

DEMONWARE

Demonware Modernizes Real-time Data Pipelines to Serve Millions of Gamers Around the World

demonware



Confluent
Open Source



Kafka
Connect

Headquarters

Dublin, Ireland
Vancouver, Canada
Shanghai, China

Industry

Online Gaming

Challenge

Needed a solution that could handle high volumes of event data, and could provide a platform for real-time data access

Solution

Deployed Confluent Platform and Apache Kafka® to handle data at scale as well as to capture and process the various types of data coming in every second from user activity

Results

- Increased reliability
- Accurate, real-time data
- Ability to process data at scale
- Faster ramp time with access to Confluent team

Demonware, a fully owned subsidiary of Activision Blizzard, has designed, developed and hosted online services powering some of the most popular game franchises in the world, including Call of Duty and Skylanders. These services touch many aspects of online gameplay. Alongside the core offerings of match-making, the services include transactions, leaderboards, data pipelines and identity management, among others for millions of users around the world.

Demonware provides the infrastructure, both on premises and in the cloud, to scale the games online. Imagine a popular new game launches at midnight, and usage spikes. Demonware manages the core infrastructure to ensure gamers don't experience downtime and get to enjoy the games they love.

Demonware deployed Confluent Platform and Apache Kafka® to scale their data pipeline infrastructure depending on usage as they capture and process the various types of data coming in every second from user activity. The data is then sent to various teams within Demonware and Activision who use it to power the various services including diagnostics and optimizing and improving player experiences, which enables designers to adjust and make improvements in the games.

"We work closely with the studios to collaborate from the early stages of development. From there, we create our own internal tools and tech to be efficient, responsive and reusable. We enjoy taking on real challenges that serve millions of players every day," said Tom Szymanski, team lead, data pipelines at Demonware.

With our Kafka-backed data pipeline, we are able to support our partners, who every year create more services, more features, more data instrumentation, and even more granular data than the year before.

—
Tom Szymanski, Team Lead, Data Pipelines at Demonware

CHALLENGES

Demonware needs to deliver smooth game launches and understand when usage is going to spike. Prior to deploying Kafka, Demonware was using a database as a data pipeline, but as the number of players increased from year to year they began running into high IO problems; different applications were all reading and writing to the same few tables and the databases could not keep up with the demand as the data volumes grew. This solution was breaking at the seams.

Demonware started looking for a solution that could handle a high volume of data at scale, and could provide a fundamental platform for real-time data access. They needed a solution to power a variety of services, including answering diagnostics-type questions: “Is this service running? Is this service healthy? Is the server okay? Is it overloaded? Is it not busy for some reason?” along with business-focused questions about the millions of people playing the games: “What kind of experiences are they having? How did they do on the last match? How many kills do they have?” etc.

“The problem was taking all these different kinds of data and figuring out how to collect it from many, many different sources, including millions of game clients playing at the same time, and funneling it through one place quickly and reliably and then delivering it to the various applications that it needs to go to,” said Szymanski.

SOLUTION

Ultimately, Demonware chose Kafka because it is the most mature streaming solution on the market, and the team believes in the power of open source technology. Kafka enables Demonware to scale, and their peak so far has been about one-and-a-half terabytes a day of data with about 3 billion messages per day.

In addition, with a real-time data pipeline, the analysts are able to produce reports and do in-depth analysis on player experience and system diagnostics. This data is then sent back to the designers to improve the game experience and allows them to validate or invalidate their hunches on weapons, level layouts, game modes, etc., ensuring players are having the intended experience when the game is released. The data also powers Analytic Services, an internal services team at Activision Publishing, to infer insights from the raw data to share with game developers.

Everywhere the process of designing and hosting games is becoming a lot more data-driven than it has been in the past. There’s millions of different kinds of instrumentation points that they collect. And that number just keeps growing. “With Kafka, there’s always more services. There’s always more features. There’s always more instrumentation. You can always look at existing features one level deeper than you did the year before,” said Szymanski.

RESULTS

By deploying Apache Kafka and Confluent Platform, Demonware has experienced significant positive results, including:

- **Increased reliability:** “Even if you start messing around with Kafka and the deployments, it still works great. If we’ve had headaches with it, it’s been user created. Reliability has been a big selling point for us,” said Szymanski.
- **Accurate, real-time data:** “Data is one of those things that people take for granted, unless it’s not there. Because teams at Demonware have experienced reliability with Confluent the last few years, the teams relying on us haven’t had reason to question it. The number of use cases and the different types of data are growing, the volume of that data is growing, and the expectations for access to it are growing. Our internal data partners depend on data access to be effective at their jobs,” said Szymanski.
- **Ability to process at scale:** Demonware’s data pipeline processes about one-and-a-half terabytes of event data a day with about 3 billion messages per day. Because Demonware can now process more events, they can track and measure more things at a more detailed level than before. The quality of the raw data is also much better, allowing for more accurate and informative analysis. “We’re most proud of the fact of just how much data goes through us. The volume of data, the volume of different types of data, the volume of different sources, are significant when you contrast that with the actual size of the team, which is actually quite small. So there’s a lot of accomplishment for a very small group of people.”
- **Faster ramp time with access to Confluent team:** “We’re a small team, and we’ve grown quite a bit. We’ve had to train up a few people from the ground up on the fly, and being able to rely on support people when we have the sort of questions that we can’t answer ourselves has actually helped expedite that process as well.”