



VESPA AT WORK

Enhancing eCommerce Success at Vinted.com

*How Vespa.ai Boosted Transactions by 1.1% and Added Over
€3.5M in GMV*

The Vinted logo, featuring the word "Vinted" in a teal, cursive-style font.

Vinted



At a Glance

- 21 European Markets Plus USA
- 120 million registered users
- 77 million monthly visitors
- 1 billion active searchable items
- 25,000 item searches per second
- 10,300 requests per second data update/remove operations.

Vespa at Work: Vinted.com



Enhancing eCommerce Success:

How Vespa.ai Boosted Transactions by 1.1% and Added Over €3.5M in GMV.

"We chose Vespa because it's a modern, reliable and performant search engine that enables advanced features like image search, recommendation, ML ranking, and more. Our choice to go with Vespa is already paying off, as our members are able to find what they need more often, which leads to an increased number of purchases and higher total merchandising value. And looking to the future, we see that Vespa will continue to be a key building block and will enable new and impactful experiences for our members." - Mindaugas Mozūras, CTO, Vinted

Introduction

Vinted.com, an online platform for buying and selling second-hand items, has built its success around a big idea: make second-hand a first choice. The platform enables users to easily list, discover, and purchase pre-owned clothing, electronics and home decoration, catering to a community-driven marketplace where buyers and sellers interact directly. Robust recommendation, search, and personalization engines deliver an exemplary online experience, maximize customer engagement and grow loyalty:

- Advanced search functionality allows users to find items that fit specific needs and preferences from a staggering 1 billion searchable items for sale.
- Recommendation engines suggest items based on individual browsing and purchasing habits, increasing the likelihood of conversions and enhancing user satisfaction.
- Personalization shapes each user's journey according to their style and shopping behaviors. With tailored suggestions, from trending brands to favored categories, Vinted effectively keeps users engaged and simplifies the discovery of relevant products.

Vinted's environmentally conscious alternative to traditional retail and superior online customer experience has led to impressive revenue growth and strategic expansion. In 2023, Vinted's revenue jumped by 61% year-over-year, reaching approximately €596 million. The platform now has a user base exceeding 120 million and offers one billion items for sale, establishing itself as a dominant player in the online resale market.

High Performance to Enable Growth

Vinted Engineering faced a challenge as rapid growth strained their existing Elasticsearch implementation for recommendation, search, and personalization, leading to high operational costs and the need for hardware upgrades. To address this, they sought a more efficient solution and, after a thorough evaluation, chose Vespa in 2023 for its scalability, high performance, automated management, and ability to support both vector and traditional search in one system. Read the Vinted Engineering blog [Vinted Search Scaling Chapter 8: Goodbye Elasticsearch. Hello Vespa Search Engine](#) for a detailed technical perspective of this evaluation and migration.

High performance in Vespa is achieved with its distributed and balanced architecture, which ensures scalability and fault tolerance by distributing data, queries, and machine learning models across multiple nodes. This allows it to scale horizontally and vertically, increasing capacity and performance. By performing computations where data is stored, Vespa reduces data transfer costs and latency, enhancing the online experience by providing visitors with faster, high-quality recommendations and boosting engagement and revenue. Vespa's highly efficient architecture also immediately impacted running costs by reducing the number of servers from 120 to 60. Search result consistency has also improved as all traffic is managed within a single Vespa cluster, reducing search latency by 2.5 times.

Vespa's balanced distributed architecture evenly spreads the workload across all nodes, preventing "hot node" bottlenecks. It automatically redistributes content when content groups are changed or added, reducing the need for performance tuning and making it easy for Vinted to optimize resources. This ensures consistent performance, even when traffic patterns fluctuate.

The improved performance efficiency has also allowed a more than threefold increase in ranking depth, enabling the algorithm to consider up to 200,000 candidate items when determining which products to display in search results.

Greater Search Accuracy

Search accuracy directly impacts user satisfaction, engagement, and, ultimately, Vinted's success. Users who receive highly relevant results are more likely to engage with the platform, make purchases, or return for future visits, driving higher conversion rates and retention.

Vinted found querying in Vespa to offer significant advantages, presenting a marked shift from their previous experience with Elasticsearch. Vespa achieves search accuracy through its hybrid search and advanced ranking capabilities. It supports vector and traditional text-based search and uses machine learning models to rank content effectively. This means Vespa can understand and match complex user queries more intelligently, considering the meaning of words, user preferences, and the search context, and do so in real-time.

With over a billion constantly changing items in its inventory, efficiently organizing and preparing data for fast and accurate retrieval is essential. This process, known as indexing, transforms unstructured information, such as text from documents or data streams, into a structured format, making it easier to deliver relevant search results and power applications like recommendation systems and AI-driven search.

This ensures that the right information can be accessed instantly, supporting business decisions and enhancing customer experiences. With Vespa's real-time indexing capability—unavailable in Elasticsearch—the delay for changes to appear in search results dropped from 300 seconds, greatly enhancing responsiveness and the user experience.

Business Impact

Vinted has commended Vespa for its pragmatic approach to problem-solving and the team's genuine commitment to support. The Vespa team's active engagement has been instrumental in enhancing Vinted's customer engagement strategy, leading to notable business outcomes. Improved search capabilities with Vespa delivered measurable results: transactions purchased rose by 1.1%, and Gross Merchandise Value (GMV) increased by 0.6%. This uplift equates to over €3.5 million in additional merchandise moving through the platform, driven by Vespa's search enhancements. Furthermore, the Total Cost of Ownership of the Vespa implementation was approximately half that of the previous system, offering a highly cost-effective solution that simultaneously boosted user engagement and conversions. This combination of reduced costs and improved search performance has directly contributed to revenue growth and enhanced customer satisfaction by making relevant products easier to find.

About Vespa.ai

Vespa.ai is a platform for building and running real-time AI-driven applications for search, recommendation, personalization, and RAG. It enables enterprise-wide AI deployment by efficiently managing data, inference, and logic, handling large data volumes and over 100K queries per second. Vespa supports precise hybrid search across vectors, text, and structured metadata. Available as both a managed service and open source, it's trusted by organizations like Spotify, Vinted, Wix, and Yahoo. The platform offers robust APIs, SDKs for integration, comprehensive monitoring metrics, and customizable features for optimized performance.

[Vespa.ai](#)