THE COMING WAVE OF LOCATION-BASED SERVICES

By Graham Neray

What if the President could view all military assets deployed in Afghanistan today in a single dashboard? What if he could visualize, manage, and analyze feeds from video, equipment sensors, and intelligence reports in real time?

What if the head of the TSA could track and monitor feeds — surveillance, security abnormalities, no-fly passengers — from every U.S. airport and leverage real-time geospatial analytics to prevent the next terrorist attack?

What if U.S. Intelligence could inspect and aggregate IP addresses, geospatial tags, and log files to identify international cyber criminals?

While the U.S. Government has been working with geospatial data for decades, a wave of innovation across wireless, sensor, and data management technologies is unearthing a slew of new applications that take advantage of this data. Declining mobile chipset costs have lowered smartphone prices, making them affordable for the masses. Sensor prices continue to drop, and they can now be embedded in anything from cars to smart watches. The growth of smartphone adoption, with over 60 percent penetration in the U.S., has created demand for innovative apps.[1] The ubiquity of high-speed connectivity over wireless networks makes it feasible to deliver bandwidth-intensive apps on a mobile device. Verizon's 4G LTE network, for instance, now covers 95 percent of the U.S. population.[2] Lastly, new data management technologies — alternatives to the relational database — that are more flexible, scalable, and can accommodate diverse data types in a single repository enable new applications on the backend.

Invasion of the Apps

Much of the innovation in this area is happening in the consumer market. Google Glass incorporates sensor and geospatial data in real time to provide augmented reality. Uber's mobile cab-hailing service is turning the taxi industry upside down. Nike's FuelBand tracks users' movements throughout the day — including running, walking, basketball, dancing, and other everyday activities — then calculates calories burned and visualizes progress in a dashboard.

The ability to incorporate location-based or geospatial data into mobile and desktop applications changes what's possible, not only for consumers, but for enterprises and government organizations as well.

The consumer market has, however, been the driving thrust behind location-based apps for several reasons. One reason is that these apps require users to transmit data over wireless networks, which in many cases are not secure enough for government entities. Perhaps more importantly, traditional relational databases — the bread and butter of government IT — are poorly-equipped to provide the rich, interactive experiences users have come to expect. This is why most location-based services are based on alternatives to relational databases, like MongoDB.

When it comes to location-based services, relational databases fall short in a number of ways.

Geospatial Data at Scale With the proliferation of smartphones and other mobile devices, sensors, and machine-generated data, the volume of information that is geocoded and suitable for location-based services is massive. Geospatial capabilities are mature in relational databases, but making these systems scale to large data volumes and numbers of users while maintaining performance at a reasonable cost has been challenging.

Data Diversity The modern location-based app must incorporate a variety of data types: unstructured data, semi-structured data, polymorphic data, and of course, geospatial data. The underlying database, therefore, must have a data model that is flexible enough to support a myriad of data types in a single data store. The traditional relational database, however, was built to support relatively static applications with highly structured data. It flattens data into two-dimensional tabular structures of rows and columns, enforced by a rigid schema, making it challenging or even impossible to support the requirements of leading-edge geospatial apps.

Continuous Global Availability The relational database of yesteryear lived in a single data center. If the database went down in the middle of the night, someone woke up the database administrator, who would fix the problem as soon as she could. By contrast, modern applications need to be available and replicated in real time across the world, and downtime is unacceptable.

How We Build and Run Apps Today

To address these shortcomings, a number of companies are foregoing the relational database in favor of MongoDB, an open-source NoSQL database.

Unlike relational databases, MongoDB is designed for how we build and run applications today. Instead of storing data in rows and columns as one would with a relational database, MongoDB stores a binary form of JSON documents (BSON). Relational databases have rigid schemas that must be defined up front. By contrast, documents in MongoDB have flexible, dynamic schemas that can vary across documents and can evolve easily over time. This model makes it possible to represent rich objects with nested documents and arrays in a way that is more natural and that mirrors the world it's intended to model. Further, the document model makes it easy to incorporate disparate types of data in a single data store and in real time, such as feeds from video, equipment sensors, and intelligence reports.

MongoDB also has native auto-sharding, which allows users to scale from single server deployments to large, complex multi-data center architectures. It also provides built-in replication with automated failover to enable high availability. Lastly, MongoDB has a rich set of geospatial capabilities, including GeoJSON support, geospatial indexes, and geospatial analytics using the Aggregation Framework with circle, line, rectangle, and polygon intersections.

Industry Adoption of MongoDB

Here are a few examples of how enterprises and government bodies are using MongoDB to transform how they use geospatial data:

City of Chicago The City of Chicago is using MongoDB to cut crime and improve municipal services by collecting and analyzing geospatial data in real time from over 30 different departments. For instance, in a given area, it might evaluate the number of 911 calls and complaints, broken lights, stolen garbage cans, liquor permits, and abandoned buildings, determining that an uptick in crime is likely (see Figure 1). It needs to marry structured and unstructured data at scale and to conduct in-place, online analysis. With legacy technologies, this would be challenging at best, infeasible at worst.

Top 5 Global Telephone Company A mobile operator is using MongoDB to monetize underused legacy data from wireless towers. Like many telephone companies, this operator collects data on the locations of its customers. Rational network investment hinges on knowing which cell sites require more capacity, or where more cell sites are needed. This operator uses MongoDB to power a new service through which a business can push location-specific offers to the company's subscribers in real time when they are in the vicinity of that business. This provides a new revenue stream for the telephone company and brings its legacy systems to life.

Top 5 Industrial Equipment Manufacturer A top industrial equipment manufacturer is using MongoDB to power a cloud-based analytics platform that ingests, stores, and analyzes readings (e.g., temperature, location) from its customers' equipment. It then presents the readings back to customers via a web interface — including visualization, key metrics, and time series

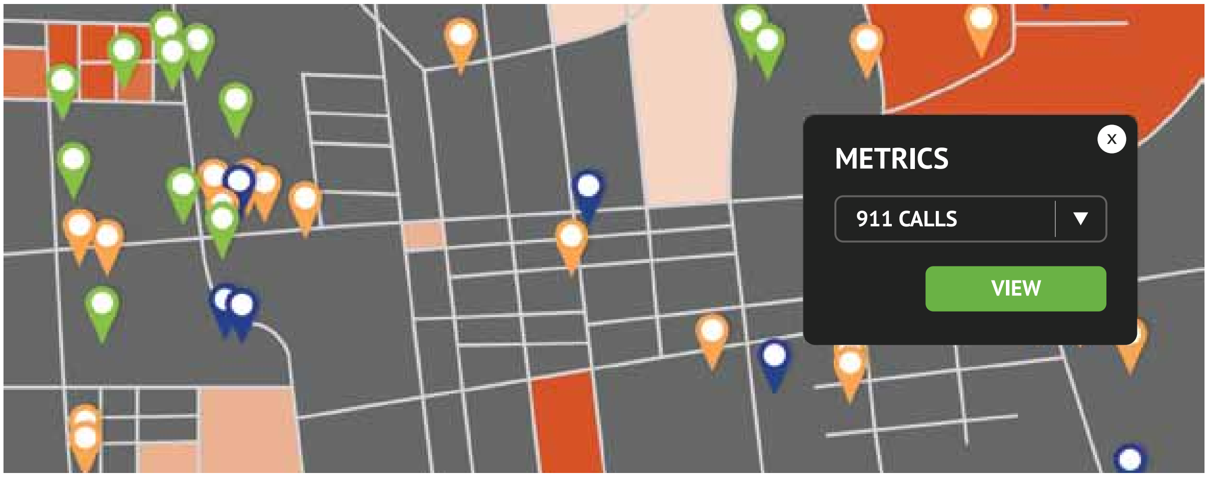

Figure 1 | The City of Chicago is using MongoDB to cut crime and improve municipal services by collecting and analyzing geospatial data in real-time from over 30 different departments.

analysis — to help them make better decisions about their businesses, such as where to provision equipment and how to increase facility efficiency. Additionally, the company uses the data to inform internal product development. The application stands out in an industry that has seen little innovation in the last half-century with the ability to drive new revenue streams.

Leveraging Location in the U.S. Government

Advances in geospatial capabilities and database scalability enable a number of unique applications for the defense and intelligence communities.

Military Asset Analytics The military owns and manages fleets of tanks, drones, Humvees, and other assets. While it may currently have the ability to see where these assets are at a given point, the military could expand its understanding of how the equipment is used, and furthermore, optimize its usage. MongoDB's scalable model and native geospatial capabilities enable the military to ingest and analyze sensor data to optimize vehicle and platoon deployment, fuel consumption, and other metrics, both in real time and retrospectively. MongoDB's flexible data model makes it possible to do all this in a single data store and to provide a 360-degree view of military assets.

Surveillance Data Aggregation Government agencies are exploring new ways to increase national security, including innovative surveillance data collection. This data is pouring in from a variety of sources and in massive volumes. MongoDB provides an adaptable and scalable platform for aggregating various surveillance feeds and making sense of the data through real-time analyses.

Crime Data Management and Analytics Legacy criminal record systems based on relational databases are often brittle and difficult to adapt to legal and regulatory changes. MongoDB makes it easy and cost-effective to adapt the technology as the law evolves and to incorporate information from disparate sources. Law enforcement organizations can deploy innovative, real-time analytics to identify offender-specific, geospatial, and other crime patterns quickly and effectively.

Government is increasingly faced with growing data volumes and shrinking budgets, new demands, and legacy infrastructure. Given the rise of innovative geospatial applications in the market and the advent of open source technologies, solutions are available that will allow organizations to make the most of their data.

__Graham Neray__ is a Product Marketing Manager at 10gen, an IQT portfolio company. Neray works closely with customers, partners, and the open-source community to articulate how MongoDB is quickly becoming the world's most popular database. Prior to joining 10gen, he was a Senior Business Analyst at CSMG, a boutique management consulting firm specializing in the high-tech and telecom industries. He graduated from Brown University with a concentration in Political Science.

REFERENCES

[1] "Mobile Majority: U.S. Smartphone Ownership Tops 60%." *Nielsen Newswire.* June 6, 2013.
[2] Fitchard, Kevin. "Verizon wraps up LTE rollout; plans all-VoIP phone launch for late 2014." *GigaOm.* June 27, 2013.