

Using Spark to Generate Analytics for  
International Cable TV Video Distribution  
Christopher Burdorf  
NBC Universal

# Problem Definition

- Gigabytes of metadata associated with international video distribution stored in Oracle and flat files needs to be processed to analyze usage
  - Store resulting data in Hbase/Hadoop – media Ids with usage over weeks, months, years, in total and on a per channel basis
  - Query Hbase tables, generate graphs and display them in a web page.
  - Generate report listing media files ordered by least recent usage – for offlining purposes
  - Generate histogram of days since last aired for media on storage cluster

# Spark usage

- Queries from Oracle database get paths to schedule files stored on Isilon
- Schedule files are read and media ID's, air times, etc are stored in Scala classes and then to RDDs
- Spark map/reduceByKey is then used to produce counts for each media ID usage per country location based on week/month/year
- Spark join used to compute list of online media files (data from Oracle) and their last usage date (from schedule files). The result is sorted in ascending order on the most recent broadcast timestamp field.

# HBase/WebApp

- Broadcast frequency count data written from Spark App to CSV format then bulk-loaded into HBase
- Java/Spring web app queries HBase generates and displays graphs to web pages on demand.
- Spark App executed and HBase updated nightly

# Production Environment

- Linux
- Oracle
- Clustered storage system
- Spark/Scala
- Java/Spring Web App

# Mesos Cluster configuration

Master / Framework 201404041522-2350266378-5050-12435-0002

**Name:** SQL MapReduce  
**User:** cburdorf  
**Registered:** 4 minutes ago  
**Re-registered:** -  
**Active tasks:** 4  
**CPUs:** 18  
**Mem:** 6 GB

### Active Tasks

ID ▾	Name	State	Host
3	Task 3	RUNNING	Sandbox
2	Task 2	RUNNING	Sandbox
1	Task 1	RUNNING	Sandbox
0	Task 0	RUNNING	Sandbox

- 18 cores, 23GB ram

# Partitions from Spark app

Added rdd\_0\_3 in memory on server1:33229 (size: 42.9 KB, free: 883.2 MB)  
Added rdd\_0\_7 in memory on server1:33229 (size: 42.8 KB, free: 883.1 MB)  
Added rdd\_0\_16 in memory on server2:45795 (size: 43.0 KB, free: 883.2 MB)  
Added rdd\_0\_15 in memory on server2:45795 (size: 42.9 KB, free: 883.1 MB)  
Added rdd\_0\_14 in memory on server2:45795 (size: 42.8 KB, free: 883.1 MB)  
Added rdd\_0\_2 in memory on server2:45795 (size: 43.0 KB, free: 883.0 MB)  
Added rdd\_0\_13 in memory on server2:45795 (size: 43.2 KB, free: 883.0 MB)  
Added rdd\_0\_10 in memory on server2:45795 (size: 42.9 KB, free: 882.9 MB)  
Added rdd\_0\_17 in memory on server2:45795 (size: 43.1 KB, free: 882.9 MB)  
Added rdd\_0\_12 in memory on server2:45795 (size: 42.8 KB, free: 882.9 MB)  
Added rdd\_0\_8 in memory on server2:45795 (size: 43.1 KB, free: 882.8 MB)  
Added rdd\_0\_6 in memory on server2:45795 (size: 43.0 KB, free: 882.8 MB)  
Added rdd\_0\_11 in memory on server2:45795 (size: 42.9 KB, free: 882.7 MB)  
Added rdd\_0\_9 in memory on server2:45795 (size: 42.8 KB, free: 882.7 MB)  
Added rdd\_0\_1 in memory on server4:43454 (size: 42.8 KB, free: 883.2 MB)  
Added rdd\_0\_5 in memory on server4:43454 (size: 42.9 KB, free: 883.1 MB)  
Added rdd\_0\_0 in memory on server3:41154 (size: 43.0 KB, free: 883.2 MB)  
Added rdd\_0\_4 in memory on server3:41154 (size: 42.9 KB, free: 883.1 MB)

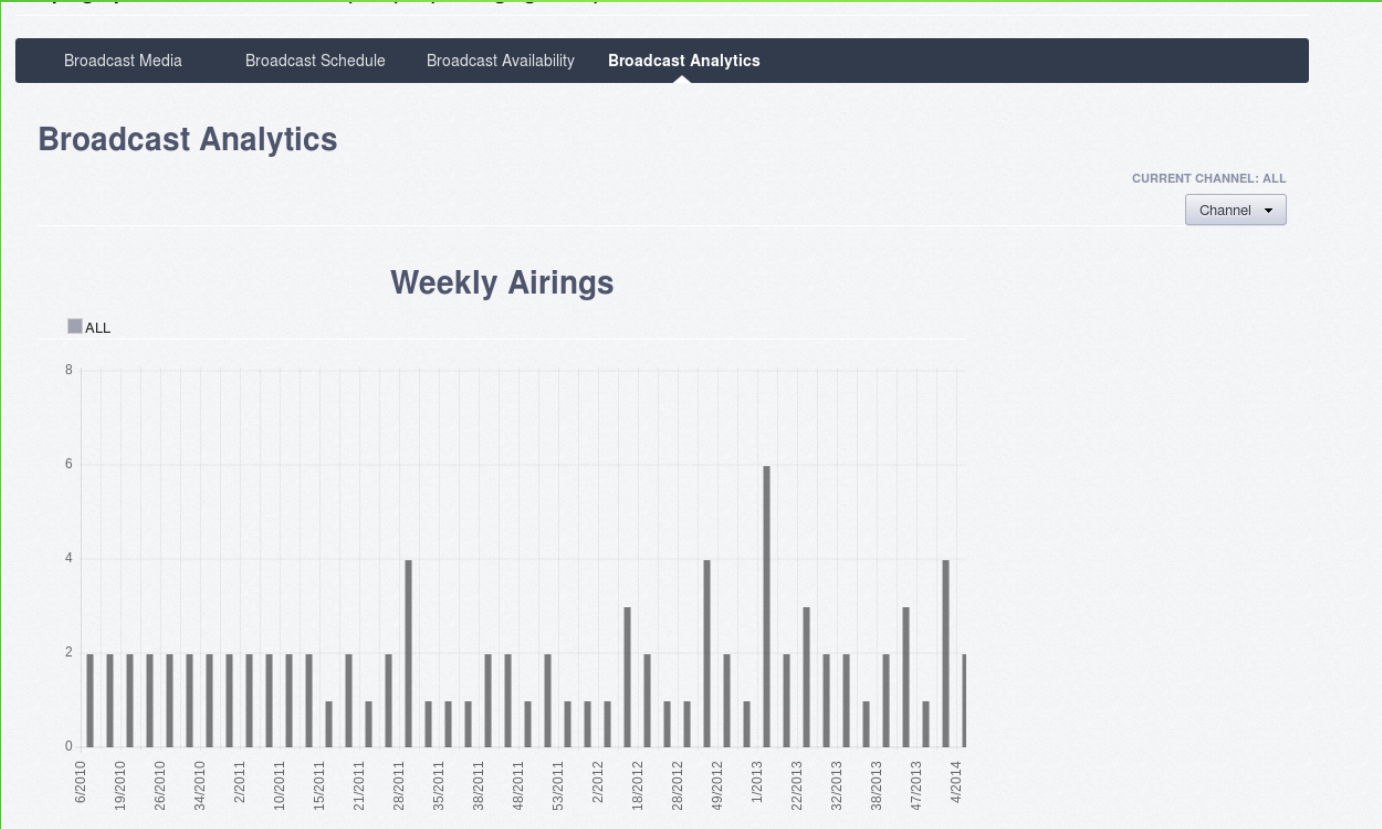
# Shuffle Spill

## Aggregated Metrics by Executor

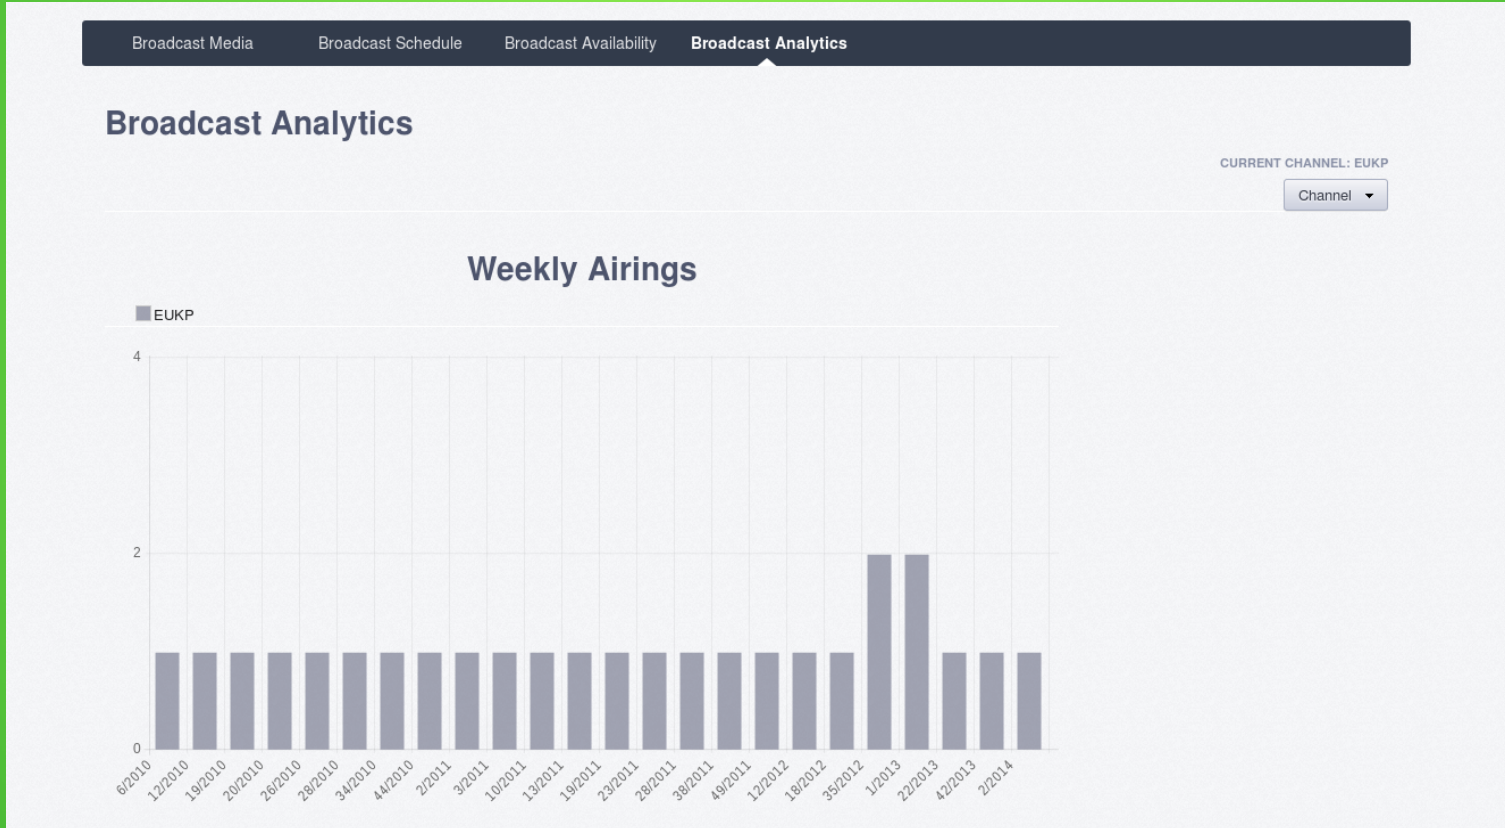
Executor ID ▼	Address	Task Time	Total Tasks	Failed Tasks	Succeeded Tasks	Shuffle Read	Shuffle Write	Shuffle Spill (Memory)	Shuffle Spill (Disk)
201404041414-2350266378-5050-10960-3		3.6 s	2	0	2	5.4 MB	4.8 MB	0.0 B	0.0 B
201404041414-2350266378-5050-10960-7		2.4 s	2	0	2	5.4 MB	4.8 MB	0.0 B	0.0 B
201404041414-2350266378-5050-10960-9		2.6 s	2	0	2	5.4 MB	4.8 MB	0.0 B	0.0 B
201404041522-2350266378-5050-12435-0		27.3 s	12	0	12	12.2 MB	28.9 MB	31.8 MB	2.4 MB



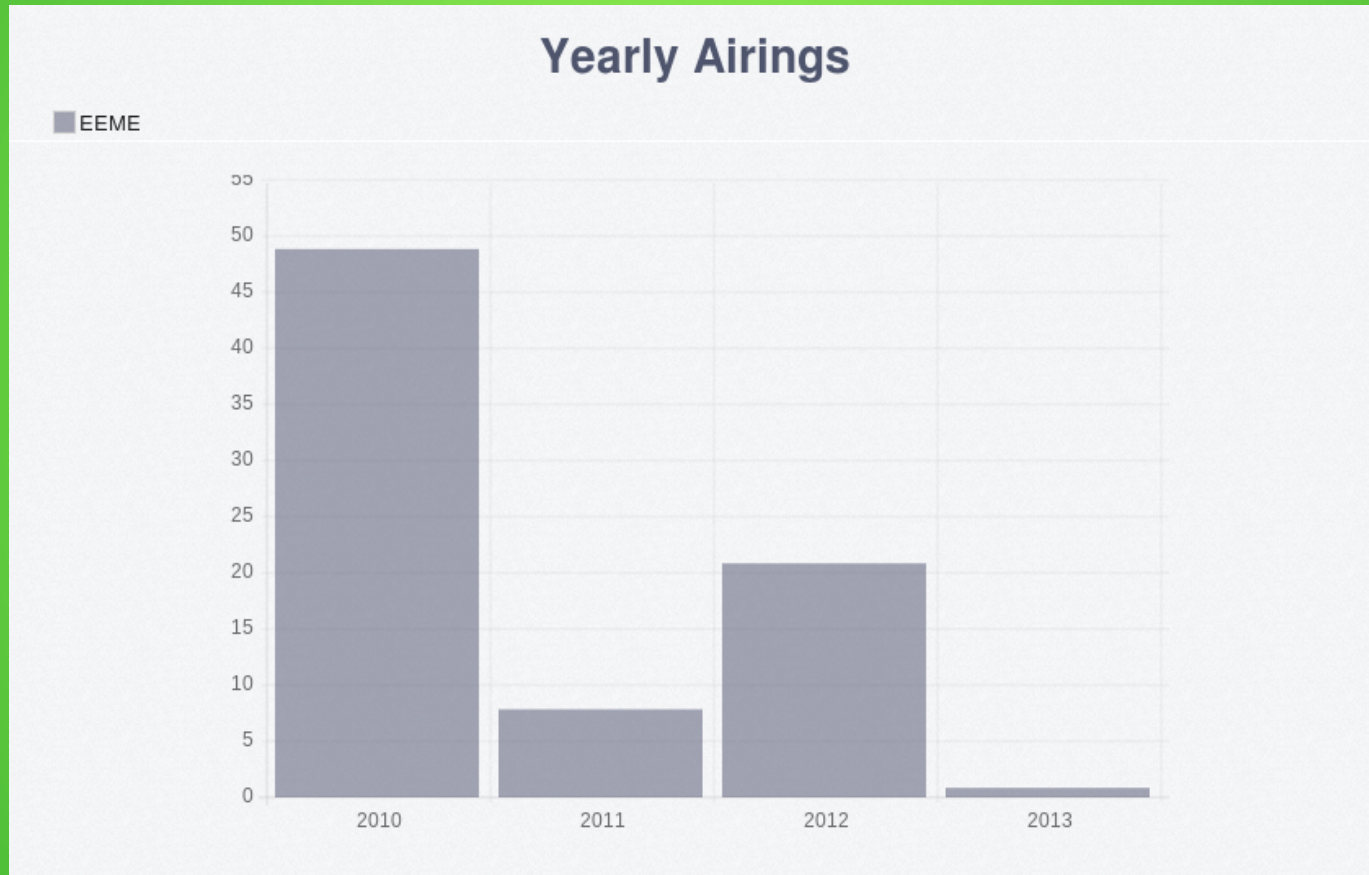
# Results



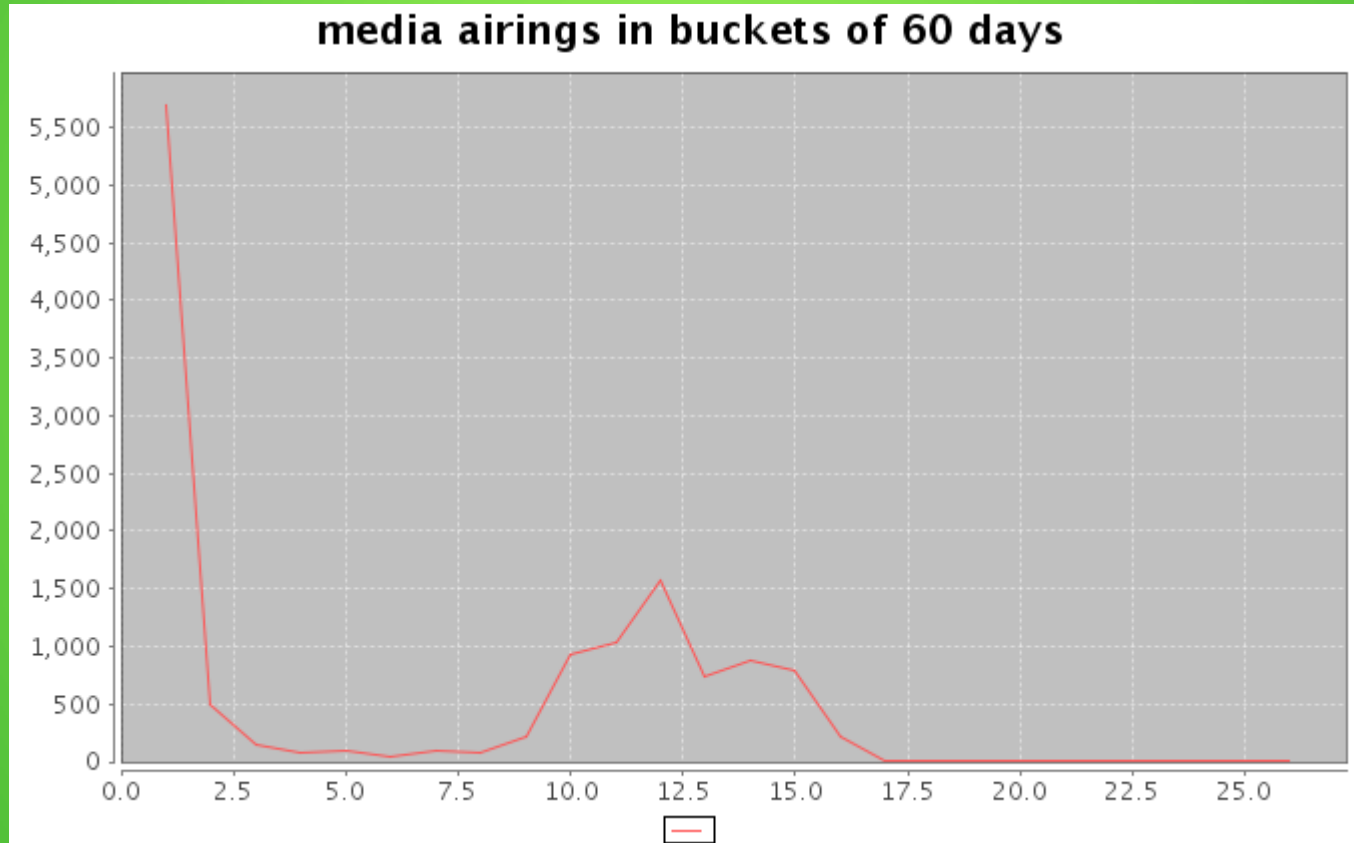
# Per channel



# Per channel



# Storage cluster utilization days since last aired per media file



# Conclusions

- Performance improvements
  - Small data: 10x faster on 18 core vs 1 core even though a good portion of the app is sequential file and DB I/O.
  - Large data: not even possible on 1 box in cluster – runs out of memory.
- Spark/Scala super fun programming env – thanks to all developers
- Mesos cluster management wonderful ease of use
- Minor issue: Mesos Mac OSX build
- Future work