



Redesigning the largest data lake in the world for an emerging generation

Next-level architectural design to handle customer intelligence for over a billion users and predict their behaviour.

Redesigning the largest data lake in the world for an emerging generation

Jio is the largest mobile network operator in India, and the third-largest in the world. They operate predominantly in Near East and South East Asia, providing a wide range of services to over 600 million users.



BACKGROUND

Jio had observed how the technology industry had struggled to innovate in this geographical and economic area, and so they developed services to fill a range of gaps identified. These services include platforms for streaming TV and films, listening to music, reading the news, video conferencing, cloud storage, communication, health management, and many more.

Jio had plans to double their user base to 1.1 billion, and simultaneously introduce smart customer analytics to their services, to help tailor user experiences. They had developed a number of use cases describing how these customer analytics would benefit the user experience and give them certain critical insight that they were lacking. They called this customer experience goal "Happy Next" — the ambition to predict customer behaviour to the extent that each next user action would be foreseen and assisted by the application, for a more enjoyable interactive experience.

A good example to illustrate this concept is a user who always watches short sitcoms on the train to and from work each day, and watches movies and sci-fi shows at home, but never mixes these types of content between environments. This would be a classic case of where the application should have the intelligence to identify this pattern and promote relevant options to the user in each case, based on location, time, and other relevant factors. Context-based predictions boost customer engagement and enjoyment, plus the data gathered is valuable as a marketing and business prediction tool.

Global heavyweight digital design agency AKQA was involved in producing the cosmetic designs for the user interface of the next-gen applications, but as they are a design-focused agency, the technical architecture required by the project was beyond the scope of what they were used to delivering. They therefore asked us for assistance, to use our expertise to guide the direction of work, advise as to the options and make recommendation, and validate technical assumptions and decisions.

Jio state that they have the largest data lake in the world. With so much content and so many areas of consumer engagement, this is entirely believable. Jio use Microsoft Azure as their primary provider, and they have dedicated Azure private cloud facilities in Microsoft's Indian data centres.

CHALLENGES

Somewhat unusually, we did not do a deep-dive into the existing solution during initial discovery, as Jio did not want us to be influenced by the direction they had already taken when proposing our ideas for their next-generation services. This provided Jio with a comparative approach, enabling them to validate their existing methodologies where they were aligned with our proposals, and assess the potential for improvement where they differed. It did however increase the challenge on our side, requiring us to establish a higher level of knowledge of the intricacies of their business than would usually have been necessary.

Although Jio's infrastructure was handling 600 million users — which is impressive and shows how well it has been built — they were concerned that it would not scale to handle twice that volume whilst simultaneously allowing the introduction of heavy realtime processing of customer experience data.

Critical to the solution design was the identification of what processes would be involved, at what stages, and in which ways. Being able to identify and predict bottlenecks based on processing demand and latency was key to being able to introduce computationally-heavy statistical analysis and machine-learning systems that would be needed to achieve the "Happy Next" customer experience that Jio was aiming for.

We needed to identify and advise on the DevOps approaches that would be needed to support the technical architecture for Jio's next-generation infrastructure.

These included pipelines for reliable software updates, automatic management and failover of systems, approaches to replicating information between services, service monitoring and health alerts, and a range of workflow requirements. We also advised on how the whole system should be maintained using modern Infrastructure as Code (IaC) tools and approaches.

The architecture we designed needed to make all content and data accessible to all Jio's users, at massive scale. We also needed to build in appropriate geographic separation and replication of services, the way they would share data and responsibility, and how closely each part of the system would need to operate with other services in order to avoid latency issues.

Additionally, due to our extensive knowledge and experience with media streaming services, we were asked to produce a detailed system design for handling streaming, broadcast, and conferencing.

SOLUTION

We produced a best-in-class architectural design for the entire infrastructure, taking into account state-of-the-art technology options and best practice, whilst being sympathetic to the specific platform and provider requirements expressed by Jio.

Central to our solution was specific identification of all the exact technologies to use, how to fit them together for the best combination of performance and resilience — and how to fully-manage and orchestrate them with automation. These ranged from databases and data stores through to queuing systems and CMS content caching options, as well as the central application processing stack.

We identified the key aspects of what would need to be tracked, processed, and stored in order to satisfy the data analytics requirements, and advised on the machine learning techniques most likely to be effective in each case. We established thresholds of acceptability for performance in order to ensure maximum responsiveness, which allowed us to also provide guidelines for required capacity in area and the behaviour in under-supply and failover scenarios. Critical to this was calculating the processing requirements with reasonable accuracy so that system load could be predicted, along with the fresh provisioning.

Every necessary automated orchestration process workflow was mapped out, with our recommendation for constructing, managing, and monitoring the facilities. We provided examples of IaC scripts and documented how to achieve full abstract and avoid close-coupling of dependencies.

We also provided detailed recommendations on the best approach to Jio's media services, including options for compressing, storing, and transcoding media; methodologies for content distribution and user-personalised recommendations, as well as advice on mixing options for audio and visual data; and guidance on format, quality pre-sets, and internet connectivity speeds, to ensure all intended devices are able to stream media.

We needed to identify and advise on the DevOps approaches that would be needed to support the technical architecture for Jio's next-generation infrastructure.

RESULTS

Jio were extremely excited by what we produced, which led to our discussing the various options in detail, with the benefit of now being told more about their existing setup. We were able to validate a lot of the approaches that they had in place, whilst also introducing new ideas and identifying areas of improvement.

As a direct result of our engagement Jio were able to finalise the design of their new platform, and introduce key practices for their DevOps team, helping them to achieve increased confidence and competency. We performed knowledge transfer sessions to ensure that their key team members fully understood the details involved, thereby helping to upskill and increase their collective abilities.

Due to the variables and metrics we had identified and established, Jio were able to model the future growth of their business and customer base, and ensure that they could accurately plan how to keep the system one step ahead.



www.dotfive.co.uk

hello@dotfive.co.uk

+44 (0) 8456 808 805

+1 (307) 222 4842

dotfive limited | frontell house, west coker hill,
west coker, yeovil, somerset, ba22 9dg
dotfive inc. | 1908 thomas avenue, cheyenne,
wyoming, 82001-3527

Dotfive is registered in England & Wales as Dotfive Limited, with company number 5445396; and in Wyoming, USA as Dotfive, Inc.

Technologies used:



AZURE AD



AZURE SQL



AZURE
COSMOS DB



AZURE CDN



AZURE BLOB

Azure | Azure AD | Azure SQL | Azure Cosmos DB | Azure CDN | Azure Blob | Azure Service Bus
AKS | Contentful | Gatsby



www.dotfive.co.uk

hello@dotfive.co.uk

+44 (0) 8456 808 805

+1 (307) 222 4842

dotfive limited | frontell house, west coker hill,
west coker, yeovil, somerset, ba22 9dg
dotfive inc. | 1908 thomas avenue, cheyenne,
wyoming, 82001-3527

Dotfive is registered in England & Wales as Dotfive
Limited, with company number 5445396; and in
Wyoming, USA as Dotfive, Inc.