



Data Revolution at Hawk Powered by Imply

About Hawk

Hawk is the first independent European platform to offer a transparent and technological advertising experience across all screens: Desktop, Mobile, CTV, DOOH & Digital Audio.

Their clients, media advertisers and agents, are on a constant lookout for more efficient digital strategies: they offer them access to the best locations on mobile and desktop screens, on digital advertising boards, audio platforms or connected TVs.

Their solution meets the needs of traders wishing to set up effective and impactful advertising campaigns:

- First, traders need to understand the available traffic, market trends, and current opportunities to set up the most efficient and performant strategies.
- Once the campaign is live, traders are very attentive to real-time results, and possible optimisations. They want to monitor their specific KPIs, and understand how they're progressing by playing on the different business levers at their disposal.
- Campaign ending goes with reporting. Traders need a holistic view of their campaign, check their KPIs, compare the performance with previous runs and draw conclusions for future campaigns.

Challenge

At every stage, Hawk account managers need precise tools to assist their users, regardless of their level of experience, by giving them relevant insights in a fast and intuitive way. But with more than 50 billion messages processed every day, providing a real-time analysis tool is a true and highly rewarding challenge.

To meet this challenge, they started with open-source Druid distribution, hosted by themselves on AWS, as the foundation stone of their data/analytics platform. But keeping a stable and performant service was a daily challenge. Druid is constantly evolving, and needs to be managed by qualified data-engineers. QoS was decent in the early years, but with extensive use by more and more users, they quickly spotted the limits of their approach: despite the mobilization of increasing resources, the sub-second commitment was an old memory.



Challenge

Needed a real-time analysis tool that could handle more than 50 billion messages processed every day

Solution

The new architecture has two key pillars:

- A real-time, Druid based analytics solution
- A cloud data-warehouse capable of storing the raw, anonymised data for heavy workloads

Results

- 50 billion messages processed every day

"The feature that we immediately loved is the Cluster Manager. With it, we can start a new Druid cluster in our AWS VPC or launch a major update in one click. Before, it was a time-consuming and dreaded operation for our data team."

Julien Lafont, CTO, at Hawk by Tabmo

Solution

Led by its CTO Julien Lafont, the company undertook a major modernisation of its big-data platform. The goal was to enable each team to produce value from its data. It had to be simple, fast and adapted to all kinds of analysis, while reducing infrastructure costs.

The modernisation had to go hand-in-hand with the significant growth of the company, especially by opening its analytics tools to external users.

The new architecture has two key pillars:

- A real-time analytics solution to ingest data from the streaming pipeline, with a rich data-mining interface.
- A cloud data-warehouse capable of storing the raw, anonymised data for heavy workloads.

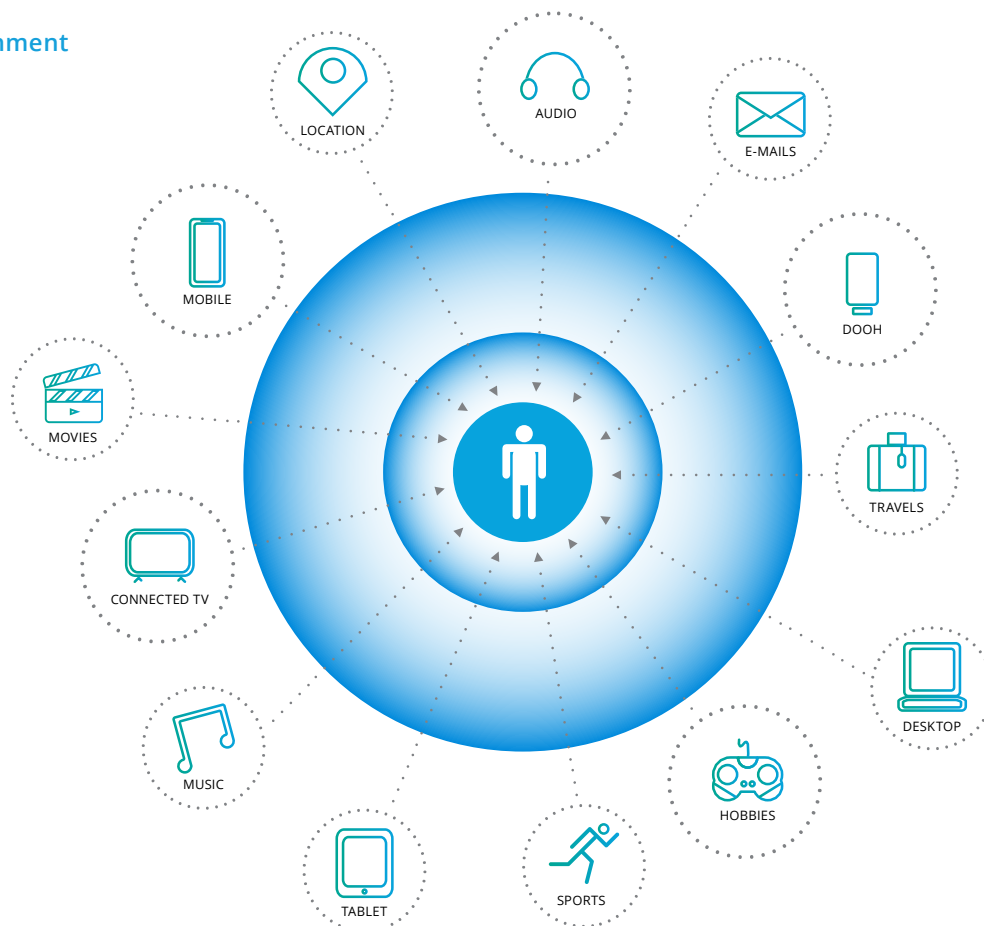
After comparing Druid to other technologies on the market, they decided to double-down on Druid, but relied on external expertise to drastically reduce its operational cost, and thus allow their data-engineering team to invest time on new projects.

Ready. Set. Go.

Imply has been a partner for the new managed Druid cluster along with other commercial solutions. Imply was selected for the expertise on the Druid solution, investment through the open-source contributions to the Druid core, and for the tools developed and made available to customers.

Pivot UI has also been greatly appreciated by their users, including the ability to create dashboards, to schedule reports or create automatic alerts.

Customer Environment



Results

The migration of their Druid cluster to the Imply Cloud solution helped them to meet several business requirements:

- Offer their customers an all-in-one solution to follow their campaign KPIs throughout their lifecycle.
- Allow their customers to build business dashboards, and visualize the evolution of their KPIs.
- Give each media-trader the capacity to set up tailored reports and send them to the right people in their organization with the automatic reporting module.
- Accelerate Time-To-Market thanks to a simplified data pipeline, with Druid directly connected to their Kafka cluster and RDS database.

Today and Beyond

In the future, they plan to replace their Aerospike NoSQL cluster with a high-performance Imply cluster that will be directly accessed through the Druid APIs. This migration will allow them to reduce the complexity of their Data infrastructure while providing new opportunities.

More Insights

Whether you are a Druid expert or data analytics driven, Imply Customer Stories will give you insights into how to build and leverage better data-driven applications.

Checkout more Imply Customer Stories at imply.io/blog



Produced from 04.30.2021 Imply blog content.

For more information visit us at imply.io

© 2021 Imply. All rights reserved. Imply and the Imply logo, are trademarks of Imply in the U.S. and/or other countries. Apache, Apache Druid, Druid and the Druid logo are either registered trademarks or trademarks of ASF in the USA. Amazon Web Services, the "Powered by AWS" logo, and any other AWS Marks are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries. Other marks are property of their respective owners.

Hawk-and-Imply-CS01-11-1-21