

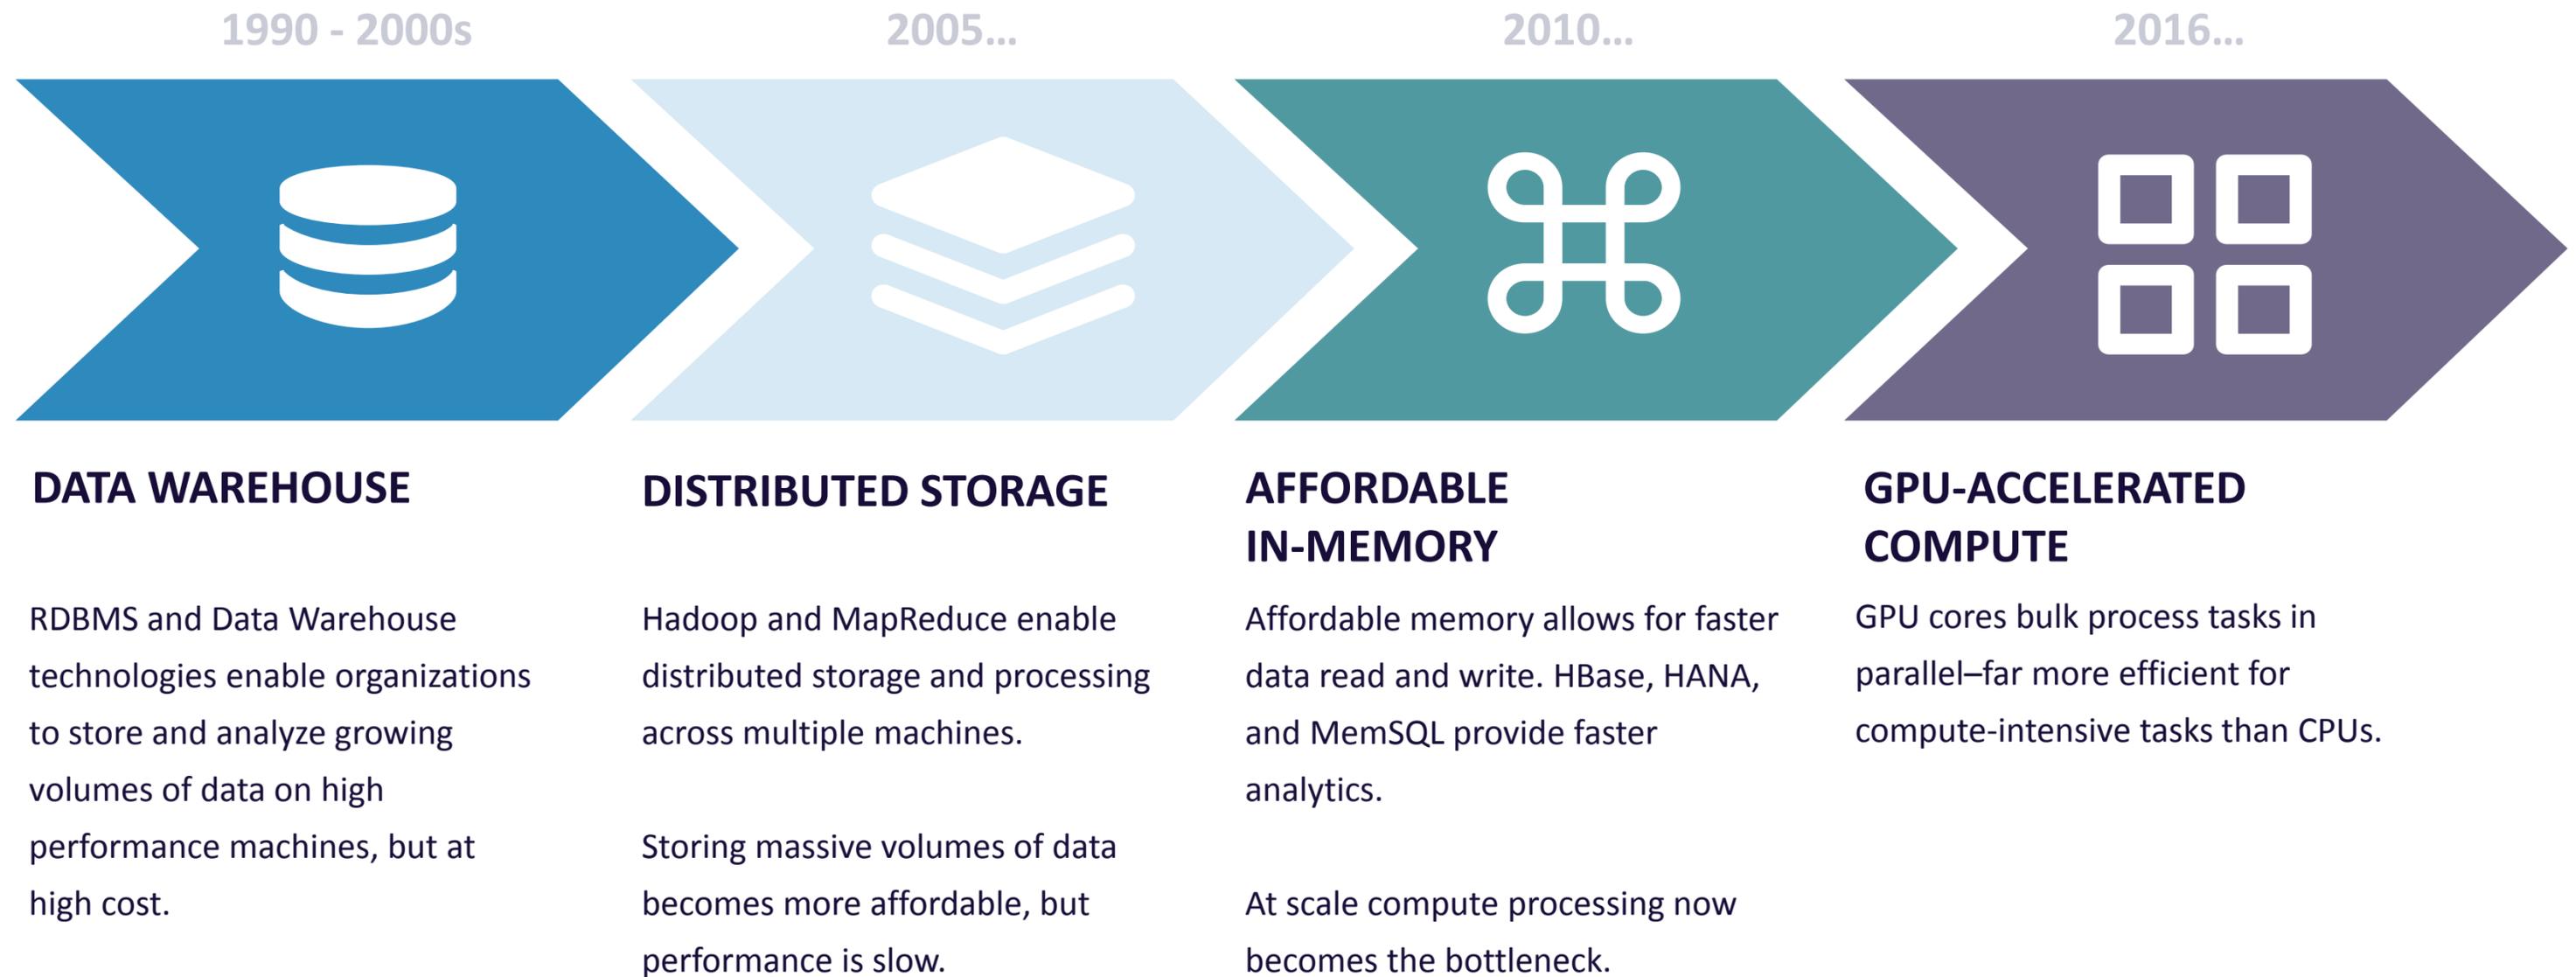


GPU-Accelerated Analytics Enable New Enterprise Solutions

Mark Johnson, Business Development

November 14, 2016

Evolution of Data Processing

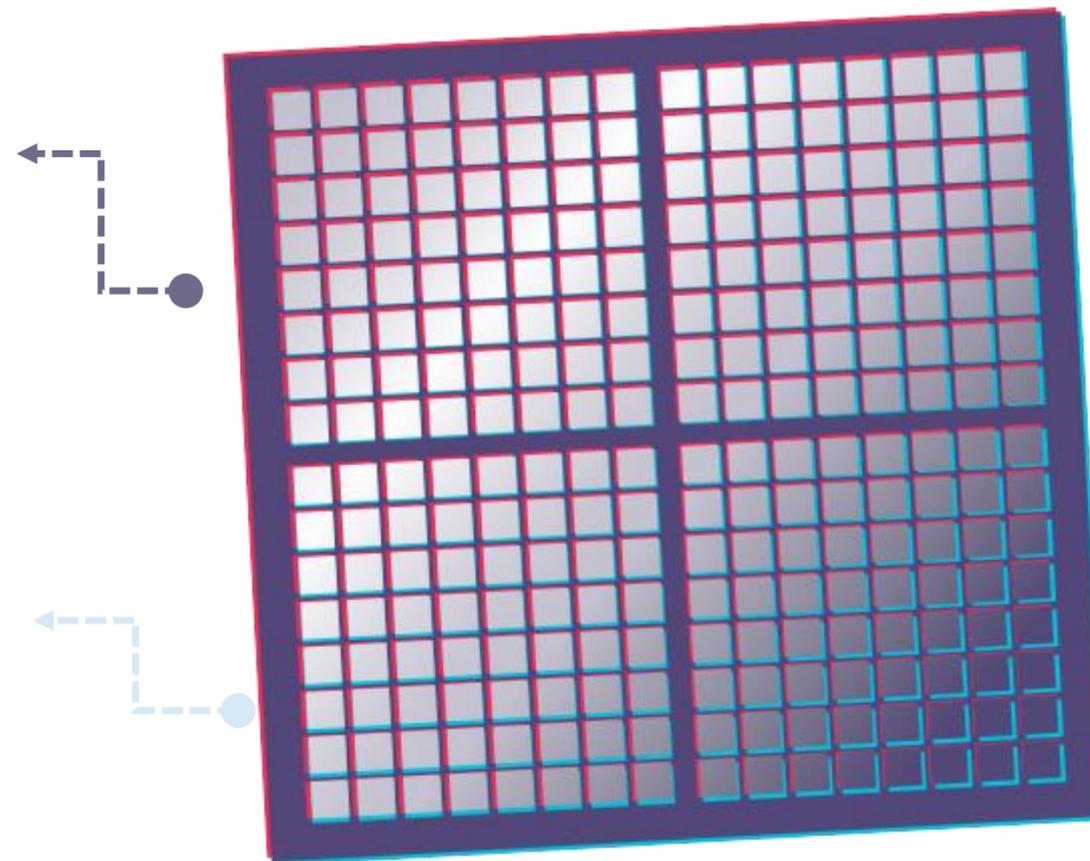


GPU Acceleration Overcomes Processing Bottlenecks

GPUs are designed around thousands of small, efficient cores that are well suited to performing repeated similar instructions in parallel. This makes them well-suited to the compute-intensive workloads required of large data sets.

4,000+ cores per device in many cases, versus 8 to 32 cores per typical CPU-based device.

High performance computing trend to using GPUs to solve massive processing challenges



Parallel processing is ideal for scanning entire dataset & brute force compute.

GPU acceleration brings high performance compute to commodity hardware

Kinetica Background

United States Army Intelligence seeks a means to assess terrorist and other national security threats.



GPUdb goes live with the US Army Intelligence.



Commercialization entered production with USPS.



No database in the market was fast or flexible enough to met their needs.



Founders Amit Vij and Nima Negahban start on the pioneering use of GPUs while building a GPU-accelerated database from the ground up.

Patent granted for “Method and system for improving computational concurrency using a multi-threaded GPU calculation engine”



Wins IDC HPC Innovation Excellence Award for work with US Army.



Wins IDC HPC Innovation Excellence Award for work with US Postal Service.

kin tica

Rebranded to Kinetica. Seed funding. Moved HQ to San Francisco. Expanded management team. Hired field team.

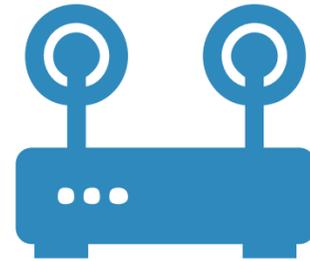
Kinetica: A Distributed, In-Memory Database



GPU-accelerated
database operations



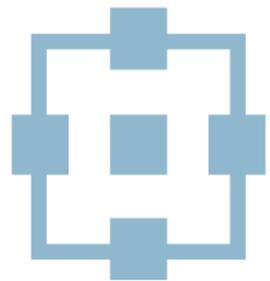
Natural Language
Processing based
full-text search



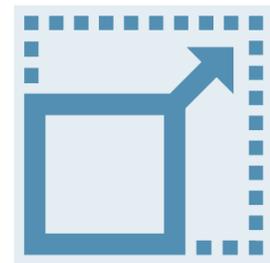
Native GIS & IP
address object
support



Real-time data
handlers to ingest
structured and
unstructured data



No typical tuning,
indexing, and
tweaking



Predictable scale out
for data ingestion and
querying



Deep integration with open
source and commercial
frameworks and applications:
Hadoop, Spark, NiFi, Accumulo,
H2O, Tableau, Kibana and Caravel



Distributed visualization
pipeline built-In



Enabling New Enterprise Solutions



Real-time Threat Intelligence

Ingests 200+ sources of streaming data, producing 200B new records per day; incorporates geo-spatial and temporal data to identify terrorist threats in real time



Logistics, Fleet Optimization

Tracking carrier movements in real time, to reallocate resources on the fly based upon personnel, environmental and seasonal data. 200,000+ USPS devices emit location once every minute, so > ¼ billion events captured and analyzed daily, with several times that amount available in a trailing window.



Cybersecurity

Fusing multiple data feeds with real-time anomaly detection to protect its financial clients against current and emerging cyber threats.



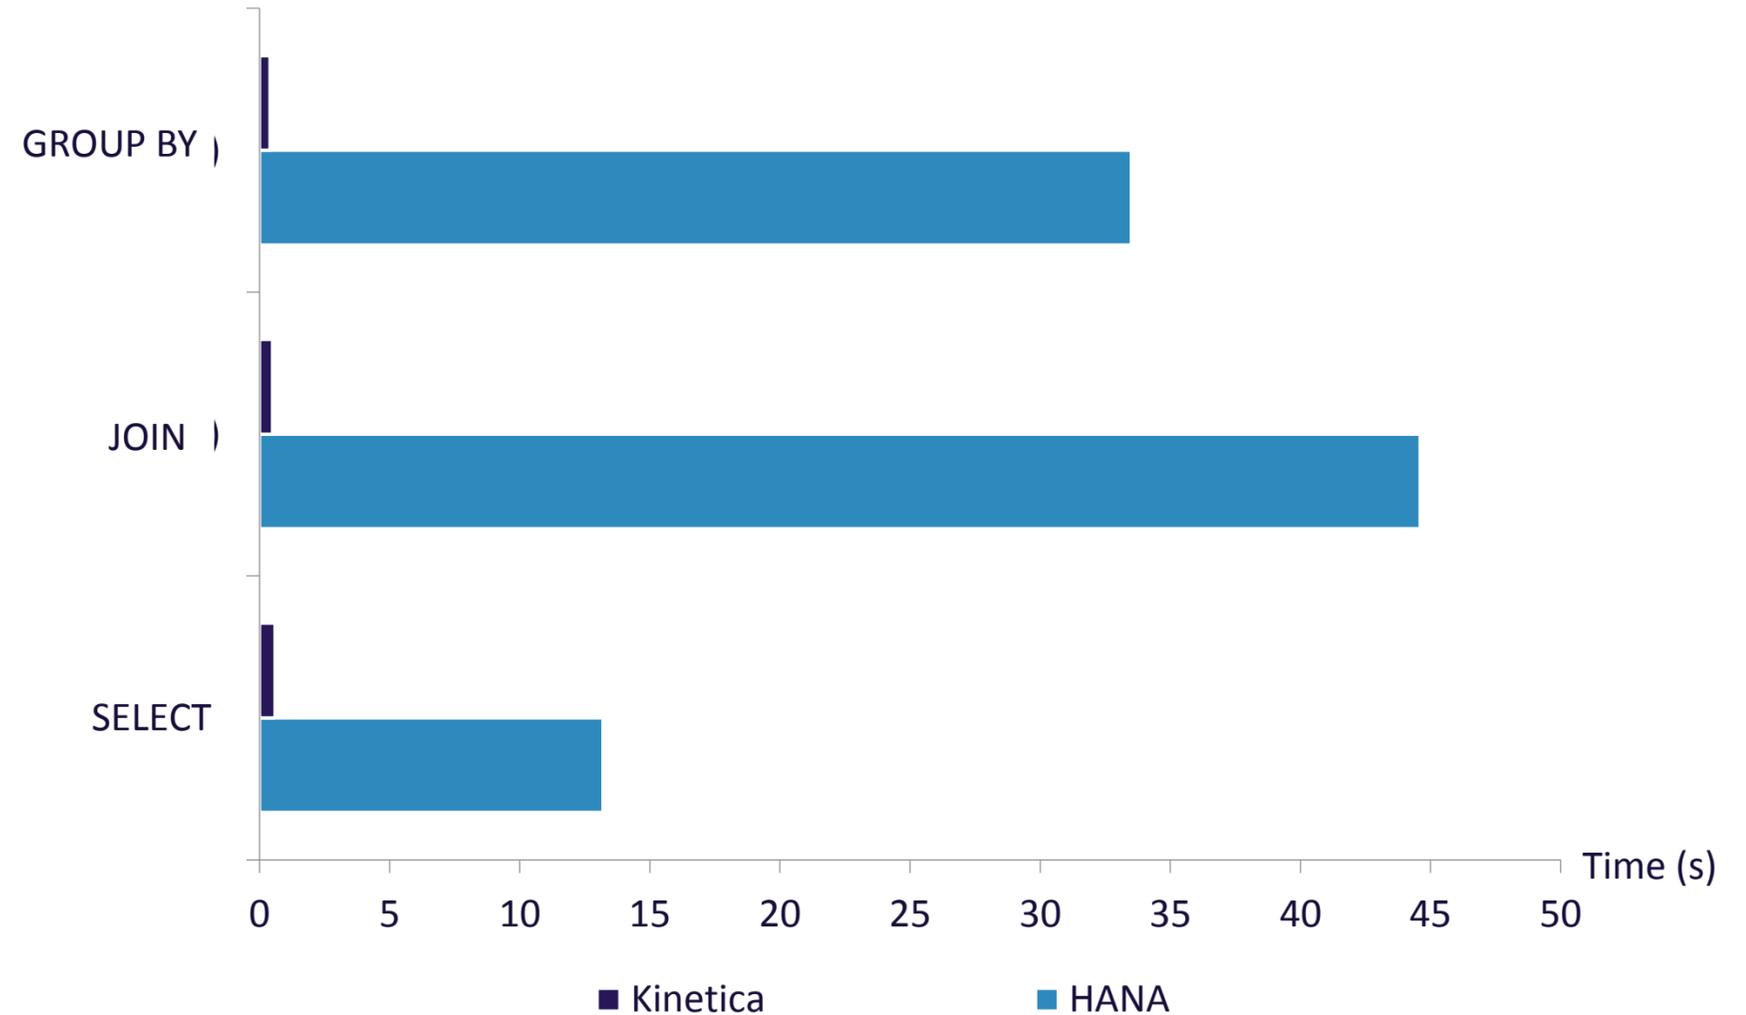
Utility: Smart Grid

Geospatially visualizing smart grid while ingesting real-time vector data streaming from smart meters. Striving for optimized energy generation and uptime based on fluctuating usage patterns and unpredictable natural disasters.

Impressive Performance vs. HANA

IN-MEMORY KNOCKOUT

30 nodes Kinetica cluster tested against 100+ node HANA cluster holding 3 years of retail data (150b rows) with a range of queries



Enabling New Enterprise Solutions



Retail: Customer 360/customer sentiment, supply chain optimization

Correlating data from point of sales (POS) systems, social media streams, weather forecasts, and even wearable devices. Better able to track inventory in real time, enabling efficient replenishment and avoiding out-of-stock situations.



Powering High Performance “Analytics as a Service” Solution

Delivering customer-focused services by leveraging all available transactional data. Currently no ability for business user to do customized analytics; IT has to. Query response times taking 10s of minutes, some over 2 hours, thus limiting ability to analyze and use data.



Fin services

Large scale risk aggregations and billion+ row joins in sub-second time (5TB+ tables choke on RDBMS joins and Hadoop is too slow).



Ridesharing/Connected Cars

View all passengers and drivers to monitor behavioral analytics. Sensor enabled predictive maintenance on connected cars.



Manufacturing IoT

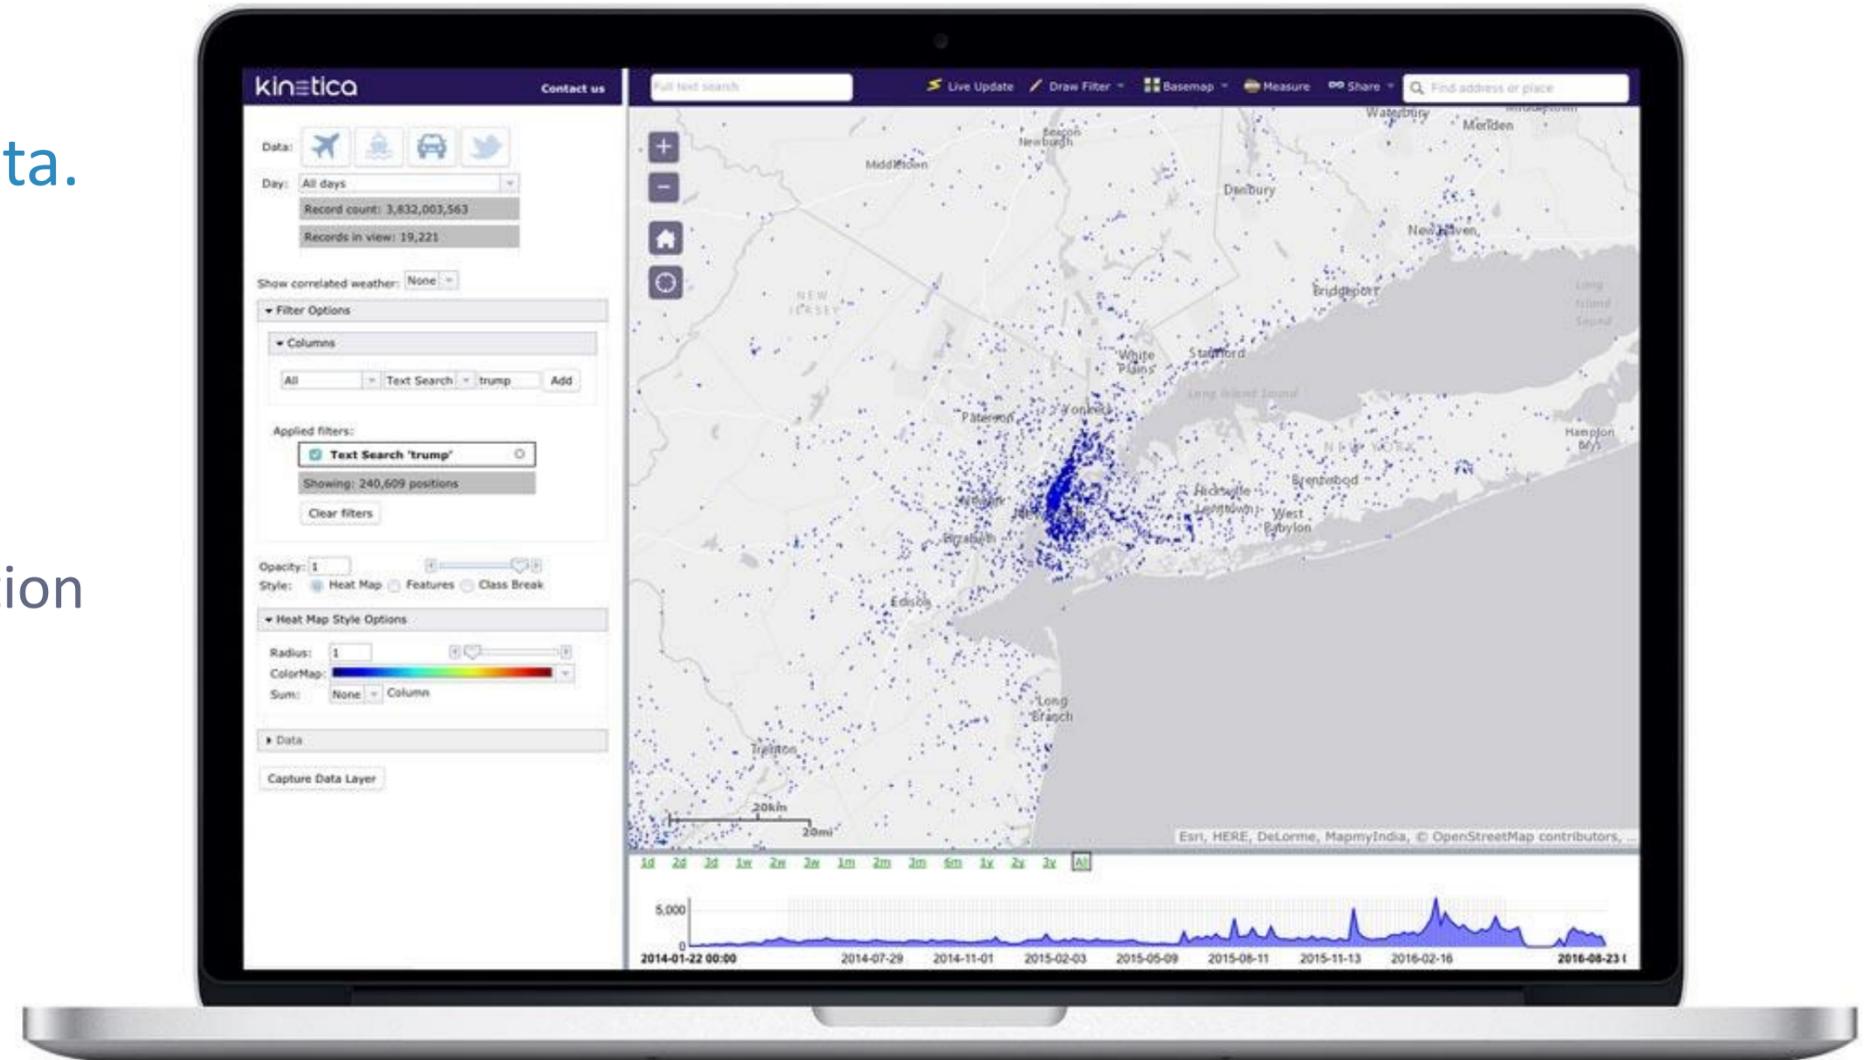
Live streaming analytics on component functionality to ensure safety (avoid failures) and validate warranty claims .

Distributed Visualization Pipeline

Kinetica native visualization ideal for fast moving, location-based IoT data.

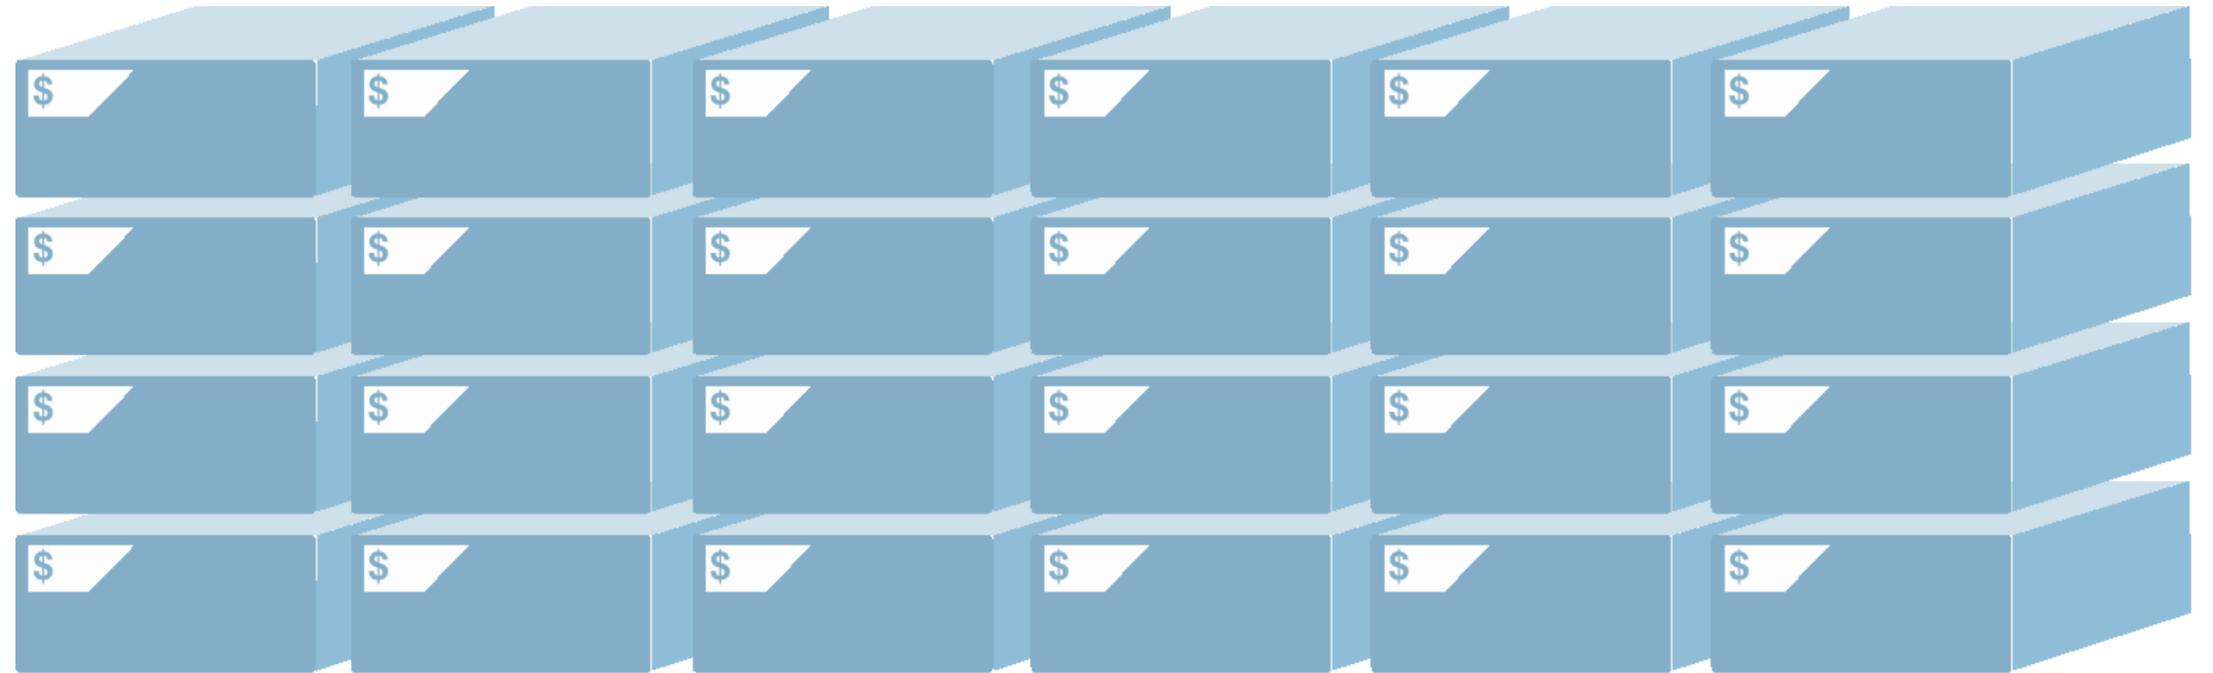
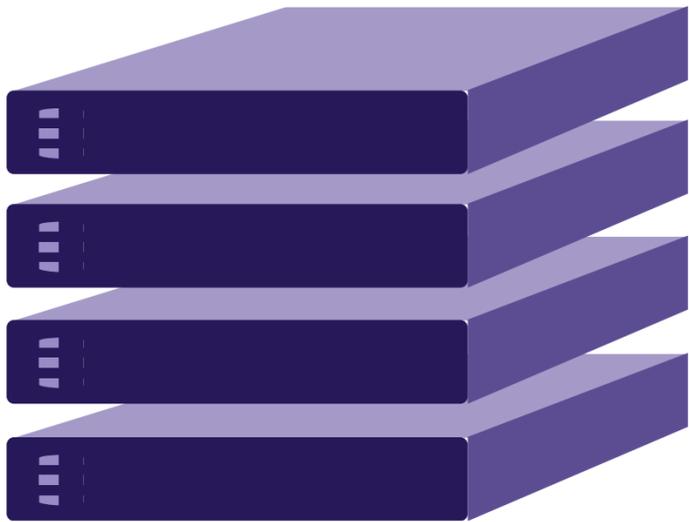
Improving visual rendering of data—
particularly geospatial and temporal data.

- Visualizes billions of data points
- Updates in real time
- Delivers full gamut of geospatial visualization
- Native GIS & IP address object support



Efficiency = Savings

Hardware costs and footprint are a fraction of standard in-memory databases.



Availability

Certified to run on premise with:



Or in the cloud:



Coming soon:



Google Cloud Platform

Accelerated by:



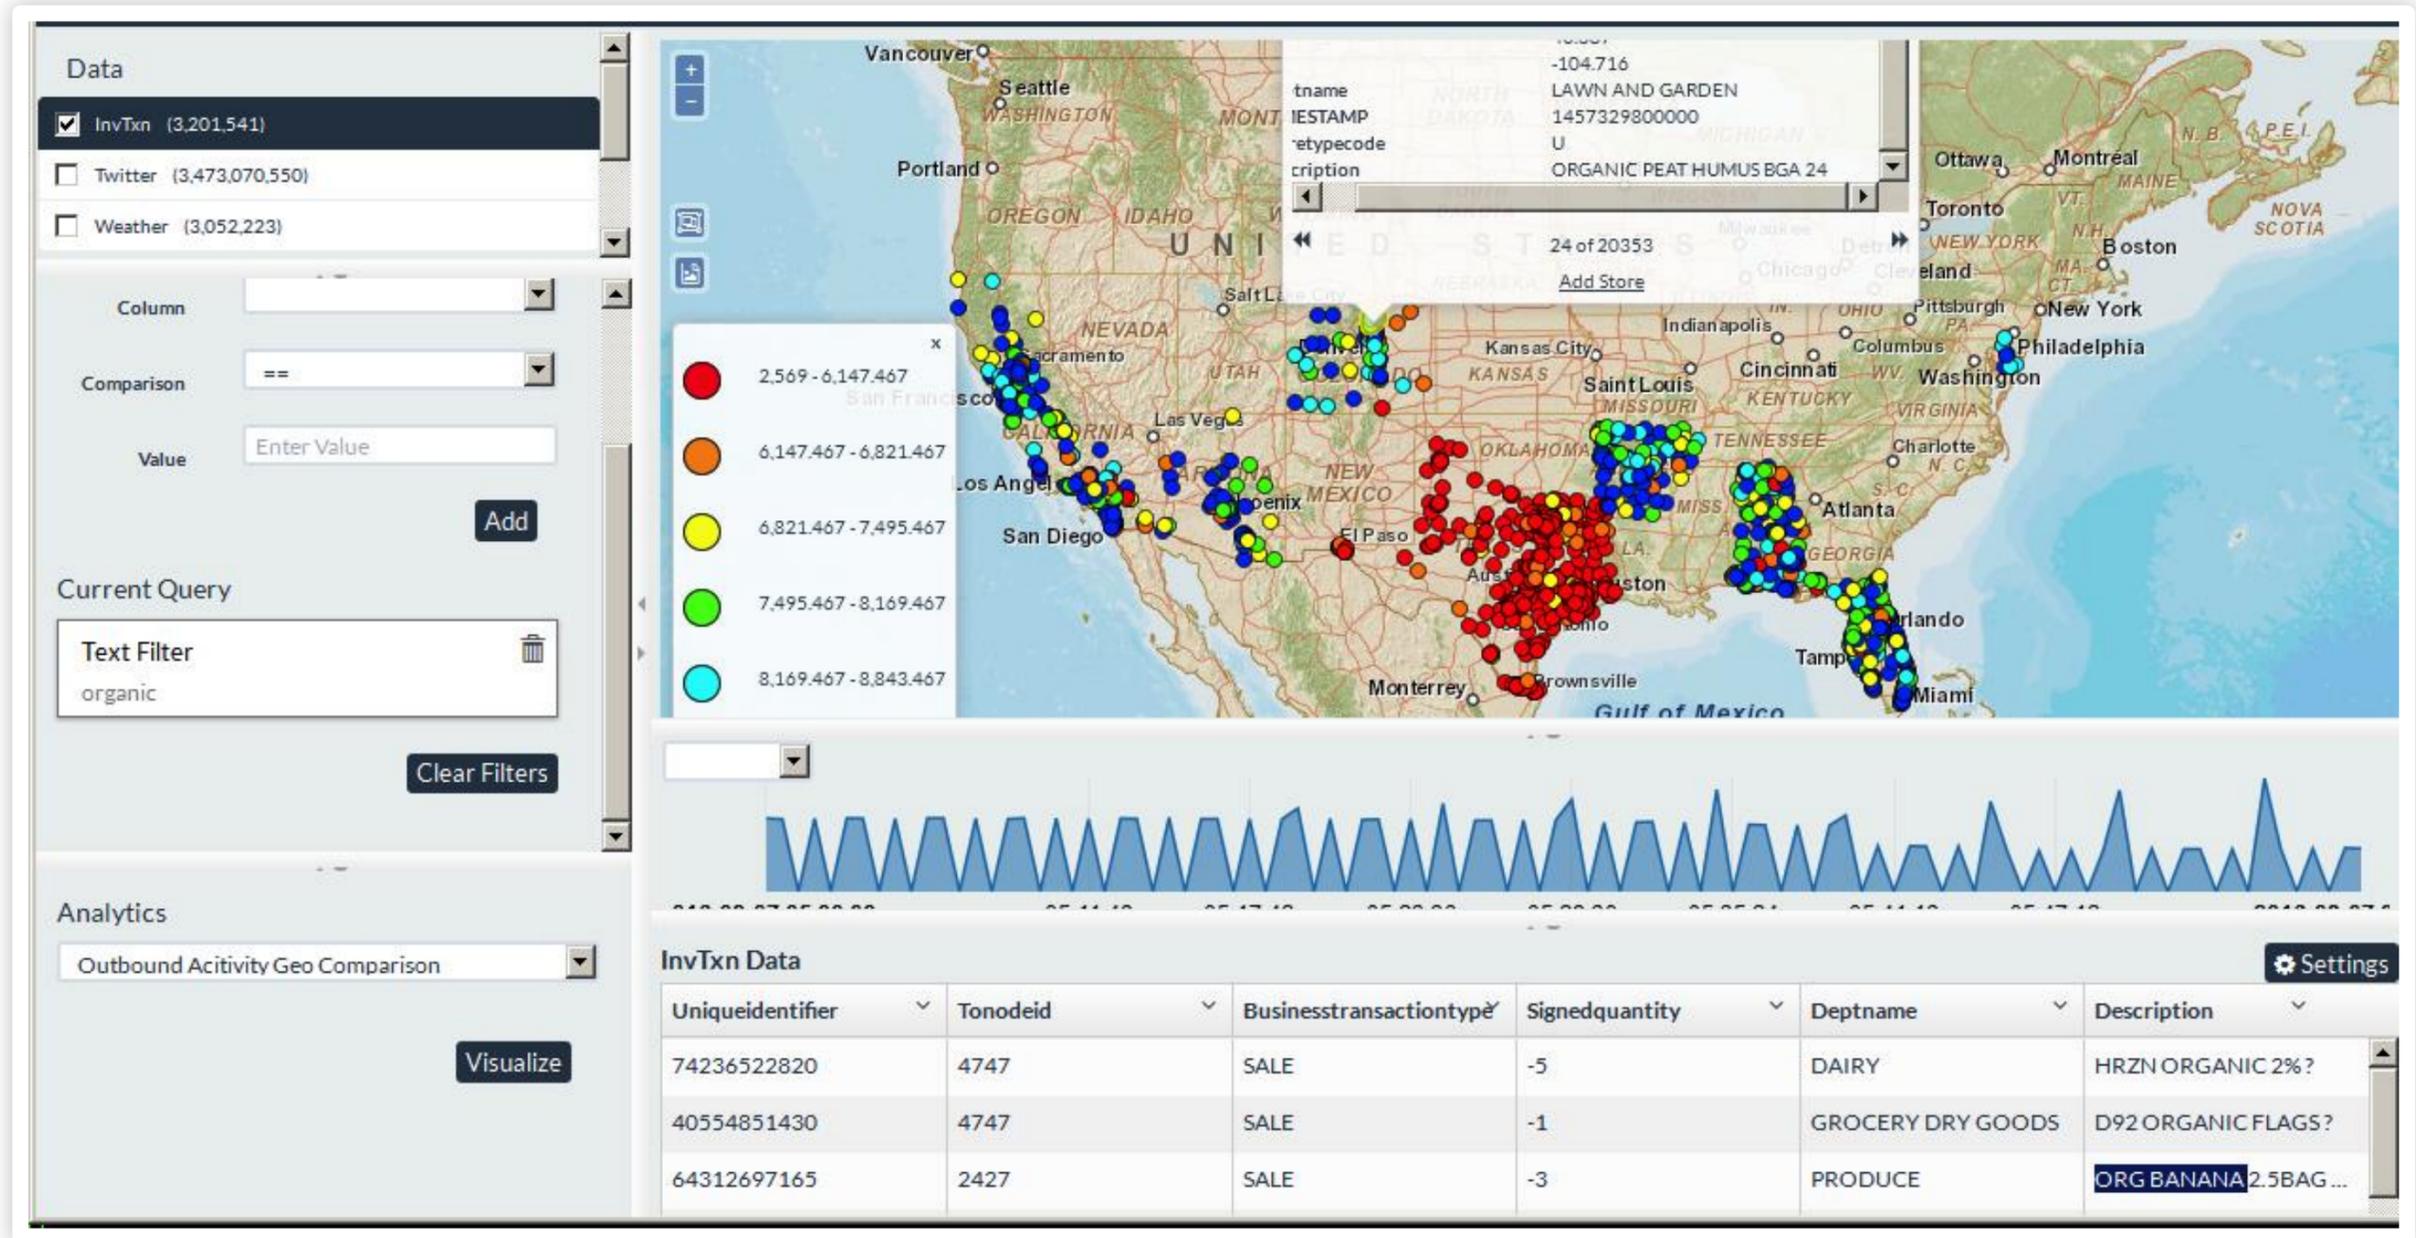
kinética

Thank you

Retail - Redesign Business Without Redesigning Your Data

Multidimensional query on “Organic” with outgoing sales on map

Outgoing sales of organic products in Texas is far less compared to California and Florida



CASE STUDY

US Army INSCOM

US Army's in-memory computational engine for any data with a geospatial or temporal attribute for a major joint cloud initiative within the Intelligence Community (IC ITE).

Intel analysts are able to conduct near real-time analytics and fuse SIGINT, ISR, and GEOINT streaming big data feeds and visualize in a web browser.

First time in history military analysts are able to query and visualize billions to trillions of near real-time objects in a production environment.

Major executive military and congressional visibility.

U.S Army INSCOM Shift from Oracle to GPUdb



The image shows a world map with numerous data points overlaid, representing geospatial data. Below the map, there are two server configurations being compared. On the left is a single GPUdb server, and on the right is a tall rack of 42 Oracle Spatial servers. The GPUdb server is significantly smaller and more compact than the Oracle Spatial server rack.

GPUdb (20ms)	Oracle Spatial (92 Minutes)
42x Lower Space	
28x Lower Cost	
38x Lower Power Cost	

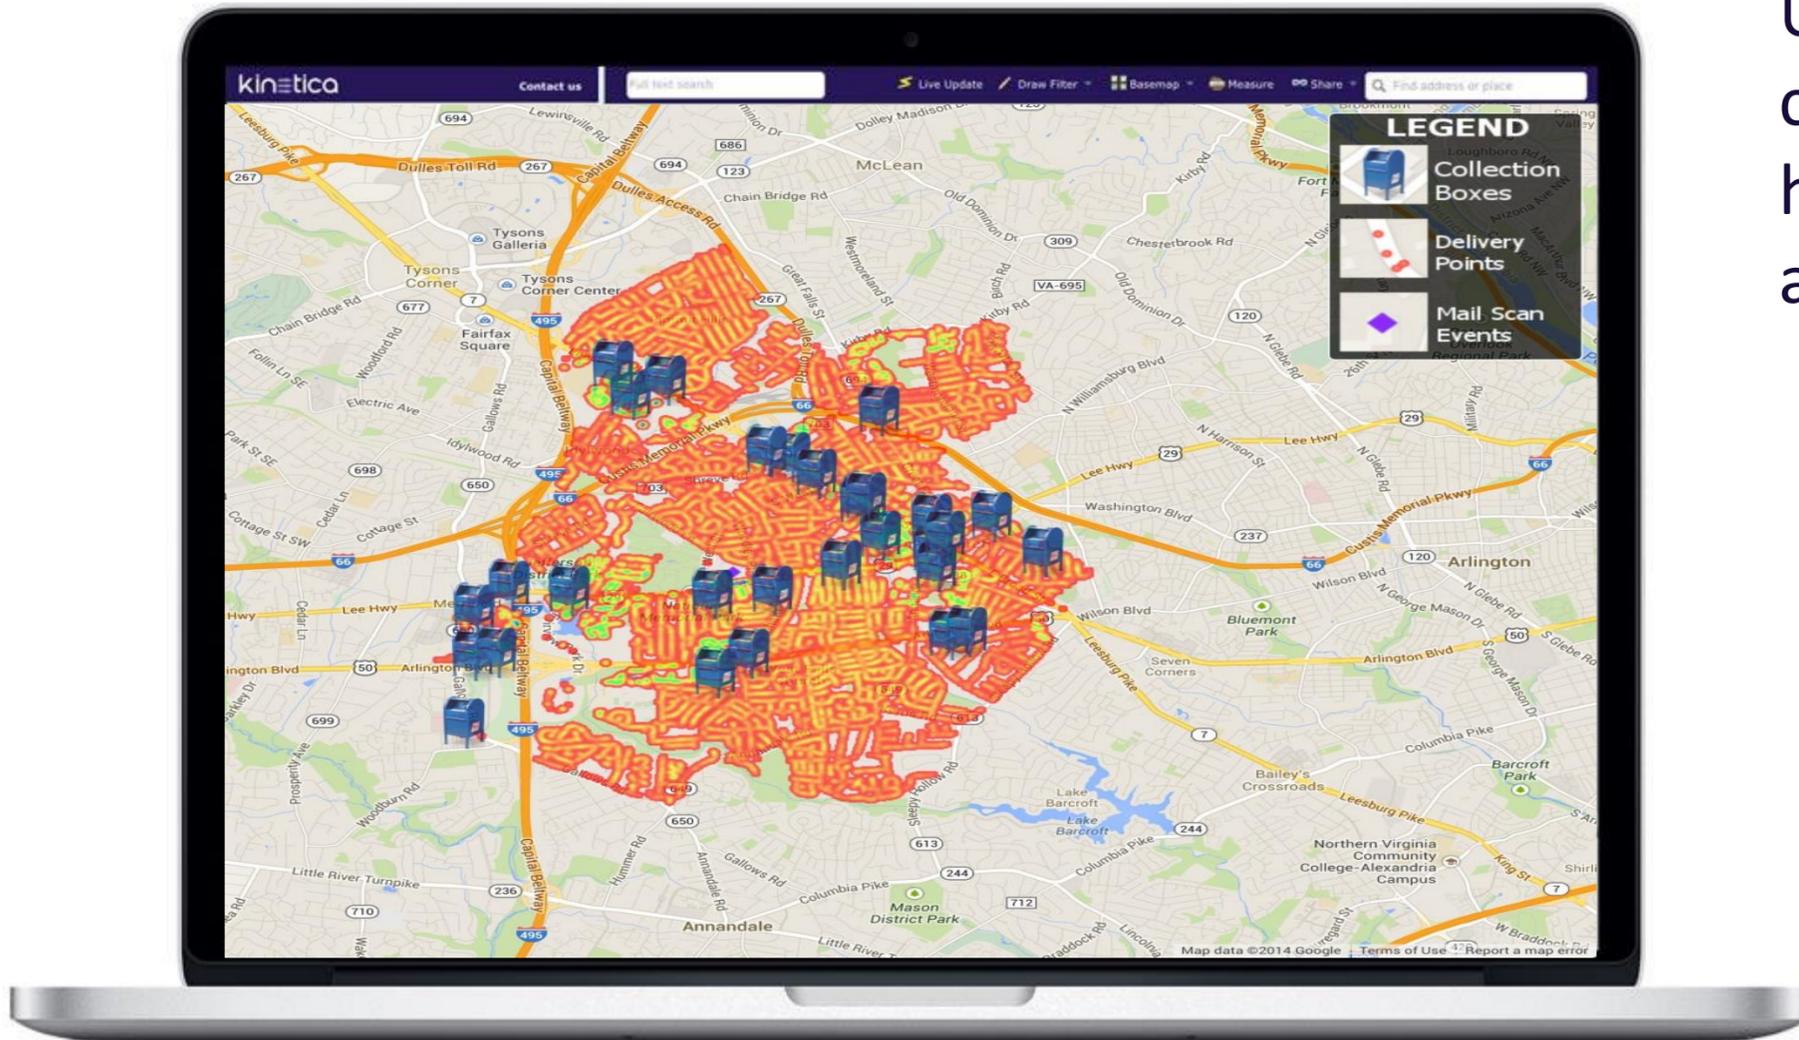
1 GPUdb server vs 42 servers with Oracle 10gR2 (2011)



Real-time route optimization



USPS is the single largest logistic entity in the country, moving more individual items in four hours than the combination of UPS, FedEx, and DHL move all year.



DISTRIBUTED ANALYSIS

USPS' parallel cluster is able to serve up to 15,000 simultaneous sessions, providing the service's managers and analysts with the capability to instantly analyze their areas of responsibility via dashboards.



AT SCALE

With 200,000 USPS devices emitting location once every minute, that amounts to more than a quarter billion events captured and analyzed daily... tracked on 10 nodes.

