



# ELASTICSEARCH STANDS UP TO COMPARISON AT COMPARE GROUP

Creating a memorable shopping experience for users.

 Company  
**Compare Group**

 Location  
**Utrecht, The Netherlands**

 Industry  
**Online price comparison**

 Products  
**Elasticsearch, Logstash, Kibana, Marvel**

 Application  
**Data indexing and publishing**

Helping hundreds of thousands of online shoppers in four countries find products and compare prices across 30 million offers is a daunting task. For The Netherlands-based Compare Group, it's all in a day's work. Compare Group is a holding company for a number of Europe's leading price and product comparison websites in Belgium, France, The Netherlands and Finland, including Vergelijk.nl, Vergelijk.be, Comparer.fr, Comparer.be and Vertaa.fi.



Compare Group's mission is to make people proud of their purchase. As Chief Operating Officer, Mark Gaasbeek, explains, *"We want to give unbiased purchase advice to people in the broad sense of the word, meaning that we provide transparency in the retail market."* To do that, the company's portals list products available both online and offline.

Anticipating the global trend of more and more people using mobile phones as their primary device, Compare Group decided to restructure its architecture. Previously, a home-grown solution based on Lucene was being used but it was difficult to configure and the databases used in Compare's search engines were getting too big for it to handle. With speed and flexibility being important criteria, Elasticsearch was chosen to power the front-end indexes searched by online shoppers as well as to provide a window into the back-end databases.



*Thanks to Elasticsearch, we are able to search for categories very quickly and directly map offers to products.*

*- Maarten Janssen, CTO of Compare Group*

---

*We are in the comparing business and with Elasticsearch we have a lot less hassle and lot more flexibility compared to what we had previously.*

*- Mark Gaasbeek, COO of Compare Group*

---

Compare's architecture comprises a combination of MySQL and MongoDB databases which are used to assemble data from a variety of sources including product offers from online and offline shops, product reviews, etc. Nearly eighty Dropwizard MicroServices running on Java – each with a data store – run a range of different tasks on the data. With data fragmented across the different stores, another microservice collects all the data and publishes it in a MongoDB database where Elasticsearch is used to access it for maintenance purposes.

*"With three thousand product categories, all of them different, a typical problem is matching the offers provided by online shops with product categories in the database", said Maarten Janssen, Chief Technology Officer at Compare Group. "Thanks to Elasticsearch, we are able search for categories very quickly and directly map offers to products."*

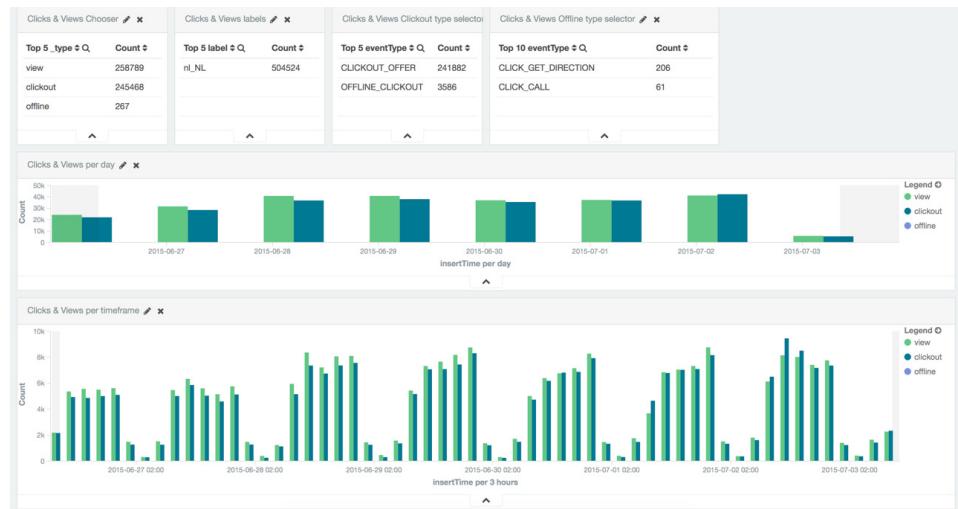
With hundreds of thousands of customers, speed is crucial to maintaining a responsive service that makes people come back time and time again. Plus, data needs to be kept continually up to date. As Mark Gaasbeek says, *"When it comes to website performance, the faster it is and the better it performs, the more users engage with the site and stay longer."*

Compare Group has a single index per label (vergelijk.nl, vergelijk.be, vertaa.fi etc), which is updated constantly during the day and night. They try to build as few indexes from scratch as possible and only do updates.

The full indexes containing around 100 GB of data and upwards 30 million records are currently swapped once a day.

It's not just at the front-end that things go faster. With Elasticsearch, the company is also saving time in the IT department. Not only are back-end processes at least twice as fast, but development times have come down significantly. Mark Gaasbeek explains: *"Development hours are scarce, so that is an indirect time saver because we can spend the time that is left on other things."* Chief Technology Officer, Maarten Janssen continues: *"As its name suggests, Elasticsearch is flexible. I don't need to worry about structures or if I can integrate something. It's just there. I know I can talk to it and do a thing and it works. I also like the stability. It's fantastically stable and downtime is reduced to almost zero, which is a huge benefit for the business side."*

To help Compare Group's marketing managers understand the data better, Kibana is being used to analyse data and identify trends to support marketing strategies. During an unexpectedly warm spell, for example, managers using Kibana noticed that there was a sudden run on air-conditioning units and were able to respond immediately, advising shops that they should feature air coolers.



**Fig: Top Clicks**

Marketing managers particularly like the ability that Kibana gives them to view all the information for each click in one place, and for each page view and then drill down and aggregate data to see which shop is the most popular in a certain category, who is doing best and which offers are getting the most views, clicks, etc.

Kibana is also being used alongside Logstash by the IT department, as Maarten Janssen explains: *"We also use Logstash for our logfiles and Kibana to see fastest response times and produced errors. If we notice, for instance, that the number of timeout errors has increased or if our response times are higher, then we know that that's something wrong and should take the necessary action."*

Before using Logstash and Kibana Compare Group only had an alerting service (site24x7 + pagerduty) which indicated something was wrong. They are still using that now since they don't use Watcher (yet), but they already benefit from the faster dashboard now.

With the incremental indexing in place Compare Group saw an improvement in the way shops see them as a platform that quickly adapt to their dynamic offer feeds and pricing. It furthermore lessened the load on their hardware because full indexing costs a lot of power for a limited amount of time.

For Compare Group, Elasticsearch has given its 5 shopping portals the power and flexibility to stay ahead of the times and meet the demands of the new mobile generation. Mark Gaasbeek comparisons are second nature: *"We are in the comparing business and with Elasticsearch we have a lot less hassle and lot more flexibility compared to what we had previously."*

To learn more about Elastic,  
contact sales@elastic.co