carahsoft The Trusted Government IT Solutions Provider®

AI Buyer's Guide **for Government**

Discover AI's Dynamic Applications and Transformative Success Stories

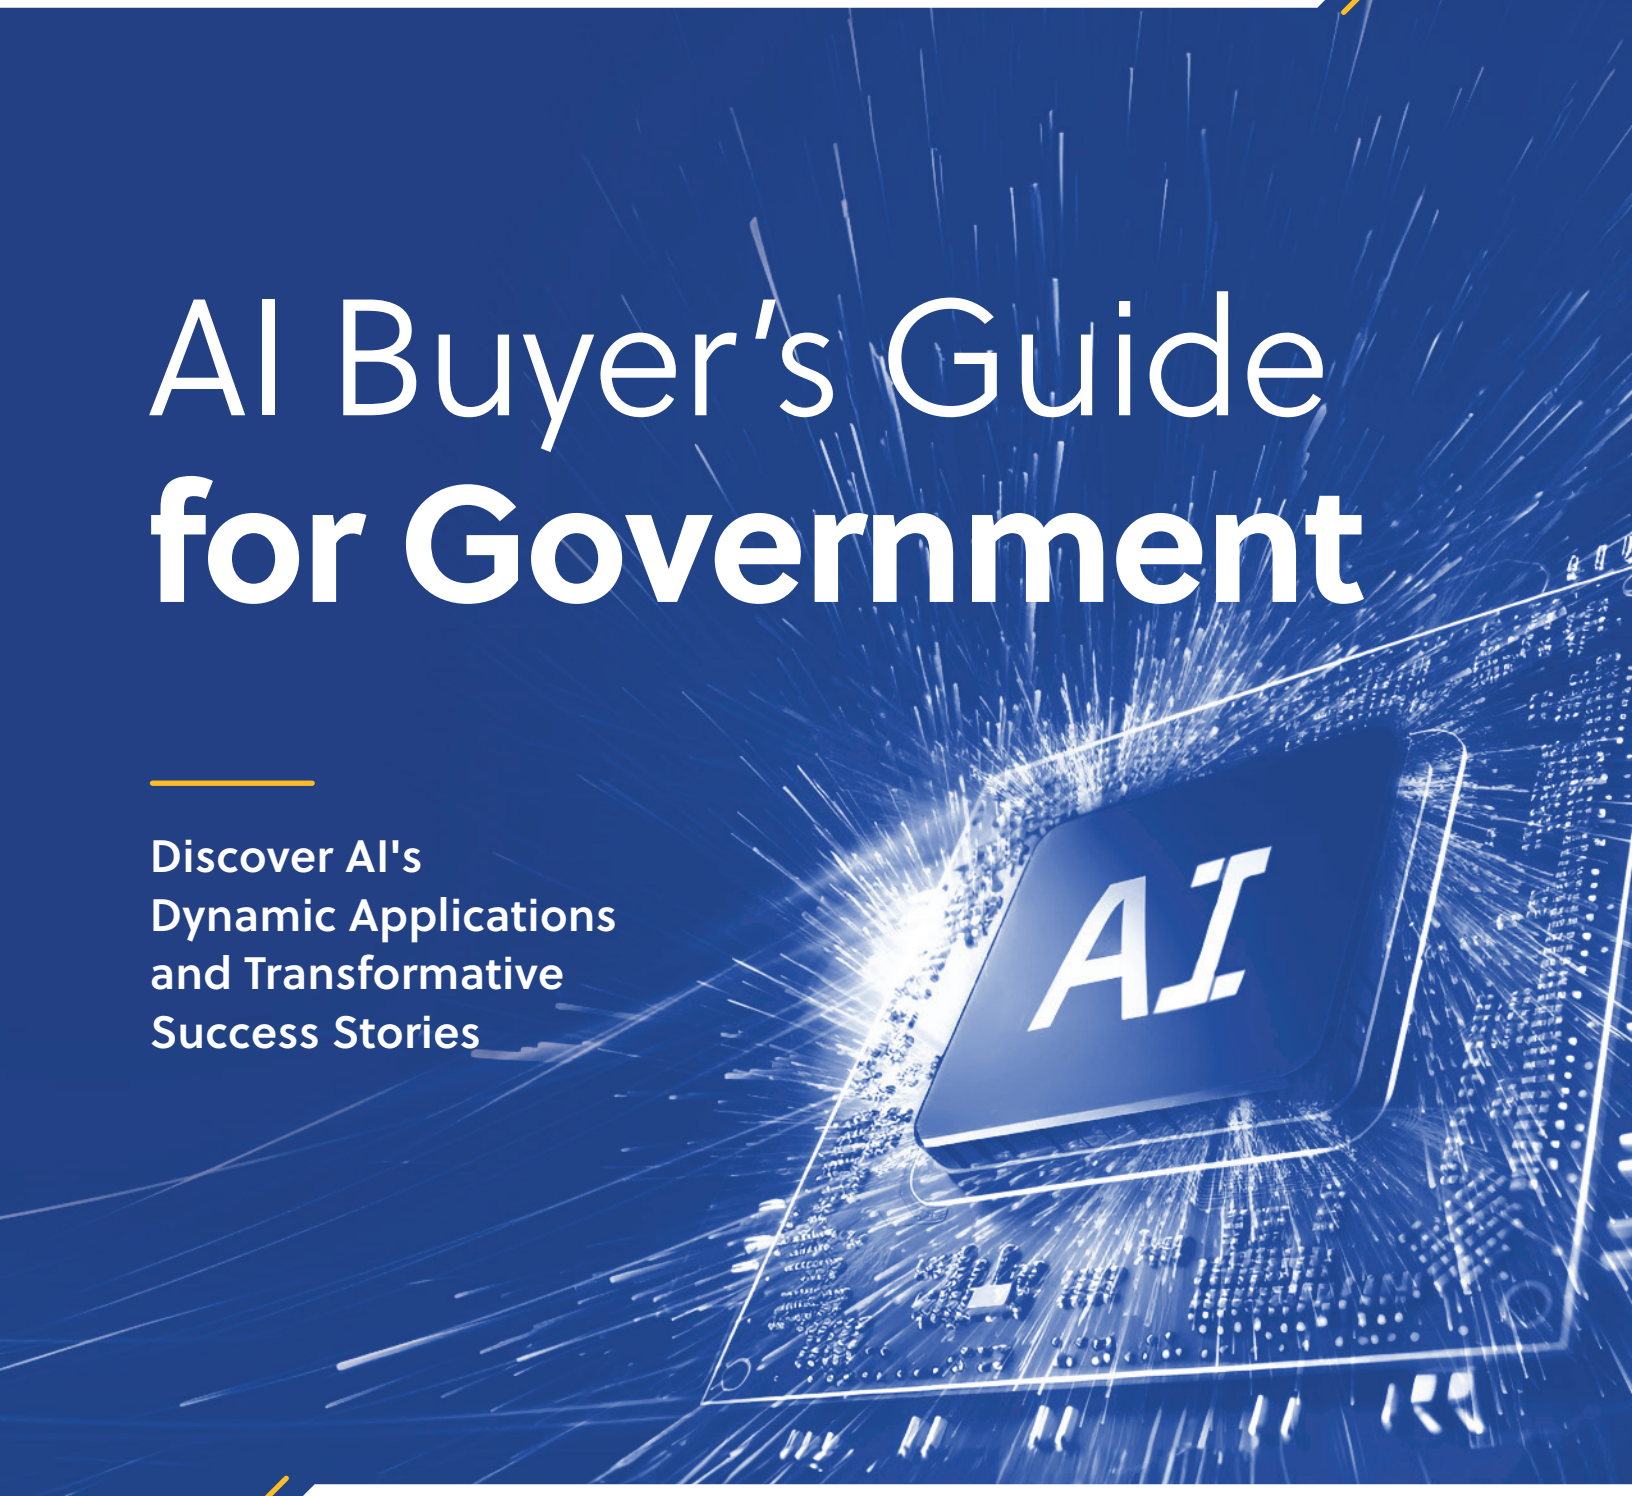

FEATURING: ***Solution Areas*** • ***Success Stories*** ***Contract Vehicles*** • ***Policies to Know*** • ***Upcoming Events***

carahsoft govloop

AI for Government Summit 2025

Innovate, Collaborate, Govern: *AI for the Public Sector*

Thursday, May 15, 2025
8:00am - 5:30pm ET

Carahsoft Conference & Collaboration Center

Carahsoft's AI for Government Summit is the premier event for exploring how AI is revolutionizing the Public Sector. This summit brings together Federal, State and Local Government leaders, industry innovators and technology experts to discuss the transformative potential of AI in addressing the nation's most pressing challenges.

Attendees will gain actionable insight into how Generative AI is transforming Government operations, how AI can streamline public services, cybersecurity and data protection, and how to leverage AI for predictive analytics and decision-making, as well as tackling ethical, privacy and compliance challenges.

Engage in expert-led discussions on key AI priorities, including:

- AI Infrastructure in Government
  Building scalable, secure AI frameworks
- Leveraging Data and AI to Drive Public Sector Innovation
  Harnessing data for smarter decision-making
- AI and Cybersecurity - Strengthening Government Safety
  Protecting critical infrastructure from AI-driven threats
- AI Policy within the New Administration
  Understanding AI governance and compliance
- Generative AI - Driving Efficiency in Government
  Real-world applications for improving agency operations

Join us as we chart the next phase of AI's role in government – scan the QR code to secure your spot!

Welcome to the AI Buyer's Guide!

Navigating AI Adoption in the New Era of Government

As the new administration takes office, the role of artificial intelligence (AI) in government is more critical than ever. With a renewed focus on modernization, efficiency, and responsible innovation, agencies at every level are exploring how AI can enhance public services, streamline operations, and drive data-driven decision-making. However, adopting AI in government comes with unique challenges—balancing security, compliance, ethical considerations, and fiscal responsibility.

This AI Buyer's Guide for Government is designed to help agency leaders, procurement officers, and policymakers make informed decisions about AI investments. Whether you are evaluating AI for cybersecurity, public service automation, or policy analysis, this guide provides a clear roadmap for assessing vendors, ensuring compliance with federal and state regulations, and aligning AI initiatives with mission-driven objectives.

With the new administration prioritizing AI governance, workforce upskilling, and digital transformation, this guide will help you navigate the evolving landscape and make strategic decisions that serve both today's needs and future challenges.

Let's build an AI-ready government that enhances efficiency, fosters trust, and delivers meaningful impact.

Mike Adams
Program Executive for AI Solutions,
Carahsoft

Table of Contents:

AI Solution Areas

Government agencies are responsible for leveraging data to make critical decisions that impact anything ranging from the daily lives of citizens to matters of national security. With Artificial Intelligence (AI), Machine Learning (ML) and High-Performance Computing (HPC) the public sector can efficiently and securely manage data while ensuring compliance is maintained.

Artificial intelligence has a wide array of use cases. For government agencies, this can range from helping process documents to hunting cyber threats on your network. Carahsoft offers artificial intelligence solutions for whatever your need is.

Scan the QR to learn more about Carahsoft's AI solutions

In this section, we cover:

- Generative AI
- Data & Analytics
- AI Infrastructure
- AI ModelOps
- Cybersecurity

Generative AI

Generative AI provides a huge boost to agency productivity by assisting in everything from the proposal writing process to aiding researchers as a personal assistant. Technologies such as Large Language Models (LLMs) will help revolutionize agencies' workflow.

Below are Carahsoft vendors who can help you achieve your mission.

Adobe	Archive360		blaize	causaLens	CITIBOT
class	CLEAR ML	codeium	COHESITY	Commvault	data.world
data iku	DISASTERTECH	Dynamo AI	dynatrace	eSHARE	edgerunner AI
eGain	elastic	fiddler	grammarly	Icertis	KnowledgeLake
MLtwist	netskope	OpenAI	permuta	PRYON	salesforce
sas	scale	servicenow	THREAD AI	tomorrow.io	YeXT
zoom					

Data & Analytics

Artificial intelligence allows agencies to harness powerful analytics to extract maximum value from their data to support their mission. Whether it is enhancing citizen service, building complex simulations, or enhancing a commander's decision-making on the battlefield, AI-enabled analytics will be crucial for any agency looking to unlock the potential of its data.

Below are Carahsoft vendors who can help you achieve your mission.

ABBYY | ALTAIR | alteryx | asana | avaap | C3.ai
causaLens | chooch | clarifai | cloudnine | COHESITY | Collibra
CORAS | CORNERSTONEX | CrowdAI | DATALOGZ | DataShapesAI | Delphix by Perforce
Deltek | denodo | EchoMark | GRANICUS | Informatica | IRON MOUNTAIN
IronNet | kyndi | netskope | NUANCE | (OCIENT)™ | pixlogic
portal26 | pyramid | qualtrics XM | SAP | SAP NS2 NATIONAL SECURITY SERVICES | sas
Semarchy | Senzing | snowflake | STRIVEWORKS | SuccessKPI | tableau from Salesforce
TORCH.AI™ | TSECOND | TUNGSTEN AUTOMATION

AI Infrastructure

AI infrastructure is the backbone of every AI deployment. Powerful GPUs and other hardware enable agencies to support custom models and run advanced analytics to achieve their mission. Cloud-based platforms unite the power of AI and cloud into one offering.

Below are Carahsoft vendors who can help you achieve your mission.

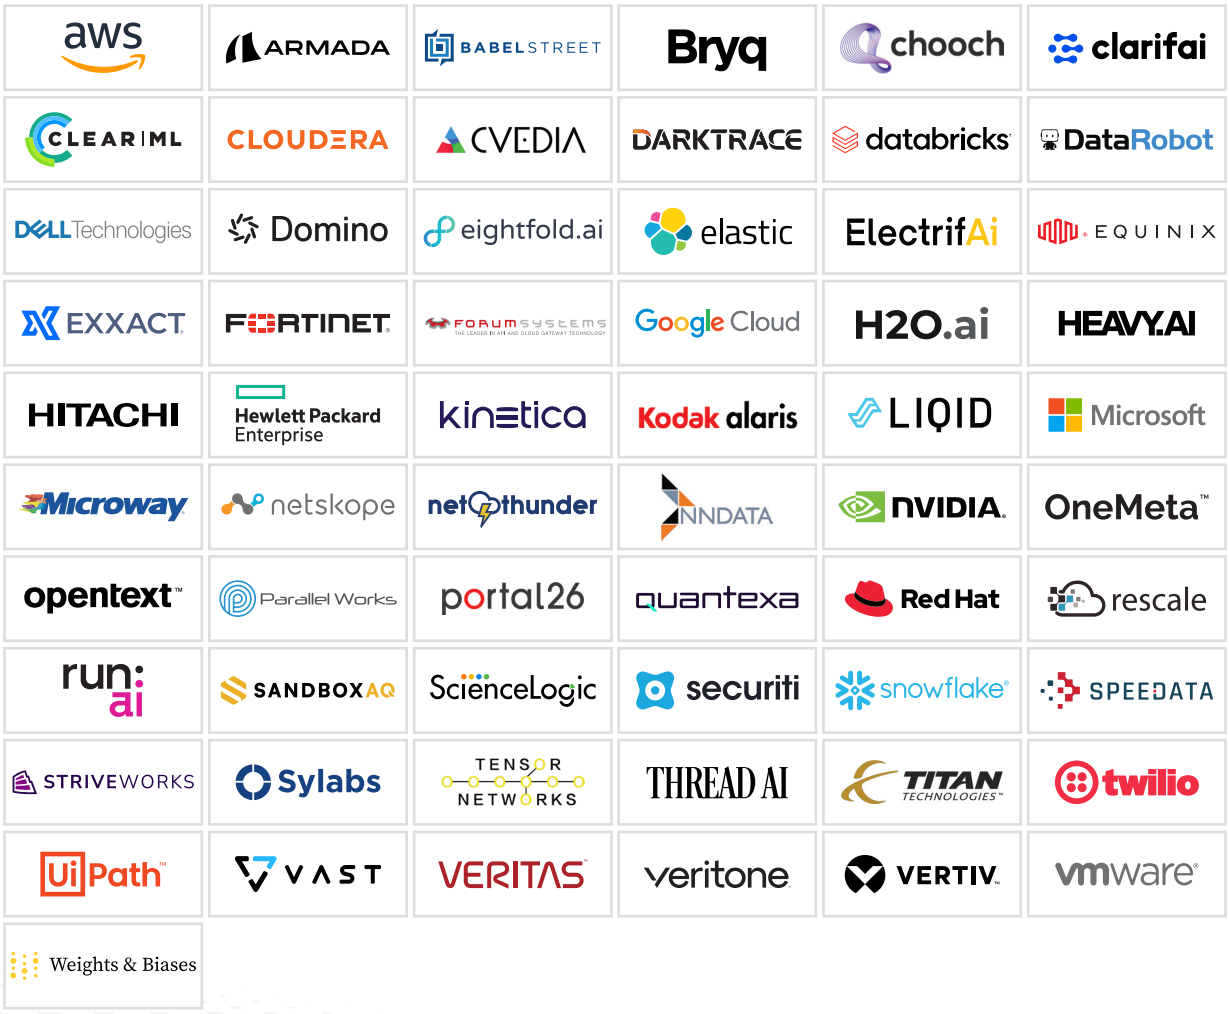

AI ModelOps

Every AI project needs to the tools to help support it. AI ModelOps helps developers and users build, manage, deploy, and monitor their models through a systems entire lifecycle.

Below are Carahsoft vendors who can help you achieve your mission.

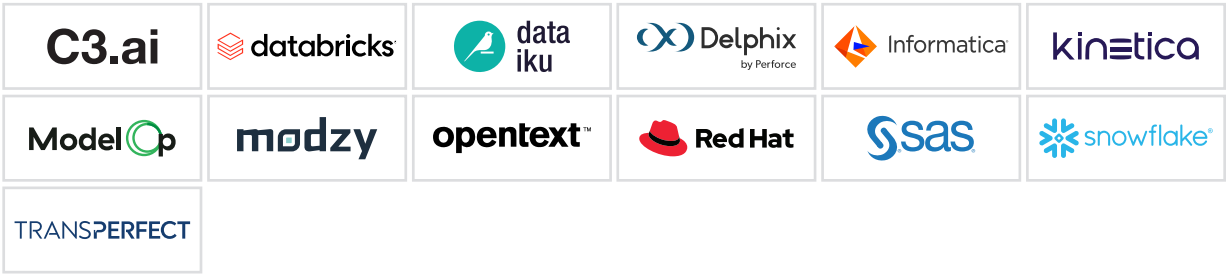

Cybersecurity

AI-enabled cybersecurity solutions allow you to identify threats before they strike, continuously monitor your network, and free up your cyber workforce to deal with more advanced threats. As malicious actors adopt AI to attack networks, government agencies will need to adopt AI-enabled security to fight back.

Below are Carahsoft vendors who can help you achieve your mission.

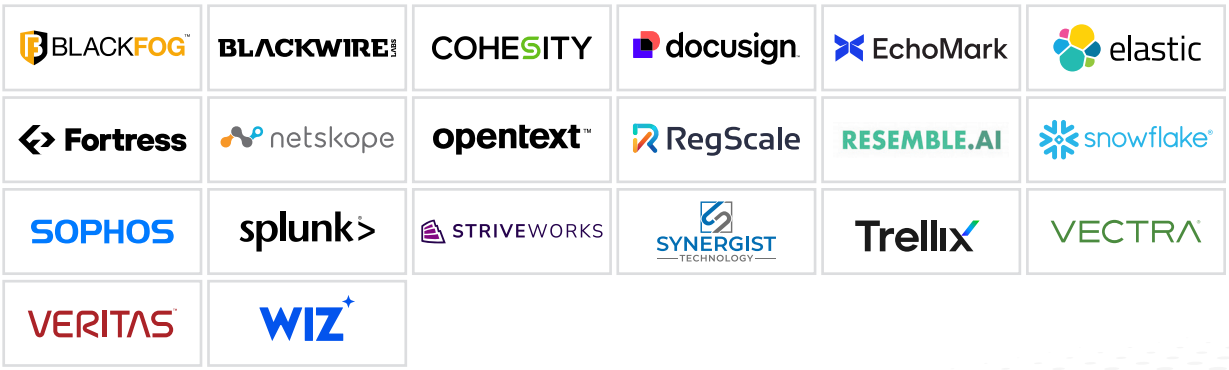

CLOUDERA

Accelerate AI with Trusted Data

Cut Generative AI production time by 50%

Deliver true hybrid for data, analytics and AI

Move data from any source to any destination

Learn More

OpenAI

ChatGPT Enterprise for Government

Transforming government operations by improving productivity, streamlining processes and enhancing service delivery

Watch the on-demand webinar to learn more

Google Cloud + Botcopy

Botcopy

The AI the Public **Deserves.**

The easy button for bringing Google AI to your website.
Our AI-powered platform ensures a satisfying UX and learns from every interaction. It's the UI of AI: ADA compliant and works with your existing IT infrastructure.

A breakthrough solution, helping you help everyone.

Your next "aha" moment starts here.

botcopy.com

Success Stories

Many agencies have found success adopting artificial intelligence to achieve their mission. Explore the stories of how other government customers have achieved their goals with the help of Carahsoft AI vendors.

How AI Helps Fight Fraud in Financial Services, Healthcare, Government and More

Companies and organizations are increasingly using AI to protect their customers and thwart the efforts of fraudsters around the world.

Voice security company Hiya found that 550 million scam calls were placed per week in 2023, with INTERPOL estimating that scammers stole $1 trillion from victims that same year. In the U.S., one of four noncontact-list calls were flagged as suspected spam, with fraudsters often luring people into Venmo-related or extended warranty scams.

Traditional methods of fraud detection include rules-based systems, statistical modeling and manual reviews. These methods have struggled to scale to the growing volume of fraud in the digital era without sacrificing speed and accuracy. For instance, rules-based systems often have high false-positive rates, statistical modeling can be time-consuming and resource-intensive, and manual reviews can't scale rapidly enough.

In addition, traditional data science workflows lack the infrastructure required to analyze the volumes of data involved in fraud detection, leading to slower processing times and limiting real-time analysis and detection.

Plus, fraudsters themselves can use large language models (LLMs) and other AI tools to trick victims into investing in scams, giving up their bank credentials or buying cryptocurrency.

But AI — coupled with accelerated computing systems— can be used to check AI and help mitigate all of these issues.

Businesses that integrate robust AI fraud detection tools have seen up to a 40% improvement in fraud detection accuracy — helping reduce financial and reputational damage to institutions.

These technologies offer robust infrastructure and solutions for analyzing vast amounts of transactional data and can quickly and efficiently recognize fraud patterns and identify abnormal behaviors.

AI-powered fraud detection solutions provide higher detection accuracy by looking at the whole picture instead of individual transactions, catching fraud patterns that traditional methods might overlook. AI can also help reduce false positives, tapping into quality data to provide context about what constitutes a legitimate transaction. And, importantly, AI and accelerated computing provide better scalability, capable of handling massive data networks to detect fraud in real time.

The Challenge:
The United States Government Accountability Office estimates that the government loses up to $521 billion annually due to fraud, based on an analysis of fiscal years 2018 to 2022. Tax fraud, check fraud and improper payments to contractors, in addition to improper payments under the Social Security and Medicare programs have become a massive drag on the government's finances.

While some of this fraud was inflated by the recent pandemic, finding new ways to combat fraud has become a strategic imperative. As such, federal agencies have turned to AI and accelerated computing to improve fraud detection and prevent improper payments.

The Solution:

For example, the U.S. Treasury Department began using machine learning in late 2022 to analyze its trove of data and mitigate check fraud. The department estimated that AI helped officials prevent or recover more than $4 billion in fraud in fiscal year 2024.

Along with the Treasury Department, agencies such as the Internal Revenue Service have looked to AI and machine learning to close the tax gap — including tax fraud — which was estimated at $606 billion in tax year 2022. The IRS has explored the use of NVIDIA's accelerated data science frameworks such as RAPIDS and Morpheus to identify anomalous patterns in taxpayer records, data access and common vulnerability and exposures. LLMs combined with retrieval-augmented generation and RAPIDS have also been used to highlight records that may not be in alignment with policies.

Key Takeaways:

Financial services, the public sector and the healthcare industry are all using AI for fraud detection to provide a continuous defense against one of the world's biggest drains on economic activity.

The NVIDIA AI platform supports the entire fraud detection and identity verification pipeline — from data preparation to model training to deployment — with tools like NVIDIA RAPIDS, NVIDIA Triton Inference Server and NVIDIA Morpheus on the NVIDIA AI Enterprise software platform.

Minnesota's Enterprise Translation Office Uses ChatGPT to Bridge Language Gaps

For government services to be truly effective, they must be accessible to all residents. Removing language barriers is essential to high quality public service. This is the founding principle of the State of Minnesota's Enterprise Translations Office (ETO), established in 2023 to provide dedicated translation services for the state's Executive Branch

The Challenge:

Over 20 percent of Minnesota's residents primarily speak a language other than English, with Spanish, Somali, and Hmong as the top three major non-English languages (opens in a new window) spoken in the state. Language gaps in government communication are a challenge for community health, safety, and trust. The ETO works to ensure equitable access to government information and resources for all, bridging linguistic gaps to ensure a stronger, more inclusive Minnesota.

Government agencies have long made consistent efforts to provide translation services, but Minnesota, like many states, faces challenges. The state's translation process was highly decentralized, with each department relying on various methods, including engaging external contractors, leading to inconsistent quality, high costs, and delays that could extend up to a month for a single request. Interpretation services were equally challenging, where limited access to approved interpreters made it difficult for people to engage with critical legal and community services effectively.

The Solution:

To address these issues, the ETO began using AI, incorporating ChatGPT as a key pillar of their solution to provide translation services faster, more accurately, and more equitably. The ETO collaborated closely with the Transparent Artificial Intelligence Governance Alliance (TAIGA), Minnesota's IT Services program, to ensure the program followed safe, ethical, and transparent AI principles. By integrating ChatGPT into the initial stage of translation with meticulous human review and quality assurance, the ETO established a more efficient workflow that sets a new standard for language inclusion in the public sector.

As Adam Taha, Director of the ETO, describes, "The balance between accuracy and speed is the crux of the problem." By incorporating AI-driven solutions and especially ChatGPT, the ETO has been able to address this challenge, enhancing both the speed and accuracy of its translation services.

Beyond translation accuracy, the ETO's ChatGPT-enabled workflow also optimizes for cultural relevance to ensure government messaging is both accessible and inclusive. For example, with Hmong and Somali, the ETO developed spreadsheets with culturally relevant terms and translations and incorporated these into their GPTs to improve the model's cultural awareness. By incorporating a process where ETO translation experts review and correct ChatGPT translations before uploading them to custom GPTs for future reference, the ETO significantly speeds up the quality assurance stage, with fewer edits required for each translation, ensuring that no mistakes or discrepancies occur twice.

The program's success led to a full rollout in July, just four months after the beta launch, with all ETO team members trained to utilize ChatGPT for translation work. Now, the ETO is supporting a limited pilot project with another state agency to test the use of ChatGPT's voice capabilities for real-time interpretation. This pilot aims to improve accessibility for non-English-speaking citizens in various live situations and provide faster, more efficient communication for those with limited English proficiency. The early results are promising. If successful, the pilot could expand to more locations, further enhancing the accessibility of public services for all Minnesotans. that could support their growing organization and new services.

Key Takeaways:

The ChatGPT-integrated translation process has been pivotal for the ETO and the communities they serve. Karine Lao Her, the department's Hmong language team lead, shares, "We anticipate rapid growth in our Hmong translation services in the coming months as we continue to refine the accuracy of our Hmong Custom GPTs. This has already made a significant difference in how our community connects with state resources."

ChatGPT has significantly enhanced the ETO's and the State of Minnesota's Executive Branch's ability to serve Minnesotans in four impactful ways:

Rapid turnaround times:
What once took weeks to complete can now be done in under 48 hours, with urgent translation requests fulfilled in as little as 2 hours. Since the beta launch, the ETO has processed and delivered over 3,000 translation requests, translating over 2,000,000 words.

Financial efficiency:
Adopting this new process has allowed Minnesota to more efficiently allocate resources, resulting in financial benefits of over $100,000 per month, which can now be redirected to support other strategic state initiatives.

Service expansion:
Beyond Hmong, Somali, and Spanish, the ETO now offers translation and quality assurance support for languages including French, Arabic, Plain Language English, and Pidgin English by using ChatGPT. The ETO also offers support for more languages by providing cost reimbursements for requesters who use ETO services in conjunction with ChatGPT and conduct quality assurance with external vendors.

More equitable access to government services:
Daha Gobdoon, the Somali language team lead, reflects on the cultural importance of knowledge, quoting the Somali proverb, "Aqoon La'aani waa Iftiin La'aan," or "lack of knowledge is like lack of light." "When communities face challenges, delivering information in their language is essential, but timely delivery is often difficult. Now, our community receives crucial government information faster than ever, with rapid turnarounds eliminating the long wait for important documents."

"Now, our community receives crucial government information faster than ever, with rapid turnarounds eliminating the long wait for important documents. "

Daha Gobdoon, Somali Language Team Lead

VAST

Enhancing Performance and Driving Groundbreaking Research

The National Center for Supercomputing Applications (NCSA) at the University of Illinois Urbana-Champaign is a premier research institution dedicated to accelerating scientific research and discovery through cutting-edge HPC infrastructure. NCSA supports a wide range of users, including National Science Foundation (NSF) researchers, University of Illinois scientists, and Fortune 100 companies. To meet the complex and evolving demands of its users, NCSA needed a modern, reliable data platform capable of handling high I/O workloads and large-scale data processing.

The Challenge:

NCSA's infrastructure comprises multiple HPC clusters, storage systems, and secure environments that serve both academic and industry users. Managing this environment, which totals 60 petabytes of storage across 18 file systems, was becoming increasingly complex and resource-intensive. The ICI Directorate, the team responsible for deploying the cyber infrastructure to support the research at NCSA, needed to simplify operations while improving performance.

As J.D. Maloney, Lead HPC Storage Engineer at NCSA, explained: "Our challenge was to find a solution that could handle the read-heavy nature of home and software directories while reducing the overall footprint and complexity of our data environment." Additionally, with NCSA expanding into GPU-heavy research utilizing its 1,400 GPUs, the need for an optimized data platform that could keep up with demanding interactive and batch workloads became critical.

The Solution:

After evaluating several options, NCSA chose the VAST Data Platform for its home and software directories under its Harbor system. VAST's architecture offered the read performance, compression, and data reduction capabilities NCSA needed to streamline its operations. "The reason we went with VAST is that it's a very read-optimized file system. For areas like home and software directories, which are hit heavily by I/O, it was a perfect match," said Maloney.

One of the key benefits of VAST's system is its ability to deduplicate and compress data efficiently. This allowed NCSA to provide its users with larger home directory quotas while keeping the physical footprint of the system manageable. "We've obtained a 3.5:1 data reduction ratio with the potential to hit 4:1 as we onboard more data," Maloney noted. "This allowed us to support more systems with a smaller footprint and provide larger home directory quotas for our users."

Key Takeaways:

Improved Performance:
Users experienced faster load times for software and home directories, which is critical for interactive HPC workloads. For instance, module loads dropped from nine seconds to just 2.5 seconds, leading to a more responsive system overall. "Our user support staff were very excited. It felt like a new system because of how responsive it was," Maloney shared.

Data Efficiency:
With VAST's advanced compression and deduplication, NCSA achieved significant data reduction, allowing them to provide larger quotas with the same physical storage capacity.

Streamlined Operations:
VAST's support for multiple fabric types (Ethernet, InfiniBand, Slingshot) and multi-tenancy capabilities allowed NCSA to consolidate its infrastructure. This reduced the number of file systems NCSA had to manage, resulting in fewer maintenance windows and more uptime for its clusters.

Scalability for Future Growth:
With the addition of GPU-optimized clusters like Delta AI, NCSA can now handle larger, more complex AI workloads. VAST's read-optimized Data Platform ensures GPU resources are fully utilized by minimizing data loading times. "The ability to launch large containers across multiple nodes quickly is critical for our GPU workloads. VAST's performance allows us to keep those expensive GPUs fully utilized, which is a huge win for us," Maloney explained.

NCSA has adopted and integrated advanced AI infrastructure to accelerate scientific research across multiple fields, helping scientists tackle complex problems faster and with greater accuracy. NCSA's focus on AI began with the deployment of Delta AI, one of its flagship systems designed to provide unparalleled computation power for large-scale AI workloads. Delta AI includes 700 Grace Hopper GPUs, making it one of the most powerful AI-focused systems in the research community. This system is aimed at supporting a broad range of AI applications, from machine learning and deep learning to computational biology, materials science, and climate modeling. It also advances NCSA's ability to process complex datasets, train AI models, and conduct high-performance simulations across various scientific domains.

Key to NCSA's AI infrastructure is the VAST Data Platform, which provides high-speed access to critical datasets and models without bottlenecks. AI workloads, especially in large containerized environments, need a data platform with extremely high read performance and low latency. VAST accelerates data access, especially for container-heavy AI workloads, maximizing GPU utilization and streamlining AI research pipelines at NCSA.

As NCSA continues to expand its HPC offerings, it is exploring further use cases for VAST Data. Maloney is excited about VAST's roadmap, particularly with upcoming features like metadata compression and enhanced write performance. "Metadata compression will allow us to increase user quotas even further. It's a huge quality-of-life improvement for our users, and we're excited to implement it as soon as possible," said Maloney.

"The ability to launch large containers across multiple nodes quickly is critical for our GPU workloads. VAST's performance allows us to keep those expensive GPUs fully utilized, which is a huge win for us"

J.D. Maloney, Lead HPC Storage Engineer at NCSA

Collibra

How the Office of the Secretary of Defense Enables Data Discovery with Collibra

The Office of the Secretary of Defense (OSD) and its Comptroller Office of the Undersecretary of Defense are responsible for US defense policy, planning, resource management, and program evaluation. In carrying out those duties, the OSD's day-to-day decisions impact a wide range of operations, including human resources, weapons acquisition, research, intelligence and fiscal policy across all of the US armed forces (Army, Navy, Marine Corps and Air Force). Steering an organization of that size means data is crucial to its decision-making processes. However, ensuring decision makers have access to the right data, can trust in its accuracy and understand its context is by no means a simple task.

The Challenge:

Recognizing that challenge, Greg Little, Director, CFO Data Transformation Office, partnered with Collibra to launch their Advana (Advancing Analytics) Data Catalog to create a centralized platform for data and analytics across the organization.

Utilizing the technology provided by Collibra, the Advana program directly responded to the need for users and stakeholders across the OSD to find, understand and trust their data, ultimately empowering and enabling better decision-making.

Little saw the Advana program as a strategic initiative with a focus on eventually expanding the program to encompass numerous use cases across the OSD. Little chose Collibra Data Catalog as the technical solution because of its enterprise capabilities and flexibility.

However, the value of the Advana program wasn't realized overnight. It's an ongoing process that began with a targeted implementation and internal champions with experience to help drive the program's success.

The Solution:

Through the Advana program, Little empowered the OSD team to identify various opportunities for data cataloging across the organization. This strategy helped the team focus the implementation on a specific use case and then expand from there. Little also worked closely with the senior leaders to understand their data governance and data catalog needs and then built out a strong data team within the OSD to bring his vision to a reality.

When Greg Little hired Matt Piester as their Data as a Service Lead for the Advana program, the program was still in its infancy. Using his past experience working in different roles across numerous governmental organizations—a submarine officer in the Navy, a Defense contractor while serving in the reserves, several years in the Department of Health and Human Services, and now at OSD—Piester saw the value of this program and identified opportunities to communicate and implement this technology across the organization.

Getting started with driving the success of Advana at OSD, Little and Piester identified some immediate hurdles to overcome. The biggest challenge? Scoping down the project to narrow in on a singular use case.

More specifically, Piester recalls, "it is important to start with a specific use case and find answers to key questions from decision makers."

At the time, the OSD team determined that the primary use case was financial management. The CFO's office needed to be able to track spending and carry out their audit function more effectively. By cataloguing all data sources in scope of the OSD's financial audit function, the team helped simplify what had been a very challenging process. Collibra Data Catalog provided a "single source of truth" for the CFO's office and enabled financial analysts to have a complete view of the available data within the department.

But preparing for financial audits is just one use case. As the CFO's office began to get more comfortable with Collibra Data Catalog, Little and Piester looked to reorganize the data in the catalog based on different lines of business such as procurement or HR. To support that goal, Piester looked to catalog various data sets in Collibra by having the data ingested into the AWS infrastructure automatically populate Collibra Data Catalog.

It is not easy to scale a data program across a large organization. As Little

and Piester began expanding the Advana program to include different use cases, they quickly discovered that different groups were at varying levels of maturity from the perspective of Data Intelligence. "It has not necessarily been a one size fits all approach, because other lines of business are at different phases in their data journey," says Piester.

But with Little and Piester's past experience, they were able to effectively communicate across the organization and found a common goal among the teams within the OSD. They understood the needs and the challenges that many of their stakeholders had on a day-to-day basis and knew that every team was trying to solve for essentially the same problem regardless of their level of data maturity—finding, understanding and trusting their data.

Key Takeaways:

With this common goal in mind, they were able to apply the structure of the financial use case to these different lines of business and make adjustments as needed to customize the use case. In that process, he recognized that this first-implemented financial use case can serve as a model for what other teams can do with a data catalog and he could leverage members of the financial management team as advocates for the Advana program across the organization.

The Advana program is still relatively nascent, but has already made significant strides in its mission, with the team tracking key performance indicators to ensure they remain on the right path and can measure their progress. The team approaches measuring the success of the program in two ways:

Quantitatively:

They categorize success as usage and uptake. More specifically, they look at how many people are currently using the data catalog, who is using it and how often across the entire organization.

"Our metrics so far have been primarily around usage and uptake within the catalog. Just eight months after launch we now have almost 1000 registered users," says Piester.

The team is also interested in how often a user accesses Collibra Data Catalog. Do they access the catalog daily, weekly or monthly? This is an important statistic because Piester considers a successful implementation and adoption as one where people rely on the data catalog for information, rather than asking a colleague or using a spreadsheet. He sees their success illustrated in the above-mentioned financial audits use case.

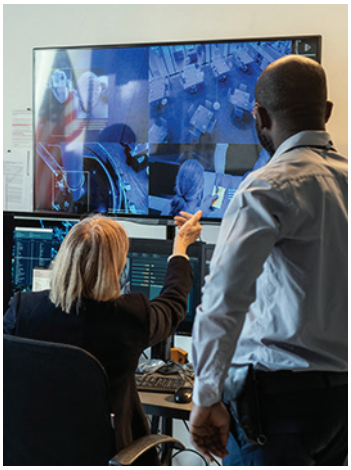

Qualitatively:

The team is currently working on a way to measure data quality. The team is putting together a data quality framework that not only aims to gauge levels of data quality but that also makes that information transparent to consumers and to help drive improvements. This data quality framework is developed and the team is planning to test its accuracy heading into 2021. To Piester and the team, this metric is crucial because it is important for users to know their data is high quality, trustworthy and accurate.

Results:

With a solid foundation in place, there are two ways in which the Advana program is set to evolve – both internally within the OSD and externally with its partners across the Department of Defense.

"Internally the focus is around increasing uptake within the community, making the catalog more comprehensive and making it part of business analysts' day-to-day workflow," says Piester. To support that goal, the team is seeking to catalog data sources beyond those currently aggregated and analyzed in Advana. Collecting pertinent metadata will help users know that data sets exist, even when accessing them may be more complex given their sensitive nature.

To succeed in that broader mission, the OSD team is working closely with colleagues across the Department of Defense, who have similar goals and who have initiated similar efforts in cataloging their data sources. Through this collaborative approach, the goal is to create a "federated data catalog" that spans key data sources across all organizations. Aligning these strategies will allow for the joint development of a strong data catalog and data management strategy across the entire Department of Defense.

"Ultimately, it should not necessarily matter who owns the data or where it is collected. I, as a user, will be able to come to Advana—or the Army, Navy or Air Force version—and know exactly what data is available, and we can ensure everyone is looking at the same universe," says Piester.

With ambitious goals in mind, the OSD team is looking forward to accelerating their data strategy with Advana in 2021 while continuing to provide value to users and stakeholders across their organization.

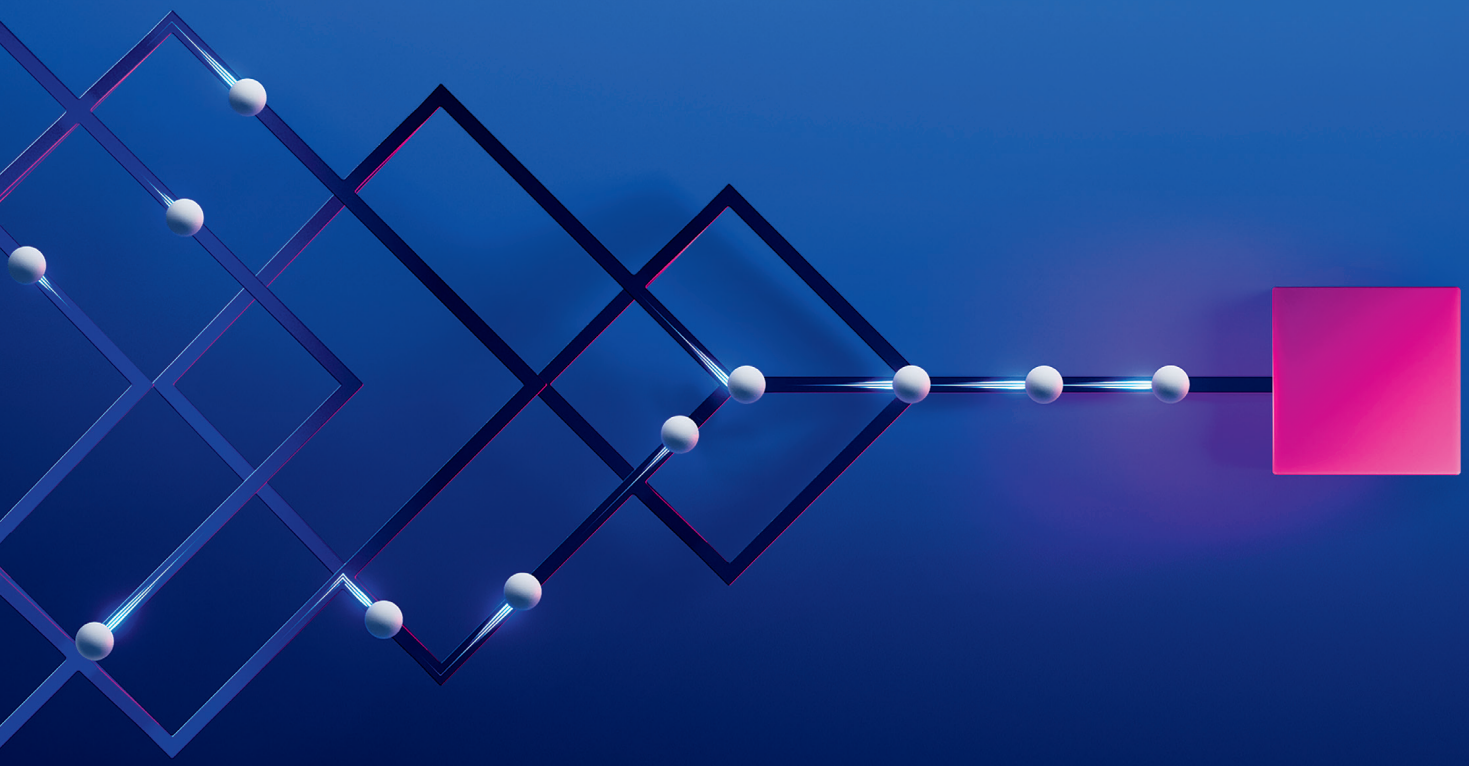

The quickest way from a billion points of data to a point of view

In a world that produces more data than ever, find answers 30x faster than ever. SAS® Viya® is the data and AI platform that understands the demands of public service, empowering agencies to make smarter, faster decisions while lowering cloud operating costs by an average of 86%.

sas viya

Start turning data into insights at sas.com/viya.

SANDBOXAQ | AQTIVE GUARD

Defend Your Data Today, Secure Your Future with AQtive Guard

- Full visibility, monitoring and control of your cryptography
- Automatic analysis and reporting on your cryptographic security posture
- Cryptographic policy management and enforcement
- Seamless upgrades to new post quantum cryptographic protocols

sandboxaq.com

Compliance with PQC mandates such as NSM-10, National Security Memorandum on Promoting United States Leadership in Quantum Computing While Mitigating Risks to Vulnerable Cryptographic Systems, OMB Memorandum Migrating to Post-Quantum Cryptography, and Quantum Computing Cybersecurity Preparedness Act

DELL Technologies

Customer Highlight: Community Services Become More Accessible with GenAI

The city of Amarillo, Texas is building a more connected and engaged community that fosters a prosperous future for its diverse citizens.

The Challenge:

- 24% of Amarillo's population doesn't speak English.
- One middle school has 62 languages and dialects.
- For more than a decade, Amarillo has had the highest number of refugees (per capita) of any city in the United States.
- Prior to 2020, half of the community did not have access to high-speed broadband.

The Solution:

Generative AI consultants from Dell Professional Services helped Gagnon take his vision to the next level.

The implementation of a digital assistant with the ability to interact in several languages will be transformational in a city with a large population of non native-speaking residents.

Residents access information through Emma, the GenAI digital assistant that lives on the city's website. The aim is to turn the interaction between resident and government into an engaging conversation.

- Designed with the city's identity and tone of voice
- Answers queries and helps people request government services
- Delivers services in 62+ languages

"We're not afraid of the future. We're embracing this wholeheartedly."

Rich Gagnon, Assistant City Manager and Chief Information Officer

Key Takeaways:

Deploying a digital assistant that delivers accurate answers for users at scale requires testing and then testing some more. The city of Amarillo uses anonymous logs during testing to get real-time feedback on what end users are asking about.

Collecting early feedback:
Using thumbs up or down reviews to get immediate reviews/feedback on conversations — and if Emma got the answers right or wrong. Helping the team build business processes that address the back-end data and getting the right team to give the right information.

Capturing questions:
Getting real-time input on what residents are asking about.

Understanding sentiment:
Using more traditional AI tools that can analyze sentiment and build dashboards and metrics. Measuring the goal of happy citizens.

Looking to the future:
Deploying Emma on the city's website serves as a blueprint for other uses within the community, including:

- Extending English as a second language assistance to help new students integrate more easily with others and bridge student/teacher communication
- Improving services in support of remote health clinics and emergency response time
- Serving as a finder guide inside government buildings

Google Cloud

How the Air Force Research Laboratory is Advancing Defense Research with AI

Through our collaboration, the Air Force Research Laboratory (AFRL) is leveraging Google Cloud's cutting-edge artificial intelligence (AI) and machine learning (ML) capabilities to tackle complex challenges across various domains, from materials science and bioinformatics to human performance optimization. AFRL, the center for scientific research and development for the U.S. Air Force and Space Force, is embracing the transformative power of AI and cloud computing to accelerate its mission of developing and transitioning advanced technologies to the air, space, and cyberspace forces.

This collaboration not only enhances AFRL's research capabilities, but also aligns with broader Department of Defense (DoD) initiatives to integrate AI into critical operations, bolster national security, and maintain technological advantage by demonstrating game-changing technologies that enable technical superiority and help the Air Force adopt to cutting edge technologies as soon as they are released. By harnessing Google Cloud's scalable infrastructure, comprehensive generative AI offerings and collaborative environment, the AFRL is driving innovation and ensuring the U.S. Air Force and Space Force remain at the forefront of technological advancement.

The Challenge & The Solution:
Let's delve into examples of how the AFRL and Google Cloud are collaborating to realize the benefits of AI and cloud services:

Bioinformatics breakthroughs:
The AFRL's bioinformatics research was once hindered by time-consuming manual processes and data bottlenecks, causing delays in moving and sharing data, getting access to US-based tools, using standard storage and hardware, and having the right system communications and integrations across third party infrastructure. Because of this, cross-team collaboration and experiment expansion was severely limited and inefficiently tracked. With very little cloud experience, the team was able to create a siloed environment where they used Google Cloud's infrastructure, such as Google Compute Engine, Cloud Workstations, and Cloud Run to build analytic pipelines that helped them test, store, and analyze data in an automated and streamlined way. That data pipeline automation paved the way for further exploration and expansion on a use case that had never been done before.

Web app efficiency for lab management:
The AFRL's complex lab equipment scheduling process resulted in challenges in providing scalable, secure access to important content and information for users in different labs. To mitigate these challenges and ease maintenance for non-programmer researchers and lab staff, the team built a custom web application based on Google App Engine, integrated with Google Workspace and Apps Scripts, so that they could capture usage metrics for future hardware investment decisions and automate admin tasks that were taking time away from research. The result was significantly faster ability to make changes without administrator intervention, a variety of self-service options for users to schedule time on equipment and request training, and an enhanced, scalable design architecture with built-in SSO that helped streamline internal content for multiple labs.

Modeling insights into human performance:
Understanding and optimizing human performance is critical for the AFRL's mission. The FOCUS Mission Readiness App, built on Google Cloud utilizes various infrastructure services, such as Cloud Run, Cloud SQL, and GKE and integrates with the Garmin Connect APIs to collect and analyze real-time data from wearables.

By leveraging Google Cloud's BigQuery and other analytics tools, this app provides personalized insights and recommendations for fatigue interventions and predictions that help capture valuable improvement mechanisms in cognitive effectiveness and overall well-being for Airmen.

Streamlined AI model development with Vertex AI:
The AFRL wanted to replicate the functionality of university HPC clusters, especially since there was a diversity of users that needed extra compute and not everyone was trained on how to use these tools. They wanted an easy GUI and to maintain active connections where they could develop AI models and test their research with confidence. They leveraged Google Cloud's Vertex AI and Jupyter Notebooks through Workbench, Compute Engine, Cloud Shell, Cloud Build and much more to get a head start in creating a pipeline that could be used for sharing, ingesting, and cleaning their code. Having access to these resources helped create a flexible environment for researchers to do model development and testing in an accelerated manner.

Google Cloud

Key Takeaways:

"Cloud capabilities and AI/ML tools provide a flexible and adaptable environment that empowers our researchers to rapidly prototype and deploy innovative solutions. It's like having a toolbox filled with powerful AI building blocks that can be combined to tackle our unique research challenges." -Dr. Dan Berrigan, Air Force Research Laboratory.

The AFRL's collaboration with Google Cloud exemplifies how AI and cloud services can be a driving force behind innovation, efficiency, and problem-solving across agencies. As the government continues to invest in AI research and development, collaborations like this will be crucial for unlocking the full potential of AI and cloud computing, ensuring that agencies across the federal landscape can leverage these transformative technologies to create a more efficient, effective, and secure future for all.

"Cloud capabilities and AI/ML tools provide a flexible and adaptable environment that empowers our researchers to rapidly prototype and deploy innovative solutions."

Dr. Dan Berrigan, Air Force Research Laboratory

Informatica

Dallas County Modernizes Public Health

More than 2.6 million citizens across 31 municipalities rely on Dallas County Health and Human Services (DCHHS) to help secure their quality of life and protect them from disease. As a public health agency, DCHHS is responsible for tracking and reporting on diseases, immunizations, environmental health, and food safety. The agency also provides essential social services, from housing aid to the county's Comprehensive Energy Assistance Program.

That's no small feat, especially at the speed and scale with which public health conditions change in today's world. "The decisions we make affect millions of people, so we can't afford to work with inaccurate or unreliable data," says DCHHS Director, Dr. Philip Huang.

The COVID-19 pandemic brought this reality to life for DCHHS, highlighting the importance of timely and accurate resident data to track, report, and manage the spread of life-threatening conditions.

When the disease hit, the agency was still using fax machines to collect citizen lab information and entering data into their systems manually. Huang's team was inundated with thousands of cases via paper-based faxes and found themselves overwhelmed by the sheer volume of data.

Driven by a sense of urgency, the team set out to meet three objectives. The first was to manage rising data volumes without sacrificing the quality of public health decision-making. Secondly, the team wanted to run large-scale contact-tracing operations and help public health workers to better control the spread of disease. And finally, they aimed to accelerate disease investigation workflows to improve DCHHS's response times.

"This was a lifechanging initiative for our organization and the citizens of Dallas," says Huang. "We needed to build more capacity, free up staff to do more epidemiological work, and combine our data into user-friendly dashboards, which could drive sound decision-making in the face of fast-shifting healthcare conditions."

The Challenge:

The top priority for DCHHS was to power its disease surveillance and response efforts with trustworthy and timely data to identify trends across communities, populations and neighborhoods. It also needed to process and combine that data more efficiently and reliably. Before the agency's transformation, staff would spend up to 15 hours per day just to eliminate redundant information and verify citizens' address data.

With the help of Informatica Master Data Management (MDM), DCHHS created a centralized data engine for its entire organization. Today, this solution is known as the Dallas Disease Surveillance and Investigation System. Seamless integration with the DCHHS Salesforce instance provides business users with a comprehensive and authoritative data record for each person to support the business. During the COVID-19 pandemic alone, the organization added new data streams, including electronic lab data, hospital case reports from electronic health records, immunization data and mortality data to its business intelligence dashboards, using SFDC and CRM functionality.

Informatica

Building a Centralized Data Engine
By managing all domains of master data in a single SaaS solution, DCHHS realized there were many advantages to making that shift, including artificial intelligence (AI) capabilities, natural language processing and data profiling. In addition to integrating and improving the quality of DCHHS' data, Informatica has improved the user experience for officials across the organization who rely on its data-driven services to power their decision-making.

"Our goal was to adopt first-rate offerings for every function in our data environment, and Informatica was selected for master data management," says Huang. "With custom matching and merging capabilities we can combine our lab, immunization, mortality, case report data and more into a golden record we can trust."

As a result, DCHHS's data engine can ingest, validate, match, and merge records automatically using business rules developed by Huang and his team. It processes more than 4.5 million data points and can easily scale to match the agency's growing data volumes. Similarly, data feeds are refreshed nightly to ensure teams work with timely and accurate information. DCHHS can increase this frequency when a public health situation calls for rapid analysis and decision-making.

"With custom matching and merging capabilities we can combine our lab, immunization, mortality, case report data and more into a golden record we can trust."

Dr. Philip Huang, DCHHS Director

The Solution:

At an operational level, DCHHS integrated several mission-critical healthcare systems, including Salesforce, the National Electronic Disease Surveillance System, the Texas immunization registry, and others into a single master data environment. By automating data ingestion, matching and merging of records from a variety of data sources, Huang and his team cut 2,000 hours of manual review per year from their workloads. To date, DCHHS has deduplicated 20% of all data involved in its contact tracing integrations and manual entries, while automated matching and merging rules have improved the reliability of its most critical data elements by 50%.

A single source of truth provides DCHHS with a 360-degree view of the public health trends of Dallas County residents, including their data history. That comprehensive understanding allows DCHHS to respond more swiftly and effectively when faced with time-sensitive public health situations.

It also drives better decision-making as the agency expands its implementation and pursues its longer-term goals. For example, when Africa experienced recent outbreaks of Ebola and Mpox in 2022, DCHHS expanded its centralized data environment to monitor return travelers from the region. The system now covers more than 120 reportable conditions, and the agency expects it to include every reportable condition by July of 2024.

Huang and his team continue to explore new features in their Informatica MDM environment. From improved data profiling to the use of artificial Intelligence and natural language processing to accelerate business decision-making, they are laser-focused on modernizing Dallas-based public health services for each citizen, and for every possible future.

Huang and his team have established DCHHS as a model of public health reporting at the state and national level. The data dashboards posted on its website during the COVID-19 pandemic became a crucial reference for the public. DCHHS also began linking these reference pages to other services and educational content to help keep the public as informed and safe as possible.

The data visualizations that DCHHS presented to the Dallas County Commissioners informed the county's response to the disease every step of the way. "There were literally trillions of dollars' worth of public health decisions made using public health data during the COVID-19 pandemic, at the federal, state and local levels. Everyone was looking at our numbers and seeing if cases were going up or down and identifying who was being most impacted. Informatica played a big role in our success," says Huang.

Today, DCHHHS is working with a team at the University of Texas at Austin, one of 13 funded CDC sites working to combine data from the country's most advanced healthcare systems to improve forecasting and help public health institutions better respond to threats.

The agency was also recognized by the Healthcare Information and Management Systems Society (HIMSS) in 2023. Their prestigious Davies Award honors public health organizations that leverage data and technology to enhance their core services and drive faster, more actionable intelligence to improve community health.

Key Takeaways:

DCHHS is now migrating its hosted Informatica environment to the cloud to support more complex forms of data analysis. For example, the ability to quickly profile data coming into its organization and create rules for data transformation and standardization will help ensure the highest standard of data quality and governance.

Huang and his team are also exploring how to leverage the AI capabilities of the Informatica system for file ingestion mapping and entity recognition. The ability to query data with natural language will make analyses more accessible to additional DCHHS epidemiologists, who won't require in-depth knowledge of SQL programming to run their numbers. The time savings delivered by AI promise to accelerate processes from data classification to discovery, as well as the gleaning of insights, making DCHHS more productive and effective.

"When people talk about healthcare, they typically think about the clinical side of the equation. But public health is just as fast-moving and complex, which is why technologies like AI will be indispensable for our organization," says Huang. "Working with partners like Informatica gives us the confidence that we can make the most of these innovations and continue improving the lives of our citizens each day."

"Working with partners like Informatica gives us the confidence that we can make the most of these innovations and continue improving the lives of our citizens each day."

Dr. Philip Huang, DCHHS Directo

Microsoft

How PNNL is Leveraging AI and ML to Enhance Infrastructure Resiliency

In recent years, the escalating frequency, intensity and cost of natural disasters have posed widespread challenges for federal, state and commercial officials. One such challenge is ensuring the resiliency of the nation's infrastructure, including the energy grid.

That concern prompted scientists at the Department of Energy's Pacific Northwest National Laboratory (PNNL) to begin exploring how artificial intelligence (AI) and machine learning (ML) could be used to predict wildfires and mitigate the impact of other potential natural disasters.

Pioneered by experts like Chief Scientist Andre Coleman, PNNL's groundbreaking work has helped modernize disaster and emergency management by using AI and ML to make sense of vast amounts of satellite and geospatial data in near real-time. This helps public safety officials, utilities and land-use managers better prepare for pending disasters.

The Challenge:

According to Coleman, PNNL began experimenting in 2014 with a vision to harness imagery from satellite, aircraft and drones to swiftly assess hazards created by wildfires, floods, hurricanes and other weather-related disasters – and their impact on critical energy infrastructure.

"The original idea was to say, 'Can we quickly make use of these sorts of imaging resources and assess a hazard? What's the extent of the hazard? Where are we seeing the damage caused from the hazard?'" says Coleman. "And how do we relate those damages to critical energy infrastructure?"

To do this initial work, Coleman says PNNL built machine learning models and tools "to take that imagery and assess it in an automated way." And then, over time, more machine learning algorithms were put into their models.

"The original idea was to say, 'Can we quickly make use of these sorts of imaging resources and assess a hazard? What's the extent of the hazard? Where are we seeing the damage caused from the hazard?'"

Andre Coleman, PNNL Chief Scientist

The Solution:

While PNNL's focus was on near real-time response at the beginning, it evolved into proactive risk assessment that allowed emergency management and operations professionals to make decisions focused on mitigating risk well ahead of a disaster, such as vegetation management in and around critical infrastructure or replacing wood power poles with steels ones so they sustain through more extreme events.

This shift in focus was catalyzed by the increasing severity of wildfire events, especially historic fires that consumed more than 10 million acres in 2017 and again in 2020, prompting the recognition of wildfires as a national emergency.

"I think there was just this realization that fire is not the same as it once was," says Coleman. "It burns more intensely; it moves faster."

Advances in satellite technology have helped propel PNNL's current approach, which involves integrating data from various satellite sensors, including passive and active sensors, to derive actionable insights. Passive sensors capture reflected energy, while active sensors penetrate cloud cover, storms and smoke, enabling continuous monitoring even in adverse conditions. By leveraging open-access satellite data and commercially collected imagery, as well as collaborating with agencies like NASA, the United States Geological Survey and the European Space Agency, PNNL can tap into a robust data ecosystem for its analysis.

In addition, the development of PNNL's Rapid Analytics for Disaster Response (RADR) represents a significant milestone. The platform combines multi-modal imagery, AI and scalable cloud computing — thanks to partnerships with major cloud providers — with an infrastructure damage assessment tool. The tool helps users better understand the current impact and risk to infrastructure from wildfires, floods, hurricanes, earthquakes and more. Currently, once a provider collects and makes new imagery available – which typically takes about 4 to 6 hours – RADR can take that imagery, process and disseminate it in 7 to 10 minutes.

However, according to Coleman, newer satellite communications and increased automation are helping to drive this latency down to under an hour. In addition, this significant advancement in data collection and increase in the amount of data available is something that the RADR cloud platform has been planning for and is well-equipped to handle. For example, newer planned satellite sensors such as FireSat will provide a 20-minute satellite revisit, requiring a highly scalable system like RADR to generate analytics for these data.

In addition, PNNL's work has caught the attention of a widening group of stakeholders, including the Department of Energy's Office of Cybersecurity, Energy Security, and Emergency Response and the Joint Artificial Intelligence Center, which have played a crucial role in supporting PNNL's endeavors. Collaboration with public and private utilities has further enhanced the research landscape, ensuring practical applications and alignment with industry needs.

Key Takeaways:

The complexity of wildfire detection and prediction necessitates continuous refinement of AI and ML algorithms. PNNL's interdisciplinary team — comprising scientists, developers and cloud architects — remains dedicated to enhancing system capabilities. With the upcoming availability of specialized earth observation satellites, like NASA's SWOT and NISAR, PNNL aims to expand its predictive capabilities further and extend its reach to a broader audience. "The idea is to keep adding and improving our algorithms, incorporate more sensors as they come online, continue to mature the capability and really work to get this out to as many end users as we can," says Coleman.

PNNL's efforts are also expanding beyond natural threats, to include a more holistic approach to safeguarding critical infrastructure, including the rise of physical as well as cyber attacks on critical infrastructure. By integrating environmental, cyber and physical security considerations, PNNL strives to build a more resilient energy grid capable of withstanding diverse challenges.

"A big role for us at the lab is to look at all these kinds of dimensions and say, 'How do we protect the grid? How do we make it more resilient as our reliance on energy continues to increase?'' says Coleman. "Through the various teams that are working on that, we have the opportunity to interface and consider how to move forward and build a better, more resilient system."

As PNNL continues to innovate and collaborate, the future holds promise for a more resilient energy infrastructure. By leveraging the power of AI and ML, PNNL stands poised to address the evolving threats posed by natural disasters and enhance grid resiliency for generations to come.

sas viya Workbench

A cloud-based, on-demand environment for development, modeling and coding

Get to know Viya Workbench

Explore features, resources and more in the environment built for developers.

Analyst viewpoint

See what IDC is saying about the latest SAS Viya offerings.

Viya Workbench documentation

Take a deeper dive into how developers can use the software.

Tech Spotlight Series
Innovation in Government

Red Hat | carahsoft

Enhancing AI Security & Deployment with Red Hat OpenShift AI

PRYON

Wingman Takes Flight: How The Department of the Air Force Digital Transformation Office Streamlines Information Delivery for Military and Civilians

The Department of the Air Force Digital Transformation Office (DAF DTO) has raised the bar for self-service excellence by introducing DTO Wingman – a public-facing chatbot powered by Pryon's trusted and secure GenAI technology.

Designed to deliver instant, trustworthy answers, DTO Wingman taps into over 150,000 pages of Department of Defense (DoD) content, redefining self-service for personnel and civilians alike.

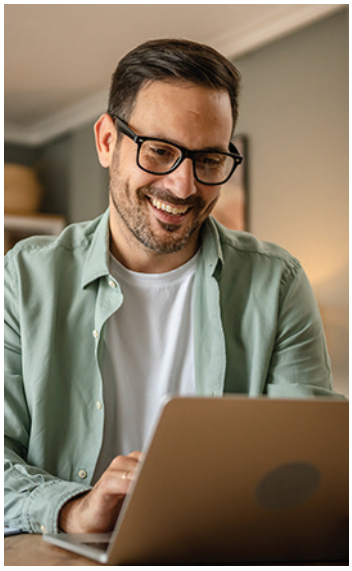

The Challenge:

- Managing an extensive library of public-facing materials, including DoD directives, instructions, manuals, and other open-source publications.
- Providing a seamless, user-friendly self-service experience for a broad audience of personnel and civilians.
- Staying ahead in AI innovation without compromising safety and security.

The Solution:

Expert Guidance:
Pryon provided valuable guidance to ensure a secure and efficient start with AI.

Data Integration:
Thousands of manuals, directives, ePubs, and other public-facing DoD documentation were ingested to build a reliable repository for accurate responses.

Custom Chabot:
DTO Wingman offers an intuitive, conversational search experience for users visiting the DAF DTO website.

Launch:
DTO Wingman launched in 2024.

Continuous Evolution:
What started as a chatbot has now evolved into a multifunctional toolset, including features like report generation, email composition, and voice-to-text capabilities.

Key Takeaways:

24/7 Access:
Around-the-clock support ensures accessibility anytime, anywhere.

High Scalability:
The system seamlessly processes over 150,000 pages, 8,000 URLs, and 2,700 documents, ensuring seamless performance at scale.

Multi-Lingual Support:
With functionality in 80 languages, Wingman ensures global usability.

Customizable user experience:
Users can personalize their experience, with customization options such as:

- Selecting specific information repositories to retrieve information from.
- Modifying response length and style.
- Receiving bullet point sources with direct links to source materials.
- Generating more creative outputs like emails or reports.

By combining Pryon's AI expertise with DAF DTO's mission-driven focus, DTO Wingman represents a significant leap forward in resource accessibility. This initiative not only simplifies access to critical DoD content but also sets a new benchmark for trust, security, and user experience in government technology.

The DAF DTO is proving that innovation is not just about technology—it's about creating real solutions for real challenges. DTO Wingman is more than a chatbot—it's a force multiplier for knowledge, decisions, and efficiency.

Rhode Island's Unemployment Insurance Team Delivers Data Transparency to Claimants Using AWS

The State of Rhode Island Department of Labor and Training (Rhode Island DLT) delivers temporary income support in the form of unemployment insurance (UI) and reemployment assistance to residents. During the COVID-19 pandemic, state unemployment agencies were challenged to deliver unemployment benefits at the speed and scale required.

The Challenge:

The Rhode Island DLT received more than 140,000 initial claims for the UI program in the first 45 days after the governor declared a state of emergency. By comparison, the state—with a population of slightly over one million people—received 107,000 initial claims during the peak of the Great Recession. This surge led to an increase in the requests for assistance from residents as they navigated a complex UI system, many for the first time.

The Solution:

In early 2020, the Rhode Island DLT used AWS to build a system that could accept the Pandemic Unemployment Assistance applications that the Coronavirus Aid, Relief, and Economic Security (CARES) Act introduced. Following the success of this deployment, the department was determined to continue innovating to transform the state's UI system using AWS. "Unemployment insurance is a black box for people who collect benefits," says Abby McQuade, senior advisor and chief innovation officer for the Rhode Island DLT. "We wanted to give people more information and increased transparency."

In May 2020, the Rhode Island DLT participated in a "Working Backwards" workshop facilitated by the AWS WWPS DI team. During these sessions, the Rhode Island DLT developed the idea for its "pizza tracker" tool—an online portal that will track the status of a claimant's UI benefits from the moment they submit their claim through final payment. Not only will this tool educate claimants about the UI process and help them track their UI claims but it will also help them quickly resolve potential issues by immediately surfacing problems. Following the "Working Backwards" workshop, the Rhode Island DLT engaged AWS Professional Services—a global team of experts that can help businesses realize their desired outcomes on AWS—and an AWS Partner to conduct the user interface design and develop the "pizza tracker."

UI Online will use AWS AppSync, a fully managed service that makes it simple to develop scalable GraphQL APIs, to respond to API calls and read from the department's legacy UI database. By making this data accessible, the Rhode Island DLT expects that the new portal will reduce the number of calls coming into its call center. "In a single day, we receive about 110,000 calls from about 6,000 individuals," says McQuade. Because many callers are only asking for UI claim status updates, UI Online will reduce the need for these inquiries, giving claimants peace of mind and easing the burden on employees who were previously spending the majority of their time managing these requests. "Any time government can make itself more transparent, that serves the public good," says McQuade.

In addition to improving the claims process, the Rhode Island DLT wanted to do more to help people who are unemployed or underemployed find good, stable jobs. "We have staff members who are experts at performing this work," says McQuade. "But we can't possibly have enough staff to help everyone at the level we want." After participating in a second

"Working Backwards" workshop facilitated by the AWS WWPS DI team, the department designed two cloud-based digital staff members in the form of chatbots. The chatbot named Hope provides answers to claimants' UI questions, and Skipper supports those searching for reemployment opportunities.

The department uses Amazon Lex, a service for building conversational interfaces into an application using voice and text, to power the chatbots' conversational interfaces.

"Using AWS, we can use technology to do the simple tasks so that our employees can use their time to work on complex problems," says McQuade. "Taking somebody from unemployment all the way through to reemployment—that's what we want to do as a department." By using the AWS WWPS DI team and AWS services, the department hopes to scale its team's ability to help thousands of workers in the state find income support and valuable opportunities.

Key Takeaways:

The next step in the Rhode Island DLT's innovation journey includes designing an internal dashboard to consolidate all of a claimant's UI information, as well as other data the state has about that resident, into a single view. Previously, the Rhode Island DLT customer service representatives had to navigate multiple applications to gain a holistic picture of a claimant to resolve issues. The time it took to look through this information—and the manual processes involved—created delays in getting claimants their benefits. With claimants' information readily available through this dashboard, customer service representatives will be able to resolve issues quickly and make sure that unemployed residents get the help they need.

Using the AWS WWPS DI team and the "Working Backwards" mechanism, the Rhode Island DLT is deploying scalable cloud-based tools to provide more holistic support for unemployed workers. Additionally, by making the UI system more transparent and efficient, the department is freeing up its staff to focus on solving more complex UI issues for claimants.

"If you work in government, you've dedicated yourself to serving the public," says McQuade. "The way we're doing that is by moving forward with the people we aim to help, using the technology that's available today on AWS to serve them better."

alteryx

Transforming University Analytics with Alteryx: A Case Study of the University of Pittsburgh

Higher education institutions are grappling with the challenge of managing and analyzing vast amounts of data. The University of Pittsburgh (Pitt), a nationally ranked Research institution, is an excellent example of how embracing advanced analytics tools can transform this challenge into an opportunity. This case study delves into how Alteryx has empowered Pitt to streamline operations, enhance decision-making, and drive impactful outcomes across the institution, from IT and research to athletics and fundraising.

The Challenge:

Pitt's Information Technology (Pitt IT) department initiated an analytics enablement program to create cross-functional subject matter experts equipped with self-service analytics tools. Previously, the manual extraction, integration, and refinement of datasets for business reporting and decision-making were time-consuming tasks for Pitt's analytics teams and data scientists. The need for a more efficient solution was clear, prompting Pitt IT to explore tools to streamline these processes and empower users across the university to handle their data needs independently.

The Solution:

Pitt IT began by trialing Alteryx Designer, which quickly proved its value in simplifying data preparation and curation. Recognizing its potential, Pitt expanded its use of Alteryx Designer to automate various time-consuming data workflows. The implementation of Alteryx Server further extended these capabilities, providing a scalable solution for the entire university. This allowed multiple departments to automate complex data analysis and business reporting processes, significantly reducing the manual effort previously required.

Key Takeaways:

The adoption of Alteryx at Pitt has led to notable improvements in several areas:

Time Savings and Automation:
Departments have reported significant time savings by automating data workflows that previously took weeks to complete manually. For instance, the Athletics Department estimated saving about 100 hours of manual work by using Alteryx to generate a stadium seating map.

Improved Decision-Making:
During the COVID-19 pandemic, Alteryx was crucial in integrating diverse data sources to provide timely and accurate insights, supporting critical decision-making processes for university leaders.

Enhanced Data Governance:
Alteryx Server's data governance features have enabled the central analytics team to manage sensitive information securely, ensuring stakeholders receive only the necessary data while protecting privacy.

Support for Research and Education:
Alteryx is now being integrated into the academic curriculum, allowing students and researchers to leverage the platform for their projects. Events like the Alteryx SparkED Datathon have provided experiential learning opportunities, improving students' analytical skills and preparing them for future careers.

The success of Alteryx at Pitt extends beyond administrative efficiencies. The platform's flexibility and ease of use have encouraged widespread adoption across various departments, supporting initiatives in research, student services, and community engagement. Alteryx has become an indispensable tool in Pitt's mission to foster academic excellence, research innovation, and community service by automating routine tasks and enabling sophisticated data analysis.

By leveraging Alteryx, the University of Pittsburgh has enhanced its operational efficiency and paved the way for future innovations in data analytics. Alteryx could be the key to unlocking your data's full potential if you want to empower your organization with advanced analytics capabilities.

databricks

Streamlining the Path to Citizenship with Data

The number of immigration- and citizenship-related applications has skyrocketed over the past decade across naturalizations, green cards, employment authorizations and other categories. With millions of applications and petitions flooding the USCIS, processing delays have reached crisis levels — with overall case processing times increasing 91% since FY2014.* Core to these issues was an on-premises, legacy architecture that was complex, slow and costly to scale.

By migrating to AWS and Databricks, USCIS adopted a unified approach to data analytics with more big data processing power and the federation of data across dozens of disparate sources. This has unlocked operational efficiencies and new opportunities for their entire data organization to drive business intelligence and fuel ML innovations designed to streamline application and petition processes.

The Challenge:

Charged with administering the nation's immigration system, USCIS is focused on enabling people to obtain work authorizations, apply for immigration benefits, and seek asylum, and it allows U.S. employers to fill critical workforce gaps. Meeting this mission requires the efficient processing of millions of immigration-related applications.

The USCIS engineering team saw an opportunity to leverage data and analytics to automate certain processes and accelerate processing times. However, their on-premises technology stack, made up of legacy systems, proved to be highly complex to manage and was overwhelmed by the scale of data.

"We were using Informatica for ETL, and the pipeline was fairly brittle," explained Shawn Benjamin, chief of data and business intelligence at USCIS. "As a result, we had a lengthy development cycle and longer workflows. It made it impossible to deliver relevant data in real time." This level of performance was unacceptable as they supported over 2,300 data analysts and data scientists, who all tried to access data across dozens of different sources.

The data opportunities at USCIS were boundless. From a data science perspective, they were looking for opportunities to use predictive analytics to answer difficult questions, such as, Which applications are being submitted the most? How many applications can they expect in future years? And what is the probability of someone not showing to an appointment? They were also looking for ways to streamline processes by digitizing applications and using NLP to better evaluate interviews. The lack of a scalable data science platform kept their data scientists from being able to deliver on these use cases.

USCIS realized they needed to modernize their data infrastructure by migrating to the cloud and adopting a unified platform that could harness all of their data for easier access and ingestion, while enabling downstream data analytics and ML.

The Solution:

Databricks provided USCIS with significant impact where it mattered most — faster processing speeds that enabled data analysts to deliver timely reports to decision makers — and that freed up data scientists to build ML models to help improve operations. Leveraging the efficiencies of the cloud and Delta Lake, they were able to easily provision a 26-node cluster within minutes and ingest tables with 120 million rows into S3 in under 10 minutes. Prior to Databricks, performing the same processes on Informatica would have taken somewhere between two to three hours.

Databricks has also served as a transformation agent to their data warehousing strategy, leveraging Delta Lake to create a lakehouse that federates all their data regardless of where it's stored for downstream consumption. In fact, they have been able to migrate 2,000 tables from Oracle to Databricks in less than a week.

Through the use of interactive notebooks, data scientists can easily collaborate with each other and across other data teams within the organization. "Notebooks allow multiple groups to work together from a single point. It eliminates having that swivel chair activity as people can work directly in the interface together." And with MLflow, they are able to easily build multiple ML projects and experiments with ease.

databricks

Key Takeaways:
Since migrating to the cloud and integrating Databricks into their data analytics workflows, USCIS is able to make smarter decisions that help streamline processes and leverage ML to reduce application processing times. These newfound efficiencies and capabilities have allowed them to scale their data footprint from about 30 data sources to 75 data sources without issue.

USCIS now has the ability to understand their data quicker, which has unlocked new opportunities for innovation. As an example, Benjamin cites that it used to take them a full business day to run a very complex query. With Databricks, they are able to run the same query in 19 minutes — a 24x performance gain. This meant that they were spending far less time troubleshooting and more time creating value.

Even processing speeds for Tableau dashboards saw a marked improvement, which was important as Databricks supports over 6,000 Tableau dashboards. Benjamin's team noticed that running some of their dashboards used to take around 15 minutes — and sometimes the queries would fail altogether due to high data load. When they ran the same queries with Databricks, they were able to return the dashboards in under 15 seconds. Faster access to data insights means smarter decision-making in near real time.

Finally, the data science team is now able to leverage all of their data to help USCIS make more informed decisions and streamline operations. For example, they have implemented eProcessing, which has replaced paper applications with electronic applications, greatly improving operations and speeding up processes. Benjamin said,

"Whether we are trying to predict the probability of a no-show to an appointment, streamline a manual process or perform sentiment analysis of survey data, the opportunities in front of us are now endless."

By liberating their data and making it easy for anyone to leverage it, the agency has been able to increase their user base by 3x. With more data, more resources and more performance, the agency has since implemented many new programs, including Electronic Immigration System (ELIS), eProcessing of applications, operational and case status reporting, fraud detection, refugee asylum, international operations (RAIO), and forecasting. With Databricks serving as a crucial factor in enabling USCIS to extract data from anywhere and feed it to whoever needs it at any time, they are well positioned to continue to drive innovation and operational efficiency in order to facilitate lawful immigration to the United States.

Domino

U.S. Navy Makes Mine Detection Intelligence More Trustworthy with Domino

For years, the U.S. Navy relied on machine learning (ML) models to support underwater target threat detection by unmanned underwater vehicles (UUV). However, it lacked a way to monitor and improve ML model performance at scale post-deployment because models were slow to adapt to changing underwater conditions or enemy tactics. Without the ability to monitor or maintain ML models, mission-critical systems risk producing inaccurate or unreliable intelligence that drives decision-making.

"One of the things that we discovered through that process was that as these unmanned vessels were out there doing what they did, [when] they came back, we needed to be able to update their software based on what they just learned," said Commander Eli Ford, Principal Assistant Program Manager, Underwater Explosive Ordnance Disposal, PMS 408 (Expeditionary Missions).

The Challenge:

In 2022, the Navy partnered with the Defense Innovation Unit (DIU) to improve ML model performance gaps by leveraging commercial technology that could be integrated into the Navy's existing systems. Using DIU's Commercial Solutions Opening (CSO) process, the team identified commercial machine learning operations (MLOps) vendors who could collectively provide a pipeline to track, modify, and redeploy underwater target threat detection ML algorithms.

"The partnership between the U.S. Navy and DIU has opened the pathways to commercial innovation and has produced a MLOps toolset that ensures our solutions adapt and evolve at the speed of tactical relevance, placing an enduring capability in the hands of the warfighter. Nowhere was this more apparent than at an Integrated Battle Problem in the Fall of 2023 where sensor data was collected, new models were trained against an expanded threat environment, and updated detection models were deployed through the AMMO ecosystem back into the hands of the operator," said Ford.

To solve this, the U.S. Navy's Expeditionary Missions team partnered with the Defense Innovation Unit to identify commercial technologies, including Domino, that deliver a potentially transformative effect in these maritime scenarios carried out by Automated Machine Learning for Mine Countermeasures Operations (AMMO). AMMO capabilities helped address the subsea/seabed warfare (SSW) mission by improving the Navy's ability to detect targets of interest in undersea environments. AMMO also allows the Navy's ML developers to focus on model creation and refinement tasks, allowing the tools to handle the critical aspects of deployment, monitoring, and governance.

Domino

The Solution:

The U.S. Department of Defense (DoD) built a successful AI model prototype to close a critical gap in underwater mine intelligence. The Domino Enterprise AI Platform helped quickly deploy and retrain ML models reliant on data collected from automated underwater mine detection vehicles, improving undersea warfare essential to U.S. success in the Pacific and other contested environments.

The Domino Enterprise AI Platform is now an integral component of a modern MLOps pipeline for mine countermeasures (MCM) AI-based model development and deployment in the cloud. The Domino Enterprise AI Platform helps the U.S. Navy deploy and retrain automatic target recognition (ATR) models dozens of times faster than before — enabling naval decision-makers to have higher confidence in field-generated insights from unmanned underwater vehicles. Automatic target recognition is the ability of an algorithm or device to recognize targets, such as mines, based on data obtained from sensors.

As a modular, open, and extensible platform hosted in AWS GovCloud (US), Domino enables the U.S. Navy's Project AMMO's distributed teams to deploy models faster with built-in observability and reduced time-to-impact. Users can leverage the best open source and commercial tools of their choice in a highly governed way, on one platform, to generate and analyze ATR field-generated insights. In addition, Domino's built-in model portability allows these teams to easily convert and export models into external formats to far-edge platforms. It also allows users to securely access and share structured and unstructured data from anywhere.

"It's highly encouraging to see Domino, the Navy, and DIU partnering to rapidly field commercial tech sector capabilities in order to deter conflict. This type of teaming, with Domino's core MLOps technology at its center, is an example of a rapid, repeatable, and scalable template that will significantly improve our nation's defenses," said Brigadier General Bobby Kinney.

Kinney is deeply familiar with both the needs of the DoD and its AI-driven missions, having previously served as the Missions Director at the Joint Artificial Intelligence Center (JAIC) under the Office of the Secretary of Defense (OSD). During his time at JAIC, he was responsible for leading AI coordinating activities and accelerating the delivery of $250M a year of high-impact AI projects across the Department of Defense (DoD).

"This type of teaming, with Domino's core MLOps technology at its center, is an example of a rapid, repeatable, and scalable template that will significantly improve our nation's defenses."

Brigadier General Bobby Kinney

Key Takeaways:

Autonomy is a key feature of next-generation warfighting, with trusted AI enabling those systems. Commercial AI solutions like Domino are critical to accelerating the DoD's ability to field autonomous capabilities across domains, particularly underwater. As an open, scalable, yet highly governable platform, Domino helps the U.S. Navy's Project AMMO deploy cutting-edge commercial tech to enable the fleet with rapidly updated models that deliver a higher level of trust in intelligence to power their missions.

With Domino, the U.S. Navy's AMMO team was able to decouple hardware and software updates to dramatically accelerate updates to ATR models at the edge. This results in more accurate imagery and sonar-based intelligence that can be trusted and actioned by naval decision-makers for next-generation warfare advantage.

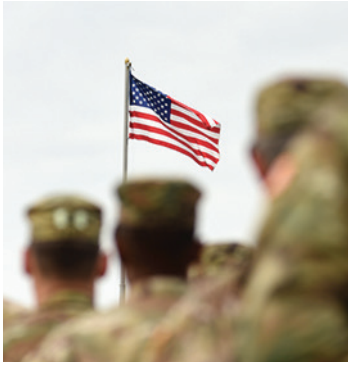

Using Domino as the factory for integrating four other commercial technologies, three different contracted teams were able to accelerate their work, reducing the time to deploy AI models at the edge — to two weeks down from six months — and retraining the models in two weeks rather than 12 months.

"Working with multiple international partners, we demonstrated the ability to deploy and update our automatic target recognition models at the speed of operational relevance, while simultaneously remaining confident in their performance. This embodiment of human-machine teaming gives our sailors a powerful tool in their arsenal and increases confidence in their systems," said Commodore Shaun Lieb, Commander, Task Force (CTF) 75.

The U.S. Navy's Project AMMO also sets the stage for the DoD to deploy and manage ML algorithms effectively at scale, through MLOPs practices that ensure the safe, reliable, and effective operation of all ML models that advance trust in new AI technologies.

"I'm proud our warfare community could be the pathfinder and exemplar for machine learning ops, tooling, and processes that will improve many other Navy and joint service warfighting capabilities in the age of artificial intelligence," said Commodore Shaun Lieb, Commander, Task Force (CTF) 75.

Domino

Mission-driven AI. Secure and governed.

The U.S. Navy uses Domino's MLOps platform to harness data science and AI for mission-critical decisions.

AMERICA'S NAVY

"I'm proud our warfare community could be the pathfinder and exemplar for machine learning ops, tooling, and processes that will improve many other Navy and joint service warfighting capabilities in the age of artificial intelligence."

—Former Commodore
CTF 75-Navy Expeditionary Forces Command Pacific

Transform Your Mission with AI

See how the U.S. Navy makes mine detection more trustworthy with Domino.

domino.ai

NUTANIX

Pennsylvania House of Representatives Democratic Caucus Boosts IT Infrastructure with Nutanix Partnership

The Democratic Caucus of the Pennsylvania House of Representatives is a legislative group composed of the Democratic members elected to the Pennsylvania House of Representatives. The caucus represents the interests and concerns of their constituents throughout the state, advocating for policies and legislation aligned with Democratic values and priorities.

With approximately 700 employees, including around 70 in the IT department, they focus on crafting and supporting laws that benefit Pennsylvania's residents, addressing issues such as education, healthcare, economic development, and public safety. The caucus utilizes technology and innovative solutions to effectively communicate with constituents, manage legislative processes, and provide responsive and efficient public service.

The Challenge:

The Pennsylvania House of Representatives Democrats serve the state's over 12 million residents, relying heavily on technology to fulfill this mission. Each of the 102 state representatives in the caucus represents about 60,000 citizens, using an advanced software application to engage with and respond to their constituents' needs. "Our network infrastructure hosts the software system that we have developed in-house to aid district offices and staff, relay the needs and wants of constituents, and assist them with services like vehicle registration, voter issues, and community issues," said Yuri Topolnicki, Senior Network Administrator at Pennsylvania House of Representatives Democratic Caucus. "This constituent service tracking software is the core of what makes our legislative functions operate."

"The Pennsylvania House of Representatives Democrats support everything from public-facing websites to internal applications that enable representatives and their staff to maximize their interactions with constituents. While these are core competencies for our team, a lot goes into the infrastructure supporting this specific effort," stated Astree.

Without a reliable platform, accessibility across all internal and external systems would disrupt services, preventing them from maintaining continuous operation. "Uninterrupted access to government services is critical, especially during a crisis," noted Astree.

The Solution:

To address these challenges, The House of Representatives Democratic Caucus chose to standardize its operations on the Nutanix Enterprise Cloud Platform. This decision was driven by the need for a scalable, reliable, and flexible infrastructure that could support their growing organization and new services.

The solution adopted by the Pennsylvania House of Representatives Democratic Caucus includes the Nutanix Cloud Platform, utilizing AHV and ESXi hypervisors, along with Nutanix Unified Storage (NUS), File Analytics, and Micro-segmentation/Flow Network Security for enhanced security and management. This robust setup supports various applications and workloads, including Active Directory, FTP servers, web servers, virtual machines for the Pennsylvania House website, and SQL workloads and clusters running in a virtualized environment.

The implementation of Nutanix's hyperconverged infrastructure was swift and effective, with the team migrating their VMware ESXi workloads and approximately 100 Microsoft Hyper-V VMs to the Nutanix AHV hypervisor in a matter of days. "Once the hardware was onsite and integrated into our network, the process was fairly painless," explained Topolnicki. "We had a fully operational hyperconverged solution in place within about a day, and we were ready to begin migrating our existing workloads to the new platform. Nutanix provided step-by-step guidance on what to do and how to do it." The Nutanix Prism management plane provides a centralized interface for all administration tasks, enhancing efficiency and reducing complexity.

> *"We had a fully operational hyperconverged solution in place within about a day, and we were ready to begin migrating our existing workloads to the new platform. "*
>
> Yuri Topolnicki, Senior Network Administrator, Pennsylvania House of Representatives Democratic Caucus

Key Takeaways:

Enhanced proactivity:
The Nutanix platform has allowed the IT team at the Democratic Caucus to be more proactive rather than reactive in their daily operations. By consolidating various infrastructure components onto Nutanix, they have reduced the number of technical issues they face, which enables them to focus on more strategic initiatives. Topolnicki explained, "Instead of having to deal with incidences, we have the ability to take those time cycles and turn them toward innovating on our applications and use cases."

Improved reliability and stability:
Nutanix has played a crucial role in improving the reliability of The Democratic caucus of the Pennsylvania House of Representatives IT environment by reducing service disruptions. This stability is crucial for ensuring that representatives can consistently serve their constituents without technical hindrances. "We appreciate the time saved and the sheer reliability of the system," said Topolnicki. "The reliability of the way the Nutanix Control VMs work together, and the rock-solid hypervisor, has given us the best of both worlds."

Time and cost savings:
The House of Representatives Democratic Caucus has generated substantial time and cost benefits from the Nutanix infrastructure. By utilizing Nutanix, they have grown their infrastructure organically by 10-15% without significant additional investments. This scalability is crucial for efficiently managing growth and maintaining service quality. Astree noted, "The investment that we're making in the Nutanix ecosystem continues to pay dividends, meaning fewer investments are needed in other areas."

Cohesive data protection and disaster recovery:
By centralizing its data on the Nutanix Cloud Infrastructure (NCI), the client has improved its data protection and disaster recovery capabilities. They no longer need to worry about whether different parts of their infrastructure are adequately backed up, as Nutanix supports seamless replication and backup across their environment. Yuri added, "We've been able to make sure that all of those workloads, all of those things that we put in the Nutanix hyperconvergence platform, we can protect, touch, and replicate, thanks to Nutanix."

Next Steps:
Looking ahead, the Pennsylvania House of Representatives Democratic Caucus plans to expand its use of Nutanix solutions. They are considering the potential for further cloud integration, additional clusters, and the use of Nutanix Unified Storage (NUS) to enhance their data center capabilities.

Additionally, they are excited about the maturation of AI integration points and the potential for advanced features like a virtual Copilot to optimize their IT operations further. "The idea of having AI do things for us and with us, like an IT partner, is very appealing," said Yuri Topolnicki, drawing a parallel to the AI assistant Jarvis from the Iron Man movie franchise.

The client remains committed to leveraging Nutanix's innovative and reliable solutions to continue serving its constituents effectively and efficiently.

fiddler

Defense Innovation Unit Issues Success Memo to Fiddler AI

We are honored to share that the Department of Defense (DoD) Defense Innovation Unit (DIU) has awarded Fiddler with a Success Memo for completing the Automated Machine Learning for Mine Countermeasures Operations (AMMO) MLOps prototype with the U.S. Navy.

The DIU in collaboration with the Navy posted Machine Learning Operations (MLOps) as an area of interest on May 31, 2022 and selected Fiddler as one of the awardees of the prototype proposal.

The Challenge:

The goal of the AMMO prototype is to build an MLOps pipeline toolset capability to rapidly retrain the Automatic Target Recognition (ATR) ML models powering the U.S. Navy's mine countermeasures (MCM) application, thereby empowering the Navy to quickly adapt to changing undersea threats in diverse theaters.

The Solution:

As part of the implementation, Fiddler was initially installed on the security-enhanced AWS GovCloud environments and then successfully validated for all three phases of Government Acceptance Testing (GAT), including functionality confirmation, security posture, and deployment automation.

In addition, the Fiddler AI Observability platform for federal agencies was recently deemed Awardable by the DoD Chief Digital and Artificial Intelligence Office (CDAO)'s Tradewind Solutions Marketplace to accelerate the procurement and adoption of AI/ML, data, and analytics at the DoD.

Key Takeaways:

The following key capabilities of Fiddler have been validated and are leveraged in the AMMO workflow:

Image Explainability:
Enterprise grade on-demand explainability to help model developers and in the future, mission operators, understand and trust the decisions of classification and object detection models.

Image Monitoring:
Patented data drift monitoring of image embeddings integrated with visual debugging of these embeddings using UMAP. This helps identify operational changes in the AI model's behavioral and perform root cause analysis of any issues quickly.

These capabilities were also validated in a live demonstration of the AMMO prototype which showed a 97% decrease in the time needed to update the ATR models with this MLOps toolset — from six months to a few days.

FinchAI

Superpowered Tools for Analysts

At Finch AI, we think like analysts. So we build tools that accelerate their workflows, that never get tired, and that dramatically improve outcomes. Staffed with data scientists, PhDs, former analysts and longtime defense and intelligence partners, every one of our innovations centers on turning human generated text into something a machine can understand quickly, accurately and at-scale – and that an analyst can trust.

The Challenge:

Our National Security customers must conduct comprehensive due diligence to detect risks within the U.S. government's supply chain. Traditional corporate risk analysis relies on structured data from state, local, and national databases, which are only updated periodically. This reactive approach can miss emerging threats, such as foreign ownership risks, regulatory non-compliance, cyber vulnerabilities, and other external factors. Identifying these risks in real-time is critical to safeguarding National Security interests.

The Solution:

Finch AI supported this due diligence use case by combining structured database information with AI-monitored, real-time, unstructured data set inferences. Using our NLP-engine, analysts surfaced people, places, organizations and their relationships from streaming news, broadcast and social media feeds. These enrichments are the foundation for our unique retrieval-augmented generation (RAG) approach wherein word embeddings are combined with entities, sentiment, temporal and geospatial search in a single query. Using semantic routers and AI agents, we determined what data stores to query and how. These agentic techniques enabled us to automatically find hidden risks. We then used our AI recipes and workflows to build rich, comprehensive and context based Insight Reports that are instantly shareable.

Key Takeaways:

By integrating our superpowered tools, National Security analysts can now proactively monitor supply chain risks, automate investigative workflows, and validate intelligence findings with traceable sources. The AI continuously scans for changes in corporate structures and risk factors, ensuring always up-to-date intelligence while optimizing efficiency — allowing the analyst to do more, with more.

dynatrace

State of Michigan Delivers $4.4 Million Efficiency Savings Over 5 Years with Dynatrace

The Michigan Department of Technology, Management & Budget (DTMB) is a centralized IT department serving 20 State of Michigan executive agencies. It is responsible for providing stable, secure, and easy access to the business, financial, and technical services that meet the needs of state employees, citizens, and customers. Its key goal is to create trusted relationships with its various customers and drive value for the department and its agency partners by providing reliable products and services at the lowest possible cost. To deliver on this, DTMB implemented a cloud strategy, enabling agency partners to deploy solutions that support their missions at the best value.

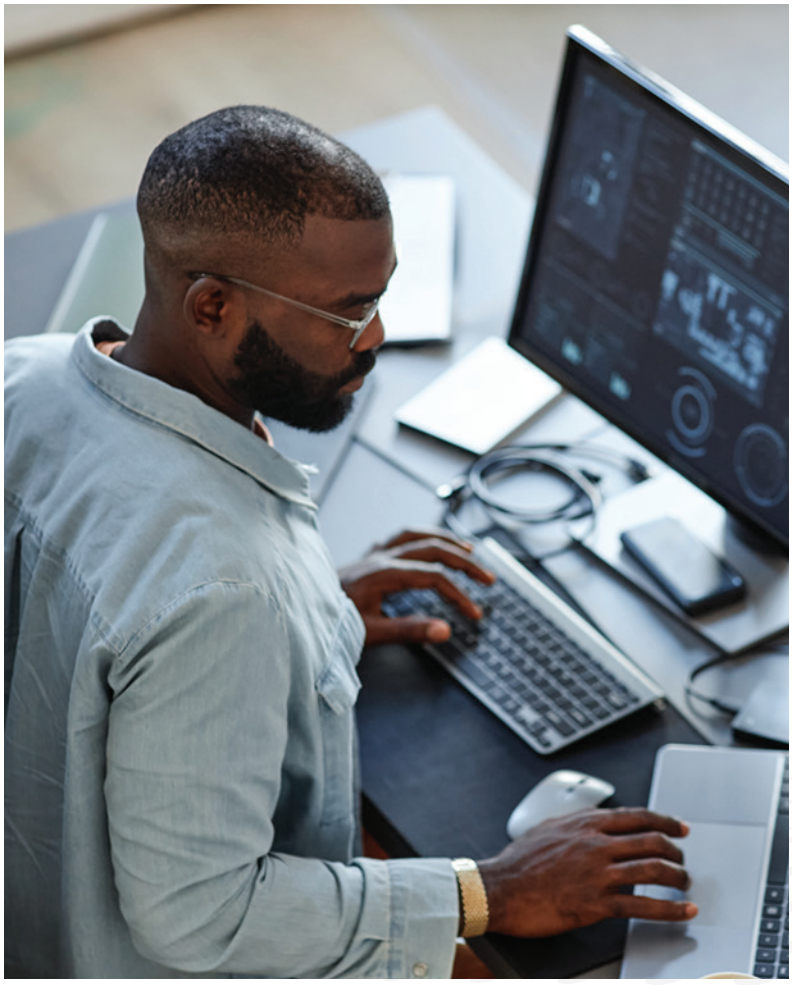

The Challenge:

As the State of Michigan increases its use of cloud services, the complexity of its technology stack is growing exponentially. Its services now span a hybrid, multicloud architecture, which makes monitoring and managing citizen and customer experiences more difficult. DTMB needed to ensure its application support teams have the insight needed to triage problems and react quickly to prevent service incidents and disruption. This is especially important when it comes to critical services, such as hunting and fishing licensing and those operated by the Secretary of State, which generate millions of dollars of revenue.

The Solution:

After evaluating the market for unified observability solutions, the DTMB team chose Dynatrace due to its highly automated approach and the maturity of its AI capabilities. For critical-user journeys, the team worked with its agency partners to set up synthetic monitoring and alerts that help them to identify potential problems before they occur. This means they can proactively work with cloud providers and software vendors to resolve issues before they result in an outage.

"Our vision is to provide IT services and products that deliver citizen outcomes so seamlessly that users hardly know IT exists," said Tammy Zbojniewicz, manager of the DTMB Enterprise Operations Center. "That became more difficult as we embraced the cloud, as we'd never been able to get a full picture across that entire ecosystem and so were reliant on service provider updates. As a result, our incident management process often involved up to 30 support staff spending hours or days searching for the root cause. Some incidents even went on for weeks. That created an unacceptable risk, as for some of our revenue-generating services, 20% of customers will not come back and try again if the system isn't working. With Dynatrace, we can show tangible value right out of the gate, by pinpointing issues and lasering in on the root cause within minutes."

"Our vision is to provide IT services and products that deliver citizen outcomes so seamlessly that users hardly know IT exists."

Tammy Zbojniewicz, DTMB Enterprise Operations Center Manager

Key Takeaways:

Doing more with less:
Dynatrace has put the DTMB team on track to deliver a total of $4.4 million in cost savings for the State of Michigan over the next five years. These savings include those achieved from the ability to use a single solution to manage all infrastructure and apps, which meant the team was able to consolidate multiple outdated solutions to save a total of $2.4 million. Dynatrace also enabled DTMB to reduce the number of contractors and staff it relies on to support its monitoring capabilities. Vacancies from retirements and new opportunities no longer need to be filled, which equates to an additional $2 million in savings on salaries.

Greater workforce satisfaction:
Dynatrace's automated approach has freed support teams to focus on more satisfying tasks, such as developing new digital services, rather than repetitive and manual work to resolve problems. This helps teams see more value in their everyday work, reducing the cost and disruption caused by employee churn.

More reliable public services:
The synthetic tests and customized alerts that DTMB set up with Dynatrace enable it to proactively identify, understand, and resolve the cause of any issues across both internal and external platforms. DTMB quickly saw the benefit of this when it used these insights to work with a vendor on resolving a configuration error that was impacting a critical service for the Secretary of State. As a result, the vendor was able to optimize the application, and the department can now provide a more reliable and consistent service for constituents.

Accelerated adoption:
Since DTMB supports 20 executive agencies, it is vital that any services it provides are fully compliant with key U.S. regulations for data security and privacy protection. Dynatrace's FedRAMP certification has made it significantly easier to drive wide adoption of an observability platform across the State of Michigan by giving agency partners confidence that the solution provides the compensating controls needed to maintain security and compliance. As a result, Dynatrace has been extended to all the state's critical services, monitoring more than 200 applications to date.

"Dynatrace enables us to do so much more with less," continued Zbojniewicz. "We've downsized our support model, reducing the need for new postings. Since we're all now looking at the same data, we can make decisions, implement change, and resolve issues faster. For example, one of our agency partners had been experiencing intermittent issues with one of its systems, with slow performance and citizens finding it unavailable. When DTMB was invited in, we deployed Dynatrace, and within five minutes, we'd pinpointed the cause of the issue, which had been eluding the agency for three weeks. That has a huge impact for everyone involved and supports our ability to create value and trust in the department and with the customers we serve."

polimorphic

AI to Power Service-First Governments

Deliver the highest quality resident services with software designed for the public sector.

AI Search & Chatbot

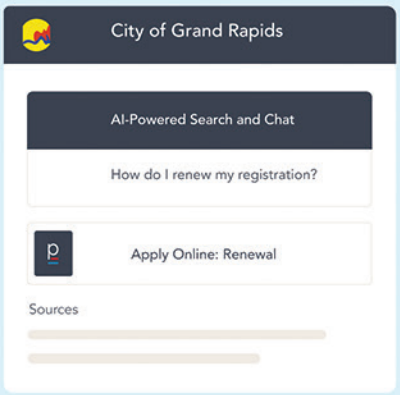

Voice AI

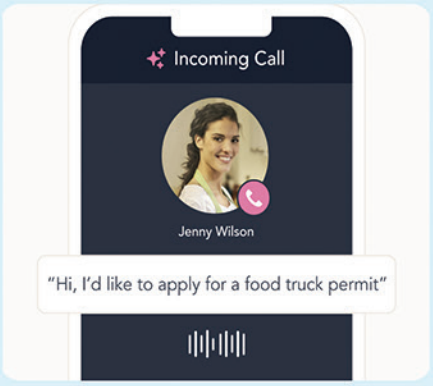

CRM & Workflows

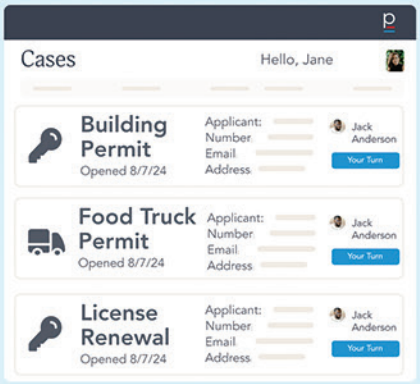

"Polimorphic has provided us a better way to get feedback from our citizens. The customizable solution gives us the opportunity to interact with citizens in a more efficient and effective way across our entire organization."

Marche Pittman, MSMIIT
County Manager, Polk County, NC

70% Reduction of staff time on paper based processes

After hours website and call support

1 week onboarding per department

Weeks saved annually by eliminating paper processes

75% Reduction in traffic to offices for resident requests

Learn More at Polimorphic.com

eightfold.ai

Precision meets purpose.

Eightfold AI-native Talent Intelligence for government

Eightfold Talent Intelligence delivers **purpose-built** AI insights and automation derived from your organizational data, market data, and user interactions.

Gain a **dynamic understanding** of your talent in the context of work — across all stages of the talent life cycle.

See the full potential of your entire talent pool. | **eightfold.ai**

eightfold.ai

Defense Innovation Unit Builds Eightfold-Powered GigEagle Platform to Promote Talent Agility Across the U.S. Department of Defense

GigEagle is a joint talent marketplace that uses AI technologies in the Eightfold Talent Intelligent Platform to match military talent with mission needs. In its first iteration, the platform helps DoD organizations address short-term operational challenges by harnessing the specialized civilian skills of the nation's Reserve and National Guard members. By facilitating an agile talent ecosystem, GigEagle can unlock the full potential of the joint force, providing a decisive military advantage in an evolving era of Great Power competition.

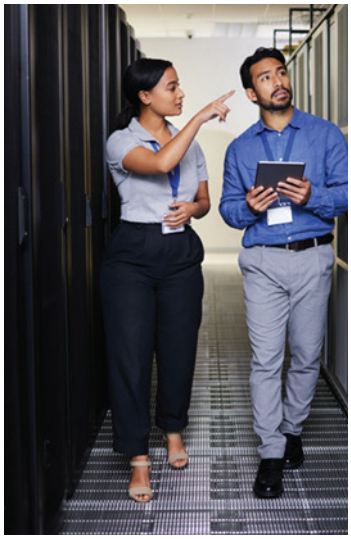

The Challenge:

The rapid evolution of modern warfare puts new pressures on the talent management practices of the United States Department of Defense (DoD). In today's national security environment, meeting fastchanging mission needs requires a highly agile and networked military—one that can leverage the full talent potential of the joint total force.

The reality is that outdated personnel management systems limit the department's ability to quickly identify and engage the best skilled talent to solve operational problems. The obstacles to talent agility include:

- The skills documented in military personnel systems don't capture the full capabilities and potential of service members.
- Tapping personal networks and soliciting resumés is highly inefficient.
- Siloed systems restrict talent searches to internal resources and hinder cross-service collaboration.

The Solution:

To optimize the use of military personnel, the Defense Innovation Unit (DIU) built GigEagle, the first-ever joint talent marketplace for the DoD. The application is powered by the Eightfold Talent Intelligence Platform to enable an agile talent ecosystem across the entire defense department.

"The goal is to make finding and deploying suitable resources to meet specific needs as frictionless as possible," Air Force Brig. Gen. Mike McGinley, director of GigEagle said.

The first iteration of GigEagle allows DoD hiring managers to leverage the specialized civilian skills of Reservists and National Guard members in temporary, short-term projects, or gigs. The expertise within this community of 750,000 service members runs deep but has been completely undiscoverable and out of reach.

In Eightfold, DIU saw an opportunity to get a running start with a commercial technology that already powers talent marketplaces for large private enterprises. The project team wanted to use Eightfold's skills-matching AI to connect service members with mission needs at unprecedented speed and scale. After a review of available technologies, "Eightfold rose to the top amongst a very large group of evaluators," Lt. Col. Chuck Kubik, GigEagle product lead said.

To prove GigEagle's viability, DIU partnered with the Air Force's AFWERX, the Space Force's Space Systems Command, the Army's 75th Innovation Command, and the Marine Innovation Unit. The mobile-enabled application is now in production following a 9-month early adoption phase and six months of prototyping.

"The goal is to make finding and deploying suitable resources to meet specific needs as frictionless as possible."

Air Force Brig. Gen. Mike McGinley, Director of GigEagle

Key Takeaways:

The DoD can better capture the skills of military personnel:

Skills that were previously undiscoverable or hard to track in military talent management systems are now captured and searchable in GigEagle.

Eightfold's AI brings to light the unique skills and capabilities of individual service members, giving the DoD an unprecedented understanding of the expertise within its own talent base.

As a result, the DoD can now harness what Kubik calls the "dual fluency" of Reservists and National Guardsmen—like the Navy Reserve sailor who is a Python coder or the Airman who runs his own business.

GigEagle has already identified a wide range of industry and technical skills among participating service members, including software engineering, cybersecurity, data science, AI/ML, program management, healthcare, drone piloting, and more.

The DoD can quickly find the right talent at the right time:

GigEagle makes the hunt for talent with specialized, in-demand skills much more precise and efficient.

DoD hiring managers can now search for talent based on the skills required to deliver a particular outcome, such as an analytics dashboard or a translated document. When a gig is posted, Eightfold's AI instantly surfaces the service members whose capabilities best fit the project's needs.

What's more, rather than listing a role on multiple service-specific job boards, hiring managers can access a joint audience with a single posting on GigEagle. Having a centralized DoD resource saves time and extends talent reach.

"It's such a transformative way for the Department of Defense to think about, How can I get the outcome I specifically need, especially when I need it on a short notice?" Col. Christoper "CJ" Johnson, U.S. Air Force Reserve – U.S. Space Force Project Lead said.

"It's such a transformative way for the Department of Defense to think about, How can I get the outcome I specifically need, especially when I need it on a short notice?"

Col. Christoper "CJ" Johnson, U.S. Air Force Reserve

The DoD can promote in-service and cross-service collaboration:

As a joint talent marketplace, GigEagle breaks down the talent silos and inefficiencies created by standalone personnel systems.

With GigEagle, DoD hiring managers can tap talent from any location or service to meet short-term operational needs. For example, when an urgent generative AI task cropped up at the Air Force Research Laboratory, GigEagle revealed two Army Reservists with the necessary data analysis and AI skills.

Kubik has demoed GigEagle for multiple senior leaders across every service. "The resounding response is, 'Holy cow, in the palm of my hand, I can look up and find the experts that live not only in my service but across all of the services,'" he said.

The DIU team is also excited about promoting "jointness" in military operations. The DoD can use GigEagle to build teams of specialists from across the total force— identifying and engaging them in a matter of hours if needed. "We talk about joint operations all the time," McGinley said. "GigEagle is putting that into effect."

The future:

The next phase in GigEagle's maturation is to scale talent agility across more DoD personnel. Pilot projects with the Air Force and Marine Corps will extend the agile talent ecosystem to include active-duty service members. In the future, McGinley envisions enabling whole-of-nation collaboration by opening up GigEagle to people in industry, academia, and other federal agencies.

U.S. military leaders have been quick to see the potential for GigEagle to address a variety of talent issues. For example, a presentation at the DoD Talent Innovation Challenge sparked the idea for GigEagle SOAR, a talent marketplace to connect military spouses with job opportunities in the public and private sectors.

DIU believes the continued development of GigEagle will benefit greatly from the strategic alignment of DIU's vision for an agile talent marketplace and the Eightfold mission to match every person with the right career. "We've grown a great partnership with these combined visions," Kubik said.

opentext™

Regional Transportation Agency

Regional Transportation Agency expands video analytics and realizes its vision of safer roads and more efficient public transportation using Knowledge Discovery, an AI platform from OpenText. Over 500 video analytics are currently running on over 200 cameras in real time with an additional 200 plus having been used in the past for completed projects. All these analytics results enable staff to respond to issues that make 1.4 million citizens safer on the road. The agency extracts about 1TB of data monthly from train CCTVs and at least 8PBs of data weekly from street cameras mounted at intersections. In addition, the Microsoft SQL Server-based warehouse holds around 3TBs of data.

The Challenge:

Making roads safe necessitates pinpointing hot spots and trends, mitigating and reacting swiftly to issues, and monitoring the performance of the entire transportation network. In addition, multiple stakeholders and partners – from police and emergency responders to third-party application developers – need actionable insight on travel activities.

A regional transportation agency was formed when the transport functions of eight former local authorities and a Regional Transport Authority were combined to oversee roads, traffic networks, and public transportation. The merger yielded five different operational centers with various technologies. A small staff monitored hundreds of older CCTV screens and tracked inputs on pedestrians, cyclists, and vehicles.

The Regional Transportation Agency needed to upgrade the city-wide surveillance camera infrastructure and implement an on-premises big data analytics platform to provide fast, real-time data to stakeholders and transform business operations. The agency faced the challenges of launching a new CCTV system, converging the units and their data.

The Solution:

Three years ago, the agency selected video recording and analytics powered by Knowledge Discovery, a data analytics solution that enables personnel to derive insights and patterns from massive amounts of real-time streaming video data. Housed in the agency's facilities, HPE ProLiant BladeSystem, HPE 3PAR StoreServ Storage, and OpenText Critical Watch support the Knowledge Discovery platform. This year, with a major new refocusing on public safety, the agency is enjoying the benefits of their investments. They have built a dedicated in-house team of video analytics specialists who work full time with OpenText Knowledge Discovery Media Server application, the highly configurable computer vision engine. This team combines analytics including Scene Analysis, Automatic Number Plate Recognition and Object Classification to build and run scores of user-story-driven analytics scenarios for surveillance and to gather evidence and accurate statistics about pedestrian safety and traffic flow all across the city.

Key Takeaways:

Exceeding Customer Expectations

The agency extracts about 1TB of data monthly from the train CCTVs and at least 8PBs of data weekly from street cameras mounted at intersections. In addition, the Microsoft SQL Server-based warehouse holds around 3TBs of data. The Knowledge Discovery data analytics platform processes CCTV video analytics, integrating that data into the incident management system and feeding the OpenText analytics database to enable data mining and higher-level alerts based on the combination of multiple video analytics events.

Over 500 video analytics are currently running on over 200 cameras in real time with an additional 200 plus having been used in the past for completed projects. All these analytics results enable staff to respond to issues that make 1.4 million citizens safer on the road. Agency personnel can detect traffic violations, congestion, and parking problems, as well as harness patterns uncovered by the Knowledge Discovery platform.

Multiple stakeholders inside and outside the agency benefit from ad-hoc investigations built with video analytics at their foundation. Whether monitoring pedestrian safety at an inner-city railway crossing, quantifying the problem of vehicles running red lights at a major intersection, detecting vehicles illegally stopping in bicycle lanes or measuring vehicle speed at key points of interest, Knowledge Discovery video analytics is enabling the agency to build a richer picture of how the regional transportation network is really used, enabling the agency to improve safety in all areas of transportation in the city.

Putting the data to work, the agency has gained an integrated ticketing system with insight on travel times, patterns, trip frequencies, and demographics.

"We can now start to tailor our messaging, especially for transport, to the stops where people are at the right time of day," says the agency's Chief Technology Officer. Instead of conducting ad-hoc surveys, city planners can use the real-time data from license plates to construct heavy transport and trucking options.

"This is very much about the planners having reliable information," says the CTO.

Enabling Strong Analytics

Current statistics and other significant volumes of data – such as the parking system – reside on the OpenText Analytics Database, which processes structured data quickly. Knowledge Discovery feeds this growing repository of relevant knowledge by producing live alerts from long-running video analytics like pedestrian counting and anonymized Automatic number Plate Recognition across the agency's CCTV network and accelerates the delivery of insights from it, linking statistics with historic corporate information and financials in the SQL warehouse.

The warehouse data integrates with the OpenText Analytics Database and API Management Solution data stores to provide integrated reports at the front end. The data analytics platform provides faster end-to-end batch file processing. Previously, batch processing took 4.5+ hours. As a result, service, bus, and ferry operators could not leverage needed information when starting their workday. Now, batch processing finishes in 2 hours.

"We have a huge performance increase," the CTO explains. "When the operators come in to work, they have yesterday's information, and they can make the right business decisions. The API Management Solution platform can manage streaming of high data volumes at high velocities."

Special Vehicle Lane Enforcement

Using Knowledge Discovery, the agency can now remotely enforce traffic rules on special vehicle lanes with video analytics. This reduces operational costs and increases compliance with bus lane regulations.

Driving Towards a Safer Future

The agency's latest success is the well-publicized deployment of a special vehicle lane use enforcement system based on analytics, which, for the first time, provides a practical and automated mechanism to identify and generate evidence to fine vehicle owners for illegally using special vehicle lanes such as bus lanes. In this way, the transportation agency will promote safer and more efficient road use for all commuters in the city.

Through its investment in big data technologies from OpenText, the regional transportation agency is building upon its "Future Cities" vision, one in which developers and officials employ all types of data, sensors, and technologies to improve products and services for their citizenry. The CTO sees the agency continuing to enhance resident experiences as more agency stakeholders and partners fully grasp the potential of big data.

"We are looking to enable them to take that data and do something with it, find patterns we don't have the expertise to find, and deliver value back to us and to everyone else. We've provided the platform and enabler. They must pick it up and run with it," he concludes.

"That has financial ROI across the medical and broader community spectrums. As a transport agency, our OpenText analytics platform is helping us exceed customer expectations and shape positive perceptions

Chief Technology Officer,
Regional Transportation agency

veritone

Escondido Police Department doubled its redaction efficiency with AI

Law enforcement agencies nationwide are facing ever-increasing audio and video footage requests. California agencies in particular are struggling to keep up with these requests, leading to a breaking point for their current redaction processes.

Such was the case with the team at Escondido Police Department, who realized their current tools couldn't help them keep up with the volume of requests without risking staff burnout and spending budget on overtime.

The Escondido Police Department had one part-time personnel working on redaction during the day shift. However, given the sheer volume of public records requests the agency was receiving, they assigned a second person, police records technician Evelyn Alva, to assist with the growing demand.

From 7:30 a.m. to 5 p.m., Alva would report to work, fulfilling all her responsibilities from her main job. Then, Alva would go home, eat dinner, nap, and return to work at 7 p.m., staying until the early hours of the morning to redact audio and video files.

The Challenge:

Billing overtime hours and risking staff burnout was a difficult yet necessary risk the command staff justified simply because the work had to get done.

Using a competitor solution that didn't have transcription capability, Alva would have to max out her desktop computer's audio volume and listen to every word spoken in the footage over and over to redact personally identifiable information (PII) and ensure no sensitive information remained in the video.

On top of this, the tool was difficult to use. Alva often had to adjust the redaction, and it wasn't intuitive if individuals left and re-entered the frame, which became even more difficult when dealing with DUI footage where minors were present. For this type of evidence, Alva had to manually review and redact every piece of audio and video footage frame-by-frame, word-by-word.

On average, an hour-long video would take five to ten hours to redact, depending on complexity. In one particularly complex instance, Alva and her coworker spent two weeks redacting a three-hour video.

It was clear that Escondido PD needed a better, more scalable solution as requests for this type of evidence continued to rise.

The Solution:

After exploring several solutions, Escondido PD adopted Veritone Redact, an intuitive, secure cloud-based solution that could remove costly manual processes, reduce human error, and improve redaction workflows.

Escondido PD was most excited about two capabilities: automatic detection and time-correlated transcriptions.

Veritone Redact automatically detects words, phrases, and objects, and allows users to define other sensitive words, imagery, and objects that need redacting. It employs head detection for redaction, which doesn't rely on visible facial features, a primary function of facial detection.

Veritone Redact also provides time-correlated transcriptions, eliminating the need to manually redact audio by listening to playback.

Key Takeaways:

- Escondido PD needed a faster and more scalable way to redact evidence as their current solution was too time-consuming and lacked key capabilities.
- The department was using valuable resources on overtime and risked burning out its small team.
- Veritone Redact has helped make the redaction of complex videos much faster and less tedious through the use of transcription-based audio redaction and automatic detection of sensitive objects.

Snorkel

Delivering Mission-Aligned AI at Scale: How Snorkel AI Helps Government Operationalize Artificial Intelligence

Many government agencies have been turning to off-the-shelf AI models, including third-party large language models (LLMs), or relying on manual data labeling to meet their goals. But these generic approaches fall short in high-stakes, security-sensitive environments like defense and intelligence, where data is siloed, fragmented, and often classified.

Across national security and civilian agencies, the mission to operationalize artificial intelligence is accelerating—but results remain uneven. While AI reshapes commercial industries, its path through the public sector is more complex. The challenge's core is not a lack of interest or ambition—it's a fundamental mismatch between generic AI solutions and the mission-specific complexity of government data, workflows, and governance requirements.

The Government AI Challenge: Off-the-Shelf Models Just Won't Cut It

Public sector agencies aren't short on AI ambition—but they face real-world constraints that commercial AI tools were never designed for:

- Highly sensitive data that can't leave the network
- Lack of labeled training data tailored to mission use cases
- Too few data scientists and overburdened SMEs
- One-size-fits-all models that don't understand policy, regulation, or real-world risk
- Requirements for full transparency, auditability, and bias evaluation

From FOIA triage at DOJ to cybersecurity classification at DoD to policy compliance at VA—government teams need more than a model. They need a repeatable, governable process to develop AI from their own unique data and SME expertise.

Snorkel Flow: AI Development Built Around Government Data

Snorkel AI's platform, Snorkel Flow, replaces brittle, manual workflows with a data-centric, programmatic approach to building and evaluating AI systems. Instead of labeling data by hand, Snorkel Flow lets users encode policy and subject matter expertise into labeling functions—logic-based rules that automatically and transparently label training data.

With Snorkel Flow, government agencies can:

- Label 100x more data, 10–100x faster
- Incorporate internal policies and SOPs directly into model development
- Achieve high accuracy without sending sensitive data offsite
- Trace model decisions back to their source for audit and oversight

The Challenge:

Case Study: A USG Cyber Organization Builds Mission-Critical Models Fast

A leading USG cyber organization deployed Snorkel Flow to build ML models for application-type classification and network attack detection. These models had to handle over 2.7 million records from internal network traffic while conforming to stringent access and governance controls.

Using Snorkel Flow, the agency:

- Programmatically labeled 280,000 records using just six labeling functions
- Built models that beat hand-labeled baselines by 7.4 F1 points
- Detected and corrected systemic labeling errors—relabelling 41% of their data with targeted analysis tools
- Delivered accurate models with zero dependency on external labeling vendors or data sharing

In a second task, the team used a pre-existing attack ontology to label 100% of a 30,000-record dataset with just two labeling functions, achieving 88.1% model accuracy without SME fatigue.

Snorkel

The Solution:

Beyond Cyber: Snorkel Across the Government

Snorkel Flow is already delivering impact across diverse federal mission areas:

- A federal audit readiness team deployed Snorkel Flow to automate contract classification and achieve 99% accuracy in under 24 hours, streamlining compliance reviews and reducing legal risk.
- A DoD lab used Snorkel Flow to triage ISR sensor data, enabling rapid deployment of mission-ready models in days instead of months while maintaining explainability and security requirements.
- In the commercial sector, a Fortune 50 financial institution used Snorkel Flow to develop a Generative AI application for news analytics, accelerating deployment by 45x and boosting model accuracy by 25 points over legacy systems. This effort highlights how custom-tuned, organization-specific GenAI applications—built on curated internal data—can outperform off-the-shelf LLMs while meeting governance and performance standards.

Key Takeaways:

Aligned with Federal AI Policy and Oversight Requirements

Snorkel Flow is uniquely suited for public sector AI governance, built from the ground up to align with the full spectrum of U.S. federal AI policy and oversight mandates. Whether your agency is operating under civilian oversight or national security protocols, Snorkel Flow enables responsible, policy-aligned AI at scale.

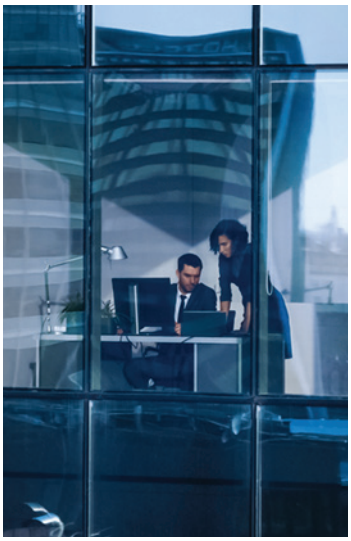

It supports and integrates with frameworks, including:

- **DoD CDAO Test & Evaluation Framework** – Snorkel Flow embeds model evaluation throughout the lifecycle, enabling testing for correctness, bias, and robustness in operational contexts.
- **NIST AI Risk Management Framework (RMF)** – Snorkel Flow enables traceability, fairness evaluation, and governance tooling essential to AI risk mitigation.
- **OMB M-21-06 & EO 13960** – Snorkel supports compliance with transparency, explainability, and performance mandates for federal AI deployment and inventory reporting.
- **Blueprint for an AI Bill of Rights** – Snorkel Flow helps agencies operationalize principles like explainability, protection from algorithmic discrimination, and data privacy via full workflow auditability.

Critically, Snorkel Flow supports the most recent federal directives:

- **National Security Memorandum on AI (October 2024)** – Snorkel Flow supports this by enabling in-network AI development on classified data with SME-in-the-loop workflows, eliminating the need to expose sensitive sources or use external labeling services.
- **Executive Order 14179:** Snorkel Flow empowers agencies to build bespoke AI aligned to their mission and data, avoiding reliance on black-box, one-size-fits-all solutions. The platform's ability to rapidly develop, test, and adapt AI in-house removes friction, accelerates deployment, and supports the broader national imperative to retain U.S. leadership in trustworthy and cutting-edge AI.

Scaling AI from Pilot to Operational Advantage

Where most government AI projects stall after proof of concept, Snorkel Flow enables sustainable, scalable AI development. Agencies can reuse labeling logic across use cases, adapt to policy shifts in real-time, and empower SMEs to build models without relying on hard-to-hire ML talent.

Snorkel AI: Trusted AI, Built on Government Data

In today's federal AI ecosystem, success isn't just about speed or sophistication—it's about trust. Trust that models are grounded in policy. Trust that data remains secure. Trust that decisions can be explained and governed. That's precisely what Snorkel Flow delivers.

If your agency is ready to turn pilot projects into production-grade AI, Snorkel Flow is the platform to make it real—securely, responsibly, and at mission speed.

CORAS

CORAS Decision Management integrated into one package. Proprietary AI and NLP Identifies Interconnections across Data Sets for Actionable Decisions, What-if Scenarios, and Trade-offs

CORAS Enterprise Decision Management SaaS is trusted to operate in the DoD's most secure environments, including IL 5, NIPR, SIPR, etc., and others as well as the FedRAMP High marketplace. Its predictive and prescriptive-AI and NLP capabilities support human leaders to take action, and dynamically drive real-time organizational decision making. Up and running in 60 days or less without disrupting current systems, it works across systems to create a single source of live truth, eliminate data calls, reporting lags, and static reporting; frequent roadblocks to mission success.

The Challenge:

The lack of integration across key systems—schedule management, financial management, portfolio management, project management, and workforce management—creates silos that waste time, energy, and resources. Leaders grow increasingly frustrated as "data in motion" consumes the majority of their workforce's time. Imagine your team spends 80% of their effort shuffling data and only 20% analyzing it. What if you could flip that ratio, empowering your workforce to focus on making data-driven decisions instead of just moving data around?

The Solution:

CORAS supported a DoD customer in leveraging Portfolio and Capability Mapping streamlining their project and portfolio management process for a $901.3M budget. CORAS linked data from many systems including contracts, financials, and systems to optimize the office's assets into a single source of truth:

- CORAS' SmartLanes feature facilitated the identification and easily viewing of relationships between contracts and projects across the portfolio to show key dependencies
- NLP similarity analysis identified duplication of effort and opportunities to better optimize spend in support of the mission
- Real-time data also allows for real-time reporting
- Programs learn over time through AI suggestion and user interaction
- Raw metadata loads and edits continuously.
- Unseen relationships between programs were revealed via similarity analysis across thousands of programs to drive complete assessment and complete insights.

Key Takeaways:

With CORAS, PMs and leaders can:

- Unify disparate systems and data into a real-time, single source of truth
- Project various scenarios and potential outcomes or historical back data in moments
- Understand the full scope of assets, weapons, readiness, funding, and resources.
- Enable dynamic data manipulation to enhance upstream or downstream communications and operational effectiveness.
- Quickly evaluate large data sets and uncover correlations, dependencies, and duplications across projects, budgets, etc.
- Anticipate problems before they arise or create automations that support workflows for efficiency and transparency.
- Get real-time data, automations, and predictions with multi-dimensional, interactive views across any system (MS, Advana, etc.)
- Configurable to adapt to unique organizational needs
- Measure performance metrics such as resources, cost, schedule, and risk with precision
- Seamless integration without disrupting current workflows or systems

CORAS is the only software that turns insights into real-time action - and anticipates and mitigates future events. CORAS doesn't just modernize your approach—it revolutionizes it. Equip your workforce with the tools they need to spend less time managing data and more time driving mission success.

salesforce

Agentforce

Build public trust and connection with Agentforce for Public Sector.

Learn how humans with AI agents can drive customer success together across public sector, higher education, and nonprofit industries.

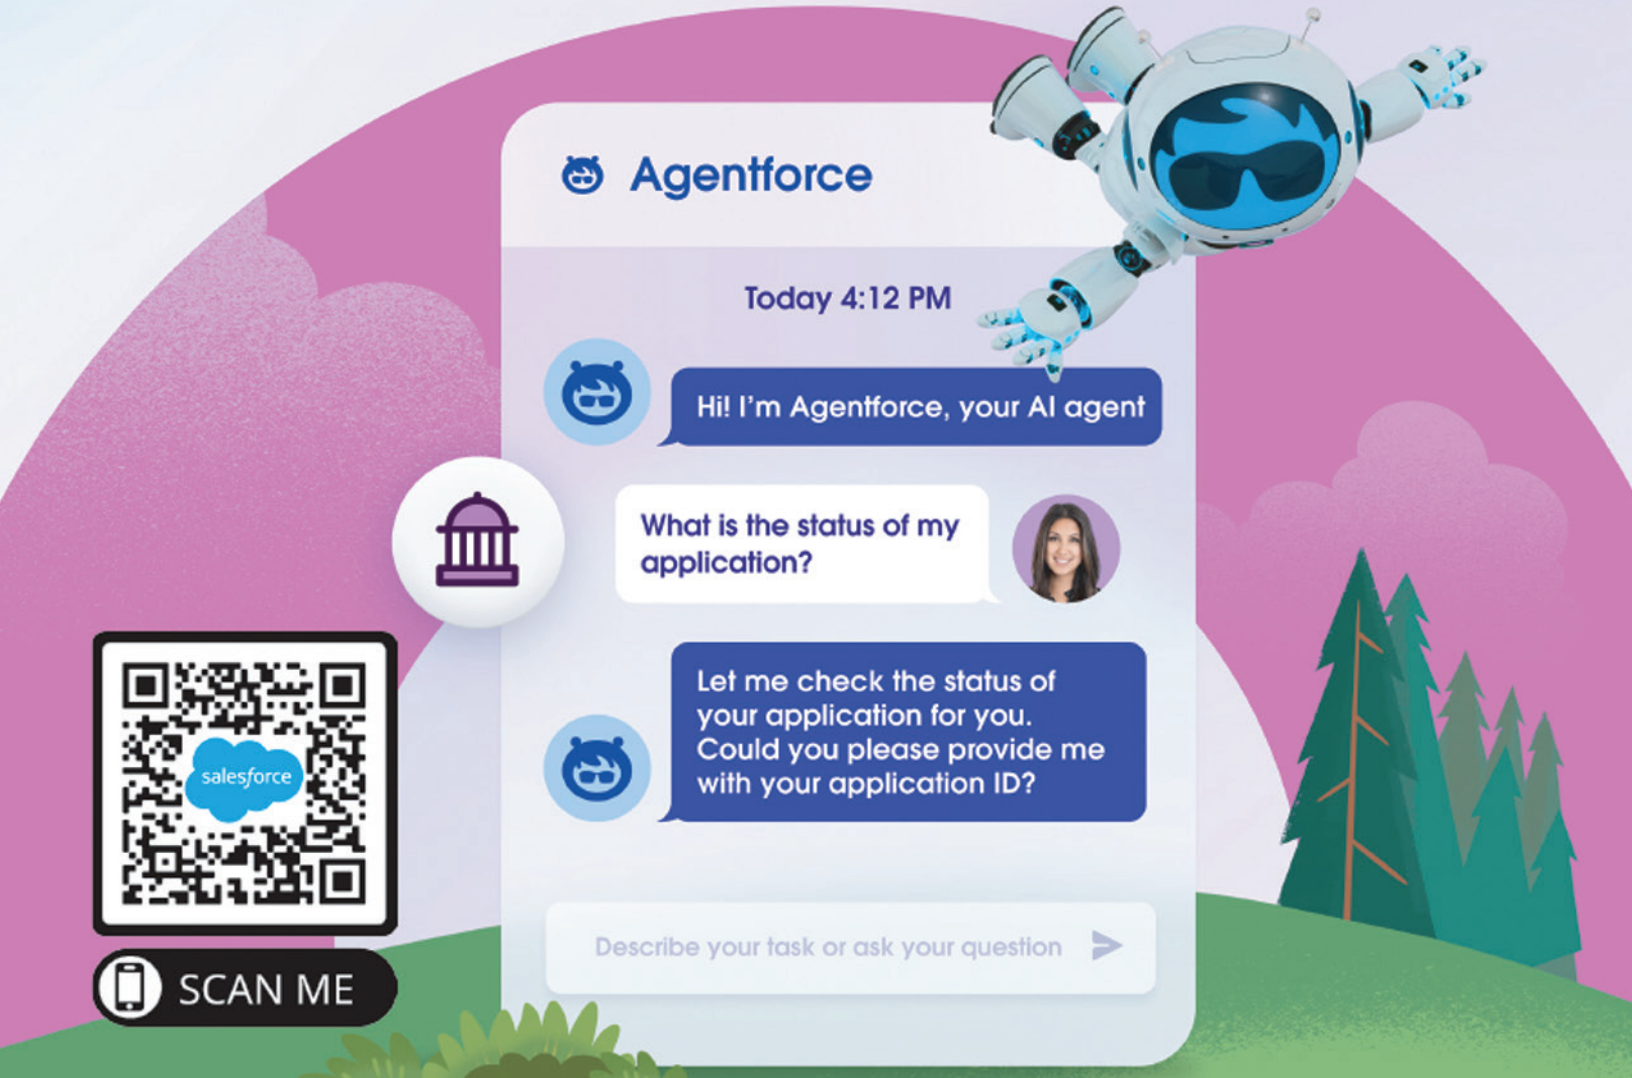

PRYON

Secure & Trusted GenAI for Government

Deploying IL2-IL7 | on-prem to cloud

- Create custom GenAI workflows
- Enable rapid decision making
- Unlock organizational data

Learn more at pryon.com

Policies You Should Know

AI policy has changed a lot over the past year. Below are some of the key policies you should be aware of to align your agency's AI to mission.

Executive Order 14179:
Removing Barriers to American Leadership in Artificial Intelligence

President Trump issued Executive Order 14179 Removing Barriers to American Leadership in Artificial Intelligence on January 23, 2025. The order prepares the foundation of the Trump administration's AI policy following the repeal of President Biden's Executive Order 14110.

Major Provisions:
Sets the goal of federal AI policy "to sustain and enhance America's global AI dominance in order to promote human flourishing, economic competitiveness, and national security."

By July 22, Assistant to the President for Science and Technology (APST), the Special Advisor for AI and Crypto, and the Assistant to the President for National Security Affairs (APNSA) must put forth a report to achieve this goal.

By March 24, OMB will revise M-24-10 and M-24-18 to come in line with this policy.

Impact on Government:
With requirements from Executive Order 14110 nullified, some government officials have looked to the executive orders from President Trump's first administration: Executive Order 13859 and Executive Order 13960. Others have also turned to the NIST AI Risk Management Framework for guidance.

With OMB M-24-18 under review, there will likely be new requirements around the procurement of AI systems that will replace those issued in the memorandum by the Biden Administration.

The White House will likely direct federal agencies to prioritize AI use cases that clearly enhance the United States' "human flourishing, economic competitiveness, and national security."

More information around the White House's AI policy will be available upon release of the Artificial Intelligence Action Plan in July.

softserve

Revolutionizing Education with AI-Powered Innovation

SoftServe is transforming the modern classroom with cutting-edge AI and Generative AI solutions, empowering educators, students, and parents like never before.
Partnering with an Independent School District in Texas, USA, we are redefining how education leverages technology to foster personalized learning, boost student engagement, and enhance collaboration between teachers and families.

Unlocking Student Potential with Gen AI

Our AI-driven methodology seamlessly integrates dispersed student data, creating a **holistic student profile** that:

- Personalizes learning experiences for every student
- Enables teachers to deliver tailored instruction with precision
- Engages parents in real-time learning progress
- Strengthens communities of growth and exploration

By combining **advanced Gen AI solutions** with robust enterprise architecture, SoftServe is driving digital transformation across the education sector - setting a new benchmark for innovation.

Experience the Future of Learning

📖 **Read our case study**

to see how Mesquite ISD is revolutionizing parent engagement with AI.

💡 **Schedule a demo with the SoftServe team**

and explore how Gen AI can reshape your education strategy.

Let's unlock the full potential of AI in education - together.

softserve

Aretec

THAT TRIUMPHANT MOMENT

—powered by Google

Feel that surge of satisfaction as every detail seamlessly clicks into place—that dopamine hit of finding exactly what you need, instantly.

Harness the power of **Google Cloud's VertexAI & DocAI** to transform your data processing into a competitive advantage, unlocking continuous moments of triumph.

No code, No chaos, Just results.

TRANSFORM YOUR WORKFLOW & CELEBRATE EVERY WIN

Aretec's Google-powered services and solutions are transforming how organizations interact with their data—turning processing bottlenecks into business breakthroughs.

By eliminating technical barriers and automating workflows, we empower teams to focus on what truly matters.

- End-to-End Governance
- Vision AI
- Intelligent Data Extraction
- Speech-to-Text
- Document Summarization
- Knowledge Management

Contact us today @ info@aretecinc.com | www.aretec.ai

Google Cloud Partner

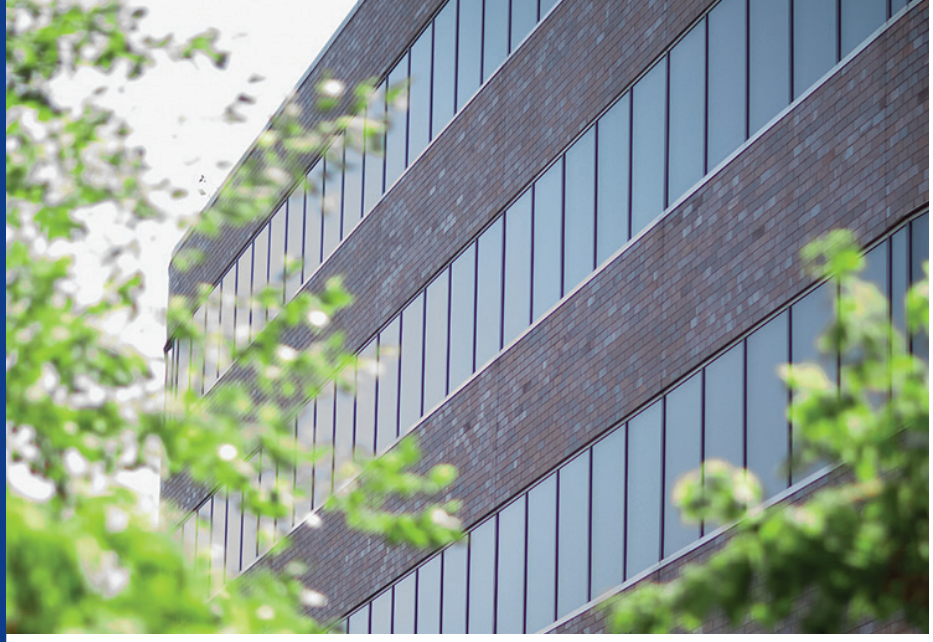

Contract Vehicles

Carahsoft offers a number of contract options for purchasing AI solutions. Our contracts offer purchasing options for civilian, defense, state, and local government customers. Customers can purchase solutions off of four major contract vehicles:

GSA Multiple Award Schedule (MAS)

Carahsoft holds a GSA Multiple Award Schedule (MAS) that allows customers to procure a wide variety of AI solutions. Carahsoft holds Contract #47QSWA18D008F and allows customers to purchase everything from AI infrastructure to advanced analytics solutions.

ITES-SW2

The purpose of the ITES-SW 2 acquisition is to support Army, Department of Defense (DoD) and all Federal Agency enterprise Information Technology (IT) infrastructure and info-structure goals by leveraging Commercially available-Off-The-Shelf (COTS) software products and maintenance in 14 product categories in addition to related incidental services and hardware.

NASA SEWP V

The NASA SEWP V GWAC (Government-Wide Acquisition Contract) provides the latest in Information Technology (IT) products and product-based services for all Federal Agencies. SEWP provides the best value and cost savings through innovative procurement tools and processes; premier customer service and outreach; and advocation of competition and cooperation within the industry.

NASPO ValuePoint Cooperative Purchasing Organization

The NASPO ValuePoint Cooperative Purchasing Organization (formerly WSCA-NASPO) provides the highest standard of excellence in public cooperative contracting. By leveraging the leadership and expertise of all states with the purchasing power of their public entities, NASPO ValuePoint delivers best value, reliable, competitively sourced contracts.

Since 1993 NASPO ValuePoint has been the cooperative purchasing arm of NASPO (the National Association of State Procurement Officials) encouraging, fostering and guiding the nation's most significant public contract cooperative. NASPO ValuePoint is a unified, nationally focused cooperative aggregating the demand of all 50 states, the District of Columbia and the organized US territories, their political subdivisions and other eligible entities spurring best value, innovation and competition in the marketplace.

Explore the benefits of how you can count on Carahsoft and our Reseller Partners

- 24x7 availability call us at 888-662-2724
- Dedicated support specializing in serving enterprise ready solutions
- Ecosystem of value-added reseller partners
- Contract Expertise: We understand your procurement needs and the outcomes you're seeking
- Quick turnaround quote: Get the IT solutions you need with the fast, accurate service you deserve
- Substantial cost savings on AI products and service portfolio from certified technology brand partners
- Advanced technology solutions including data & analytics, generative AI, AI infrastructure, and cybersecurity

Upcoming AI Events

Artificial intelligence (AI) has grown from simple automation and robotics in the mid-20th century to modern-day advanced technology, experiencing major evolution within the last several years. The power of AI and increased usage across many Public Sector markets has sparked a new wave of innovation surrounding development of tools, solutions, policy, ethics and more. Whether showcasing new technologies or networking with AI tech industry experts, dive into Carahsoft's recommendations below for the top upcoming AI events to attend this year as we explore all these topics and what is to come for AI.

AI Week Presented by AI Scoop

April 21-25 | Washington D.C.
Hybrid Event

AI has the potential to transform our world, and at AI Week 2025, attendees will get to experience this shift firsthand. Join thousands of C-suite leaders from the Government, tech industry and education communities across the country online and in-person for a week-long festival celebrating present modernization and future innovations of AI. Participate in lightning talks, networking opportunities, interactive sessions and community events to gain valuable insights on the latest in AI.

Sessions to Look Out for:

AI TALKS: Explore the Transformative Potential of Artificial Intelligence in the Public Sector Space.

Carahsoft partner, Salesforce, is the Underwriting Sponsor for AI Week. Join them along with our other partners Broadcom, Microsoft and Red Hat as sponsors of AI Talks for discussions around sustaining and enhancing the nation's leadership in AI "to promote human flourishing, economic competitiveness, and national security.

Microsoft AI Tour - Washington, D.C.

May 6 | Washington D.C.
In-Person Event

Microsoft AI Tour is a free one-day, in-person event happening around the world. The Washington DC tour stop is designed for Federal government business leaders, partners, practitioners, and the developer community to come together to learn about the latest technology and innovations from Microsoft to turn your AI vision into reality.

Carahsoft is excited to be an exhibitor sponsor of the event. Stop by our expo to learn how Carahsoft can support you and your agency with your Microsoft needs.

GovCIO AI FedLab

May 13 | Reston, VA
In-Person Event

The GovCIO AI FedLab is a unique event, bringing together Government and industry in a workshop environment to learn about AI challenges that span across all agencies. Attendees of over 50 senior-level Federal IT executives will participate in working sessions on problem-solving, brainstorming and creating meaningful engagement for participants and speakers.

Carahsoft is the leading sponsor and host of GovCIO's AI FedLab. Join us at the Carahsoft Conference & Collaboration Center for the second year running of this cutting-edge collaborative workshop.

Carahsoft AI for Government Summit

May 15 | Reston, VA
In-Person Event

AI has the potential to transform our world, and at AI Week 2025, attendees will get to experience this shift firsthand. Join thousands of C-suite leaders from the Government, tech industry and education communities across the country online and in-person for a week-long festival celebrating present modernization and future innovations of AI. Participate in lightning talks, networking opportunities, interactive sessions and community events to gain valuable insights on the latest in AI.

Sessions to Look Out for:

- Government Progress on White House Executive Order 14110 and the NSM on AI
- How the Trump Administration is Handling Prior White House EO's on AI
- Update on Government Modernization through the Data Center and AI Infrastructure
- What Efficiencies Government Has Gained Using Generative AI and Where it is Headed
- Enhancing Customer Experience (CX) in the Public Sector with AI
- What Use-Cases are in Production Today

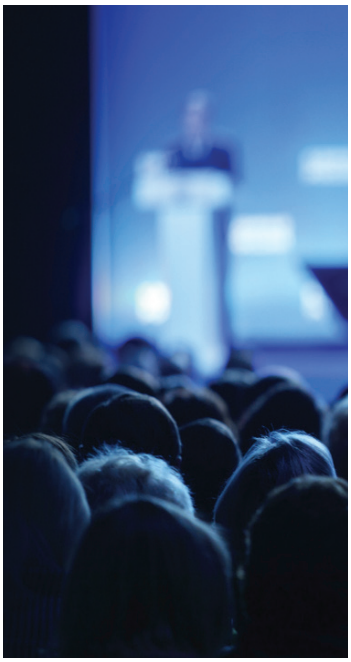

Carahsoft is the proud host of the AI for Government Summit. This was a milestone event in 2024 with speakers from NVIDIA, OpenAI, Microsoft, Google, AWS, Dell, HPE and others. Join us at the Carahsoft Conference & Collaboration Center to learn about all things AI within the Public Sector.

GEOINT 2025

May 18-21 | St. Louis, MO
In-Person Event

The GEOINT Symposium explores the critical role of geospatial intelligence in building a secure future. Engage with Government leaders, industry pioneers and thought leaders to uncover how GEOINT is addressing pressing challenges and opportunities in today's complex global environment. From land and sea to cyberspace and outer space, geospatial intelligence is revolutionizing every domain, empowering decision-makers to respond effectively to critical needs.

As an annual sponsor of the GEOINT Symposium, Carahsoft will be joined by more than 70 of our technology providers on the tradeshow floor in our pavilion space. Additionally, Carahsoft will be hosting a partner-sponsored networking reception at the Post Building. Carahsoft is also the exclusive lanyard sponsor of GEOINT 2025 and attendees will have a front seat to more than 200 speakers across a variety of educational sessions, meaningful networking opportunities and technology demonstrations showcasing modern solutions to present day challenges.

Billington Cybersecurity Summit

September 9-12
Washington, D.C.
In-Person Event

A long standing and experienced event, the Billington Cybersecurity Summit features an extensive array of cyber topics, speakers, sessions and interactive breakouts for attendees to truly immerse in the world of today's emerging cybersecurity solutions and trends. In its 16th year running, this leading Government cybersecurity summit promises an exceptional line up of Government presenters, an invaluable leadership luncheon, an all-attendee networking reception and over 100 vendor booths featuring strategy development and technology demos.

For a sneak peek into what you can expect at the summit, topics covered during last year's event included:

- Zero Trust
- Ransomware
- Advancing cyber diplomacy
- Learning how to use proactive defenses
- Engineering AI into cybersecurity platforms
- Implementing an effective risk management approach
- Protecting critical infrastructure

Carahsoft is looking forward to sponsoring this year's event and will feature a booth to engage with attendees throughout the week. We will also be hosting a large partner pavilion and encourage attendees to stop by and learn more about our partners and their technology solutions! Check out the events tab on our website for more details closer to the event!

SC25: International Conference for High Performance Computing

November 16-21
St. Louis, MO
Hybrid Event

Supercomputing25 (SC25) is the premier global conference for high-performance computing (HPC), networking, storage and analysis, tailored to address the needs of Government, defense and research organizations. This year's event explores the transformative impact of HPC technologies on solving critical challenges, advancing national security and driving innovation across scientific and governmental missions.

At SC25, there will be a wide array of programming including presentations on new research, showcasing innovative work or practices and teaching and guiding the next generation of HPC students and professionals.

Carahsoft will once again host a large pavilion space and is an exhibiting sponsor of SC25 along with many of our partners at the forefront of high-powered computing, including AWS, Broadcom, Cloudian, Dell, Google, Groq, HPE, IBM, Intel, Microsoft, Microway, NVIDIA, Oracle, Red Hat and VAST Data. Building. Carahsoft is also the exclusive lanyard sponsor of GEOINT 2025 and attendees will have a front seat to more than 200 speakers across a variety of educational sessions, meaningful networking opportunities and technology demonstrations showcasing modern solutions to present day challenges.

carahsoft

Contact Us:

(571) 591-6040
AITeam@Carahsoft.com

11493 Sunset Hills Road, Suite 100
Reston, Virginia 20190

carahsoft.com/solve/ai