elastic

Elastic in action in Public Sector

How organizations are using Elasticsearch to drive mission outcomes in government, defense, education, and healthcare.

Table of contents

Using the power of Elastic to serve the public

One of the most frequent questions we get when we visit public sector customers or attend trade shows is whether we have any examples of creative ways that their peers are using Elasticsearch in their organizations. And there are: Our public sector community around the world is using Elasticsearch in so many innovative ways to serve the greater good. (Of course, because of the sensitive nature of many public sector missions, we unfortunately can't share them all.) But this guide provides eight examples of how organizations are using Elastic's search, analysis, and data-visualization capabilities to move their missions forward.

The Elastic Public Sector Team

Improving patient outcomes through data patterns: **CogStack at NHS**

About CogStack:
CogStack is a data storage and analysis platform, powered by Elasticsearch and Language AI, that is used by several leading UK hospitals that together have handled more than 200 million clinical documents.

The challenge:
Patient records are stored in different databases and different formats (handwritten clinician notes, image scans, numerical data, just to name a few), making it difficult for clinicians to have a complete view of patient data – and impossible to detect larger trends in population data.

The solution:
A team at King's College Hospital, London, built the CogStack platform, which consolidates patient information from multiple sources into a single data warehouse. An advanced search layer, powered by Elasticsearch, enables clinicians to extract meaningful information from data, no matter what format it's in. Clinicians "can search in the way that they write and talk, rather than having to use coding or specialist data skills," says Professor James Teo.

The outcomes:
By using Elasticsearch as a key component of the CogStack platform, clinicians are now able to:

- Study decades of data in order to identify and correlate health and treatment patterns
- Have a comprehensive view of a patient's past information in order to prescribe accurate medication
- Prioritize at-risk patients for immediate medical care
- Quickly respond to large-scale medical alerts, such as medication recalls and drug interactions

In the case of a major medical safety issue, CogStack and Elastic enabled us to find all the relevant medical records in just one afternoon, when previously it would have taken months or even years. As a result of our efforts, thousands of women were able to avoid taking a medication that could have harmed their unborn babies.

Professor James Teo
Clinical Director of AI and Data, Professor of Neurology
King's College Hospital and Guys & St Thomas Hospitals

Read the full story

Securing distributed data via a unified data layer: **CISA's CDM program**

About CISA:
In the US, the Cybersecurity and Infrastructure Agency (CISA) operates a central dashboard to uncover and share cybersecurity vulnerabilities across over 100 federal agencies. This dashboard, known as the Continuous Diagnostics and Mitigation (CDM) dashboard, handles massive quantities of data across localized data systems.

The challenge:
With cybersecurity threats on the rise, CISA wanted insight into patterns and threats, so it could provide support, insight and remediation to agencies when needed. CISA needed access to sensitive agency data without owning it – or moving it or replicating it from its original location with individual agencies.

The solution:
CISA worked closely with ECS, an Elastic Premier partner, to build out a centralized dashboard, which includes Elasticsearch as an essential component. One of the critical capabilities that Elastic provides is Cross-Cluster Search, or the ability for data to be searched no matter where it's located – even across cloud and on-prem environments. This way, CISA can quickly access agency data when needed, without agencies having to replicate or give up control of their internal data.

The outcomes:
CISA is meeting its CDM adoption goals among federal agencies. According to CISA Executive Assistant Director for Cybersecurity Eric Goldstein, CISA is able to "gain real time visibility into every asset, every vulnerability, every misconfiguration, across every federal civilian network continuously."

The CDM toolset has come in handy during every cyber threat and incident agencies have faced over the last five years. Whether it was the WannaCry ransomware attack or Log4J or any number of threats, agencies and CISA can turn to the dashboard from Elastic to discover more complete data more immediately.

Federal News Network, 2022

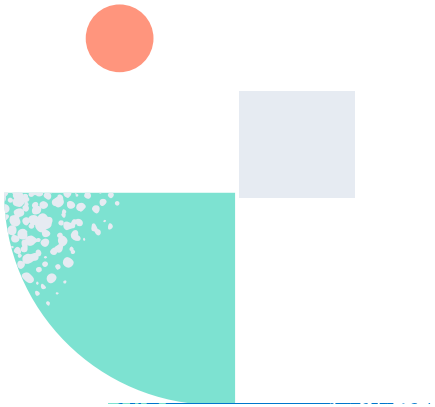

Read the full story

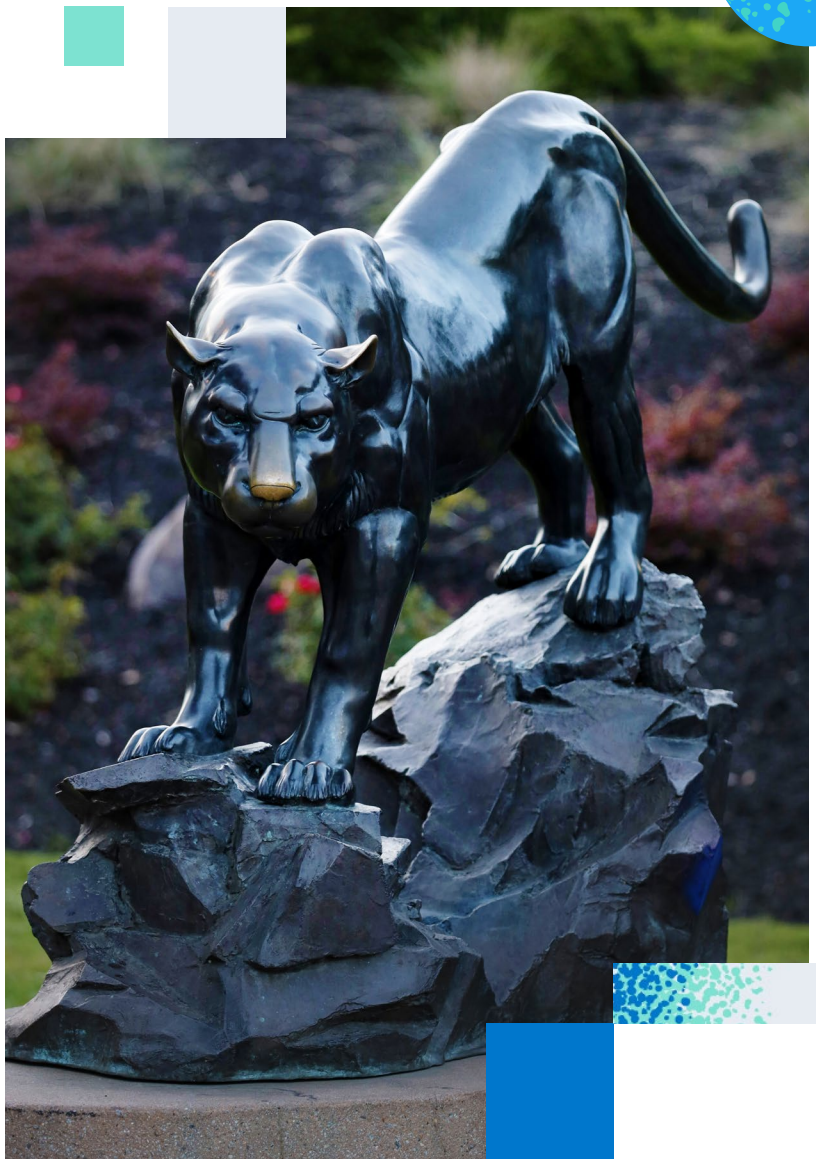

Connecting students to financial aid with generative AI: **Georgia State University**

About Georgia State University:
Georgia State University is a US public research university in Atlanta, Georgia. Founded in 1913, it is the largest institution of higher education by enrollment based in Georgia and one of the largest in the US, with a student enrollment of around 50,000, including approximately 33,000 undergraduate and graduate students at the main campus downtown.

The challenge:
The Georgia State IT team wanted to begin using generative AI at the university. They chose financial aid as its pilot use case, due to the large volumes of financial aid data available. In order to apply for financial aid, students had to comb through many disparate resources, in different formats and from multiple sources, that often didn't meet their specific situations.

The solution:
The team used Elasticsearch to set up a proof-of-concept for generative AI at the university, using the financial aid use case as its pilot. Financial aid data from public sources was ingested into Elastic, turned into vectors (numerical values), allowing it to be searched semantically (or, in a user's natural language instead of through a specific keyword).

The team then used Elastic to set up a retrieval augmented generation (RAG) workflow that adds critical context to the publicly available financial aid data via the university's own proprietary data. This data supplies the generative AI model with contextual information as well as accompanying instructions to the gen AI model on how that information should be factored into the output.

The outcomes:
Now, when students visit the university's financial dashboard, they see personalized content reflecting their current status as well as information or alerts about external financial aid resources relevant to their status. This success of this pilot is now serving as a viable AI model for the university to build upon for additional use cases.

Read the full story

The financial aid process can be complicated as students try to understand whether they qualify for scholarships and grants. Students who need assistance but can't find funding may be less likely to complete their college degrees.

Jaroslav Klc
Director of Strategic Initiatives and Development for IT
Georgia State University

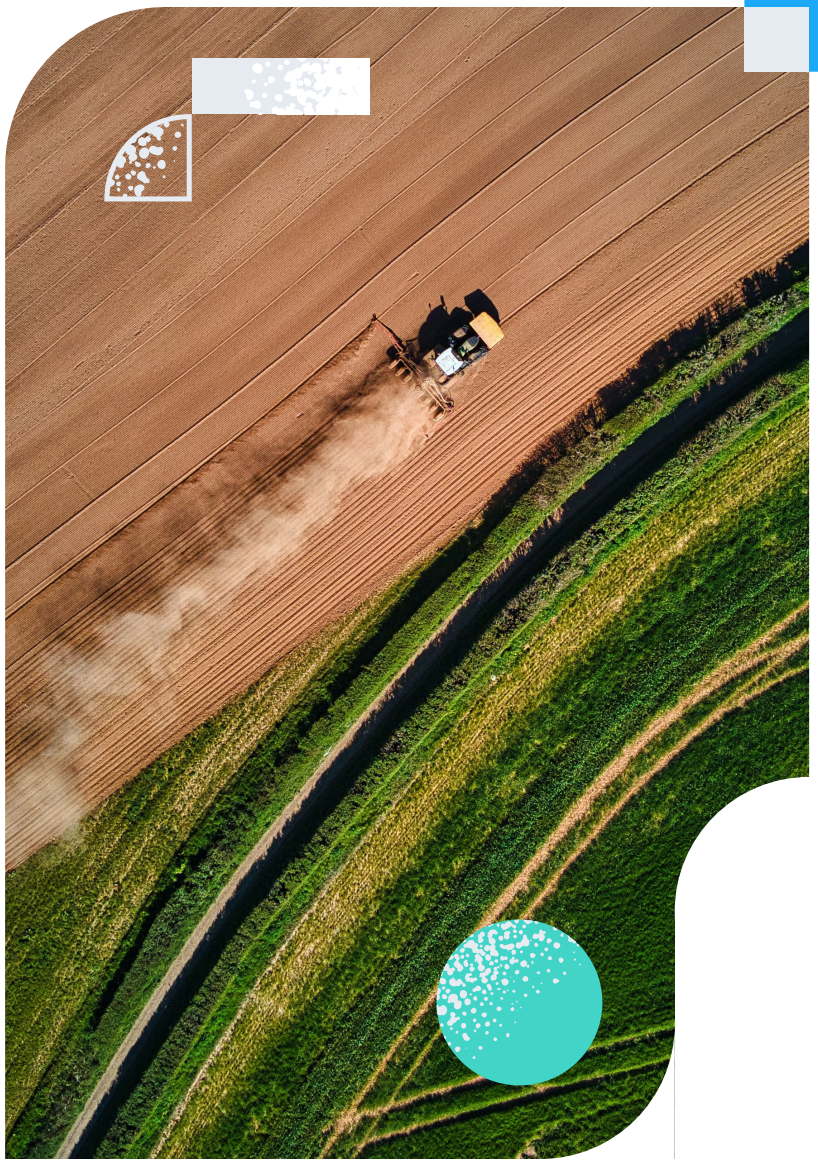

Harnessing real-time insights for commercial fishing: **The French Ministry of Agriculture**

About the French Ministry of Agriculture & Food:
Within the French Ministry of Agriculture and Food, a team of architects in the Methods, Support and Quality office (BMSQ) aims to improve the quality, speed, and precision of how they collect and analyze large volumes of data connected to the commercial fishing industry — such as declared fish hauls, harbor exit manifests, and GPS data tracking vessel locations.

The challenge:
The team needed the speed and flexibility to provide up-to-date data that was verified, complete, dynamic, and that could be viewed in different ways depending on the target audience.

The solution:
The Elasticsearch platform ingests all data types and renders visualizations so the team can share and use real-time data about commercial fishing activity. With this integrated data, the team can pinpoint quantities and species fished, track registration and equipment, take enforcement action, stop illegal fishing, and negotiate fishing rights with neighboring countries.

The outcomes:
Leveraging Elasticsearch, the French Ministry of Agriculture can:

- Easily visualize data, with the ability to refine the granularity, quality and format to ensure the fishery data is presented in the most suitable format for the audience, especially a non-technical audience
- Track GPS-enabled boats, while at the same time indexing and monitoring that data into Elasticsearch
- Store and access 10+ years of historical data, amounting to 135 million records in Elasticsearch
- Use real-time data by region to immediately remedy any law infringements and react to media controversy related to protected species

Read the full story

Elastic handles ingesting all of our data and render visualizations to make it easy to share and use real-time information about commercial fishing activity. With this integrated data, we can take enforcement action, stop illegal fishing, and negotiate fishing rights with our neighboring countries.

Sébastien Arnaud
Exchanges & Data Architect
French Ministry of Agriculture

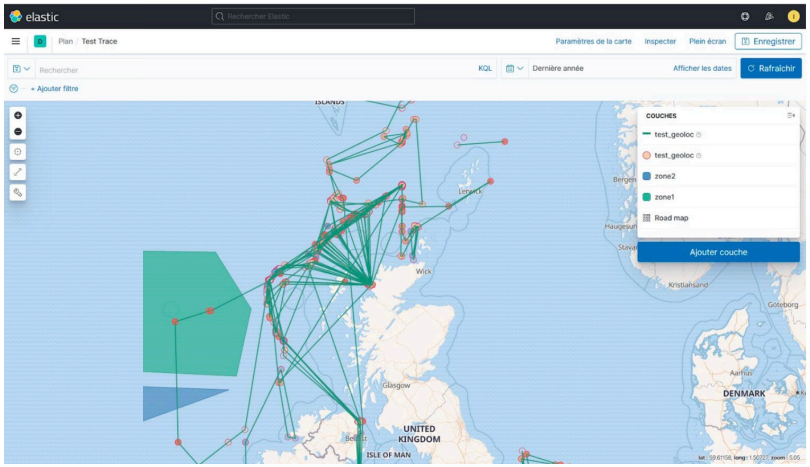

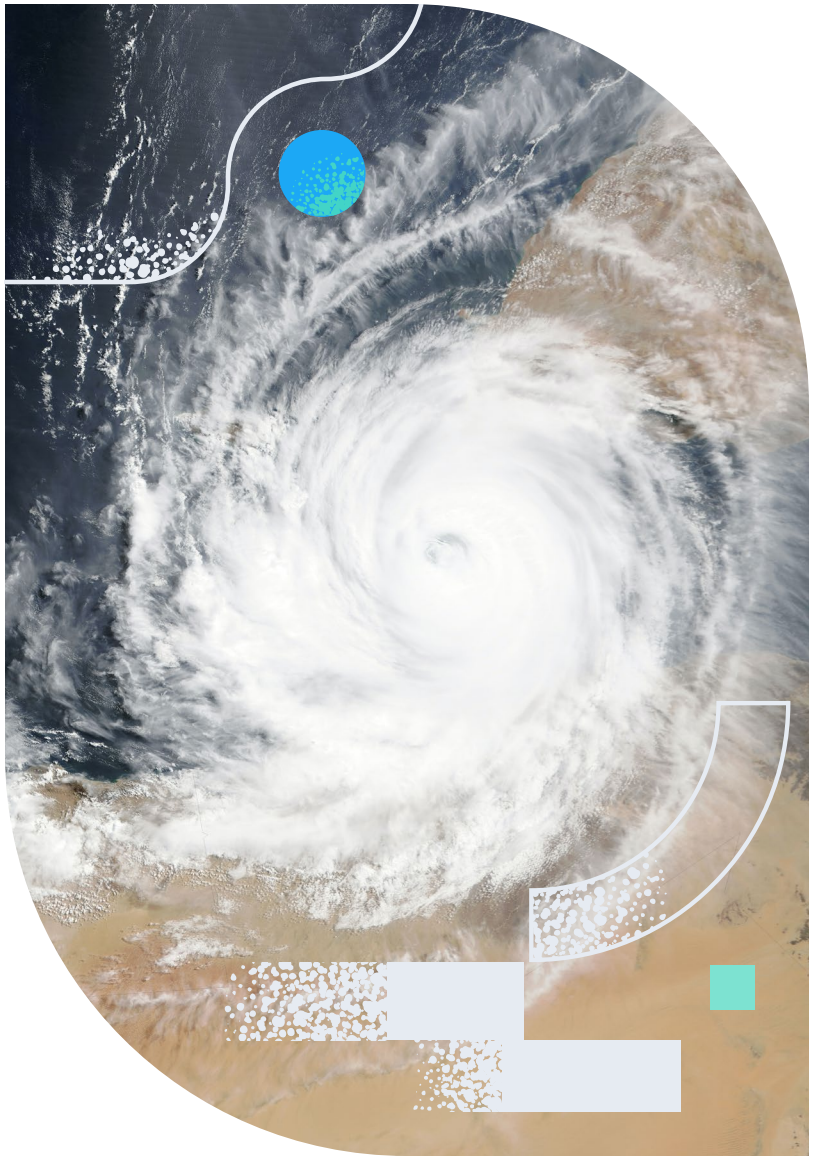

Consolidating tools and data for more accurate weather: **The Met Office**

About The Met Office:
The UK's Met Office uses more than 10 million weather observations and an advanced atmospheric model to create 3,000 tailored forecasts and briefings every day. To make its forecasts, the Met Office uses some of the most advanced technology on the planet, including one of the world's most powerful supercomputers. It also runs a mix of on-premise, cloud, and software-as-a-service platforms.

The challenge:
The Met Office was gathering logs from multiple systems in many locations. This was a fragmented process with different teams managing their own logging systems and storing data in multiple formats.

The solution:
The Met Office created a unified platform where different teams could access the data that they needed from one location. Elastic Cloud (running on AWS and Azure) is an essential part of this platform, allowing the Met Office to scale to over 2 billion data logs each day.

The outcomes:
With Elastic powering their platform, the Met Office can:

- Run searches and correlate insights that speed up the resolution of network issues
- Get a clearer view of its systems, including its National Severe Weather Warning System, which issues alerts to regions affected by extreme conditions
- Build simple search queries to find impactful events on multiple systems
- Boost system security by providing a complete view of system activity and suspicious behavior

The Met Office is a huge consumer of data. By adopting Elastic Cloud running on AWS and Azure we can scale the system to more than two billion data logs every day.

John MacGrillen
Solutions Architect, Core Services
Met Office

Read the full story

Deploying role-based data access: **Enterprise Data Center at Hill Air Force Base**

About HEDC at Hill Air Force Base:
Hill Enterprise Data Center (HEDC) provides hosting services for more than 100 information systems of the US Air Force logistics center at Hill Air Force Base (AFB), Utah. Their private cloud can be linked across more than a dozen separate sites, which they use to communicate with bases around the US – resulting in a lot of data transfer and aggregation.

The challenge:
HEDC frequently dealt with data silos and the inefficacy of “paper compliance.” They needed a way to connect disparate data sources and share that data securely based on clearance level and location.

The solution:
The HEDC team leverages Elasticsearch to ingest all data types and provide unified data visualization and analytics. Elasticsearch's Role-Based Access Control (RBAC) capability provides a consistent, cloud-shared UI that allows remote teams to seamlessly access and corroborate the data available to their specific role and location.

The outcomes:
With Elasticsearch deployed, HEDC can:

- Utilize a streamlined ingestion process, secure data transfer and permissions protocol, and a (compliance mandated) correlated sequence of record
- Customize user privileges by role, providing certain levels of information system access for an engineer, application administration access, or create an access role with specific functions for the information system security manager — a DoD requirement
- Extend RBAC into their PaaS infrastructure deployment. Now analysts in SecOps, operations, hyper_infra, etc. can all have different levels of access to HEDC's vast data store

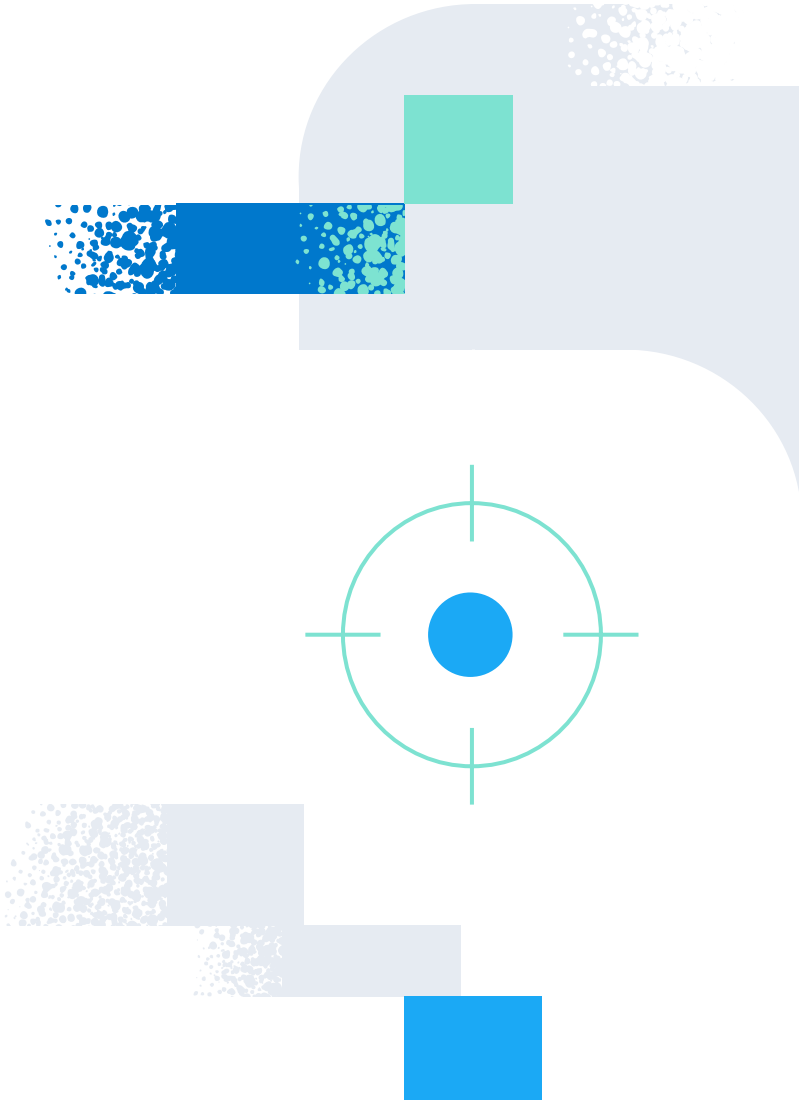

Read the full story

Reducing environmental impact through IoT: **The University of Edinburgh**

About the University of Edinburgh:
The University of Edinburgh, Scotland, is a large institution with expertise in informatics and AI research. The university's IoT Research & Innovation Service focuses on using sensor technologies and multi-modal communications technologies that facilitate efficient operations on its Smart Campus.

The challenge:
With multiple campuses and 600 buildings embedded throughout the city of Edinburgh, the IoT Research & Innovation Service team needed a solution to help them ingest data and maintain data governance from various data sources. The team needed to provide actionable insights to help the university maximize space, meet environmental regulations, and provide a healthy environment for students to learn in.

The solution:
The University of Edinburgh uses the Elasticsearch platform to gather real-time data from its network and systems, including IoT sensors. Their team is able to visualize this data through Elastic's Kibana interface. The university maintains all control of its data and data governance processes, from data ingest (via sensors, routers, etc.) through data processing, to data access.

The outcomes:
With Elasticsearch, the University of Edinburgh can more efficiently perform functions of a Smart Campus and has the ability to:

- Monitor carbon dioxide levels in campus buildings in real-time, detecting patterns that could be detrimental to student learning and long-term health
- Track library occupancy to better connect students to open study spaces
- Map student flow and resource utilization across the university's multiple campuses

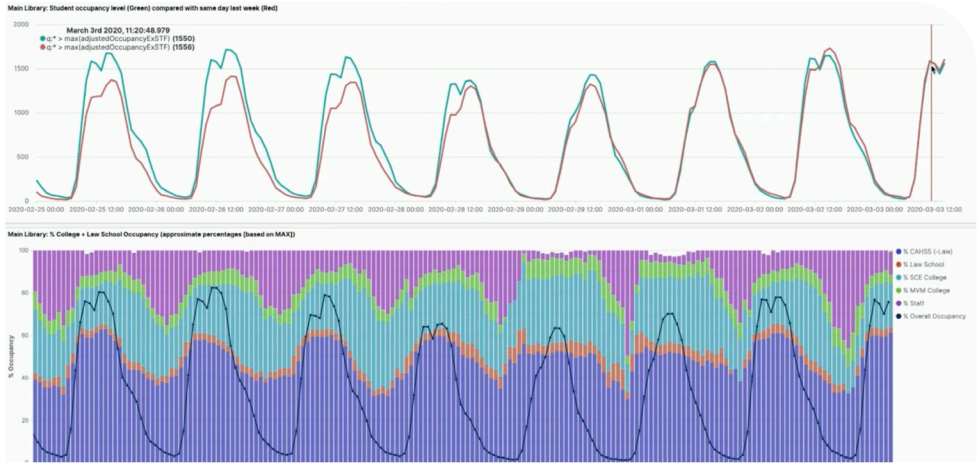

Hear the full story

Helping law enforcement accelerate criminal investigations: **Bluestone Analytics**

About Bluestone Analytics:
Bluestone Analytics, a CACI company, is an international leader in dark web analysis. Its technology suite, DarkBlue Intelligence, enables clients, including national security and intelligence teams, to search open-source intelligence (OSINT) and unveil the identities of criminals operating on the dark web.

The challenge:
Law enforcement organizations tasked with investigating activities such as drug trafficking and weapon sales need access to information that's only found on the quickly-evolving dark web. This data is difficult to access, and directly browsing the dark web exposes a user to risks such as malware, as well as to potentially disturbing content.

The solution:
With the DarkBlue Intelligence suite, powered by Elasticsearch, users can search information from the dark web in a safe, text-based environment without having to download an actual dark web browser. The search, pivot, and indexing capabilities in Elasticsearch provide the ability to link many anonymous personas to one individual with just one click. “It turns a single breadcrumb into a trail that leads to your target,” says Jason Nack, Head of Technology, Bluestone Analytics.

The outcomes:
Using Elasticsearch to power its DarkBlue Intelligence suite, Bluestone Analytics has the ability to:

- Archive data indefinitely and quickly access historical data for a complete view into investigations
- Quickly ingest and incorporate new data sources, ensuring long-term scalability and agility
- Provide potentially life-saving insights to the right people, regardless of how the dark web evolves

The beauty of Elastic is that it saves hours of time that we would otherwise spend installing a specific set of software and connections for different applications. Being able to add documents and data sources with minimum effort is a killer feature for us.

Jason Nack
Head of Technology
Bluestone Analytics

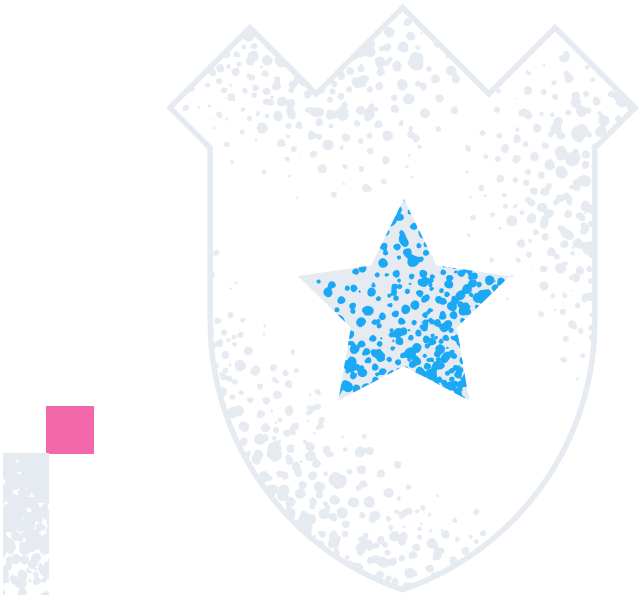

Read the full story

Integrating data into a geographic information system (GIS): **US Army Corps of Engineers**

About the US Army Corps of Engineers:
The NavPortal project team at the US Army Corps of Engineers is tasked with providing safe, reliable, efficient, and environmentally sustainable transport systems for US vessel traffic, which is essential to national security, commerce and recreation. They maintain over 12,000 miles of inland and intracoastal waterways, 13,000 miles of coastal waterways, and 400 ports. Dredge 220 million cubic yards dredged since 2010. Even small innovations can yield tremendous savings (for taxpayers) since working on such a large scale.

The challenge:
Because they operate on a project or mission basis, siloed data and tools are a problem. The sheer volume of data can be difficult to keep track of, including years of historical data and different data types coming from external organizations. The complexity of geospatial data also posed a challenge, especially because there was no framework to compare the disparate data.

The solution:
The Army Corps of Engineers is using the Uber H3 global grid system and integrating it with Elasticsearch, which serves as a seamless platform where all data types can be visualized in one place – and that can leverage APIs to connect to other teams and projects.

The outcomes:
Using Elasticsearch and its geospatial data capabilities, the Army Corps of Engineers can:

- Use hexbins to study and compare geospatial data over time, regardless of data type
- Inform local entities, district offices, and headquarters when there are problems in channels
- Predict issues before they occur by inferring volume in a channel based on data

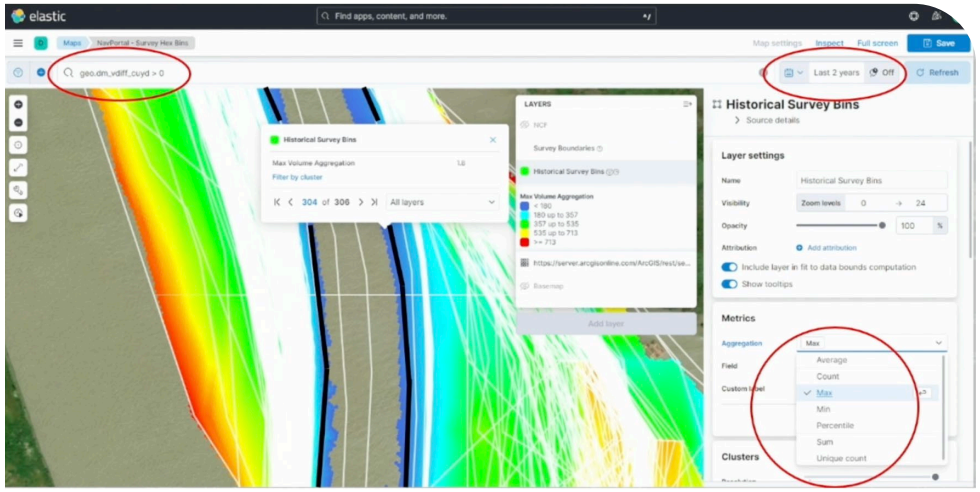

Hear the full story

Summing Up

Elastic's flexible architecture and open foundation lend themselves to the organic growth of a collaborative user community committed to sharing ideas, expertise, and use cases with one another. Because of the agile nature of the Elasticsearch platform, once the technology is implemented and you're ingesting data, you can easily scale to additional use cases without grappling with yet another disparate tech solution.

To learn more about Elastic for public sector, visit our website: elastic.co/industries/public-sector, or contact us directly at elastic.co/contact/publicsector.

We'd love to hear your story.