



Tourism Media Automates Travel Content Generation with Neo4j Aura at a Massive Scale

INDUSTRY

Travel / Media

USE CASE

Knowledge Graph

GOAL

Automatically generate custom content for hundreds of thousands of travel destinations

CHALLENGE

Existing relational database couldn't model the ragged hierarchy of geographic data

SOLUTION

Used Neo4j Aura, a fully managed graph database to model complex relationships and add new data sources

RESULTS

- Automated content generation performs as well or better than human generated content
- Worry-free cloud deployment offers scalability and frees team to focus on new product features

Tourism Media generates travel content for destinations all over the world but relational databases couldn't keep up with changing model and scale requirements. With Neo4j Aura, they could model geographic data for 200,000 cities in the cloud, increase context, and automate content generation at a massive scale.

The Company

Tourism Media is an SMB with a lean, in-house team and an extended collection of editors, writers, and photographers around the globe. Seeking to bring better travel content to the online space, the Tourism Media team offers diverse services to their clients – ranging from original videos, photography, and written content to entire content management and delivery systems. The company's clientele includes travel industry giants.

The Challenge

Tourism Media produces dynamic travel content for hundreds of thousands of destinations across the world. Travel content is highly reusable for similar places, but each piece has to be customized and refreshed regularly for both geographical and seasonal reasons. And when it comes to data, geography is highly unstructured and difficult to model.

The Tourism Media team tried to fit their geographical data into a relational data model, but its predefined, columns-and-rows structure made it difficult to navigate that data and produce relevant, meaningful content in a scalable way.

"If we're generating content around a ragged hierarchy like geography, how do we do that in a way that makes sense and scales?" said Terry Franklin, a freelance consultant specializing in Neo4j who worked with Tourism Media on its implementation. "For example, if we produce photos for the country Singapore, how do we make sure that they also appear on pages that understand Singapore to be a city?"

Tourism Media wanted a way to ask complex questions of data that is highly connected and hierarchical without lots of JOINS and long-running queries. They wanted agility for handling challenges such as handling route mapping for cruises around the world, translating local content into multiple languages, and automatically serving up top rated local hotels – in short meeting needs of travelers today and tomorrow.

The Strategy

While looking into ways to solve this problem, the Tourism Media team came across graph databases. They built a graph data model in Neo4j to vet how well it reflected their geographical data.

“Graph databases immediately became a much better fit, because they store your data in a way that human beings can intuitively relate to it much more readily,” Franklin said. As their use case began to scale, the Tourism Media team needed to expand its use of Neo4j. The release of Neo4j Aura with its pay-as-you-go pricing model and elastic scalability was perfect for the team’s needs.

The Solution

With the help of the Neo4j support team, Tourism Media migrated their data to the cloud, and it was quickly ready for production. The Tourism Media graph is flexible and expanding, with new data sources added to enable new products. Standard content templates use aggregated data coupled with simple, repeatable queries to create content for hundreds of thousands of destinations. Relationships map geographical hierarchies and allow context to be added on top of each template. Ontology-based weighting is applied to tags to provide relevant recommendations for similar places. Using only GraphQL, Tourism Media can now ask complicated questions of their knowledge graph.

“Neo4j, GRANDstack, and APOC really fit well together, because they allowed us to import all the data into the graph in a way that truly represents geography and ask the questions in GraphQL in the exact same shape that we expect the answer back,” Franklin said.

The Results

Tourism Media’s knowledge graph now handles highly complex queries much faster. As Tourism Media continues to add new data sources, Neo4j Aura enables them to scale their business and meet the needs of their clients.

The Tourism Media team is able to generate valuable content for clients at a much faster rate as well as refresh and regenerate those pieces at a useful cadence.

“There are about 200,000 cities in the system; those builds took the relational database seven or eight days to complete,” said Franklin. “When we switched to a graph database, it took less than a day.”

Responses are “very positive from the clients,” according to Franklin. Reports show that content that is automatically generated performs as well as human-written content in terms of SEO – in some cases, even better.

The flexibility of the graph data model enables Tourism Media to move quickly to add new features including accommodation options, activities, attractions and flights – including whether it’s a direct flight or layover, how long it takes to get from a city to the destination, which planes are available, and more.

Tourism Media chose Neo4j Aura for multiple reasons. “It’s the reliability, the confidence that it’s always available and the fact that it’s managed,” said Franklin. “There is a good support team behind it, and it’s continuing to move forward. That makes Aura very compelling, I think.”

“There is about 200,000 cities in the system; those builds would take seven or eight days to complete. When we switched to a graph database, it took less than a day.”

–Terry Franklin, Freelance
Consultant Specializing in Neo4j

Neo4j, the Graph Database & Analytics leader, helps organizations find hidden relationships and patterns across billions of data connections deeply, easily, and quickly. Customers leverage the structure of their connected data to reveal new ways of solving their most pressing business problems, from fraud detection, customer 360, knowledge graphs, supply chain, personalization, IoT, network management, and more – even as their data grows. Neo4j’s full graph stack delivers powerful native graph storage with native vector search capability, data science, advanced analytics, and visualization, with enterprise-grade security controls, scalable architecture, and ACID compliance. Neo4j’s community of data leaders comprises a vibrant, open-source community of more than 250,000 developers, data scientists, and architects across hundreds of Fortune 500 companies, government agencies, and NGOs. Visit neo4j.com.