



Optimizing Quality of Service and Coverage with Advanced Geospatial Analytics

- *Advanced geospatial workloads supporting gap analysis, previously impractical using industry standard tools*
- *Never-before visual fidelity of every building and population to inform strategic network build decisions*

T-Mobile US provides wireless voice, messaging, and data services in the United States. There are several key factors T-Mobile competes on, including coverage and quality of service. Increasingly, the ability to analyze massive geospatial data sets to help make the best possible, highest ROI, network improvement and site buildout decisions, requires more advanced capabilities and more processing power.

Challenges

Many of the popular cloud and previous generation databases haven't kept up with the geospatial opportunities and requirements of modern-day big data, causing T-Mobile to seek a new solution for the type of large scale geospatial, predictive and interactive rendering capabilities it needed to support real-time advanced analytics. Given the complexity and growth of the data involved, T-Mobile knew the problem would only get worse at an increasing pace, especially with the network density of 5G and the growing availability of IOT data harvesting.

When it looked to replace the industry-standard solution it had been using, open source PostGIS spatial database extender for PostgreSQL object-relational database, with a solution that could scale to process billions of data events, including cell phone, weather, and more, T-Mobile knew they'd be charting new territory. Prior generation databases were never designed to handle location-enriched data for modern

geospatial analytics, resulting in excessive costs, needless latency, and lost data insights. T-Mobile identified several key limitations of its current environment, including significant waste and analytical limitations in its current data processing; and not extracting the full value of its data or providing sufficient rationale for investment prioritization.

Additionally, T-Mobile was suffering from a resource management standpoint. Because its prior open-source solution required continuous optimization that was unsustainable for the variety and ad-hoc nature of their initiative, T-Mobile had to provide more monitoring and vertically scale resources to match the data set. This caused a lot of outage problems over time because of the explosive nature of how downloading data can spike CPUs. It needed a solution that would lighten the load, resulting in better performance to run operations and expedite the flow and load balancing against its current architecture.

Selection

After a thorough and stringent evaluation, Kinetica, the database for time and space, was selected. In its research, T-Mobile found no other database rivaled Kinetica's real-time analytical and processing power of geospatial and time-series sensor data. When the team set-out to further justify its investment, it found only Kinetica met these never-before achieved requirements:

- *Analyze* massive geospatial data sets and add more processing power of billions of records and apply complex geospatial and predictive modeling from a single unified database that is ANSI SQL 92 compliant
- *Gain* powerful geospatial processing that goes far beyond the very limited geospatial workloads supporting network gap analysis, for example, that could take months or even years to complete
- *Predict* network build or coverage ROI to rationalize spend prioritization decisions
- *Provide* interactive, real-time analysis in a context the business already understands
- *Add* real-time data feeds and dashboards to leverage aggregated data from larger historical datasets
- *Utilize* predictive analysis and advanced analytics more broadly to generate deeper network and subscriber insights
- *View* actual network coverage in real time, vs planned coverage
- *Leapfrog* competitors already using legacy big data analytics

Working with Kinetica in an initial trial project, T-Mobile processed 90 billion spatial object records, turning lat/longs provided by Google into atomic routes which could be used for coverage planning and ROI projections for network/retail build out. Previously, attempting to execute Google Route analysis took several weeks and significant technical churning to map

two streets in San Francisco. On the Kinetica database, T-Mobile was able to map all streets in the entire state of California in under an hour.

"We had geospatial workloads supporting network gap analysis that could take months or even years to complete on the previous data stack," said James No, engineer, software, T-Mobile.

"Today, we're blending more big data sets leading to better performing models and more accurate ROI prediction," explained Jake Holman, senior manager, software development, T-Mobile.

T-Mobile was also spending a significant amount of time and effort managing a hodgepodge of big data solutions that offer poor agility. Its big data datasets required special tooling and techniques to analyze and manage. Tools like Hadoop and Spark handled some workloads well but the stack was not very cohesive and the setup overhead involved for these tools was substantial and the functionality provided was limited.

"In some cases, especially for geospatial workloads, these tools were difficult or impossible to scale," said Jeny Bhimani, senior product manager, technical, T-Mobile. *"Kinetica gives us a unified database platform that saves substantial time and provides functionality not available elsewhere."*

"We had geospatial workloads supporting network gap analysis that could take months or even years to complete on the previous data stack."

James No
Associate Software Engineer, T-Mobile

Results

Since the trial, Kinetica has been primarily used for large network data sets where layer rendering and visualization has been a significant use case. For example, the 'Buildings Layer' data set entails visualizing the in-building coverage experience for each building area nationally in the US. The processing and visualizing of the data set is where Kinetica comes in. For the first time, T-Mobile has the visual fidelity of every building and can experience their coverage in a visual format.

"Our analysis must be predictive driven," explained No. "It's no longer sufficient to only use a static historical perspective to inform network planning. Predictive modeling using rich streaming location data is becoming an analytical baseline in telecom. With Kinetica, for the first time ever, we have an analytical and visual representation of our coverage in every building in the country."

T-Mobile also gained huge time savings processing massive joins of data sets of tables with several billion rows with fact tables with tens of million records going from several seconds to milliseconds with Kinetica. The Buildings Layer is combined with a Census Block Layer that together is used to determine

"Our analysis must be predictive driven. It's no longer sufficient to only use a static historical perspective to inform network planning."

James No
Associate Software Engineer, T-Mobile

"In some cases, especially for geospatial workloads, these tools were difficult or impossible to scale. Kinetica gives us a unified database platform that saves substantial time and provides functionality not available elsewhere."

Jeny Bhimani
Sr Technology Product Manager, T-Mobile

coverage and planning network builds by both the business and home Internet groups. Kinetica displays critical coverage and population data that informs these strategic decisions.

Kinetica is the geospatial analytical database in T-Mobile's cloud stack that includes other best of breed technologies in the company's overall cloud on-premise architecture.

