

The background of the slide features a complex network diagram. It consists of numerous small, light-gray circular nodes scattered across the white space. These nodes are interconnected by thin, light-gray lines, forming a dense web of connections. The overall appearance is that of a large-scale network or data structure, possibly representing a telecommunications network or a data processing pipeline. The nodes and lines are more prominent in the center and fade slightly towards the edges.

Real-Time Analytics for Time-Series & Spatial Workloads in Telecommunications

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

James No 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, senior product manager, T-Mobile

“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.”

James No Software Engineer, T-Mobile

What our clients have built with Kinetica


Real-time streams of cell phone signals **fused on spatial dimensions** with buildings and roads to detect and address 5G coverage weak spots **resulting in** better cell service for customers*

*“T-Mobile wins the 5G Download Speed award for the sixth time in a row, increasing once again its lead on Verizon and AT&T.” – OpenSignal, July '22



Kinetica's Superpower: NATIVE VECTORIZATION

Kinetica scales out like all MPP databases, plus parallelizes the data within each node, known as vectorization. This force multiplier results in order of magnitude performance improvements on a smaller compute footprint, and removes the complex data engineering required by other databases to make up for their inefficiencies.



In the Lab

Startups like ClickHouse and established vendors like Databricks are using native vectorization to leapfrog traditional databases, but TPC-DS 100 benchmark results on identical CPUs shows they can't come close to matching Kinetica.

13X Faster than ClickHouse


8X Faster than Databricks

In the Field	Replaced	with Kinetica
Top Wall Street Bank	700-nodes of Spark	16-nodes of Kinetica
Top Global Retailer	100-nodes of Cassandra	8-nodes of Kinetica
Big Pharma	88-nodes of Impala	6-nodes of Kinetica

Best in Class Location and Temporal Intelligence

Adding location and time dimensions to data makes it significantly more valuable, but much harder to process and analyze.

Kinetica solves this problem with a native vectorization that operates on millions of data points with a single instruction.



Spatial Functions

- ST_Geometry suite
- Entity Tracks
- Geo-joins
- Server Side Visualization
- **And more**



Time Series

- Aggregations
- Window functions
- Interpolations
- AS-OF joins
- **And more**



Graph

- Adjacency
- Centrality
- MSDO
- Shortest Path
- **And more**

kinetics

The Database for Time & Space