



Use Case: Product Catalogs & Playlists

A product catalog is an organized collection of products, services, inventory or other assets that is constantly modified, searched and analyzed by both internal and external user communities. Playlists refer to user-defined queues of items like songs, movies, games, gift registries and lessons.

Industry

Digital Media

Challenges

- Scale challenges created by a fast growing user base
- The need to ensure 100% up time of the application
- Small team managing a large system

Solution

- Apache Cassandra for linear scale
- Multi-data center replication and no single point of failure for 100% uptime
- Operational simplicity

Results

- More than doubled subscriber-base in two years
- Primary Cassandra cluster consists of 32 nodes split between two data centers
- Stores several billion rows with approximately 1TB of unreplicated data per data center

PRODUCT CATALOGS + PLAYLISTS

HULU RELIES ON CASSANDRA TO STREAM VIDEO TO MORE THAN 6 MILLION SUBSCRIBERS WITHOUT INTERRUPTION

Hulu is an online video service that offers a selection of hit TV shows, clips, movies and more on the free, ad-supported Hulu.com service, and the subscription service Hulu Plus. Hulu has more than 6 million paid subscribers and is accessible on over 400 million internet-connected devices including connected TVs and Blu-ray players, gaming consoles, set-top boxes, mobile phones, and more.

The Challenge

The massive increase in the usage of multiple devices, particularly mobile, to provide real-time access to content has fueled the growth of video streaming services like Hulu. As the number of devices and subscribers to Hulu's service increases, so does the volume of user-generated data, which adds enormous stress on an application infrastructure not built to support this level of data load and velocity of scale.

Andres Rangel, a Senior Software Development Lead at Hulu explained, "Our previous system was having trouble expanding to support our fast-growing subscriber base. We needed something that could scale quickly and would be easy to maintain because we had a very small team."

Specifically, the problem Hulu faced was that they had trouble scaling their writes as millions of users interacted with the application. The application follows a user from device to device, saving the previous session and starting up a show where it was left off. This type of experience requires the ability to handle a high velocity of writes and Hulu's systems simply couldn't keep up with the growing database demands.

"Scaling to support this level of load adds a lot of complexity and requires resources to handle it," Rangel explained. "We had a small team, so we needed something that would scale, but wasn't overly complicated to manage."

The Solution

The Hulu team realized they needed a database solution that could support their scaling needs while ensuring that their service never goes down, critical in the highly-competitive media industry. In a three-way evaluation between Apache Cassandra, HBase and Riak, Hulu turned to Cassandra as the database of choice to ensure linear scale and 100% uptime, so they could deliver the most engaging customer experience possible.

"Our experience with Riak was that it was able to scale, but performance wasn't as good as Cassandra. The ability to handle range queries was also important and Riak couldn't do this at the time," said Rangel. Finally, with a small team tasked with immense responsibilities across massive scale, Rangel concluded that they could not work with a technology that required a steep learning curve. "Another problem was that nobody on the team was experienced with Erlang, the programming language native to Riak."

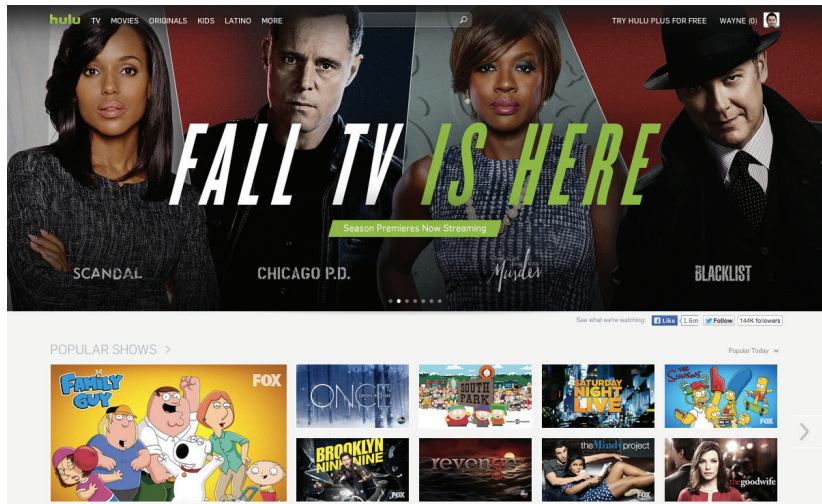
When Hulu evaluated HBase, they found that Hadoop instances took a lot of work to set up. Due to its master-slave architecture, Hbase, which runs on top of the Hadoop Distributed File System (HDFS), has a single point of failure. In fact, in their tests, Hulu noted concerns as they saw cascading failures take down all region servers.

“With Cassandra, it managed to handle the load, it’s very reliable, it allows range queries without limitations, and it’s easy to maintain. It’s night and day compared to HBase.”

- Andres Rangel
Senior Software Development Lead
Hulu

With Cassandra, Hulu found a match in its ability to manage large workloads and ensure 100% uptime through its masterless architecture that eliminates a single point of failure. “With Cassandra, it managed to handle the load, it’s very reliable, it allows range queries without limitations, and it’s easy to maintain,” said Rangel. “It’s night and day compared to HBase.” Overall performance also improved dramatically since Cassandra is optimized for SSDs. Rangel also said that Cassandra was better at replication, which was critical to ensure availability across multiple data centers.

Today, Hulu depends on Cassandra’s distributed database technology to store and provide real-time access to subscriber watch history. Using Cassandra, Hulu users can pause one episode on their TV and then resume viewing the same show at the same spot on any other device, such as a smartphone, when they leave the house.



Cassandra allows Hulu to scale across multi-data centers without fear of downtime, ever.

The Results

Hulu’s primary Cassandra cluster consists of 32 nodes split between two data centers, one on the east coast and the other on the west coast. It’s watch history keyspace contains several billion rows with approximately 1TB of unreplicated data per data center. When a Hulu subscriber watches a video, uses Hulu for recommendations, or saves sessions across various devices, these actions are handled by Cassandra.

“We have been able to more than double our subscribers in two years and it has had no trouble handling the load,” he explains. “We only need one engineer to maintain it and it takes less than 5% of his weekly time.”

After finding immense success in current production environments, Hulu will deploy Cassandra to support four other services. These include handling social data from users, messaging, allowing user-controlled viewing on smart TVs, and gaming consoles via smart phones.