

Executive Summary: Real-Time Recommendation Engines

“You may also like’ is a deceptively simple phrase: In offering tailored suggestions, businesses **maximize value by providing highly targeted, real-time product recommendations to their online consumers.**”

[Learn More](#)

Who Needs a Recommendation Engine:

- Ecommerce retailers
- Media and content publishers
- Offline retailers with online channels
- Internal data discovery applications
- Search engines
- Sales enablement tools
- Social networks or social-component applications
- Academic course catalogs

Common Challenges for Recommender Systems:

- Connecting a vast array of various siloed systems
- Ensuring results with a relevancy of 99% or above
- Including results from a user’s current session or shopping cart
- Avoiding expensive JOINS or overnight batch processing
- Delivering results in **real time** with zero latency

Challenge: Distilling (and Delivering) the Relevance from the Noise

Marcos Wada realized that a relational database wasn’t going to cut it. As a software developer for the Walmart ecommerce Brazil group, Wada had one goal: Give shoppers at Walmart.com the best web experience. This meant nth-degree personalization.

However, he was struggling to deliver *real-time*, highly relevant product recommendations because of the complex and siloed nature of Walmart’s buyer and product data.

“A relational database wasn’t satisfying our requirements about performance and simplicity, due to the complexity of our queries,” Wada said.

Of course, his was also a scale problem: Walmart has nearly 250 million weekly customers across 11,000 stores in 27 countries – not to mention retail websites in 10 countries.

Amidst all this data noise, Wada had to find a way to distill each user’s preferences, purchase history, product similarities and current shopping cart data. Then, using this distilled analysis, he had to deliver the most accurate product recommendations in real time at Walmart scale.

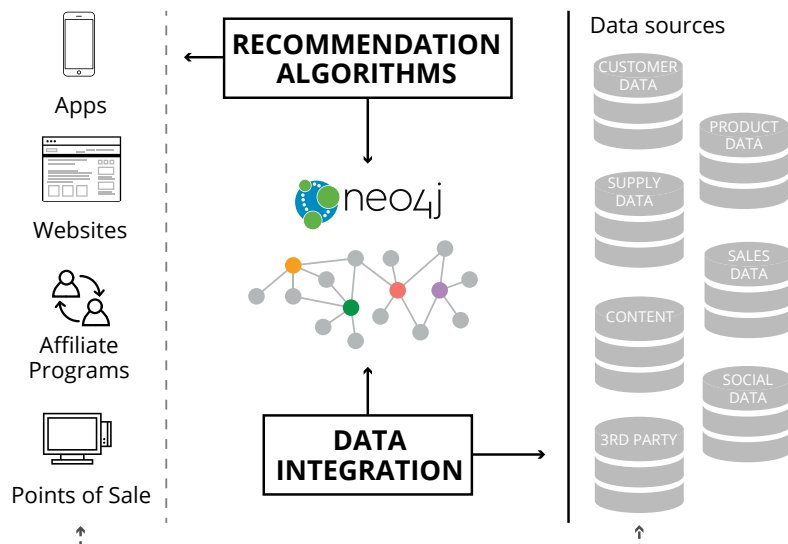
Why User Personalization Requires a New Approach

Today’s users expect finely tuned personalized experiences, whether they’re shopping online, using a search engine or browsing the news. However, this level of relevance requires data hops at depths that RDBMS can’t handle or deliver within a meaningful time window.

As a result, many recommendation engines resort to the following sub-optimal tactics:

- They resort to batch processing to pre-compute recommendations
- They traverse no more than three levels of depth
- They don’t flexibly accommodate new data sources

No matter the trade-off, the result is the same: A loss of context and less relevant results.



Executive Summary: Real-Time Recommendation Engines

Neo4j Graph Platform Benefits

Neo4j Database

Deliver real-time user personalization with the power of [a native graph database that maintains rapid performance](#) even as your dataset grows.

Cypher Query Language

Rewrite your 100-line SQL queries with just seven lines of Cypher ([like eBay did](#)) and increase efficiency while reducing tuning and debugging times.

Data Integration Tools

Don't rip and replace your current data stores: Integrate them seamlessly with Neo4j for recommendations that [draw in data from every source and silo](#).

Data Discovery & Visualization

Easily communicate and visualize the efficiency of your recommender systems to non-technical peers [using Neo4j Bloom for data visualization](#).

Graph Analytics

Test and optimize your recommendation algorithms with [powerful, offline graph analytics](#) and feed those insights seamlessly into your operational database.

Solution: Real-Time Personalization for an Infinite Audience of One

Wada first encountered Neo4j at a QCon conference and instantly recognized its potential. His team experimented with the graph database technology and the results were positive.

Using Neo4j, Wada's team **substituted a heavy batch process with a real-time graph database to make sense of online shoppers' behavior**.

"With Neo4j we could substitute a complex batch process with a real-time recommendation system featuring low-latency queries," Wada said.

Walmart's recommendation engine now swiftly sifts through mountains of data to optimize and cross-sell major product lines in core ecommerce markets. As a result, Walmart delivers real-time recommendations that personalize an online shopper's experience as if they were an audience of only one.

Even better, Wada's team achieves that effort at a potentially infinite scale.

"As the current market leader in graph databases, and with enterprise features for scalability and availability, Neo4j is the right choice to meet our demands," Wada said.

Read the rest of the story at <https://neo4j.com/case-studies/walmart/>

Building Real-Time Recommendation Engines with Neo4j

Neo4j is a native graph database that takes a connections-first approach to data. These relationships are persisted as first-class entities equal to other data points – rather than disappearing after each analysis or being prohibitively expensive to query.

With Neo4j, you continually improve relevant, real-time recommendations by accommodating new data sources and types – without rewriting your data model.

Using graph tech, recommender systems conduct connected data analysis across users, products, social, shopping cart and supply data to make the perfect suggestion at scale.

Real-Time Recommendation Applications Built on Neo4j

- **Increased Revenue:** Recommendations done right have direct impact on revenues.
- **Higher Engagement:** Improved personalization and content recommendations lead to higher user engagement.
- **More Accurate Searches:** Recommendation-based search engines are essential to sorting out irrelevant information and delighting end users.

[Learn more about using Neo4j to deliver real-time recommendations](#)

Neo4j, Inc. is the graph company behind the leading platform for connected data. The Neo4j graph platform helps organizations make sense of their data by revealing how people, processes and digital systems are interrelated. This connections-first approach powers intelligent applications tackling challenges such as artificial intelligence, fraud detection, real-time recommendations and master data.

More than 250 commercial customers, including global enterprises like Walmart, Comcast, Cisco, eBay and UBS use Neo4j to create a competitive advantage from connections in their data.

Questions about Neo4j?

Contact us:

1-855-636-4532

info@neo4j.com