

AUTO

SCOUT 24



The Challenge:

How do you cost-effectively serve 100 million searches per day without sacrificing performance?

The Solution:

By using Elasticsearch to improve search query and indexing performance, while simultaneously reduce capacity requirements

Challenge

- 100 million search queries per day
- Thousands of requests per second
- Latent results being delivered in a user-facing, revenue-generating, e-commerce platform

Solution

- Utilize Elastic training and support for efficient, seamless rollout
- Achieve 10x faster query response time
- Lower TCO by reducing capacity requirements up to 5x

Serving 100 million search results a day

German company AutoScout24 is the largest online car market in Europe. In addition to the more than 40,000 European dealers using the site, almost 20 million individuals offer more than 2 million new and used vehicles for sale every year.

AutoScout24’s web applications get more than a billion page views per month from real browser users, and four times that number in bot-generated traffic. The site receives about 100 million search queries per day.

“Our main business case is meeting demand with offer,” explains Florentin Zorca, Lead Architect Search, AutoScout24. “A very important part is searching for the best offer. This is where good search functionality makes a difference.”

“Before Elasticsearch we used Endeca,” he recalls. “With peaks measured in thousands of requests per second, latency is the enemy – and we had query latency problems with Endeca. These latency problems forced us to apply a lot of performance tricks, which in turn brought more complexity. This latency was the main reason to switch the search engine to Elasticsearch.”

Elasticsearch now powers AutoScout24’s classifieds search – one of the company’s most important products – and other features, such as car dealer search. All new AutoScout24 products requiring search functionality use Elasticsearch and the company’s flexible index building pipeline.

Utilizing Elastic training and support

By using Elastic’s recommended rollout methodology – training, development support, and production support – AutoScout24 achieved an extremely smooth deployment and has had exceptional performance.

“The whole rollout was almost completely incident-free and ran even better than expected,” Zorca states. “I found Elastic on-site training to be a very useful start. It provided an extensive overview of the Elasticsearch capabilities.

“During the Elastic migration project, we often used Elastic development support with success,” Zorca adds. “We still use production support for the live platform – mostly for new features, as we rarely have issues with Elastic. AutoScout24’s experience with Elastic support is great. Our issues were quickly solved by very helpful Elastic contacts, and the response time is short.”

Handling 1 million daily updates and gaining 10x performance

“Latency of query response time was the main reason to migrate to Elasticsearch,” Zorca confirms. “Before Elasticsearch, we had an average of 200 milliseconds per query. Elasticsearch responds at least ten times faster for the same data and functionally-equivalent queries.”

Indexing speed is another important advantage AutoScout24 gained from Elasticsearch. AutoScout24 indexes around 2.2 million documents for the classifieds search which represents one of our main Elasticsearch use cases. Around one million updates to the classifieds happen daily at AutoScout24, and some classified ads are updated multiple times a day. AutoScout24 registers daily high peaks of 12,000 updated records per minute in Elasticsearch.

“All of the database updates eventually land into an Elasticsearch index,” Zorca explains. “We capture the database changes and update the Elasticsearch index every two minutes. We also build a whole index from scratch every time we release changes to the index for new product features at AutoScout24. Updating the index runs smoothly on our installation, and Elasticsearch has no problem coping with our update volume and frequency.”

Reducing capacity needs and saving money

Compared with the previous Endeca-based solution, AutoScout24 has reduced the number of virtual machines needed to handle their search queries by a third, the number of dedicated CPU cores by 5x, and the reserved memory by 6x. With the current setup, they still maintain capacity for more than 5x the current peak traffic.

“By using Elasticsearch for our main search functionality, we got the performance we needed and enough capacity for steep traffic growth,” Zorca says. “In addition, Elasticsearch enabled us to cut costs, simplify the architecture, speed up the time-to-market, and even enable new search and analytics products.”

“Hard figures show reduced costs with Elasticsearch, due to lack of license costs and reduced usage of computer resources,” Zorca continues. “With a capacity largely exceeding our current peaks, we do not need to plan for additional solution costs in the near future.”

Streamlining the release pipeline

For the migration from Endeca to Elasticsearch, AutoScout24 focused on automating and speeding up the release process. They now have a release pipeline to roll out a new Elasticsearch index anytime and, if needed, multiple times a day.

“We did everything to ensure a reliable, deterministic, automated, and flexible release pipeline for Elasticsearch indices,” Zorca says. “We benefit from the automated release with each change rollout and we have successfully reused this pipeline for various projects. Elasticsearch helps with easy installation and an excellent REST API, which every developer quickly finds familiar.”

“With its easy setup, quick document indexing and automatic cluster management, Elasticsearch simplifies the system architecture of the AutoScout24 search funnel. With its low search latency, Elasticsearch renders the earlier complex performance optimizations obsolete, further simplifying the overall solution architecture.”

The greatest advantage Elasticsearch has over the rest of the search engines is developer friendliness, according to Zorca. Every developer at AutoScout24 can quickly download, install, and interact with Elasticsearch.

“Being able to quickly set up an Elasticsearch cluster and release changes and new features provide big time savings,” says Zorca. “Additionally, with Elasticsearch we now have the possibility to create new products based on more advanced Elasticsearch features, such as free text analysis, the percolator, analytics, and scoring functionality.”

AutoScout24's benefits using Elasticsearch

- ✓ **Seamless implementation**
Architect Elastic in the best way possible by utilizing Elastic training, as well as development and production support.
- ✓ **10x faster query results**
Elasticsearch increased query response speed for AutoScout24 by at least 10x.
- ✓ **Reduced capacity needs and costs**
By reducing resource needs by up to 6x, all while increasing capacity by 5x, Elasticsearch saves AutoScout24 money.
- ✓ **Accelerating new product releases**
AutoScout24 streamlined its release pipeline with Elasticsearch, enabling the company to reduce time-to-market on new features and products.

Elastic believes getting immediate, actionable insight from data matters. As the company behind the three open source projects — Elasticsearch, Logstash, and Kibana — designed to take data from any source and search, analyze, and visualize it in real time, Elastic is helping people make sense of data. From stock quotes to Twitter streams, Apache logs to WordPress blogs, our products are extending what's possible with data, delivering on the promise that good things come from connecting the dots.