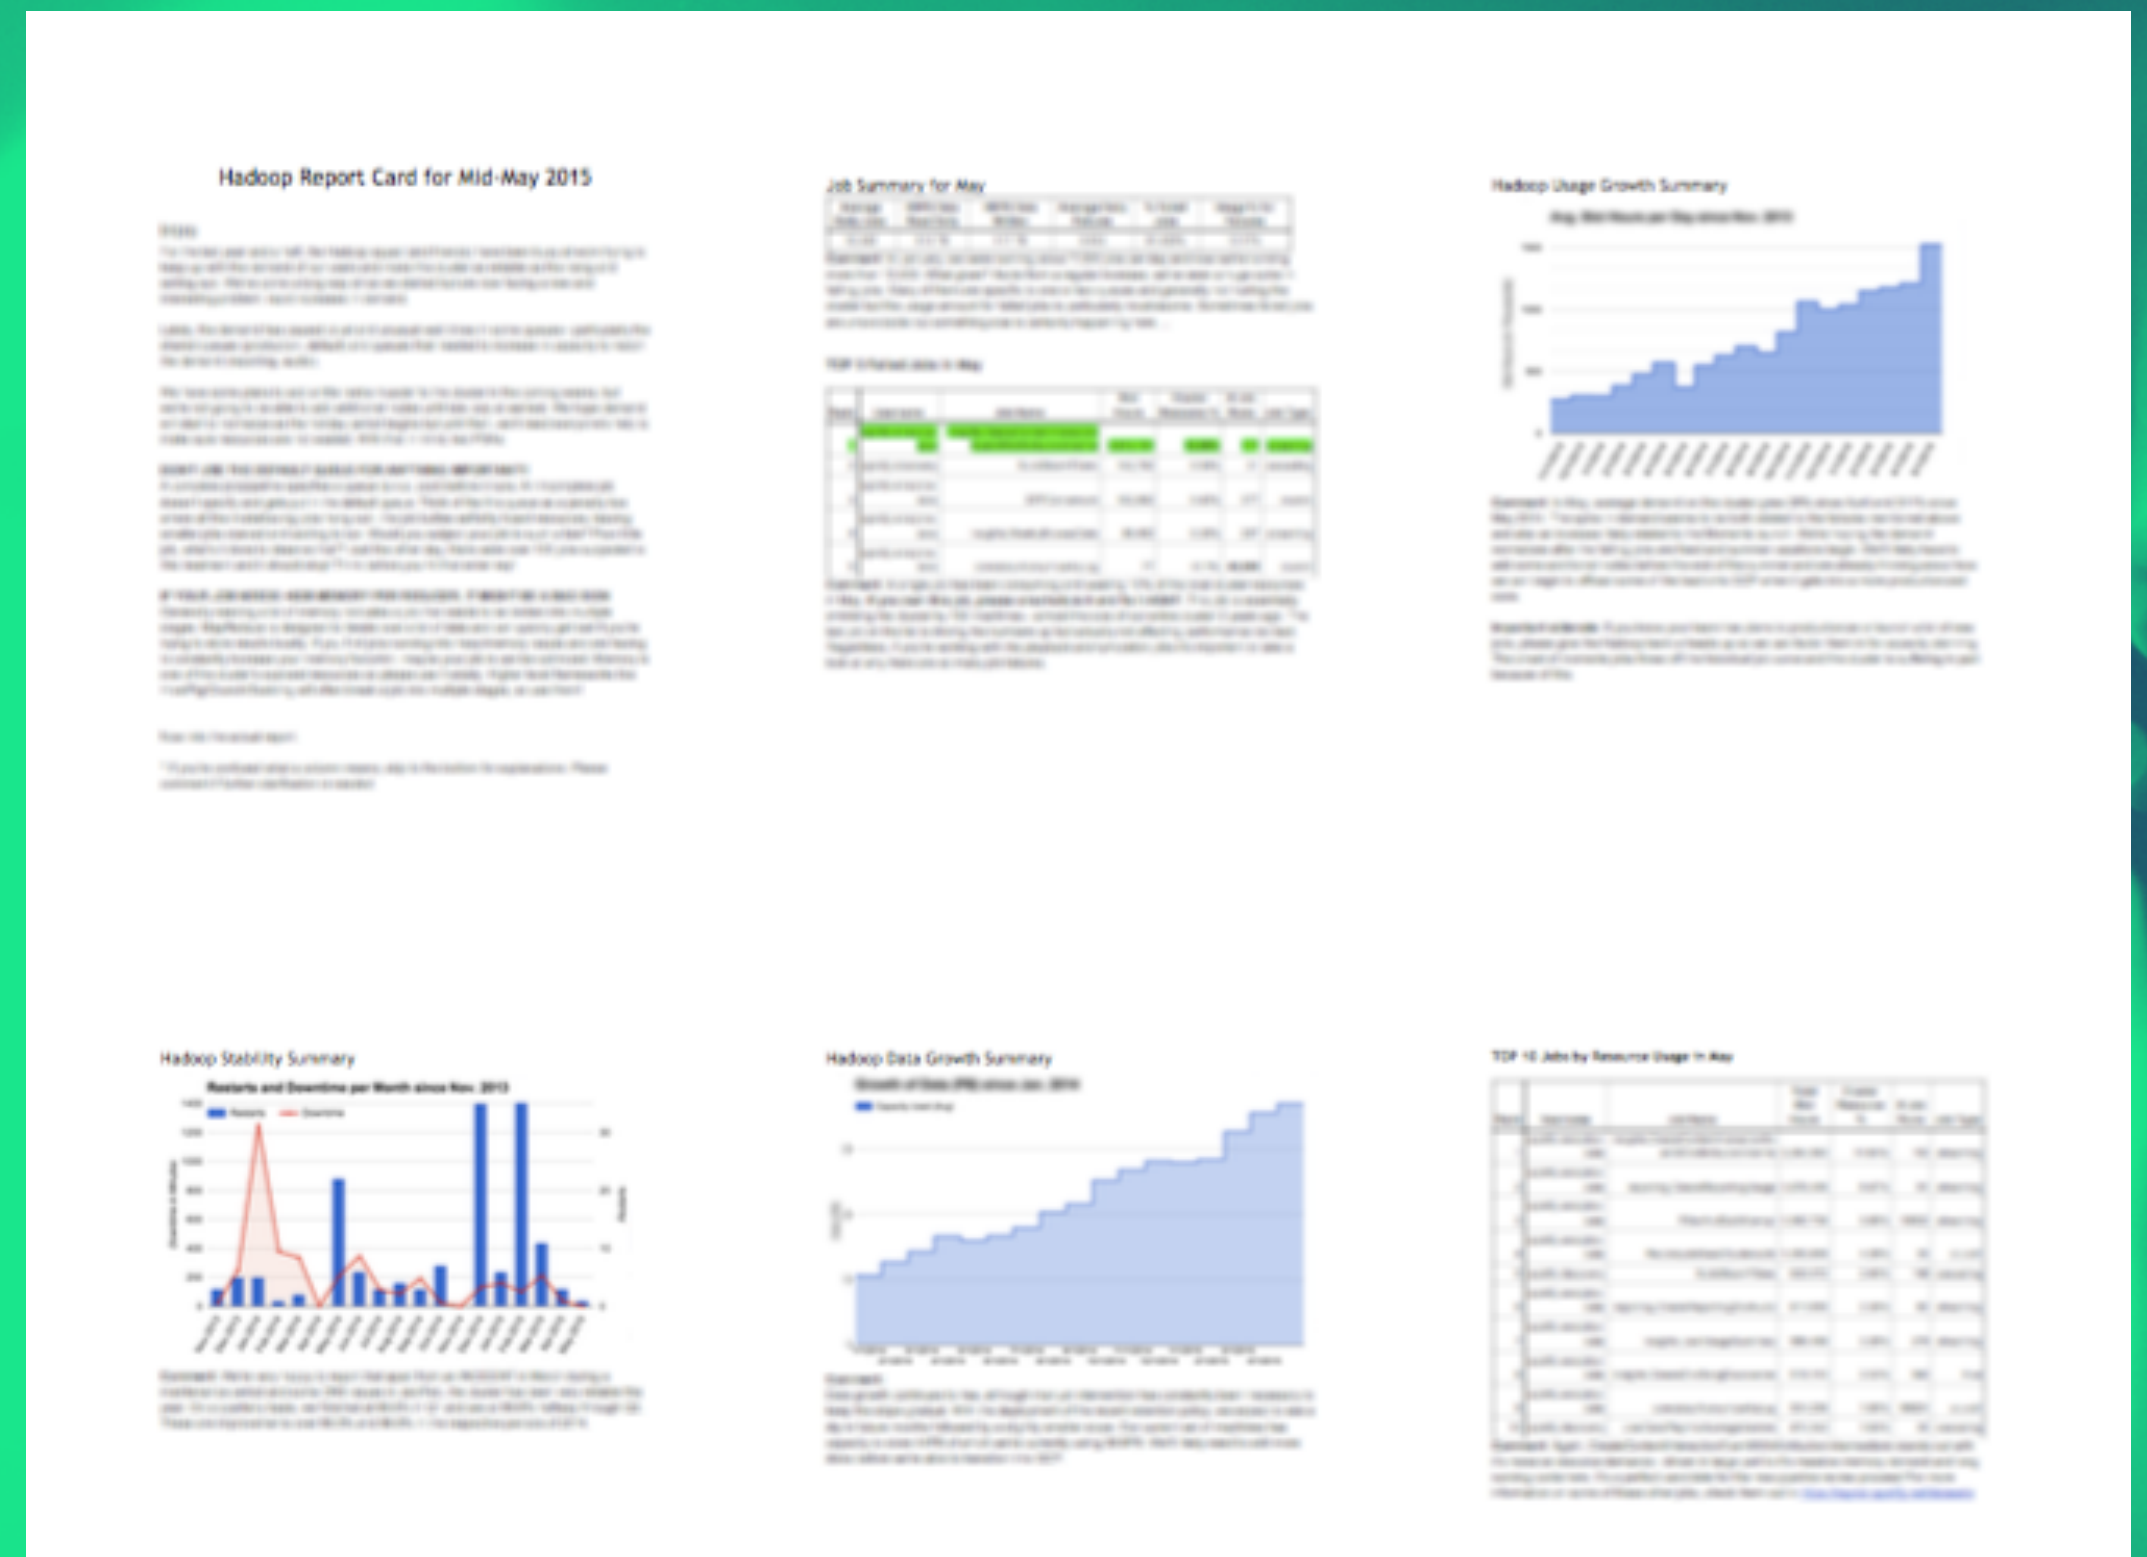
Inviso



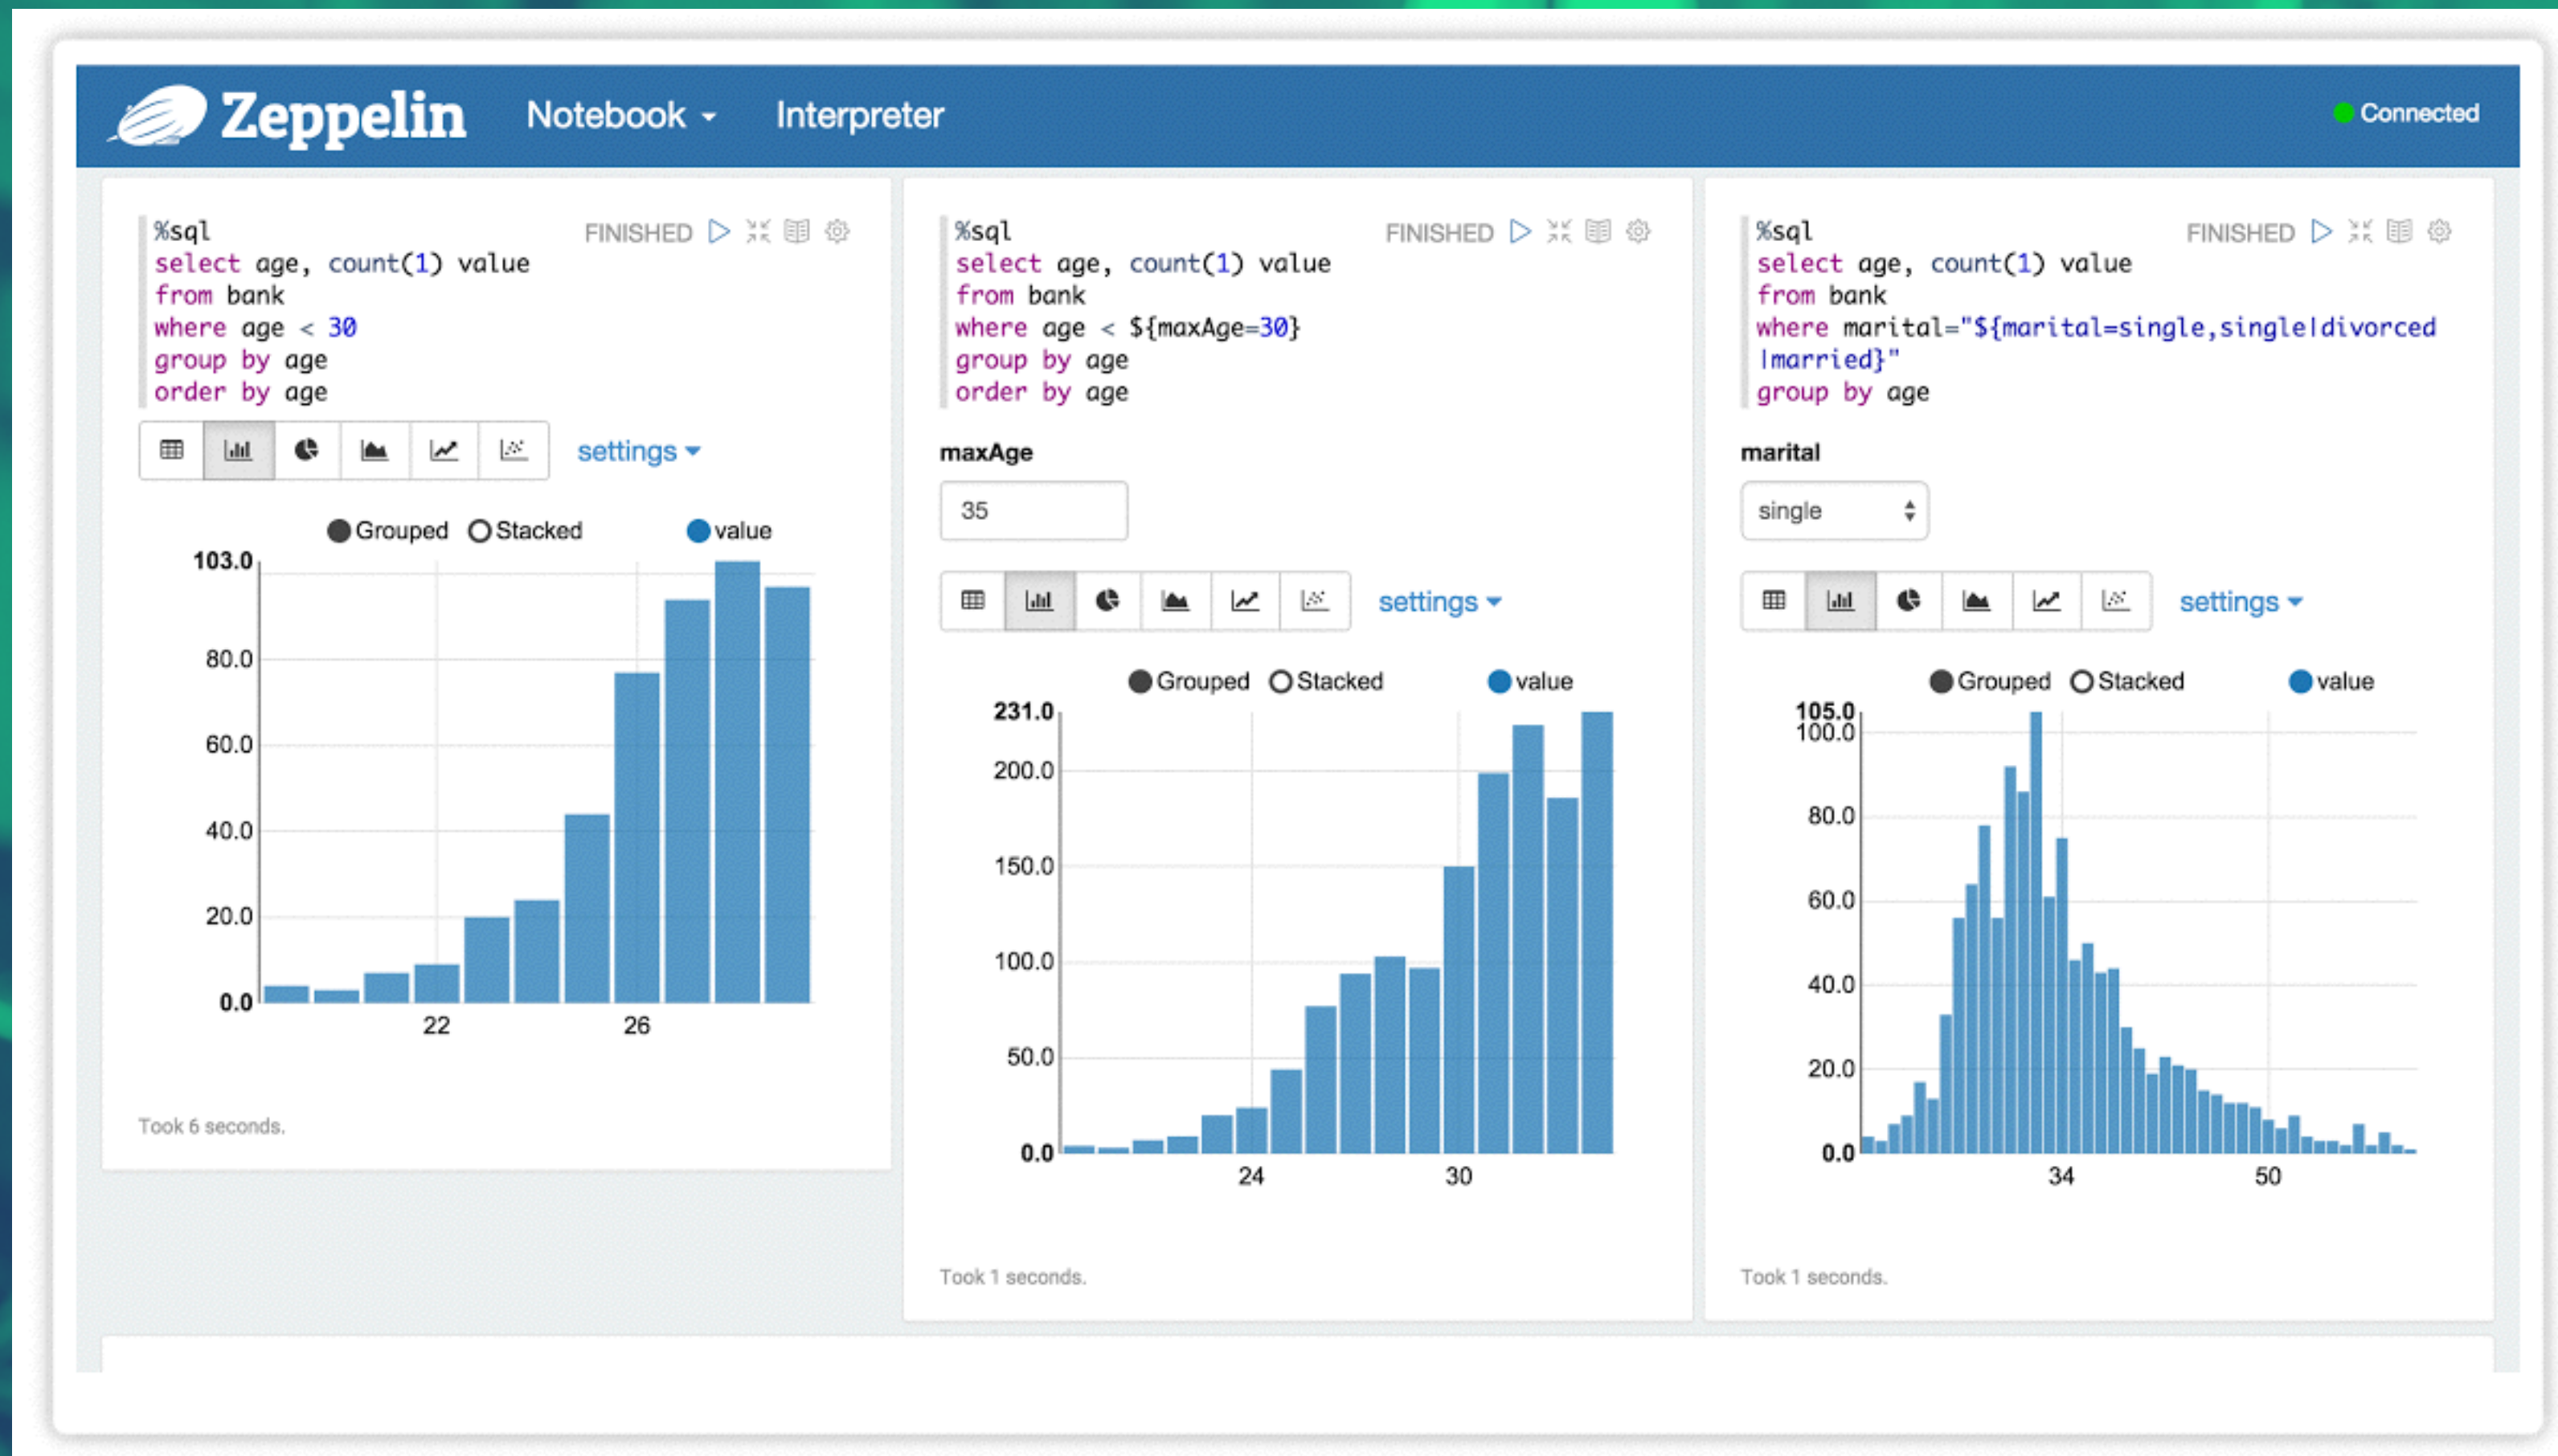
Developed by Netflix: <https://github.com/Netflix/inviso>

Hadoop Report Card

- Contains Statistics
- Guidelines and Best Practices
- Sent Quarterly



Apache Spark with Zeppelin



Takeaways

- There's no golden path
 - No perfect solutions, only ones that work now!
- Big Data is constantly evolving
 - Don't be afraid to rebuild and replace!

A vibrant, high-contrast image of a modern house at night. The scene is illuminated by a bright yellow filter. In the foreground, a swimming pool is visible. To the right, a woman in a bikini is jumping or dancing near a large glass door. In the bottom left, a group of people are gathered, some with their arms raised in celebration. The background shows a modern house with large windows and a balcony. String lights are visible across the scene.

Join The Band!

Engineers needed in NYC,
Stockholm

spotify.com/jobs



Please

**Remember to
rate this session**

Thank you!



COPENHAGEN

INTERNATIONAL
SOFTWARE DEVELOPMENT
CONFERENCE 2015

goto;

conference

 Join the conversation #gotocph

Conference: October 5-6 // Workshops: October 7-8, 2015

