



# Data-driven marketing



## The world of real-time marketing

The business world is changing rapidly – new products, new ideas come and go. So it happens in the marketing environment. Marketers must give up on gut feeling as the tool of the trade and learn to use customer data to make the best decisions.

Also, it's become a standard to select recipients precisely. Gone are the days when customer target groups were selected offline based on the data from the previous week.

New-generation marketing tools help marketers meet these new demands:

- **Detect important changes in the customer journey.** In the case of the telecom industry, this could mean making more phone calls, running out of credit, approaching their operator's point of sale.
- **Check the customer's potential interest with the offer.** By analysing their behaviour history, it's possible to send a personalised message according in compliance with GDPR regulations.
- **Act fast.** Catching the optimal moment to reach out can be a matter of minutes or even seconds, e.g. when the customer is visiting the operator's physical store. If such conditions are tracked in real time, marketers can run truly dynamic campaigns involving SMS, push messages or customer service phone calls.

## Working in the fast lane

Keeping up with the new conditions requires changes not only in terms of the approach to campaign creation, but also to the implementation process. The usual way of creating online marketing campaigns looks as follows:

- The marketing team draws up a new idea.
- The concept and essential requirements are handed to the development team.
- Development starts, followed by tests, staging, refinement, and so on.

The first feedback after the implementation comes after days or weeks, even with competent agile teams. The delay caused by the involvement of multiple teams can be costly both in terms of time and money. We also have to verify if the implemented idea works. For that, we need:

- **Monitoring** to track customers' response in real time

- **A/B testing** in simple versions or multi-armed bandits, depending on the campaign size and complexity

What if marketers were able not only to create ideas for marketing campaigns and their components, but also handle the deployment, testing and monitoring?

TouK has successfully delivered a real-time marketing solution for one of the largest Polish mobile operators. The tool is based on two key components:

- **Apache Flink** – leading stream processing engine
- **TouK Nussknacker** – a unique GUI for analysts and business people to design, test and monitor processes deployed on Flink

## Apache Flink

Apache Flink is one of the leading stream processing platforms. Its main features are:

- **Scalability and throughput** – it can process 10k events per second per one core, which means that even a small cluster (2-3 machines) can handle 100k of events per second.
- **Very low latency** – our experiences show that typical event processing latency is less than 100ms.
- **Ability to easily define and efficiently process various aggregations** – e.g. time windows

While not as widely known as Apache Spark, Flink is by no means unforeseeable or immature. It has many complex deployments worldwide, including such industry giants as Alibaba, Netflix or Uber.

Flink is also used in the financial sector with ING and Capital One are widely recognisable examples, and in telecom companies, such as by Bouygues Telecom, Play and Ericsson. You can read about other use cases on the [Apache's Flink product page](#).

## Nussknacker

One of the most significant challenges in introducing stream processing platforms is giving analysts and business people (e.g. Revenue Assurance team) the ability to change the logic, test new behaviour, monitor and analyse performance.

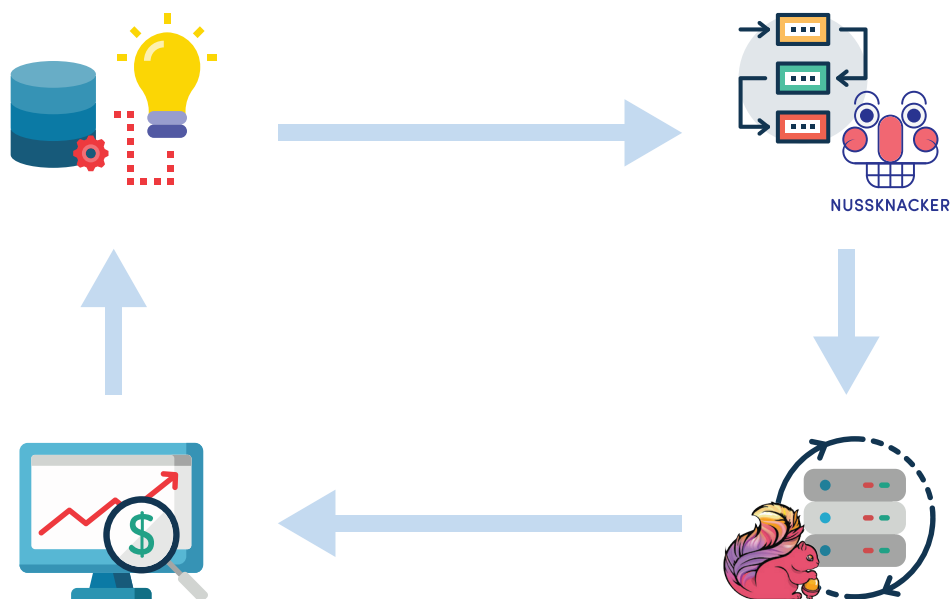
These features are especially important for companies that do not have an internal software development team. In such cases, each change that requires development from an external vendor consumes time and increases costs.

TouK Nussknacker is a unique tool that gives analysts access to the power of streaming with Apache Flink. It has a user-friendly interface built with the following general principles:

- **Business power users can author and configure processes themselves** - with no programming skills required – minimal knowledge about SQL will be enough.
- **Testing and prototyping are easy** – the users have to feel empowered to try out their ideas, and they have to be able to deploy and control the process in the production environment.

Please note that **some coding and configuration may still be needed to develop specific models and integrations** unique to a particular deployment.

What we want to enable is closing the feedback loop shown on the diagram below.



We want to enable analysts to perform all steps:

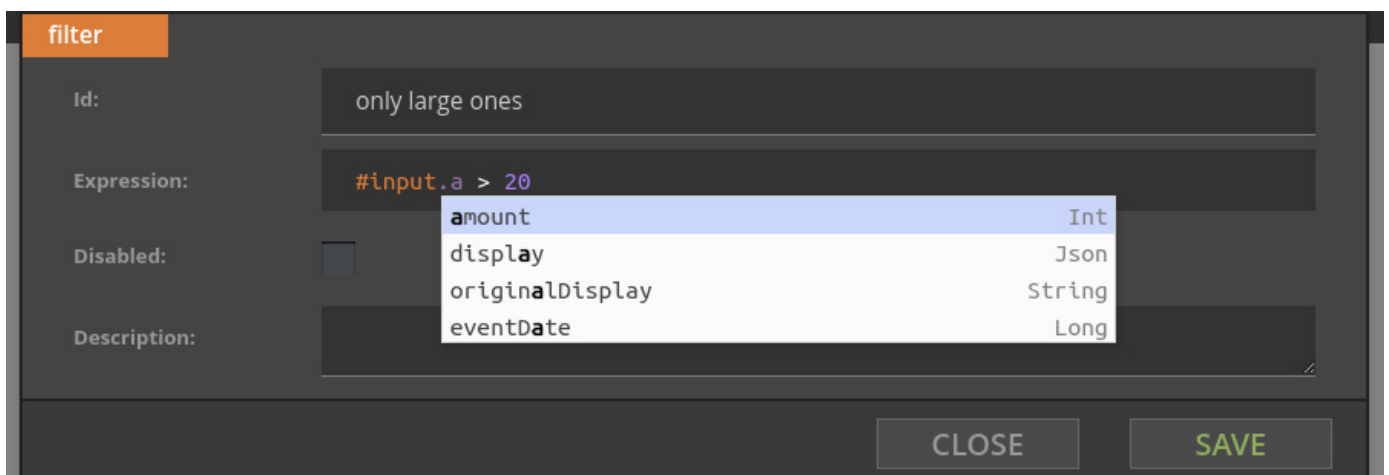
1. Design required events
2. Prototype and test the logic in Nussknacker sandbox
3. Run on real data on a high-performance Flink cluster
4. Analyse outcomes

Below you can see a sample process in Nussknacker:



The diagram defining data flow consists of some generic parts, like filters, variable definitions, aggregate definitions, and those specific for a given use case or deployment, e.g. enriching events with additional customer data or defining specific actions. In marketing scenarios, these would be contacts with customers.

Let's have a look at a sample filter and aggregate definitions:

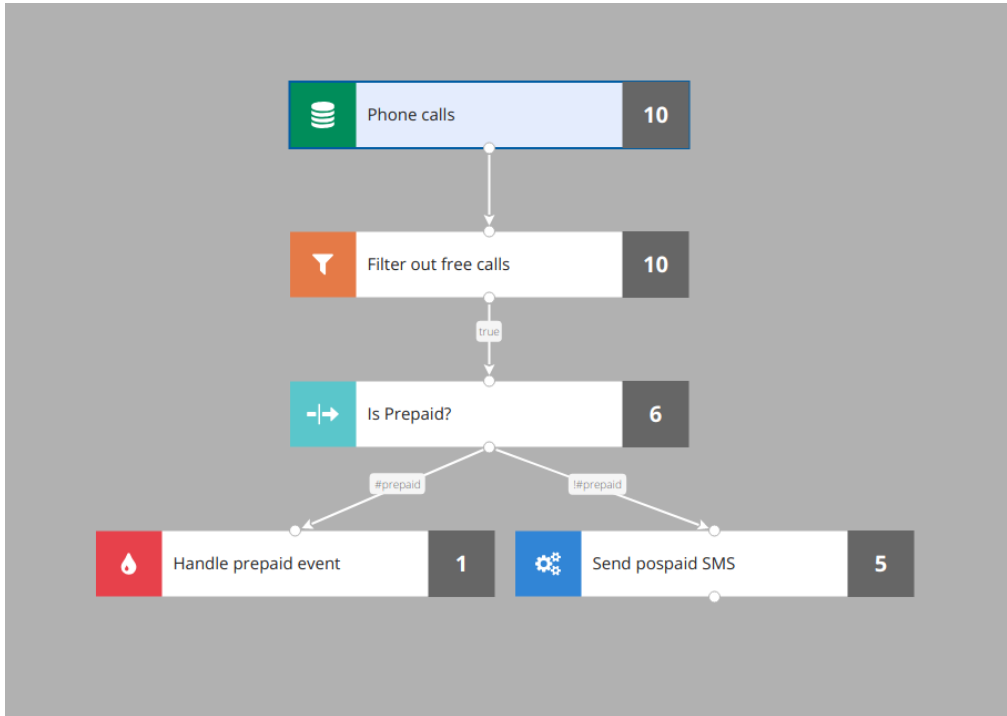


The rules are defined with a simple expression language, in many ways similar to SQL. It's expressive enough to handle most cases with no coding skills required.

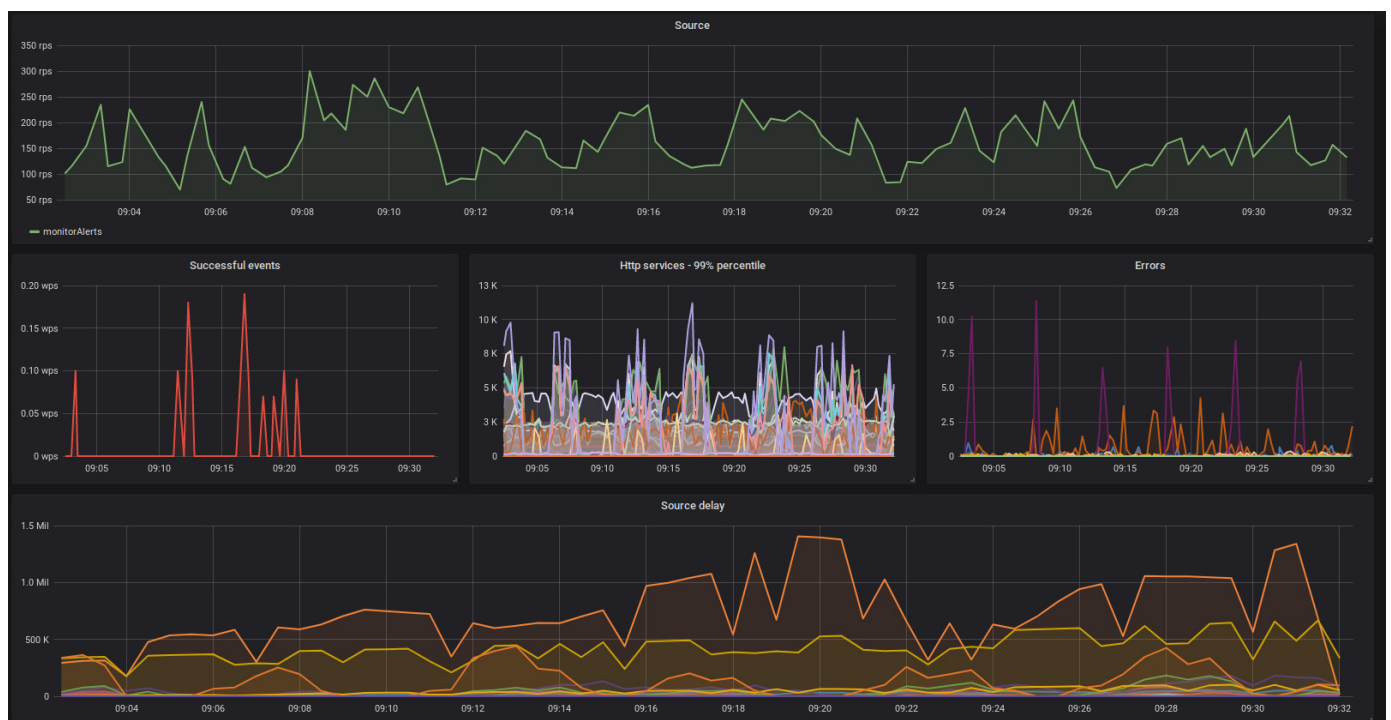
Designing a diagram is just the beginning. It's even more important to test it properly. Nussknacker can help with it in two ways:

- **Sandbox testing.** With easy-to-prepare test cases based on real data, the user can run a process in a sandbox environment and observe how it behaves.

- **Acceptance testing.** It's done in a dedicated environment, which has all relevant data but does not trigger real actions. The architecture of the platform (especially Apache Kafka features) simplifies cloning production data while Nussknacker monitoring solutions help to determine whether the process is ready for production. Then, migration is just one click away.



After deploying the process to the production environment, monitoring is crucial. We observe how the process behaves in two ways. The first one is monitoring based on metrics. Each process is displayed on an individual dashboard:



where you can see:

- **Throughput** - how many events per second are processed
- **Latency** - how long does it take from event to processing
- **Errors**, detected frauds and so on.

This dashboard gives a general view whether the process performs well and indicates if any problems occurred.

To go deeper into details, Elasticsearch and Kibana come in handy. We analyse detected events and perform basic analysis and search to find out if all conditions are fine and check whether the analysed case was real fraud or only false positives.

## Other possible use cases

Apart from real-time marketing, Flink and Nussknacker can be used in many other areas which involve business rules that change dynamically. One of the most important is fraud detection, e.g. in banking industries ([ING Bank for fraud detection](#)).

Flink is also a great match for solutions for the telecom industry, like real-time Quality of Service monitoring ([Bouygues Telecom](#)). For one of the largest Polish mobile operators, TouK successfully introduced Nussknacker and Flink as a fraud detection solution – see our [case study](#) for more information.

This form of use does not limit to telecom businesses – [VoD services](#) like Netflix are also an area worth investigating. TouK also has experience in the VoD field – we used Flink in a monitoring platform for one of the VoD providers.

## Conclusion

Stream processing gives excellent possibilities in terms of real-time marketing. However, meeting the needs of all company departments involved with the process can be complicated, especially when software development teams need to work on the same tools as business departments. Apache Flink gives powerful data processing while Nussknacker transforms this power into a handy tool that will support team collaboration.

## Learn more:

<https://touk.pl/esp/>

<https://touk.pl/portfolio>

<https://github.com/TouK/nussknacker/>

<https://touk.github.io/nussknacker/>

