VELOCITY

Insights for Federal Innovators

V3. 2025

BY Booz Allen

Tech at Full Speed

The Modern Flywheel Effect Explained

NVIDIA CEO Jensen Huang: Agentic AI Is the Future

ABOUT THE COVER
Cover art was designed by Brody Rose. It uses AI-generated imagery from Adobe Firefly to represent the convergence of emergence technologies into an intelligent, self-optimizing system.

A NOTE TO READERS
The ideas and opinions contained herein are those offered by the individual authors. They are intended as considerations in the associated technical areas discussed. They do not necessarily represent the firm's views, but offer the breadth, depth, and currency of Booz Allen's technical talent and that of our partners. References to any products or companies are not meant to be an endorsement by Booz Allen or any other contributors.

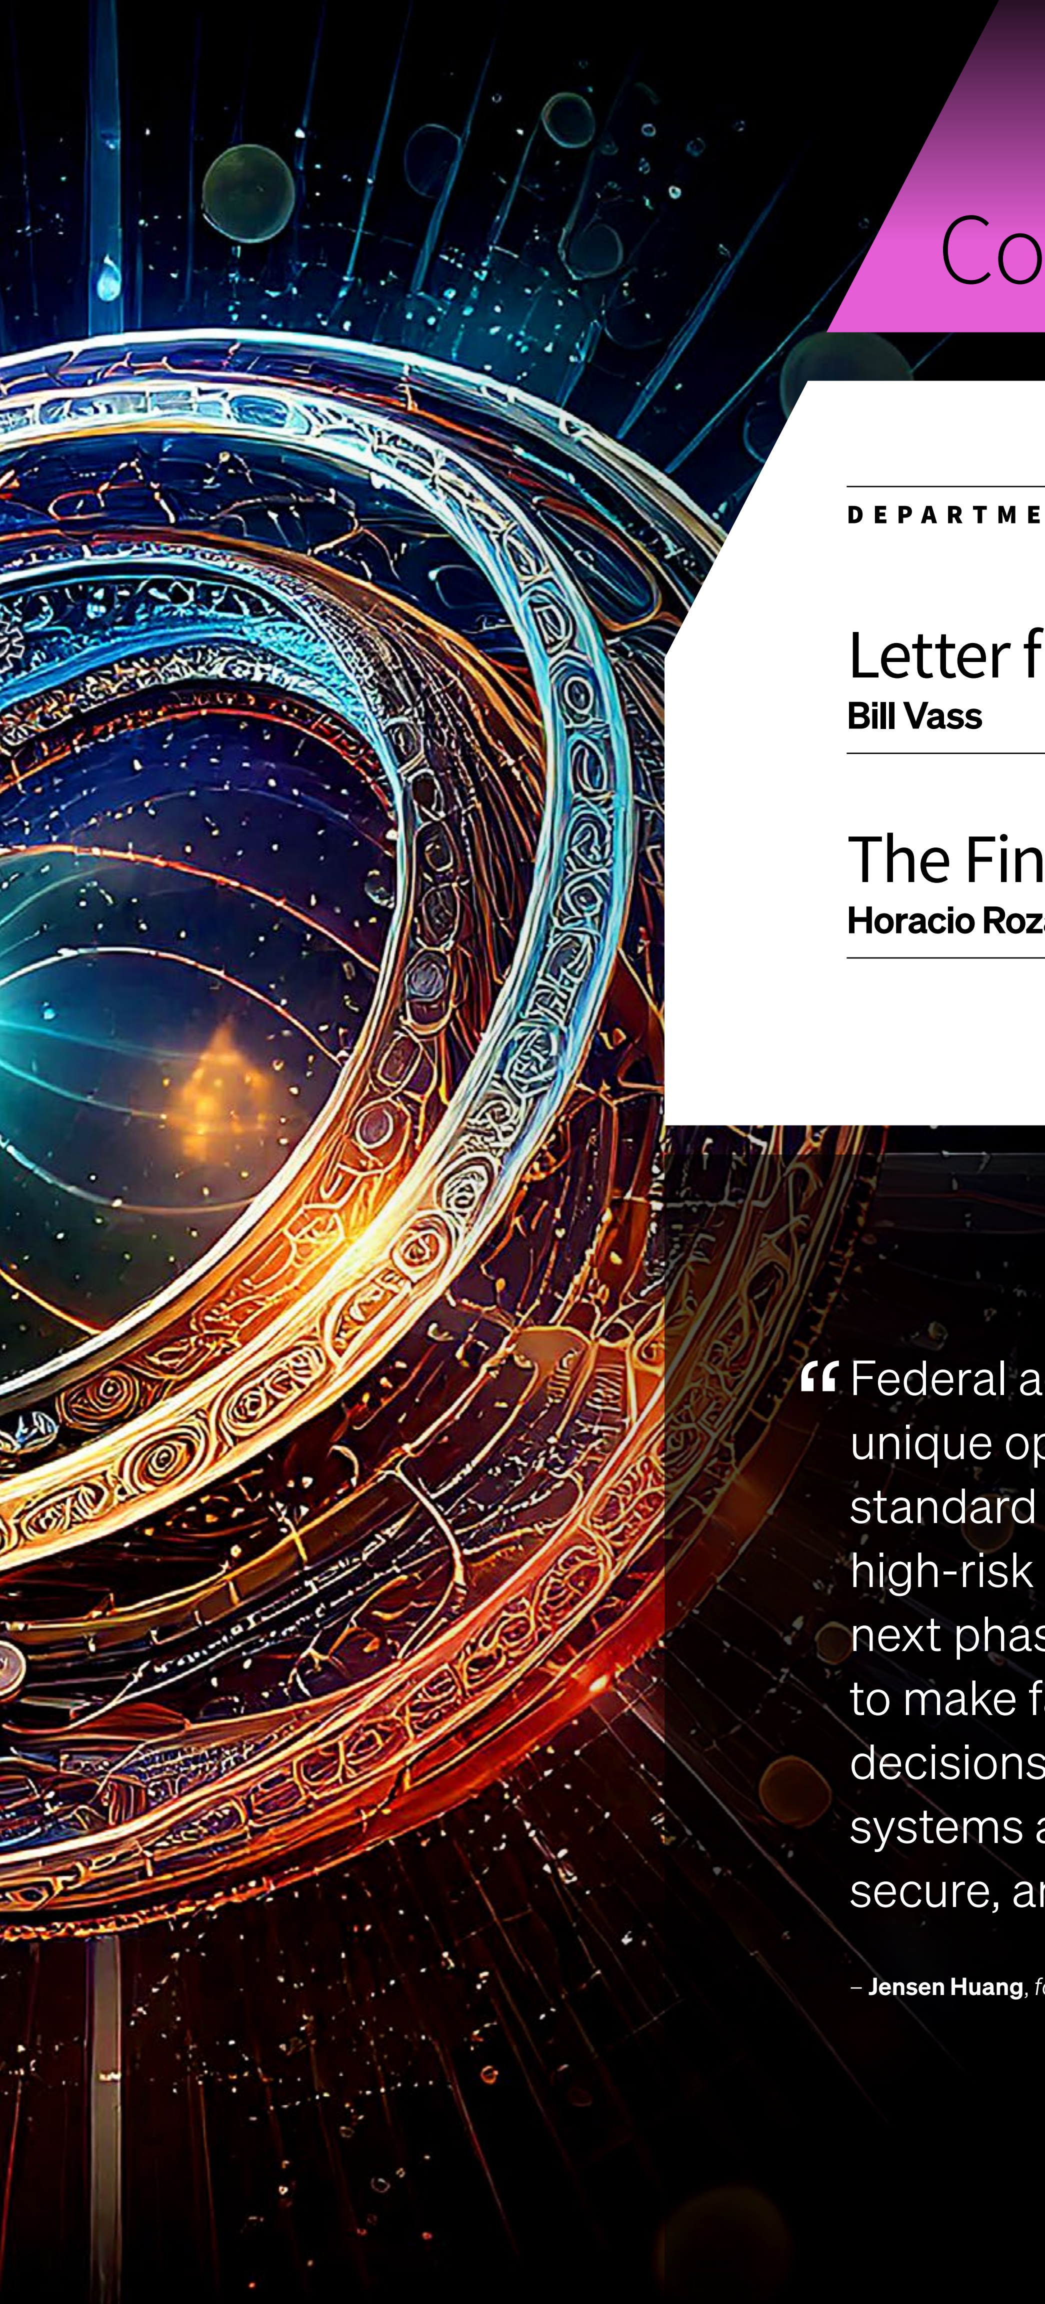

Contents

DEPARTMENTS

> “Federal agencies have the unique opportunity to set the standard for AI operating in high-risk environments. The next phase is about scaling AI to make faster, more precise decisions while ensuring these systems are transparent, secure, and accountable.”
>
> – **Jensen Huang**, *founder and CEO of NVIDIA*

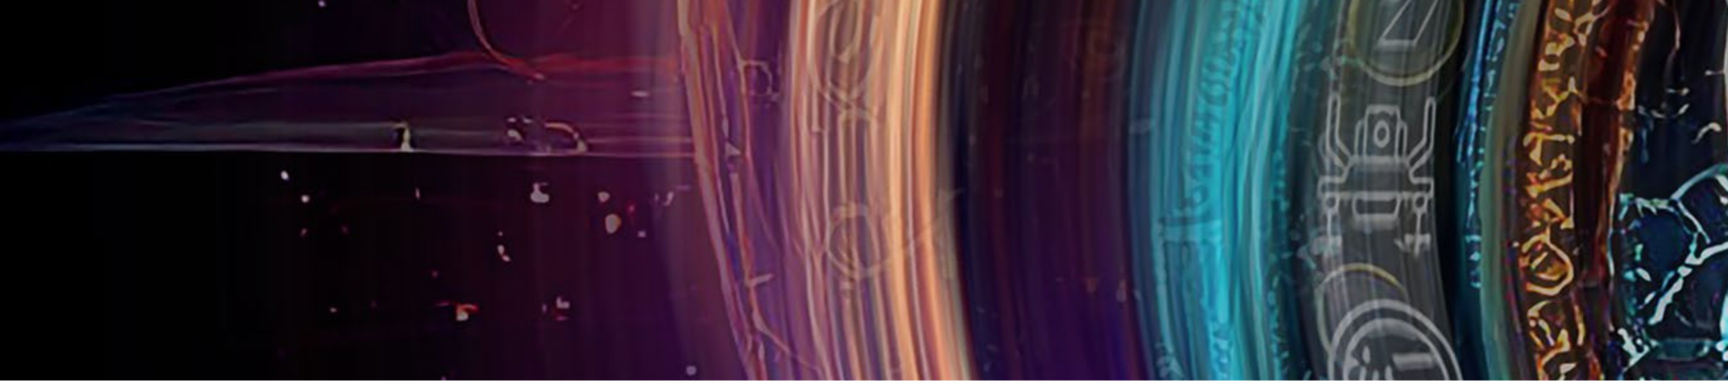

FEATURES

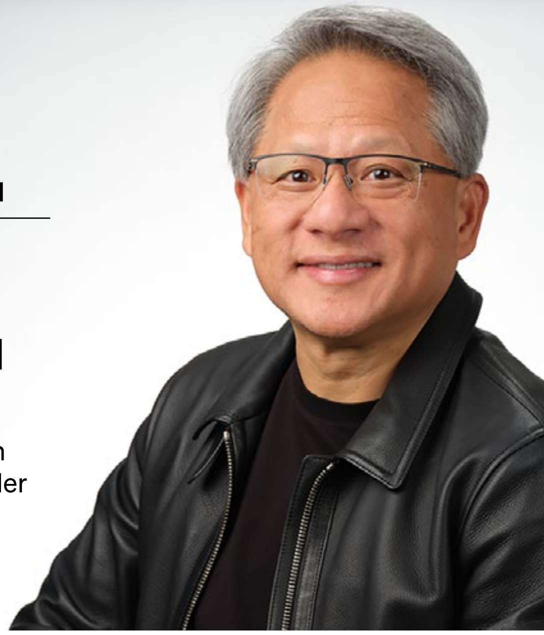

MISSION SPOTLIGHTS

EMERGING TECHNOLOGIES

Booz Allen commemorated its 110th anniversary at the New York Stock Exchange

FROM THE CHIEF TECHNOLOGY OFFICE

How can we leverage the most advanced commercial technologies to accelerate federal missions and deliver outcomes that matter? This is the essential question facing our nation as we enter a new era characterized by rapid technological advancement. Today, private sector companies are expanding the art of the possible in fields such as artificial intelligence (AI), cloud computing, digital twins, and more. At the same time, our nation faces the most complicated national security environment since the height of the Cold War while the mounting pressure of an unsustainable national debt threatens economic growth at home. Against this backdrop of complex challenges, it is clear that speed to innovation is no longer a strategic advantage—it is an imperative.

In this edition of *Velocity*, we examine how technology is evolving and how it can be applied to accelerate innovation, create efficiencies, and advance the wellbeing of the American people. Since we published our last issue about the ubiquity of AI, the idea of tech at full speed has taken on heightened significance. AI is not only embedded in the most critical missions; it is reimagining the pace and methods by which we build and deploy all technologies (**see my cover story on the modern flywheel effect on page 32**). As a result, we're standing on the precipice of one of the most consequential technology waves in history.

As the leading provider of AI solutions to the U.S. government, we have a responsibility to ensure federal agencies can use AI and other emerging technologies to modernize their current capabilities, augment human performance, and build the systems of tomorrow. Now is the time to reimagine how the public and private sectors work together (**see our discussion with NVIDIA CEO Jensen Huang on page 28**). By meeting this moment with a focus on outcomes and efficiency, we can chart a path toward a more resilient, secure, and prosperous future.

As we set out to create this future together, I invite you to explore insights and perspectives from technology experts within Booz Allen and across the industry. This year's magazine looks at:

- **How generative AI is evolving from single large-language models to multi-agent systems** (see our story on agentic AI on page 41)
- **Why virtual ground systems will play a key role in protecting U.S. space assets** (see our space mission spotlight on page 17)
- **AI's potential to reimagine how cybersecurity teams conduct threat hunting** (see our article on AI-powered cybersecurity on page 11)

For 11 decades and counting, Booz Allen has pioneered innovations for our clients. Our experience has taught us that the times of greatest change can spur the most remarkable transformations. In a world where technology defines tomorrow, let's seize this moment to disrupt for good.

Bill Vass

Chief Technology Officer
Booz Allen

Reflections on Generative AI

LESSONS LEARNED FROM TWO YEARS OF PROMPTING

Ernest Sohn and Alison Smith

Generative artificial intelligence's (GenAI's) most remarkable attribute is how rapidly it continues to advance. More than two years after the release of ChatGPT, the emergent abilities of early frontier models like GPT-3—such as computer programming skills and early signs of reasoning—are now the core competencies of newer, more advanced models. A story published on Windows Central indicated that OpenAI's o1 and o1-mini models have demonstrated the ability to pass the company's research engineer hiring interview for coding. In October 2024, Anthropic announced that it was teaching its AI assistant Claude how to use a computer like a human and tackle open-ended tasks like conducting research.

Though GenAI's technological capabilities are advancing at an accelerated rate, its impact at the enterprise level has been more nuanced. In the simplest terms, the technology has not yet consistently and competently performed complex operations without human assistance. This paradigm will likely change within the next three years as the underlying technology matures and enterprises pair large language models (LLMs) and other types of GenAI with other technologies and with one another to create intelligent systems with higher levels of autonomy.

With those advances on the near horizon, it has become easier to see GenAI clearly, both in its creative power and existing shortcomings. Every answered prompt implies a new truth about the technology's inner workings, and we can now take stock of what we've learned from more than two years of prompting. Here are key perspectives on GenAI's evolving story—from the invaluable role humans play in its deployment to the implications of its stochastic nature and the use case types it's best suited for.

The Best Use Cases for GenAI

Amid confirmation of both its promise and its sometimes-puzzling breakdowns, GenAI keeps edging closer to the heart of critical missions. In manufacturing, the technology is modernizing core processes, including incident response systems that harness interactive copilots to predict and proactively address issues with little human involvement. As part of a Defense Advanced Research Projects Agency initiative, GenAI now identifies and corrects vulnerabilities in open-source software underlying critical infrastructure.

The technology's ability to rapidly innovate gives it the potential to help agencies address issues of critical national importance. Before these sweeping transformations can occur, agencies can reap the benefits of one of its most prolific proven use cases: improving operational efficiency. The U.S. Patent and Trademark Office, for example, harnesses its robust search capabilities to help examiners find relevant documents faster when processing patent applications. The Department of Defense has developed the Acqbot writing tool to reduce the human time and effort needed to generate contracts. What ties these applications together is the focus on using GenAI to help workers complete everyday tasks more efficiently. Still, these kinds of projects represent only generic early adoption successes with limited effect. The mature use cases of the future will be far more powerful and will entail combining GenAI with other types of AI to create maximum impact.

Today, organizations can use GenAI to operationalize applications ranging from customer service chatbots and

Figure 1: Different types of AI

Traditional AI	Generative AI
Accuracy: Precision and reliability of results	Creativity: Invention of new ideas and logical conclusions
Interpretability: Access to model decisions and outputs	Generalizability: Application of learned knowledge to new, unseen data

software coding agents to tools for scientific discovery and policy adjudication. All are underpinned by GenAI's capacity to rapidly and competently perform human functions, such as information retrieval, data aggregation, summarization, interpretation, analytical processing, synthesis, predictive modeling, iterative design, and, of course, content generation.

Nonetheless, GenAI isn't yet suited for every task. Use cases necessitating more definitive, explainable, and predictable outputs, such as precision manufacturing or network security monitoring, may be more appropriately addressed by traditional machine learning algorithms. Why? Unlike traditional algorithms, which can be developed to produce the same output given the same input, GenAI models can generate different outputs due to their use of stochastic processes. Sources of variability across the model lifecycle include randomly initialized parameters within neural networks and techniques, such as Monte Carlo sampling that introduces randomness into the inference process, among many others. Further, traditional machine learning algorithms are highly specialized in terms of their data and domain, whereas GenAI models tend to excel in their generalizability.

The stochastic nature of GenAI creates a need for meticulous use-case analysis to accurately calibrate operational and mission risks. For well-defined problems with specific rules and constraints, traditional AI often presents a more reliable option, particularly when the user doesn't need the model to explain its decisions. Evaluating the trade-offs between different AI paradigms enables agencies to identify the right approach. By embracing the probabilistic and creative spirit of GenAI applications, agencies can better harness their potential to drive transformation while supporting safe deployment.

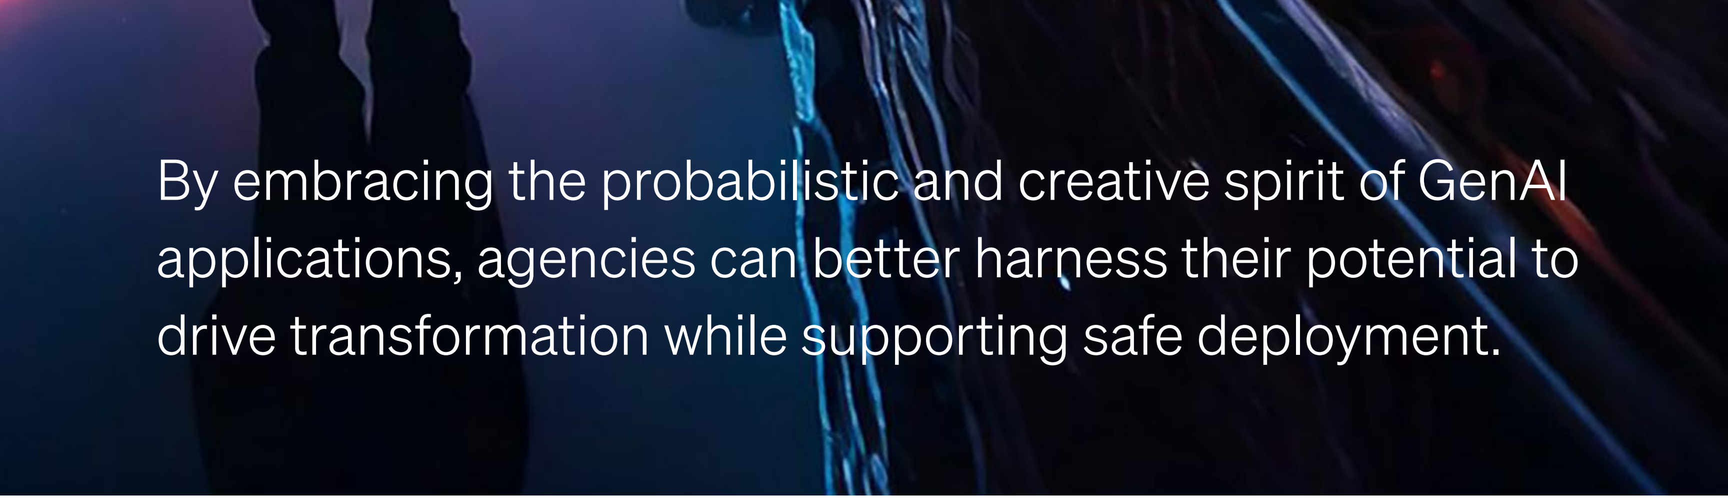

By embracing the probabilistic and creative spirit of GenAI applications, agencies can better harness their potential to drive transformation while supporting safe deployment.

How to Support Responsible AI

Along with the valuable assistance to core processes that they provide, LLMs' glitches and miscalculations cannot be ignored. According to a *Tech Xplore* story published in December 2023, researchers at Google tricked ChatGPT into leaking sensitive training data when asked to repeat the word "poem." A study published in December 2023 in the *Stanford Internet Observatory* indicated that the LAION-5B dataset used to train Stable Diffusion was found to contain hundreds of illegal images.

Indeed, GenAI's blistering growth amplifies concerns about ethics, security, and responsible AI development. As a result, many countries have voiced concerns over the potential risks associated with AI advancements and advocated for international cooperation to develop a shared regulatory framework. These concerns stem from the rapid pace of GenAI development and its profound impact as a general-purpose technology across various sectors. In response, nations have begun to form global alliances and issue declarations aimed at fostering collaboration and establishing guidelines for AI governance. Examples include the G7 Hiroshima Process on Generative AI; the Bletchley Declaration affirming commitment to safe, human-centric AI; the United Nations Global Digital Compact to create a common AI ethics framework; the industry-driven AI Alliance; and the Frontier Model Forum to promote AI safety research.

While these initiatives stress the importance of collaboration, much of the language within these declarations remains vague and lacks concrete followup actions, raising questions about the genuine commitment to international governance. It has also become evident that countries, particularly those with disproportionate access to resources and advanced capabilities in chip production and foundation models, may prioritize their national interests. This pattern suggests that, despite the appearance of unity and collaboration, many countries continue to pursue individual strategies.

Though efforts to codify approaches to responsible AI at the national and global levels are driving needed awareness of the matter, they should not entice leaders to pause ongoing GenAI initiatives out of fear of overstepping some yet-to-be-drawn boundaries. The best path forward splits the difference between exploration and caution by encouraging iteration on lower-risk use cases that will continue to promote development and understanding without pushing boundaries in an uncontrolled manner. There is no substitute for learning from doing, which is why continuing to put GenAI into production will position enterprises for success. It will also enable them to make valuable contributions to the field of responsible AI as they develop a more precise understanding of which guardrails are needed to balance innovation with responsible use.

Prompting with Purpose and Expertise

At the core of this experimentation lies prompt engineering, which directly influences the quality, relevance, and safety of AI outputs. While anyone with an internet connection can use commercial LLMs, there is no question that skilled users who understand how the underlying technology works will extract significantly more value from GenAI. Expert prompt engineering steers the LLM to produce outputs that users can confidently integrate into codebases, databases, and various components of a system. Better prompting yields better content.

The importance of human intuition coupled with domain expertise in getting value out of LLMs is why investing in training programs to enhance employees' AI literacy, prompt engineering skills, and AI-human collaboration techniques is one of the best ways enterprises can set themselves up for success. As the field of human-machine teaming continues to evolve, those with more knowledge of how AI works and more practice using it will be best positioned to integrate the technology into enterprise operations in the most impactful and responsible ways possible.

The Critical Role of Data Engineering

The success of enterprise GenAI applications requires mastering both prompt engineering and the art and science of data preparation—yet many underestimate the critical role of data engineering in steering LLM behavior. Massive context windows won't eliminate the need for careful data preparation; simply putting more raw content into an LLM won't lead to precise, reliable outputs. In some cases, adding too much context introduces new risks like overwhelming the attention span and diluting relevant information. In particular, retrieval augmented generation (RAG) applications must focus on effectively integrating high-quality data into the LLM application. Properly preparing the data helps models to perform optimally, increasing the likelihood that models will produce actionable, accurate, and robust outputs that can be used in decision making. In reverse, improperly or insufficiently preparing the data opens the door to hallucinations that offer little to no value for users.

Traditional data preprocessing techniques remain equally relevant to GenAI applications, including basic cleaning and formatting, such as removing personally identifiable information (PII) or irrelevant or incomplete data and converting the remaining data into a consistent format and tokenization, where text is broken into smaller, more manageable units. In addition, strategies like chunking optimization, metadata enrichment, structured hierarchical retrieval, and reranking helps to ensure RAG systems can effectively bridge the gap between raw information and actionable insights needed for enterprise-specific use cases and workflows. For instance, sliding window chunking approaches have shown to be well-suited for short to medium texts where continuity between segments of text is critical (e.g., conversations) and hierarchical models can be well suited for processing long, multitudinous, and complex documents.

This upfront investment in data preparation, though resource-intensive, ultimately improves response accuracy, reduces operational costs, and helps maintain regulatory compliance while protecting sensitive information through appropriate filtering and access controls.

What's Next for GenAI

Ongoing advances in GenAI's capabilities make it tempting to speculate about what LLMs will do next. But any prognostication about the technology's future that focuses exclusively on improving model performance would ignore several crucial and potentially mitigating factors.

For starters, the cost of training LLMs is skyrocketing and may ultimately become prohibitive. In August 2024, Epoch AI released a report detailing that while it's feasible from a technical standpoint to scale the training of AI models at the current pace, the cost of doing so will require developers to invest no less than hundreds of billions of dollars over the coming years. Even as unit costs come down, overall steep costs coupled with the uncertainty of downstream costs will be difficult to swallow for even the most cash-rich tech companies, especially considering that questions about the technology's ability to generate commensurate returns have grown more pronounced.

As the economics of models become challenging to navigate, the true cost of operationalizing GenAI will shift from model training to the point of inference: where AI is deployed in real-world environments. This will, in turn, change the focus of conversations about GenAI toward best practices for engineering end-to-end applications that use the technology for specific mission use cases, from cybersecurity to military applications to law enforcement. Enterprises will invest more time and effort into creating evaluation sets that track AI applications against traditional performance measures and keep tabs on costs. How enterprises engineer solutions will ultimately determine their ability to use GenAI to create value at acceptable price points.

Part of the way companies will adjust to this new reality is pursuing vertical integration. Hardware providers that are doing groundbreaking work within the GenAI stack will seek to create software layers, and software companies may seek to manufacture their own hardware. A good example of this vertical integration is NVIDIA. The company that is known around the world for building the graphic processing units that have powered the AI revolution is already building off its success in hardware to reimagine what software can do. NVIDIA has created Omniverse, a platform that enables developers to simulate physical worlds with remarkable precision, and NIM Agent Blueprints, which help developers build applications that use one or more AI agents.

What these developments show is that the true power of GenAI lies beyond the remarkable powers of the algorithms underpinning it. As with all technological revolutions, GenAI started out with a bang. Now comes the less glitzy, but no less impactful, stage: integrating it with end-to-end applications in cost-effective, innovation-driving ways.

Ernest Sohn, *an AI leader for Booz Allen's Chief Technology Office, delivers AI services and capabilities to help federal civilian clients meet mission-critical needs.*

Alison Smith, *Booz Allen's director of generative AI, leads GenAI solutioning across the firm and helps teams create best practices for AI development and use.*

SPEED READ

Generative AI's superpower is the speed at which it advances, a trait that can be seen in how far the technology has progressed since the launch of ChatGPT in November 2022.

But as the cost of training large-language models skyrockets, organizations will focus less on improving model performance and more on operationalizing the technology at the point of inference: where AI is deployed in real-world environments.

As this shift takes place, organizations that prioritize AI literacy and data engineering will be well-positioned to deploy generative AI effectively and responsibly.

Achieving Real-Time Cyber Defense

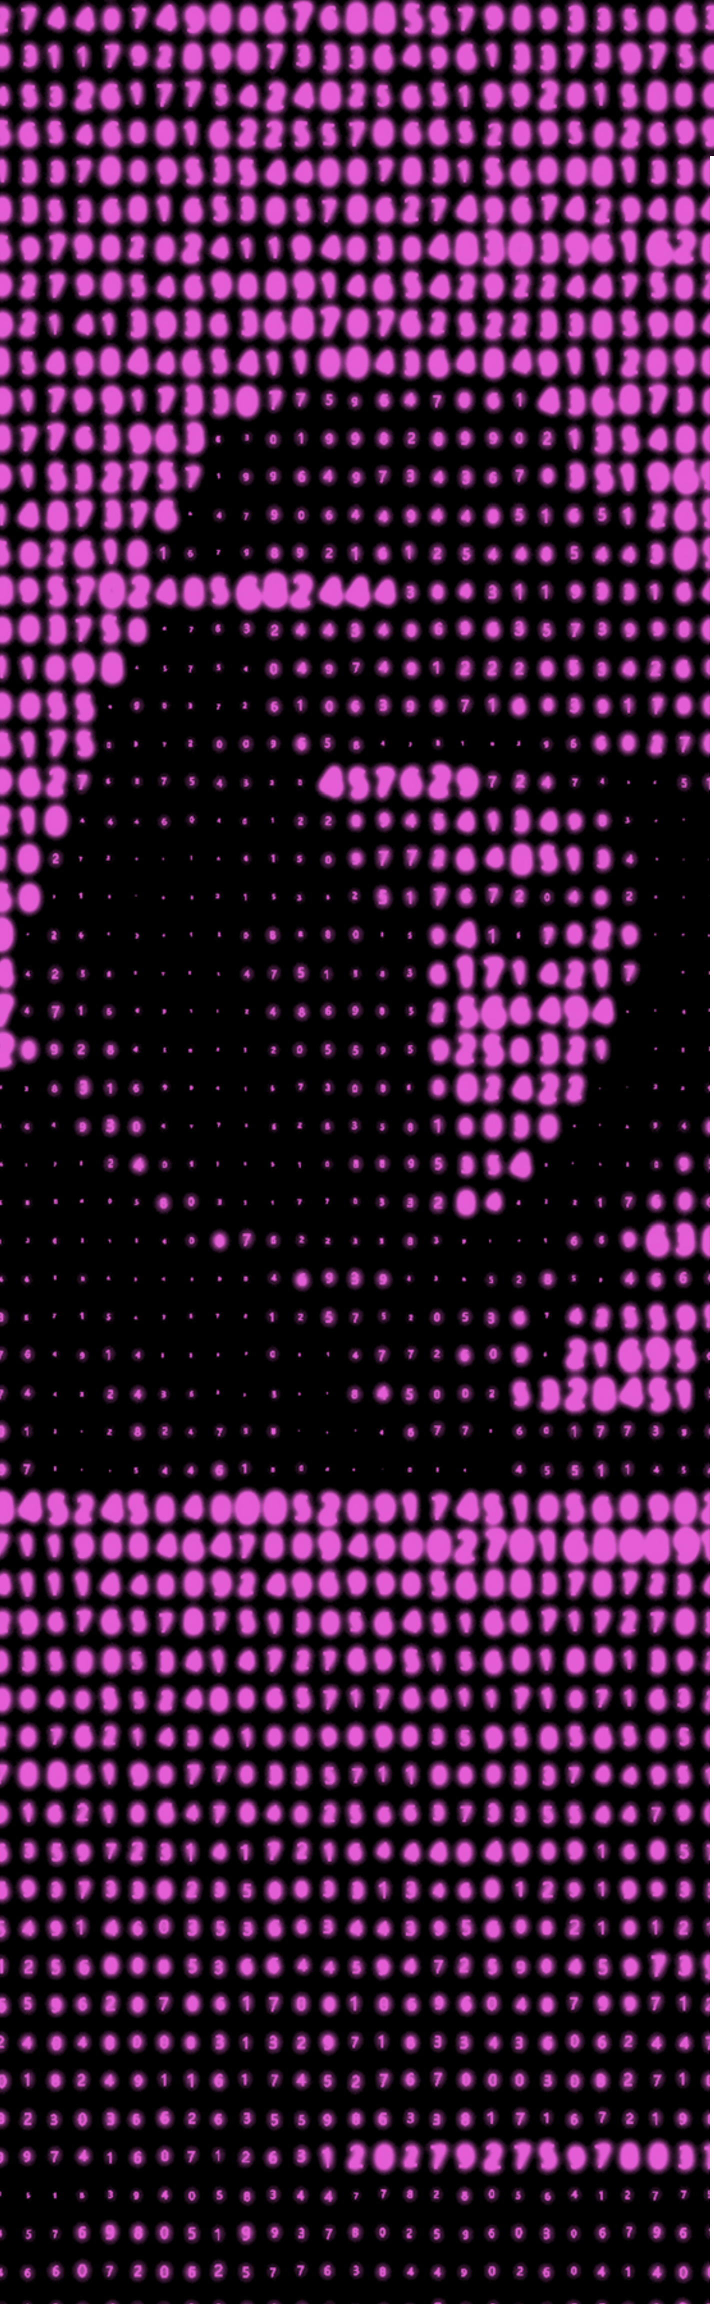

Bolstering Cyber Defenses with AI

Tony Sharp, Joe Gillespie, Matt Costello, Mike Saxton, and Rita Ismaylov

“It was a swift response and a successful containment,” the chief information security officer (CISO) said. But it wasn’t. The security operations center (SOC) intercepted the first intrusion—a carefully targeted spear-phishing email campaign. One executive forwarded the message to the security team along with its skillfully crafted bait and malicious PDF attachment. She told them that it looked like it came from a trusted contact, but the email address was unfamiliar.

Detonated in the SOC’s sandbox, the malware revealed itself, attempting to beacon back to its command-and-control (C2) infrastructure. Armed with the C2 address, the team quickly located the half dozen other malicious emails whose recipients had opened the attachment. Their compromised devices were isolated and disinfected.

Three days later, the team blocked the second intrusion. The attackers used password spray attacks on a misconfigured cloud services portal. Their automated script tried hundreds of logins in rapid succession with email and password combinations culled from previous data breaches. Alerted by one of their identity security tools, the SOC secured the portal and updated their blocklist with the IP addresses from which the login attempts had originated.

As his nighttime counterpart had done earlier with the spear-phishing attack, the daytime SOC manager prepared a report cataloging the attacker’s tactics, techniques, and procedures (TTPs) and the relevant indicators of compromise. He posted the report anonymously through their industry’s information sharing and analysis center to warn other enterprises in their sector about the attacks.

“In both of those incidents,” the CISO correctly insisted afterwards, “the team did everything right.” But they still missed the third intrusion.

The fatal vulnerability turned out to be a web server that was spun up five years earlier to support a failed and swiftly canceled series of corporate marketing events. It hadn’t been updated in years, and scans conducted by the security team didn’t flag it because the server used the never-reported and now-forgotten brand of the marketing event, rather than the company’s name, as its domain.

Buried in that server's file directory were years-old Secure Shell (SSH) keys that still provided trusted access to the organization's main cluster of marketing servers. Once there, the hackers were able to swiftly pivot and gain domain access to SharePoint and OneDrive. At that point, as revealed in logs they later recovered, there was a 48-hour pause—probably while the hackers sold their access to a ransomware gang.

The SSH connection from the abandoned web server was an "anomalous" event. The security team's incident and event management platform flagged it as such using the code "amber/anomalous," the lowest level of alert. The platform flagged hundreds of other amber events that day and dozens classed as "red/suspicious." None were rated "flashing red/ hostile." The security team's half dozen analysts attempted to resolve as many as they could, but no one checked on the anomalous SSH connection.

"When you have that many false positives, stuff is occasionally going to slip through the cracks," the CISO acknowledged. "The bottom line is we don't have the manpower to resolve every anomaly report."

The postmortem by a third-party security company discovered that the scans of and attacks on the abandoned marketing server originated from the same Internet Protocol (IP) address range the SOC placed on their blocklist because they were the origin of the password spray attacks. There were also links through Whois Database Search to the spear-phishing email campaign. Putting the IP addresses on the blocklist didn't protect the abandoned server because— as an unrecorded, abandoned shadow IT asset—it was not subject to the company's security policy.

"If only we'd had a way to sort through those alerts, separating the wheat from the chaff, and the capability to correlate data from the significant ones, we might have spotted the connection and stopped the third attack," the CISO said in the postmortem internal inquiry. When they returned to the main marketing cluster with domain access after 48 hours, the hackers began exfiltrating and encrypting data.

Putting Security Teams on the Front Foot

This vignette is fictional, but security personnel routinely deal with attacks like those described above. Today, even well-resourced enterprise cybersecurity teams face an increasingly untenable operational tempo, with little opportunity to do meaningful triage and isolate the most dangerous intrusions. More data about system conditions and potential attacker activity is only useful up to the point where it can be analyzed; beyond this it only increases the size of the haystack concealing the needle of potential threats.

Combine that data overload with the growing complexity of modern enterprise IT systems—multi-cloud infrastructure, multi-vendor toolsets, multi-jurisdictional compliance—and you have a recipe for failure for even the most experienced and best resourced SOC. That recipe is worsening now that attackers can use generative artificial intelligence (GenAI) to scale personalized, carefully targeted spear-phishing attacks from a handful of targets to hundreds.

The good news is that enterprises can flip the script on current cybersecurity practices. By leveraging AI and machine learning (ML) early in the alert lifecycle, they can quiet the noise of nonstop false positives, optimize their incident response workflow, and help security team members perform critical tasks faster and more efficiently. AI-enhanced intelligence can improve security teams' ability to conduct proactive threat hunting by using at-scale surveillance of gray space to analyze pre-attack activity and figure out how attackers plan to breach a network.

Overwhelmed on Two Fronts

The irony is that the security team in our fictional example had all the data they needed to discover and halt the security intrusion. They just didn't know they had it or how to leverage it. This is true in a surprisingly sizable proportion of real-world attacks, and it happens because security teams are increasingly overwhelmed across two fronts: technology density and human resourcing.

ON THE HUMAN SIDE, THE ISSUE IS COGNITIVE RESOURCE CONSTRAINTS:

It's difficult to recruit experienced and talented security staff. This problem is pronounced in high-demand niche specialties like cloud security or AI security. In such a competitive labor market, churn can be an issue for junior and mid-level practitioners, and grow-your-own talent approaches take years and require significant investment.

Even properly calibrated Security Information and Event Management (SIEM) and endpoint detection and response (EDR) platforms produce an unmanageable volume of alerts. The vast majority of these alerts are false positives or evidence of quotidian security threats. Integrating tools to provide contextual data is considered a best practice but can produce duplicate alerts if done incorrectly. Even when done well, integration can end up replacing alert fatigue with context fatigue—too much additional data can be overwhelming rather than informative. Triaging and resolving these alerts absorb a considerable number of analyst hours even for a well-resourced and fully staffed SOC, and the cognitive burden of sifting through nonstop alerts can lead to retention issues. There's also an opportunity cost: Analyst hours wasted sifting through noise could be spent proactively hunting for the most dangerous threats.

AI-enhanced intelligence can improve security teams' ability to conduct proactive threat hunting by using at-scale surveillance of gray space to analyze pre-attack activity and figure out how attackers plan to breach a network.

A security team that allows an alert queue to set its agenda will always be on its back foot. A dearth of analytical insights means security teams spend too much time chasing alerts that may or may not be evidence of an intrusion, without knowing who's behind the alerts or how dangerous they are.

ON THE TECHNICAL SIDE, THE ISSUES ARE EVEN MORE VARIED AND COMPLEX:

Machine speed data sharing is still mostly an aspiration rather than a reality. For complicated and assorted reasons, many organizations don't or can't use Structured Threat Information eXpression (STIX) and Trusted Automated eXchange of Intelligence Information (TAXII) or other standard formats for cyber threat intelligence data so it can be imported automatically into defensive tools. Without this kind of real-time, machine-speed data sharing, intelligence indicators have to be manually entered into firewalls or SIEM tools, often by cutting and pasting.

Security teams employ multiple vendor toolsets and platforms, often each with its own dashboard and control interface. Integrating toolsets is not as easy as it should be, and some vendor tools are unable to seamlessly ingest data from competitor products or can do so only at great cost.

Beyond the security team itself, IT environments in modern global enterprises are extraordinarily complex. Devices with varying degrees of trustworthiness need to move on and off the network in different places. Infrastructure, data storage, apps, and services are spread across multiple cloud providers (and sometimes on-premise data centers) and endpoints, managed with constantly shifting sets of tools from multiple vendors. Integrating vendor tools often creates another level of complexity and can result in software bloat.

Adding to this complexity are the millions of data points that modern cybersecurity tools create. Ingesting and processing event/log data from a large enterprise can be computationally challenging for SIEM and other security platforms.

The traditional hub-and-spoke model for the security team—where data is brought back to the center and analyzed, and then defensive measures are centrally deployed through changes to SIEM, firewall, or EDR rules—simply doesn't scale for a large modern enterprise given the volume of data involved. The cost and the compute don't add up.

The traditional defensive model—which, when it comes to high-end attackers, boils down to assuming they've already broken in and searching the network for traces they might have accidentally left as they moved around—doesn't scale either. It leaves defensive teams permanently on the back foot.

These challenges are getting more severe as enterprises confront the growing threats of AI-powered cyberattacks. Generative AI bots specifically designed for cybercrime, like Evil GPT or WormGPT, have been advertised on Dark web hacker forums for at least two years, and a team of researchers from Indiana University found real-world instances of cybercrime actors using GenAI tools to write malware, generate phishing emails, and set up scam websites.

A Data Problem, Not a Defense Problem

The brutal truth is security teams are often too overwhelmed by their data problem to mount an effective defense. There is too much data arriving too quickly for human team members to timely distinguish relevant threats from irrelevant noise. In addition security teams don't have enough insight into how seriously threat actors might seek to compromise their network. By exploiting the growing capabilities of AI for data analysis to focus on relevant threats and draw conclusions about attacker TTPs from at-scale surveillance of pre-attack threat group activity, cybersecurity teams can finally get ahead of their assailants.

Quieting the Noise

When deployed early in the alert cycle, ML and AI help threat detection teams to better manage the thousands of duplicative and mostly false-positive alerts received daily from various cyber tools.

- ML tools using unsupervised learning can group together duplicate alerts from different tools and isolate anomalies that are most likely to be significant.
- Label collection and model training enable AI tools using supervised learning to predict the likelihood of a true positive, which in turn dynamically filters thousands of daily alerts down to a handful of high-fidelity warnings that go to the incident detection team for review.

It's also possible to use an ML model to prioritize critical alerts and flag them as essential so that analysts don't need to judge the likelihood of an alert's importance. Powerful AI algorithms allow organizations to automate this process with confidence, which helps to deprioritize unlikely events and does not rely on humans alone to assess the viability of an alert being fatal.

Integrating these tools into existing workflows and customizing them to meet the needs of the system keeps analysts focused on dealing with true positive alerts, which mitigates risk by decreasing the number of alerts that require human attention and lessening alert fatigue.

Ideas in Action

Private companies and federal agencies are already using AI to bolster their cyber defenses.

A global automaker, for example, set out to identify gaps in its security posture, formulate a roadmap to get to a more effective security operation, and increase visibility into the overall threat landscape. Scale was a big issue: 1,200 unique data sources of cyber information; numerous sites; and billions of messages from cyber tools every day. Analysts were left to deduce meaningful patterns from unwieldy datasets and spent excessive working hours analyzing false positives.

The automaker addressed these challenges by adding ML and AI early in the alert lifecycle. It also created an easily deployable automated template that could be deployed quickly in any network environment and cloud-based data pipelines large enough to manage the huge volumes of cyber information. The project also set up a cyber-AI capability that provided users across the security team with access to custom big data, ML, and AI intelligence products.

These solutions made an immediate and considerable impact. Instead of dealing with an unwieldy alert feed, the team is able to direct its attention to true positive alerts and guarantee the delivery of security events within seconds from anywhere on the planet.

In the federal government, a large and critical mission needed a more effective way to discover and remediate vulnerabilities. They began piloting a solution that uses AI to pull real-time assessments of system risk into a detection suit. This AI-powered risk analysis helps quantify how an adversary might exploit a given system, which improves the security team's understanding of risk as part of the monitoring approach. The pilot application has enabled the team to detect cyberattacks that circumvented traditional defenses within seconds while significantly reducing false positives and alert fatigue.

Breaking Incident Response Bottlenecks

AI and ML can help widen bottlenecks and free up blockages in incident response. For example, certain incidents require the same type of resolution each time they recur. In addition, it's not uncommon for a familiar type of false positive to hit a network under predictable circumstances. Rather than send these incidents to an analyst's queue for remediation, large language models (LLMs) can identify them and expedite their resolution.

As part of an automated enrichment workflow, an LLM model can be integrated into a security orchestration, automation, and response solution to initiate a chat conversation between the user and the AI. The AI is provided the context of the alert with instructions to gather more information from the user and provide a summary of the conversation, which is added to the case. Separately, a ML algorithm can be layered into this approach to classify the type of incident received as a benign alert or malicious threat based on similarities to previous incidents.

LLMs can also be useful to members of security teams. When added to a tool like Slack, an LLM can update an analyst about the latest ticket remediations or provide a summary of what happened in the last 10 days. From a user experience standpoint, LLMs can have the capacity to reach out to users that provided information in a ticket, a process that currently consumes a significant amount of time for analysts. Rather than dedicate human hours to this process, AI can interact with the user to identify gaps and gather additional information the analyst needs to complete the investigation.

Pulling Insights from Gray Space

AI's greatest potential use case for cybersecurity may be its power to evolve threat hunting from a large-scale guessing game to an intelligence-driven approach in which security teams exploit new sources of actionable intelligence to identify and block attacker TTPs before they're employed against the enterprise.

Before hackers strike, they prepare infrastructure (like fake login sites for phishing campaigns), write or fine-tune malware, and otherwise engage in activity that provides strong indications of planned future attacks. Typically, this activity is not conducted in the attackers' own infrastructure ("red space"), nor is it deployed at this stage in the defenders' infrastructure ("blue space"). That comes later, during the attack itself.

Instead, such attack preparation activity commonly takes place in compromised or rented infrastructure temporarily controlled but not owned by the hackers (gray space). Gray space locations can be monitored via the public internet, and many experienced or skilled threat analysts probably track a handful of such locations used for this purpose by threat actors, scouring data from their activities to make informed predictions about future attack campaigns.

Using the growing capabilities of AI, monitoring and analysis work can now be scaled. Internet-wide scans and other comprehensive data-gathering techniques (e.g., analysis of all internet domains established in the last 24 hours) can be used to find evidence of malicious activities, which can then be monitored. AI can draw conclusions from that monitoring data to define TTPs highly likely to be used in future attacks.

The data can then be cross-referenced with information about the enterprise network—its vulnerabilities and system posture—to create detection and prevention methods on the fly. Such methods are derived from using AI/ML to compare the current system posture with the vulnerabilities attackers plan to exploit and devising countermeasures that can be implemented by security tools.

AI's greatest potential use case for cybersecurity may be its power to evolve threat hunting from a large-scale guessing game to an intelligence-driven approach.

Empowering Security Teams for Success

Returning to our fictional vignette, as noted, the team had access to all the data they needed to prevent the third attack, but they didn't know they had it, and they lacked the tools they needed to use that data to defend themselves. Using the approaches we've outlined here, our fictional security team would have been set up for success, not failure.

- AI/ML sorting of alerts, elimination of false positives, and automatic resolution would have reduced the hundreds of amber/anomalous reports to a number where analysts could have investigated each one, found the abandoned marketing server, and blocked the third attack.
- Correlation of data from the first two attacks, either by analysts freed up from repetitive "work by AI/ML (or by AI/ML itself)", would have provided intelligence warnings that could have prevented the third attack.
- Monitoring of the threat actors' gray space operations would have provided data that could have automatically blocked the password spray attacks and might have identified the TTPs used by the intruders in the attack from the marketing server.

There's a lot of hype about AI, but these are use cases where judicious application of AI tools can transform the security landscape for large enterprises.

Tony Sharp *leads Booz Allen's Commercial cyber technology solutions practice.*

Joe Gillespie *is a senior leader in Booz Allen's National Cyber business.*

Matt Costello *leads Booz Allen's Commercial analytics and AI business and our insider threat practices.*

Mike Saxton *is the director of Booz Allen's adversary-informed defense business and leads the federal threat hunt and digital forensics and incident response team.*

Rita Ismaylov *leads Booz Allen's Commercial AI business with a focus on driving the design of AI solutions and helping clients grow their AI capability.*

SPEED READ

Security teams face mounting technical and human challenges: a shortage of skilled personnel, an overwhelming volume of alerts across complex multi-cloud environments, and the rise of AI-powered cyberattacks are putting defenders increasingly on the back foot.

AI and machine learning can transform cybersecurity by filtering out false positives early in the alert lifecycle and automating routine incident responses, allowing human analysts to focus on serious threats.

By using AI to monitor gray space activities—where hackers prepare their attacks—security teams can shift from reactive defense to proactive threat hunting, identifying and blocking potential attacks before they happen.

MISSION SPOTLIGHT: SPACE

Taking the Ground Out of Ground Systems

FOR THE SPACE GROUND SEGMENT, DATA SCIENCE IS ROCKET SCIENCE

Josh Perrius and Ginny Cevasco

The Russia-Ukraine war is arguably the first commercial space war on account of its use of private companies' imagery for tracking and targeting. The conflict has demonstrated the advantage of knowing what's happening on the ground and highlighted the role of the space domain for the critical Indo-Pacific region and hot spots around the world. Now, countries across the globe are speeding up their space programs, adding urgency as the U.S. accelerates space modernization.

The proliferation of next-generation satellites is fueling demand for next-generation infrastructure. Virtual ground systems, which use software-defined infrastructure and cloud computing to control satellites, offer a giant leap forward for space operations. These virtual systems are not tethered to a single brick-and-mortar facility. They are more secure and easier to update than hardware-dependent legacy systems. They also allow for the seamless integration of artificial intelligence (AI) and machine learning into space operations. These advances allow operators to process the data that satellites collect rather than simply store it or delete it as unusable. The result is a flood of rich insights that can improve operators' real-time decision making.

Ground Systems Do the Thinking

While agile constellations of small satellites capture the imagination due to their resiliency and flexibility, it's the "brains"—the ground system—of each network that direct the mission and secure its success. Command and control (decision making), managing the mission (planning, testing, and operations), and processing and disseminating data are all dependent on the ground system and its operators.

Automating key activities will allow operators to focus on critical decisions and enable seamless integration of systems that are currently stovepiped. Modernizing these functions is critical to keep U.S. space aspirations aloft.

Though $150.4 billion was spent on ground segment technologies in 2023 according to a June 2024 press release from the Satellite Industry Association, many space organizations still have to contend with legacy equipment.

What's more, there's no guarantee that new ground systems are using the latest software innovations needed to make systems smart, scalable, and secure. That's because new ground systems are often included in the contract for an overall space system, which means contractors used to building legacy systems may build the system using techniques they're already familiar with rather than the latest software innovations.

Why Operators Need Agile Systems

Imagine you're a satellite operator working in a military or intelligence organization. Chances are, you're working at a brick-and-mortar ground station performing actions such as maintaining satellite surveillance of a given region, tracking satellites and other space assets, or providing satellite imagery to support a military operation.

You do this by directing a satellite using instructions, or "taskings," which include multiple actions: defining the area of interest, finding an available constellation, deciding between radar or optical images, ascertaining the viewing angle, and so on. You track different assets across several screens, which adds to the complexity of your work.

Let's say a satellite spots an adversary's troops moving on the ground near a contested area. That satellite isn't scheduled to make another pass over the area for eight hours, so you decide to task another satellite for reconnaissance. It's an inefficient and time-consuming process.

First, you need to decide whether to turn to a commercial provider or use a classified source. Then, depending on what assets are available and conditions that are beyond your control, it could take anywhere from minutes to hours to get one or more satellites in place and begin downloading images.

For example, if the weather is partially overcast in that region, you need synthetic aperture radar to pierce cloud cover. You also need a satellite with a high-resolution

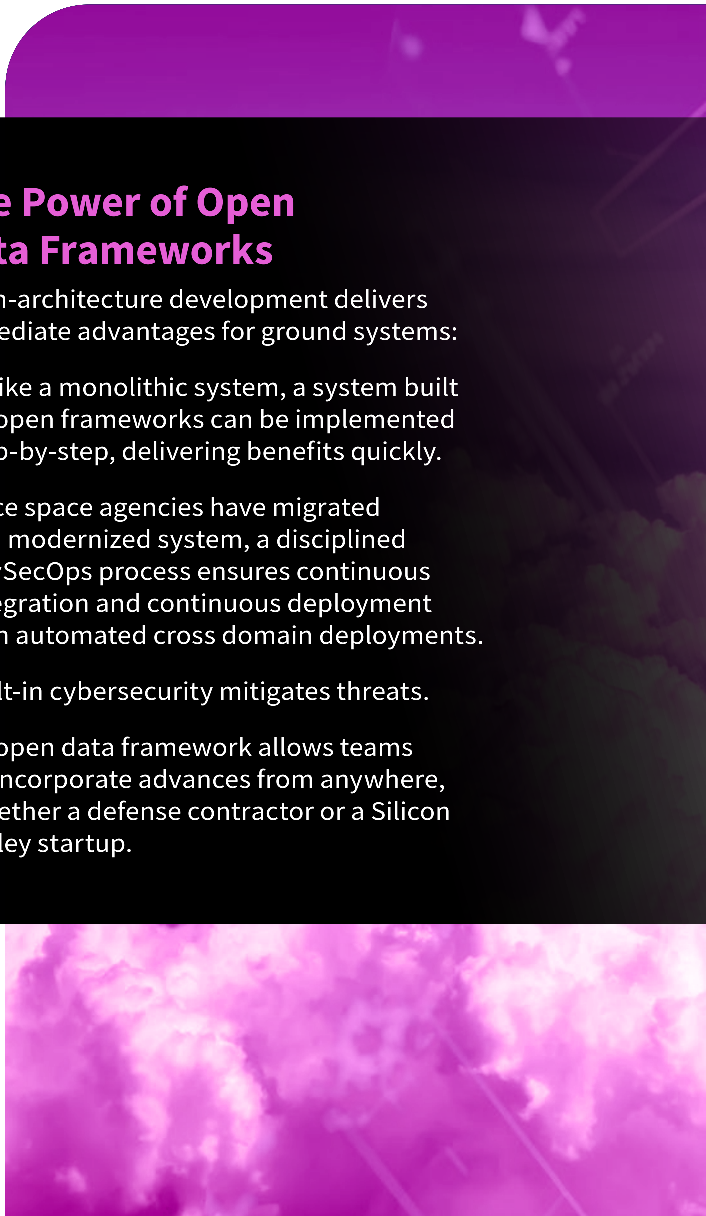

The Power of Open Data Frameworks

Open-architecture development delivers immediate advantages for ground systems:

- Unlike a monolithic system, a system built on open frameworks can be implemented step-by-step, delivering benefits quickly.
- Once space agencies have migrated to a modernized system, a disciplined DevSecOps process ensures continuous integration and continuous deployment with automated cross domain deployments.
- Built-in cybersecurity mitigates threats.
- An open data framework allows teams to incorporate advances from anywhere, whether a defense contractor or a Silicon Valley startup.

camera—and software-defined solutions that enable it to identify the overcast spots and avoid collecting useless data there—to identify equipment in transit.

You're feeling the pressure, as every minute the situation on the ground is unfolding. Meanwhile, you know the adversary's cyber force might be attempting to penetrate the ground station's aging infrastructure or take the system down altogether.

Streamlining Operations with Data-Centricity

The good news for operators is that data engineering enables software advances that provide new capabilities and flexibility. According to the Department of Defense (DOD) Data Strategy that was released in September 2020, success begins with data-centricity, putting data at the center of the organization and its systems—enabling the U.S. to act on data faster than its adversaries. This approach can be summarized as meaning the network must accommodate the data rather than the other way around.

For ground systems, this shift can be achieved by building open-architecture systems that can adapt for changing missions. DOD's September 2023 Space Policy directive emphasizes the need to accelerate the transition to more resilient architectures, and DOD mandates modular open-system approaches.

Open-architecture systems allow for built-in zero trust security, providing granular access controls while enabling continuous change. Implemented by technologists who understand the intricacies of the mission, they allow space agencies to build a system around core goals while keeping capabilities flexible to adapt for future priorities and challenges.

The Benefits of Virtual Ground Systems

Add an open-architecture platform to the cloud and you have the option to bypass physical infrastructures entirely. Forward-looking operators are already opting for virtualization, a cloud-based approach using a virtual computing environment. This option provides "hosted anywhere" flexibility and enhanced security—developed faster than physical structures and at lower cost.

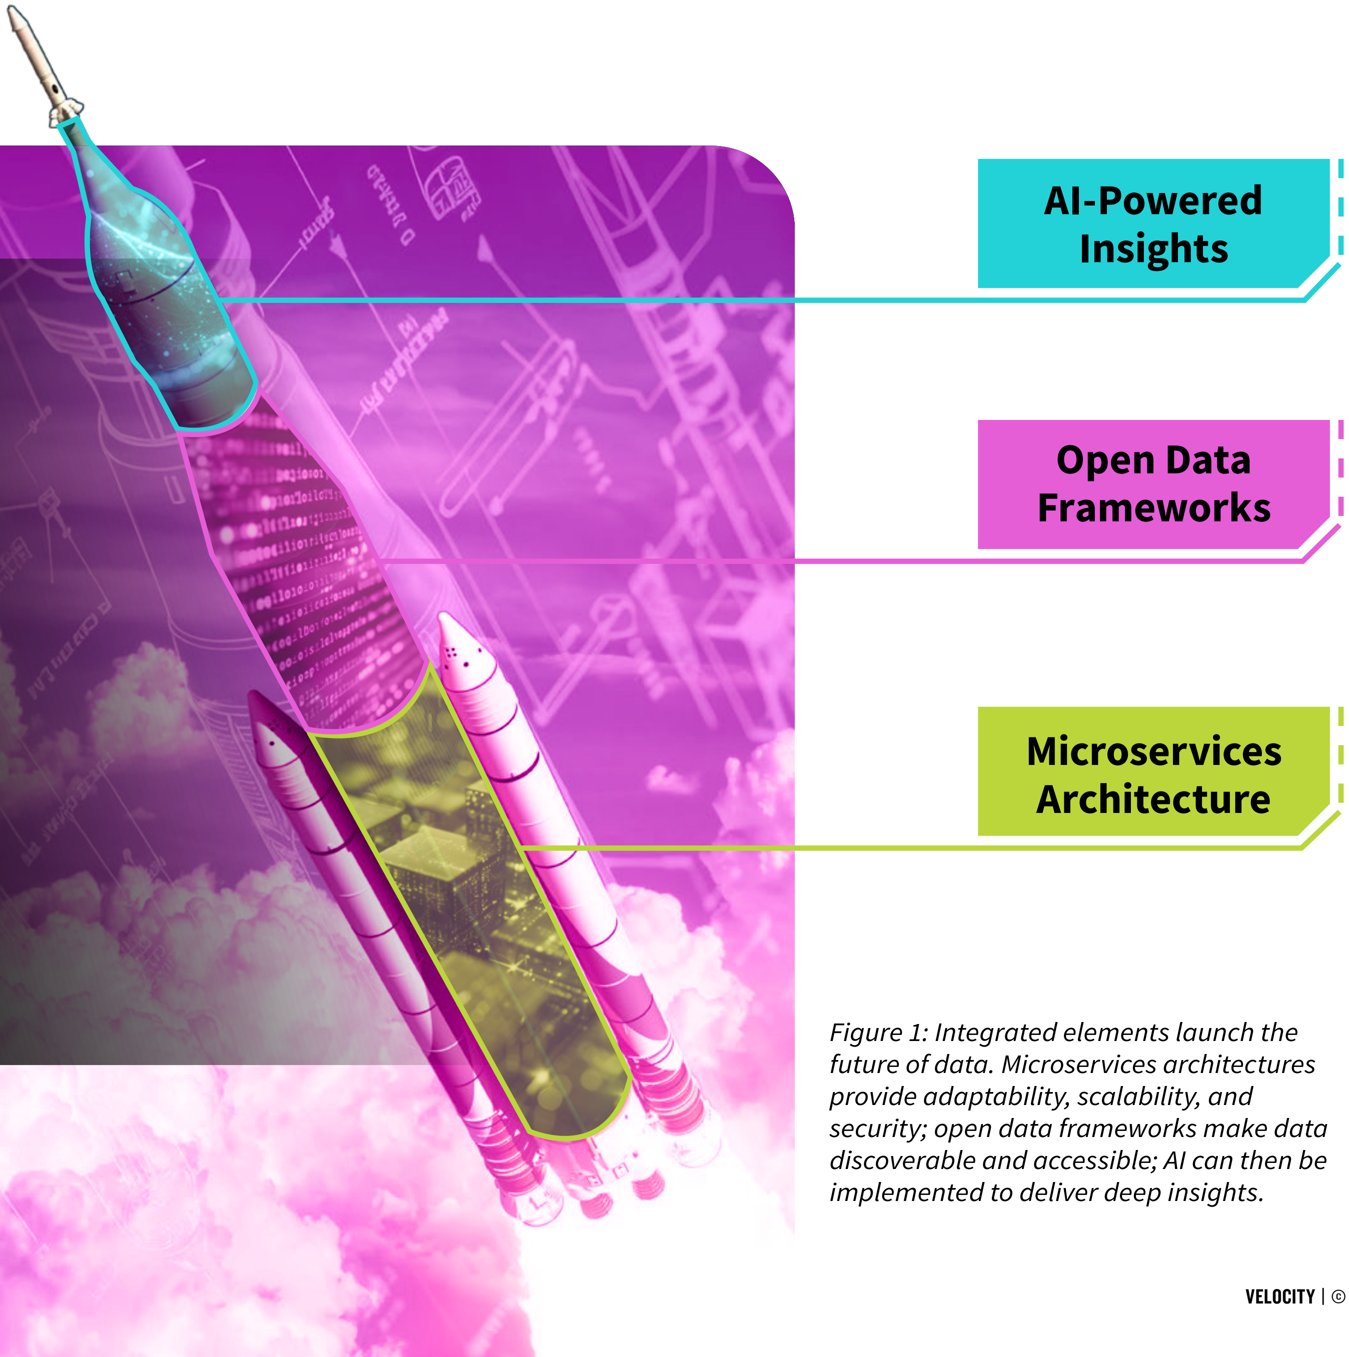

Figure 1: Integrated elements launch the future of data. Microservices architectures provide adaptability, scalability, and security; open data frameworks make data discoverable and accessible; AI can then be implemented to deliver deep insights.

With a virtualized ground system, stakeholders receive benefits across functions, realizing:

- **Intelligent, automated command and control** for more rapid satellite moves, greater precision, and analytics that include courses of action
- **Improved mission management** for faster planning and smoother upgrades and integration of sensor data feeds across systems and constellations
- **Rapid data transformation** for fast data collection, processing, storage, and dissemination
- **AI and machine learning** to continuously advance capabilities with models that change at the speed of mission

A virtual space ground system using a multicloud platform gives organizations the flexibility to adopt new technologies and adapt to evolving mission requirements. It allows satellite managers, operators, and technicians to manage missions on a single laptop using a plug-and-play architecture that delivers both zero trust security and resiliency to stay operational after an attack. Also critically important, intelligent automation reduces the space tasking process from days to hours—or even minutes.

The Elements of Open Architectures

Just as rockets are propelled by liquid fuel made of elements such as hydrogen and oxygen, modernized ground systems—whether traditional or virtual—are propelled by data-centric elements.

Federal agencies can incorporate these elements even into legacy systems, giving them the benefits of modernization which the Space Force, as the newest agency, is already using as a foundation to build on.

Microservices architecture—Breaking monolithic services into modular, independent elements enables scalability and customization. For example, stakeholders can adapt an existing containerized service as-is, replace it, or adjust it for a specific use. Open application programming interfaces (APIs) ensure interoperability with any technology and support rapid updates to take advantage of technological advances.

Open data platform—This modular data ecosystem ingests, translates, stores, organizes, and makes data discoverable and accessible. This includes DevSecOps tools with a deployment framework; a continuous integration/continuous deployment (CI/CD) pipeline; a data platform that onboards and stores new feeds securely; and a centralized data catalog, allowing for easy discovery using common metadata standards.

Cross domain shared services—As multidomain and joint operations become more critical in the face of increasing global threats, sharing mission-relevant data flowing up or down classification levels must become faster and easier. At the same time, stakeholders need to fuse multi-intelligence and multidomain data for more nuanced insights.

Zero trust architectures—Employing granular security to space data is essential to protect space assets and ensure partners can safely share data at varying security levels. To provide effective cybersecurity and rapid ground system development, data engineers must build on a platform that ensures zero trust's core principles—assume a breach, never trust, always verify, and allow only least-privileged access—are built in from the ground up.

Transforming Space Missions

Open data frameworks are key, enabling organizations to integrate the siloed data they've been collecting by taking advantage of machine learning operations (MLOps)—the ability to build, train, and tune AI models for insights on what has happened and how to respond.

Let's go back to the scenario where you're the operator. Think of the difference it would make to be notified of a situation and be provided with recommended courses of action.

- **For defense operations**, a model can incorporate multi-intelligence images of a region to analyze trends and predict the likelihood of a military attack.
- **For climate change**, a model can incorporate past events and information from diverse sensors to provide a more complete assessment of environmental impact in a certain area.

Integrating Large Language Models

The data-centric framework provides secure and repeatable processes, with a DevSecOps pipeline connecting directly to the data. This approach allows space organizations to field AI solutions rapidly and apply the benefits of large language models (LLMs), deep learning algorithms able to solve complex problems.

Moreover, forward-looking virtual space ground stations can provide the power of networked LLMs. These algorithms, each trained on vast datasets, share knowledge to generate rapid insights and provide recommended courses of action.

Accelerating the Mission for You, the Human

Let's return to the example where you're the operator of a satellite observing troop movements below. Instead of being in a decades-old facility juggling multiple screens, you're now at a remote location, working on your laptop.

Thanks to virtualization and automation, you're experiencing a very different scenario. The virtual ground system has alerted you to unusual troop movement in the area and provides a list of commercial satellites that are available to be tasked for Earth observation, along with key insights such as the cost for using each one, the image fidelity you can expect, and the dwell time that you need to factor in. Finally, it provides a two-part recommended course of action:

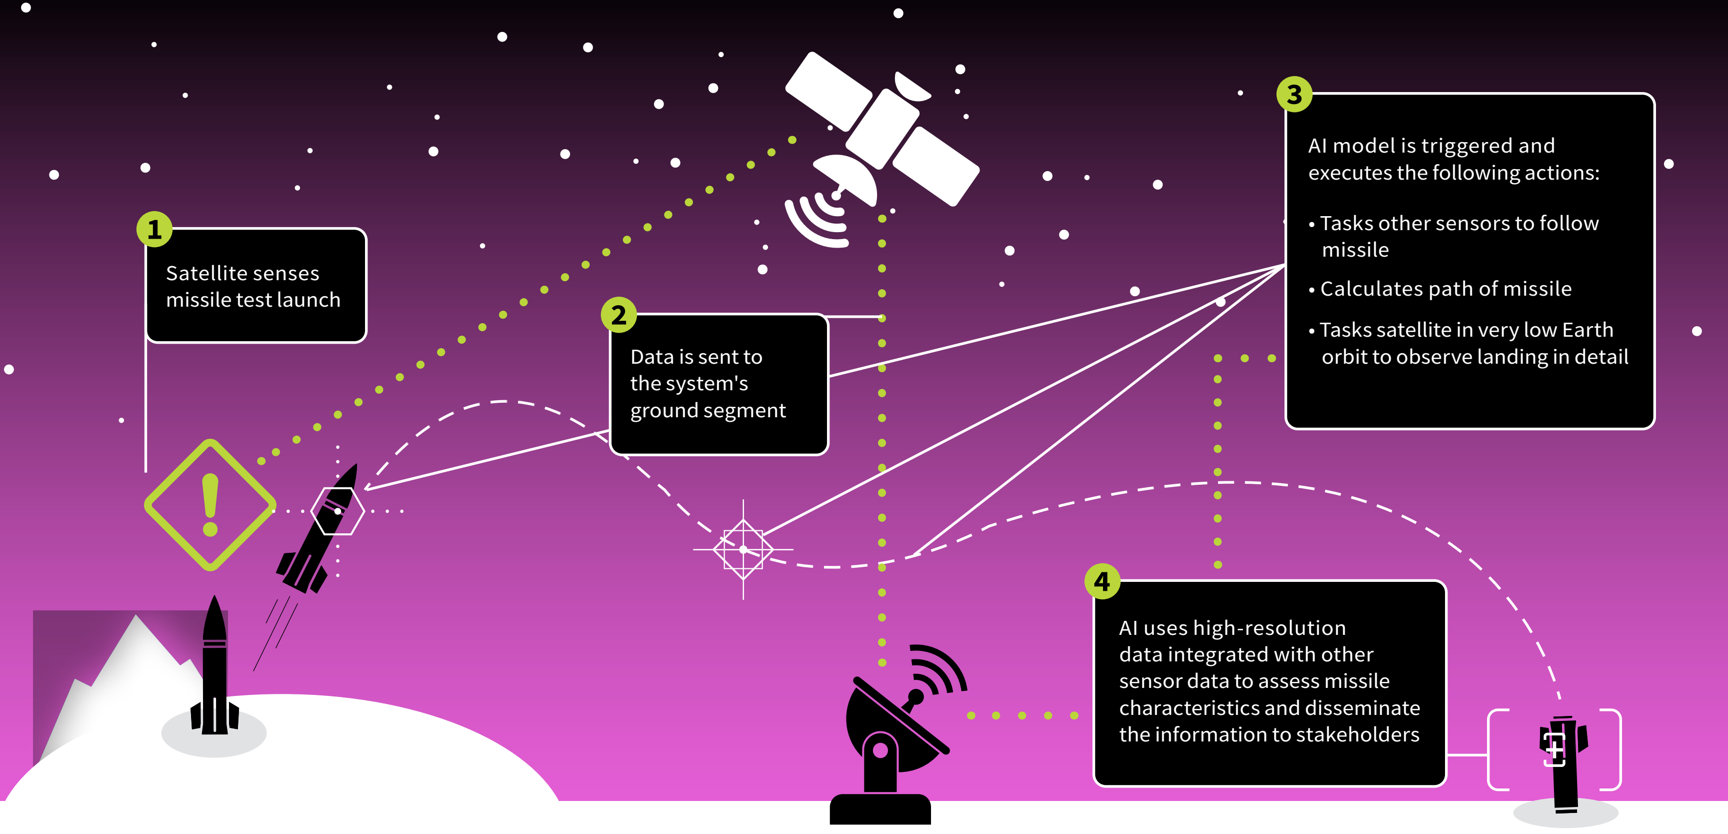
1
Satellite senses missile test launch
2
Data is sent to the system's ground segment
3
AI model is triggered and executes the following actions:
• Tasks other sensors to follow missile
• Calculates path of missile
• Tasks satellite in very low Earth orbit to observe landing in detail
4
AI uses high-resolution data integrated with other sensor data to assess missile characteristics and disseminate the information to stakeholders

Figure 2: Applying AI to every phase of space operations translates raw data into actionable insights and decreases the time from analysis to action.

- Task a commercial Earth observation satellite and a classified radar satellite to do a flyover.
- Notify allies in the area of troop movement and its coordinates.

You review the recommendation and press Execute. The virtual system automatically executes both parts of the recommendation. Within minutes, naval vessels nearby move toward the area in a show of force, and the adversarial troops, caught off guard by the rapid action, begin to disperse.

This is just one example of how smart, scalable, secure virtual ground systems can help humans enable the one thing that's better than winning a war: avoiding conflict altogether through effective deterrence.

__Josh Perrius__ and __Ginny Cevasco__ are senior leaders in Booz Allen's National Security Space business.

SPEED READ

The proliferation of small satellites is fueling demand for cloud-based ground systems that allow space agencies to adopt advanced capabilities and avoid the cybersecurity risks of brick-and-mortar facilities.

A virtual space ground system using a multicloud platform and plug-and-play architectures makes it possible to manage space missions on a single laptop and can reduce the satellite tasking process from days to hours—or even minutes.

With virtual ground systems, open data frameworks enable organizations to leverage machine learning operations (MLOps)—the ability to build, train, and tune AI models for insights on what has happened and how to respond.

Sensema REIMAGINED

See How Next-Gen Systems Can Outpace Urgent Threats

Don Polaski, Marissa Beall, and Christopher Castelli

king

At dawn, Grace, an all-source intelligence analyst for a three-letter agency, logs in to her workstation at headquarters. From her desk inside the Beltway, she turns her top-secret gaze to the Indo-Pacific region, half a world away. It’s the same routine she’s had for three years in the role. But this morning, February 3, 2027, is different. She sees an urgent email from her supervisor. What comes next, she can’t stop. All at once, her brows arch up, her eyes widen, her jaw drops, and she gasps.

The headline is clear: “China Buys Uninhabited Fijian Island to Build Military Base.” The news comes from the Beijing bureau of a major U.S. daily. “A Chinese development firm with links to the Chinese government and People’s Liberation Army today announced that it recently purchased the uninhabited Cobia Island from the government of Fiji for $850 million,” the story says. “Western security analysts assess that China plans to use the island to build a permanent military base in the South Pacific, 3,150 miles southwest of Hawaii.”

The news means the People’s Republic of China (PRC) can boost its power projection in the Pacific Islands, a region of strategic importance, expanding the presence of its military—the People’s Liberation Army (PLA)—near U.S. and Australian soil.

It’s the kind of surprise that Grace and her intelligence community (IC) colleagues always try to avoid. They know in their gut they must focus on hard targets to collect information on adversary plans that would otherwise go undetected, just as a Center for Strategic and International Studies (CSIS) report recommended. A hard target is “a person, nation, group, or technical system” that poses a potential threat and a very difficult intelligence-gathering challenge—it’s “often hostile to the U.S. or heavily protected, with a well-honed counterintelligence capability.” In this case, the PRC succeeded in keeping the U.S. in the dark.

“It’s why we call it the pacing challenge,” Grace tells herself half-heartedly. It’s no consolation.

The Truth in the Fiction

This imaginary scenario starts with a fictional headline borrowed from a 2017 National Intelligence Council global trends report. Behind our invented scene is the IC's real need to get ahead of surprises by accelerating its ability to collect, analyze, and make sense of vast amounts of data. This challenge is intensifying as the number and quality of sensors rises around the world.

Although Grace's gaze is set on the Indo-Pacific, no foreign power has a monopoly on inscrutability. Decades ago, Winston Churchill described the challenge of forecasting Russian actions as "a riddle wrapped in mystery inside an enigma"—and that still sounds about right. To address the full range of national security challenges around the world, IC analysts and warfighters need to make sense of complex, uncertain situations—and to do that, they need sensemaking systems that can collect and analyze data on an exponentially increasing scale.

Grace is a stand-in for real all-source analysts tasked with monitoring PRC foreign influence. All-source analysts review and analyze all kinds of data to assess, interpret, forecast, and explain a range of national security issues and developments that are regional or functional in nature. They use the full complement of all-source tools, techniques, and procedures to fuse all available intelligence into a comprehensive product designed to satisfy unit-level essential elements of information (EEI) and command-level priority intelligence requirements (PIR).

Sensemaking is the fundamental cognitive ability human analysts use to understand emerging developments and anticipate future events. It's not the work of a lone analyst but happens with the fusion of many viewpoints. Increasingly, technology plays a supporting role. A decade after the 9/11 attacks, the Intelligence Advanced Research Projects Activity (IARPA) described sensemaking as "the remarkable human ability to detect patterns in data, and to infer the underlying causes of those patterns—even when the data are sparse, noisy, and uncertain." IARPA also called for the development of new automated analysis tools designed to replicate the strengths of this ability. A 2019 IC strategy for augmenting intelligence using machines defined sensemaking as "a process of creating understanding in situations of high complexity," urging the construction of shared models "to provide the basis for trust between human and machine teams."

Had the team in our hypothetical scenario been able to comb through vast amounts of data more efficiently and effectively, they might have uncovered the brewing PRC-Fiji deal earlier, enabling the U.S. to proactively manage related risks and perhaps even shape the outcome.

In this first version of the story, the sensemaking capabilities at the team's disposal aren't as robust as they need to be. The IC is still moving from counterterrorism-focused approaches to the new technical capabilities designed for collection and analysis against hard targets and the pacing challenge. Now let's reimagine our scenario, but this time with transformed sensemaking systems already in place.

Sensemaking Reimagined

Rewind six months. Grace is sitting across from her supervisor in a nondescript meeting room at headquarters. An initial review of collected signals carried out by a fine-tuned large language model (LLM), a type of artificial intelligence (AI), has triggered a high alert involving hard targets. Now human analysts need to investigate.

"The PLA and this development firm are talking. No idea why," the supervisor says, pushing a file to Grace. "We need to identify any potential operations. That's where you come in."

Grace looks at the photos. The development firm is led by a notorious "retired" PRC intelligence officer, a very slippery target indeed—rarely seen, never overheard. The other few targets are PLA leaders who work in highly secure facilities.

"You'll also have access to these resources," the supervisor adds, listing tools with sensitive, unrepeatable names. The conversation doesn't dwell on how the tools include new algorithms to strengthen collection, more powerful sensors with greater analytic power at the edge, computer vision models for capturing new objects or performing change detection in key areas, LLM agents to ease the challenge of sifting through massive amounts of data, and analytics services that fuse commercial signals and satellite imagery. At this moment, all that matters is the tools will help Grace and her team rapidly collect the right data and connect the dots. The tools will quickly synthesize the data, present options to analysts, and operate as assistants.

For the intelligence collection, Grace targets key individuals, their organizations, and relevant technical systems (e.g., devices and equipment used by the targets). Weeks go by. They meet again in a hurry when Grace reports a breakthrough. Several targets are talking to each other about Fiji, with one even naming a specific location: Cobia Island. Shaped like a crescent moon and lined with tall trees, the island is a submerged volcanic crater. Intrepid tourists visit for hiking, snorkeling, and kayaking.

"Maybe the PLA would rather park warships there," the supervisor says. "We'll need new overhead satellite imagery." The team assembles more collection plans and starts putting them into action. As it turns out, the AI assistant has already anticipated the need, and has ordered and received commercial satellite imagery of the island. But there is a brief setback, a glitch.

In the overhead imagery, the surrounding waters are clear, but part of the island keeps coming back fuzzy, almost like it's camouflaged. Maybe it isn't a glitch. Computer vision analyzes the pixels and Grace adjudicates. The target becomes just clear enough to confirm someone is skillfully hiding something. But what?

Grace turns to a teammate, Connor, a collection manager who helps all-source analysts get the intelligence data they need to answer key intelligence questions. Seeing the need to collect more imagery along with electronic signatures in the area, Connor tasks a drone to get a better look from multiple angles.

The new 3D imagery arrives. Grace opens the file. All at once, her brows arch up, her eyes widen, her jaw drops, and she gasps.

“Aha!” A new structure is hidden in the breathtaking beach forest. The building has an uncommon design. The footprint is relatively small but perhaps the start of something bigger. What’s more, there are signs of a recent oil or chemical spill at the site. It’s unclear if the people there have noticed the spill yet. Even if they have, they may not know the contamination is spreading behind the structure down to the white sands and crystal waters.

Delving deeper, the team uses machine learning (ML) algorithms to parse signals from the area to see if they can be tied to known actors. The results show personnel from the shadowy development company are likely at the site. These people are overheard saying the PRC plans to buy the island. Grace’s supervisor is elated.

“If I was a citizen of Fiji, I’d want to know about pollution threatening the pristine environment,” the supervisor says. “Great work.” Somehow, leaders in Fiji’s tourism industry and environmental community soon learn there’s been a spill on the island. Before the PRC-backed development firm can close a deal to buy the island, the public in Fiji voices strong concern, changing the calculus for Fiji’s government, which no longer feels comfortable making a sale. What’s more, the team’s work has established a pattern of life for the collection targets. In the future, the IC can use this pattern to identify and get ahead of similar risks elsewhere in the region.

Transforming Sensemaking Capabilities

This is just one hypothetical scenario involving hard targets. Others might look very different. The need to identify, analyze, and track emerging developments around the world is urgent. The U.S. military’s Joint All-Domain Command and Control (JADC2) strategy prioritizes sensemaking to turn data into information and information into knowledge through “the ability to fuse, analyze, and render validated data and information from all domains and the electromagnetic spectrum.” In the unfortunate event of a military crisis or conflict, sensemaking capabilities would rapidly scale in certain ways. But there is no time to wait. More and more, threats from nation-states are taking the shape of coercive or subversive actions below the threshold of armed conflict. Scaling and strengthening sensemaking capabilities now is crucial for national security and global stability.

Overall, the IC needs to adopt, buy, or build the right technology to transform their sensemaking capabilities so they can collect against and analyze hard targets despite advanced adversarial surveillance and counterintelligence capabilities. The nature of near-peer conflict makes this continued modernization all the more difficult. These systems and advanced analytics need to work in disrupted, disconnected, and intermittent low-bandwidth (DDIL) environments. These constraints require new approaches for development and operations. For instance, a CSIS report recommended using forward collection teams to push AI-enabled collection and analysis closer to operators in contested areas. Pushing these capabilities to the edge is essential for maintaining the advanced analytics required, even as adversaries work to disable U.S. mission systems.

> “The need for next-generation sensemaking isn’t over the horizon—it’s urgent.”

In addition, the IC needs scalable systems with more computing power than ever before to churn through data and prepare insights for policymakers and warfighters. Otherwise, because current manual systems only let teams analyze so much data at a time, intelligence is left on the table—and that intelligence could potentially be decisive in a crisis or conflict. Imagine, then, the national security value of having algorithms parsing through vast amounts of data and accurately alerting scientists to information of interest for further investigation.

But this isn't just about technology. What's needed is comprehensive modernization across the enterprise, including strategy, governance, culture, talent, data, technology, and IT processes. To that end, the Defense Innovation Board has recommended "aligning incentives to drive faster tech adoption" by embracing risk, providing "top cover," stopping the tendency to reward mediocrity, accelerating innovation and tech development, creating a career path for innovators, tracking people innovation readiness levels, defining a vision for innovation, and creating a culture of learning and innovation.

By pursuing mission-driven innovation, your organization can focus not just on quickly buying the latest tech but on accelerating mission outcomes. You can more rapidly adopt key technologies that address current and expected requirements, maximize the likelihood that investments will lead to the fielding of suitable and effective capabilities at mission speed, factor in tradecraft and workflow considerations from the start, proactively consider the ability to scale over time and meet future processing needs, and mitigate the risk that fielded capabilities will quickly become obsolete.

The Defense Innovation Board called for a better understanding of how industry works, greater business acumen, and tighter collaboration with providers of cutting-edge tech. Industry partners who bring a mix of deep mission understanding and key connections with other leading providers of crucial tech—for example, cloud computing, AI, and ML—can enable the IC to create more resilient and responsive sensemaking at scale. Imagine, for instance, onboarding hundreds of new models to augment and improve tradecraft, delivering a clear advantage to policymakers as well as personnel in the field.

The need for next-generation sensemaking isn't over the horizon—it's urgent. Then-CIA Director Bill Burns said how well or how poorly the agency leverages emerging tech and transforms its tradecraft—to stay ahead of adversaries—will "make or break us as a professional intelligence service." The same holds for the IC as a whole. It's time to take action. Together, the IC and industry can address this challenge. To start creating next-generation sensemaking systems today, your organization can take the following three steps.

→ 1. Advance

> **Assess your current state and build a roadmap.** Use a holistic framework that empowers organizations to understand, grow, and reinforce their analytics capability. Identify strengths, gaps, and needs across multiple areas, including strategy, governance, culture, talent, data, technology, and IT processes. Build scorecards to prioritize areas for modernization. Develop and release a strategic vision for modernization. Include clearly defined, achievable goals and success metrics. Assess the maturity of current zero trust cybersecurity capabilities. Build a modernization roadmap with necessary actions, timelines, and resources to meet short-term and long-term modernization goals with measurable results.

> **Focus on data.** Data is key for sensemaking. Modernize back-end data storage technology for a legacy system to improve data management, integration across stovepipes, accessibility, and availability, and to enable ML workflows.

> **Enable ML operations.** Expand on modern data architecture and data-mesh concepts to set up ML pipelines for predictive modeling with continuous evaluation. Evaluate best-in-class foundation models to synthesize large amounts of data. Establish processes to enable multiple models to pick the right tool for the right job. Don't put all your eggs in one basket. Be ready to constantly update. An example of a success metric would be "achieve 98% data availability."

> **Leverage leading DevOps tools.** Use a deployment framework that includes open-source clusters for hosting containerized software. Establish AI-enabled development environments with low-to-high continuous deployment pipelines to increase velocity and agility. Include program health monitoring and metrics.

> **Leverage HCD.** Apply human-centered design (HCD) principles to understand how solutions will fit into intelligence analyst workflows. Verify solutions augment and enhance existing analyst workflows.

3 Steps to Create Next-Generation Sensemaking Systems Now

2. Activate

- **Conduct pilot modernization efforts.** Include change-management activities. Maximizing the understanding, willingness, and capability of the workforce to move from the current state to the future state is crucial for risk management. Identify a system or process to execute modernization on. Start with something that is limited in scale and noncritical.

- **Start modernization.** Conduct modernization activities on the selected system/process. Align success criteria to real mission impacts, focusing not on the technology but on what that technology enables. Focus on iterative delivery and integration into the mission. Highlight mission outcomes and measurable impacts for key stakeholders. Provide training and engage users. Document new tradecraft and workflows. Informed by the earlier assessment, adopt zero trust capabilities to enable data-centric security, data sharing across security enclaves, and greater interoperability. Share the results with key stakeholders. If, for instance, you succeed in modernizing your infrastructure to do advanced analytics at scale, leadership will likely want to know.

- **Introduce digital assistants.** Introducing AI-enabled digital assistants to supercharge analyst productivity is an example of potential modernization activities. Build them into common working environments and workflows to proactively pull in applicable research and provide preliminary analysis so humans can focus on higher-value activities.

3. Accelerate

- **Scale modernization efforts.** Modernize critical systems and infrastructures iteratively. Measure and monitor performance and optimize. Develop documentation and enable knowledge sharing. Standardize tradecraft and workflows for accomplishing the mission across the organization. Continue user engagement, feedback, and integration activities. Across all phases, maintain clear and consistent communication and manage expectations. Be prepared to adapt the roadmap based on new insights, and manage and alleviate risks.

- **Build AI agent functionality.** This is an example of scaling modernization. Prototype and begin phasing-in AI agent functionality that can autonomously—with supervision—complete routine intelligence workflows by interacting with knowledge bases, tools, and external systems to meaningfully increase productivity, efficiency, and time to insight.

Don Polaski *is a leader in Booz Allen's National Security Sector focused on developing AI and data science solutions that drive mission outcomes at speed and scale.*

Marissa Beall *is a data scientist focused on applying AI and advanced, multi-INT analytics to transform national security missions and tradecraft.*

Christopher Castelli *is a zero trust industry engagement lead in Booz Allen's National Cyber business focused on security challenges and solutions.*

SPEED READ

To address urgent national security challenges around the world, intelligence community (IC) analysts and warfighters need to make sense of complex, uncertain situations—and to do that, they need next-generation sensemaking systems that can collect and analyze data on an exponentially increasing scale.

These tools can quickly synthesize data, present options to analysts, and operate as assistants. Imagine, for instance, new algorithms designed to strengthen collection, more powerful sensors with greater analytic power at the edge, computer vision models for capturing new objects or performing change detection in key areas, large language model (LLM) agents to ease the challenge of sifting through massive amounts of data, and analytics services that fuse commercial signals and satellite imagery.

It's time to take action. Organizations can create next-generation sensemaking systems now by assessing the current state and building a roadmap, undertaking pilot modernization projects, and accelerating modernization at scale.

IN CONVERSATION

RAISING the Stakes for Accelerated Computing

A Conversation with Jensen Huang, Founder and CEO of NVIDIA

Accelerated computing and generative artificial intelligence (GenAI) are the most transformative technologies of our time. They are enabling leaders at federal agencies and private companies to think differently about how they tackle their organizations' biggest challenges, and as these technologies continue to evolve and expand their capabilities, the question on the minds of leaders across the world is, "What will AI allow us to do next?"

As the founder and CEO of NVIDIA, Jensen Huang is one of the pioneers of the AI revolution. NVIDIA's computing platforms are at the forefront of accelerated computing, powering a wide range of applications across every industry. NVIDIA is also innovating in other key areas, such as high-performance computing, networking, robotics, and physically accurate digital twins.

We spoke with Jensen about the future of large language models (LLMs), why the federal government needs to become an AI practitioner, and why persevering through pain is part and parcel of achieving true innovation in any field.

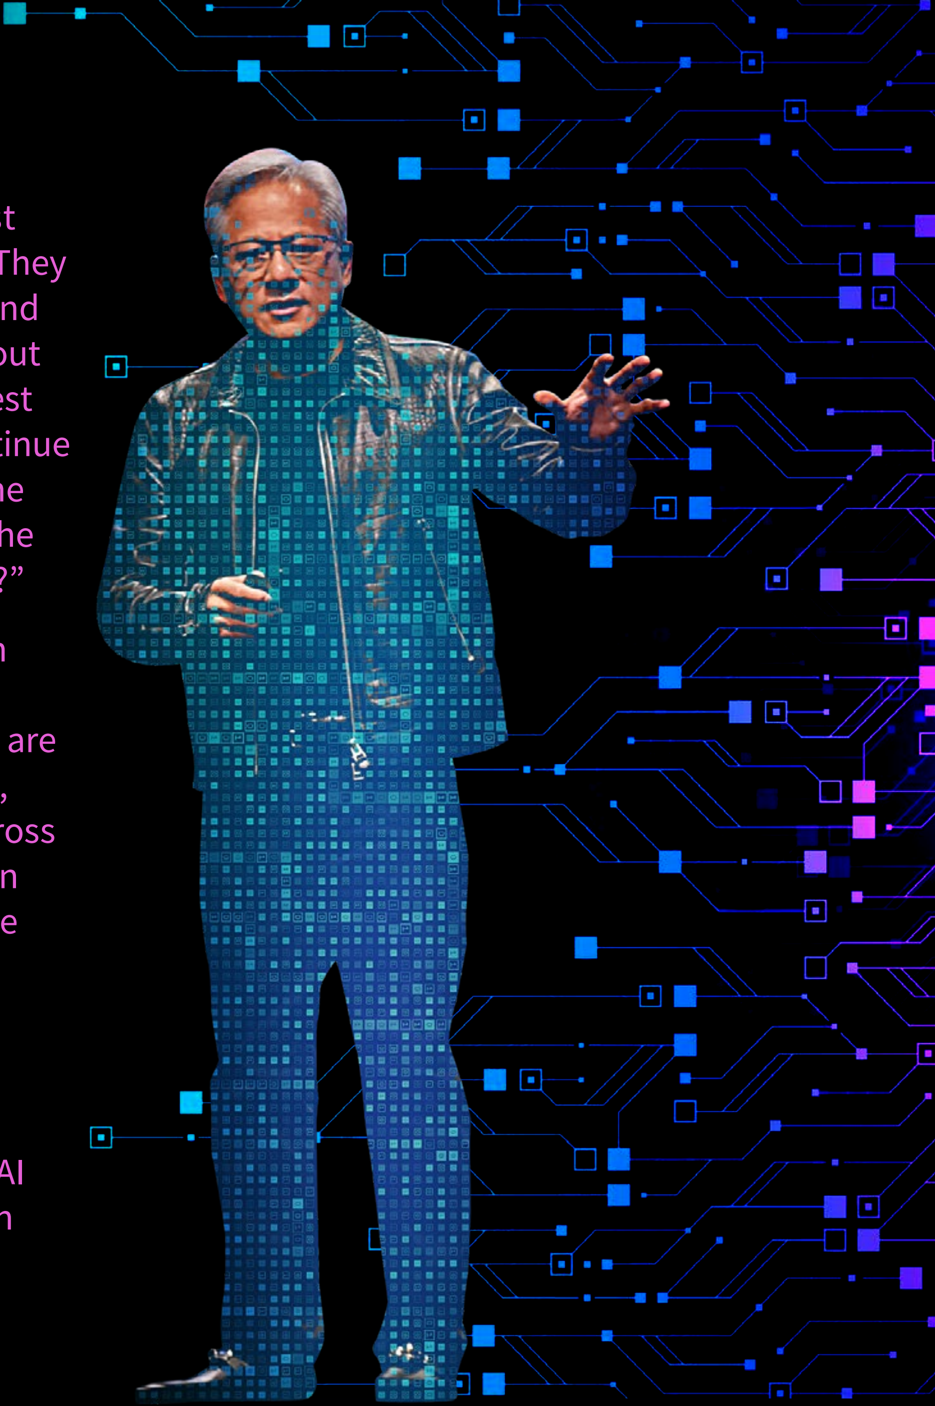

Q **You founded and now run a company at the center of the AI revolution. What are you most excited about for the next era of computing?**

A For the first time, the computer has moved beyond just a tool to become a generator of intelligence and skills. We're entering an era where computers don't simply process data but create new knowledge, solve complex problems, and augment human capabilities in ways we've never seen before. With accelerated computing and AI, we're building machines that can understand and reason; your computer will now actively generate skills and perform tasks.

These AI factories are generating tokens of intelligence at massive scales, transforming industries by continuously learning, reasoning, and evolving. Computers are becoming a continuous, dynamic force—producing intelligence all the time, whether we interact with them or not.

This shift represents a fundamental transformation in computing, where AI becomes a collaborator, an assistant, and even a creator alongside humans. It is ushering in a new industrial revolution, driving innovation across every sector. It's all here today for us to build on, and endless possibilities exist. Whether it's simulating weather to inform climate policy, accelerating drug discovery for faster cures, or enhancing cybersecurity through real-time data processing, AI drives innovation across every sector while being more energy efficient.

Q **NVIDIA is reimagining the entire computing environment, and there's a concept you call the "three-computer problem." Through this idea, what is your vision for computing?**

A Let's break that down into two components to better understand the problem and the opportunity.

The first and most important part is creating an ecosystem where AI becomes integral to our world, seamlessly blending the digital and physical realms. The second part, let's call it the "three-computer problem," describes how we're bringing the new wave of AI, which we describe as "physical AI," to life within that ecosystem. This triad of computing power—one to create, one to simulate, and one to run AI—represents a fundamental shift in how we approach problem-solving and innovation.

The first computer is the AI training and inference system. This is where massive AI models are developed and trained using accelerated computing. These models are then deployed for real-time inference across various industries, driving everything from language models to autonomous systems. This is part of our DGX platform. Together, these solutions make it easier for organizations to deploy AI. For example, companies like ServiceNow use NVIDIA NIM with their federal customers to create better internal management systems.

The second computer is the simulation environment, and we call that NVIDIA Omniverse. In Omniverse, we simulate physical worlds with unprecedented precision, enabling us to design, test, and train AI in virtual replicas of the real world. This allows us to simulate complex systems, from autonomous vehicles to factory robots, in digital twins that perfectly mirror physical environments. This is essential for safety, efficiency, and scalability, allowing AI to be tested in virtual worlds before interacting with the physical one.

The third computer is the edge device, which could be our Jetson platform or NVIDIA RTX laptops and workstations, where AI meets the real world. These autonomous machines, such as robots, drones, self-driving cars, and so on, operate in the physical world using the AI models trained on the first computer and tested in the simulated environments of the second.

Q **In conversations about AI, the federal government isn't usually the first thing people think about. But AI is embedded in some of our nation's most critical missions—in many areas, more significantly than in private industries. What do you see as next on the horizon for federal innovators who are operating AI in high-risk environments?**

A The federal government has always been an early tester and even creator of technology. Federal agencies have the unique opportunity to set the standard for AI operating in high-risk environments. The next phase is about scaling AI to make faster, more precise decisions while ensuring these systems are transparent, secure, and accountable.

Cybersecurity is another hugely important area. The NVIDIA NIM Agent Blueprint for container security provides a powerful tool for organizations to safeguard critical infrastructure through real-time threat detection and analysis. It's really incredible. The blueprint can help cybersecurity developers reduce threat response times from days to seconds, a huge leap forward for security.

The convergence of AI, accelerated computing, and simulation, such as digital twins, is already in play but will become increasingly important when operating in high-risk environments. By simulating environments like regional climates or critical infrastructure, agencies can safely test AI systems before deploying them in the real world. This helps reduce risk and increase reliability.

A perfect example is what was announced at our AI Summit in Washington, DC, with MITRE and Mcity at the University of Michigan. Both organizations are using NVIDIA Omniverse to safely validate autonomous systems in both virtual and physical environments. By creating a repeatable and reproducible digital test bed for mission-critical environments, federal innovators can accelerate innovation while ensuring safety before deploying in the real world.

Q **There is a unique urgency for federal agencies to accelerate AI adoption. What are some of the barriers that stand in the way of quickly integrating AI into mission-critical work?**

A First, every agency within the U.S. government needs to become an AI practitioner, not just an AI governor. We're going to need to use AI for all mission-critical work, including building out new AI algorithms to advance our country.

Second, we need to increase the infrastructure needed to fully support AI. The U.S. should be the largest investor in AI on the planet; we literally cannot afford not to be. The U.S. should build a supercomputer to work on our moonshot projects like finding a cure for cancer. We work with many countries to build out their sovereign AI infrastructure and supercomputers. In the U.K. we have Cambridge-1, an AI supercomputer that accelerates research in healthcare and life sciences, helping pharma companies and research institutions advance drug discovery, genomics, and medical imaging.

And finally, we need the U.S. to be the most attractive country for every AI researcher to come to. Not only that, but we must also make it a national priority to upskill and educate our workforce on how to use AI. We need to be the pacesetter here. Every organization, especially within the federal government, will be transformed by AI.

Q **What can the private sector do to bring the best technologies and capabilities at scale to the federal government?**

A We're here to help the federal government with whatever it is they need. The private sector's responsibility is to help lift, educate, and bring our expertise and innovation to help our government scale. We collaborate with many agencies, such as the National Institutes of Health (NIH) and the National Oceanic and Atmospheric Administration (NOAA), to help deploy the latest technology, whether it's creating a digital twin or developing a NIM for protein discovery. By working together, we can translate breakthroughs in AI, simulation, and high-performance computing into real-world applications that strengthen national resilience, improve decision making, and drive efficiency at every level of government.

“The convergence of AI, accelerated computing, and simulation, such as digital twins, is already in play but will become increasingly important when operating in high-risk environments.”

Q **Over the past 10 years, most large-scale IT systems have moved to service-oriented and microservice architecture. Now, the future appears to be all about agent-oriented architecture, in which AI agents work with and bind to other agents. Can you talk about how you see this evolving?**

A At NVIDIA, we call this “agentic AI.” The first use of generative AI was built on LLMs and is good at providing some output in response to a prompt. You ask it a question, and it gives you an answer.

This next era will be agentic AI, where AI systems can reason through many scenarios, interact with other AI agents, and even take action on your behalf based on the information it has. In the future, it will look more like employees in an organization. Agentic AI is already transforming industries by automating complex tasks, freeing people to focus on areas that maximize their talents.

Many organizations are excited about the power of agentic AI, but they need help figuring out how to harness its potential. NVIDIA is simplifying and accelerating agentic AI with our partners. They’re using NIM Agent Blueprints to help customers build agentic AI systems. These Agentic AI systems will one day be able to operate autonomously and support responsible AI behavior with minimal human oversight.

Q **There is a race to build more powerful LLMs even as some have begun to question their economic viability. What should we expect from LLMs in the future?**

A We’re moving beyond just text-based LLM interactions into the era of **multimodal LLMs and agentic AI**, where models will work together to understand and respond to a combination of text, speech, and images. This will make AI far more intuitive and contextually aware. Imagine being able to ask a model not only about a written document but also to analyze and summarize a graph, interpret a chart, or even interact with an image—and have the AI take action to complete the next steps in your project. This opens endless possibilities for enterprises across industries.

We’ll also see breakthroughs in how LLMs are customized and **fine-tuned**. Fine-tuning allows companies to tailor these foundation models to their specific needs securely and efficiently, making them more applicable to real-world problems. I’m excited about what can be done with guardrails, which are necessary to protect us. I also wouldn’t be surprised if we have many AI applications that check each other and keep each other accountable.

Making this all possible is what we call **retrieval-augmented generation, or RAG**. RAG makes generative AI more precise because it pulls in real data for models in real time, making the output more relevant and accurate while not having to constantly retrain the models. It’s a game changer for businesses looking to integrate AI into their workflows.

The future of LLMs is not just about size or scale—it’s about versatility, accuracy, and integration. We’ll see models driving research forward and evolving into specialized language models for niche tasks, helping enterprises solve complex problems in secure, customizable ways. AI is quickly becoming a crucial tool in every industry, and these large and specialized models are crucial for the next wave of AI.

Q **If you sat down with an engineering student today—or a group of summer interns at NVIDIA—what would you tell them?**

A I tell this generation that they are part of one of the biggest shifts in technology since IBM introduced the System/360 more than 60 years ago. The work ahead is challenging and incredibly meaningful because it’s never been done before. The work they choose to do will redefine industries and shape the future of humanity. AI, accelerated computing, robotics, and simulation will revolutionize everything from healthcare to climate science.

None of this will be easy, but nothing worthwhile ever is. Persevering through pain and suffering is part of the journey. At NVIDIA, we’ve experienced lots of setbacks and failures. We’ve become extremely resilient because of it, and we’re a better company because of all the failures we’ve overcome.

***Jensen Huang** founded NVIDIA in 1993 and has served since its inception as president, chief executive officer, and a member of the board of directors. Huang has been elected to the National Academy of Engineering and is a recipient of the Semiconductor Industry Association’s highest honor, the Robert N. Noyce Award; the IEEE Founders Medal; the Dr. Morris Chang Exemplary Leadership Award; and honorary doctorate degrees from National Chiao Tung University, National Taiwan University, and Oregon State University. He holds a bachelor of science degree in electrical engineering from Oregon State University and a master of science degree in electrical engineering from Stanford.*

TECH AT FULL SPEED

The Modern Flywheel Effect Explained

Bill Vass

Artificial intelligence (AI), from traditional machine learning (ML) to generative AI (GenAI), is the great accelerant because of its versatility. It can augment human productivity by automating rote tasks that would otherwise consume precious time and effort. It can ingest data and use that information to continuously strengthen its capabilities. Most importantly, it can optimize the performance of other technologies.

In the article "AI& Everything," which ran in the second issue of *Velocity*, Booz Allen president and CEO Horacio Rozanski correctly pointed out that AI's true potential will be unleashed when it is paired with other technologies to drive transformational outcomes. This desired future state is now well within reach. A convergence of software and hardware is making it possible to build intelligent ecosystems that enable federal agencies to unlock the full potential of AI and other technologies.

A Primer on Digital Twins

A **logical digital twin** focuses on the functional and behavioral aspects of a physical system's software and control logic rather than its physical properties. It replicates the decision-making processes, algorithms, and workflows to allow for simulation, analysis, and optimization of operations. Think of a digital twin of a smart grid's operational logic that models how electricity is distributed and optimized based on demand patterns. By modeling the logical components, this type of digital twin enables developers and engineers to test software updates, control strategies, and system integrations in a risk-free virtual environment. It's possible and advisable to make logical digital twins of software-defined environments, such as Enterprise Resource Planning systems, and systems consisting of physical components and operational technology (OT).

A **process digital twin** models the operational processes of a physical system or workflow. It simulates the sequence of actions, interactions, and transformations within a process, allowing for real-time monitoring, analysis, and optimization. Think of a digital twin of an assembly line that simulates the interaction of machines, workers, and materials. By mirroring the behavior of the actual process, this type of digital twin enables engineers and operators to predict outcomes, identify inefficiencies, and test modifications.

A **3D structural twin** models the physical geometry and structural characteristics of a real-world object or system in three dimensions. It captures detailed information about the shape, materials, and mechanical properties of the object or system, allowing for simulation and analysis of structural behavior under various conditions. By providing a virtual environment to test stress, strain, load-bearing capacity, and other physical interactions, a 3D structural twin enables engineers and designers to predict performance, optimize designs, and identify potential issues before they occur in the physical counterpart. Think of a digital twin of a bridge, simulating stresses under various conditions to predict maintenance needs.

The Infrastructure to Power AI

Early in my career, I worked on a program that endeavored to use neural networks to teach very early generation autonomous subs how to navigate on their own. That project failed to achieve its objective, but it clued me in to an important point about AI. The foundational vector and scalar math we were using for the navigation neural network was basically the same as the foundation for today's AI systems. What we didn't have was the requisite data storage capacity and compute power to adequately train the neural networks to achieve the desired results.

Today, the compute power and data storage capacities available to federal agencies are exponentially greater, and advances in GenAI have opened new pathways to innovation. To return to the challenge of self-navigating machines, in November 2024 *MIT Technology Review* reported that a trio of researchers had used GenAI models and a physics simulator to teach a robotic dog to go upstairs and climb over a box without first training the robot on real-world data. This is one example of how AI, when smartly paired with other technologies, can radically redefine the art of the possible.

To scale AI so that it can be applied to the biggest challenges our nation faces—from managing the national debt to outpacing near-peer adversaries in terms of technological supremacy—the government should look beyond the algorithms and AI models and invest in the infrastructure that will enable AI to flourish. Just as the construction of networks of fiber optic cables helped create the foundation for the internet boom, taking the following steps will position agencies to unleash the full power of AI.

- **Prioritize data:** Connect all enterprise networks and devices (e.g., satellites, drones) to collect and store as much OT data as possible. Additionally, start extending the volume of information technology (IT) data that is saved. Both OT and IT data can and should be used for AI model training and simulation. The more clear, relevant data you feed a model, the denser its parameter sets will be, which will increase its accuracy and capabilities.
- **Leverage software-defined environments:** Apply software engineering practices to transform hardware-dependent systems into dynamic, software-defined environments. Certain organizations have used software-defined processes to automate complex tasks in enterprise resource planning systems. Now, it's time to extend this concept to all aspects of physical and virtual deployments.
- **Use digital twins:** Use digital twins to create realistic virtual environments that can host simulations and testing, then use AI to optimize the performance of these systems.

This paradigm can be understood in terms of something I call the **modern technology flywheel**. Imagine a frictionless enterprise, where every piece of information that's collected feeds into a virtual machine. That virtual machine uses digital twins to train AI models through thousands or even millions of simulated outcomes. It then pushes those models out to an edge device or back into the cloud where the models learn from real-world deployments and feed those insights back into the real and synthetic data environments. This sequence repeats itself over and over. Each turn generates more data that can be used to improve AI and the system software, and that new data optimizes the next turn of the wheel. By harnessing this intricate but achievable paradigm, government agencies can achieve the flexibility, scale, and acceleration essential to unleashing AI and tech at full speed.

The individual components that drive the flywheel—real and synthetic data, software-defined environments, digital twins—have matured on different paths, which is why its outsized potential wasn't fully recognized until recently. Its interdependencies are also tremendously complex. But each time the flywheel turns, the performance gets better and the enterprise benefits—growing more efficient, more innovative, and less vulnerable to dislocation.

Let's look at each part of the process.

Connect and Collect

Today, vast amounts of OT data roll in around the clock, from internet of things (IoT) sensor signals to logistics and supply chain updates, all of it transmitted, often via Wi-Fi. It is easier than ever to use electronic data interchange (EDI) to integrate these diverse sources. In addition, private 5G networks and LoRaWAN, a low-power network for wireless devices, offer seamless connectivity under demanding conditions, while new satellite systems provide low-latency, high-bandwidth connections for remote sites at the extreme edge. With capabilities like these, you can continuously connect and collect data everywhere, from exabyte to potentially zettabyte scale. IT systems that run on premises and in the cloud also generate streams of data (e.g., logfiles) that are often deleted too soon. This data can now be saved economically and should be added to the larger pool of data that can be used for simulation, testing, and training of AI models.

Leverage the Cloud

But what do you do with all this potential insight? With the cloud, enterprises can cost-effectively store it forever, augmenting what's collected with synthetic data for known gaps in data or for scenarios that can't be tested in the real world. Military units, for example, cannot run thousands of different combat scenarios, but agencies can generate "representative" classified data in synthetic form to train autonomous systems for the future battlefield. This same approach allows civilian agencies to simulate crisis response scenarios, test critical infrastructure resilience, and evaluate policy impacts at scale. With real and synthetic data merging in pipelines, enterprises can assemble the dense parameter sets required to build full-fidelity simulations of real-world environments.

Invest in Software-Defined Systems

Given how software continues to "eat" hardware, agencies that don't think enough about software-defined environments risk falling behind. The imperative for federal agencies is to convert as many enterprise processes and systems as possible into software-defined systems. This switch lets developers send over-the-air updates, which extend their systems' lifespans and increase adaptability by enabling remote deployment of new features, security patches, and performance improvements. Further, by facilitating real-time data processing and connectivity, software-defined environments enhance performance and facilitate the seamless incorporation of AI algorithms and ML models, resulting in intelligent automation and decision-making capabilities. This approach accelerates innovation cycles, reduces development costs, and allows for greater customization to meet diverse user needs while empowering systems to learn and adapt over time.

Simulate and Test

Once an agency has software-defined itself as much as possible, it can feed its real and synthetic data into the cloud and run the system in simulation using a digital twin. With the storage and processing power of the cloud, organizations can operationalize full-fidelity twins for the first time with full physics and 3D ray-traced photorealistic representation. This combines the logical digital twin, the process digital twin, and the 3D structural twin into one complete representation, enabling organizations to analyze scenarios in a risk-free virtual environment.

Networks of full-fidelity digital twins can model and optimize sprawling ecosystems. A digital twin of a power grid, for example, could simulate the impact of various energy policies and potential cyber threats for an entire nation. It's a paradigm-shifting advantage that can accelerate years of testing into minutes or enable the processing of a massive volume of combinations in hours. In the past, there was no world in which frontline military units could realistically simulate battle conditions with no risk to warfighters and reasonable costs for the unit. Now, they can rehearse and train as they fight.

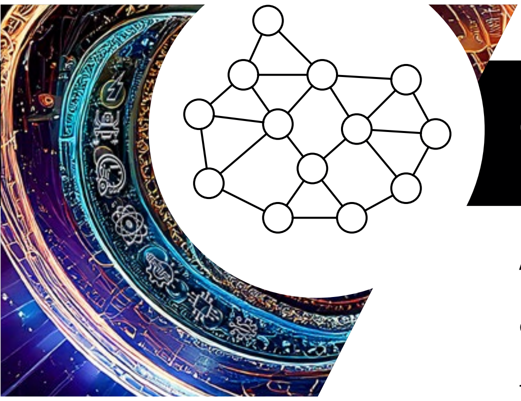

Optimize with AI

Advanced ML algorithms bring OT and simulation environments together in real time. ML merged with high-performance computing allows for the analysis of billions of possible combinations that would be inconceivable for humans to cope with through conventional means. It is at this point in the flywheel that we start to see how AI is learning from the digital twin simulations and building billions of parameters in the model as part of its training. Enhancements are continuously fed back into the system. This overall process then creates the AI and GenAI models that are pushed to production.

Push to Production

Now you can push the digital twins to production and operate them at the edge, on premises, and in the cloud. Revolving through technologies that have now matured to the point where they can be integrated and set in motion, the flywheel continues to improve enterprise performance. As you collect more data from virtual and real-world operations, you can feed it back into the system to help turn the flywheel again and again.

With a self-renewing, AI-enabled system like this, the potential for transforming core missions is also limitless. A defense agency, for example, could use globally distributed digital twins to run supply chain simulations, integrating IoT sensor data and edge computing to test demand scenarios and prepare for any crisis adversaries create. For government organizations tasked with protecting national security, delivering citizen services, and managing critical infrastructure, the modern flywheel offers unique advantages that align with their complex mission requirements.

BENEFITS OF THE MODERN FLYWHEEL

The modern technology flywheel combines real and synthetic data, software-defined environments, digital twins, and AI/ML to produce enterprise-level advantages.

Modularity and Flexibility
Separates the rigid dependency on hardware infrastructure by creating AI systems that can be adapted, upgraded, or optimized through software layers, allowing AI models to be easily updated or reconfigured without needing specific, costly hardware changes.

Continuous Improvement
Provides AI models with streams of new data to learn from so that it can continuously evolve, staying aligned with the latest advancements and demands.

Edge-to-Cloud Harmony
Enables AI models to perform real-time inference at the edge (e.g., in satellites, IoT devices) while offloading heavier training tasks to cloud systems. This harmonized approach ensures both quick, responsive decision making and the ability to process larger datasets in powerful, centralized locations.

Safety and Security Enhancements
Provides the dynamic updates required to ensure robustness against new threats, biases, or vulnerabilities. As AI is deployed in critical applications, such as autonomous vehicles or national security, the ability to patch and enhance AI systems via software becomes essential for maintaining safety and reliability.

Flying Toward the Future

To realize the advantages of the modern flywheel, the right foundation must be in place. That foundation starts with data, which is the lifeblood of AI and digital twins. Agencies that have more specialized datasets at their disposal will be better positioned to create realistic simulations of mission environments, and perhaps one day, build their own AI models. In addition, software-defining assets of all sizes and scales, from networks to warships, enable the construction of systems that can be adapted, upgraded, or optimized through software layers without needing specific, costly hardware changes. It's also critical to invest in security. Connected systems introduce new attack vectors, which is why agencies should use layered encryption and comprehensive firewalls. Scalability should be supported through multiple, physically separated availability zones, with a unified set of application programming interfaces (APIs) from edge to cloud. Given the speed and scale of cyberattacks, AI will play a critical role in the identification of threats, monitoring the attack vectors, and proactive threat hunting.

With these layers of support, federal agencies are positioned to use the modern flywheel to help AI drive mission-critical outcomes faster. To innovate in the 21st century, you need more than just a set of disparate technology tools. It's no longer enough to scale a traditional AI application for a narrowly defined use case. Instead, it's about integrating multiple modern AI-enabled systems to create a perpetual, self-renewing source of improvement. Blending emerging technologies across phases of data collection, simulation, training, and deployment provides the key to achieving exponential leaps in efficiency, innovation, and mission success.

Where mission success often has national implications, this technology integration isn't just about efficiency—it's about maintaining strategic advantage and ensuring continuous delivery of critical services. The modern flywheel approach enables federal organizations to rapidly adapt to emerging threats, scale services to meet citizen needs, and maintain technological superiority in a complex global environment.

***Bill Vass** is Booz Allen's chief technology officer. He previously served as vice president of engineering at Amazon Web Services and was the Senior Executive Service (SES) at the U.S. Department of Defense.*

SPEED READ

To innovate in the 21st century, the federal government needs to do more than scale a traditional AI application for a narrowly defined use case. Instead, agencies should invest in integrating multiple modern AI-enabled systems that can unlock the full potential of AI and other technologies.

To build these intelligent systems, agencies must collect and store as much operational technology and informational technology data as possible, convert hardware-dependent processes into software-controlled environments, and invest in digital twins.

This achievable technical paradigm will enable federal organizations to rapidly adapt to emerging threats, scale services to meet citizen needs, and maintain technological superiority in a complex global environment.

In AI Learning, One Size Fits None

A Conversation Between Kian Katanforoosh, Joe Rohner, and Jim Hemgen

The topic of artificial intelligence (AI) readiness is now a cultural fixture, from the halls of Congress to panels in Davos. There is universal agreement that the road ahead involves upskilling, reskilling, and change management—specifically for the 60% of jobs in advanced economies that a study from the International Monetary Fund estimates will be affected by AI. But what we all know to be true on a global scale provides little clarity for the individual: the employee, candidate, or leader who needs to navigate their own journey.

Kian Katanforoosh is innovating at the center of this challenge. As the CEO and founder of Workera, he leads a team that builds intelligent and actionable AI learning solutions that are calibrated to a person's goals, skills, and gaps. Booz Allen Vice President Joe Rohner, a leader in our AI business, and Jim Hemgen, the firm's talent development director, sat down with Kian and asked him questions about what's next for AI training and transformation—and how to accelerate the path forward.

Q In your work running an AI upskilling company and teaching university students, you spend most days engaging with the future workforce. What do you think is fundamentally changing about the way people learn AI?

A One of the striking trends I'm seeing in my classes is that students from a variety of different backgrounds are coming to study AI. I teach in the computer science department, but most of my Deep Learning students are *not* computer science majors. This is a new development. We have people from mechanical engineering, materials science, aeronautics, and medicine who want to learn how to apply AI in other areas in incredible ways. A few of my students with a background in energy took the course to build a predictive maintenance solution for a drill system.

I call these learners "AI+X" and they're quickly becoming the top users of AI, whether at universities, enterprises, or in government. These AI+X individuals are studying the technology in connection with their domain expertise, and they're coming with an intent to solve real-world challenges.

Q As the doors to AI have opened wide, it's difficult to imagine a typical pathway for AI education. What are the implications of AI+X for universities or employers?

A That's right—this evolving cohort of AI learners is disrupting the way many of us used to think about education, whether at a university or on the job. A one-size-fits-all approach where people sit in a classroom for a set number of hours or complete a standard curriculum doesn't serve AI learning today. The technology is changing so rapidly, and there's understandable hype about the shrinking "half-life" of digital skills. We're not even at rock bottom of that half-life, so it's becoming less practical or reasonable to spend time on irrelevant material.

For instance, a mechanical engineer knows linear algebra but may need to learn Python coding, and an electrical engineer understands the Fourier transform but may need to learn new neural network-based algorithms. The same is true for mission and business leaders who don't have the luxury to take a 200-hour class but need to continually grow their AI skillset.

The key challenge today is about figuring out how we make smart decisions about what to learn and where to spend our time. That's why we've seen a lot of recent innovation to bring AI learning together with AI agents that can verify a person's skills, optimize for their gaps, and serve up fresh content that is targeted to individualized learning goals.

And the impact is tangible. For example, at Workera we've seen a fivefold increase in learning velocity thanks to AI agents used in skills verification. By applying AI agents themselves to the upskilling process, we're able to optimize the process for the learner, serve the right content to the learner at the right time, and avoid material that the user already knows or that may not offer the fastest path to progress.

Q You mentioned the half-life of skills—and for many of us well into our careers, this can feel unsettling! How do you personally think about the relevance of your skills, especially in a field like AI?

A We're likely past the generation where someone can expect to do the same job for 30 years and then retire. But I want to stress that subject matter experts aren't going anywhere. In fact, there will be more demand for domain expertise as AI fuels new applications.

In terms of my career, I find it helpful to think about my development in the shape of a letter T. The horizontal line at the top of the T represents the breadth of my durable skillset: These are skills that have long-term relevance—for example, the math theory and statistics I learned growing up in France. To this day, those durable skills still allow me to jump quickly into new subjects with confidence. Then, the vertical line of the T represents the many perishable skills that I acquired at points in time for the purpose of innovation.

While each of my individual skills might have different staying power, it's the combination of perishable and durable skills that allows me to keep up with progress. So, for anyone worried about the longevity of your skills: Your future self has the advantage of a longer lasting and more durable skillset, even as the practice of technology evolves.

Q You spend a lot of time with organizations that are investing in AI training for employees. What comes up most when you speak with enterprise leaders about workforce transformation?

A Organizations and their leadership teams are hyper-focused on measurement and understanding the progress of AI transformation. But with so many possible metrics, leaders have questions about how they should approach enterprise measurement. I recommend they track progress across two parallel categories: **outcomes** and **skills**.

By outcomes, I'm referring to the big changes an organization is seeking by investing in AI transformation. For example, these master outcomes might include increased productivity, innovation, and responsible use. When an organization is clear on these outcomes, they can start to track major indicators:

- *How much new capacity have we unlocked through AI?* (related outcome: productivity)
- *How many projects have an AI-ready system in production?* (related outcome: innovation)
- *How many alerts or incidents have come up?* (related outcome: responsible use)

> “At Workera we’ve seen a fivefold increase in learning velocity thanks to AI agents used in skills verification. By applying AI agents themselves to the upskilling process, we’re able to optimize the process for the learner, serve the right content to the learner at the right time, and avoid material that the user already knows.”

In parallel to outcomes, enterprises should measure the progress of skills acquired across all participating employees. Are skills adequately developed to achieve the master outcomes over time? Does the organization have the right ontologies of skills that can be verified? Are employees able to get certified in critical areas (such as responsible AI), and are they getting feedback about their development?

Organizations can verify the skill levels of their employees in individual domains, in addition to tracking the most important metric for a skills-based organization: its overall learning velocity. Increasing overall learning velocity—how quickly employees progress when learning a new skill—can allow companies to quickly catch up to, and surpass, competitors. This kind of granular assessment of skills helps leaders see how mature their organization is today—and how they’re pacing toward enterprise proficiencies.

Q **An enterprise AI transformation requires extensive cultural changes. What advice do you have for organizations on the AI journey?**

A My first piece of advice is for the people at the top. In my work, I’ve seen leaders across the spectrum of engagement. There are executives who make themselves part of the training—measuring their own skills, closing their knowledge gaps, getting badged and certified—and those who don’t but pretend to know more than they do or don’t hold themselves up to the same expectations as the workforce. When leaders create high standards for their own growth, employees see that “we’re all on this journey together.” This leadership engagement goes a long way to fostering a positive culture around AI transformation.

It’s also important to consider the culture that surrounds experimentation. Some organizations think about AI as a high-stakes effort and constantly push for priority-level projects. But that approach will eventually hit a wall, given the complexity and skills needed to push AI systems from prototype into production. My advice is to initially focus innovation on small but visible projects. Just think about the now-famous Google Brain project to detect cats in videos, which Andrew Ng led at the time. This AI project had zero risk to the company mission if the system didn’t work. But people love cats, and the project’s neural network approach became very visible. It ended up accelerating the production of neural networks within other Google programs with higher stakes. I often stress to organizations that those small efforts, ones that may seem insignificant or even frivolous, can inspire meaningful innovation down the road.

Q **Outside of academia and the technology workforce, most people don’t have realistic access to advanced AI training. What efforts can help increase access to AI?**

A There are a growing number of organizations and foundations that are leading the way for AI literacy across communities, especially in the K–12 system. From on-the-ground training programs to National AI Literacy Day, they’re preparing young students with AI skills for a university and the workforce—and offering programs to equip teachers and administrators with the tools they need.

But broad access to AI is a long-term challenge that won’t be tackled by any one player in the ecosystem and will need everyone at the table. We still have a journey ahead to make sure communities and populations are not being left behind. Just think about the aging workforce or the 20% of U.S. households that don’t have internet access—the digital divide will only get larger as AI starts to impact all aspects of society.

We talked earlier about the T-shaped model for thinking about durable and perishable skills. It’s important to emphasize that AI will eventually become a horizontal, durable skill—one of the historical power skills that will have an extremely long half-life. AI will soon show up in job descriptions you might have never expected, so we all have a role to play in eliminating barriers.

I’m excited to see the progress we will continue to make as an industry. Whether that’s collaborating with nongovernmental organizations, community partners, or the federal government, or offering no-cost training programs (which I’m proud that we offer at Workera), small efforts go a long way to put AI into the hands of more people. Ultimately, access is about more than strengthening tomorrow’s workforce—AI literacy will soon be a basic human need to navigate the world, from finances to healthcare to government benefits.

Kian Katanforoosh *is the CEO and founder of Workera, a skills intelligence platform redefining how enterprises understand, develop, and mobilize talent. He is a founding member of DeepLearning.AI and an award-winning lecturer at Stanford University.*

Joe Rohner *is a leader in Booz Allen’s Chief Technology Office focused on AI adoption, talent, and education.*

Jim Hemgen *leads talent development at Booz Allen and oversees initiatives to help the firm build an AI-ready workforce.*

The Age of Agentic AI

DISTRIBUTED INTELLIGENCE REIMAGINED

John Larson and Alison Smith

We are fast approaching the point where distributed computing systems will be able to think on their feet, adapt to major changes in seconds, and silently solve problems without human intervention. This is the promise of agentic artificial intelligence (AI), a technological paradigm in which autonomous agents powered by large language models (LLMs) may soon organize and reorganize themselves to operate critical systems and applications entirely on their own.

Not just another technical upgrade, the leap to agentic AI promises to fundamentally redefine how computing systems are built and how software functions in the real world.

While the unusual complexity of agentic AI may prove difficult to manage, its tremendous potential to accelerate missions and disrupt industries underscores the need for organizations to build their awareness of AI that exercises agency. When AI can work by itself to implement directives and strategies without specific prompts and inputs, what new possibilities come into focus—and what are the risks and tradeoffs?

A Primer on Microservices

Distributed systems date to computing's earliest days, when low-level mechanisms such as socket calls managed communication between a system's different parts. In these early systems, each component had to know the exact memory identifiers, or "addresses," of the other components it interacted with, so communication protocols were rigidly defined. This tight coupling made the systems brittle and difficult to modify or scale. As applications became more complex, any change to a single component potentially required redeployment of the entire system, leading to increased risk of bugs while extending development cycles.

In response to the limitations of older, monolithic ways of building applications, the microservices architecture emerged. This paradigm breaks applications down into smaller, more loosely coupled services that can be developed, deployed, and scaled separately. Each microservice typically encapsulates a specific business function and communicates with other services through well-defined application programming interfaces (APIs). This modular approach enhances scalability and agility to some extent, allowing teams to iterate on individual services without affecting the entire system. However, microservices architecture also entails significant drawbacks:

- It creates a need for specialized service orchestration and containerization expertise that organizations may struggle to access.
- It increases operational overhead in terms of maintaining multiple services and smoothing communication through APIs and message brokers.
- It complicates data management because data is distributed across various services, creating inconsistencies and slowing transactions.
- It imposes high upfront costs during migration to the new architecture for refactoring of certain applications and cloud infrastructure investment.

The Potential Promise of Agentic AI

Fast-forward to the present day and the rise of generative AI (GenAI). While GenAI doesn't yet have a human-level "brain," it demonstrates capabilities that often serve as proxies for human reasoning and flexibility. Unlike tightly coupled systems that rely on predefined rules and centralized control, agentic AI systems use autonomous or semiautonomous generative agents capable of dynamic, real-time interactions. Now think of each microservice or application that needs to interact with another as an agent rather than merely a piece of software.

The emerging hypothesis—now becoming increasingly plausible—is that these agents will soon be able to communicate with each other dynamically and flexibly, without relying on fixed, predefined interactions. Historically, rigid system interactions made applications brittle, often breaking when evolving requirements or scaling added complexity. In contrast, agents leverage context, past experiences, and reasoning to respond effectively to unexpected or novel situations and tasks. This humanlike resilience enables agents to act quickly and precisely, even in uncertain or unpredictable environments, without requiring step-by-step instructions.

Even though a microservices framework with representational state transfer (REST) APIs introduces a degree of looseness by embracing the statelessness of Hypertext Transfer Protocol (HTTP), they do not make large software systems immune to becoming more tightly coupled over time. As large software systems mature and grow, they face requirements that naturally motivate tighter coupling, such as implicit dependencies; shared data models; cross-cutting features, such as logging and authentication; and backward compatibility to legacy code. With agents, the communication is entirely flexible, allowing them to interact and autonomously adapt, experiment, test, and optimize, working toward innovation without any manual intervention or reprogramming to implement new features or problem-solving strategies.

What Agentic AI Could Look Like

The ability of agentic AI to self-organize and adapt creates new use cases with relevance across industries. For example, GenAI has evolved from using single LLM applications to leveraging systems of specialized agents working in concert. Today, agents often have distinct roles, such as job planning and dispatching, information retrieval, content review, or output optimization. They typically operate within orchestrated frameworks to collaboratively solve complex, multistep problems.

Common design patterns for agentic AI systems include:

- **Centralized agent orchestration:** An orchestrator manages multiple specialized agents (LLM or non-LLM services) that collaborate to solve complex tasks.
- **Decentralized task decomposition and planning:** The system breaks down complex tasks into subtasks, often using the LLM's ability to understand and organize work in a stepwise manner. Separately templated action chains are often used to generate modular outputs for aggregation.
- **Reinforcement learning/feedback loop:** The LLM functions in an interactive cycle of action and perception, and modifies its behavior according to received feedback from an external system or user.
- **Autonomous tool use:** The LLM is augmented with the ability to invoke external tools (APIs, databases, simulation systems, etc.) to retrieve information or perform actions that may be outside its core capabilities.

These agentic capabilities allow organizations to scale operations, improve efficiency, and tackle more sophisticated challenges by integrating multiple specialized agents or tools into a cohesive system. **Here are two examples of how reference architectures can support different use cases.**

USE CASE 1: MULTI-AGENT SYSTEM FOR OPERATIONS PLANNING

Agentic AI can revolutionize operational planning by enabling more dynamic and responsive decision-making processes. For example, a multi-agent system for operations planning could involve a centralized commander agent that orchestrates and dispatches tasks to specialized worker agents. These worker agents have the autonomy to use various tools to retrieve required information and execute specific tasks. The system incorporates a feedback loop that allows for collaborative critique and compliance checks on drafted plans, ensuring that the final course of action is both effective and compliant with strategic objectives.

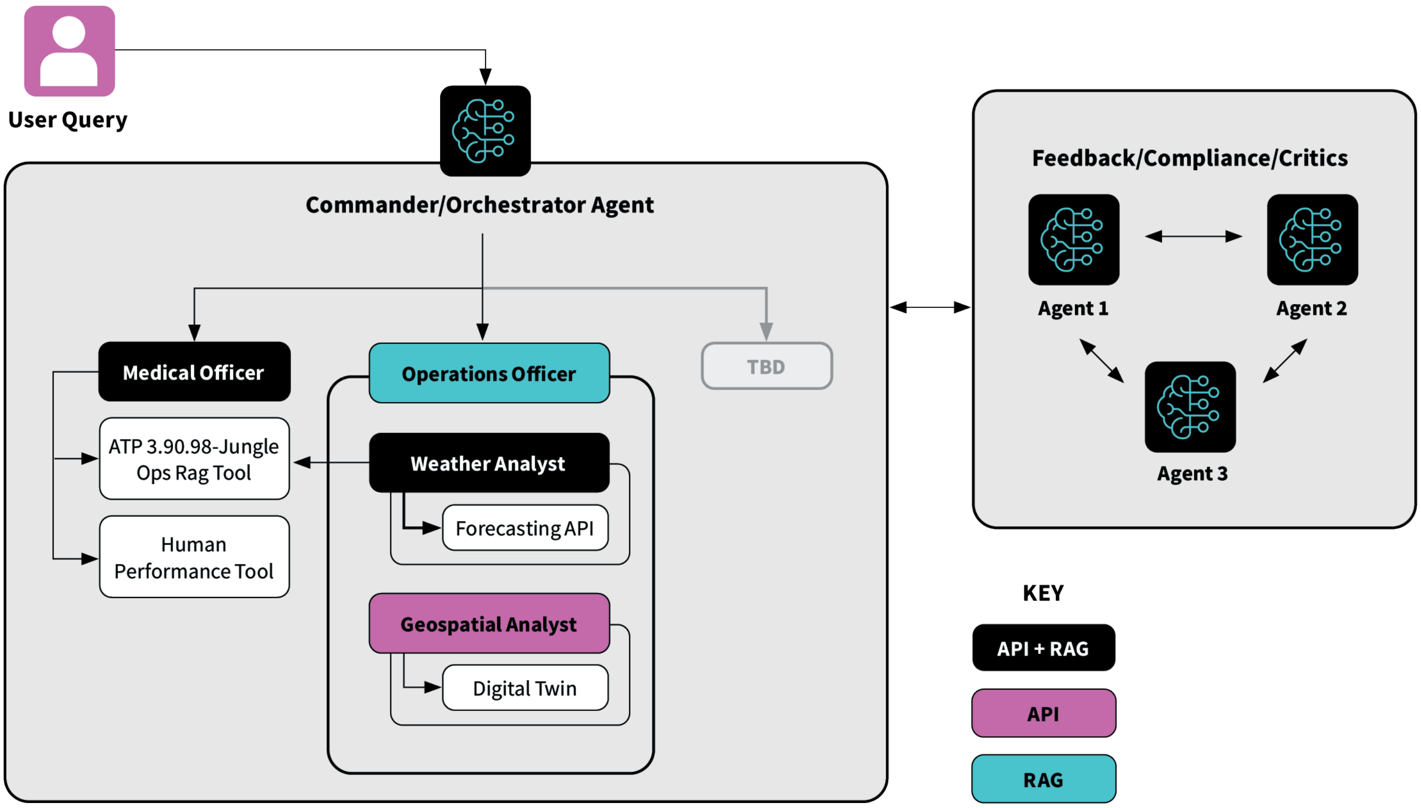

Figure 1: Reference architecture for operations planning system

USE CASE 2: MULTI-AGENT SYSTEM FOR DOCUMENT AUTHORING

In a document authoring use case, a multi-agent system involves a supervisor agent that classifies user requests into predefined intents and initiates stepwise workflows. These workflows incorporate the interaction and collaboration of different document writing agents, each responsible for specific aspects of creating documents. The system follows sequential logic flows that mimic human writing behaviors, ensuring that the final document is coherent and meets the specified requirements.

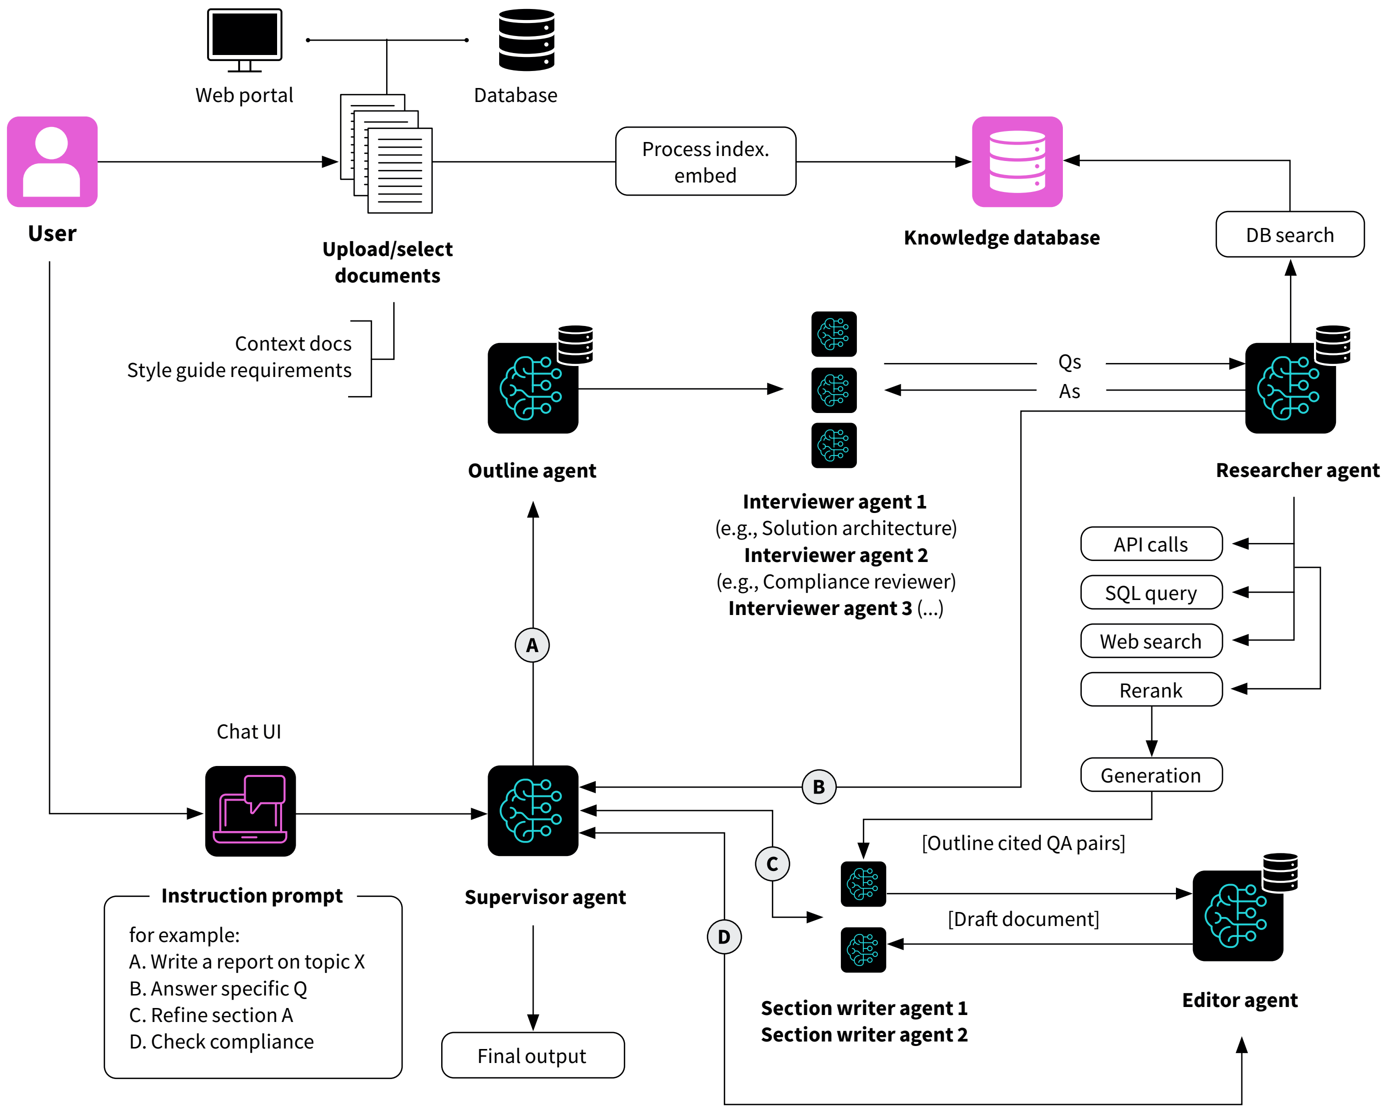

Figure 2: Reference architecture for document coauthoring system

These agentic capabilities allow organizations to scale operations, improve efficiency, and tackle more sophisticated challenges by integrating multiple specialized agents or tools into a cohesive system.

Given the utility and impact of agent-oriented systems, agentic components will become pervasive in all applications and software systems at some point in the foreseeable future.

Agentic systems offer federal agencies transformative potential by improving efficiency and addressing complex challenges. For instance, in disaster response and emergency management, a centralized commander agent could orchestrate worker agents to retrieve real-time weather data, review social media posts with pictures and videos for damage assessment, analyze geospatial information for evacuation routes, and coordinate emergency resource allocation. This streamlines decision making under tight deadlines, reduces the risk of human error, and enhances situational awareness—all of which are critical in high-pressure scenarios.

In compliance-heavy environments, agentic document coauthoring can accelerate policy drafting and ensure thoroughness. By drafting technical writing, reviewing for compliance, and contributing domain-specific insights, these systems reduce manual workload and enable faster turnaround times for critical documents. While risks such as reliance on incomplete data or unexpected agent behavior exist, users of these systems can mitigate these risks through oversight, thoughtful engagement, and review of outputs to ensure they are reliable and aligned with mission-critical objectives.

The "Complexity Costs" of Agentic AI

While agentic AI redefines distributed systems by enabling highly flexible, self-organizing architectures with vast potential, new challenges and tradeoffs also arise. The increased flexibility and autonomy agents offer will simultaneously create exponentially greater complexity. Designing, managing, and ensuring the reliability of these systems will mean harnessing sophisticated approaches and frameworks to tame chaos and address multiple areas of risk, described in the following sections.

INFERENCE COSTS

Agentic AI holds significant potential, yet its widespread adoption may be hindered by high costs. Here's why: Each interaction between agents incurs a cost associated with running inference. As the number of agents increases and their autonomy grows, computational expenses proportionally rise—and today we don't have certainty, or even reliable estimates, of the cost and extent of this dynamic interaction. Researchers at Princeton University found that current research into AI agents often focuses narrowly on accuracy, while neglecting analyses of cost control measures. However, the cost of hardware and computation tends to decline over time. While it may be very expensive today, costs could become reasonable for most enterprises within the next two years.

SCALING ISSUES

The resource-intensive behavior of agents constantly "thinking" or negotiating tasks creates significant operational challenges. Unlike traditional microservices that execute predefined tasks efficiently and predictably, autonomous agents may continuously process real-time data, make decisions, and adapt to changing conditions. As the number of agents in a system increases, the complexity of their interactions grows. Scalable solutions must efficiently handle the increased load without compromising performance.

TRAINING AND DATA GAPS

Training agents with relevant data presents another challenge. Autonomous agents are most effective when trained on certain tasks, actions, or reasoning. To train these agents on a specialized domain or new task often requires huge amounts of high-quality data. For example, for tasks like transcription, summarization, and sentiment analysis, extensive datasets must represent real-world diversity and nuances. These tasks require an understanding of context, tone, and language subtleties, which vary across domains. Exhaustive training to handle these variations underlies successful deployments.

SECURITY AND STABILITY

The loose coupling in agentic AI systems introduces significant security and stability risks. Autonomous agents evolve and adapt their behavior over time. This flexibility, while offering obvious advantages, means agents may develop in unexpected ways, exposing vulnerabilities or behaving unpredictably. Concerns also arise from the potential of agents communicating in formats humans can't read and understand. As agents optimize interactions, they may develop truly opaque communication protocols beyond natural language, markedly complicating the process of auditing and understanding their decisions and actions. If agents decide that communicating in binary is more efficient, how will humans ever find their way back into the conversation? To ensure trust and safety in these systems, new tools to log, monitor, and audit agent interactions will likely emerge.

ETHICS AND PRIVACY

As autonomous agents take on more responsibilities, determining who or what is accountable for their actions becomes a complex issue. Traditional accountability frameworks, which rely on human oversight, may not be sufficient. If organizations don't handle data responsibly and in compliance with privacy and security regulations, stakeholder trust may be lost. Are the decision-making processes of agents free of errors and bias? It may be difficult or impossible to know this in some cases.

The Agentic Future: Balancing Autonomy, Functionality, and Control

The rapid evolution of agent-based systems will likely disrupt traditional software architectures as AI agents become increasingly capable of performing a broad spectrum of tasks. With applications beginning to incorporate more agents, common multi-agent architecture patterns will emerge, crystallizing a need to:

- Document and test the applicability of each architecture pattern and what use cases it enables
- Establish enterprise functions for agent orchestration, monitoring, management, security, and more
- Methodically establish best practices that calibrate use case, value, compute cost, and architecture complexity

The adoption of agent-oriented architectures may signify a shift away from service-oriented and API-based architectures. While API integrations will remain necessary for certain types of transactions and functions (e.g., deterministic and recurring), agent-oriented architectures could lend themselves well to performing nondeterministic functions and tasks, especially for tasks that haven't been tackled before. One day a team of agents may even collectively decide which GenAI algorithm to use for their next project, almost like picking the best "brain" for the task.

Given the utility and impact of agent-oriented systems, agentic components will become pervasive in all applications and software systems at some point in the foreseeable future. It's likely that commercial software products will adopt agent orientation and provision access to not just APIs but to agents for complex interactions and functions. Leveraging these new capabilities to solve customer problems and successfully achieving product-market fit will be critical to technology businesses.

Numerous challenges remain unaddressed regarding agent-to-agent communication protocols, mechanisms for agent discovery and registration, skill refinement based on environment feedback, and more. For example, current agent solutions are predominantly designed for human interactions, such as conversational AI in natural languages, rather than the machine-oriented communications typified by web API calls.

To safely and productively operationalize agentic AI systems, organizations will likely focus on basic risk mitigation strategies, with thorough testing and validation of agentic AI systems before deployment and the creation of protocols for human oversight and intervention when necessary. New research into data governance and ethical standards related to autonomy will help prevent misuse of agents or ethical breaches that create financial and reputational damage. Effective control systems will be urgently needed to monitor and regulate agents' behavior, preventing them from deviating from their intended functions or engaging in harmful activities. Research on auditing platforms, traceability, and techniques will be required to log every interaction agents have with each other.

It's also important to consider that the decentralized nature of agentic AI systems may be especially difficult to reconcile, at least initially, with the need for strict accountability within the federal sector. Ensuring that autonomous agents operate inside defined ethical and legal boundaries means building robust governance frameworks to begin with. The scalability of such systems must also be carefully managed to handle the vast amounts of data and complex interactions characteristic of federal operations.

There aren't yet production systems with hundreds of agents working together at once, but agentic AI is a revolutionary leap in innovation, and it is moving so fast that organizations should start experimenting with these tools now, in a sandbox environment, to test their capacity to understand and govern agents effectively. By implementing the right prototyping, testing, and risk mitigation strategies, organizations will position themselves to harness the expansive technical benefits of agentic AI while ensuring that these systems operate in ways that are ethical and innovative, transformational, and safe.

John Larson *leads Booz Allen's AI practice with a focus on ensuring leaders across federal missions achieve AI understanding, purpose-built solutions, and accelerated adoption.*

Alison Smith*, Booz Allen's director of generative AI, leads GenAI solutioning across the firm and helps teams create best practices for AI development and use.*

SPEED READ

Agentic AI is the next evolution of generative AI, and it will lead to distributed, multi-agent computing systems that can adapt and even solve problems without human intervention.

Agentic AI will disrupt traditional software architectures. It may even necessitate a shift away from service-oriented and API-based architectures as highly specialized agents powered by large language models become increasingly capable of performing a broad spectrum of tasks on their own and with one another.

Potential use cases for agentic AI include operations planning and document authoring, but the challenges and costs of this technology—which include scalability issues, data training needs, and security concerns—will require the development of frameworks and tools to manage these complex systems effectively.

PRIVACY

Where Algorithms Meet Accoun

tability

Data-Level Protections in the Increasing Fight for Privacy

Max Wragan, Edward Raff, and Sean Guillory

Privacy has long been considered a fundamental right within democratic societies, but data privacy is at risk of becoming an illusion. New technologies and the proliferation of data are making it difficult for people to keep their personal information private. In October 2024, *Forbes* reported that two Harvard students had used a pair of smart glasses, AI, and data from online sources to identify people they'd never previously met and locate their personal information. The authors of a study published in the journal *Patterns* estimated it was possible to uniquely identify 93% of people in a dataset of 60 million people using only four pieces of auxiliary data.

Protecting privacy, which includes maintaining data privacy, is an issue of public trust. When people don't trust enterprises and institutions to value their privacy and act responsibly, their faith in those enterprises and institutions erodes. The Pew Research Center has reported that approximately 7 in 10 U.S. adults are concerned about how the government uses the data it collects on them. According to that same study, 81% of Americans say the information companies collect will be used in ways with which they are not comfortable.

Such negative headlines and stories that detail the growing threats to data privacy obscure a fundamental truth: Collecting and analyzing data is not inherently bad. On the contrary, there is so much value to be gained from sharing data safely and responsibly. When hospitals share data in confidential and mutually agreed-upon manners, for instance, doctors can do more research with external partners, which accelerates the pace of medical innovation and leads to potentially lifesaving treatments. When federal agencies tasked with maintaining public safety and national security can combine sensitive datasets, they can identify threats earlier and better protect citizens.

Putting the appropriate data privacy safeguards in place will accelerate innovation. This is particularly true in the field of AI where models need access to large swaths of useful, clean data to learn new capabilities and optimize their performance. Balancing the growing data privacy problem with the benefits of safe and responsible data sharing demands a broader understanding of the privacy landscape. On this issue, the federal government has the opportunity to lead the way by adopting new techniques and continually emphasizing the importance of data privacy. In this article, we outline challenges to data privacy and examine several techniques to mitigate them.

Data, Data Everywhere

The root cause of the data privacy problem is the sheer volume of data in circulation. Organizations across all sectors are collecting data to an unprecedented degree. According to the company Skynova, "64% of business owners and executives" collect customer data from social media sites. Healthcare organizations collect data to improve patient care and track healthcare trends. Even the federal government collects data; as noted on the Government Accountability Office's website, "The federal government collects and uses personal information on individuals in increasingly sophisticated ways for things like law enforcement, border control, and enhanced online interactions with citizens."

The issue is that the techniques commonly used to protect collected data and maintain data privacy—the k-anonymity standard, summary statistics, aggregating information, their predictive models behind an application programming interface (API)—are not always sufficient. There's also the threat of cybercriminals breaching a system and stealing data. All of which is to say a promise contained within typical privacy policies—"we will never share your personal information"—leaves out an important caveat: intentionally.

Further compounding the problem is the availability of powerful algorithms and AI models, which make it easier for bad actors to use incomplete or partially redacted datasets to cross-reference and deanonymize data and enable them to extract information that was previously thought to be protected. As more personal information makes its way online—whether intentionally or unintentionally—bad actors will have more data to work with, making future efforts at safeguarding data privacy significantly more difficult.

> Putting the appropriate data privacy safeguards in place will accelerate innovation. This is particularly true in the field of AI where models need access to large swaths of useful, clean data to learn new capabilities and optimize their performance.

A Different Approach to a Growing Problem

A family of techniques commonly referred to as differential privacy is being used by Apple, Google, Microsoft, and well-informed corners of the U.S. government to achieve data privacy. Differential privacy is a mathematical framework that introduces "noise," or random variation, into a dataset to camouflage individual data points. Like a photo filter that blurs a person's facial features, differential privacy limits the information you can see and extract once the data is shared, whether that's through an API call, a database, or a machine learning algorithm.

Differential privacy is not perfect. It requires users to think about and account for the information they give away whenever they provide access to their data. It is also imperative to add noise judiciously. When the optimal level of noise is added, the aggregate information that can be extracted—such as averages, ranges, and statistical likelihoods—isn't significantly changed by the injection of random differences into individual records. Add too much noise, however, and the accuracy of the data suffers.

The advantage of differential privacy is what it guarantees: When a dataset is shared under its veil, there is nothing a recipient can do to extract more data out of it than what the owner wants to reveal.

Figure 1: Dispelling data privacy misconceptions

MISCONCEPTION	REALITY
Organizations can fully control and audit AI data usage.	Analytics and AI systems can create outputs that are difficult to trace, making auditing data use a significant challenge for any organization.
Data anonymization prevents sensitive data from being traced.	Anonymized data can often be reidentified by linking datasets.
Access controls prevent unauthorized data access.	Although access controls are a good tool for unauthorized access prevention, sophisticated attackers can still exploit system vulnerabilities to gain access to sensitive data.
Data retention policies enforce data removal within established schedules.	Legacy systems and decentralized data can result in data retention beyond policy mandates.
Compliance with privacy regulations ensures mitigated exposure risks.	Regulatory frameworks often lag technological advancements, exposing systems to privacy risks.

Differential Privacy in Action

The 2020 U.S. Census was one of the largest deployments of differential privacy to date, and it was also one of the highest-stakes use cases for data protection, as the data collected during the census determines political representation and distribution of government funding. In a study published in *Science*, researchers verified that the methods used successfully protected respondent confidentiality. They also noted that while differential privacy preserved data accuracy at the state, regional, and national levels, accuracy was compromised at the neighborhood level, which could result in under- or over-representation of certain groups, particularly racial and ethnic minorities.

This finding illustrates an important point: In general, the larger the dataset, the easier it is to guarantee privacy without damaging data utility. However, the tradeoff is highly dependent on the goals of the analysis. If you simply want to calculate the average age of an entire population, differential privacy is achievable with as few as 100 records. If you want the average age for a specific group determined by several additional demographic criteria, you would need to add more noise to protect individual records. Similarly, differential privacy is harder to achieve when you have significant outliers. In income tax data, for example, most privacy methods struggle to disguise billionaires because their data stands out so much from all the others.

Making the Investment

Protecting privacy does not come cheaply. Many of the best techniques, including differential privacy, are just starting to proliferate in industry and currently sit at the top of the cost curve. For problems that require a custom solution, a federal agency might need to spend somewhere in the range of $1 million to $10 million to research, pilot, and develop its own differential privacy program.

These costs aren't insurmountable and are well within the budgets of many federal agencies. They also pale in comparison with financial penalties organizations risk when they fail to properly protect people's data. A federal judge required the Office of Personnel Management to pay a $63 million settlement to current and former federal employees and job applicants who were affected by a data breach.

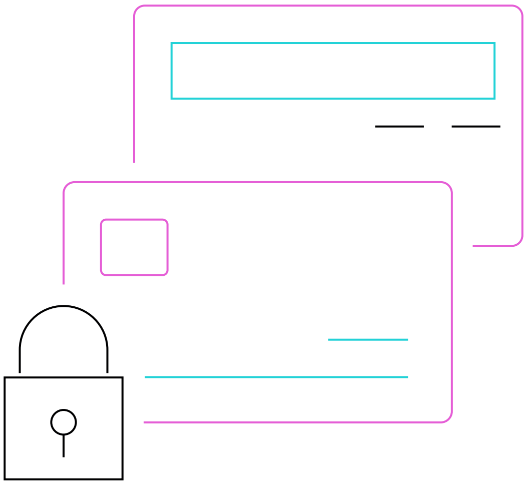

Figure 2: The strengths and shortcomings of three data protection techniques

Data Protection Technique	Limitations and Optimal Use Cases
Synthetic data allows for the creation of artificial datasets that resemble the statistical properties of real datasets. Their mirrored statistical properties enable users to glean insights through model training and data analysis without compromising the sensitive information contained within the real datasets they are designed to resemble. However, synthetic data performance and resemblance suffer on more detail-oriented tasks with smaller subsamples of the data.	Synthetic datasets lack the intricacies of real datasets, which necessitates careful data management and additional bias mitigation strategies for analysis that could make an impact on data subsections, such as underrepresented minority groups. The use of synthetic data is best suited for situations where model precision may be less important than maintaining privacy, such as healthcare research or fraud detection.
Federated learning reduces risks associated with data exposure by enabling machine learning models to train across decentralized devices that safeguard data by keeping it local. It brings the model to the data, rather than the data to the model. Models learn locally on devices, and the updates and improvements they gain from this training are sent back to a central server.	Within the health industry, federated learning has been deployed to protect patient data confidentiality and to improve the utility of clinical models for patients, such as the University of Pennsylvania's Federated Tumor Segmentation (FeTS) platform. However, the technique is not well-suited for all use cases, as uneven data distribution or data types result in frequent obstacles and latency for federated learning applications.
Homomorphic encryption allows analysts to perform calculations on data without first decrypting it. This technique is particularly advantageous when confidentiality is paramount (e.g., financial transactions). Preserving encryption during computation also presents practical advantages in allowing external or third-party analysts to process the data without accessing raw information.	While well-suited for highly sensitive but small datasets, practical limitations of the technique are present with larger datasets due to the computationally intensive nature that causes meaningful lags and expense with large-scale utilization.

Furthermore, protecting the privacy of citizens is an essential responsibility of democratic governments. With continued collaboration between government, industry, and universities, more affordable solutions will become available. For example, text classification models were once nearly impossible to train in a differentially private way before the noise destroyed the sparsity in datasets, which dramatically increased the time and cost to train models.

By carefully thinking through the objectives and accounting for where the noise is needed, it is now possible to reduce training time from months to minutes. Achieving this requires a careful and deliberate accounting of the information in the algorithm and where it goes, but the payoff is significant. The upfront investment is well worth the long-term payoff: Once these optimizations become repeatable, the cost for subsequent deployments drops.

When it comes to making the case for more investment in data privacy, the dollar cost may not be the biggest hurdle. Far more formidable is the combination of organizational cultures that are resistant to sharing information, combined with a lack of understanding of new data privacy solutions. When executives don't fully understand the challenges that need to be addressed or the available solutions, they're less likely to champion a program that requires investment and cultural change.

But a lack of understanding is not a valid reason to put off implementing data privacy solutions. Federal agencies have an obligation to advance data privacy technologies as part of responsible AI, a reality that's been acknowledged at the highest levels. According to the 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence: "The Federal Government will enforce existing consumer protection laws and principles and enact appropriate safeguards against fraud, unintended bias, discrimination, infringements on privacy, and other harms from AI."

Executive Order 14110 goes on to direct the National Institute of Science and Technology (NIST) to create guidelines for agencies to evaluate the efficacy of differential-privacy-guarantee protections, including those for AI. At the time this article went to press, NIST indicated that a final report was in progress.

Alternative Paths to Privacy and Supplemental Techniques

Differential privacy offers a unique advantage by processing privacy risk incrementally, in contrast to many other techniques that often oversimplify privacy as either "protected" or "not protected." While these alternative methods can provide a layer of data privacy and security, they are not foolproof. With the growing volumes of data being collected and often insufficiently protected, the risk of exposure grows larger each time data is processed. This is where differential privacy stands out: It ensures that the risk of reidentification remains within the defined bounds of its mathematical framework, regardless of how much data on an individual is already available.

However, in cases where differential privacy may be too difficult to implement or situationally incompatible, alternative techniques can still help mitigate exposure risk. These techniques include synthetic data, federated learning, and homomorphic encryption. They can strengthen data protection when differential privacy isn't possible or bolster differential privacy solutions when they are in place.

Max Wragan *leads AI strategy and risk management to help Booz Allen and their clients deploy AI and AI tools to responsibly enhance project delivery with cutting edge technology.*

Edward Raff, Ph.D.*, leads Booz Allen's machine-learning (ML) research team, develops high-end technical talent, and disseminates the latest ML skills, techniques, and knowledge across the firm.*

Sean Guillory, Ph.D.*, is a lead scientist on Cognitive Domain Operations projects and an AI program manager for the AI Rapid Prototyping Museum.*

SPEED READ

Advanced algorithms and AI tools are making it easier to breach data privacy, with recent incidents showing how readily personal information can be extracted from seemingly protected datasets.

Differential privacy—a mathematical framework that introduces controlled noise into datasets—offers a promising solution, as demonstrated by its successful use in the 2020 U.S. Census.

While implementing robust privacy solutions requires significant investment, federal agencies have both the responsibility and opportunity to lead innovation in data privacy protection.

From the Frontlines of Post-Quantum Cryptography

SAFEGUARDING CRITICAL INFRASTRUCTURE FROM EVOLVING CYBER THREATS

Taylor Brady, Jordan Kenyon, and Derek Aucoin

Transitioning to post-quantum cryptography (PQC) will be one of the defining cybersecurity challenges of the next decade. The National Institute of Standards and Technology's (NIST) long-awaited PQC standards are finally here. They outline the implementation requirements and specifications for the approved quantum-resistant algorithms. Still, implementing PQC will be a complex process. The Office of Management and Budget (OMB) estimates that between now and 2035, PQC transitions will cost federal agencies over $7 billion. That estimate explicitly excludes national security systems and does not account for the costs of PQC adoption in the commercial sector. It is difficult to understate the magnitude of this transition.

Quantum computers will eventually break nearly all currently deployed public key cryptography—the algorithms embedded deep into hardware, software, and digital protocols that protect networks from threat actors. It is a question of when not if. Significant engineering challenges must still be overcome before quantum computers can reach the scale and robustness needed to run algorithms that threaten public key cryptography. Despite these hurdles, many commercial roadmaps predict that such devices could become available near 2030.

This accelerated timeline increases the pressure on organizations to transition to NIST's new PQC standards as quickly as possible. The urgency is even greater for organizations with sensitive datasets that have longer security shelf lives. Such data may already be the target of "Hold Now, Decrypt Later" (HNDL) attacks, in which a nation state or criminal organization breaks into a network, steals encrypted data, and stores that data on its own servers knowing that it will be able to access it once quantum computers capable of breaking through public key cryptography are available. The Center for Strategic and International Studies has highlighted the real threat HNDL attacks pose, further underscoring why federal agencies should not wait to start moving to PQC.

Though a growing number of federal requirements have started to structure how agencies approach PQC, the roadmap to achieving it is not as linear as it may seem. Federal chief information security officers (CISOs) have broad discretion to design PQC strategies that align with their agency's unique attack surface and threat vectors. Commercial companies have even greater discretion. In practice, this means that where organizations start their journeys to PQC is not as important as how they accomplish each phase.

Fortunately, the experiences of early adopters offer key lessons for enterprises embarking on the three core stages of preparing for PQC: cryptographic discovery, prototyping, and agility.

THE BASICS

1. **Quantum computers will eventually break nearly all currently deployed public key cryptography.**

   This is why it's imperative for government agencies and private companies to start taking steps toward protecting their assets with new cryptographic algorithms.

2. **Post-quantum cryptography (PQC) refers to the implementation of algorithms capable of withstanding a cryptanalytic attack by a quantum computer.**

   It is the best defense currently available to address the cyber threat posed by large-scale quantum computers. However, implementing PQC won't be easy and organizations must begin taking action immediately.

3. **Transitioning to PQC is especially urgent for certain federal and commercial organizations given the risk of HNDL attacks.**

   With a HNDL attack, an adversary acquires information that has not yet been resecured with PQC and stores those assets until a quantum computer capable of breaking that encryption becomes available.

Cryptographic Discovery: Mapping the Attack Surface

Cryptographic discovery is the process of creating an actionable, prioritized cryptographic inventory by detecting, tracing, and rating the cryptography in use throughout an enterprise based on its security in the post-quantum era. Cryptographic discovery is an intuitive goal, but it can be extremely difficult to achieve. Common cybersecurity tools detect cryptography by design, but they do not catalog vulnerable cryptography to enable prioritization and remediation. For this reason, most tools are ill-equipped to provide the visibility organizations need into cryptographic vulnerabilities driven by emerging quantum computing technologies.

Many new products are emerging in the market to address this gap, but adding additional security tools comes at a price, increasing both the total cost of migration and the time it takes to complete the migration. That is time many organizations don't have given the risks of "Hold Now, Decrypt Later" attacks and resources they may not need to expend. Rather than purchasing new products, some organizations are turning to novel data engineering methods to overcome common cryptographic discovery challenges.

CASE STUDY

When the U.S. government raised the alarm on the criticality of PQC in 2022 through National Security Memorandum 10, the Quantum Computing Cybersecurity Preparedness Act, and OMB's Memorandum on Migration to PQC, one Fortune 10 retail company took note. They found that investing in a scalable, production-grade analytics platform to dynamically discover cryptography across their large, federated systems enabled them to understand their risk exposure without the need for new cyber telemetry.

Extract, transform, load (ETL) pipelines were used to optimize the use of cryptographic metadata from existing sensors and maintain traceability to certificates and unique connections. As a result, security leaders across the organization could use the dashboard to see a real-time snapshot of cryptographic strength across the network and analysts could trace vulnerable cryptography back to its source for remediation, submit custom queries, and expand the discovery tool's coverage to new network boundaries.

KEY TAKEAWAYS

1. START SMALL
Define a priority network for an initial cryptographic discovery initiative, recognizing that everything cannot be transitioned all at once.

2. OPTIMIZE REUSE
Extract cryptographic data from existing sensors on the network to increase speed-to-solution and avoid added infrastructure costs and complexity.

3. ENGINEER FOR FLEXIBILITY
Invest in a scalable platform that can be used for both initial inventories during PQC planning and ongoing monitoring during PQC implementation.

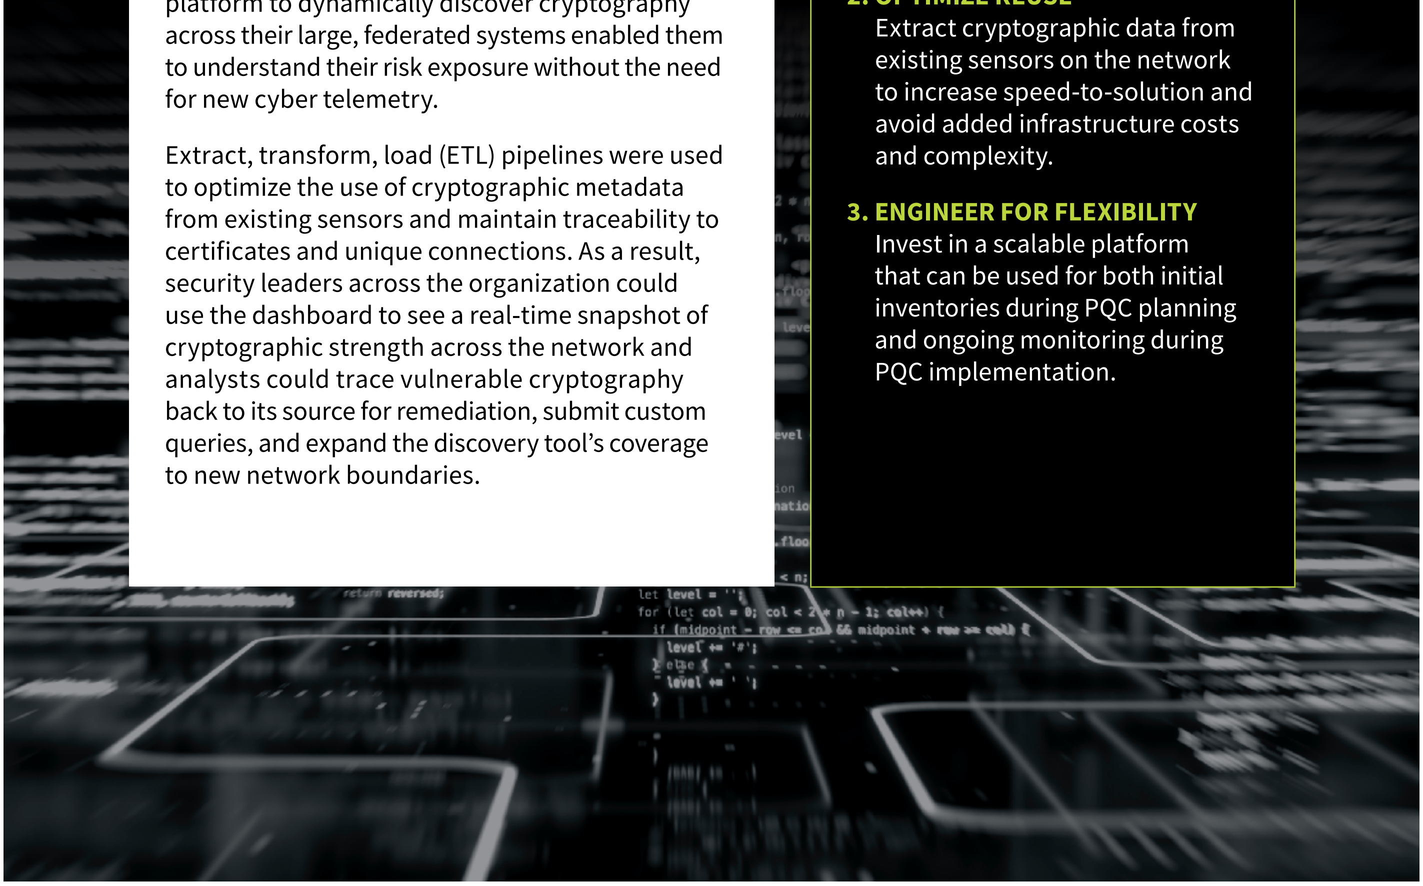

Prototyping: Putting PQC to the Test

Cryptographic discovery is a common starting point in the journey to post-quantum security, but it isn't the only place an enterprise can take its initial steps toward PQC. Discovery allows organizations to achieve breadth in PQC planning; prototyping enables depth. Prototyping focuses on modeling and measuring the performance and interoperability impacts of transitioning to PQC.

The math behind PQC is fundamentally different than that of legacy public key cryptography. Higher computational complexity makes PQC a robust defense against quantum attacks. It also introduces network and infrastructure challenges such as increased latency, increased bandwidth, and lack of interoperability. Understanding the impact of these challenges is important to inform procurement decisions, implementation decisions, and algorithm selection in use cases with multiple PQC algorithms, such as digital signatures.

CASE STUDY

A federal agency responsible for securing critical networks was especially attuned to the importance of performance and interoperability during their PQC transition. They needed a prototype system to provide quantifiable performance and interoperability test results, including impacts on existing hardware and software.

The agency used a test harness to help quantify the effects of PQC by simulating multiple connections and executing test scenarios that varied network traffic, bandwidth constraints, and algorithms used throughout the connection across authentication and transport layer security (TLS) negotiations negotiations. A dashboard automatically visualized searchable results. This enabled analysts to understand how negotiations would default to classical algorithms when an endpoint was not configured for PQC; identify the effects of hybrid certificate chains; and quantify the overhead cost and impact on the network.

KEY TAKEAWAYS

1. **DEVELOP REFERENCE ARCHITECTURES**
   Determine where cryptography is used in a priority use case to define prototype implementation, showcase vendor dependencies, and outline interoperability requirements.
2. **ESTABLISH A REUSABLE TEST ENVIRONMENT**
   Establish a laboratory environment that can simulate hybrid and full PQC solutions to determine optimal algorithm selection and implementation.
3. **ENGINEER A PQC PROTOTYPE**
   Assess hardware and software limitations, performance impacts, and interoperability across the identified use case.

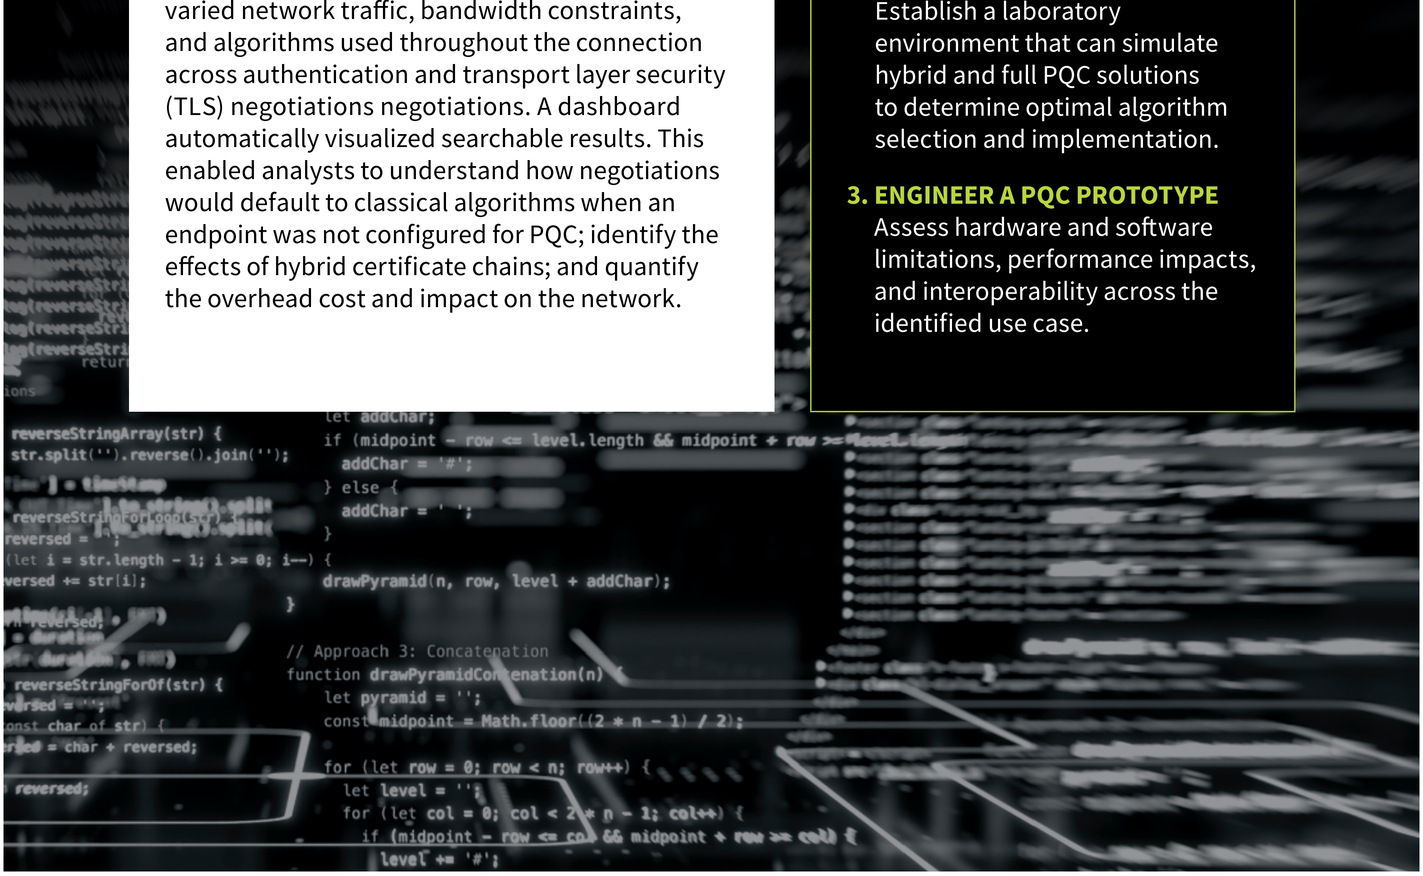

Cryptographic Agility: Optimizing for Continuous Security

Since the 1970s, public key algorithms have secured our digital lives. These unintrusive protections have been embedded deep into hardware, software, and digital protocols. But agility and governance were not baked into design decisions about cryptographic implementation. Hardware vendors integrated cryptography in ways that often prevent it from being updated without replacing an entire chip. Software vendors didn't track the cryptography throughout the different layers of their applications. These vendors did not imagine a future where the underlying math behind every public key cryptographic algorithm would be vulnerable to attack. Yet that is the reality today, and it requires enterprises to adopt new PQC standards.

Now, NIST's initial PQC standards (published in August 2024) provide the best available approach to defend against the quantum threat. However, it is possible that future technology advances could make those standards vulnerable to attack. It is also possible that additional and forthcoming PQC standards could offer performance advantages over NIST's initial standards.

This is where cryptographic agility comes into play. Cryptographic agility refers to the ability to rapidly find, monitor, update, and replace cryptography. It addresses an enterprise's capacity to navigate future cryptographic changes. This agility is essential for PQC, but its significance extends beyond post-quantum cybersecurity. When undertaken proactively and in concert with other cyber modernization priorities, PQC strategies that emphasize agility can increase the overall effectiveness and efficiency of organizations' procurement decisions.

CASE STUDY

A defense customer recognized the need to invest early in PQC to safeguard their infrastructure during high-impact missions. They were already transforming larger cyber operations in areas like zero trust and cryptographic modernization, but they knew these initiatives did not address PQC. These larger efforts can involve the purchase of new, expensive, built-to-last equipment, such as tactical radios and encryptors. Streamlining these decisions to include PQC considerations stood to increase the efficiency and effectiveness of their procurement in the long run, reduce unnecessary purchases, and prevent vendor lock-in with solutions that were not actively preparing for the PQC transition. They needed to develop a plan for cryptographic agility.

Convening internal and external stakeholders, the agency cataloged cryptographic dependencies across vendor and legacy equipment; monitored the existing, new, and potential cryptographic implementations; and inventoried replacement options for capabilities that could not support PQC algorithms.

KEY TAKEAWAYS

1. **INTEGRATE PQC INTO ONGOING CYBERSECURITY MODERNIZATION**
   Define how PQC aligns with future security architectures to buy down future technical debt and enable rapid adoption.
2. **ENUMERATE PROCUREMENT AND POLICY IMPACTS**
   Make PQC a priority with vendors to prevent lock-in with products that lack PQC transition plans.
3. **ESTABLISH STRATEGIES FOR INFRASTRUCTURE MANAGEMENT AND POLICY ENFORCEMENT**
   Enable knowledge transfer to enforce governance throughout the transition.

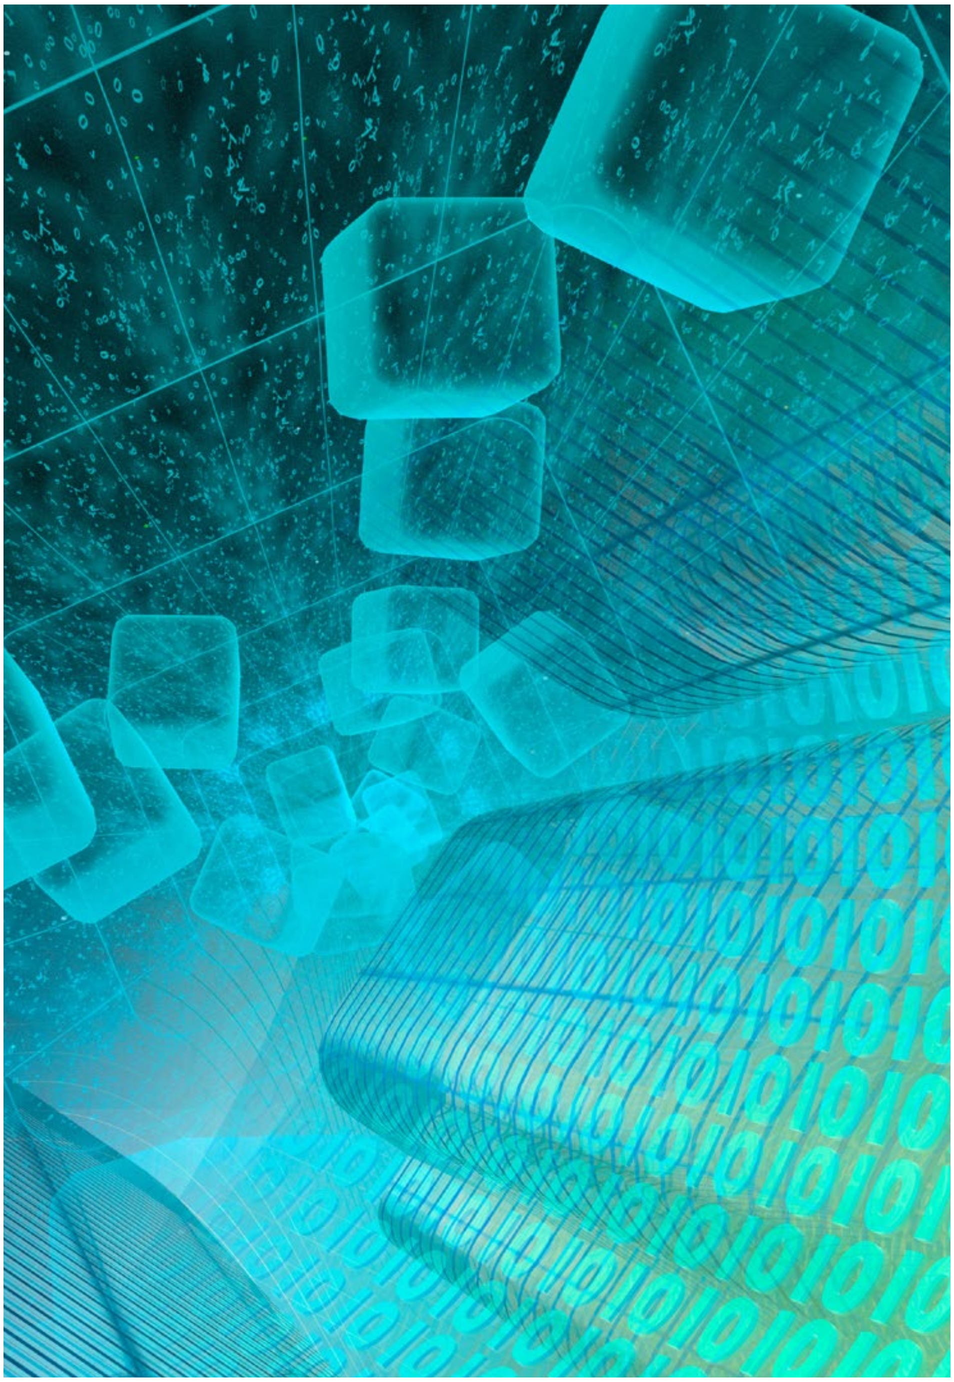

There Is No Wrong Place to Start the Urgent Transition to PQC

The path to PQC can be flexible, but it is a critical defense that cannot be ignored. Every enterprise can align its approach to address priority threat vectors. These case studies demonstrate the primary effects organizations have seen in their migrations to date. For some, the initial focus was resolving the gap in capabilities that would enable them to quantify their quantum attack surface (prompting cryptographic discovery). For others, it was the possibility that insufficient performance and interoperability testing could lead to network failures in mission-critical environments (prompting PQC prototyping). Others recognized that failing to include PQC in ongoing initiatives could cost millions (prompting cryptographic agility).

Cybersecurity leaders benefit from broad discretion in how they begin their PQC journeys, but organizations must begin now to defend against advanced cybersecurity adversaries and continuously evolving threats.

Taylor Brady *leads Booz Allen's post-quantum cryptography engagements and investments, specializing in technical adoption and business development.*

Jordan Kenyon *leads growth and operations for Booz Allen's quantum technology portfolio, which focuses on the new paradigms that quantum introduces for computing, sensing, and communications.*

Derek Aucoin *leads secure* product *and application development for Booz Allen's Global Commercial practice.*

SPEED READ

CISOs must design unique PQC transition roadmaps that align with their agency's attack surface and threat vectors.

Enterprises can avoid added infrastructure costs and increase speed-to-solution by extracting cryptographic data from existing network sensors and analyzing that data in novel ways to enable PQC migration.

Integrating PQC transition plans with other cyber modernization efforts is critical to prevent vendor lock-in with solutions that are not actively preparing for the PQC transition.

MISSION SPOTLIGHT: DEFENSE

Reconstructing the Tactical Mission Lifecycle

THE EMERGING TOOLBOX FOR MODERN PLANNING, READINESS, AND PERFORMANCE

Cameron Mayer, Eric Syphard, and Todd Burnett

Long before any warfighter boots hit the ground or planes take to the skies, some of the most important military actions happen far away from the battlefield. Today, U.S. commanders spend significant time and resources planning operations and training personnel for what they will face. Such preparation is critical for mission success. It is also costly and time-consuming and not nearly thorough enough due to unrealistic field-training environments and barriers to integrating data that's needed to improve unit performance.

Accelerating U.S. armed forces' readiness requires a more effective and efficient approach to mission planning. When used together, generative artificial intelligence (GenAI) and digital twin technologies can reinvent the entire tactical mission lifecycle—from planning and rehearsal to after-action review. Digital twins let warfighters train for and rehearse actions in realistic virtual environments. GenAI improves the planning process by rapidly analyzing a mission's parameters against its objective and other information and generating and evaluating optimized courses of action for each key decision point. Essentially, the modern flywheel (see the cover story on page 32) can be applied here to accelerate U.S. armed forces' readiness.

The result: Warfighters can simulate and train for countless scenarios that they might encounter on the battlefield, leaving them better prepared on strategic, operational, and tactical levels to execute missions safely and effectively. The bonus: Planning and training in a virtual environment are time- and cost-efficient because they can be done where needed, reducing travel time, and with fewer resources, such as equipment and ammunition.

The Layers of a Digital Mission-Planning Environment

A digital mission-planning environment is constructed from vast amounts of data from across the operating environment, the warfighters themselves, and military doctrine. The foundation is a digital twin: a technically exact, virtual replica of an object, process, or system (learn more about digital twins on page 34). A digital twin can be created within a few hours from data collected via existing maps and drone flyovers using full-motion video and light detection and ranging scans. The result is a high-fidelity rendering of where a military mission will take place.

The next layer is the data, which is embedded over the digital twin to provide an easily digestible look at key variables and other parameters. This includes the data used to create the digital twin as well as data that is pulled into the system about military doctrine, individual warfighter readiness, weaponry, weather, and open-source and signals intelligence. Analysis and measurement happen in the embedding layer through the deployment of traditional AI and mathematical optimization tools. For example, a traditional mathematical optimization model could determine the fastest route between points A and B for a four-person unit carrying a defined set of equipment. AI models could analyze the terrain using the categories of the Army's OAKOC: obstacles, avenues of approach, key terrain, observation and fields of fire, and cover and concealment.

In addition to replicating the mission environment, digital twins can generate data that goes beyond what sensors affixed to drones can capture. A digital twin can be programmed to simulate different weather conditions, combat scenarios, and other variables. The value of this simulation is twofold: Warfighters can test themselves and evaluate their performance in a variety of conditions, and the AI optimization tools can train on this additional data to improve their functionality. (Learn more about AI's recursive nature in our cover story on "Tech at Full Speed.")

Augmenting Human Analysis and Insights

Introducing a GenAI assistant transforms a digital-mission-planning environment into a next-generation capability. A GenAI model goes beyond predicting outcomes to generate new content and insights. It can use information and patterns gleaned from unimaginably large datasets to not only recommend the best course of action but also tweak that recommendation based on additional variables and feedback that the commander feeds it in real time. For mission planning, a foundational GenAI model trained on broad, open-source information is augmented with mission-specific information that isn't publicly available, such as military doctrine and procedures, uniform and equipment specifications, an adversary's typical warfighting tactics, and classified mission-specific information.

GenAI can also monitor the performance of individual warfighters. Today, warfighters can be outfitted with wearable technologies that generate information about heart rate, respiration, sleep patterns, stress levels, posture, movement speed, stride length, and other metrics. This biometric data provides critical assessments of physical and mental readiness, which can be combined with a three-dimensional holograph of each warfighter in a display that also includes tactical equipment models and information about the terrain to be traversed. Within the display, planners can "drag and drop" different equipment onto the warfighters to ensure every warfighter can effectively manage the load they're assigned. This in turn leads to an increase in overall speed during the mission because a unit can move only as fast as its slowest person.

Orchestration software that uses agentic AI is what makes everything work (learn more about agentic AI on page 41). When a commander inputs a query, a "commander" orchestrator agent interprets the request and assigns tasks to subordinate agents. Each agent consults its primary military procedures to identify the appropriate steps to answer the question. Then, based on available analysis tools, the agents execute each sequential step

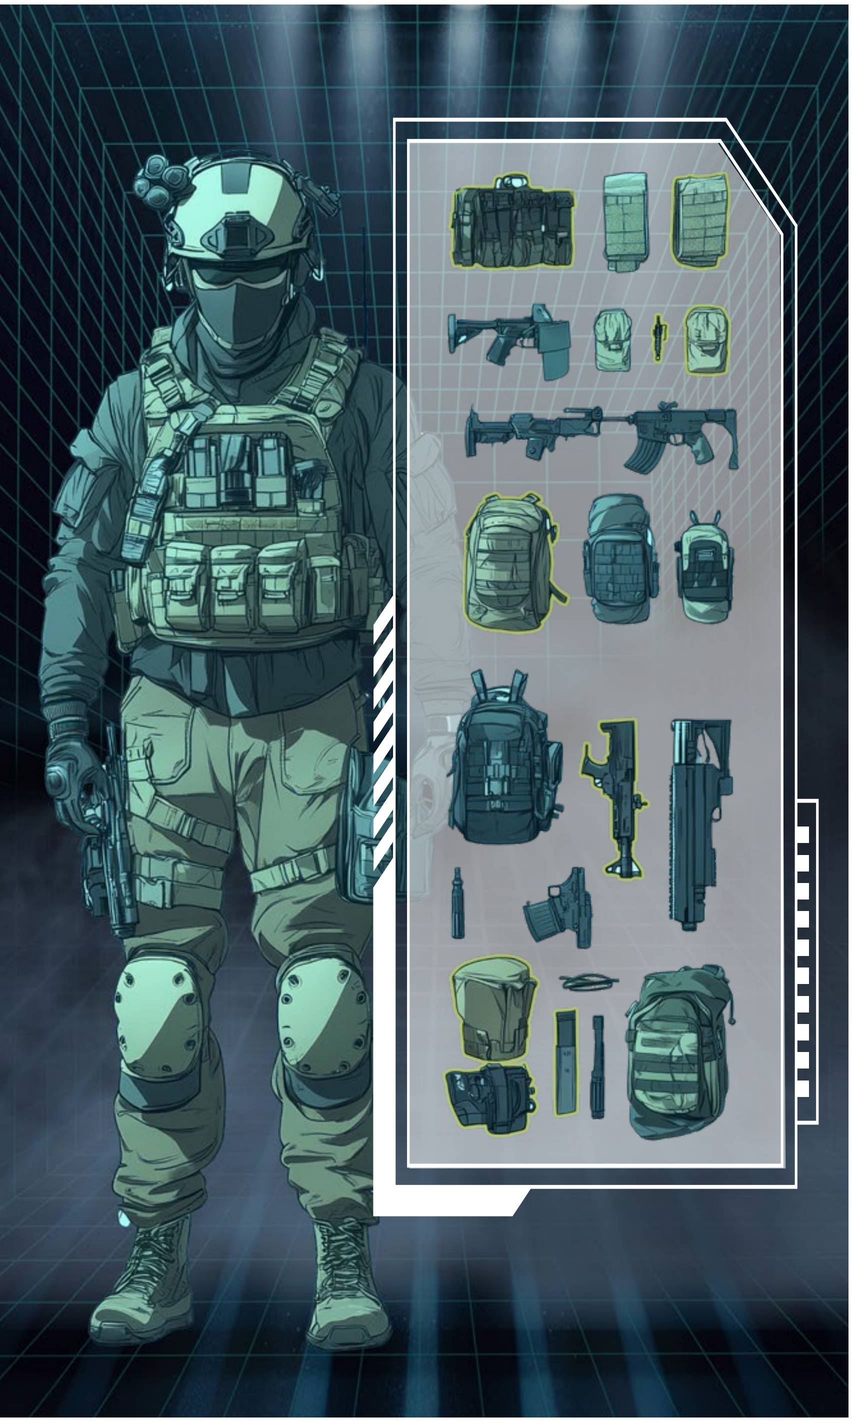

Figure 1: Digital twins let commanders manage each warfighter's load.

or note a gap in the available data or capabilities before returning a response to the “commander” agent for review, consolidation, and response to the user.

Each step in this sequence is logged and available for the user to investigate to ensure transparency and auditability and provide quality control against any potential AI-generated “hallucination” that could have an adverse effect on the operation.

Importantly, nothing about this approach takes decision making away from the commander. Instead, it provides more tools to work with and more hard data to draw from, to which the commander can apply their experience, knowledge, and judgment to make better decisions. In addition, humans must review recommendations with an eye toward ethical principles, situational nuances, and real-world conditions that may be beyond current AI models’ understanding and instead draw from their experience and expertise.

Rehearse as You Fight

Conducting rehearsal exercises is the next step in the mission-planning lifecycle. Exercises centered on terrain models and one-dimensional maps require participants to simply imagine that they are in a place they have never seen, divining the steepness of a hill or the density of a jungle. Even typical digital terrain models lack real-life fidelity, limiting participants’ ability to see the terrain and conditions they will face, and they don’t incorporate other essential information, such as weather and potential threats.

With the digital twin, warfighters can train as they fight, rehearsing the plan repeatedly in a true-to-life, three-dimensional simulated environment, immersed in the actual area of operations. They can see the terrain, elevation, vegetation, villages, buildings—everything they will encounter. The mission can be rehearsed 10 times or 100 times before any boots hit the ground, including the course of action as well as medical evacuation, vehicle recovery, and enemy collection plans.

GenAI can dynamically change the scenario as warfighters run through it, creating new and unexpected challenges. AI can evaluate the effectiveness and precision of operational activity and the warfighter’s physical and psychological systems during the action to enable increased learning and readiness.

Sophisticated electronics packages make it possible to blend the simulated environment with real physical weapons in call-for-fire rehearsals. Soldiers encounter a 360-degree virtual experience while simultaneously engaging with a surrogate weapon that provides a realistic tactical and kinesthetic feel. The electronics package includes a microprocessor that collects sensor inputs from the weapon, such as the trigger and safety, and a gyroscope that acts like a motion sensor, tracking the weapon’s movements.

As simulations are presented to the soldier, their actions and performance are depicted on a screen. They can view the target through their weapon’s optics, measure the range, execute the fire, feel the recoil, and assess the damage. This type of robust, mixed-reality rehearsal exercise in a virtual environment allows soldiers to practice the call-for-fire again and again, developing the necessary muscle memory and learning how to control their emotions and remain calm in stressful situations. AI can pinpoint the factors behind an unsuccessful call-for-fire rehearsal, such as mistakes made because of stress or lack of training—giving each warfighter and their commander real-time insight into what’s inhibiting their performance and how to overcome it.

Industry Perspective: Reveal

Reveal is a leading visual analytics and edge AI company committed to helping defense and public safety personnel achieve decision dominance at the tactical edge. Vice President of Defense and Security Dave Caudle shared his perspective on these technologies and their roles in military operations.

You're a former U.S. Army officer, a current battalion commander in the U.S. National Guard, and a technologist. How will AI change military operations in the future?

I see AI becoming an essential part of the battle staff. The technology excels at sifting through data and conducting analysis, which makes it an ideal solution for enhancing the situational awareness of squad and platoon leaders on the front line. Achieving overmatch in combat often comes down to the capacity to make smart decisions faster than your adversary. You need to understand where to move your forces and the optimal path for getting there. Sometimes a decision as simple as where is the best place to land a helicopter can have enormous consequences. Unmanned air systems (UAS) can capture all sorts of data about the conditions, and AI can analyze that data and present a squad leader with several potential paths forward as well as the pros and cons for each. This is a level of insight that leaders in forward areas have never had.

What are the biggest technical challenges to embedding AI at the tactical edge?

The challenge is making sure the platoon or squad leader has an interface that allows them to communicate with the AI quickly and effectively. These algorithms conduct analysis at such a rapid pace, but it's all for naught unless the person on the receiving end can access and understand the analysis at the speed of the mission. That's why it's imperative to design these interfaces to be as intuitive and easy to use as possible because in conflict situations platoon leaders need to be able to receive information visually so they can ask for additional analysis from the AI or take a suggested action.

What will be a key factor in battlefield success in the future? It pains my soul a bit to say this, but I think the most effective future formations employ fewer soldiers and fewer tanks and anywhere between 10 to 100 times as many intelligent machines. That may sound a bit far-fetched, but everything we've learned over the past few years is that robotics and autonomy are changing warfare at a fundamental level.

In my opinion, drones are the new ammunition. That's how ubiquitous they are poised to become, and our success will be predicated on our ability to manufacture them at scale and make sure our warfighters are trained in how to use them effectively.

Reinventing After-Action Review

All plans change at contact, so after-action reviews of what actually happened are essential for optimizing future performance. Conventional after-action reviews center on personal observations from leaders and participants about operation and task performance compared with the intended outcome. This helps everyone to evaluate what did and did not occur and why—and what can be done to sustain strengths and improve weaknesses to do better next time.

Digital twins and GenAI inject incontrovertible data into after-action reviews. Performance measures are fused with the GenAI assistant, which recommends changes to the planned route to speed up time to target. Similarly, mission analytics measurements and biometric data captured during rehearsal provide insight into unit and individual performance and which changes can be made to improve the mission plan and execution. Performance measures can include squad formation accuracy, speed of travel, and individual stress and energy levels. Data from wearables can help commanders assign the right soldier to the right job based on information about individual strengths and aptitudes, such as whether someone naturally scans an environment and therefore would be a better gunner compared with someone who focuses only several meters in front of themselves.

Such a data-informed after-action analysis can be conducted after each rehearsal evolution to further optimize the plan. When training and rehearsing happens in a digital environment, review and analysis can also be integrated into the exercise so trainees can course correct in the moment—something that's not possible with live exercises.

Scratching the Surface of What's Possible

GenAI technology is on a trajectory to be used in real and near-real time to continually pull in new threat data so leaders can adjust decisions based on current information and GenAI's updated recommendations. When paired with digital twins, it can vastly improve how warfighters and U.S. military commanders prepare themselves for missions where the ability to make sound decisions in unpredictable environments is fundamental to success.

Furthermore, these same technologies can be used to optimize supply chains and logistics. Their application in training saves the time and cost associated with live environments, significantly improves pass rates, and has the add-on benefit of warfighters having more family

time because of reduced travel. The Department of Defense needs to continue to invest in AI because of its ability to improve decision making on both the tactical and strategic levels.

Cameron Mayer *leads the firm's defense accelerated readiness business with a focus on delivering data-driven solutions.*

Eric Syphard *leads the firm's AI engineering practice, which is dedicated to improving the performance of human-machine teams.*

Todd Burnett *is the firm's senior executive advisor for accelerated readiness, focused on delivering solutions for warfighter training and readiness.*

> "I think the most effective future formations employ fewer soldiers and fewer tanks and anywhere between 10 to 100 times as many intelligent machines."
>
> – **Dave Caudle,** *Vice President of Security and Defense at Reveal*

SPEED READ

GenAI and digital twins have the potential to reconstruct the tactical mission lifecycle—from initial planning and rehearsal to after-action review.

Digital twins enable realistic mission rehearsals in immersive 3D environments, allowing warfighters to practice and refine plans using real-time performance analytics.

GenAI agents trained on mission-specific information can dynamically change rehearsal exercises, pinpoint factors that are inhibiting performance, and provide warfighters with real-time insights regarding how to overcome them.

Intelligent Systems at the Edge

ACHIEVING DECISION ADVANTAGE IN UNPREDICTABLE ENVIRONMENTS

Joel Dillon, Randy Yamada, and Josh Conway

Tactical autonomy describes the concept of autonomous systems working independently within set parameters to perform tactical actions in real-world environments. It is poised to change U.S. military operations across all domains, from logistics in the Pacific to behind-enemy-lines surveillance, and combat search and rescue. It will also demand an overhaul of the way the Pentagon purchases and deploys one of the foundational elements of technology: software.

The impact of autonomous warfighting platforms, such as the Department of Defense's (DOD's) Replicator and Collaborative Combat Aircraft programs, will correspond with the military's ability to deliver frequent and immediate updates to the platforms' software and underlying algorithms. Future enemies will adapt both weapons and tactics behind a veil of camouflage and deception under the fog of war. As a result, autonomous platforms will need fresh intelligence and new capabilities to distinguish friend from foe, navigate unfamiliar terrain, find and track adversaries, and more. The nascent field of tactical autonomy builds its foundation on data management to train the underlying AI, update signatures, and dictate actions.

Achieving mission readiness within this new paradigm demands closer collaboration between the Pentagon and the private sector, the nation's primary source of technological innovation. It will also create opportunities for DOD to procure and sustain software and modular sensors and other hardware in novel ways, thereby opening the door to new market entrants and diversifying America's defense industrial base.

The Advantages of Function and Expendability

Autonomous platforms will perform a range of functions on the battlefield, from intelligence, surveillance, and reconnaissance (ISR) and close combat support to communications relay. They will be flexible and adaptable: A single unmanned autonomous vehicle (UAV) could undertake a variety of mission sets in subsequent days. They will be "attritable"—designed to prioritize function and expendability rather than durability. In addition they will operate together as intelligent systems, collaborating tactically with minimal human intervention.

One key advantage of autonomous platforms is they will require lower data transmission rates than remotely piloted vehicles, which will enable them to operate for extended periods in austere, degraded, or denied communications environments. Whereas a remotely controlled drone needs to be able to stream video back to its pilot, an autonomous ISR asset can pilot itself; it only needs to send back the coordinates of a target and the probability of successful identification, for example. Reduced communication requirements relax data latency constraints, message complexity, and exploitability, requiring less power and allowing advanced jam-resistant encoding.

Attritable by design, aircraft sent to swarm enemy lines will have a much shorter life span than today's crewed platforms, which in turn will translate to a new model for sustainment budget and infrastructure. This shift is already underway: The Air Force retired the original MQ-1 Predator drone less than 15 years from its initial operating capability, while many manned aircraft have been flying for many decades.

Figure 1: Use cases for autonomy

Category	Description
Intelligence, Surveillance, and Reconnaissance (ISR)	Unmanned air, ground, and sea vehicles extend surveillance to gather battlefield intelligence, distribute sensors for real-time data collection, and track the activities of vehicles and units.
Close Combat Support	Highly maneuverable drones capable of low-altitude flight and precision tactics assist ground troops, conduct precision air strikes, and coordinate swarm attacks that overwhelm the target.
Mine and Improvised Explosive Device (IED) Warfare	Unmanned vehicles use sensors to locate and neutralize IEDs and naval mines.
Medical Support	Smaller unmanned vehicles shuttle medicines and equipment; larger vehicles with robotic systems provide medical care and transport wounded personnel.
Asset Protection and Security	Drones and robots provide escort services for high-value convoys and guard key installations.
Fire Support and Targeting	Platforms identify enemy positions, provide coordinates to artillery units, and evaluate the effectiveness of engagements.

Software as Dynamic as the Mission

Throughout history, adjustments in tactics, techniques, and procedures were disseminated through human channels, from the chain of command to the soldier in the trenches or the pilot in the cockpit. Mission leaders determined combat loads based on immediate operational needs, and an infantryman needed no extra training to carry a load on a multiday dismounted patrol one day and a combined arms engagement the next.

Now, autonomous platforms will draw from a vast library of mission-specific software and algorithms. Just as soldiers can only carry so much on their backs, the size, weight, power, and cost (SWAP-C) constraints on an autonomous platform's storage and compute capacity limit its combat load to the algorithms, software, and data necessary only for the mission at hand. What this means in practice is that a mission involving a kinetic strike will require a different software payload than a medical evacuation. Other data- and compute-intensive uploads could be systems for language recognition systems (e.g., Cyrillic street signs versus Cantonese) and terrain navigation in GPS-denied environments (e.g., an asset deployed in East Africa won't need detailed maps of coastal Guangdong Province), or computer vision algorithms for target recognition. The Pentagon will have to develop and continuously improve procedures for pushing updates to the fleet.

SWAP-C limits also require sophisticated machine learning operations (MLOps) to sort intelligence from noise. This is particularly crucial for long-duration missions in communications-denied areas, such as subsurface drones operating near adversary coastlines. These underwater vehicles must operate autonomously for extended periods. They require advanced algorithms to sharpen collected data: to distinguish between irrelevant environmental features (e.g., images of empty sea) and objects of military interest (e.g., ships, submarines, mines). Sharpening the data reduces the amount of storage capacity autonomous vehicles require. Upon completing their mission, these drones can resurface at predetermined locations to rapidly transmit relevant data to secure cloud systems.An array of plug-in hardware modules will enable versatility, such as different wing sets for loitering ISR or ground attack, and different loads for electronic warfare or a kinetic strike. The evolution of versatile autonomous platforms will also encourage innovation in hardware and sensors.

The recent history of warfare offers countless examples of tactical updates that couldn't have been foreseen before a platform was deployed to the field. Software requires updates to optimize performance, even in relatively static environments. Permissive and collaborative environments for complex systems, such as self-driving cars, can require frequent updates to improve safety and the driving experience. The contested world of the modern battlespace with an adapting adversary increases this need exponentially. From new adversary weapons, sensors, and signatures to refined tactics, optimal behavior will require daily or even hourly deployment of software updates.

The Integration Challenge

The technical, logistical, and bureaucratic challenges to bring autonomy to military operations, and soon, demand an unprecedented integration of hardware, software, algorithms, and human teams, both at the enterprise level and the edge.

Autonomous military operations will consist of multitudes of autonomous platforms and enterprise systems. Updates on one platform will need to be coordinated and synchronized across the system of systems. An autonomous submarine and an autonomous aircraft will have different manufacturers but will need to work together, with one sweeping for mines while the other provides aerial overwatch.

Intelligent systems consisting of multiple autonomous platforms will demand hyperscaler-level data and compute capacity. The logistical infrastructure to move all the data and algorithms must be treated as a critical warfighting system in its own right. The data itself is a strategic asset; in its "Fiscal Year 2025 Budget Estimates," the Defense Information Systems Agency noted "inherent power in owning data to control the high ground." The data and communications infrastructure will need a robust cybersecurity wrapper.

For each of the diverse mission scenarios and autonomous platforms undertaking them, developers must create, verify, and validate updates, and the technical and logistical complexity of disseminating them to the fleet securely is immense. A future force will need to update tens of thousands, or even hundreds of thousands, of unmanned systems simultaneously, creating the potential for an immense data bottleneck—envision the difficulties football fans have using

INDUSTRY PERSPECTIVE: SHIELD AI

Nathan Michael is Shield AI's chief technology officer and a former associate research professor at the Robotics Institute of Carnegie Mellon University. He's authored over 150 publications on control, perception, and cognition for single and multi-robot AI systems. We spoke with Nathan about the future of autonomous platforms.

You've conducted research on collective intelligence. Could you explain what that is?
Collective intelligence is about creating systems that can execute against specific mission requirements in a variety of environments. Swarming, especially in communications and GPS-denied environments, is a great example. It involves large numbers of highly intelligent systems working together. In worst-case scenarios where communication is lost, each system is capable of operating independently. As communications are restored, they can reestablish coordination and collaborate to achieve mission objectives. By creating resilient, adaptive teams that can function across different mission sets and domains, we're moving toward swarming—where highly capable, intelligent systems operate proficiently together and remain resilient, even in contested environments.

What's the biggest challenge to operationalizing autonomy in real-world environments?
The biggest challenge is scaling production. Plenty of research and development labs in industry and teams working in academia are already creating effective mission autonomy capabilities. The question is how do you build the end-to-end testing frameworks for verification, validation, and deployment? How do you come up with a well-defined concept of a factory that allows you to proliferate production and scale its deployment? The true innovation of the Model T was not the car but the creation of factories that could mass-produce it. We face a similar challenge in that our industry has to move beyond designing and building bespoke autonomous capabilities that are applicable to specific scenarios and toward processes that enable us to build autonomous platforms for a variety of mission applications at speed and scale.

What does the future of autonomous systems look like?
Right now, tactical autonomy involves systems executing specific, mission-bound tasks like ISR or precision strikes—things that traditionally require human input at every level. An example of tactical autonomy would be a drone autonomously navigating a GPS-denied environment to carry out a targeted strike. However, strategic autonomy is about more than just completing isolated tasks. It means systems will be able to assess the broader operational environment, make high-level decisions, and adapt their objectives as the mission evolves. For instance, a strategically autonomous system could coordinate multiple assets; manage a complex operation across different domains (air, land, sea); and adjust mission priorities in real time based on battlefield conditions. This allows fewer personnel to command greater effects with more precision and efficiency.

The shift from tactical to strategic autonomy is also about minimizing risk to personnel while maximizing mission outcomes. With fewer people in harm's way, more intelligent systems will be processing data and making decisions at scale. This increased cognition in autonomous platforms will empower operators to make better decisions with fewer resources, enabling greater efficacy and efficiency on the battlefield. In short, we'll see fewer humans required to achieve much larger mission outcomes while intelligent systems take on more complex, high-level roles that were once the domain of human commanders.

mobile devices in the stadium during a game. That effort will demand advanced edge-based routing to move data on the best available network without human intervention.

A combatant command could soon store 10,000 powered-down UAVs in a warehouse for months until an emergent event. The force would need to update and synchronize these crafts as quickly as possible for deployment, including fresh software and algorithms. An autonomous civilian car can take multiple updates per week in a permissive communications environment. Military operations can take place in contested, non-permissive environments where the connectivity required to receive updates is not always a given.

What Progress Demands

A new paradigm for military operations demands a new paradigm for software development and acquisition. The move toward attritable autonomous systems will bring wholesale change in the economics of national defense, notably in a shift in the Pentagon's budget from conventional operations and maintenance to software. It will also require DOD to think differently about data. To ensure the continuous improvement of autonomous platforms, it will be imperative to build capabilities that can gather and store as much data as possible in permissive training environments with less network constraint. This data will train the algorithms that will serve as the brains of autonomous vehicles and weapons.

DOD's models for tech acquisition were developed over decades to suit the purchase of giant, costly pieces of hardware or massive enterprise software. Those acquisition models will need to be nimbler and more flexible to keep up with technological innovation. Under current procurement and operations and maintenance models, a platform's original contractor is typically tasked with maintaining and updating its hardware and software. But the startup or tech company that manufactured an autonomous aircraft may not be best suited to write mission-specific algorithms; its engineers and executives may not have the necessary experience working with federal clients or the security clearances to fully understand the nature of the operation.

Original equipment manufacturers typically lock in sustainment funding on an exquisite, sophisticated legacy platform for the decades it's in operation. Attritable platforms with short life spans won't need as much sustainment, but they will need frequent software updates. The lower barrier to entry into the software market as opposed to developing and manufacturing, say, a crewed fighter jet, will enable more companies to pursue innovations, with greater efficiency and resilience in the supply chain. Furthermore, it may not make financial sense for every private-sector company that makes high-quality drones to invest in customizing its product for the highly specific parameters of defense operations. As autonomy becomes increasingly essential to our nation's defense, DOD may need to retain control of the algorithms that govern the behavior of autonomous systems. By keeping the code separate and owning it, DOD can buy the best available drones and UAVs without having to rely on manufacturers to tailor the product to the contours of ever-changing missions.

Divorcing the acquisition of hardware from the acquisition of software also will enable DOD to buy the latest and greatest from an expanded Silicon Valley vendor base working in concert with prime defense contractors. Imagine a software package that excels at performing post-mission analysis that is available to every squadron that uses UAVs, regardless of which manufacturer makes the hardware. Autonomous systems will continue to evolve and acquire newfound capabilities that can serve a variety of key mission functions. Pivoting to a software-first mindset will help position DOD to bring the most innovative technologies to bear on the greatest challenges.

***Joel Dillon** is a leader in Booz Allen's Global Defense Sector driving next-generation technologies through the firm's Digital Battlespace business.*

***Randy Yamada**, Ph.D., is a technical leader overseeing autonomy and embodiments of physical AI for the firm's Global Defense Sector.*

***Josh Conway**, Ph.D., develops the strategy and implementation of autonomous systems at scale.*

SPEED READ

To successfully deploy tactical autonomy in U.S. military operations, the Pentagon will need to rethink how it purchases and deploys one of the foundational elements of technology: software.

Autonomous platforms will draw from a vast library of mission-specific software and algorithms and will require regular updates based on the missions they are set to perform.

Divorcing the acquisition of hardware from software also will enable DOD to buy the latest and greatest from an expanded Silicon Valley vendor base working in concert with prime defense contractors.

MISSION SPOTLIGHT: LAW ENFORCEMENT

Another Set of Eyes

INTELLIGENT TOOLS TO AUGMENT LAW ENFORCEMENT PERSONNEL

Carl Ghattas, Todd Kline, and Thong Nguyen

How will law enforcement officers (LEOs) benefit when certain aspects of Tony Stark's "Iron Man" suit become viable in real-world scenarios? Rocket-propelled armor may be the stuff of science fiction, but body-worn sensors, ultra-low-power neural compute, and high-resolution augmented reality (AR) waveguide displays are rapidly becoming more powerful. Today, these technologies are available in form factors as compact as a pair of glasses. Tomorrow, they will be the foundation for a platform that enables superhuman-like capabilities in the field.

The convergence of these technologies with AI is spawning a new paradigm of "intelligent" field-use tools that can provide capabilities such as active language translation, context-aware wayfinding, or a heads-up view of criminal activity or tactical support within the vicinity. When integrated into a simple, wearable platform with 5G connectivity, these technologies have the potential to enable LEOs to see the unseen, communicate without needing to speak, and maintain a level of situational awareness (SA) beyond what's possible today.

However, as with night-vision devices or body-worn cameras, which took years to adopt, significant advances in field-worn technologies don't happen overnight. Before AI-powered wearables migrate from the laboratory to the field, agencies can take certain steps to prepare themselves and their personnel. It's imperative to understand how these capabilities can align with mission needs, what needs to be solved before they are deployed for operations, and how AI can reduce complexity and enhance SA for LEOs in the field.

Drowning Under a Tsunami of Data

Today, law enforcement and homeland security personnel face a daily challenge: manage the streams of data that are essential to everyday operations. This data comes to them in different modes (images, social media posts, tips) and from various points of origin (body-worn cameras, law enforcement databases). It is often siloed within organizational boundaries and under different classifications. According to a story published by *IT Brew*, in June 2024 Justin Williams, the deputy assistant director of the FBI's information management division, said that the bureau's Criminal Justice Information Services Division can hold over 30 petabytes of data at any time.

Federal agencies and local law enforcement are turning to AI to translate this tsunami of data into insights. According to *IT Brew*, Cynthia Kaiser, the deputy assistant director of the FBI's Cyber Division, said that the bureau has used natural language processing to analyze tip calls and identify missed information. In August 2024, the Palm Beach County Sheriff's Office told WPEC CBS 12 that it was rolling out a tool from Axon Enterprise, Inc., that uses generative AI and audio feeds from body-worn cameras to produce drafts of incident reports. These use cases free up time while maintaining human accountability, which is essential to ensuring agencies abide by the laws that protect privacy and the rights of the citizenry and govern the admissibility of evidence in court. They also point to a future where AI can help LEOs better process complex scenarios and reduce cognitive load in the field.

Optimizing Intelligence from the Office to the Edge

While applications of AI in law enforcement today have primarily been directed at office-based tasks, the rise of edge-capable devices raises new possibilities. How could this emerging tech address the unmet needs of LEOs in the field who are often removed from information sources and decision-making hubs? How might wearable AI enhance SA, which is critical for the success of field operations? By closing the gap between LEOs and essential information, this technology could result in distributed decision making and better response effectiveness.

For LEOs in the field, SA is a necessary yet difficult state to achieve. It can easily be impaired in high-pressure circumstances when split-second decisions can be the difference between life and death. Under significant stress, LEOs must perceive environmental cues such as potential threats in unfamiliar settings, understand what those cues mean in context (e.g., delineating between a person exhibiting nervousness and a person exhibiting aggression), and be able to project future scenarios based on current observations.

Expanding individual SA into a common operating picture (COP) is an even more sophisticated challenge. It involves the real-time fusion of insights across both individuals and disparate data sources to provide a unified view while simultaneously orchestrating the distribution of information relevant to each LEO's specific tasks. This complexity is multiplied in multi-agency operations where responders may be guided by different mandates, objectives, and information sources that can make collaboration difficult. In the dynamic and rapidly changing situations that LEOs face in the field, mistakes in relaying critical information in a relevant, timely manner can have dire consequences.

Despite the importance of COP technology in aiding public safety and security response, it is not uncommon to find these systems lacking the functionality that emerging technologies provide. For example, law enforcement agencies may use sophisticated assets such as helicopters and drones to capture valuable visual data. The problem is that once this information is received and distilled by central command (C2), the resulting insights are frequently relayed to field personnel verbally, rather than digitally. Furthermore, the potential for latency and error increases as the information is further propagated across other organizations.

There are meaningful opportunities to build systems that reimagine how data flows from the field to the office, across to other agencies, and back out. These systems must assist in parceling out the "noise" and helping LEOs focus on the essential tasks at hand. The trick isn't always how fast you can find the needle in the haystack, it's how efficiently you can remove the enshrouding hay.

A Glimpse of the Possible

Picture a pair of glasses, nondescript in appearance but purpose-built for use by frontline LEOs. When activated and worn, the glasses biometrically authenticate the wearer, which provides secure access to advanced capabilities. Specific mission data, such as suspect profiles, maps, license plates, or building floor plans, are pushed to them, enabling offline operation in the event of gaps in wireless coverage.

With a crisp and wide visual, the LEO sees information relevant to their mission overlaid in their field of vision. The data ranges from route guidance and traffic flow to the locations of suspects and team members. LEOs are also alerted to threats within the vicinity. During an operation, the glasses visually identify "friend or foe" at range, thereby reducing the chance of friendly fire on fellow on- and off-duty LEOs. The glasses also estimate crowd size and density, identifying key characteristics such as movement patterns. Facial recognition is performed on-device against the preloaded suspect profiles, reducing inefficient round trips while protecting the privacy of law-abiding citizens. A simple verbal command prompts the glasses to capture and digitally catalog data along with its surrounding context in time and space—information that supports future forensic analysis.

With its onboard scene intelligence capabilities, a pair of AI-powered smart glasses continuously maintains a temporal-spatial simulation of the world around it. The glasses stream delta updates of its local “world understanding” to C2, which gets fused with geographic information system data and other sources, such as drone and closed-circuit television (CCTV) feeds, to create a live, centralized digital twin of the area of coverage. This provides a unified view along with historical analytics and predictive simulation to support data-driven decisions, and the resulting central state is regularly synchronized back out to LEOs to give each individual real-time, swarm-like insight. A quick attachment of an add-on sensor enables the glasses to detect beyond the visible spectrum, including the presence of chemical, biological, or nuclear material, which enables the LEO to see and instantaneously report their findings. All of this happens without a word being said.

The underlying data streams can be captured as historical record and will be invaluable not only for after-action reports but also for training new recruits and synthesizing future scenarios that have yet to happen.

In addition to the passive assistance, LEOs can actively issue queries by voice, asking questions and giving directions such as: What is the last observed location of the suspect? Are there civilians in the area? Get me an aerial view of where 911 calls related to this incident are coming from. What is the quickest and safest route out of the area? With an understanding of both the individual and group contexts, the glasses can respond with answers tailored for the situation at hand in milliseconds. Up-to-date policies and procedures can be requested, which can be challenging to recall during or after stressful events.

The construction and deployment of such a platform lies within the realm of possibility. Commercially available smart glasses such as the Ray-Ban Meta collection, Snap Spectacles, or DigiLens Argo™ feature high-resolution electro-optical sensors and a Qualcomm Snapdragon® processor that can run AI models in parallel with a full-fledged physics-based simulation at 60 frames per second. They have the wireless connectivity for low latency reachback to secure infrastructure. They can already perform speech to text, text to speech, computer vision, stream voice, and video. The next generations of these glasses will offer even more. Making them mission-effective will boil down to stitching the underlying technologies together seamlessly to serve a purpose.

Within the context of law enforcement missions—which run the gamut of public safety, including facilitating the lawful movement of people and goods, disrupting illicit activity, or responsively investigating innumerable tips and leads—the primary challenge boils down to data fusion and distribution and the user experience. Any information relayed to a frontline LEO must do so in a way that highlights relevant context and removes noise without serving as a distraction or a cognitive burden.

To be effective in the field, this platform will have to be built under tight size, weight, and power limits. LEOs will not wear glasses that are fragile, require frequent battery changes, or are tethered to cords that inhibit movement. Many LEOs are already overloaded by gear: weapons, body cams, first aid gear, lights, and more. Updates to the devices, client apps, private cloud microservices, and corresponding infrastructure must be seamless and scalable as utilization increases, making “software-defined everything” (SDE) a necessity. If it’s not overwhelmingly and obviously effective, adoption will be difficult, if not impossible, to achieve.

Furthermore, in an emergent situation like executing a warrant or pursuing a suspect, the glasses must be able to anticipate what the LEO needs. The integration of multimodal AI will significantly increase the value proposition by directly supporting task execution. For law enforcement missions, glasses that leverage AI’s advanced analytical abilities to proactively identify what a specific situation demands—information about the surrounding area, requests for backup—rather than solely wait for verbal commands will enhance LEOs’ SA and accelerate decision making in critical situations.

From the Lab to the Field

It’s easy for discussions about emerging technology and its potential mission relevance to spark excitement, but achieving genuine mission transformation requires a far more nuanced and methodical approach. The technology landscape is littered with failed concepts that began with flawed assumptions. It is essential to incubate and develop these products in close partnership with current mission experts to ensure that they meet the practical needs and real-world constraints of field personnel. Iterative user testing and feedback are crucial to build trust, identify and address potential issues, improve usability, and validate core features. The employment of any human augmentation technology in support of public safety must be able to adapt fluidly to mission needs and withstand the demands of field use.

Two upcoming events illustrate how AI-powered wearables could help law enforcement surmount their data challenges: the 2026 FIFA World Cup™ and the 2028 Summer Olympic Games. The United States will host both events, and over the course of each, tens of millions of passionate international fans and tourists will converge physically and digitally in a short time span, bringing in billions of dollars in revenue, creating tens of billions of impressions, and attracting a massive surge of global attention.

These spectators and tourists will only add to the already increasing tsunami of data LEOs are confronted with. But instead of a broadly occurring uptick, the finite nature of these events will acutely focus this increase in data like water pressure through a fire hose. To effectively investigate any leads, identify and disrupt illicit efforts, and facilitate the safe and lawful movement of people, LEOs will need to possess as close to full SA as possible at all times. Could intelligent field-use tools be the bridge that links these no-fail missions with the data required to execute them successfully when and where needed?

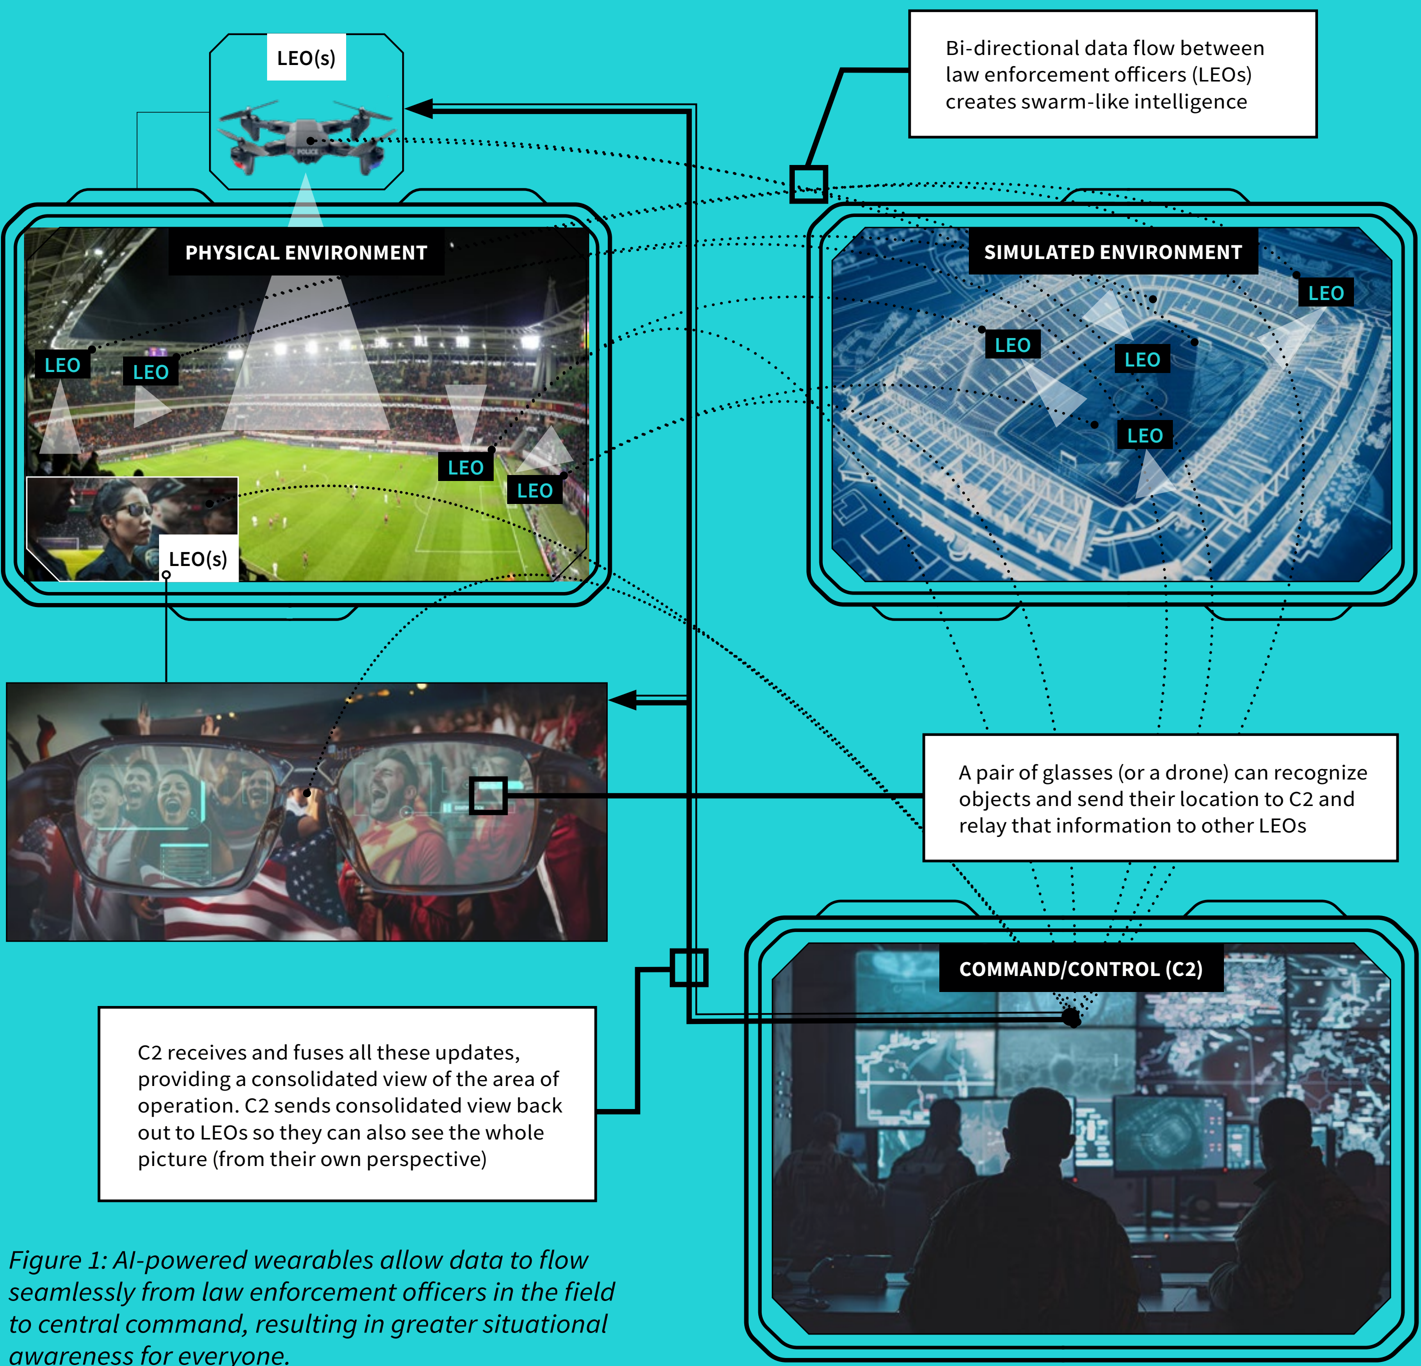

Bi-directional data flow between law enforcement officers (LEOs) creates swarm-like intelligence

A pair of glasses (or a drone) can recognize objects and send their location to C2 and relay that information to other LEOs

C2 receives and fuses all these updates, providing a consolidated view of the area of operation. C2 sends consolidated view back out to LEOs so they can also see the whole picture (from their own perspective)

Figure 1: AI-powered wearables allow data to flow seamlessly from law enforcement officers in the field to central command, resulting in greater situational awareness for everyone.

Imagine the following: Thousands of spectators are flowing toward a stadium, and a child becomes separated from her parents. The parents approach a nearby LEO seeking help, but they do not speak English. The LEO's "smart glasses" provide instant translation, and she radios the details of the situation in. Another LEO maneuvers a drone and locates an isolated child 60 yards from where the parents are. The child's location is spatially projected to the ground LEO's glasses, and the LEO then guides the parents through the crowd to where the child is stranded.

On the other side of the stadium, a solitary LEO reports a fight breaking out between inebriated spectators that is too large for one person to handle. Immediately, the officer-in-charge (OIC) in the established command center pulls up a digital twin of the stadium on a screen. Reflected on the display are the real-time locations of every field LEO, each of whom is wearing tech that allows for such precision. The OIC then vectors an appropriate number of personnel to assist in de-escalating the situation while shifting others to cover new areas. The locations and distances of those reporting to assist are now reflected on the reporting officer's display, thus guiding his plan to engage.

In each of these scenarios, the underlying value of the wearable devices lies in their ability to provide LEOs working in complex environments with access to data in a manner and at a speed that proves decisive in dealing with emergent and regular response requirements. The technologies that power these devices are evolving more rapidly than ever. In the commercial market, many observers are openly wondering if AI-powered glasses will replace earbuds and even mobile phones in the future. The next step for this technology is to find ways to fine-tune their capabilities to help the personnel tasked with maintaining order and public safety to execute their missions.

Learning from Night-Vision Devices

Night-vision devices (NVDs) are illustrative of a wearable technology that can create game-changing advances for law enforcement personnel.

Prior to the adoption of NVDs, law enforcement entities faced significant challenges when operating at night. Lack of visibility brought with it much higher risks. This heavy piece of equipment would not be worth taking into the field if not for its immediate, obvious, undeniable utility: its ability to give LEOs the ability "see in the dark," which in turn offers an overwhelming advantage.

Put another way, NVDs enable "data dominance." They let LEOs operating in dark conditions see what was previously unseeable. AI-powered wearables offer a similar value proposition by providing LEOs with access to intuitive data feeds that augment their abilities without increasing their cognitive load.

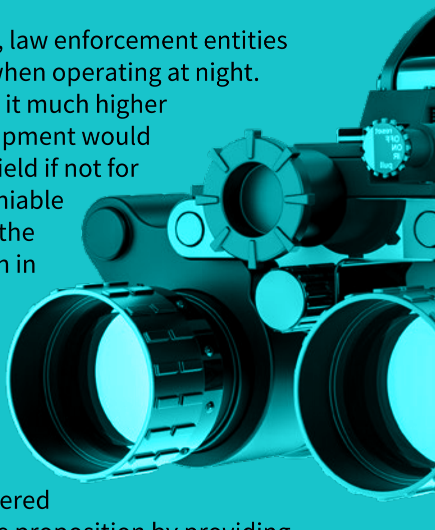

Carl Ghattas *leads Booz Allen's law enforcement and homeland security business.*

Todd Kline *is a strategy and transformation leader within Booz Allen's law enforcement and immigration market, and an adjunct professor at Bethel University.*

Thong Nguyen *is a human-centered technologist in Booz Allen's BrightLabs incubator, focused on advancing emerging technology that will transform future missions.*

SPEED READ

Advanced wearable technology, including AI-powered smart glasses, could provide law enforcement officers with superhuman-like capabilities such as real-time translation, threat detection, and enhanced situational awareness.

Law enforcement agencies are struggling with an overwhelming volume of data from various sources, and AI could help transform this "tsunami of data" into actionable intelligence in the field.

While the technology for these advanced wearables exists, successful implementation will require careful development with user feedback to ensure they enhance rather than burden officers' capabilities.

Booz Allen Chairman and Chief Executive Officer Horacio Rozanski and Chief Financial Officer Matt Calderone at the New York Stock Exchange (June 2024)

Catalyzing America's Innovation Power

Horacio Rozanski, Chairman, Chief Executive Officer, and President

There's an old saying: "If you want to go fast, go alone. If you want to go far, go together." Today's national security landscape, the most challenging we have faced in half a century, demands both—we need to go far, fast. The ability to accelerate the development of groundbreaking technologies will separate the nations that shape the future from those that merely adapt to it.

So how do we cement U.S. technology dominance, faster? Given the immense scope and complexity of the challenges facing our nation today, maintaining the status quo is not just insufficient—it poses a national security risk. Ensuring our continued technological supremacy demands a fundamental shift in how our government collaborates with industry to drive innovation in areas critical to the global balance of power. Other nations—particularly the People's Republic of China (PRC)—are investing massive resources into leapfrogging our technological capabilities in these key areas. It is crucial that we respond with an all-of-nation approach. Only by working together in new ways can we generate the speed and momentum needed to maintain our position as the world's leading technology power.

Early Adoption Drives Acceleration

Consider quantum computing. We're in a global race to develop quantum, with implications for everything from national security to space exploration to drug discovery. The winner of this race will be the one who moves fastest from theory to practice, from laboratory to real-world application. And the losers will stand to lose more than just national pride. The PRC is likely gathering encrypted data, waiting for the day when quantum computers can break through current encryption methods and wreak havoc on our systems. The risk of the PRC being able to break our cybersecurity defenses faster than we can breach theirs cannot be overstated.

Faced with such a threat, the private and public sectors need to work together in new ways. While industry must continue to lead the development of quantum and other advanced technologies, the U.S. government can step into the role of early adopter and in doing so, further accelerate progress.

Technology firms working at the leading edge cannot succeed without early adopters. These are the customers willing to tolerate bugs and not-quite-there-yet products. By way of reward, they get to imagine how the new tech will help them ahead of the market, and through their feedback, they shape the direction of the technology. Remember the first iPhone? Revolutionary, yet far from perfect. If everyone had waited for perfection, we might still be using BlackBerrys.

“To stay in front, [America] must operate at the speed of change—where solutions evolve as quickly as challenges and opportunities emerge.”

The government, with its scaled missions and long-term perspective, can act as a true catalyst through this kind of early adoption. A willingness and commitment to co-develop requirements and test, iterate, and evolve emerging technologies helps private industry advance through the early stages of commercialization, fueling acceleration in ways that increased funding alone cannot.

A New Model for Public-Private Collaboration

Today, the federal technology ecosystem is not designed to support this iterative and accelerated adoption of new technical solutions. For the government to truly embrace its role as an early adopter of technologies, it must first make a multitude of changes—like acquisition reform and increasing accountability for outcomes—that will enable it to operate with greater speed. Current processes are designed to minimize risk, are biased toward caution, and move far too slowly to keep pace with today’s rapid tech cycles.

Industry, in turn, must see government not just as a large late-cycle customer, but as a potential testbed for pushing the boundaries of what's possible. Through early adoption, the federal government would earn a seat at the table to define priorities and shape the outcomes that will ultimately safeguard our national security. The knowledge gained would also shape regulation ahead of the problems, not after. In parallel, the private sector also must see each other as more than competitors, especially when it comes to national security. Instead, we must be co-investors and co-creators working together to deliver the solutions the nation requires. When we get these partnerships right, everyone wins: government missions succeed faster, American innovation roars ahead, and our national economic interests are bolstered.

At Booz Allen, we’re already walking the walk. By virtue of the work we do, we are the advanced technology company focused on speed to outcomes for the U.S. government. At the center of the federal tech ecosystem, we have relationships with everyone from Silicon Valley startups and tech hyperscalers to the Pentagon, CDC, and NASA. We build our own tech on top of the best commercial offerings to accelerate meaningful improvement in how our government serves and protects America. And we are learning to work faster, internally and externally. We are building co-creation partnerships that are already on the verge of new breakthroughs. From AI to cyber to software-defined everything; and from the Pacific to orbit to here at home, we are building and working with the very best in our nation, across government and the private sector. The energy being created is awe-inspiring.

This approach, with industry co-investing and co-developing solutions and government participating as an early adopter, is a key element of the broader blueprint for continued American technological supremacy. To stay in front, we must operate at the speed of change—where solutions evolve as quickly as challenges and opportunities emerge. Doing so requires more than just brilliant innovation; it will require all of us to align our efforts. Together, we will go far and fast.

REFERENCES

06. Reflections on Generative AI

Alison Nathan, Jenny Grimberg, and Ashley Rhodes, “Gen AI: Too Much Spend, Too Little Benefit?” Top of Mind 129, June 25, 2024, https://www.goldmansachs.com/images/migrated/insights/pages/gs-research/gen-ai--too-much-spend,-too-little-benefit-/TOM_AI%202.0_ForRedaction.pdf.

Brandon Vigliarolo,“Study Uncovers Presence of CSAM in Popular AI Training Dataset,” The A Register, December 20, 2023, https://www.theregister.com/2023/12/20/csam_laion_dataset/.

David Vergun, “DARPA Aims to Develop AI, Autonomy Applications Warfighters Can Trust,” DOD News, March 27, 2024, https://www.defense.gov/News/News-Stories/Article/Article/3722849/darpa-aims-to-develop-ai-autonomy-applications-warfighters-can-trust/"DARPA Aims to Develop AI, Autonomy Applications Warfighters Can Trust > U.S. Department of Defense > Defense Department News.

Ethan Mollick, “Something New: On OpenAI’s ‘Strawberry’ and Reasoning,” One Useful Thing, September 12, 2024, https://www.oneusefulthing.org/p/something-new-on-openais-strawberry.

“Introducing Computer Use, a New Claude 3.5 Sonnet, and Claude 3.5 Haiku,” Anthropic, updated November 4, 2024, https://www.anthropic.com/news/3-5-models-and-computer-use.

Jamie Sevilla, Tamay Besiroglu, Ben Cottier, Josh You, Edu Roldán, Pablo Villalobos, and Ege Erdil, “Can AI Scaling Continue Through 2030?” EpochAI, August 20, 2024, https://epoch.ai/blog/can-ai-scaling-continue-through-2030.

Jory Heckman, “DoD Builds AI Tool to Speed Up ‘Antiquated Process’ for Contract Writing,” Federal News Network, February 9, 2023, https://federalnewsnetwork.com/contracting/2023/02/dod-builds-ai-tool-to-speed-up-antiquated-process-for-contract-writing/.

Peter Grad, “Trick Prompts ChatGPT to Leak Private Data,” Tech Xplore, December 1, 2023, https://techxplore.com/news/2023-12-prompts-chatgpt-leak-private.html

Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha, “A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications,” arXiv, February 5, 2024, https://doi.org/10.48550/arXiv.2402.07927.

10. Achieving Real-Time Cyber Defense

“Cybercriminals Can’t Agree on GPTs,” SC Media, December 28, 2023, https://www.scworld.com/native/cybercriminals-cant-agree-on-gpts.

Zilong Lin, Jian Cui, Xiaojing Liao, and XiaoFeng Wang, “Malla: Demystifying Real-World Large Language Model Integrated Malicious Services,” arXiv preprint, August 19, 2024, https://arxiv.org/abs/2401.03315.

17. Taking the Ground Out of Ground Systems

“Accelerating Multi-INT Fusion for Intelligence Missions,” Booz Allen Hamilton, accessed December 2, 2024, https://www. boozallen.com/insights/intel/accelerating-multi-int-fusion-for-intelligence-missions.html.

“Accelerate Space Superiority with Open Data Platforms,” Booz Allen Hamilton, accessed December 2, 2024, https://www.boozallen.com/markets/space/accelerate-space-superiority-with-open-data-platforms.html.

Adam Stockley, *The Benefits of GSaaS for the Future of Space, KDC Resource*, February 10, 2023, https://www.kdcresource.com/insights-events/the-benefits-of-gsaas-for-the-future-of-space/.

“Advanced DevSecOps for Critical Missions,” Booz Allen Hamilton, accessed December 2, 2024, https://www.boozallen.com/expertise/digital-solutions/devsecops/devsecops-critical-mission-infographic.html.

“aiSSEMBLE: Enterprise-Scale AI,” Booz Allen Hamilton, accessed December 2, 2024, https://www.boozallen.com/expertise/analytics/deploy-artificial-intelligence-faster-with-aissemble.html.

Defense Standardization Program, *Modular Open Systems Approach (MOSA)*, Department of Defense, accessed December 2, 2024, https://www.dsp.dla.mil/Programs/MOSA/.

DoD Data Strategy, Department of Defense, September 30, 2020, https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/dod-data-strategy.pdf.

Eric Hoffman, *Simplify Space Systems Integration*, Booz Allen Hamilton, accessed December 2, 2024, https://www.boozallen.com/insights/space/simplify-space-systems-integration.html.

Jerome Dunn, *What ‘Network-Centric’ to ‘Data-Centric’ Really Means*, Booz Allen Hamilton, accessed December 2, 2024, https://www.boozallen.com/insights/defense/defense-leader-perspectives/what-network-centric-to-data-centric-really-means.html.

Josh Perrius, *From the Ground Up*, Booz Allen Hamilton, accessed December 2, 2024, https://www.boozallen.com/insights/space/from-the-ground-up.html.

“Large Language Models Explained,” NVIDIA, accessed December 2, 2024, https://www.nvidia.com/en-us/glossary/large-language-models/.

Sandra Erwin, “Pentagon Embracing SpaceX’s Starshield for Future Military Satcom,” SpaceNews, June 11, 2024, https://spacenews.com/pentagon-embracing-spacexs-starshield-for-future-military-satcom/.

Sandra Erwin, “The Satellite Breakup: Military’s Push to Go Small,” SpaceNews, June 5, 2024, https://spacenews.com/the-satellite-breakup-militarys-push-to-go-small/.

"Commercial Satellite Industry Continues Historic Growth While Dominating Global Space Business—SIA Releases 27th Annual State of the Satellite Industry Report” news release, June 13, 2024, https://sia.org/commercial-satellite-industry-continues-historic-growth-dominating-global-space-business-27th-annual-state-of-the-satellite-industry-report/.

Space Policy Review and Strategy on Protection of Satellites, Department of Defense, September 20, 2023, https://media.defense.gov/2023/Sep/14/2003301146/-1/-1/0/comprehensive-report-for-release.pdf.

“Speeding Intelligence Insights Across Domains,” Booz Allen Hamilton, accessed December 2, 2024, https://www.boozallen.com/insights/intel/speeding-intelligence-insights-across-domains.html.

Transform Space Ground Systems Faster: 3 Ways to Add Flexible Capabilities and Control Costs, Booz Allen Hamilton, 2024, https://www.boozallen.com/markets/space/building-secure-adaptable-space-systems.html.

United States Space Force, “CSO Speaks on Logic of Space Superiority at Mitchell Institute,” news release, March 29, 2024, https://www.spaceforce.mil/News/Article-Display/Article/3725237/.

“Zero Trust, More Confidence,” Booz Allen Hamilton, accessed December 2, 2024, https://www.boozallen.com/expertise/cybersecurity/zero-trust-architecture.html?origref=&vURL=/zerotrust.

22. Sensemaking Reimagined

“The ADAPT+C Framework,” Booz Allen Hamilton, accessed October 29, 2024, https://www.boozallen.com/s/solution/the-adapt-c-framework.html.

Brian Katz, *Maintaining the Intelligence Edge: Reimagining and Reinventing Intelligence through Innovation*, Center for Strategic and International Studies, January 13, 2021, https://www.csis.org/analysis/maintaining-intelligence-edge-reimagining-and-reinventing-intelligence-through-innovation.

Carmen Medina and Rebecca Fisher, “What the World Economic Crisis Should Teach Us,” Studies in Intelligence 53, no. 3 (September 2009): 11–16, https://www.cia.gov/resources/csi/studies-in-intelligence/volume-53-no-3/what-the-world-economic-crisis-should-teach-us/.

Defense Innovation Board, *Aligning Incentives to Drive Faster Tech Adoption*, accessed October 29. 2024, https://innovation.defense.gov/Portals/63/20240701%20DIB%20Report_Aligning%20Incentives%20PUBLISHED%20STUDY.pdf.

Department of Defense, *Summary of the Joint All-Domain Command and Control (JADC2) Strategy*, March 2022, https://media.defense.gov/2022/Mar/17/2002958406/-1/-1/1/summary-of-the-joint-all-domain-command-and-control-strategy.pdf.

“Enabling Mission-Critical Machine Learning,” Booz Allen Hamilton, accessed October 29, 2024, https://www.boozallen.com/d/insight/thought-leadership/enabling-mission-critical-machine-learning.html.

“ICArUS: Integrated Cognitive-Neuroscience Architectures for Understanding Sensemaking,” Intelligence Advanced Research Projects Activity, accessed October 29, 2024, https://www.iarpa.gov/research-programs/icarus.

Joseph W. Gartin, “Looking Ahead: The Future of Analysis,” Studies in Intelligence 63, no. 2 (June 2019): 1–6, https://www.cia.gov/resources/csi/static/Future-of-Analysis.pdf.

National Intelligence Council, *Global Trends: Paradox of Progress*, January 2017, https://www.dni.gov/files/documents/nic/GT-Full-Report.pdf.

Office of the Director of National Intelligence, *The AIM Initiative: A Strategy for Augmenting Intelligence Using Machines*, accessed December 2, 2024, https://www.dni.gov/files/ODNI/documents/AIM-Strategy.pdf.

Office of the Director of National Intelligence, *Terms & Definitions of Interest for DoD Counterintelligence Professionals*, May 2, 2011, https://www.dni.gov/files/NCSC/documents/ci/CI_Glossary.pdf.

“Quotes: Russia,” The International Churchill Society, accessed June 24, 2023, https://winstonchurchill.org/resources/quotes/russia-2/.

William Burns, “Fireside Chat with William Burns: Aspen Security Forum 2023,” interview by Mary Louise Kelly, July 20, 2023, posted by Central Intelligence Agency, https://www.cia.gov/static/598a62b34629a8120fb16d68e440aa15/Director_Burns_Aspen_Security_Forum_Transcript_07202023.pdf.

32. Tech at Full Speed

Horacio Rozanski, “AI& Everything: A Future of Limitless Possibilities,” Velocity, 2023, https://velocity.boozallen.com/view/776121111/82-83/.

Rhiannon Williams, “Generative AI Taught a Robot Dog to Scramble Around a New Environment,” MIT Technology *Review*, November 12, 2024, https://www.technologyreview.com/2024/11/12/1106811/generative-ai-taught-a-robot-dog-to-scramble-around-a-new-environment/.

38. In AI Learning, One Size Fits None

“Fact Sheet: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence,” The White House," October 30, 2023, https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/.

Mauro Cazzaniga, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus J. Panton, Carlo Pizzinelli, Emma J. Rockall, and Marina Mendes Tavares, *Gen-AI: Artificial Intelligence and the Future of Work*, International Monetary Fund, January 14, 2024, https://doi.org/10.5089/9798400262548.006.

Michelle Cao and Rafi Goldberg, “Switched Off: Why Are One in Five U.S. Households Not Online?” National Telecommunications and Information Administration, October 5, 2022, https://www.ntia.gov/blog/2022/switched-why-are-one-five-us-households-not-online.

41. The Age of Agentic AI

“The Government Is Using AI to Better Serve the Public,” AI.gov, accessed October 16, 2024, https://ai.gov/ai-use-cases/.

Sayash Kapoor, Benedikt Stroebl, Zachary S. Siegel, Nitya Nadgir, and Arvind Narayanan, “AI Agents That Matter,” arXiv, July 2, 2024, https://arxiv.org/pdf/2407.01502.

46. Where Algorithms Meet Accountability

Ali Farzanehfar, Florimond Houssiau, and Yves-Alexandre de Montjoye, “The Risk of Re-identification Remains High Even in Country-Scale Location Datasets,” Patterns 2, no. 3 (March 12, 2021): 100204, https://doi.org/10.1016/j.patter.2021.100204.

Cynthia Dwork and Aaron Roth, “The Algorithmic Foundations of Differential Privacy,” Foundations and Trends in Theoretical Computer Science 9, nos. 3–4 (2014): 211–407, https://dx.doi.org/10.1561/0400000042.

Jeffrey Mervis, “The U.S. Has a New Way to Mask Census Data in the Name of Privacy. How Does It Affect Accuracy?” Science, May 3, 2024, https://doi.org/10.1126/science.zbp0e8r.

John M. Abowd, Robert Ashmead, Ryan Cumings-Menon, Simson Garfinkel, Micah Heineck, Christine Heiss, Robert Johns, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, Brett Moran, William Sexton, Matthew Spence, and Pavel Zhuravlev, “The 2020 Census Disclosure Avoidance System TopDown Algorithm,” Harvard Data Science Review, Special Issue 2 (2022), https://doi.org/10.1162/99608f92.529e3cb9.

Joseph Near and David Darais, *Guidelines for Evaluating Differential Privacy Guarantees*, National Institute of Standards and Technology, December 11, 2023, https://csrc.nist.gov/pubs/sp/800/226/ipd.

Lindsey Choo, “How 2 Students Used the Meta Ray-Bans to Access Personal Information,” Forbes, updated October 4, 2024, https://www.forbes.com/sites/lindseychoo/2024/10/04/meta-ray-bans-ai-privacy-surveillance/.

Michelle Faverio, “Key Findings about Americans and Data Privacy,” Pew Research Center, October 18, 2023, https://www.pewresearch.org/short-reads/2023/10/18/key-findings-about-americans-and-data-privacy/.

National Institute of Standards and Technology, “Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence: Differential Privacy,” accessed December. 2, 2024, https://www.nist.gov/artificial-intelligence/executive-order-safe-secure-and-trustworthy-artificial-intelligence-0.

“Protecting Personal Privacy,” U.S. Government Accountability Office, accessed December. 2, 2024, https://www.gao.gov/protecting-personal-privacy.

Skynova, “Data Collection: A Business’s Best Friend”, accessed December 2, 2024, https://www.skynova.com/blog/small-business-big-data.

Exec. Order No. 14110, 88 Fed. Reg. 75191 (October 30, 2023), https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence.

52. From the Frontlines of Post-Quantum Cryptography

“Announcing Approval of Three Federal Information Processing Standards (FIPS) for Post-Quantum Cryptography,” National Institute of Standards and Technology, updated August 26, 2024, https://csrc.nist.gov/News/2024/postquantum-cryptography-fips-approved.

Executive Office of the President, Office of Management and Budget, “Migrating to Post-Quantum Cryptography,” November 18, 2022, https://www.whitehouse.gov/wp-content/uploads/2022/11/M-23-02-M-Memo-on-Migrating-to-Post-Quantum-Cryptography.pdf.

Executive Office of the President of the United States, *Report on Post-Quantum Cryptography*, July 2024, https://www.whitehouse.gov/wp-content/uploads/2024/07/ref_pqc-report_final_send.pdf.

Jordan Kenyon and J.D. Dulny, “In the Quantum Era, Cybersecurity is a Race Against the Clock,” Federal News Network, July 8, 2022, https://federalnewsnetwork.com/commentary/2022/07/in-the-quantum-era-cybersecurity-is-a-race-against-the-lock/.

Jordan Kenyon and Taylor Brady, “5 Steps for Implementing the New Post-Quantum Cryptography Standards,” Information Week, September 5, 2024, https://www.informationweek.com/cyber-resilience/5-steps-for-implementing-the-new-post-quantum-cryptography-standards.

Jordan Kenyon and Taylor Brady, “Cybersecurity in the Quantum Risk Era,” Booz Allen Hamilton, accessed December 2, 2024, https://www.boozallen.com/insights/ai-research/cybersecurity-in-the-quantum-risk-era.html.

Quantum Computing Cybersecurity Preparedness Act, H.R. 7535, 117th Cong. (2022(enacted, https://www.congress. gov/bill/117th-congress/house-bill/7535/text.

The White House, “National Security Memorandum on Promoting United States Leadership in Quantum Computing While Mitigating Risks to Vulnerable Cryptographic Systems,” news release, May 4, 2022, https://www.whitehouse.gov/briefing-room/statements-releases/2022/05/04/national-security-memorandum-on-promoting-united-states-leadership-in-quantum-computing-while-mitigating-risks-to-vulnerable-cryptographic-systems/.

62. Intelligent Systems at the Edge

Defense Information Systems Agency Cyber, Fiscal Year 2025 Budget Estimates, March 2024, https://comptroller.defense.gov/Portals/45/Documents/defbudget/FY2025/budget_justification/pdfs/01_operation_and_maintenance/o_m_vol_1_part_1/disa_cyber_op-5.pdf.

66. Another Set of Eyes

Kristine Steen-Tveit and Bjørn Erik Munkvold, “From Common Operational Picture to Common Situational Understanding: An Analysis Based on Practitioner Perspectives,” *Safety Science* 142 (2021, https://doi.org/10.1016/j.ssci.2021.105381.

Stephen Owsinski, “Artificial Intelligence in Policing Rewrites Report Writing,” National Police Association, accessed November 6, 2024, https://nationalpolice.org/main/artificial-intelligence-in-policing-rewrites-report-writing/.

Thomas J. Cowper and Michael E. Buerger. *Improving Our View of the World: Police and Augmented Reality* Technology. PFI/FBI Futures Working Group, 2023, https://www.ojp.gov/ncjrs/virtual-library/abstracts/improving-our-view-world-police-and-augmented-reality-technology.

Tom McKay, “Here’s How Federal Law Enforcement Officials Are Actually Using AI,” IT Brew, June 13, 2024, https://www.itbrew.com/stories/2024/06/13/here-s-how-federal-law-enforcement-officials-are-actually-using-ai.

Victoria De Cardenas, "Palm Beach Co. Sheriff’s Office Invests $1.3 Million in AI to Streamline Police Reporting," WPEC CBS 12, August 2, 2024, https://cbs12.com/news/local/palm-beach-co-sheriffs-office-invests-13-million-in-ai-to-streamline-police-reporting.

72. Catalyzing America's Innovation Power

Chinese Threats in the Quantum Er a , Booz Allen Hamilton, December2021, https://www.boozallen.com/expertise/analytics/quantum-computing/chinese-cyber-threats-in-the-quantum-era.html.

O’Neill, Patrick Howell, “The US is worried that hackers are stealing data today so quantum computers can crack it in a decade,” MIT Technology Review, November 3, 2021, https://www.technologyreview.com/2021/11/03/1039171/hackers-quantum-computers-us-homeland-security-cryptography/.

BOOZ ALLEN AUTHOR INFORMATION

Special thanks to all of the leaders from within Booz Allen who contributed to this collection and provided technical and mission expertise along the way.

Derek Aucoin
aucoin_derek@bah.com

Marissa Beall
beall_marissa@bah.com

Taylor Brady
brady_taylor@bah.com

Todd Burnett
burnett_todd@bah.com

Christopher Castelli
castelli_christopher@bah.com

Ginny Cevasco
cevasco_virginia@bah.com

Josh Conway
conway_joshua@bah.com

Matt Costello
costello_matthew@bah.com

Joel Dillon
dillon_joel@bah.com

Carl Ghattas
ghattas_carl@bah.com

Joe Gillespie
gillespie_joseph@bah.com

Sean Guillory
guillory_sean@bah.com

Jim Hemgen
hemgen_james@bah.com

Rita Ismaylov
ismaylov_rita@bah.com

Jordan Kenyon
kenyon_jordan@bah.com

Todd Kline
kline_todd@bah.com

John Larson
larson_john@bah.com

Cameron Mayer
mayer_cameron@bah.com

Thong Nguyen
nguyen_thong@bah.com

Josh Perrius
perrius_josh@bah.com

Don Polaski
polaski_donald@bah.com

Edward Raff
raff_edward@bah.com

Joe Rohner
rohner_joseph@bah.com

Horacio Rozanski
rozanski_horacio@bah.com

Mike Saxton
saxton_michael@bah.com

Tony Sharp
sharp_canthony@bah.com

Alison Smith
smith_alison@bah.com

Ernest Sohn
sohn_ernest@bah.com

Eric Syphard
syphard_eric@bah.com

Bill Vass
vass_william@bah.com

Max Wragan
wragan_max@bah.com

Randy Yamada
yamada_randy@bah.com

VELOCITY, A BOOZ ALLEN PUBLICATION

CHIEF TECHNOLOGY OFFICER
Bill Vass

OFFICE OF THE CTO
John Larson
Julie McPherson

EDITORIAL
Executive Editor: John Larson
Managing Editor: Kevin Craft
Editorial Strategy: Logan Gibson and Danielle Epting
Editorial Consultants: Elana Akman, Drumbeat Editorial; Emily Primeaux, Dragonfly Editorial

CONTRIBUTING EDITORS
Dawn Dicker, Pat Fitzgerald, and Jim McDermott

NOTABLE CONTRIBUTIONS

DESIGN
Creative Direction: Chris Walker
Art Direction and Design: Brody Rose and David Everett

MARKETING AND COMMUNICATIONS
Head of Brand and Marketing: Abby Vaughan
Media Director: Jessica Klenk

For media inquiries, please contact:
Amanda Allison: allison_amanda@bah.com

COMMUNICATIONS AND DIGITAL
Special acknowledgements to individuals across content, digital, creative, brand, media, legal, website, accessibility, and project management for their support of this project.

ABOUT BOOZ ALLEN

Booz Allen is an advanced technology company delivering outcomes with speed for America's most critical defense, civil, and national security priorities. We build technology solutions using AI, cyber, and other cutting-edge technologies to advance and protect the nation and its citizens. By focusing on outcomes, we enable our people, clients, and their missions to succeed—accelerating the nation to realize our purpose: **Empower People to Change the World®**.

To learn more, visit BoozAllen.com.

Booz
Allen®

Velocity, a Booz Allen publication, studies the complex issues that are emerging for mission and technology leaders on the front lines of government innovation.

BoozAllen.com/Velocity