BY Booz Allen Insights for Federal Innovators V4 2025

From Sim to Field

Physical Al That Trains, Tests, and Proves Systems



FROM THE CHIEF TECHNOLOGY OFFICE

here's no stopping the **technology flywheel** that I wrote about in the previous issue of *Velocity*. A continuous cycle of real-world data collection, synthetic data generation, simulation, model training and optimization, and system deployment is profoundly changing how we operate—and we're experiencing a fundamental shift from temporal processing to a live-streaming mission.

But as chief technology officer, the flywheel of change isn't the only thing on my mind.

As I reflect on Booz Allen's impact this year, I come back to the big things that happened because of purposeful convergence in areas like hypersonics, additive manufacturing, and the use of visual language and world models to propel spatial computing and physical Al. Pressing missions called for the creative blending of capabilities and domains that were otherwise separate and distinct.

In this issue, we explore the power behind convergence: Why missions of national priority depend on technical fusion for speed, value, and outcomes. Our cover story spotlights key efforts across government sectors to architect Software-Defined Everything for frictionless transitions between virtual and physical environments (page 24).

What else becomes possible by connecting the right dots? How will technical convergence change the game for border security, critical infrastructure, space dominance, and more? And how can federal enterprises use contemporary environments to unleash convergence at scale? I invite you to engage with the range of perspectives in this issue, rooted in Booz Allen's ongoing work to navigate these questions with our customers. You'll find practical analysis and insights that include:

- What's next in software-defined autonomy, with a co-founder of Shield AI (page 4).
- How silicon photonics are quietly reshaping geopolitics and accelerating Al across mission domains (page 10).
- A look ahead at how Al at the edge will change what's possible in field operations (page 16).
- Why **agentic developers** are the future of software development and what this means for the core systems driving government (page 34).

I'm excited about this lineup because we showcase cutting-edge technologies that, frankly, I couldn't have imagined early in my career. But even more so, because each chapter points to a tenet that I believe will accelerate momentum for us all: that one plus one equals three.

As our chief financial officer, Matt Calderone, shares in this issue's closing comments: "America needs to use new technology to solve our collective challenges and ensure our leadership on the global stage."

As the federal innovation ecosystem continues to advance software-defined missions from the enterprise to the edge, let's find opportunities to break more technical silos, think broadly across development domains, and accelerate the modern flywheel through meaningful convergence.

A final note: This is the debut of a new format for *Velocity*, as we are now publishing several issues annually to keep pace with the tremendous changes in technology. I hope you find it valuable and insightful.

Bill Vass

Chief Technology Officer Booz Allen

Contents

01

FROM THE CHIEF TECHNOLOGY OFFICE

Convergence Is Everywhere

Bill Vass

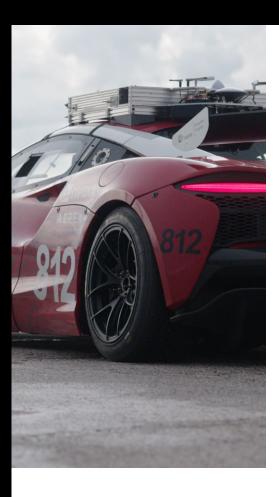
04

NEWSMAKER

Autonomy at the Tactical Edge

In conversation with Brandon Tseng Randy Yamada, Ph.D.





24

COVER STORY

From Sim to Field: Physical Al Trains, Tests, and Proves Systems

Conquering the trust gap for AI systems deployed in the real world

Munjeet Singh, Greg Kacprzynski, Randy Yamada, Ph.D., and Drew Massey 34

MISSION SPOTLIGHT

Agentic Development Decoded: Hard-Won Lessons From the Trenches

How federal IT teams can compress the software development lifecycle, accelerate delivery, and break the legacy transformation logiam

Bassel Haidar, Ryan Stone, Jennifer Sheppard, and Michael Martoccia 10

TECH WATCH

Traveling Light

Silicon photonics and the fight for device dominance Meghan Hauser, Ph.D.

and David Perry

16

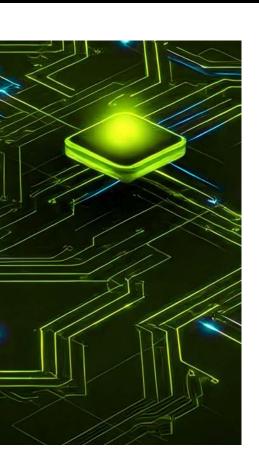
EMERGING TRENDS

AI-RAN and the Race to 6G

Driving missions at the edge with commercial innovation

Chris Christou, Jennifer Congdon, and Andrew Bolka



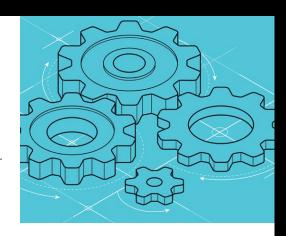


44

LESSONS FROM THE EDGE

How We Built a Better Innovation Factory

Haluk Saker



50

FINAL THOUGHTS

Accelerating
Mission
Outcomes Through
Tech Investment

Matt Calderone





Autonomy at the Tactical Edge

In Conversation With Brandon Tseng, Co-founder and President of Shield Al

ass has always been a defining principle of warfare. But what happens when mass is no longer limited by the number of available pilots? Advances in autonomy are allowing militaries to confront challenges like pilot shortages and increasingly contested environments. Swarms of intelligent, collaborative systems are moving from concept to reality, promising new levels of operational scale and greater safety for the warfighters who carry out those missions in the most dangerous environments.

At the center of this transformation is Brandon Tseng, co-founder and president of Shield AI. Drawing on his experience as a Navy SEAL, Tseng launched the company in 2015 with a singular mission: protect service members and civilians with AI-powered autonomy. Today, Shield AI is pioneering technology that can fly and fight in denied environments, scale from single aircraft to swarms, and—through its partnership with Booz Allen—help the U.S. field the intelligent mass needed to compete with peer adversaries.

In this conversation with Dr. Randy Yamada, Ph.D., Booz Allen's vice president for autonomy, Tseng discusses the origins of Shield AI, the bold decisions that shaped its growth, and the future of AI-powered warfighting.



You co-founded Shield AI to protect service members and civilians with artificially intelligent systems. Tell us more about what problem these systems solve for the warfighter.

Since our founding in 2015, that has been our mission, and it comes from my time in the U.S. Navy SEAL teams. On deployments from Afghanistan to the Pacific to the Persian Gulf, I saw firsthand the need for this capability. At Shield AI, we are building the world's best AI pilot. The easiest way to think about an AI pilot is self driving technology. In a military context, it enables unmanned systems to operate without GPS, without communications, and without a remote pilot. It also unlocks the concept of "teaming" or "swarming," which, with mass being a fundamental concept of warfare, introduces the warfighting concept of unlimited mass.

Was there a defining moment when you realized that scaling this software wouldn't just improve individual missions, but could fundamentally transform the way wars are engaged?

The first problem I wanted to solve was close quarters combat. As a Navy SEAL, that's one of the most dangerous missions we conducted: going room to room inside buildings. It looks exciting on TV, but in reality, there's nothing more terrifying than being in a gunfight in such tight spaces. I began talking to people about putting self driving technology onto a quadcopter to help with those missions.

One of those conversations was with a Naval Academy classmate of mine—a mechanical engineer and F 18 pilot. He told me that the problem I was trying to solve with quadcopters indoors was the same one they were facing in the air warfare community. When you fly a drone into a building, you lose GPS and communications. At that time, during operations in Syria, our forces were beginning to experience the same thing on a larger scale: jamming that disabled GPS and communications, grounding Predators and Reapers.

He explained that fighter pilots were now facing another threat: proliferated, integrated air defense systems. Against those systems, survivability rates were extremely low. You don't send pilots into those areas because they simply won't come back. He told me the odds of survival were similar to stepping on an IED. That was my "aha" moment: realizing that the autonomy we were developing for small quadcopters could also solve a life or death problem for advanced aircraft in contested environments.

Building a company like Shield AI is no small task, and timing, technology, and opportunity all play a role. Looking back, what helped you take Shield AI from an early-stage startup to a company with more than 1,000 employees today? And what advice would you offer entrants trying to cross that so-called "Valley of Death" between prototype and program of record?

I'd break it into two parts. First, get ready for the pain train. Building a company is full of suffering, sleepless nights, and stress in every direction. I often tell veterans who start companies that one advantage they bring is knowing how to endure hardship and still accomplish the mission. That mindset helps.

Second, you have to be bold. Every entrepreneur will face moments when the right decision feels uncomfortable or even terrifying. At Shield Al, we've had to make several of those "bet the house" calls, and they've been critical to our growth. The reality is, if you're not making bold decisions and taking big risks, you're not really growing.

And what would you say was Shield Al's boldest bet?

Without question, it was acquiring Martin UAV and Heron Systems at the same time in 2021. That was an incredibly uncomfortable decision. My brother [Ryan Tseng] and I had talked about "climbing the aviation food chain," but I always assumed we'd do it organically. At the Series C stage, you didn't see many venture backed startups making acquisitions of that scale.

We ultimately convinced the board that to have the impact we wanted, we had to grow inorganically. Looking back, had we not made those two acquisitions, Shield Al's story in 2021 would have been very different and much less positive. Within three months, we had raised \$225 million and closed both deals. I'll never forget one of our board members saying, "Holy crap, I can't believe what just happened." Those acquisitions gave Shield Al the platform to continue growing into what it is today.

Another large part of that bold growth story is V-BAT, powered by Hivemind, which has been part of Shield Al's portfolio for several years. What set it apart years ago and what continues to differentiate it today?

Two things set V BAT apart. First is the hardware architecture. Its design creates a uniquely small logistics footprint and makes it easy to operate and transport. Features like the ducted fan give it advantages such as a payload to weight ratio roughly twice that of other vertical takeoff and landing aircraft. In practice, that means V BAT can carry the same payload while being half the size and weight, making it a much more agile and versatile platform.

The second differentiator is software. Our autonomy is what truly makes V BAT stand out. We've repeatedly demonstrated operations in GPS- and communications denied environments, proving mission autonomy on the platform. We've also shown multi agent mission autonomy, and we continue to expand those capabilities. Increasingly, software is becoming the defining factor for unmanned systems.





Hardware innovation is valuable, but it's hard to sustain differentiation there. Autonomy and mission software create a far greater competitive edge.

A good comparison is Tesla: The company disrupted hardware by committing to electric vehicles, but its real long term advantage is in software. If you can differentiate on both hardware and software, the effect is powerful. That's how we think about V BAT today.

Adoption is always a challenge for emerging defense technologies. What has it been like to see V-BAT move from demonstration into a program of record, and how have you watched it change the way missions are executed?

It can be a painful journey. You start with innovators in the Department of Defense (DOD) who see the potential and are willing to prove it in the field, but those early efforts don't always translate into scale because of the way the acquisition system works.

What really changed the conversation was Ukraine. The conflict revealed just how devastating electronic warfare can be for legacy systems—thousands of drones lost each month, and even advanced platforms like HIMARS [High Mobility Artillery Rocket System] seeing dramatically reduced effectiveness in jammed environments. When V BAT showed that it could operate for 10 plus hours, fly hundreds of miles into denied territory, and still deliver mission effectiveness, it created a global recognition that operations could be conducted differently.

That lesson is resonating in the U.S. After seeing expensive unmanned systems shot down by cheap missiles, the Army is rethinking its reliance on large, runway dependent aircraft. V BAT's small logistics footprint and distributed deployment model offer a fundamentally different way to execute ISR [intelligence, surveillance, and reconnaissance] and targeting missions. It's helping customers reimagine how to solve problems in contested environments, moving from a requirements driven acquisition mindset to one that starts with the operational problem and works backward.

Building adoption is one piece, but building trust is another. Beyond demonstrations and marketing, how do you ensure users have confidence in the system?

Trust doesn't come from a single demo. It comes from working hand in hand with users as they adopt and scale the system. At the end of the day, the question isn't "does the product work?" It's "is the customer achieving the outcomes they need?" That requires a lot more than delivering hardware or software it means standing alongside them to make sure they're successful.

That's why we forward deploy engineers, operators, and flight personnel. When our teams are embedded with users whether in Ukraine or with U.S. partners they're right next to the problem, helping solve it in real time. It shows we're in the trenches with them, adapting the technology to the mission, and ensuring it delivers results. That proximity builds trust.

Speed is also everything in today's environment, whether we're talking about technology competition or mission outcomes. I'd like to talk about your vision for an "autonomy factory." When you look at the market, there hasn't really been a professional-grade autonomy stack—it's been mostly open source and community-driven. What pushed you to invest in Hivemind as a product?

Strategically, we believed we had built the world's best Al pilot. Shield Al was the first company to put an Al pilot on an unmanned system and deploy it on the battlefield starting with a quadcopter in Afghanistan in 2018. We were also the first to execute denied environment operations with V BAT and to demonstrate long endurance ISR and targeting in Ukraine. More recently, we were finalists for the Collier Trophy after flying an F 16 autonomously.

So, the question became: How do we proliferate this capability? How do we put a million Al pilots in the hands of customers in the next 10 years, and 100 million in the next 20? We realized Shield Al couldn't do that alone. To scale, we had to enable the broader defense industrial base with the same developer tools, infrastructure, and pipelines we use internally to build autonomy. That's what led to Hivemind Enterprise.

I want to pivot for a moment to something we're hearing a lot about at the Pentagon: "commercial first." As a former SEAL, what does that concept mean to you?

To me, it reflects a recognition that more R&D is happening in the commercial sector than in the defense sector. The idea is to lean on commercial solutions that already exist or can be built rapidly, rather than standing up major development programs that take years to deliver. That's not to say large strategic systems like fighter jets or submarines won't require government investment. But for the vast majority of capabilities, commercial products are available or can be quickly adapted and we should bias toward those.

I often tell program officers: I'd rather you buy a product and let Shield AI invest our own R&D to adapt it, rather than fund us directly. When the government puts development dollars into a program, it tends to move ten times slower than what a company can do with its own dollars.

And when you were in uniform, did you already think that way? Would you rather buy something that wasn't perfect and figure out how to use it than wait three years for a requirements-driven solution?

Absolutely. As a SEAL, I was fortunate to work with an incredible range of aviation assets, from Predators to C 130s to attack helicopters. I rarely felt like I lacked capability. But whenever there was a gap, the mentality was never, "Let's spin up the acquisition system and wait." It was always, "What can we buy now to solve the problem in front of us?" That urgency has stuck with me.

Before we wrap up, I want to touch on Shield Al's partnership with Booz Allen. From your perspective, how has that collaboration accelerated the mission?

It's been fantastic working with Booz Allen. Randy, you've been a great partner and champion of what autonomy can mean for the world, and our visions are very much aligned.

One of the biggest realities we face is scale. Ukraine has openly talked about building a million drone army. China will do the same. Inevitably, the U.S. will have to field its own force of millions of autonomous systems. But you can't train enough human pilots to operate them. Even with today's limited fleet, there's already a pilot shortage. To unlock the power of intelligent mass, autonomy is the critical enabling technology.

That's where the partnership with Booz Allen matters. You bring deep engineering expertise, a clear eyed view of the technical challenges, and an ability to integrate solutions across the defense enterprise. Together, we're building the pathways to scale it across the DOD. That's the journey ahead, and Booz Allen has been a fantastic partner in making it real.

Lightning Round

One myth about autonomy you wish the Pentagon would retire this year.

That autonomy is easy. It's not. I'd put it on par with landing a rocket from outer space in terms of technical difficulty. It takes massive expertise, resourcing, and capital to get it right and very few companies can actually deliver.

One acquisition rule you would change tomorrow and the impact it would have on time to field.

I'd move from a requirements based system to a problem based one. Too often, novel solutions die in the "valley of death" simply because no requirement was written for them. If you start with the problem say, persistent ISR over a 100 mile area and let companies bring different solutions to the table, you'll get faster, more innovative outcomes.

Finish this sentence: "Three years from now, autonomy will be judged successful if it has _____."

Three years from now, autonomy will be judged successful if it has proliferated at scale because it has proven it can reliably achieve mission objectives on the toughest, most contested operational battlefields.

SPEED READ

Autonomous technologies from companies like Shield Al are revolutionizing warfare, enabling unmanned systems to operate in GPS- and communications-denied environments, increasing operational scale and enhancing warfighter safety.

Shield Al's rapid expansion to over a thousand employees underscores the impact of bold decisions, such as acquiring Martin UAV and Heron Systems, which accelerated its capabilities and growth in the defense tech industry.

Successful deployment of autonomous systems requires building trust through hands-on collaboration with military personnel, ensuring these technologies meet operational needs and proving their effectiveness in contested environments like Ukraine.

Counter China's Cyber Warfare and

SEIZE THE ADVANTAGE

How to outpace, outsmart, and disrupt the PRC cyber threat

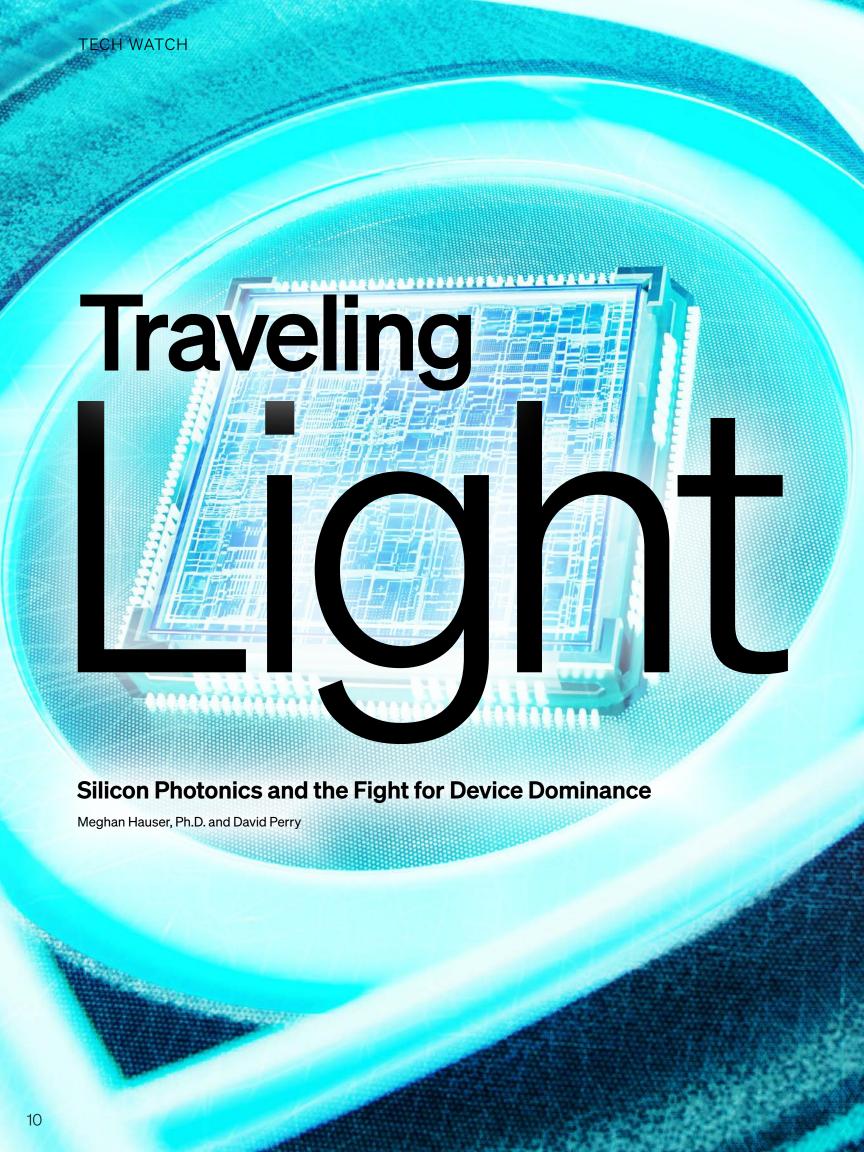


using AI to accelerate cyber operations and expand its geopolitical influence. In Booz Allen's latest report, we provide a strategic roadmap for America and its allies to overcome and prevail in cyberspace. Explore actionable insights and how the nation must move faster and innovate smarter to set the pace for future cyber operations in our publication, Breaking Through: How to Predict, Prevent, and Prevail Over the PRC Cyber Threat.

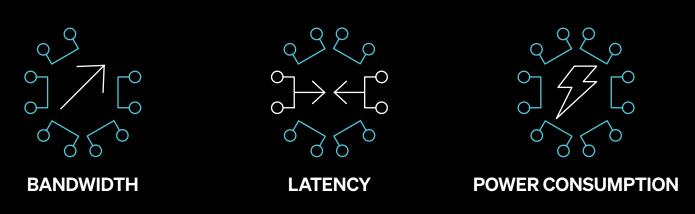
Read Now at BoozAllen.com/PRCreport

Booz Allen

<its_in_our_code>



odern computing, sensing, and communications are hitting an interconnect ceiling: Chips generate more data than copper wires can carry efficiently. As speeds climb, electrical links lose signal over short distances, require power-hungry cleanup circuits, and create heat and congestion at connectors. The result is that bandwidth, energy, and reach of the links now limit system performance more than the processors do. This sets the stage for silicon and integrated photonics (SiPh), which moves bits using light instead of electrons to break through those limits in new and scalable ways. This seemingly small hardware innovation changes everything we've known about:



Consider China's billion-dollar SiPh industry. From Peking University to Guangdong Province, Chinese national labs and corporations are operating at full speed to leapfrog advanced electronics.

In 2022, General Secretary Xi Jinping's remarks about SiPh were explicitly clear. It's an area, he said, "in which our country has the conditions to achieve breakthroughs ahead of others." Since then, China has made aggressive investments in this technology that is quietly reshaping the compute landscape—and becoming a focal point of U.S. national security and geopolitical advantage.

In the context of pressing U.S. defense and intelligence applications, it's no surprise that the U.S. House of Representatives' select committee on the Strategic Competition Between the United States and Chinese Communist Party warned in 2024 that "[s]ilicon photonics represents the next front in our semiconductor competition with the PRC, and the United States is currently not winning this competition." Early movers will lead in Al, quantum computing, sensing, and space-based systems—so active engagement today offers the chance to shape supply chain security and trusted hardware development in the face of adversarial investment.

Over the next 2 to 5 years, silicon and integrated photonics will become central to the future mission, with the power to enable underlying technology across virtually every national system. With analysis and insights

from the Booz Allen Tech Scouting and Intelligence team, we examine where the market trends are heading and immediate considerations for federal government leaders and the U.S. innovation industry.

Technology at a Glimpse

In 2024 research, our Tech Scouting and Intelligence team discussed how photonics is "an attractive new approach to data transfer—both on- and off-chip—given its added benefits of speed, bandwidth, and low heat generation." Since that original research was published, the SiPh market has grown rapidly and has shown the potential to transform a myriad of technologies as the Al industry demands increased compute performance at lower cost and higher efficiency.

SiPh enables the integration of optical components onto a single silicon (or silicon-based material like silicon nitride or silicon-on-insulator) chip leveraging standard complementary metal oxide semiconductor processes. It leverages existing semiconductor manufacturing techniques to simplify this evolution.

The resulting devices use light (photons) instead of electrical signals, which overcomes many limitations of traditional copper- and electron-based devices. These photonic integrated circuits (PIC) are highly versatile and can efficiently generate, manipulate, and transmit light at high speeds and low power, making them valuable across many applications.

Key Benefits of SiPh vs. Traditional Electronics

Data Rates

SiPh greatly expands the amount of information that can be moved at once, making it possible to handle the growing demands of modern computing and communications. It enables much higher data rates with wavelength division multiplexing a technique where multiple signals are transmitted simultaneously on a single waveguide by using different wavelengths (i.e., colors) of light. As a result, while high speed electrical interconnects have plateaued at data rates of hundreds of gigabits per second, SiPh is on track to scale up to tens of terabits per second by the end of this decade.

Latency

Electrical interconnects are constrained by resistive capacitive delays and require signals to be constantly boosted, converted, and cleaned up—all of which accumulates latency as data moves across a system. Photonic links avoid these bottlenecks by transmitting information optically, where waveguides experience negligible dispersion and do not require frequent regeneration. While the speed is still shaped by how fast the light signals can be turned on and off, SiPh reduces the overhead introduced by electronic signal processing, enabling lower end to end delay and faster communication across increasingly complex systems.

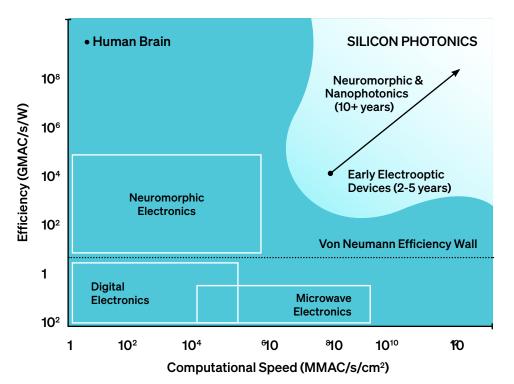
Signal Integrity/Interconnect Reach

Electrical interconnects lose their signal strength over increasingly long distances and suffer crosstalk and interference. On the other hand, photons don't suffer from resistive losses, parasitic effects, or electromagnetic interference, maintaining signal integrity over long distances. Optical interconnects also reduce signal propagation delays, as photons experience less dispersion compared with electrical signals.

Energy Efficiency

Heat matters: The more heat, the more wasted energy, cooling costs, stress on components and the less space there is to add processors and memory without the risk of melting. SiPh devices can be more energy efficient since using light, rather than electrical currents, reduces resistive heating and power loss. Additionally, photonic interconnects can achieve high speed data transfer with lower energy consumption by minimizing the need for power hungry repeaters and buffers.





SiPh goes beyond Moore's Law to break the efficiency wall that limits all-electronic devices.

Sources: Bhavin J. Shastri, Mitchell A. Nahmias, Alexander N. Tait, Thomas Ferreira de Lima, Hsuan-Tung Peng and Paul R. Prucnal. "Integrated Neuromorphic Photonics." Proceedings Volume 10721, Active Photonic Platforms X; 107211M (2018) https://doi.org/10.1117/12.2322182 (reproduced courtesy of SPIE), and Paul R. Prucnal and Bhavin J. Shastri. "Neuromorphic Photonics." CRC Press – Taylor & Francis Group (2017) https://doi.org/10.1201/9781315370590.

Spotlight Application: Al Acceleration

This technology can advance compute-intensive use cases across a number of domains. One of the most compelling is using SiPh to accelerate the performance of Al systems.

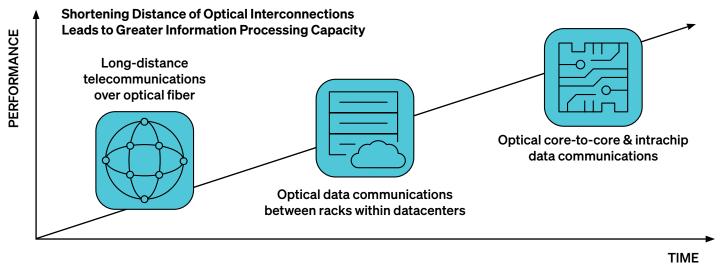
Challenge

Today's AI computations are fundamentally limited by the "memory wall," which is the gap between the amount of memory access needed and the amount available. This gap is growing exponentially as models become larger and more complex, but memory bandwidth is stagnating. While data center applications can address these limits through scale-up architectures, edge devices face strict limits on size, weight, and power (SWaP), meaning they cannot simply add more chips or energy-hungry cooling systems to keep up. These combined bottlenecks create mounting pressure for new approaches to move data faster and more efficiently within AI hardware.

Application

The bandwidth, signal integrity, and power advantages of SiPh directly address the memory wall by easing the flow of data between compute and memory, reducing one of the biggest bottlenecks in Al systems. By integrating on-chip optical interconnects, photonics also makes it possible to bring more compute power closer to edge sensors, enabling real-time Al inference in devices like drones, autonomous vehicles, and industrial robots.

Beyond raw bandwidth, the ability of photonics to deliver deterministic, low-latency communication supports tight synchronization across distributed sensor networks, which is critical for autonomy and situational awareness. In this setup, data is still processed electrically but transmitted optically, with PICs embedded alongside electronic integrated circuits to maximize reach, network bandwidth, and efficiency. This co-location not only supports memory disaggregation and composable Al architectures but also scale-out and scale-up systems that break the limits of traditional copper interconnects.



Dynamics We're Watching

#1: The geopolitical chess match is on

Export restrictions on lithography and advanced electronics have heavily contributed to China's SiPh indigenous development drive. Moving forward, China and the United States will compete for supremacy in the geopolitical microelectronic and SiPh race. China is aggressively investing in SiPh as part of its broader effort to secure self-sufficiency in advanced compute, sensing, and communications technologies.

China's supply chain dominance in germanium, lithium niobate, and other rare and critical metals strengthens their hold over the silicon wafer (SW) market, needed for SiPh, and accounts for 50% of global SW production. The top five global SW producers are all Chinese, including the Shanghai Xinkehui New Material Co.

#2: Private investment is on the rise

In 2024, private investment in SiPh grew rapidly after a year of protracted funding. The large decrease in capital allocation and stagnant deal volume in 2023 was likely due to macroeconomic factors and longer runways created from the previous two years of large, later-stage deals.

#3: U.S. federal priorities are gradually aligning

The U.S. government still maintains a majority investment in electronic-only technologies. But there are significant efforts to accelerate development of SiPh, given its heightening role in national security priorities. The Defense Advanced Research Projects Agency (DARPA) and the Defense Sciences Office are the major drivers of Department of Defense (DOD) SiPh funding, with DARPA being a top catalyst for SiPh research and development over the last two decades.

Over the coming years, the SWaP-C metrics of SiPh will improve sufficiently to accelerate edge Al computations and other critical edge applications.

LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) LiDAR for robotics, drones, and industrial sensing SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/ Advanced biomedical/	Clonger distance, between systems/facilities)		Near Term (0-2 Years)	Medium Term (2 5 Years)	Long Term (5+ Years)					
Spatial computing/cloud XR rendering Description of the properties of the propertie	Spatial computing/cloud XR rendering Description of the properties of the propertie	Spatial computing/cloud XR rendering Description of the properties of the propertie	Spatial computing/cloud XR rendering Description of the properties of the propertie	Spatial computing/cloud XR rendering Description of the properties of the propertie	Spatial computing/cloud XR rendering Description of the properties of the propertie	Spatial computing/cloud XR rendering Description of the properties of the propertie	ters	(longer distance, between	for data center/HPC AI	(rack-scale optics)
Spatial computing/cloud XR rendering Ditical networks-on-chip (NoC) LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul Telecom/	Spatial computing/cloud XR rendering Ditical networks-on-chip (NoC) LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul Telecom/	Spatial computing/cloud XR rendering Ditical networks-on-chip (NoC) LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul Telecom/	Spatial computing/cloud XR rendering Ditical networks-on-chip (NoC) LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul Telecom/	Spatial computing/cloud XR rendering Ditical networks-on-chip (NoC) LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul Telecom/	Spatial computing/cloud XR rendering Ditical networks-on-chip (NoC) LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul Telecom/	Spatial computing/cloud XR rendering Ditical networks-on-chip (NoC) LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul Telecom/	ata Cen	Al and HPC interchip optical	Datacom between quantum	computing
LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networks Intrachip optical comms for edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms Optical networks-on-chip (NoC) Autonomous navigation with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networks Intrachip optical comms for edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms Optical networks-on-chip (NoC) Autonomous navigation with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networks Intrachip optical comms for edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms Optical networks-on-chip (NoC) Autonomous navigation with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networks Intrachip optical comms for edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms Optical networks-on-chip (NoC) Autonomous navigation with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networks Intrachip optical comms for edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms Optical networks-on-chip (NoC) Autonomous navigation with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networks Intrachip optical comms for edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms Optical networks-on-chip (NoC) Autonomous navigation with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	LiDAR for robotics, drones, and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networks Intrachip optical comms for edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms Optical networks-on-chip (NoC) Autonomous navigation with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	۵		Spatial computing/cloud	
and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference Telecom/long-haul networking and industrial sensing edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference Telecom/long-haul networking and industrial sensing edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference Telecom/long-haul networking and industrial sensing edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference Telecom/long-haul networking and industrial sensing edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference Telecom/long-haul networking and industrial sensing edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference Telecom/long-haul networking and industrial sensing edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks	and industrial sensing IoT with optical sensors Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference Telecom/long-haul networking and industrial sensing edge Al accelerators SG/6G edge comms (optical fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms with photonic sensor fusion EW with photonic front-ends Rad-hard computing devices for space XR glasses with PIC-based sensing and compute Distributed quantum networks			XR rendering	Optical networks-on-chip (NoC)
Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early biomedical wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Fronthaul) SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms				
wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul petworking SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms for space XR glasses with PIC-based sensing and compute Distributed quantum networks	wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul petworking SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms for space XR glasses with PIC-based sensing and compute Distributed quantum networks	wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul petworking SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms for space XR glasses with PIC-based sensing and compute Distributed quantum networks	wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul petworking SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms for space XR glasses with PIC-based sensing and compute Distributed quantum networks	wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul petworking SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms for space XR glasses with PIC-based sensing and compute Distributed quantum networks	wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul petworking SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms for space XR glasses with PIC-based sensing and compute Distributed quantum networks	wearables/diagnostics and environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul petworking SIGINT/ISR systems Space comms and sensing Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms for space XR glasses with PIC-based sensing and compute Distributed quantum networks	<u>•</u>	loT with optical sensors		EW with photonic front-ends
environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul	environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul	environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul	environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul	environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul	environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul	environmental sensors Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Telecom/long-haul	Edg			
Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Early quantum sensors (e.g., atomic clocks) Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Early quantum sensors (e.g., atomic clocks) Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Early quantum sensors (e.g., atomic clocks) Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Early quantum sensors (e.g., atomic clocks) Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Early quantum sensors (e.g., atomic clocks) Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Early quantum sensors (e.g., atomic clocks) Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	Early quantum sensors (e.g., atomic clocks) Secure optical links less prone to electromagnetic interference 5G comms Telecom/long-haul networking Early quantum sensors (e.g., atomic clocks) Advanced biomedical/environmental sensing Advanced quantum sensors Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms				
Secure optical links less prone to electromagnetic interference 5G comms	Secure optical links less prone to electromagnetic interference 5G comms	Secure optical links less prone to electromagnetic interference 5G comms	Secure optical links less prone to electromagnetic interference 5G comms	Secure optical links less prone to electromagnetic interference 5G comms	Secure optical links less prone to electromagnetic interference 5G comms	Secure optical links less prone to electromagnetic interference 5G comms	4		Advanced biomedical/	sensing and compute
interference 5G comms Telecom/long-haul networking Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	interference 5G comms Telecom/long-haul networking Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	interference 5G comms Telecom/long-haul networking Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	interference 5G comms Telecom/long-haul networking Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	interference 5G comms Telecom/long-haul networking Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	interference 5G comms Telecom/long-haul networking Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms	interference 5G comms Telecom/long-haul networking Integrated photonic transceivers for RF spectrum sensing, 6G, or optical comms				Diotributou quantum motivom
Telecom/long-haul sensing, 6G, or optical comms										
networking		5G comms								
								Telecom/long-haul	sensing, 6G, or optical comms	

What's Next for the Technology

Projected timeline of development

It's no accident that SiPh has reached this inflection point. Decades of research and development by both the private and public sectors, including more than 20 years of DARPA investment, have strategically accelerated SiPh applications. Already dominating data communications in the data center market, SiPh is now on the cusp of disrupting several other key fields in the next 2 to 5 years.

Reinforcing U.S. leadership

For long-term dominance in advanced computing, communications, space, and other technology verticals, the United States must prioritize the adoption of photonics and optoelectronics. They are the next generation of capabilities to break the limits of all-electronic solutions.

In the coming years, the federal government can leverage SiPh for:

- Secure, high-bandwidth communications that are both compact and resistant to electromagnetic interference.
- Tactical and strategic networks where low SWaP is critical, such as space environments.
- The foundational platform to manipulate quantum states of light, making them key enablers for scalable quantum systems.
- Support of next-generation sensors used in intelligence, surveillance, and reconnaissance (ISR) and electronic warfare. Applications include LiDAR, hyperspectral imaging, and quantum radio frequency sensors, among many others.

A building block for global competition

Silicon photonics will enable and transform a diverse array of industries by providing a scalable platform for manufacturing advanced devices. By marrying light waves with the ubiquitous semiconductor manufacturing used for modern electronics, SiPh technology overcomes traditional barriers in cost, performance, and scalability that have historically limited photonics adoption.

For continued progress, it will be crucial to:

- · Prioritize domestic R&D and manufacturing investment into photonic and optoelectronic technologies to ensure competitive positioning in SiPh for the United States and our allies.
- Explore supply chain alternatives for key SiPh components and critical mineral resources currently controlled by China.
- Keep U.S. capital and SiPh intellectual property from China, as with other advanced semiconductor technologies crucial to the Tech Cold War.

Today's leaders in emerging software solutions like Al must take hardware more seriously than before. Key innovations in SiPh design, materials, integration, packaging, and system-level design will unleash exponential growth in computing, sensing, quantum information technologies, telecommunications, and many more applications—some of which we have yet to discover.

SPEED READ

Silicon photonics (SiPh) is set to revolutionize data transfer, offering higher bandwidth, lower latency, and improved energy efficiency by addressing limitations in traditional copper and electronic interconnects.

China and the United States are competing for geopolitical dominance in silicon photonics technology, with China making significant investments and the United States recognizing its strategic importance for national security and advanced computing.

.....

The United States must prioritize SiPh adoption and innovation, with support for domestic R&D and manufacturing, to maintain leadership in advanced technologies such as Al, quantum computing, and secure communications.

About the research:

The Booz Allen Tech Scouting and Intelligence team scans the market for novel technologies with the potential to disrupt U.S. government critical missions. Delivering research, trends spotting, and deep dive technical analysis, the team is hyper-focused on getting the best technology from the private sector into the hands of the warfighter and civilian federal workforce.

Al-RAN and the Race to



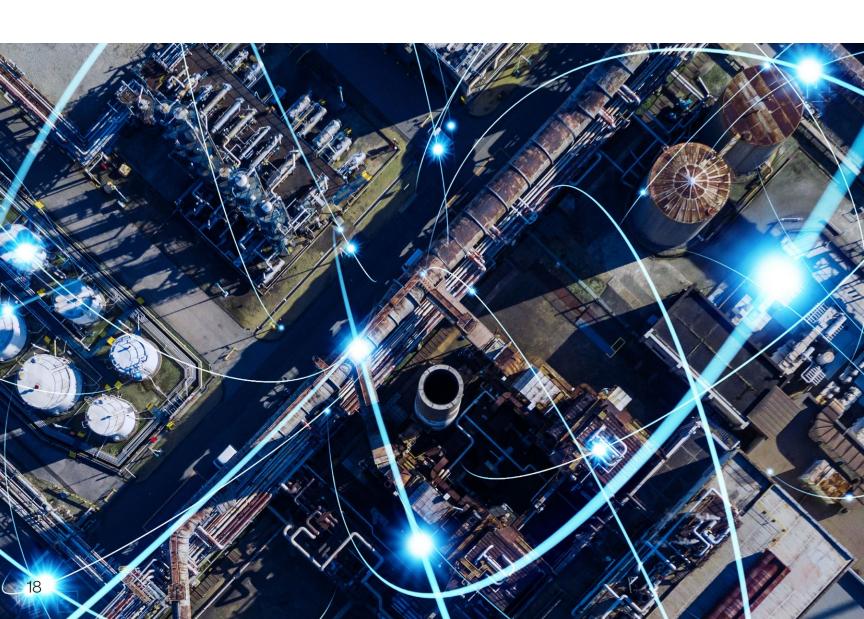


But this network operates differently. Within seconds of threat detection, machine learning (ML) algorithms running on the installation's private 5G infrastructure begin autonomously reconfiguring transmission parameters. The unit's network monitors the tactical environment and automatically adjusts power levels across base stations. The network's coverage area contracts and reshapes itself, maintaining some connectivity while reducing its electromagnetic signature enough to evade enemy sensors.

This scenario illustrates the emerging reality of artificial intelligence-radio access networks (AI-RAN), a technological breakthrough that transforms wireless networks from passive data pipes into intelligent, adaptive platforms. This evolution represents the technical foundation for future networks to unlock innovative new edge applications and devices.

What makes AI-RAN so powerful is its ability to solve two strategic challenges with the same technology. Specifically, AI for RAN optimizes network performance by automatically adjusting settings that normally require manual intervention, allowing base stations to detect problems and fix themselves without central control. AI on RAN enables AI applications to run directly on the same equipment as the RAN itself instead of on remote servers, reducing delays and enabling new capabilities in austere environments.

key technology for national security and economic competitiveness, with the United States and its allies committing to "support open, free, global, interoperable, reliable, resilient, and secure connectivity" as a means to promote stability and peace.



As the telecom industry maps sweeping transformations in the coming years, Al-RAN will also become the foundation of the future 6G stack. This means the geopolitical stakes for the development of these advanced networking capabilities could not be higher. 6G has become a key technology for national security and economic competitiveness, with the United States and its allies committing to "support open, free, global, interoperable, reliable, resilient, and secure connectivity" as a means to promote stability and peace.

While previous wireless generations were mostly static, 6G networks will be software- and Al-based and easier to rapidly reconfigure. The Al-RAN platform's architectural flexibility will enable nations with advantages in software and Al expertise to better access the innovation this critical infrastructure promises to unlock. Consequently, Al-RAN offers a strategic opportunity for the United States to reclaim leadership in the global telecom arena. By moving quickly and leveraging the disaggregated, software-defined nature of Al-RAN architectures, America can play a pivotal role in building the Al-native wireless standards and technologies that will define connectivity's next generation and compete more effectively in the global marketplace.

Why AI-RAN Matters Now

Several converging trends make 2025 pivotal for Al-RAN. Across the Department of Defense (DOD), the intelligence community, and civil agencies, growing needs for enhanced mobile broadband capabilities and more robust security have driven deployment of an initial set of private 5G networks. As adoption expands, agencies will identify new use cases spurring demands for more agile edge networks.

Executive Order 14320, Promoting the Export of the American AI Technology Stack; mandates within the National Defense Authorization Act around 5G; and national spectrum policy initiatives from the National Telecommunications and Information Administration (NTIA) point to a federal push to secure U.S. telecom technology stacks. These moves underscore current urgency and create a framework to guide AI-native 6G network deployments.

Al-RAN represents a broader shift toward greater ecosystem-driven innovation. Multiple component providers can participate through open application programming interfaces (APIs), enabling network apps and edge Al applications, similar to how mobile app ecosystems transformed consumer technology. In addition, the rapid expansion of the Al-RAN Alliance to more than 100 members signals industry alignment around a common vision for this innovation and change. The coalition includes technology companies and academic institutions all working toward the integration of Al with RAN, with the goal of advancing RAN performance and enabling new capabilities and use cases.



America's Al-RAN Opportunity

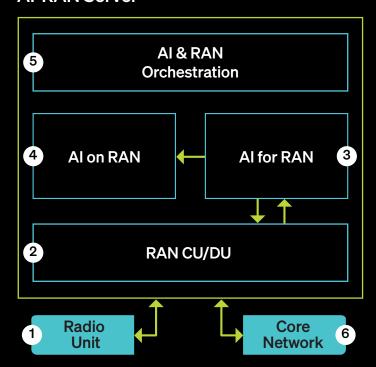
China has significant mobile network market share globally, from 5G infrastructure to handset manufacturing. This position extends beyond deployment to standard setting, pricing leverage, and geopolitical influence over spectrum policy and infrastructure exports.

China's centralized decision making has enabled rapid scaling and aggressive global expansion of its telecom model. For America, ceding leadership in wireless communications erodes both economic competitiveness and the resilience of critical communications infrastructure. Spectrum represents a finite national asset. Whoever wields it most efficiently gains outsized geopolitical and commercial leverage.

Al RAN presents a unique opportunity for America to reverse this dynamic. By combining U.S. strengths in advanced semiconductors, cloud infrastructure, and Al innovation, the nation can leapfrog static approaches and pioneer intelligent, adaptive allocation. Al RAN transforms wireless networks from a constrained, zero sum resource into a dynamic, software defined asset, enabling networks that are faster, more efficient, and more resilient than centralized alternatives.

This approach allows America to reassert control over spectrum optimization, set the pace for global standards, and demonstrate innovation driven leadership that safeguards national security and fuels the next wave of digital growth.

AI-RAN Server



- 1. Radio Unit (RU) transmits & receives radio frequency signals from mobile devices.
- 2. Al-native Centralized Unit (CU)/Distributed Unit (DU) uses AI to process and transmit wireless signal via the RU to & from mobile devices.
- 3. Al for RAN Applications dynamically optimize network performance.
- 4. Al on RAN Applications enable Al inferencing at the edge.
- 5. AI & RAN Orchestration is a software control layer that coordinates, manages and optimizes radio access networks (RAN) using Al and automation.
- **6. Core Network** is the central backbone connecting the AI-RAN platform to the broader internet and other networks.

Understanding AI-RAN

Every second, millions of radio transmissions use a spectrum that extends from 600 megahertz to beyond 100 gigahertz. These transmissions can include voice calls, data packets, video streams, sensor readings, GPS coordinates, encrypted military communications, and other media. To manage this profusion of traffic, Al RAN integrates artificial intelligence directly into the radio access network, which is the infrastructure layer connecting end user devices from the edge to the broader telecom network.

This integration operates through the deployment of Al models that continuously optimize power control, beamforming, interference mitigation, and anomaly detection, with the ability to continuously train these models based on observed traffic patterns and network conditions. Traditional cell towers function as sophisticated relay stations; AI RAN changes them into thinking machines.

As mentioned, AI RAN will transform networks in two fundamental ways. First, the application of AI to wireless networks will drive new levels of spectral and power efficiency and make the network more programmable to adjust to multiple types of scenarios and use cases. This programmability will reduce the need for manual configurations while making the network more adaptable to changes in demand, interference, and outages.

In addition, the open programmability of AI RAN will enable new network features. For example, industry and government have been developing new integrated sensing and communications (ISAC) capabilities, a key feature that will be incorporated in future 6G standards. With ISAC, the network now becomes a sensing system with the ability to detect objects in its vicinity, unlocking the potential for a plethora of novel applications.

Second, AI RAN will allow for the same infrastructure

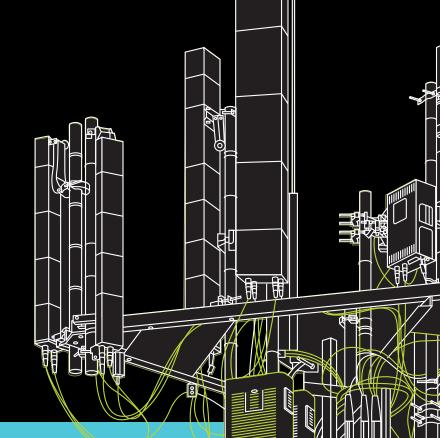


management. Beyond routing traffic, AI RAN systems use neural networks that learn from patterns and anticipate congestion, allowing for precise, instant adjustments.

The AI RAN systems of the future will run on models that understand the deep physics of radio wave propagation and the tactics of electronic warfare, all rolled into a single AI system that can adapt to shifting scenarios. Technical implementation includes network apps that execute AI algorithms in sub 1-millisecond loops. These models range from optimization algorithms to sophisticated neural networks capable of processing computer vision or natural language tasks.

The Road to 6G

Looking ahead, the current evolution to AI native 6G follows an accelerated cycle that departs from traditional telecom transitions. Software defined network architectures compress development timelines by reducing hardware



FEDERAL MISSION USE CASES

For federal agencies, the time is now to start pilots so that decision makers can evaluate the new technology and influence 6G standards and technology development if needed. Here are potential areas of focus:



Autonomous systems represent one the most compelling federal use cases for AI-RAN because they require split-second coordination among multiple types of sensors during operation in contested network environments. Today's autonomous vehicles, drones, and unmanned maritime systems operate largely in isolation, making independent decisions based on locally processed sensor data, which limits operational effectiveness.



For autonomous intelligence, surveillance, and reconnaissance (ISR) platforms, AI-RAN enables unmanned aerial and ground vehicles to stream and analyze sensor data locally while adapting communication links to conditions on the ground. Instead of losing capability when traditional links fail, platforms become more resilient as the network intelligently adapts.



Pop-up disaster response networks address critical needs for Federal Emergency Management Agency, Department of Homeland Security, and Coast Guard operations requiring rapidly deployable communication infrastructure in areas without traditional networks. Al-RAN enables portable cell sites as converged infrastructure offering both advanced communications and mission-focused edge applications.



Border security and critical infrastructure monitoring benefit from AI-RAN's edge inference capabilities on video and thermal sensor feeds. Rather than streaming raw surveillance data to centralized processing centers, network infrastructure identifies anomalies and automatically allocates additional bandwidth to priority alerts.



Other high potential use cases include smart cities and intelligent transportation networks, air traffic control, and energy optimization across large facilities and campuses.

How to Prepare: Agency Action Plan

Federal agencies can successfully deploy Al RAN within their environments by following several key principles:

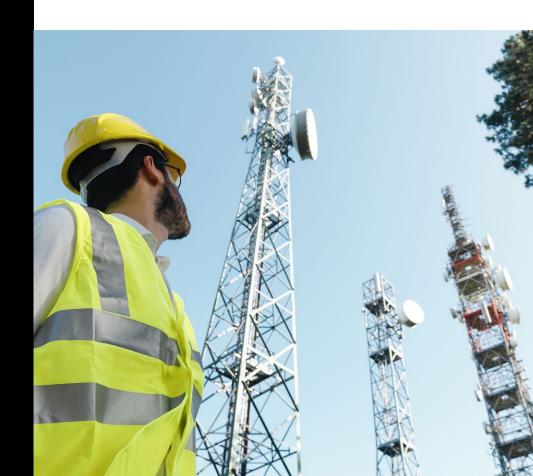
- Establish a cross-functional tiger team. Bring together network engineers, spectrum planners, Al specialists, cybersecurity experts, and acquisition professionals. Al RAN's interdisciplinary nature requires this integrated expertise.
- Launch pilot programs immediately.
 Pilots provide safe environments to discover vulnerabilities and failure modes, streamlining preparation for future deployment.
- Embrace open, modular architectures.
 Align with existing Open Radio Access Network (O RAN) frameworks and open APIs to prevent vendor lock in while enabling rapid AI model updates as algorithms improve.
- Build in security and governance from day one. AI RAN introduces new attack surfaces through AI models and edge processing. Implement zero trust architectures for RAN elements and establish model risk management procedures.
- Shape policy and spectrum strategy.
 Ensure that mission critical frequency bands and spectrum sharing frameworks allow optimization using AI. Shape NTIA and FCC regulatory environments rather than adapting to them.
- Plan for sustained operations and talent. AI RAN requires ML operations capabilities, radio frequency engineers who understand AI, and continuous testing ranges for model updates. Build these capabilities in parallel with technical deployments.
- Consider alternative procurement mechanisms. Given Al-RAN's hybrid nature (integrated hardware, software, and network) and emerging status, tools like other transaction authorities or commercial solutions opening pathways may be more effective for acquiring an initial Al RAN technology stack.

dependencies across standards development, testing, and compliance cycles. This enables more rapid iteration than previous hardware-centric generations.

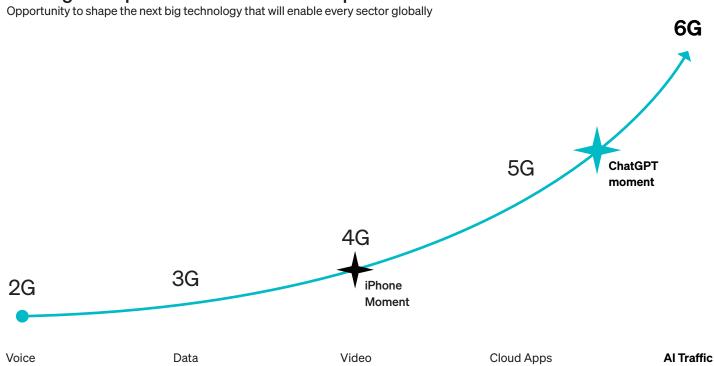
Developed by a consortium of seven leading telecom standard development organizations, 3GPP Release 18-19 (5G-Advanced), finalizing through 2024–2026, is establishing Al/ML service enablers within RAN elements as the technical foundation. Release 20 (2026–2028) features dedicated 6G studies, with Release 21 codifying core specifications for commercial availability around 2030–2031.

Helping drive industry momentum is the NVIDIA-led AI-WIN project. This project, announced in March 2025, brings together Booz Allen and NVIDIA along with T-Mobile, MITRE, Cisco, and ODC to develop an AI-native wireless network stack for 6G on the NVIDIA AI Aerial platform. Our collaboration focuses on creating AI-native hardware, software, and architectures that improve services for potentially billions of users.

To support this effort, Booz Allen is developing Al-RAN algorithms as well as approaches to secure the Al-native 6G wireless platform. Through our NextG lab, a dedicated, secure test environment including a leading-edge carriergrade network, we will conduct functional, performance integration and security testing to ensure the resilience and security of the platform against the most sophisticated adversaries. We will also lead field trials for advanced use cases such as autonomy and robotics. The Al-WIN approach promises to compress traditional telecom development cycles by embedding Al functionality from the start rather than retrofitting existing architectures, positioning America to lead 6G development.



The Urgent Imperative for U.S. Leadership in 5G



2010

Act Early, Lead Boldly

1990

AI-RAN represents a fundamental shift toward intelligent, adaptive infrastructure that optimizes itself and hosts mission-critical Al services at the edge. This dual capability makes it central to 6G architecture and essential for future federal operations.

2000

America has a small window to reassert telecom leadership through investment in open, Al-native architectures. Success requires moving beyond traditional procurement toward active participation in standards and technology development, aggressive pilot programs, and policy alignment.

The convergence toward Al-native 6G will continue accelerating with or without U.S. leadership. American agencies and companies can shape this transformation or be forced to adapt to architectures developed elsewhere.

With proven deployment experience, strong industry partnerships, and clear mission requirements, federal agencies are positioned to lead this transition with decisive action.

2020

2030

But agencies shouldn't wait for 6G's anticipated deployment before acting. Current 5G-Advanced infrastructure provides the foundational compute and edge processing power to implement AI-RAN solutions today. Moving ahead now will enable agencies to gain strategic advantages in network resilience, spectrum efficiency, and operational intelligence while building knowledge to streamline future 6G integration.

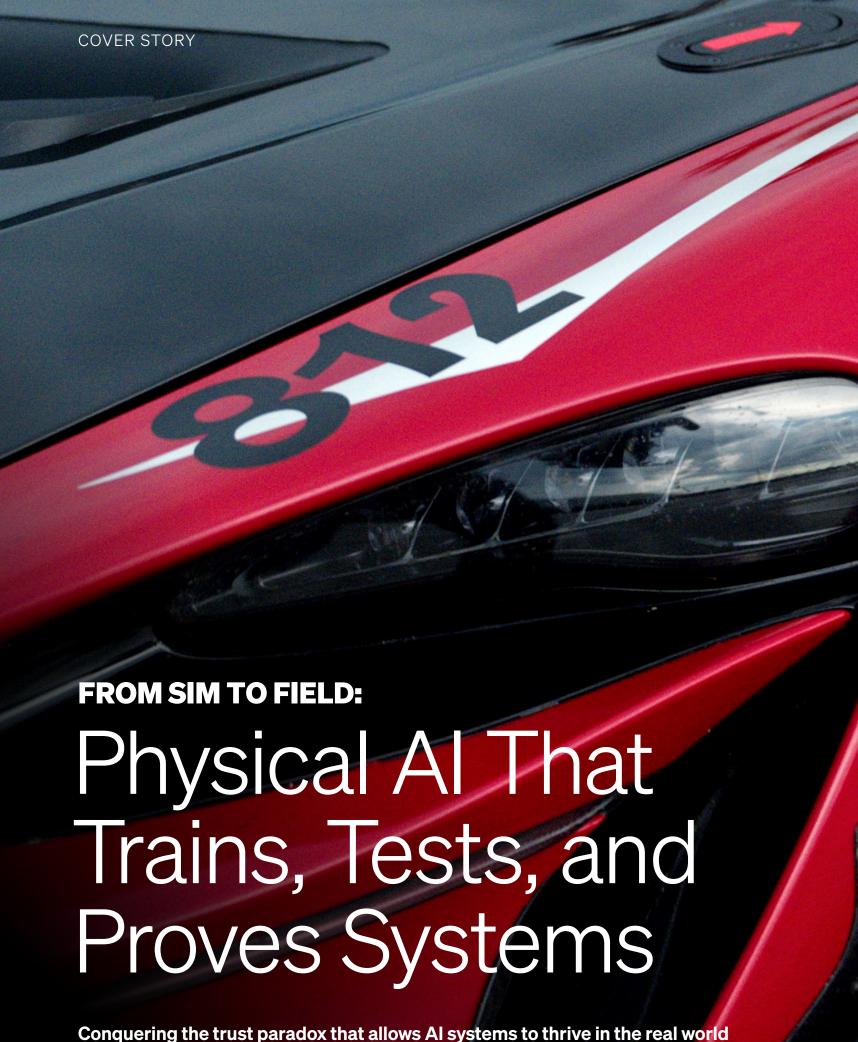
The path forward demands immediate steps: select pilot sites, join industry alliances, align spectrum policies, establish governance frameworks, and build talent pipelines necessary to own advanced capabilities. In the race to 6G, intelligence wins. The future will belong to nations that teach their networks to think.

SPEED READ

AI-RAN transforms wireless networks from passive data pipes into intelligent, adaptive platforms, optimizing network performance and enabling edge applications using the same infrastructure.

AI-RAN is the foundation for the future 6G stack, with its technical development offering the U.S. an opportunity to reclaim leadership in the global telecom arena by leveraging a wealth of differentiating software and AI expertise.

Current 5G-Advanced infrastructure supports Al-RAN, allowing agencies to begin benefiting from the technology today while developing readiness for the upcoming 6G era.



Munjeet Singh, Greg Kacprzynski, Randy Yamada, Ph.D., and Drew Massey



ess than a decade ago, experts predicted that 10 million self-driving cars would be on the road by 2020. This kind of optimism about autonomy—not just in vehicles, but across a slew of commercial use-cases, from factory production lines to healthcare—was widely shared at the time. As recently as 2016, Stanford University's One Hundred Year Study on Artificial Intelligence, or Al100, predicted "We will see self-driving ... delivery vehicles, flying vehicles, and trucks," as well as cars by 2020.

Five years later, in its 2021 report, Al100 acknowledged that this prediction had been "overly optimistic." In 2025, there are only a few thousand fully autonomous vehicles operating in the U.S. (most with human minders), out of 285 million cars and trucks on the road.

What happened? Self-driving cars have proven increasingly capable, performing as or more safely than human drivers in certain situations, like highway driving in good light conditions. But the public remains skeptical—only 13% trust the safety of autonomous vehicles—and has set high expectations.

For other use cases, the bar might be even higher: What is an acceptable error rate for autonomous defense or weapons systems? Or healthcare robots?

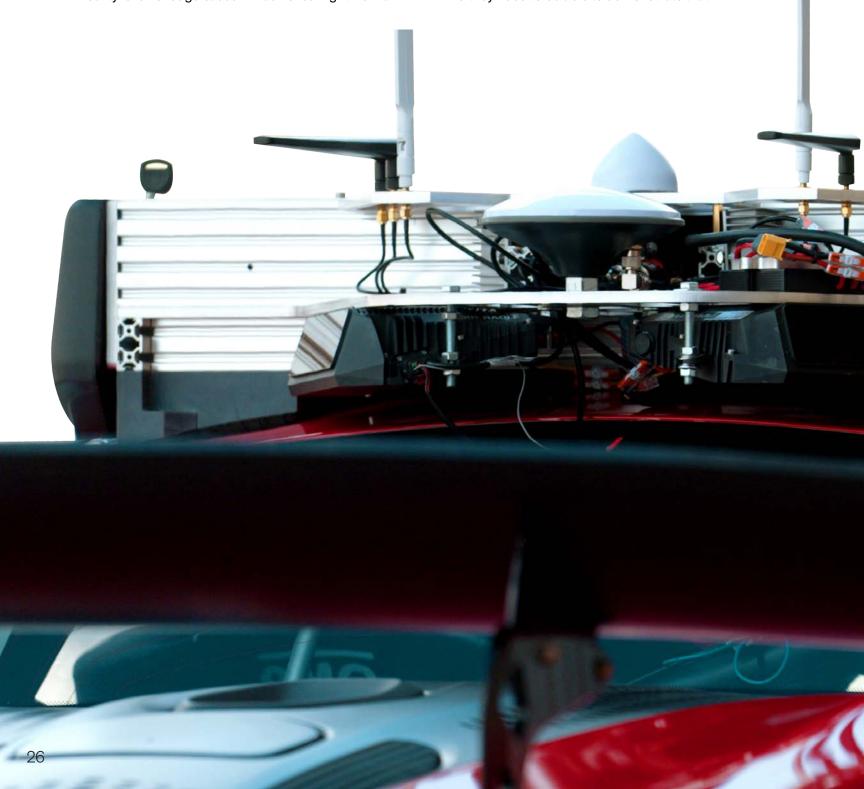
The hard truth is when Al-enabled systems are operating in the physical world, there's a lot that can go wrong. Reality is full of edge cases: A flash of sunlight from an

unexpected angle, a unique architectural feature, a pedestrian attired or moving unusually. Deployment in a factory where the environment is controllable to a degree is very different from deployment on city streets or a battlefield.

There's a paradox here: In some cases, autonomous systems can already surpass human performance. However, that's still not good enough given the heightened regulation and scrutiny many missions and applications face.

In the physical world, one mistake can wreck infrastructure and put lives at risk.

Many federal missions must strive to achieve reliability, performance, and safety standards closer to the same five-nines (99.999%) required for critical infrastructure like telecommunications and in commercial aerospace. And they need to be able to demonstrate that



performance against benchmarks defined in law and regulation.

Advances in physical AI offer a solution: Physical AI integrates virtual environments that replicate real-world conditions with high-fidelity, randomized simulations of actual scenarios. This combination allows for training, testing, improving, and deploying AI-powered systems at scale, through extended exposure to modeled environments, edge cases, and unique scenarios.

And these same digital world models and virtual environments can be used to address the extensive and often time-consuming validation & verification (V&V) testing needed to confirm safe and trustworthy performance of autonomous systems.

Physical Al is both a force multiplier and risk mitigator for Al-powered systems operating in the physical world—designed to push boundaries, reduce uncertainty, and enable innovation.

PHYSICAL AI: AT-A-GLANCE

Physical Al is a branch of Al focused on enabling machines and smart sensors to perceive, understand, and perform complex actions in the physical world. Training physical Al models requires either real or synthetic data that accurately reflects real world conditions, such as lighting, mass, motion, and contact dynamics. While virtual environments are often preferred for their cost effectiveness and scalability, the "sim to real gap" the discrepancy between simulation and reality remains a challenge. Fortunately, the increasing fidelity of modeling and simulation capabilities helps create more accurate world models. Additionally, techniques like domain adaptation and transfer learning are advancing to minimize errors and shrink the sim to real gap. This progress leads to more reliable digital replicas, or digital twins, of the real world, which serve as "digital proving rounds." These digital twins help de risk and accelerate the operationalization of physical Al.

VELOCITY | © 2025 BOOZ ALLE.

THE LOOP THAT CHANGES EVERYTHING: A GUIDE FOR ADOPTION

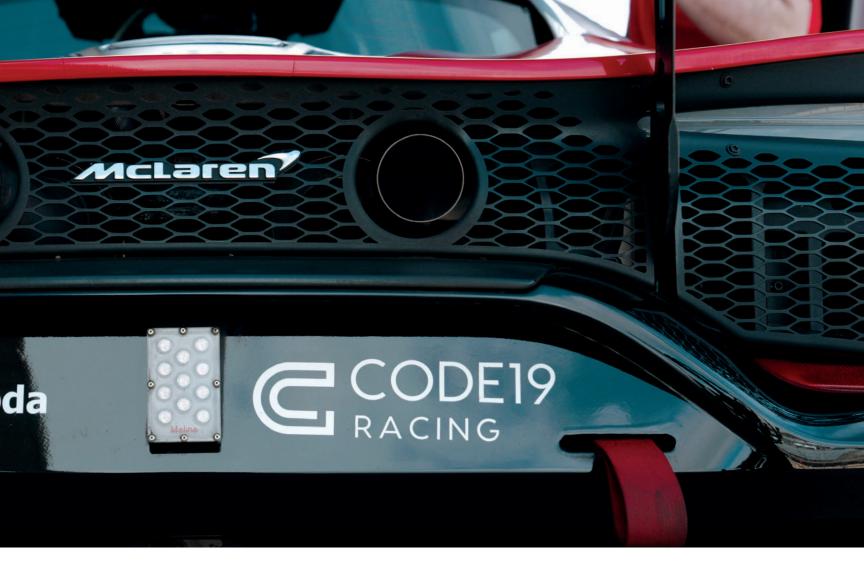
In the last issue of *Velocity*, Booz Allen CTO Bill Vass introduced the concept of the modern technology flywheel. As he explained, this operating model combines real-time data, machine learning, and digital twins within a software-defined environment to create a positive feedback loop driving continuous improvement and performance optimization. Physical Al is at the heart of this model.

A physical AI stack takes existing simulation and digital twins to the next level. Instead of waiting for reality to reveal edge cases, it manufactures them. Instead of learning from operational failures, it fails thousands of times in simulation before succeeding in the field. Instead of hoping the next deployment goes better, it knows exactly which scenarios have been mastered and which remain risky.

In practice, a mature physical Al loop should function like a well-run factory for producing validated and verified Al. The process is continuous, disciplined, and focused on measurable improvement.







TRAIN ON REAL AND SYNTHETIC DATA

The training process pulls from two complementary streams. The first is the real world, providing the essential grounding for the models, including field logs, annotated failures, near-misses, and detailed operator feedback. This data is precious because it reflects the unscripted reality of the operating environment.

The second stream is the synthetic world. This is where you achieve scale. Using digital replicas of terrain, facilities, vehicles, and sensors, you can recombine variables into thousands of edge cases. This allows you to specifically target the known weaknesses of your models. Perception models learn what matters under varied conditions like rain, snow, and fog. Planning and control policies learn how to trade off speed, safety, and mission goals in complex situations.

PROVE IN HIGH-FIDELITY SIMULATION

Before any physical Al model is deployed, you rehearse it across a comprehensive catalog of scenarios, each tied to a specific mission risk. You can sweep through thousands of variations of lighting, weather, traffic, and sensor faults. Crucially, you can also include adversarial tactics, testing the system's resilience against intelligent opponents.

For each scenario, you record performance and confidence. The result is a body of test evidence that government evaluators can reproduce. This is a critical step. You now have a safety case that is more than a narrative; it's a set of numbers with traceable origins. This evidence becomes the bedrock of trust between the program office, the developer, and the end user.

DEPLOY, MONITOR, AND CLOSE THE LOOP

Only when the system clears these measurable gates does the model get packaged with guardrails and versioned. It lands on a fixed compute edge devicean NVIDIA Jetson, a core processor, or a mixed architecture that includes safety controllers.

The work doesn't stop at deployment. Updates are gated by the same scenarios used in simulation, ensuring regression testing is continuous and comprehensive. Telemetry returns to the factory: successes, failures, near-misses, and rich context about the environment. This data is triaged and prioritized, and the next training cycle focuses on what fleets—a group of autonomous systems or devices working collectively—found hard. When this becomes routine, your programs have something they rarely enjoy: A predictable way to get better, quickly and safely.

The magic isn't in the simulation technology itself. Commercial game engines have been photorealistic for years. The breakthrough is in the systematic approach to synthesizing physics-based scenarios that matter, validating that simulated failures predict real failures, and creating a continuous loop between field operations and synthetic training in an Al-native environment.

PHYSICAL AI AT THE CUTTING EDGE

Let's revisit self driving automobiles but at an even faster speed autonomous race cars operating on a road course.

The demands placed on the cars and teams make the world of auto racing one of the most complex environments to navigate autonomously. It forces perception, planning, and control to work under extreme time pressure, dynamically changing conditions, and against intelligent opponents. That makes it an ideal testbed for physical Al: We can train on real and synthetic data, rehearse thousands of scenarios in high fidelity simulation, then field improvements on track and feed telemetry back into the loop in an environment where we can readily track and assess performance.

Booz Allen is working with Code19 Racing, the only U.S. based team in the Abu Dhabi Autonomous Racing League (A2RL), which runs at the Yas Marina Circuit. The cars are Dallara EAV24 machines derived from Super Formula, built for autonomy with added sensors, actuators, and onboard compute. Working with Code19 offers a repeatable, measurable way we can prove behaviors before race day and then harden them after every session.

Racing highlights the "sim to real" challenge in concrete and complex ways. Consider tire temperature and grip. Cold tires slip. Track evolution (from temperature variations to rubbering and marbling) changes braking points and demands robust prediction under uncertainty. With integrated telemetry providing real time feedback on these dynamics to the simulation, it can learn and better replicate the real world driving experience. While the industry still has a long way to go to address the challenges within the "sim to real" gap, integrating simulated and real environments such as those that Code19 competes on provides an opportunity to explore tactics and techniques to close this gap.

A2RL has also staged human versus autonomy demonstrations at Yas Marina. So far, the human driver has prevailed, which is precisely why the race environment is valuable. It exposes the last eight seconds of performance gap that matter to mission risk and makes them measurable, fixable, and repeatable in the next turn of the flywheel.

Bottom line: High-speed autonomous racing combines a high performance platform, a world class circuit, and a rigorous physical Al loop to turn wicked hard operational problems into tractable engineering work faster, safer, and with evidence leaders can trust.

EVERYTHING CHANGES WHEN YOU BRING THE LOOP INTO THE REAL WORLD

Consider three domains where physical AI can make a mission critical difference.

Expeditionary autonomy. Improving navigation in GPS-denied areas used to require a lot of complicated, time-consuming testing, often involving extensive simulations of jamming and spoofing. Traditional methods like hardware-in-the-loop and software-in-the-loop testing are important but slow.

Now, with the integration of a physical Al loop, we can rehearse thousands of degraded communication sorties in a simulation environment. This speeds up the process and helps identify the most critical scenarios for real-world testing. For example, simultaneous localization and mapping (SLAM), a key method visual-inertial odometry uses for GPS-denied navigation, can be assessed in a digital proving ground. Consequently, fewer physical tests are needed, allowing for more frequent updates. This continuous learning approach enhances the readiness and reliability of navigation systems in GPS-denied situations, without waiting for large-scale exercises.

Smart manufacturing. On high-mix production or assembly lines, the constraint is variability. A part is slightly out of tolerance, a surface is oily, a bin is half empty, lighting shifts, a tool wears. Robots that ace a demo can stall in production. With a physical Al loop, you turn computer-aided design (CAD), process plans, and line telemetry into a digital workcell. You train perception and force control across millions of synthetic picks, insertions, torques, and welds, then verify with instrumented floor trials. You release when coverage, first-pass yield, and cycle-time thresholds are met. Real faults and E-stops expand the scenario catalog so each shift gets steadier.

Smart environments. Borders, ports, bases, and substations rely on multi-sensor networks. Static rules often produce either alarm fatigue or missed signals. With a loop, you train fusion models on rare multi-sensor events. You test how a video anomaly, a vibration pattern, and a radio observation interact under different weather and lighting. You don't need to wait for a real emergency to learn. When the model hands off to an autonomous platform or triggers a barrier, you already know how it will behave in conditions you would not accept as a live trial.

The theme across these examples is not novelty. It is confidence—at mission speed. You are not skipping safety; you are automating it.

THE LOOP WORKS FOR TESTING AS **WELL AS TRAINING**

Across the federal space, policy is moving from "simulation as adjunct" to digital-first, evidence-based testing and evaluation (T&E). The Department of Defense set the tone with DoDI 5000.97 requiring new programs to incorporate digital engineering across the lifecycle and explicitly includes digital T&E. It calls for program managers to plan, resource, and govern models, digital twins, data, and artifacts as part of their acquisition strategy.

Other federal agencies are pointing the same way. NASA-STD-7009B (2024) is a mature government standard for model credibility, V&V, uncertainty quantification, documentation, and configuration control, widely referenced as a best practice. The National Institute of Standards and Technology's (NIST) AI Risk Management Framework (AI RMF 1.0) emphasizes T&E V&V throughout the Al lifecycle and encourages the use of controlled testbeds and synthetic data where appropriate. In regulated aviation, the FAA's AC 20-115 recognizes modelbased development and verification as an acceptable means of compliance, which in practice relies heavily

on simulation to demonstrate software assurance. The Department of Homeland Security also signals the role of modeling and simulation in T&E planning and maintains centers of expertise to support operationally realistic evaluation.

Bottom line for programs. More advanced modeling and simulation can and should fulfill these needs, but only when it is planned from the start, tied to missionrelevant scenarios, and backed by V&V evidence. Digital engineering policy now expects programs to adhere to the following guidelines:

Treat simulations, twins, and data as governed acquisition artifacts.

Use distributed and hybrid testing to combine sim, software-in-the-loop, hardware-in-the-loop, and limited range events.

Present auditable coverage and performance metrics to decision authorities.

That is the policy foundation for a physical AI "digital proving ground" that reduces risk, speeds learning, and supports continuous accreditation.

"Motorsport is a crucible where challenging engineering problems are solved but where failures are both public and unforgiving. Consequently, Motorsport is also the benchmark for closing the 'sim-to-real' gap. By tightly integrating simulated and real-world race conditions, we can validate Al-driven systems at a pace that traditional development cycles simply can't match."

- O.G. Wells, Cofounder, Code19 Racing



TEST AND EVALUATION: TURN THE LOOP INTO A DIGITAL PROVING GROUND

Physical AI does more than teach systems to handle edge cases. It also gives programs a rigorous T&E engine that replicates real-world performance in software, then verifies and validates it before deployment.

You define a mission-relevant scenario catalog, tie each scenario to a risk budget and requirement, and rehearse thousands of controlled variations across weather, lighting, traffic, sensor faults, and adversarial behaviors. The same builds then run through software-in-the-loop, hardware-in-the-loop, and limited range trials, creating reproducible evidence that evaluators and safety authorities can trust. Leaders get an auditable package. not anecdotes: scenario coverage and pass rates, confidence and calibration curves, latency and stability margins, disengagement and safe-state entry metrics, and a traceable safety case that blends simulation with instrumented field data. Because scenarios, assets, and telemetry schemas are standardized, you can replay the identical test battery across multiple engines and test sites, compare apples to apples, and keep regression testing continuous after deployment.

The result is faster developmental testing (DT) and operational testing (OT), clearer readiness decisions, and a path to continuous accreditation where each release ships when coverage and risk thresholds are met, not when the next test window opens.

WHY NOW? HARDWARE AND SOFTWARE ADVANCES CREATE NEW OPPORTUNITIES

Today's robotics, autonomous systems, and smart infrastructure are powered by tech stacks that sense, understand, decide, and act, often inside a smart environment that is instrumented with connected sensors and compute.

Perception and Localization. Cameras, LiDAR, radar, IMUs, RF signals, and maps place the system in context.

Decision and Control. Planners and policies balance mission goals, constraints, and safety while coordinating across agents.

Connectivity and Compute. Edge processing supports low-latency reflexes; cloud and high-performance computing (HPC) support training, analytics, and fleet learning.

But this stack is poised for disruption as ongoing advances in both inputs, like camera and sensors, and outputs, including actuators and propulsion systems, enable higher fidelity, more agile performance. For these systems to operate at their full potential, more decisive reasoning is needed, which physical AI can foster.

PHYSICAL AI: SHIP FASTER, SAFER, WITH EVIDENCE

Even with solid building blocks, programs still struggle to generalize to the long tail of edge cases, to produce auditable evidence of readiness, to maintain data integrity and provenance, and to remain resilient when networks degrade or adversaries intervene. This is where physical AI shines. Physical AI is not just a tool. It is an operating model that teaches, tests, and proves behaviors before fielding, then keeps improving them after. The payoff is a repeatable, validated process cycle that automates both improvements and safety testing.

Treat the loop as a digital proving ground for testing and evaluation. Define a mission-relevant scenario catalog, rehearse thousands of controlled variations in high-fidelity environments, and get reproducible evidence your evaluators can trust. Use the same scenarios for software-in-the-loop, hardware-in-the-loop, limited range trials, and post-deployment regression so each release ships on demonstrated coverage and risk thresholds, not on calendar windows.

To stay agile as the ecosystem evolves, own the interfaces and artifacts: The scenario catalog, policy bundle, telemetry schema, and evaluation reports. This preserves flexibility across engines and vendors while accelerating accreditation and scale from pilot to program to portfolio.





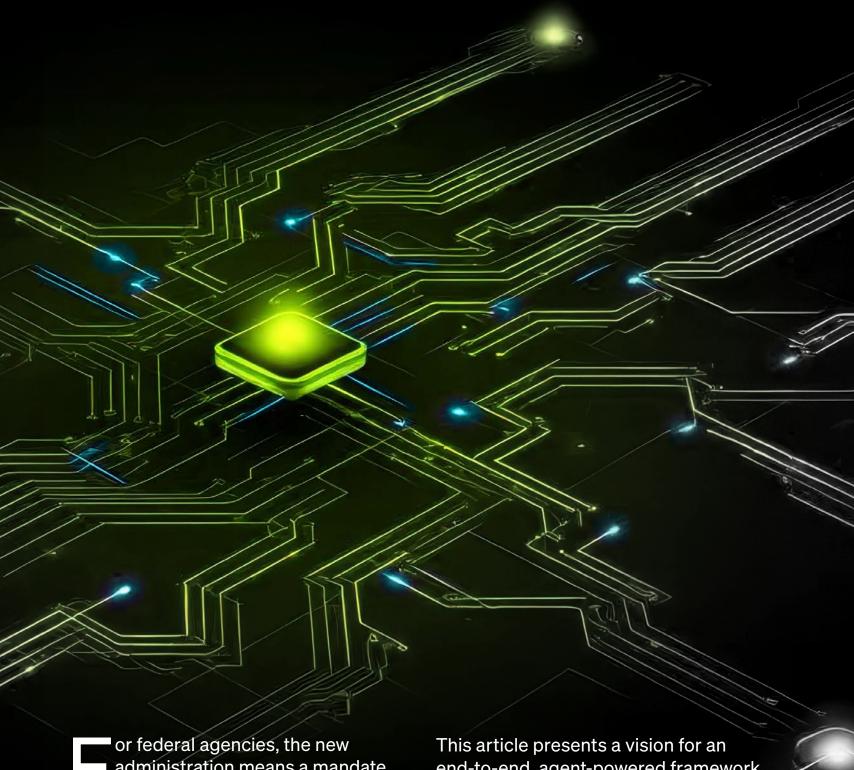
The integration of physical Al enables continuous learning and improving operational readiness through systematic and reproducible testing in both simulated and real world environments, reducing the need for extensive physical trials and allowing for quicker, safer deployment of Al powered systems.

Physical Al also offers a rigorous test and evaluation process that replicates real world performance in software and verifies it before deployment, producing auditable evidence that accelerates accreditation and supports continuous updates. It is an essential tool for mission critical applications in diverse fields such as autonomous navigation, smart manufacturing, and security.

Agentic Development Decoded: Hard-Won Lessons From the Trenches

How federal IT teams can compress the software development lifecycle, accelerate delivery, and break the legacy transformation logiam.

Bassel Haidar, Ryan Stone, Jennifer Sheppard, and Michael Martoccia



or federal agencies, the new administration means a mandate for accelerated enterprise transformation. Current methods are often too slow to meet evolving mission demands, and too resource-intensive for agency budgets. While Al-powered code generation offers a significant tactical boost, the true transformative potential of Al lies in automating nearly every phase of the entire software development lifecycle (SDLC).

This article presents a vision for an end-to-end, agent-powered framework that is holistic, human-centric, and toolagnostic. We explore how Al agents and intelligent platforms interconnect to automate everything from requirements analysis to operational monitoring, using a real-world case study to illustrate the path from legacy systems to future-ready, user-centered applications that can finally move at the speed of mission.

The Weight of Legacy and the False Dawn of Al Coding

In an era of accelerating peer competition and rising citizen expectations, fast delivery of secure and reliable software isn't just essential to the mission; it is the mission. For federal leaders, the core challenge is clear: The current approach to the SDLC is too slow, too brittle, and too burdened by manual compliance to keep pace. The distance between a new requirement and a deployed capability is measured in months or years, a delay we can no longer afford.

A new generation of federal leaders with backgrounds in technology or business are determined to reshape the organizations they lead. Agencies and departments are expected to do more with less and are looking to use technology to square that circle.

When it comes to modernization, the message is clear: No more band aids. There's a preference for radical, transformative solutions and greenfield builds.

Against this backdrop, first generative AI and now agentic AI are transforming the development of software, advancing over the past year at a pace that has far exceeded our most optimistic expectations (visit *The Age of Agentic AI* in the previous issue of *Velocity* for a primer on these technologies). For example, the extraordinary speed of improvement of the large language models (LLMs) powering the Claude Code, Cursor, and Github CoPilot tools means that any performance numbers we could include in this article would be out of date by the time it was printed.

It's hard to overstate the significance of these changes. All has moved from a clever coding accelerant to a broader change in how software gets made. The biggest gains don't come from drafting snippets in an integrated development environment (IDE); they come from shortening the distance between intent and impact. The goal is no longer just improving individual productivity; it's about building an Al-powered software factory.

We believe there's synergy between these vast ongoing changes: The emerging Al-powered software development ecosystem, properly managed, will make the transformation of federal agencies—or any other coding project—possible at a speed and scale unimaginable just a year ago.

Indeed, it is already starting to do so.

As we write, in September 2025, AI tools are highly effective for automating certain discrete tasks in the software development process. Over the next two years, this paradigm shift will only accelerate and grow.

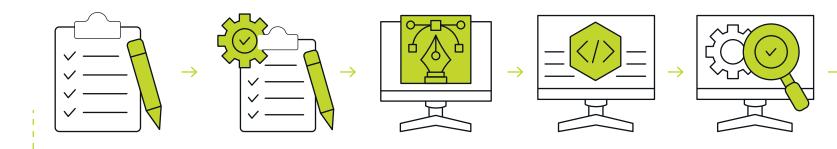
What We've Learned: From Early Pilots to Durable Platforms

Early pilots and projects across federal agencies and mission partners tell a consistent story. The center of gravity has shifted from "Al coding" to "Al-enhanced delivery." Agencies are experimenting with assistants in IDEs, but the real leverage appears when Al improves upstream clarity and downstream assurance.

The immediate wins are clear and compelling:

- Code understanding and navigation: Developers use
 Al to summarize unfamiliar modules, trace dependencies,
 and generate scaffolding and documentation that would
 normally take hours.
- Test generation and maintenance: Al accelerates unit, integration, and regression tests, including self-healing scripts tied to selectors and application programming interfaces (APIs).
- Knowledge capture: Teams use AI to generate living runbooks, design notes, and architectural decision records from sources, tickets, and wikis.
- Operational toil: Incident analysis, log summarization, and safe rollout plans benefit from AI that is connected to continuous integration (CI)/continuous delivery (CD) pipelines and observability platforms.

```
EXPORTSYMBOL(groupsalloc);
              void groups_free(struct group_info *group_info)
              Void groups_free(struct group_info *group_info)
                   if (groupinfo->blocks[0] != group_info->small_block) {
                   if (groupinfo->blocks[0] != group_info->small_block) {
                             = 0; i < group_info->nblocks; i++)
      True
                       freepage((unsigned long)groupinfo->blocks[i]);
for (i = 0; i < group_info->nblocks; i++)
_z = False
"MIRROR_Z":
                            reepage((unsigned long)groupinfo->blocks[i]);
x = False
                   kfree(groupinfo)
y = False }
                   kfree(groupinfo);
_z = True
= modifier
the end -add back the do
               EXPORTSYMBOL(groupsfree);
=1
objects.active = modifier
const struct group_info *group_info)
tatic int groups_touser(gid_t _user *grouplist,
                            const struct group_info *group_info)
                   int i;
                   unsigned int count = groupinfo->ngroups;
                   int í;
                   unsigned int count = groupinfo->ngroups;
                   for (i = 0; i < group_info->nblocks; i++) {
                       unsigned int cpcount = min(NGROUPSPERBLOCK, count);
                   for (i = 0; i < group_info->nblocks; i++) {
                      unsigned int len = cpcount * sizeof(*grouplist);
unsigned int cpcount = min(NGROUPSPERBLOCK, count);
                      unsigned int len = cpcount * sizeof(*grouplist);
                      if (copyto_user(grouplist, group_info->blocks[i], len))
                                      ULT;
(qrouplist, group_info->blocks[i], len))
                                                                                            37
```



However, real gaps remain, meaning results were too often uneven:

- Whole-app generation: All can draft simple apps and boilerplate, but complexity, security rules, and legacy interfaces limit end-to-end automation.
- Context gaps: Assistants without access to the right repositories, APIs, and policies deliver generic answers that look polished yet require rework.
- Brittle integrations: Promising pilot agents become fragile in production when they rely on ad hoc scripts or cannot survive dependency changes.
- Change management: Productivity gains fade when teams do not update code review norms, branching strategies, and release gates for Al-touched work.

These early successes share the same profile. They are high-volume, high-friction, low-risk activities where the definition of "good" is clear and verification can be automated. This is exactly where federal programs should start:

- Guardrails: Policy, privacy, and provenance built into the platform.
- 2. **Grounding:** Access to the right code, data, and documentation—at the right time.
- 3. **Granularity:** Agents scoped to well-defined tasks with clear acceptance criteria.
- 4. **Good measures:** Delivery and quality metrics that show where the gain actually is.

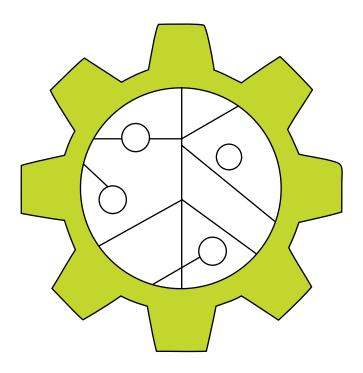
Select external signals confirm the direction. Adoption is rising, agentic patterns are emerging, and the SDLC is compressing in places. At the same time, organizations report uneven results when pilots skip governance, skip measurement, or overreach for full automation.

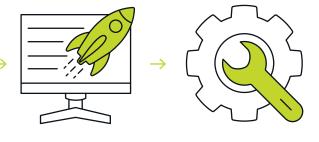
How AI Compresses the Full SDLC

Let's start with a global perspective. Gartner® predicts that:

"By 2028, 90 percent of enterprise software engineers will use AI code assistants, up from less than 14 percent in early 2024. The role of developers will shift from implementation to orchestration, focusing on problem solving and system design, and ensuring AI tools deliver high-quality outcomes. To succeed, teams must balance automation with human oversight, considering business criticality, risk and workflow complexity."

We believe this mirrors the increasing use of Al-powered approaches to compress the current SDLC that we are witnessing. A starting point is the use of generative Al to automate rote coding tasks. However, agentic Al is increasingly being used to execute more complex assignments, building upon autonomous, outcome-driven Al agents that continuously learn and adapt to achieve goals with minimal human intervention.





The situation is fast-changing, fluid but clear: Traditional manual development is rapidly transforming into humanguided, Al-powered execution. Digging deeper, we find this transformation evolving as follows.

· Upstream: Intent to Architecture

Al agents turn unstructured inputs into a living backlog, cluster related needs, flag conflicts, and link requirements to acceptance tests. Design agents propose patterns, enumerate trade-offs, and stresstest options against nonfunctional constraints such as privacy, safety, and performance. Human architects still decide, but they decide faster and with better evidence.

· Midstream: Build to Assurance

Leading teams use a curated workbench with approved APIs, software development kits (SDKs), and templates that encode security and compliance by default (see more in the *Velocity* article "How We Built a Better Innovation Factory"). Shift-left becomes real as agents expand test coverage, generate edge cases, and keep tests current as code changes. Policies and security rules are expressed as code and enforced as automated pre-merge checks.

• Downstream: Operations to Learning

In production, agents assist site reliability engineers with anomaly detection, impact analysis, and rollback planning. Connected to playbooks, logs, and topology maps, they recommend the next best action during incidents and feed lessons back into requirements and tests. The aim is safer changes, fewer pages, and faster recovery.

A key differentiator is orchestration across the lifecycle. Software teams add outsized value by mastering these tools, knowing their limits, and governing them well.

Cyber Teams Embrace Agentic Engineering to Defeat the Threat

Agentic AI is emerging as a gamechanger across many facets of software engineering. For example, the Booz Allen Vellox Reverser™ product is an AI first, cloud native malware reverse engineering product designed to radically speed up threat response. It leverages agentic AI and machine learning to automate an exhaustive analysis of complex malware, delivering deep insights and actionable reports with comprehensive countermeasures in just minutes. AI agents orchestrate dynamic submissions across various tools and sandboxes to quickly assess and adapt defense strategies based on performance and behavior in real time.



The Agentic Developer: A New Blueprint for IT Transformation

It is time to rethink the end to end SDLC. Over the past 20 years, we've evolved from waterfall to agile to DevOps to today's loosely coupled world of cloud based containers, web services, and APIs. Al is rendering this world obsolete faster than we probably realize.

We can no longer draw lines around existing, often human based functions. Yes, there are compliance or other reasons to wall some functions off. However, these barriers are under constant attack from technology's relentless advances.

Instead, we should be guided by principles: Foundational pillars that reflect our values and shape our strategy:



Holistic: We look at the entire software development lifecycle end to end. And we have a master plan to orchestrate every stage and persona of the process. If you don't have all the puzzle pieces, you can't see the big picture.



Human-Centric Design: The software end user is at the center of everything we do, because the users are the ones doing the mission, and the mission is why we are here. But the human centric approach also emphasizes automation and reusability to free developers from the drudgery of repetitive, non creative work.



Tool-Agnostic: Avoiding vendor lock in is essential. We must be able to utilize tools from any trusted supplier, because the competition to be "best of breed" is only going to get more intense, and, in an emerging technology field, that contest can shift swiftly in unexpected ways.

When we look at the transformative impact these AI tools are going to have, it's important to break down the SDLC and examine each stage. Different parts of the process will be impacted differently:

Phase 1: Make Knowledge Computable

Adapting to the AI powered SDLC starts with making knowledge computable. Consolidate requirements, constraints, and decisions in a single, queryable knowledge base that cites its sources. A requirements assistant sits on top of this record and answers questions, highlights contradictions, and proposes acceptance tests that trace back to stated needs. When teams can interrogate the truth, they find gaps earlier, align faster, and prevent rework before code exists.

Phase 2: Institutionalize Design Intelligence

Design intelligence comes next. Institutionalize agents that compare architectural options against concrete targets for cost, performance, privacy, and scalability. Treat nonfunctional requirements as living guardrails, not slideware. Fitness functions continuously test architectures as they evolve, while an architectural decision record captures the why behind each choice. Human architects still decide; they do so with stronger evidence and less cycle time.

Phase 3: Build the Intelligent Workbench

Developers need a smart workbench, not a bag of ad hoc tools. Stand up a browser based portal that curates approved APIs, SDKs, templates, and code libraries with clear usage guidance. Include examples, ratings, and guardrails so teams can move fast without creating risk. Instrument the portal, learn what gets used, retire unsafe patterns, and expand what actually accelerates delivery.

Phase 4: Make Assurance Continuous

Assurance becomes continuous when testing shifts from a late gate to a built in practice. Use Al agents to generate unit, integration, property based, and adversarial tests, and to keep those tests current as code changes. Express policies and security rules as code, then enforce them before merge and before deploy. Standardize canaries, automatic rollbacks, and post incident learning loops so production teaches development in near real time.

Phase 5: Manage the Agent Ecosystem

Manage the agent ecosystem with the same rigor you bring to software supply chains. An agentic Al studio curates and secures domain trained agents, sets permissions and data scopes, and defines where a human must approve actions. An agentic Al mesh orchestrates handoffs across requirements, planning, design, coding, testing, and operations. Start with supervised autonomy, expand privileges as evidence of reliability grows, and instrument everything. Maintain audit trails, model bills of materials, bias and performance monitors, and clear incident playbooks. The result is a governed system of work where Al accelerates and humans remain accountable.



Case Study: Reimagining Modernization

Imagine a federal grant program in the near future. The core client server application is more than two decades old, making the system more brittle and costly to maintain. What's more, the associated business processes and data are often locked in place as well. This means that the agency is reliant on manual processes to report performance and ROI. Requirements for updates are communicated via email. The grant approval process takes three or four months, and data from the app shows that many grant applications are abandoned before being submitted, but with no granularity as to why.

Engaged to modernize the process and the app, an Al-assisted Booz Allen development team deploys. Within days, they can generate artifacts: Requirements, business logic analyses, software pathways. The legacy codebase can either be analyzed and translated into business logic by our specialized code-to-logic engine

to prepare it for refactoring or transposed directly into modern code languages by a specially trained code translator Al. Within a week, the team has a detailed roadmap, laying out these options and others for the modernization process.

Within a few months, the modernization is complete. The new system is live. The new, more user-friendly interface has a much lower level of abandoned applications. The app uses AI to flag high-risk applications for additional review, reducing fraud. But despite these additional layers of review, grant approval time is down to 30 days or fewer. An operational agent from our prebuilt agent catalogue now provides real-time, public-facing dashboards on fund distribution, fulfilling transparency and congressional reporting requirements automatically, which frees up the program's staff for other duties.



The Human-Centric Imperative: The New Federal Technologist

It might seem counter-intuitive to say so, but the human developer is front and center of this coming agentic Al transformation. Because of the important limitations of Al coding assistants and these other tools that we've discussed, they still need humans to manage them, to check their work, and to add vital context to their output. Developers need to transition from directly employing their coding skills to leveraging them to assign and assess the work of Al agents. They need to move from being coders to mission engineers, orchestrating and managing the work of trained and specialized Al personas.

Because of the centrality of human software developers, and the way their skillsets need to evolve, workforce issues are key for federal agencies seeking to ride this transformational wave. Currently, adoption of these new tools is spotty, according to Booz Allen internal research. There's uncertainty about their value, but also resistance stemming from broader cultural narratives. Some of that is due to fears about job losses, and some to concern about the costs, benefits, and risks of Al—its growing use of power and water for instance, or the perceived dangers of empowering non-human intelligence.

Whatever the source, management has to confront these doubts and fears head on. These tools won't replace developers; they'll empower them, freeing them from repetitive and uncreative drudge work. The expertise of senior developers will be more important than ever. Only those with deep mission knowledge and advanced coding skills will be able to ask the right questions, and successfully select and direct an ecosystem of different Al tools to solve complex problems. Highlighting the successes of early-adopter teams, and leading by example—with management publicly using the tools in high-profile projects—are two ways of getting this message across.

Even as AI agents become more capable and the agentic AI mesh enables them to interact autonomously with each other, that human expertise will still be essential. The best developers and engineers won't write (or won't need to write) the best code any more, but they will train the best AI agents. In the future, the intellectual property represented by those exquisite prompts and curated training data sets will be one of the enterprise's major assets, just as its codebase is now.

Above all, the message for the workforce should be about how much better these tools make us at doing our job: Enabling the mission. It is about speed, but it's not all about speed. It's all about being able to deliver better, more performant, more reliable software—and do it faster.

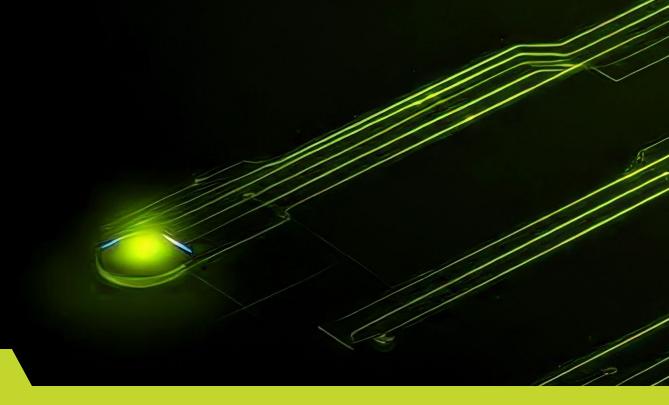
Assistive Today, Integrated Tomorrow: Your Roadmap to Al-Powered Software Development

As we write, in September 2025, Al tools are highly effective for automating certain discrete tasks in the software development process. Over the next two years, this paradigm shift will only accelerate and grow. In the commercial sector there will be vendors promising (although not always delivering) end to end automation "out of the box." In the federal context, supervised Al agents will be executing entire chains of activity: Upon receiving a requirement, they will begin drafting user stories, suggesting architectures, generating code, and designing tests, and eventually deploy to a staging environment with minimal human input beyond approval gates. The agent mesh enables software developers to empower "Al teams" chains of agents each executing on the tasks for which they're specially trained and passing off to the next.

Here are specific steps that federal IT leaders can take to prepare their organizations for this brave new world:

1. Establish Your Baseline: You cannot justify an investment without proving an ROI. And for that, you need data. Begin by deploying a software engineering intelligence platform to measure your current processes, identify bottlenecks and blockages, and start to build your business case.

- 2. Launch a Lighthouse Project: Select a high visibility, low risk project. Maybe a legacy modernization, maybe greenfield: Whatever is available and fits the required risk/visibility profile. The point is not the project; it's the public testbed for new Al tools and the agentic developer concept. A successful lighthouse project demonstrates their utility and builds organizational buy-in.
- 3. Set Realistic Expectations: The initial ROI is efficiency. The ultimate, strategic ROI is mission agility. The goal is not just to deliver faster, not even to deliver faster and better. Rather, the goal is to build an organization that can create, adapt, and deploy new digital capabilities at the speed of need, not at the pace of a multi year procurement cycle.
- 4. Create an Al Center of Excellence (CoE): This group will be responsible for governance: Interpreting policy and setting standards. They will evaluate and vet available tools, curating a digital workbench of Al tools to ensure developers have a consistent and secure toolset.
- 5. Invest in Your People: Begin training programs now to upskill your workforce, focusing on Al literacy, data analysis, data science, and the new software development and engineering skills required to manage and orchestrate Al systems.



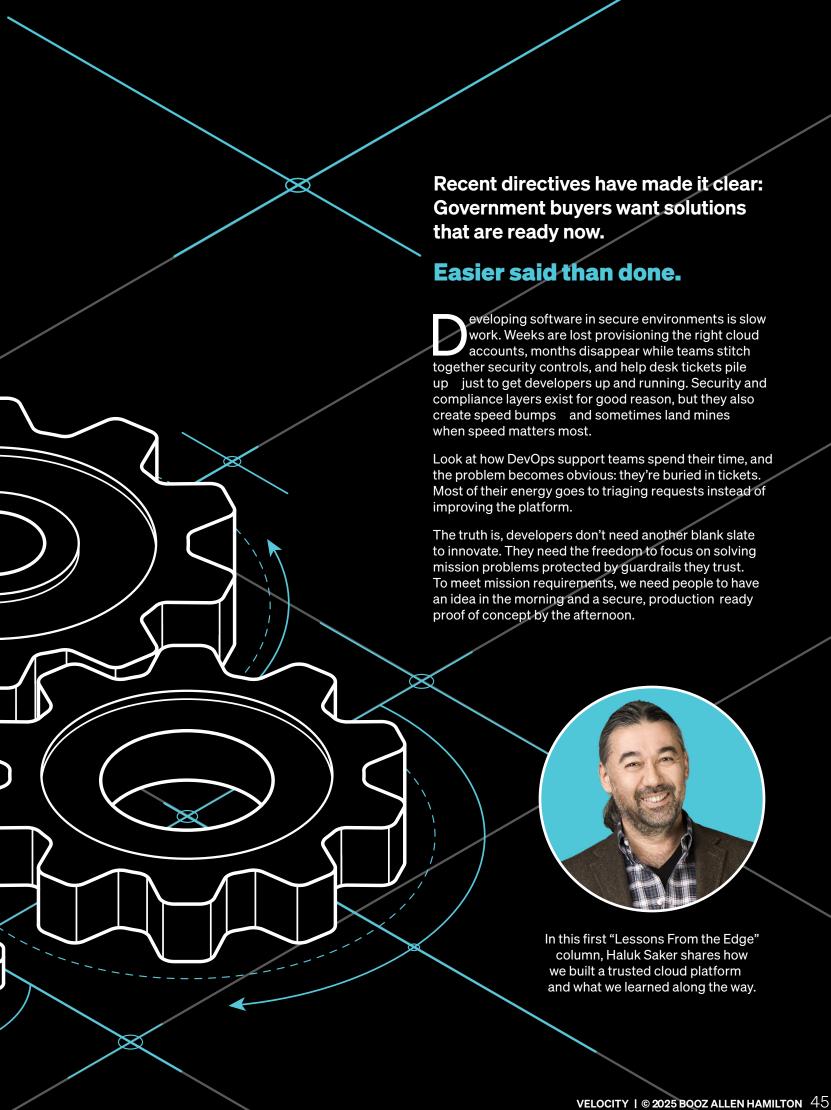
SPEED READ

Al agents can automate nearly every phase of the software development lifecycle (SDLC), enabling federal IT teams to accelerate delivery and modernize legacy systems at unprecedented speed.

A holistic, human-centric, and tool-agnostic approach is crucial for leveraging AI to transform federal software development, ensuring that AI empowers developers rather than replacing them.

Federal agencies can achieve faster, more reliable, and mission-centric software delivery by integrating Al-powered tools, fostering a collaborative innovation environment, and upskilling their workforce.





Breaking the Innovation Bottleneck

Like many companies, Booz Allen has embraced a DevSecOps approach for years. I even co-wrote the Enterprise DevOps Playbook back in 2016. We built centers of excellence across the company that improved quality and productivity. But we still struggled to move innovations seamlessly and readily into client production environments at the accelerated pace we wanted.

We tried being "cloud-agnostic." That was impractical. Government projects struggled to capitalize on the native features of a single cloud—let alone multiple environments—and ended up with the worst of all worlds. We tried DevSecOps factories. They delivered quick wins, but only for one team or one program and couldn't be reused. We tried finding "unicorn engineers" who knew cloud, security, and DevOps end to end. That wasn't sustainable or economical.

We built dozens of effective tools that solved specific problems. But getting them to work together often took more effort than it saved. We needed to scale both horizontally across the company and vertically across delivery.

Finally, we realized: more tools weren't the answer. Sometimes you don't need better tools—you need a better environment.

A New Platform for Innovation

We stopped treating the platform like a backdrop and started treating it like a product.

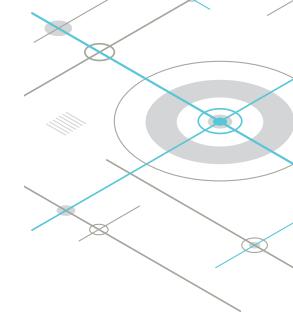
The vision was simple. We imagined a pre-vetted environment developers could step into: controls pre-mapped, tools vetted, pathways cleared. We wanted pre-loaded templates and reusable modules for authorization, data handling, and monitoring to eliminate redundant work. We imagined that developers could stop submitting tickets for setup and instead start coding immediately.

Nobody thought it would work.

Every team had its own portals, contracts, and silos. They were tied to their own ways of doing things. People said, "One size doesn't fit all."

And they were right, but we weren't building a one-size tool. We were building one framework. One that could support many sizes—in our case, defense, national security, civil, and commercial—while still inheriting the same foundation.

Instead of focusing on small differences buried deep in requirements, we designed our platform around shared expectations like our security posture. We all asked, "What does every developer need? And how can we get it to them faster?"



Self-Service Speeds Innovation

Shared platforms don't happen by decree. We had to build bridges across silos and give people a framework to contribute. For us, that meant creating a marketplace model, like an internal app store. Teams could grab what they needed—or add their own modules—as long as they passed our compliance standards.

For developers, the change was night and day. Onboarding that once took three weeks can now be done in 70 minutes. Developers can be coding in a new, fully stocked environment faster than you can bake a loaf of bread. And they can focus on solving mission problems—not configuring buckets or rebuilding the basics.

There are no shortcuts. Compliance is built in, with documentation that makes integration into client environments smoother, safer, and faster. We're talking about production deployments in days, not months, and with far lower risk.

This isn't just theory. We proved it with Booz Allen Vellox Reverser™, an Al-first product to reverse-engineer malware. Our timeline from concept to production was about six months—with no ATO delay.

Because the Vellox Reverser product's components—authorization modules, data-handling blueprints, and monitoring frameworks—now live inside the platform, they can be reused by other teams. Its success didn't just deliver one product; it accelerated the next wave of innovation.

For agencies under pressure to deliver more capabilities, more quickly, and with fewer resources, this approach changes the equation. It clears compliance and infrastructure bottlenecks so breakthrough ideas can move at mission speed. When new Al tools and dev accelerators are being released weekly, waiting six months to start writing code isn't an option.

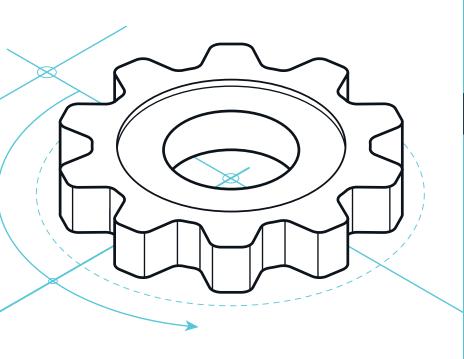
Now we can build as fast as clients and developers can dream.



Looking Back, Looking Forward

Giving up some individual control was uncomfortable at first, but it unlocked a much richer and more sustainable platform. Solutions don't live and die within a single project, so breakthroughs stop being one-offs—they become repeatable.

Every component built in our environment is available in our marketplace for the next project, creating an innovation flywheel. The more we build, the faster innovation spins.



Our platform is designed in accordance with the NIST 800-53, 800-171 and other security control families and accredited standards. Security and compliance are built in from the start, so developers can generate code, troubleshoot, and deploy without worrying about compliance.

Built for developers

Instead of a blank cloud account, developers step into a curated workbench with pre-vetted tools, available through a self-service portal. If a developer needs to open a help desk ticket to get a tool, we consider it a failure. Their time should be spent innovating.

Reusable

Pre-approved golden paths, reusable services, and shared components cut costs and time from every project. Monitoring, alerting, and scaling come standard. Rapid prototyping blueprints let teams spin up experiments quickly—without compromising compliance, even in highly regulated production environments.

Al-enabled

Al works two ways in our trusted platform: developers benefit from Al-assisted coding, pre-configured CI/CD pipelines, and faster troubleshooting. At the same time, embedded Al makes the platform itself smarter and more efficient with every use.

Innovation in Action

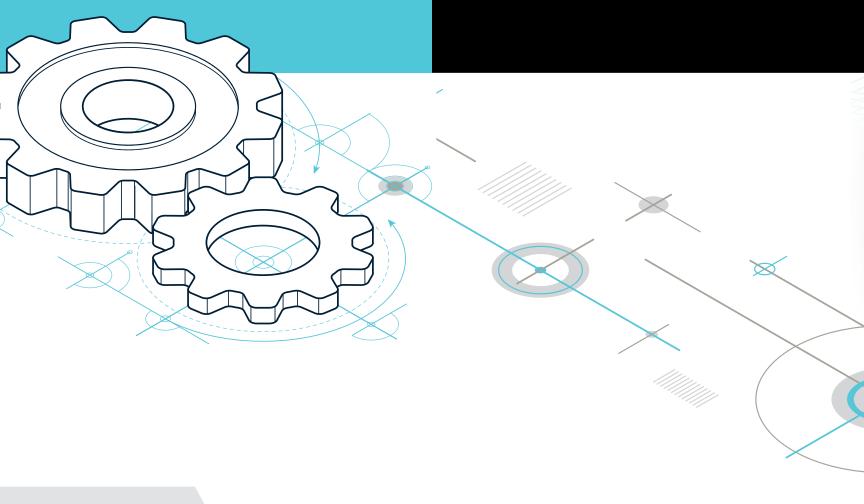
Sound too good to be true? It's not. We used this approach to build and launch **Vellox Reverser**, an Al-first tool that reverse-engineers malware. And the results speak for themselves:

- Deployed in 6 months
 From idea to mission-ready deployment
- Zero Authority-to-Operate delays
 Continuous compliance eliminated lead time
- Accelerated reuse
 Marketplace makes code reuse real
- Endless possibilities
 Every innovation seeds the next

About the Author

Haluk Saker has spent over three decades modernizing digital systems for government clients. He led the redesign of Recreation.gov, transforming it into one of the federal government's most-used platforms, and now heads innovation efforts at Booz Allen.

The author wishes to recognize Shums Hoda, Steven Terrana, Stephen Mao, Matt Lisco, Andrew Smallwood, Martin Folkoff, Madhu Ravi, Lizzie Salita, Goran Osim, and Regina Trapilo whose contributions were invaluable to this effort.



SPEED READ

The platform is the product. Once we treated the platform like a core product—with its own roadmap, support team, and governance—we gained the discipline to keep it secure, modern, and optimized for every project.

Quality isn't negotiable—it's engineered. Moving fast doesn't mean cutting corners. Our marketplace model ensures every contribution is security-reviewed and compliance-tested before it's available. That lets developers move quickly without sacrificing any rigor.

The platform's marketplace model fosters collaboration across teams, enabling code reuse and creating an innovation flywheel, which accelerates the delivery of mission-critical solutions like the AI-first Vellox Reverser.

Vellox Reverser A Booz Allen® Product



Outpace the threat. Reverse your unknowns.

Vellox Reverser™ is an Al-first, cloud-native malware reverse engineering product designed to radically speed up threat response. Today's malware doesn't wait. It evolves in real time—gaining ground before cybersecurity teams can react. Spending weeks to counter threats is not an option. Booz Allen Vellox Reverser changes the game. Vellox Reverser automates the process, completing an exhaustive analysis of complex and evasive malware to deliver deep insights and actionable reports with comprehensive countermeasures in just minutes.

With Vellox Reverser, security teams move faster, respond smarter, and outpace the most advanced threats.



Get started at www.boozallen.com/vellox-reverser.

Accelerating Mission Outcomes Through Tech Investment

Matt Calderone, Chief Financial Officer



has been the biggest technology topic this year. It's at the center of our national security, infrastructure, workforce, and innovation ecosystems. The White House's Al Action Plan emphasized the global strategic importance of American Al leadership. It's clear that Al is the biggest tech wave of the next few decades and that the promise and capabilities of this tech have captured our collective imagination.

Booz Allen has been preparing for this moment since 2014. We knew the potential impact Al would have on everything from defense tech and public safety to critical mineral discovery and the space domain. We invested in ourselves and became the leading provider of Al to the federal government.

We also broadened the aperture of how Booz Allen views tech and found ways to use this wider field of vision to ensure American Al leadership. The applications for Al may well be limitless, and we are eager to combine our passion for bringing the best tech to market and to mission while pushing the boundaries of discovery.

Our perspective has driven the three focus areas that have helped us stay ahead of the surge in demand: engagement with the tech ecosystem, collaboration with commercial tech companies, and investments that fueled promising solutions.

Tech Partnerships

First, we knew no one company could drive Al innovation alone, which is why we've spent decades deeply engaging with companies that share our commercial-first approach. We've worked with leading companies—including NVIDIA, Meta, and Databricks—and we've delved into the many applications for Al, as you saw in the article "Agentic Development Decoded: Hard-Won Lessons from the Trenches."

We've also made investments in promising tech with real-world applications. Booz Allen Ventures supported Reality Defender and their development of Al tools that combat identity-based fraud and deep fakes. Our company has a longstanding and robust relationship with Shield Al, including combining our mission engineering expertise with their Hivemind platform to advance autonomy in defense tech. There are many ways Al can be applied—next-generation communications networks, new breakthroughs in intelligence gathering, and simulations of the physical world—and we're committed to the partnerships that fuel that innovation.

Commercial Collaboration

Next, we've helped companies of all sizes build, develop, and implement faster. We know that co-creation and collaboration bring greater speed to outcomes, and we're working with everyone from startups to hyperscalers. Combining Booz Allen tech and expertise with that of emerging tech companies has led to game-changing, positive impacts for our customers—and for the vital missions they serve.

Sometimes we come across exceptional tech that's ready for immediate implementation, such as Hidden Layer's cybersecurity platform for AI/ML models. Other times, the tech needs to be tuned or combined with other solutions to get it ready for delivery. That's when we open our doors,

"Looking ahead, Booz Allen plans to take these efforts further and faster than ever before. In July, we announced that we are increasing our corporate ventures commitment to \$300 million."

share our tech talent and mission insights, and leverage our intellectual property to complement third-party offerings. Working collectively means that we innovate faster, strengthen the entire ecosystem, and get the right tech into the hands of those who need it most.

Booz Allen Ventures

Finally, in some cases, we bring in our investment dollars to fuel companies with great promise. Booz Allen Ventures uncovers these companies, helps them mature their offerings for new markets, and accelerates their solutions' time to impact. We currently have 17 companies in that program, and we're seeing great results. We're helping Albedo get their very low earth orbit satellites into space, and those satellites will provide new sources of ultra-high resolution commercial imagery. We're also working with Second Front to get their DevSecOps platform into national security missions. Over time, we build on those relationships; we just expanded our relationship with SEEQC to use their chipsets in quantum software development.

Looking ahead, Booz Allen plans to take these efforts further and faster than ever before. In July, we announced that we are increasing our corporate ventures commitment to \$300 million. We'll invest in crucial areas like Al, cyber, quantum, defense, and deep tech while also exploring opportunities in more specific priority areas, such as robotics, energy resilience, and reindustrialization. We're also leveraging our strong relationships and thinking in a more expansive way about how we can engage with the venture-backed tech ecosystem. As we invest alongside leading companies and expand our awareness of the amazing tech that will be coming to market over the next decade, we're looking for ways to combine our efforts and deliver the tech superiority that keeps America safe and prosperous.

America needs to use new technology to solve our collective challenges and ensure our leadership on the global stage. The mission is expanding, and it's our responsibility to build, develop, and deliver solutions that matter. Our part in that mission is to keep innovating at the edge of this ever-evolving tech landscape and keep partnering across the ecosystem. That is how we will ensure that our company, our partners, and our nation thrive.



BOOZ ALLEN AUTHOR INFORMATION

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

Andrew Bolka

bolka_andrew@bah.com

Chris Christou

christou_chris@bah.com

Jennifer Congdon

congdon_jennifer@bah.com

Bassel Haidar

haidar_bassel@bah.com

Meghan Hauser, Ph.D.

hauser_margaret@bah.com

Greg Kacprzynski

kacprzynski_gregory@bah.com

Michael Martoccia

martoccia_michael@bah.com

Drew Massey

massey_james@bah.com

David Perry

perry_david2@bah.com

Haluk Saker

saker_haluk@bah.com

Jennifer Sheppard

sheppard_jennifer@bah.com

Munjeet Singh

singh_munjeet@bah.com

Ryan Stone

stone_ryan@bah.com

Randy Yamada, Ph.D.

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:** John Conley

Editorial Strategy: Logan Gibson and

Danielle Epting

Contributing Editors: Pat Fitzgerald, Jim McDermott, Suzanne Sandberg,

and Elana Akman

Promotion: Jennifer Bundschu

ART & DESIGN

Creative Director: Chris Walker

Art Director: Brody Rose

Graphic Designers: Sarah Bell and

Matt Dunham

Video: Angshuman Mojumder

WEB DEVELOPMENT

Creative Director: Shauna Hidlebaugh Web Design and UX: Brianne Cochran

and Ashley Lowe

MARKETING & COMMUNICATIONS Head of Brand & Marketing:

Abby Vaughan

Media Director: Jessica Klenk

For media inquiries, please contact:

media@bah.com

ACKNOWLEDGEMENTS

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 the 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, customers, and their missions to succeed, accelerating the nation to realize our purpose:

Empower People to Change the World®.

To learn more, visit BoozAllen.com.

GE*SECURITY*DEFENSE_TECH*AGENTIC_AI*ENGININ THIS COUNTRY, WE DON'T FOLLOW. WE LEAD. WE DON'T WAIT. EERING*LEADERSHIP*PASSIONATE_SERVICE*CODE WE MOVE FIRST. WE DON'T WATCH HISTORY HAPPEN. WE MAKE IT. UTION*DIGITAL_TWIN*COLLECTIVE_INGENUITY* WHILE OTHERS TALK ABOUT AI...WE PUT IT UNDER THE OCEAN. OUTCOMES*PATRIOTISM*ADVANTAGE*PIONEER*ON THE FRONTLINES. AND UP IN SPACE. THAT'S RIGHT—SPACE. *TRANSFORMATION*FUTURE*FORWARD*BREAKHOW DO WE SOLVE OUR NATION'S BIGGEST CHALLENGES? SPEED. THROUGH*CO-CREATION*CHAMPION'S_HEART*MIPRECISION. HONOR. NO MANUAL. NO ROAD MAP. NO PROBLEM. DEVELOPING 5G VR MILITARY TRAINING. SECURING BORDERS WITH ARTIFICIAL INTELLIGENCE AND TRAILBLAZING TECH. BECAUSE REAL LEADERSHIP? IT'S ABOUT BUILDING WHAT NOBODY ELSE CAN AND CODING SO WE CAN'T LOSE. ENGINEERS. CODERS. SPECIAL OPS. WINNING EVERY TIME. AND MAKING AMERICA STRONGER. SAFER. FASTER. IT'S WHO WE ARE. IT'S WHAT WE STAND FOR. IT'S WHAT WE'RE MADE OF. BOOZ ALLEN <ITS_IN_OUR_CODE>



