

Disney+ Hotstar Powers Explosive Growth with Real-Time Event Streaming on Confluent Platform



Headquarters
Mumbai, India

Industry
Tech

Challenge
Support rapid growth while improving the experience of 300 million customers on the largest OTT platform in India.

Solution
Use Confluent Platform and Apache Kafka to build an enterprise-scale event streaming platform that powers real-time analytics, increases customer engagement, and enhances customer experiences.

Confluent Products

- Schema Registry
- Confluent Control Center
- Confluent Auto Data Balancer

Results

- Time to market shortened
- Ability to scale with Confluent Auto Data Balancer
- Large-scale deployment and high availability operations supported
- Data compliance concerns mitigated

In 2018, 3.1 million viewers watched the Super Bowl via live stream in what was touted as the most live-streamed Super Bowl ever. That same year, however, more than three times as many viewers watched the Indian Premier League (IPL) cricket finals on the Disney+ Hotstar over-the-top (OTT) streaming service. And just one year later, 25.3 million Disney+ Hotstar customers used the service to watch, and engage with, the ICC Cricket World Cup.

Disney+ Hotstar is the largest OTT platform in India, with 300 million active users. The recent explosive growth of Disney+ Hotstar, which includes quadrupling its user base in a little over a year, has been driven by several factors. Increased smartphone penetration in the market and access to inexpensive data has helped. A larger factor is the platform's content offering which includes not only live sports and events, but also original series as well as HBO and Disney content. Also among the biggest growth drivers are innovative features that let Disney+ Hotstar customers actively engage with the content and with each other in real time. During a live IPL telecast, for example, Disney+ Hotstar introduced a real-time cricket prediction game in which over 33 million unique users answered more than two billion questions—about 100K requests per second.

From a technology perspective, the company's growth has been underpinned by real-time event streaming with Confluent Platform and Apache Kafka®. On an average day, 10 to 15 billion messages flow over Confluent Platform and Kafka, with peaks of up to five times higher. This event-driven architecture serves as a foundation for real-time analytics, increased customer engagement, enhanced customer experiences, and rapid, enterprise-wide growth.

At Disney+ Hotstar, event streaming with Confluent Platform has democratized data, enabling teams to share data and make it easily

Confluent has been instrumental in the company's ability to scale. It's not unusual to have 25 million concurrent users for a live match, and on a daily basis seven to eight terabytes of data come into the platform. According to Disney+ Hotstar, powering so many use cases at this scale was made possible because of the real-time event streaming infrastructure that was built on Confluent Platform and Kafka.

accessible to others for any number of real-time use cases. This has greatly simplified integration and accelerated innovation—teams no longer worry about how their systems will connect because everyone either reads from or writes to Kafka for interservice communication. Confluent has been instrumental in the company's ability to scale. It's not unusual to have 25 million concurrent users for a live match, and on a daily basis seven to eight terabytes of data come into the platform. According to Disney+ Hotstar, powering so many use cases at this scale was only possible because of the real-time event streaming infrastructure that was built on Confluent Platform and Kafka.

Disney+ Hotstar leveraged Confluent support and services during the initial phases of adoption and its use of event streaming has rapidly increased. Confluent engineers helped Disney+ Hotstar get up and running quickly, identified and resolved potential operational challenges, and provided guidance during design reviews as the company built its own in-house event streaming expertise. Going forward, Disney+ Hotstar plans to make use of ksqldb for real-time stream processing in compliance and security applications, and employ cluster federation as the company continues its expansion in North America, Europe, and worldwide.

Business Results

- **Time to market shortened.** For Disney+ Hotstar, one of the biggest advantages of event streaming with Confluent Platform is time to market. Disney+ Hotstar stated that because the data is now democratized, teams throughout the company can pick it up and immediately work with it. That makes teams more agile and significantly reduces delivery times, because developers are not wasting time figuring out API contracts or setting up new infrastructure.
- **Ability to scale with Confluent Auto Data Balancer.** When there's a cricket match or other live event, the platform can easily have 20 times the number of users on it than the day before. With Confluent ADB, Disney+ Hotstar engineers don't have to figure out how to scale to handle large increases in traffic; data coming into the platform is balanced automatically to level the workload across the cluster.
- **Large-scale deployment and high availability operations supported.** Disney+ Hotstar emphasized the importance of Confluent support in realizing the company's event streaming vision. Confluent engineers helped with the initial steps, and have continued to partner with the Disney+ Hotstar team build their own knowledge and gain the

expertise to be more self-reliant. In addition, Confluent engineers have helped to eliminate downtime, maximize availability, and ensure Disney+ Hotstar runs at an optimal scale, with real-time speed and responsiveness.

- **Data compliance concerns mitigated.** Today, Disney+ Hotstar no longer relies on a third party for analytics; all customer data remains in-house. As the company launches in new countries, keeping all of that data in its own ecosystem with Confluent Platform and Kafka is critical for compliance with data protection regulations. The company is currently exploring additional Confluent Platform security capabilities, including Role-Based Access Control, Structured Audit Logs, and more.

Technical Solution

At Disney+ Hotstar, the vision for creating a single platform for data exchange became a reality a few years ago with the launch of Knol, a real-time streaming platform built on Apache Kafka. Knol served as a proof of concept for Kafka, albeit a large and critical one as the platform was soon handling multiple terabytes of data each day, including ad impressions, clickstream data, customer support data, social data, and more.

Building on the success of this initial effort, Disney+ Hotstar engineers began moving the company's analytics off of a third-party provider's service to an in-house data ingestion platform. The Disney+ Hotstar team worked with Confluent engineers together on the design and technical reviews of the architecture and its implementation.

The new platform, built on Confluent Platform running on AWS, soon became the basis for interservice communication, logging, and numerous projects aimed at improving the customer experience. Disney+ Hotstar noted that the personalization it offers subscribers would be impossible without first bringing analytics in-house with Confluent Platform and Kafka. Accessing data in near-real time with low latencies enables the company to power new and better customer experiences. When external partners provide content to Disney+ Hotstar, the data is ingested via a secure API and then flows over to Kafka, where it becomes available along with other data on the platform for post-processing, analysis, and dashboard visualization.

Disney+ Hotstar teams are taking full advantage of Schema Registry to help ensure compatibility between producers and consumers. They've also used Confluent Control Center to manage and monitor their Kafka ecosystem and serve as a quality assurance tool to see how their data is flowing through Kafka.

In addition to relying on Confluent Auto Data Balancer for rebalancing partitions when brokers are scaled, Disney+ Hotstar engineering teams are heavy users of the Kafka Connect Amazon S3 sink connector, which they employ to export data from Kafka topics to S3 objects.

The personalization Disney+ Hotstar offers subscribers would be impossible without first bringing analytics in-house with Confluent Platform and Kafka. Accessing data in near-real time with low latencies enables the company to power new and better customer experiences.

The widespread, cross-industry use of Kafka has further added to its value at Disney+ Hotstar, particularly in hiring new engineers, many of whom are already familiar with Kafka and want to continue working with the latest technology.

Learn More About Disney+ Hotstar

hotstar.com