Carnegie Mellon University
Software Engineering Institute

2023 YEAR IN REVIEW

Always focused on the future, the Software Engineering Institute (SEI) advances software as a strategic advantage for national security. We lead research and direct transition of software engineering, cybersecurity, and artificial intelligence technologies at the intersection of academia, industry, and government. We serve the nation as a federally funded research and development center (FFRDC) sponsored by the U.S. Department of Defense (DoD) and are based at Carnegie Mellon University (CMU), a global research university annually rated among the best for its programs in computer science and engineering.

> The *2023 SEI Year in Review* highlights the work of the institute undertaken during the fiscal year spanning October 1, 2022, to September 30, 2023.

A Message from the Director and Chief Executive Officer

With the emergence of artificial intelligence, our society faces a time of great possibility but also great risk. The latest innovation is generative AI, which has appeared nearly everywhere people produce information digitally. The technology has tremendous potential, significant flaws, and few norms around its usage.

Like other challenges we've tackled, AI involves software, and software has faced other uncertain times. When the SEI was founded in 1984, software engineering lacked the necessary rigor, so we developed best practices for engineering dependable systems. When unsecured networked systems threatened the Department of Defense and businesses around the world, we formed the CERT Division to launch new defenses against cyberattacks.

Now, with AI, we don't yet have a functional understanding of its end state or how we can best use the technology. Gaining this understanding is our next big challenge as we drive toward knowing the roles of AI, the rules that govern them, and how those roles apply to the DoD.

This challenge is not unlike the first ones we took on in the 1980s, and in the past few years we have been applying a similar approach with AI engineering. More recently, we've been testing generative AI's initial usefulness and trustworthiness in different contexts important to our DoD sponsor (p. 7). Over the past year, we expanded our capability to tackle AI risks more broadly with the AI Security Incident Response Team (AISIRT, p. 5) and the Center for Calibrated Trust Measurement and Evaluation (CaTE, p. 19).

AI may get headlines, but the SEI's other core competencies in cybersecurity and software continue to support our national security mission. Our roadmap for software engineering research and development has led to important community discussions on U.S. leadership in software and AI engineering (p. 13). Cybersecurity research has enhanced vulnerability prioritization at the Cybersecurity and Infrastructure Security Agency (p. 8) and brought zero trust practices to Army tactical networks (p. 14). Our involvement with architecture and systems modeling languages led to a breakthrough in the development of safety-critical software systems (p. 24).

AI surely won't be the last uncertain technology we face in our unique mission to advance the art of software engineering. We'll keep adapting to emerging software developments as they arise, drawing on our dedicated, driven researchers and collaborating with government, industry, and academia to find the best path forward and continue delivering impactful solutions for the DoD.

Paul Nielsen

Execution Strategy

The SEI facilitates the transfer of research results to practice in Department of Defense (DoD) programs, the Office of the Secretary of Defense's science and technology initiatives, and non-DoD U.S. government organizations where improvements will also benefit the DoD. In doing so, we gain deeper insight into mission needs—insight that forms the basis for new research. In addition, we transition matured technologies more broadly to defense industrial base organizations and others in the DoD supply chain.

We execute applied research to drive systemic transition of new capabilities for the DoD. Our deep understanding of DoD needs and of the state of the art inform our selection of challenges in software, cybersecurity, and artificial intelligence.

To validate research and development concepts, we rapidly iterate with the research community and select mission partners. The results typically impact a single agency. We then scale the concept to multiple agencies and domains by iterating with additional mission partners based on their timing and needs. Finally, we engage policy agencies and industry partners and build the DoD's awareness of and capacity for the solution to create DoD-wide capability.

Our multidisciplinary approach informs prototype tools, innovative solutions, and input for our sponsor's policy decisions about software and related technologies. Through ongoing work and communication with customers, the SEI identifies priority areas for further research and development. We combine our body of knowledge with external material and systems engineering to deliver quantitative impact to U.S. government organizations, DoD organizations, and DoD end users.

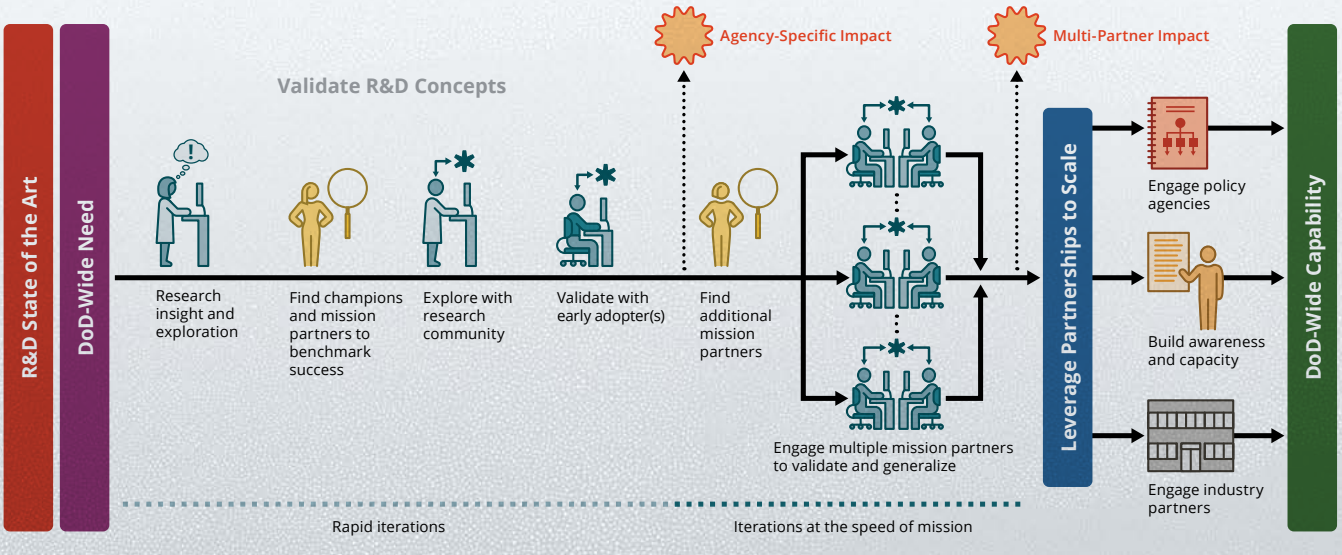

Funding Sources

In fiscal year 2023, the SEI received funding from a variety of sources in the DoD, civil agencies, and industry.

Table of Contents

A Message from the Director and Chief Executive Officer . . . 1

Execution Strategy . . . 2

News Briefs . . . 4

National Security Demands Rigor, Not Rush, for Generative AI . . . 7

CISA Adapts Innovative SEI Approach to Transform Vulnerability Management Landscape . . . 8

CMU Collaborations Enhance Outcomes for U.S. Government . . . 9

Reducing the Risk of UEFI's Hidden Security Challenges . . . 10

SEI Open Source: From Research to Community . . . 11

Workshop Identifies Critical Needs for U.S. Leadership in Software Engineering, AI Engineering . . . 13

Bringing Zero Trust Practices to Army Tactical Networks . . . 14

SEI Support for Long Range Standoff Program Spurs New Engagements . . . 16

Assuring Trustworthiness of AI for Warfighters . . . 19

Pathfinding Project Explores Large Language Models for the Intelligence Mission . . . 20

Supporting the Human and Technical Elements of Responsible AI for National Defense . . . 22

Extending SysML V2 with AADL Concepts to Support Engineering and Certification of Safety-Critical Systems . . . 24

Professional Organization Participation Boosts SEI Staff and Mission . . . 26

Leadership . . . 28

Key Publications . . . 32

2023 Featured Research Teams . . . 38

NEWS BRIEFS

SEI Team Leads First Independent Study on Technical Debt in Software-Intensive DoD Systems

Section 835(b) of the 2022 National Defense Authorization Act (NDAA) mandated an independent study on technical debt in software-intensive Department of Defense (DoD) systems. The Office of the Under Secretary of Defense for Acquisition and Sustainment chose the SEI, a recognized leader in the practice of managing technical debt, to lead this work. A team of SEI researchers conducted the independent study and produced a report on its findings and recommendations.

The study included a literature review, a review of SEI reports developed for program stakeholders, and deep dives on program data from SEI engagements with DoD programs. Researchers also interviewed 16 organizations—11 from the DoD, 4 from industry, and a federal government agency—using 10 study elements specified in Section 835(b).

The study's broad-ranging findings give a first-ever snapshot of the state of technical debt in DoD programs. Though many actively manage technical debt, the practice is often inconsistent and deprioritized. The study's recommendations cover best practices, policy and guidance, training, metrics, financial management regulation, tools, and further research. The final report was delivered to Congress in December 2023.

> Download the report at **insights.sei.cmu.edu/library/congressional-report-section-835-technical-debt-cmu-sei-2023-tr-003/**.

President's Cup Competition Expands Access to SEI Cybersecurity Simulations

For four years, the Cybersecurity and Infrastructure Security Agency's (CISA's) President's Cup Cybersecurity Competition has identified the federal government's best cyber professionals and teams. In 2022, the SEI once again developed challenges for the event and orchestrated it in partnership with CISA, this time adding a new multiplayer video game called *Cubespace* and making it available after the event.

Cubespace uses an immersive science-fiction story to present highly realistic cybersecurity simulations for a variety of cyber work roles such as cyber defense, incident response, and software developer. "We build the challenges to reflect real-world situations," said the CISA President's Cup Cybersecurity Competition lead, Michael Harpin. "The video game is a fun and unique addition to the President's Cup."

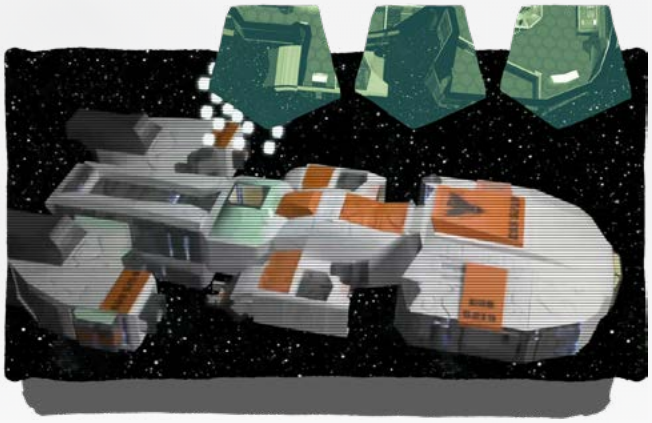

The fourth competition saw more than 1,200 participants from the Department of Defense and other federal agencies. Since then the event has offered its training to more cyber professionals than ever. Federal employees can now hone their cyber skills with nearly 200 challenges from the past President's Cups in an online Practice Area.

> *Cubespace*'s source code is freely available at **github.com/cisagov/prescup-challenges**.

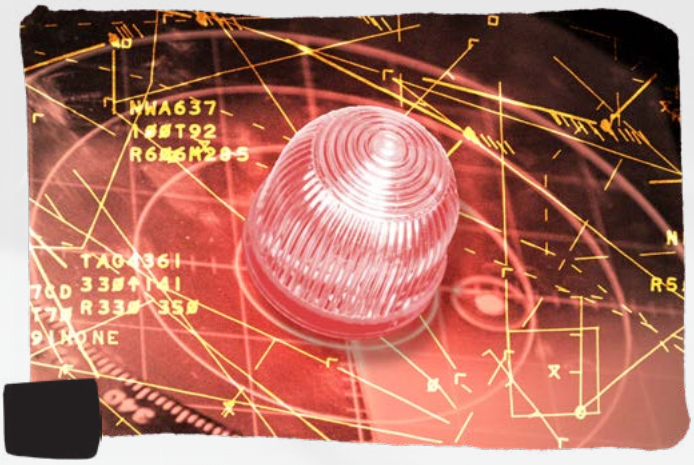

SEI Establishes First AI Security Incident Response Team

Artificial intelligence (AI) systems can perform amazing feats, but improper deployment or deliberate misuse of AI presents great risks. In 2020, SEI experts published the first machine learning (ML) vulnerability note as well as guidance for managing ML vulnerabilities. Three years later, they started evaluating a new crop of AI security incidents. Amid a rapid proliferation of AI in 2023, the SEI leveraged its expertise in cybersecurity and AI to field the first AI Security Incident Response Team (AISIRT).

AISIRT analyzes and responds to AI and ML threats and security incidents and researches incident analysis, response, and vulnerability mitigation. The team will also coordinate with Carnegie Mellon University experts to research new techniques that assure the security of AI platforms. AISIRT has a broad scope that covers AI systems for all purposes, from consumer applications to defense, national security, and critical infrastructure.

"Our research in this emerging discipline reinforces the need for a coordination center in the AI ecosystem to help engender trust and to support advancing the safe and responsible development and adoption of AI," said SEI Director and CEO Paul Nielsen.

> **AI attacks or vulnerabilities in AI systems may be reported to AISIRT at kb.cert.org/vuls/report/.**

SEI Quantum Experts Join Pittsburgh Computing Research Organizations

Two researchers in the SEI's AI Division joined Pittsburgh computing research organizations in 2023. Senior researcher Jason Larkin and software developer Daniel Justice became members of the Pittsburgh Quantum Institute (PQI) and the Pittsburgh Supercomputing Center (PSC).

Founded in 2012 and funded by the University of Pittsburgh, the PQI organizes outreach activities such as seminars and poster sessions, sponsors graduate students, and provides financial support to quantum research led by university affiliates. Larkin and Justice joined a network of more than 130 PQI faculty and staff researchers.

The PSC is a high-performance computing and networking center run jointly by Carnegie Mellon University and the University of Pittsburgh. As courtesy researchers, Larkin and Justice work with PSC to build infrastructure capable of meeting the resource requirements needed for quantum computer research.

Larkin researches quantum advantage, or what applications and problems quantum computing is better at than classical computing. Justice studies software architecture to support quantum-classical hybrid systems. Their PQI and PSC participation helps the SEI fulfill its mission to scout new technology for the Department of Defense.

National Security Demands Rigor, Not Rush, for Generative AI

EDITORIAL *Tom Longstaff*

From chatbots to generated images and voices, generative artificial intelligence (AI) permeated the general consciousness in 2023. The fusion of large language models (LLMs) with user-friendly interfaces has manifested in applications for everything from entertainment to specialized professional domains. Whether you fear it or embrace it, generative AI is already changing how we live.

The SEI's mission, as the nation's only federally funded research and development center dedicated to software, is to establish and advance software as a strategic advantage for national security. On behalf of the U.S. Department of Defense (DoD), we are conducting research and development on generative AI not simply because it is software, but because it affects how our nation makes software, secures its cyber systems, and grows AI overall.

Generative AI is one of several emerging technologies that we enable the DoD to harness in a trustworthy, responsible, timely, and affordable manner. Last year we conducted a foundational study describing and categorizing exemplar LLM use case archetypes, concerns, and remedies. The resulting paper, *Assessing Opportunities for LLMs in Software Engineering and Acquisition*, provides a novel, structured way for decision makers in organizations with large software operations to assess the fitness of LLMs for addressing software engineering and acquisition needs.

Another SEI study, *Demonstrating the Practical Utility and Limitations of ChatGPT Through Case Studies*, explores how generative AI might enhance data science processes, training and education, literature reviews, and organizational strategic planning. Some of the researchers also experimented with ChatGPT to detect malicious code and vulnerabilities in source code, simulated realistic network activity using LLM directives, and explored deepfake video detection.

The researchers of the SEI project *A Retrospective in Engineering Large Language Models for National Security* established a test LLM to investigate the feasibility, cost, and trustworthiness of using generative AI in highly sensitive intelligence environments. The SEI also hosted a workshop that discussed DoD use cases, challenges, and needs for generative AI.

In our collective excitement to create new AI tools, generative or not, we must continue to maintain discipline at the end of the development lifecycle. Testing and evaluation remain essential for the safe and effective deployment of AI in defense platforms, as exemplified in a DoD-sponsored study of the Department of the Air Force that I had the honor to co-chair. The SEI is also researching rapid assurance of large-scale software systems, which can address potential new risks to system integrity and availability introduced by AI components. These efforts are steps on the roadmap laid out in the SEI's *National Agenda for Software Engineering Research and Development*, particularly in three of its focus areas: AI-augmented software development, assuring continuously evolving systems, and engineering AI-enabled software systems.

Throughout even the most narrowly scoped studies, the SEI keeps the big picture in mind. Our ongoing work in AI engineering and responsible AI will continue to inform our generative AI research. Blog posts, podcasts, and live webcasts from our experts have offered high-level takes on LLMs: hype versus reality in software engineering and development, harnessing their power for economic and social good, leveraging tools, and critically assessing outputs.

It is still early days for generative AI, and its place in our lives will change as this technology rapidly evolves. But it certainly has a role in maintaining our nation's strategic advantages in security and the global economy. As the national security enterprise is realizing generative AI's potential to transform its mission, the SEI's dedicated researchers are approaching generative AI with the same rigor, expertise, and collaboration we have brought to software engineering and cybersecurity for nearly 40 years.

CISA Adapts Innovative SEI Approach to Transform Vulnerability Management Landscape

Most organizations struggle to prioritize responses to the tens of thousands of cyber vulnerabilities discovered each year. In a recent push to transform the vulnerability management landscape, the Cybersecurity and Infrastructure Security Agency (CISA) adapted and promoted the SEI's Stakeholder Specific Vulnerability Categorization (SSVC) approach.

The SEI applied nearly four decades of experience researching vulnerability response when it developed SSVC in 2019. This conceptual tool for prioritizing vulnerabilities emphasizes stakeholder perspectives, which are often missing from vulnerability data. Using SSVC, vulnerability analysts gather human input to incorporate an organization's particular attributes and values rather than rely on stakeholder-agnostic indicators such as the long-standing Common Vulnerability Scoring System (CVSS) base score.

SSVC captures risk owners' perspectives on vulnerabilities before analysts address them, enabling analysts to process more vulnerabilities, a benefit that drew CISA's attention in 2020. The SEI and CISA developed a custom SSVC decision tree to help CISA better support its U.S. federal civilian executive branch; state, local, tribal, and territorial governments; and critical infrastructure stakeholders. These organizations can also use SSVC themselves to efficiently decide the best responses that align with stakeholder values and justify decisions that affect other government bodies.

In CISA's SSVC decision tree, an organization leverages external vulnerability data and knowledge of its own environment to evaluate a series of decision points about a given vulnerability: exploitation status, technical impact, automatability, mission prevalence, and impact on public well-being. The answers lead to a vulnerability prioritization recommendation—track, closely monitor, attend to, or act on—that considers the organization's risk appetite and other attributes.

In November 2022, CISA announced critical steps that organizations should implement to help them manage the number and complexity of cyber vulnerabilities. The use of SSVC, which CISA supported by releasing an SSVC web page, guide, and online calculator, was one of these three steps.

"With these advances," wrote CISA's executive assistant director for cybersecurity Eric Goldstein in a blog post about the campaign, "we will make necessary progress in vulnerability management and reduce the window that our adversaries have to exploit American networks."

Learn more about the SEI's version of SSVC at **certcc.github.io/SSVC**.

Photo: CISA

"With these advances we will make necessary progress in vulnerability management and reduce the window that our adversaries have to exploit American networks."

ERIC GOLDSTEIN, Executive Assistant Director for Cybersecurity, CISA, U.S. Department of Homeland Security

CMU Collaborations Enhance Outcomes for U.S. Government

Tackling the nation's toughest challenges in software engineering, cybersecurity, and artificial intelligence requires a rare breadth and depth of knowledge. As an integral part of Carnegie Mellon University (CMU), the SEI has access to a talented and experienced pool of domain experts. Research and development for the Department of Defense and other U.S. government agencies was enhanced by the SEI's collaboration with CMU faculty, staff, and students on the following projects in 2023:

- **Automatic Detection of Malicious Code— Ruben Martins,** assistant research professor in the Computer Science Department, developed large language model approaches to automatically determine if system application programming interface (API) functions could leak sensitive information, thereby easing the burden on security analysts.
- **Building a Security Operations Center (SOC) Knowledge Base and Ontology— Travis Breaux,** associate professor of computer science and director of the Requirements Engineering Laboratory, guided the development of a first-of-its-kind ontology of SOC expert knowledge and the description logic to turn that ontology into a functioning tool for operationalizing new SOC capabilities.
- **Risk-Aware Adaptive Moving Target Defense (MTD)— Ehab Al-Shaer,** distinguished career professor in the Software and Societal Systems Department, investigated appropriate risk metrics and defense vectors for a capability to defend systems against advanced persistent threats proactively.
- **Portable High-Performance Inference on the Tactical Edge (PHITE)— Tze Meng Low,** associate research professor with the Department of Electrical and Computer Engineering, and his graduate students Upasana Sridhar, Nicholai Tukanov, and Elliott Binder, helped develop Software for Machine Learning Libraries (SMaLL), an open source machine learning (ML) framework for low-power devices. The PHITE project enables tactical edge devices to use ML more efficiently.
- **Automatic Detection of Stakeholder Assumption Mismatches in ML System Development— Christian Kästner,** associate professor and director of the software engineering PhD program, and his doctoral students Nadia Nahar and Chenyang Yang, helped identify stakeholder collaboration challenges and tools to test model production readiness.
- **2022 President's Cup Cybersecurity Competition—** Faculty and students of the Center for Transformational Play (CTP) and Entertainment Technology Center (ETC) developed the video game *Cubespace*, based on challenges designed by the SEI, for the finals of the Cybersecurity and Infrastructure Security Agency's annual contest for federal cybersecurity practitioners.

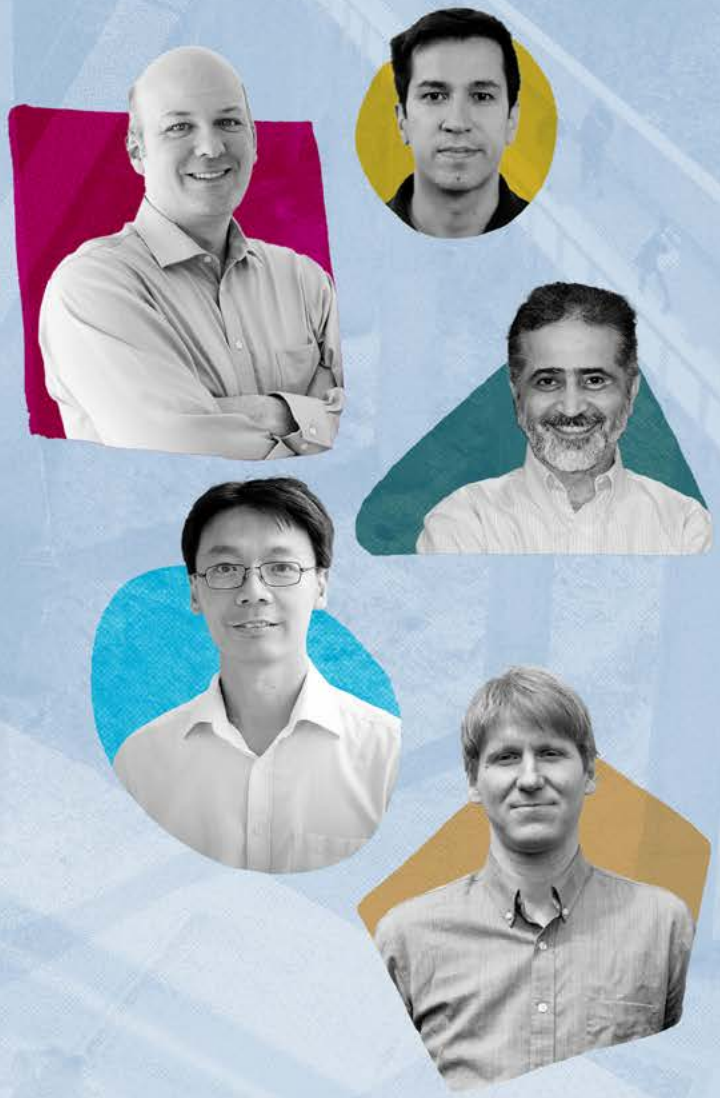

Reducing the Risk of UEFI's Hidden Security Challenges

Memory management vulnerabilities fundamentally impacting more than 280 million computers were discovered in late 2022 due to various implementations of software built using the Unified Extensible Firmware Interface (UEFI) standard. SEI research has informed both industry and government responses to this deeply rooted problem.

UEFI-based software, installed as part of firmware, is crucial for initializing computer hardware at startup and managing the ongoing interaction between hardware and the operating system (OS). UEFI-based software, often invisible to users, is an appealing target for attackers. Those who can exploit UEFI software vulnerabilities can establish persistence and remain invisible to most security software and often even to the OS.

After UEFI-exploiting malware such as BlackLotus was confirmed in early 2023, SEI senior information security architect Vijay Sarvepalli led a study of UEFI software security. Leveraging the SEI's experience in coordinating identified vulnerabilities and establishing secure coding standards, Sarvepalli produced five recommendations for securing the UEFI ecosystem, detailed in a 2023 white paper.

Both the Cybersecurity and Infrastructure Security Agency (CISA) and the National Security Agency (NSA) have cited Sarvepalli's work to enhance their research and improve the UEFI ecosystem. The research has had the largest impacts on improving UEFI memory management vulnerability and UEFI digital signature and revocation maturity.

Exploits like BlackLotus continue to target UEFI software. However, the SEI's work with the diverse stakeholders in UEFI's supply chain has improved UEFI security throughout the vulnerability management lifecycle. The CERT Coordination Center has engaged with vendors and researchers on mitigations across the UEFI ecosystem, such as memory management, digital signature and revocation, privilege separation, supply-chain security, and automated patching.

The SEI's continued research and outreach through supply-chain channels has been in the spirit of Executive Order 14028 on improving the nation's cybersecurity. As a federally funded research and development center, the SEI is in a unique position as a neutral third party to raise public awareness and impartially influence vendors to resolve this problem.

"The promise of this work is to continue to make this type of hidden software more secure," said Sarvepalli. "We can shine light on the security needed not only for UEFI, but other critical, invisible software."

"We can shine light on the security needed not only for UEFI, but other critical, invisible software."

VIJAY SARVEPALLI, Senior Information Security Architect, SEI CERT Division

SEI Open Source: From Research to Community

SEI researchers often produce software whose utility extends beyond its original scope and application. The SEI frequently offers these tools to the community and has made 124 repositories freely available through its GitHub site.

Popular and Recently Released Open Source Repositories

- The **CERT Kaiju** static binary analysis framework is an extension to the National Security Agency's Ghidra reverse engineering platform. Kaiju implements many features found in the Pharos framework, which facilitates the automated analysis of binary programs and detection of operating system paradigms. Kaiju implements Pharos's object oriented analysis, function hashing, and malware analysis tools; provides additional tools to perform binary path analysis; and includes several integrated utilities and services to support reverse engineering tasks in Ghidra.
- **GHOSTS** simulates what anyone might do at a computer: create documents, browse websites, and download files. While GHOSTS was originally designed for cyber training and exercises, it is now also used for many other scenarios in which realistic activity on a computer is needed.
- The **Source Code Analysis Laboratory (SCALe)** helps source code analysts audit source code for security flaws and enables them to combine results from multiple tools into one interface.
- The **TEC Machine Learning (ML) Mismatch Detection Tool** helps developers of ML-enabled software capture key information about ML system elements from stakeholders in a set of descriptors, compares information in these descriptors, and flags any mismatches or missing information to help resolve problematic differences early in development.
- The **Juneberry** platform facilitates ML experimentation by helping users train and compare ML models that may have different architectures, data sets, and/or hyperparameters. By automating training and evaluation, Juneberry can improve robustness and security, qualities foundational to artificial intelligence engineering.

Beyond Open Source: SEI Installable Software Packages

The SEI also makes useful installable software packages freely available to the community. These packages include the System for Internet-Level Knowledge (**SiLK**), a collection of traffic analysis tools for large network analysis, and the Linux Incident Response and Forensics Tools Repository (**LIFTeR**).

> To learn more about the SEI's open source repositories and installable software packages, visit insights.sei.cmu.edu/software-tools/.

"Now that fast-moving AI technologies are beginning to affect all aspects of society, it is more important than ever to bring together cross sections of academia, industry, and government to inform a community strategy for building and maintaining U.S. leadership in software engineering."

ANITA CARLETON, Director, SEI Software Solutions Division

Workshop Identifies Critical Needs for U.S. Leadership in Software Engineering, AI Engineering

Artificial intelligence (AI) and advancing technologies in software development will bring great change to future software systems, and the software community needs new approaches to navigating these opportunities and challenges. In June 2023, the SEI cohosted a workshop with the Networking and Information Technology Research and Development (NITRD) program of the White House Office of Science and Technology Policy (OSTP). The workshop identified five critical needs and priorities for building and maintaining U.S. leadership in software engineering and AI engineering.

Workshop participants included thought leaders from federal research funding agencies, research laboratories, mission agencies, and commercial organizations. They explored areas critical to multidisciplinary research for the future of software engineering.

Participants identified a need to invest in the right technical domains and to improve mechanisms for collaboration among academia, industry, and the federal space. Five major themes emerged:

1. AI is transforming the software engineering process and how we engineer software systems. The increasing symbiosis of humans and machines is transforming every phase of the software development lifecycle.
2. While generative AI has reached a level of sophistication that may seem to resemble human intelligence, it is considerably harder to determine the level of trust that should be placed in the outputs.
3. Redefining the discipline of software engineering to encompass the use of new technologies (including but not limited to generative AI) is imperative, along with rethinking the associated curricula, tools, and technologies. This effort is key to designing and building, evolving, and evaluating trustworthy software systems in a responsible, ethical way.
4. New technologies, including generative AI, seem to hold the promise of making almost everyone a programmer. As a result, AI literacy and the development of new skills are needed throughout the workforce.
5. The use of AI tools such as large language models can mask the tradeoffs being made between the functionality of software systems and their safety and security. Research is needed to identify and make explicit the key engineering tradeoffs being made during the design, development, training, testing, and authorization of systems that include AI components.

These themes align with three areas of future research—AI-augmented software development, assuring continuously evolving systems, and engineering AI-enabled software systems—recommended in the 2021 SEI report *Architecting the Future of Software Engineering: A National Agenda for Software Engineering Research and Development*. The workshop marked a milestone for the SEI's National Agenda, which calls for the software engineering community to come together around rapidly changing challenges.

"Now that fast-moving AI technologies are beginning to affect all aspects of society, it is more important than ever to bring together cross sections of academia, industry, and government to inform a community strategy for building and maintaining U.S. leadership in software engineering," said Anita Carleton, SEI Software Solutions Division director and the workshop's co-organizer.

> **Download the summary of the *U.S. Leadership in Software Engineering & AI Engineering: Critical Needs & Priorities Workshop* at insights.sei.cmu.edu/library/us-leadership-in-software-engineering-ai-engineering-critical-needs-priorities-workshop-executive-summary/.**

Bringing Zero Trust Practices to Army Tactical Networks

The tenets of zero trust cybersecurity remove implicit trust within the network and shift security from network perimeters to network users, assets, and resources. This approach is a set of best practices that were initially focused on enterprise networks. Though the tactical networks of warfighters in the field are very different, these personnel will soon be required to follow zero trust principles. The SEI is helping the U.S. Army prepare to implement zero trust in a tactical environment for the first time.

In late 2022, the Department of Defense (DoD) released its *Zero Trust Strategy,* which envisions the implementation of a department-wide zero trust cybersecurity framework by fiscal year 2027. The DoD enterprise settings implicated in this strategy have stable infrastructure and network connections, so it is reasonable to authenticate network users, assets, and resources—one of the main zero trust techniques.

Soldiers in the field, however, cannot always enter a password or scan their fingerprint. They also often operate in denied, disconnected, intermittent, or limited (D-DIL) network environments, making authentication data difficult to pass. The U.S. Army Program Executive Office (PEO), Command, Control, and Communications-Tactical (C3T) approached the SEI in 2023 for help implementing zero trust principles in its tactical networks.

"In a tactical situation, warfighters are used to pushing a button and having things work," said Tim Morrow, the SEI's situational awareness technical manager and technical lead on the Army PEO C3T engagement. "Zero trust is about assessing the risk before you take an action." Any zero trust implementation must balance additional security against rapid capability.

Because zero trust is a set of institutional practices, it cannot be accomplished by a device or even a single cybersecurity vendor. This reality was highlighted at the 2022 SEI event Zero Trust Industry Days, where vendors proposed zero trust solutions to federal government representatives. This event—plus the SEI's experience in cybersecurity, software engineering, and defense software acquisition—later helped convince Army PEO C3T leaders that the SEI had the right expertise to research and develop their zero trust implementation.

Morrow and his team are partnering with Georgia Tech Research Institute and Johns Hopkins Applied Physics Laboratory to understand the Army's current and envisioned enterprise and tactical cyber infrastructure and to determine what the Army needs to develop or acquire to implement zero trust principles. Along with documenting the business drivers, technical drivers, and quality attributes of a future Army combined network, the SEI team has initially developed several mission threads to capture the mission environments and what information and services soldiers in tactical environments might need.

A second SEI team is creating a schema to score the risk of different identity and access management techniques since deployed soldiers cannot always authenticate on the network.

The SEI's mission engineering approach will help the Army develop the contextual awareness it requires to move from virtual-machine-based applications to cloud services, which will enable an improved understanding of the Army's cybersecurity needs. This shift further complicates zero trust practices in D-DIL environments, especially when cloud services from multiple providers interact, as with combatant commands directing joint and coalition missions. The SEI is setting up a cloud-agnostic cyber testbed to trial different zero trust concepts.

The SEI's engagement with the Army PEO C3T is just in its first year, but already the program is planning to open the work to other projects in its division.

Photos: U.S. Marines

"In a tactical situation, warfighters are used to pushing a button and having things work. Zero trust is about assessing the risk before you take an action."

TIM MORROW, Situational Awareness Technical Manager and Technical Lead, SEI CERT Division

SEI Support for Long Range Standoff Program Spurs New Engagements

In its ongoing engagement with the Long Range Standoff (LRSO) program, the SEI has contributed its technical acumen and in-depth experience in novel software verification analysis techniques to advance the mission of LRSO. The verification techniques the SEI helped establish have already provided rapid feedback to the government for its developers, ensuring swifter iterations and improvements to a highly complex software baseline.

"It's exciting to be the technical lead of a high-performing group whose impact spans multiple programs within the strategic domain," said David Walbeck, the SEI's LRSO technical lead within the Advanced Deterrents Group. "The team's tireless efforts have led to groundbreaking collaborative opportunities with the Air Force's Safety Center, establishing a first of its kind Summit for Nuclear Certification and placing the SEI at the forefront of verification and validation research in highly regulated environments."

The SEI's collaboration with the LRSO program has also spurred engagements with other advanced deterrent programs. These partnerships serve as a testament to the SEI's capabilities and its potential for transformative impact. These opportunities promise a continuum of influence and advancements for years to come.

"The mission is urgent, and we have so much more to do," said Stephen Beck, Advanced Deterrents Group lead for enabling mission capability at scale at the SEI. "I know that 2024 is going to be an incredible year for the program as we advance the verification methodologies for this critical domain and help to ensure the program remains on track for initial operating capability and for the nation."

As an outgrowth of the engagement with LRSO, the SEI created a broad design of experiments grounded in a model problem: operation of vehicles in GPS-denied environments. The experimental design allows the SEI to do research with academia on sensitive national problems. The design has already enabled two research projects, one on large language models and another on formal arguments for large-scale system assurance. Future research based on the design could catalyze advancements in multiple domains of software engineering, such as artificial intelligence and machine learning, navigation and timing, and heterogeneous computing, to name a few.

Photos: U.S. Air Force (top right), U.S. Army (full spread), U.S. Navy (bottom left)

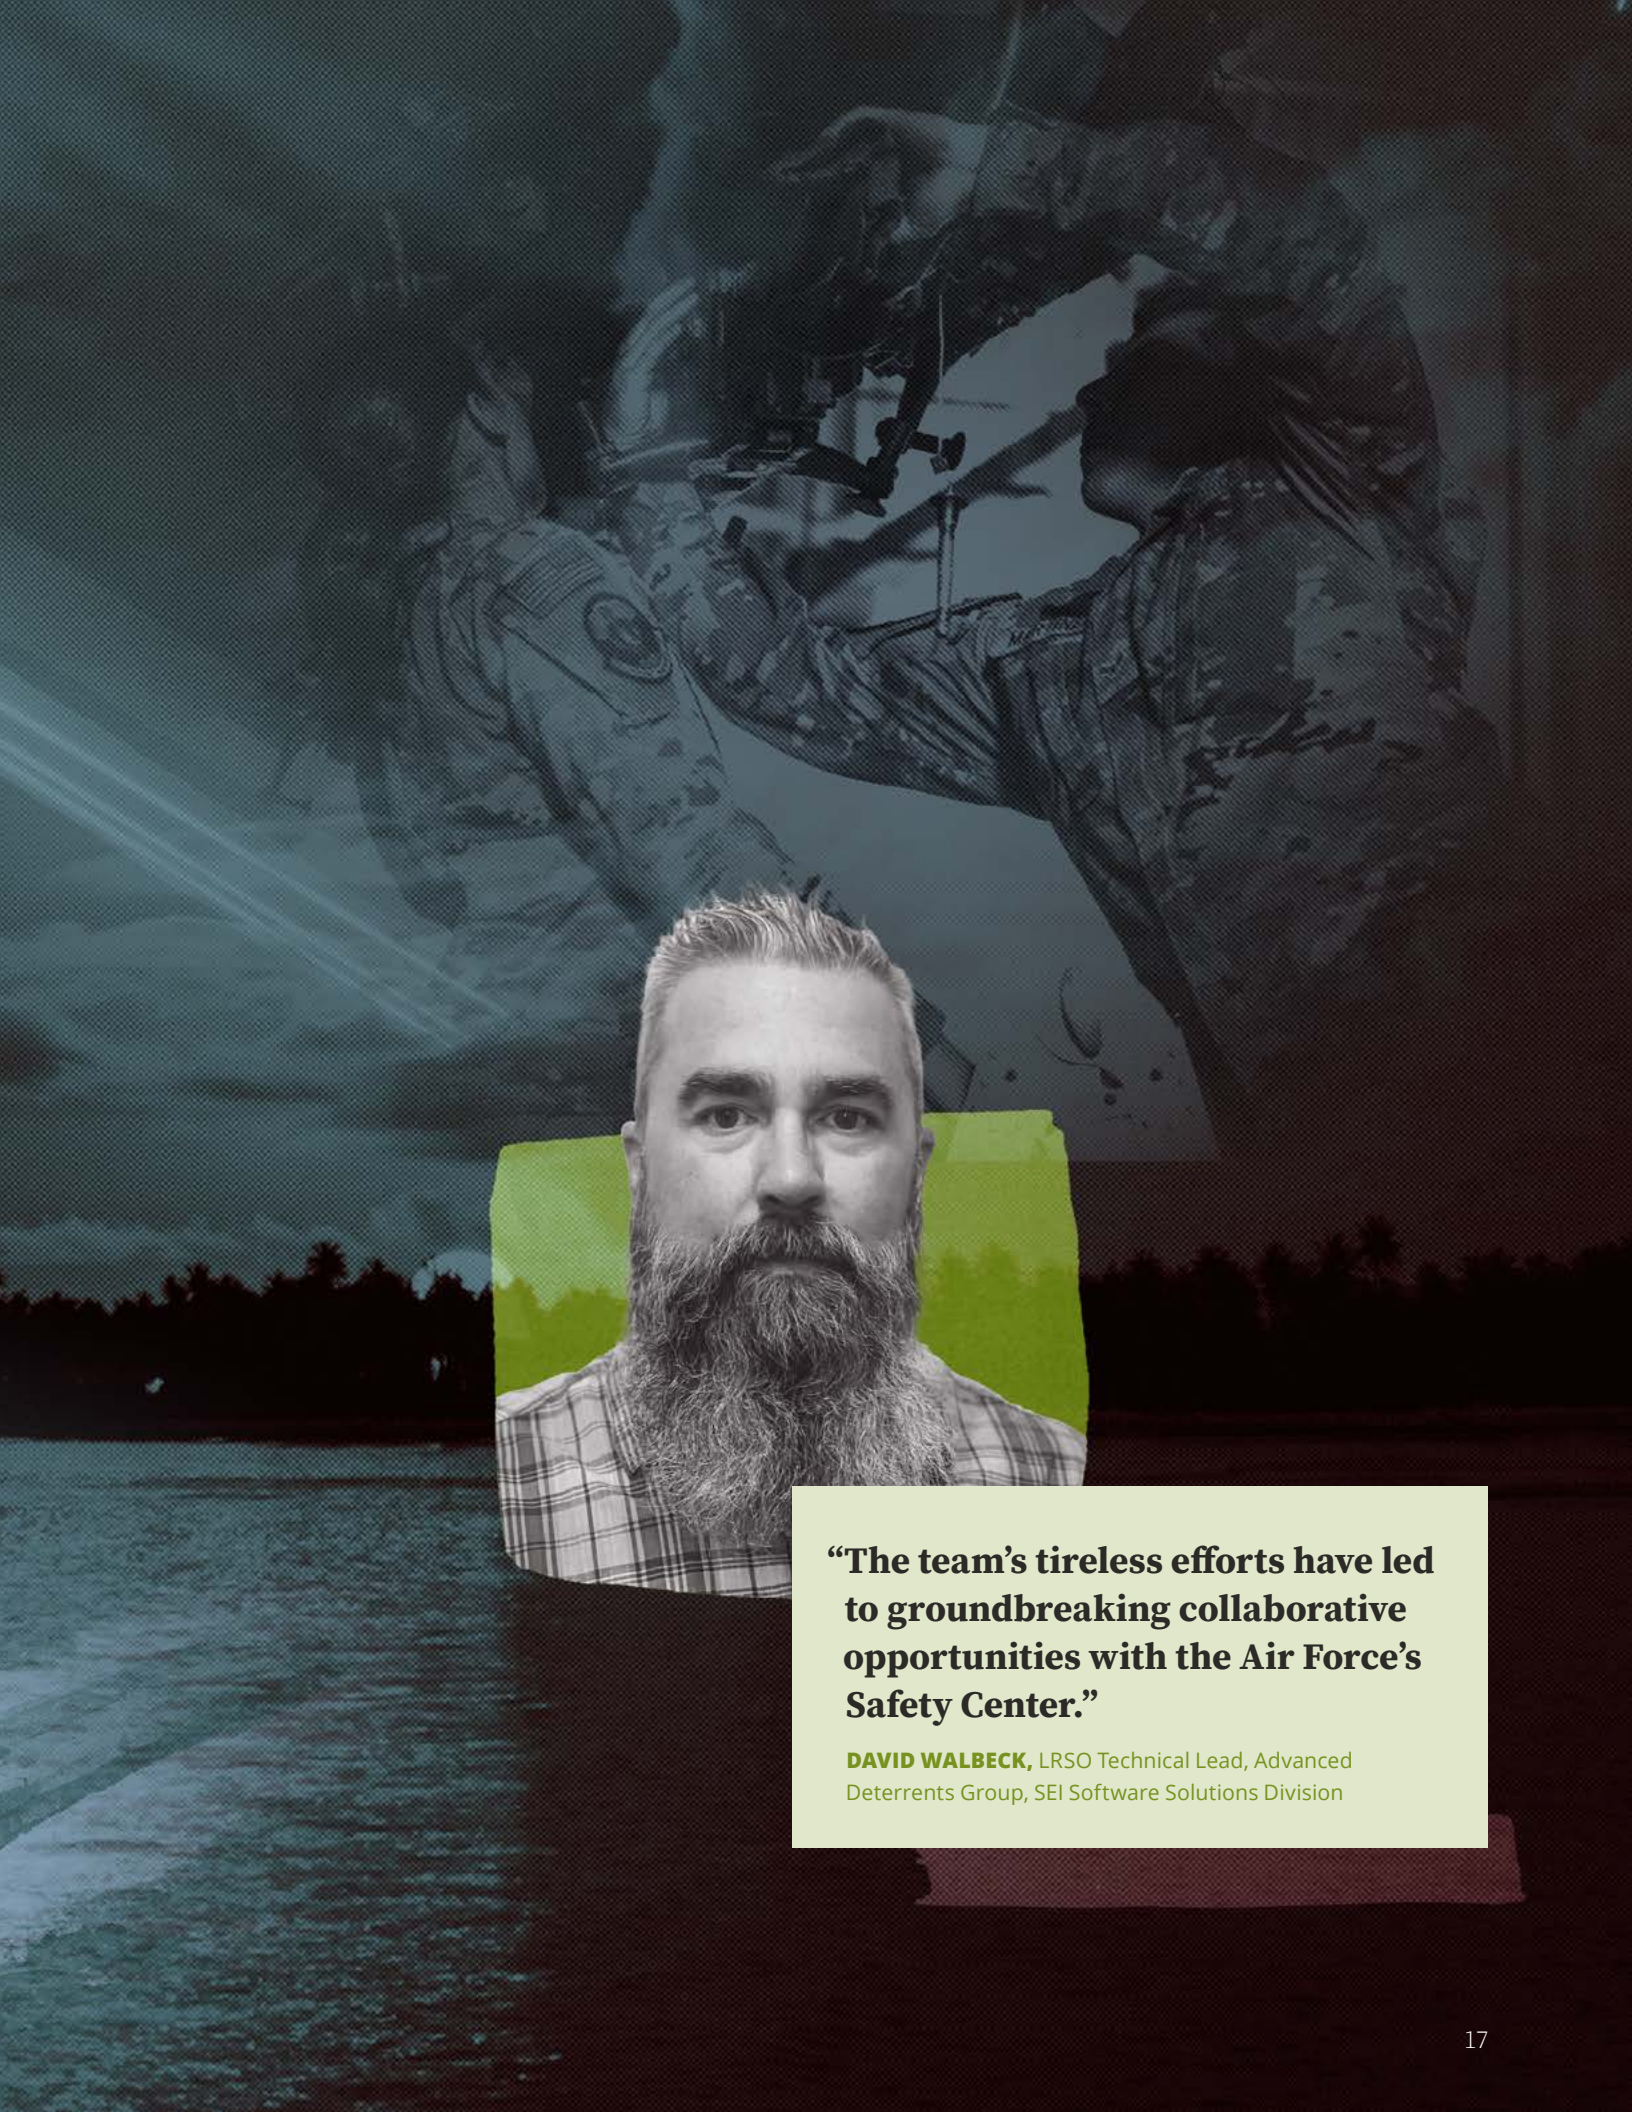

“The team’s tireless efforts have led to groundbreaking collaborative opportunities with the Air Force’s Safety Center.”

DAVID WALBECK, LRSO Technical Lead, Advanced Deterrents Group, SEI Software Solutions Division

"The human has to understand the capabilities and limitations of the AI system to use it responsibly."

KIMBERLY SABLON, Principal Director, Trusted AI and Autonomy, OUSD(R&E), Department of Defense

Assuring Trustworthiness of AI for Warfighters

Advances in artificial intelligence (AI), machine learning, and autonomy have created a proliferation of AI platforms. While these technologies have shown promise for advantages on the battlefield, developers, integrators, and acquisition personnel must overcome engineering challenges to ensure safe and reliable operation. Currently there are no established standards for testing and measuring calibrated trust in AI systems.

In 2023, the Office of the Under Secretary of Defense for Research and Engineering (OUSD(R&E)) and the SEI launched a center aimed at establishing methods for assuring trustworthiness in AI systems with emphasis on interaction between humans and autonomous systems. The Center for Calibrated Trust Measurement and Evaluation (CaTE) aims to help the Department of Defense (DoD) ensure that AI systems are safe, reliable, and trustworthy before being fielded to government users in critical situations.

Since launching, CaTE has embarked on a multi-year project addressing the complexity and engineering challenges associated with AI systems while utilizing software, systems, and AI engineering practices to develop standards, methods, and processes for providing evidence for assurance and developing measures to determine calibrated levels of trust.

"The human has to understand the capabilities and limitations of the AI system to use it responsibly," said Kimberly Sablon, the principal director for trusted AI and autonomy within OUSD(R&E). "CaTE will address the dynamics of how systems interact with each other, and especially the interactions between AI and humans, to establish trusted decisions in the real world. We will identify case studies where AI can be experimented with and iterated in hybrid, live, virtual, and constructive environments with the human in the loop."

CaTE will be a collaborative research and development center and will work with all branches of the military on areas such as human-machine teaming and measurable trust. It is the first such hub led by a non-governmental organization. Carnegie Mellon University (CMU) has been at the epicenter of AI, from the creation of the first AI computer program in 1956 to pioneering work in self-driving cars and natural language processing.

"Developing and implementing AI technologies to keep our armed forces safe is both a tremendous responsibility and a tremendous privilege," said CMU President Farnam Jahanian. "Carnegie Mellon University is grateful to have the opportunity to support the DoD in this work and eager to watch CaTE quickly rise to the forefront of leveraging AI to strengthen our national security and defense."

Together with OUSD(R&E) collaborators and partners in industry and academia, SEI researchers will lead the initiative to standardize AI engineering practices, assuring safe human-machine teaming in the context of DoD mission strategy.

"When military personnel are deployed in harm's way, it's of the utmost importance to give them not only the greatest capability but also the assurance that the AI and autonomous systems they depend on are safe and reliable," said Paul Nielsen, SEI director and chief executive officer. "Because of our work to define the discipline of AI engineering for robust, secure, human-centered, and scalable AI, the SEI is uniquely positioned to support this effort."

For more information about the SEI's AI engineering research, visit sei.cmu.edu/our-work/artificial-intelligence-engineering.

Photo: U.S. Marines (top left), U.S. Department of Defense (portrait)

Pathfinding Project Explores Large Language Models for the Intelligence Mission

Large language model (LLM) applications such as ChatGPT seem a beneficial fit for the data-heavy intelligence community (IC). But IC agencies cannot expose their sensitive information to public models, and LLM output cannot always be trusted. A 2023 SEI study explored how IC organizations might establish their own trustworthy LLM.

In early 2023, the Office of the Director of National Intelligence (ODNI) began investigating use cases for LLMs within the IC. "Technologies like LLMs have the potential to greatly enhance current mission workflows but can also reveal new insights in our existing and future data sets that can't necessarily be derived with legacy approaches," said Bob Lawton, chief of mission capabilities in the ODNI Office of Science and Technology.

ODNI turned to the SEI, which has been researching AI engineering for the agency since 2020. The resulting project asked, for the first time, how the IC might set up a baseline, stand-alone LLM; customize LLMs for intelligence use cases; and evaluate the trustworthiness of LLMs across use cases.

SEI researchers focused on two hallmark LLM use cases: question answering with source attribution and document summarization. "Intelligence analysts frequently need to query data sets, review large corpora of documents, accurately distill the important information, and report it out for different audiences," said Shannon Gallagher, the SEI's AI engineering team lead and the project's principal investigator.

The most cost-effective method of building a domain-specific LLM is to adjust an existing foundational model. One way is to augment the model with external tools at inference time. Another way—more permanent but costly—is fine tuning, which further trains the foundational model on custom data.

The SEI researchers tried both approaches. The solution had to be scalable, one of the SEI's three pillars of AI engineering, so they stood up LLMs of four sizes in both on-premises and cloud environments and fine-tuned them on a custom set of documents. "We benchmarked actual resources that would be needed, like cost, data, compute cycles, and time," said Gallagher.

The results, detailed in a September report, showed that using unclassified infrastructure for the LLM could be affordable if the fine-tuning data set is small and unclassified. The mix of fine tuning and augmentation would vary across intelligence agencies, though the report recommends using the quicker, cheaper augmentation until models can be fairly compared.

Assessing LLM performance is an open area of research. "There's a limited set of metrics for evaluating LLM performance, especially for national security applications," said Gallagher. The SEI is starting to develop quantitative metrics for LLM trustworthiness, security, and reliability. "We need to know these attributes before LLM systems are deployed in any automated function, even with humans in the loop. If intelligence analysts are to use these tools, they have to trust them, or at least know their limitations."

Attributing answers to sources is a human-centered AI principle the SEI followed to help users trust the responses of the project's test LLM. But the system's hallucinations, biased data sources, and high sensitivity to prompt wording led the researchers to conclude that, for high-stakes intelligence tasks, LLM output is not trustworthy without expert review.

ODNI plans to use the project results to inform IC senior leadership about the potential uses, limitations, and implementation considerations of LLMs and in forthcoming AI policies and standards for the IC, including those prescribed in the recent Executive Order on Safe, Secure, and Trustworthy AI.

Download the SEI report *A Retrospective in Engineering Large Language Models for National Security* **at insights.sei.cmu.edu/library/a-retrospective-in-engineering-large-language-models-for-national-security.**

"If intelligence analysts are to use these tools, they have to trust them, or at least know their limitations."

SHANNON GALLAGHER, AI Engineering Team Lead, SEI AI Division

Supporting the Human and Technical Elements of Responsible AI for National Defense

Since the Department of Defense (DoD) adopted ethical artificial intelligence (AI) principles in 2020, the SEI has engaged with multiple defense agencies to support their responsible AI (RAI) implementations.

The Chief Digital and Artificial Intelligence Office (CDAO) leads DoD implementation of RAI policy and guidance and has engaged the SEI to create training for AI workforce development. The SEI has extensive knowledge and experience in the field of RAI and how it extends across different types of complex systems.

The SEI created two RAI curricula: one for future specialists in RAI system development work and another for those just getting familiar with RAI. The curricula helped inform a revision of the needed knowledge, skills, abilities, and tasks for the DoD's data workforce and AI workforce. RAI is critical to many of the roles identified in the DoD Cyber Workforce Framework, and the CDAO is using the curricula to develop a stand-alone course on RAI principles and techniques for selected data and AI roles.

Education in assuring safe, ethical, and responsible AI is challenging. While previous software systems were somewhat static, AI dynamically combines data sets and connects systems, bringing new potential risks. The course provides relevant training and materials about these risks and many other RAI challenges.

The Defense Innovation Unit (DIU) is another agency concerned with RAI. It is the only DoD organization focused exclusively on fielding and scaling commercial technology across the U.S. military at commercial speed and scale. The SEI has been providing technical advising and support at all phases of the DIU pipeline as the organization reviews and evaluates potential vendors to address DoD mission needs, including the need for RAI solutions as stated in the DoD's *Ethical Principles for Artificial Intelligence*.

Part of this work supports DIU's operationalization of *Responsible Artificial Intelligence Guidelines in Practice*, co-authored by SEI researchers. These guidelines include worksheets to help vendors better plan, develop, and deploy AI tools. Completing the worksheets enables vendors to develop their own test metrics and facilitates the SEI's independent testing and evaluation of developed tools.

"Companies developing solutions often have not thought about these very ethically driven questions, such as harms modeling," noted Sumanyu Gupta, machine learning engineer and team lead in the SEI's AI Division. The SEI has been directly engaging AI solution vendors to consider ethical insights in everything from tool

> **"Companies developing solutions often have not thought about these very ethically driven questions, such as harms modeling."**
>
> **SUMANYU GUPTA,** Machine Learning Engineer and Team Lead, SEI AI Division

requirements to roadblocks. Using the worksheets can even suggest features that the vendors had not previously considered.

The SEI integrates RAI principles, AI fundamentals, software engineering and acquisition practices, and workforce development expertise to address the technical and human obstacles faced when planning, developing, and deploying AI systems. The work is also informed by the institute's relationship with Carnegie Mellon University (CMU), such as collaborative research on explainable AI and fairness of AI systems. SEI AI researchers Carol Smith and Matt Hale participated in the 2023 CMU-organized workshops on the National Institute of Standards and Technology (NIST) *Artificial Intelligence Risk Management Framework*. Smith is also on the Advisory Council and Interim Leadership Team of the Block Center CMU Responsible AI Initiative. These ongoing connections enrich the SEI's engagements with the CDAO and DIU, which support the U.S. military's legal, ethical, and policy commitments to be responsible, equitable, traceable, reliable, and governable in its adoption of AI.

Extending SysML V2 with AADL Concepts to Support Engineering and Certification of Safety-Critical Systems

Department of Defense (DoD) program offices and defense industrial base vendors employ model-based systems engineering (MBSE) practices to engineer complex embedded systems architecture and reduce safety and security risks through early analysis. Both the Architecture Analysis and Design Language (AADL) and the Object Management Group's (OMG's) Systems Modeling Language (SysML) support this effort by providing engineers with the ability to develop architectures, conduct reviews, and perform analysis. SysML Version 1 focuses on systems engineering, traceability, and system decomposition and refinement. AADL focuses on the precise evaluation of performance and safety metrics.

Defense developers need to use both SysML and AADL in their MBSE efforts, but using two languages is challenging and entails significant, unsustainable costs in both training and maintenance. With the upcoming release of SysML Version 2, the SEI is spearheading the effort to combine SysML's capabilities to describe complex systems and AADL's analysis capabilities. The work will bridge the gap between the two languages through a refinement of SysML V2 concepts that align with AADL V2, delivered as a SysML V2 library.

Jérôme Hugues, principal architecture researcher at the SEI, noted, "The expected impact is in having a single way to model systems by reducing the number of languages required, while preserving the semantics and analysis capabilities of AADL."

This work will also reduce the effort required for the DoD to design safety-critical systems using SysML V1. Limited by SysML V1's semantics, the DoD currently uses SysML profiles and translation tools to refine SysML models into AADL ones and continue design activities.

Photo: Alice de Casanove

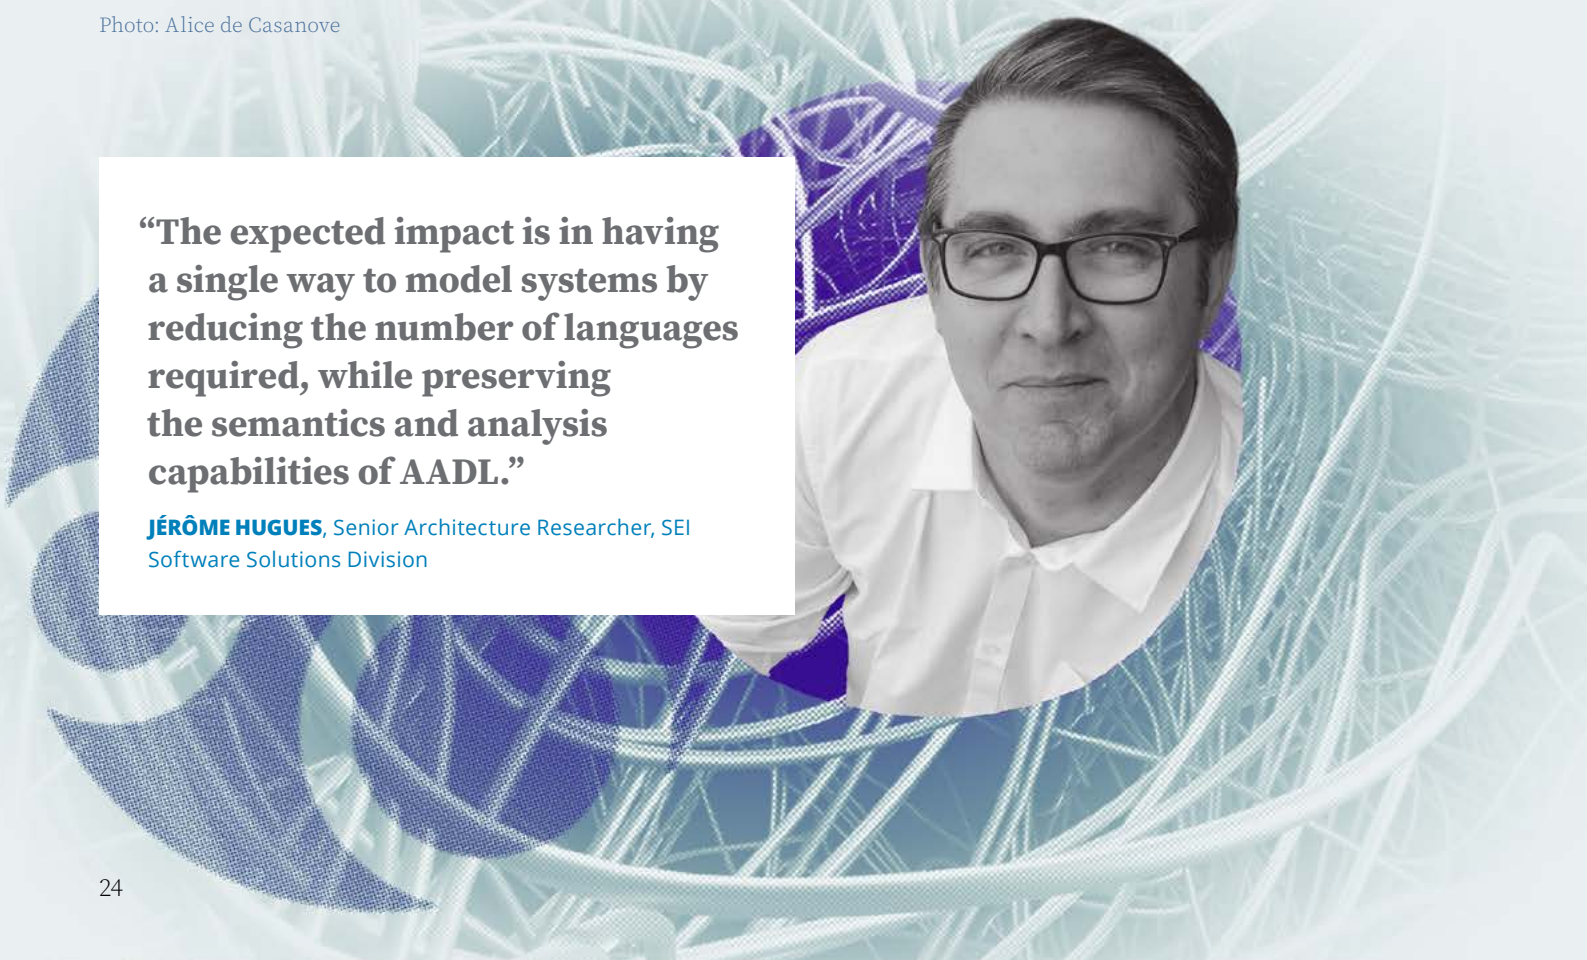

"The expected impact is in having a single way to model systems by reducing the number of languages required, while preserving the semantics and analysis capabilities of AADL."

JÉRÔME HUGUES, Senior Architecture Researcher, SEI Software Solutions Division

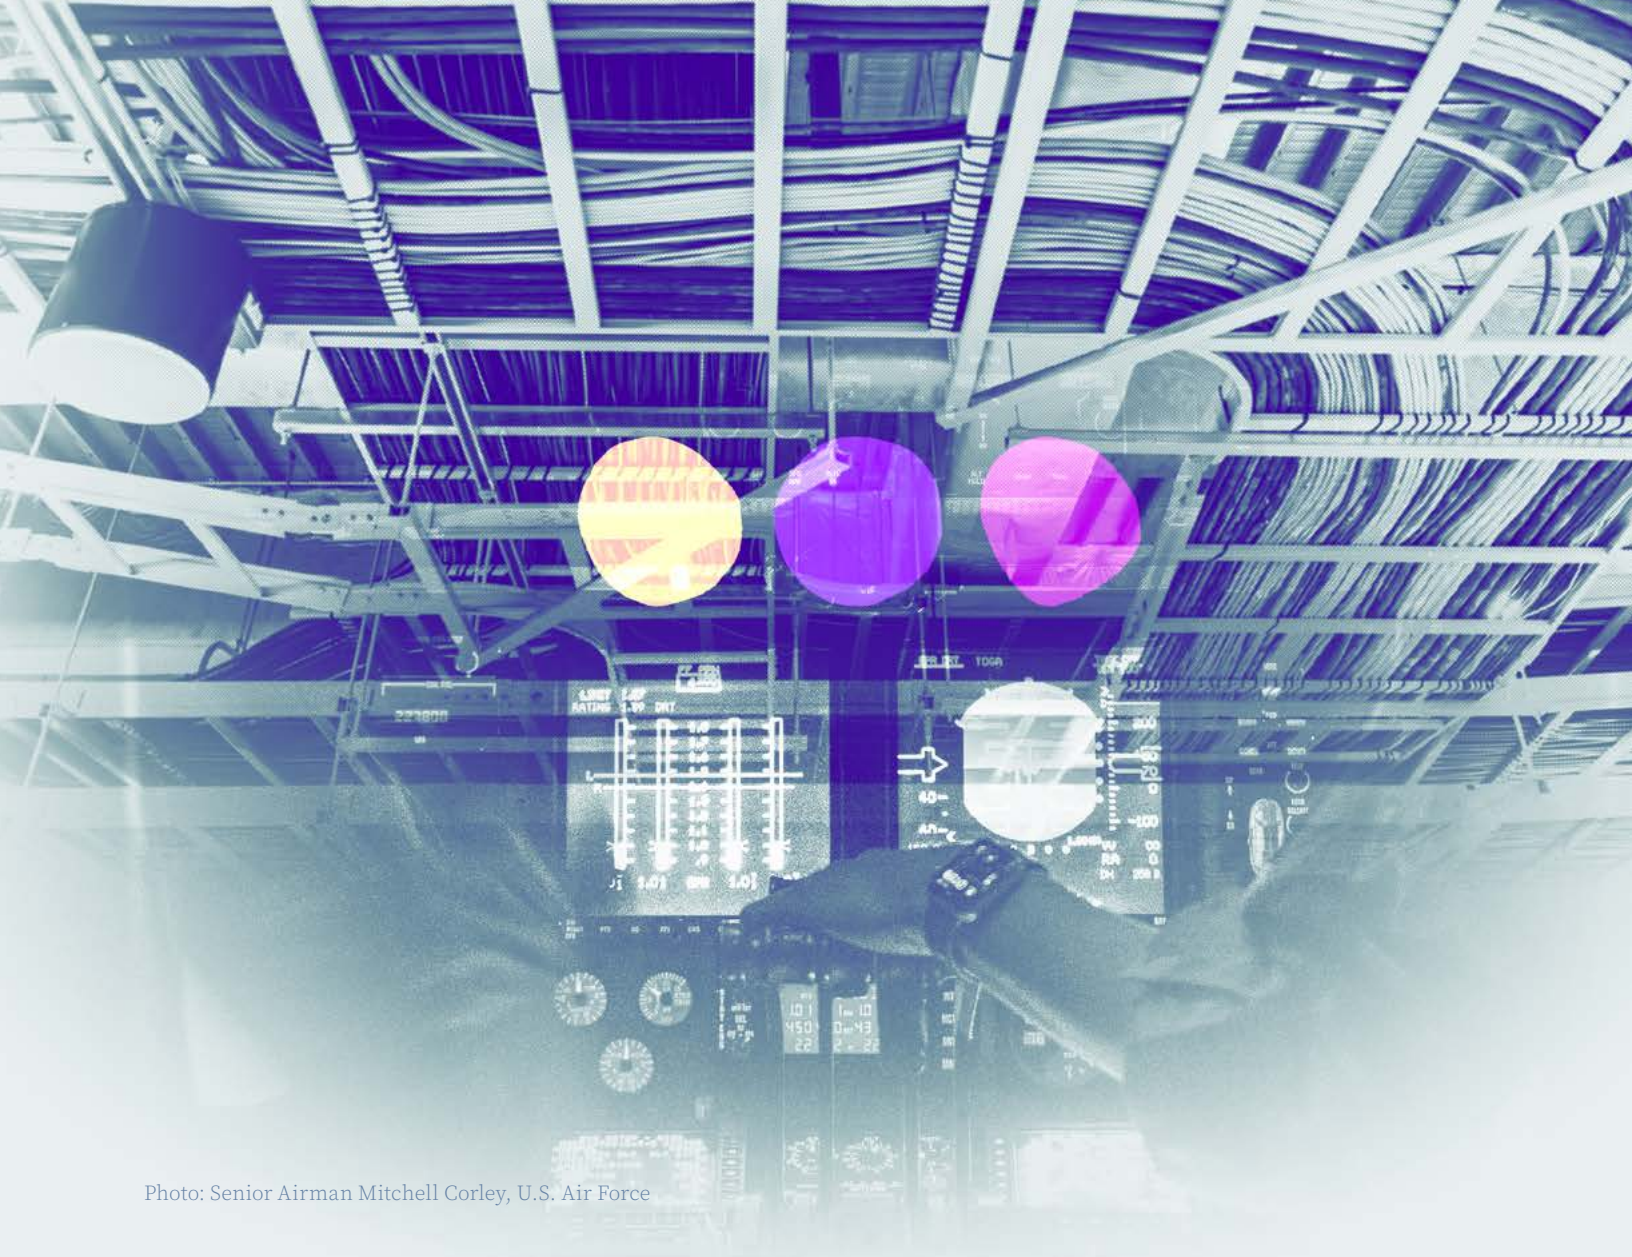

Photo: Senior Airman Mitchell Corley, U.S. Air Force

But the approach of using multiple tools and iterations is slow and error prone and requires engineers to have additional expertise. With leaner syntax and semantics, as well as extensibility mechanisms, SysML V2 provides an opportunity to better integrate AADL concepts directly into a single language, which will ease the engineer's learning curve.

The SEI has developed a proof-of-concept extension of SysML V2 with AADL V2 capabilities, leveraging both SysML V2 and KerML libraries. Hugues and Gene Shreve, a systems engineer with Integration Innovation, Inc., have also recently established a working group as part of the newly formed OMG Systems Modeling Community (SMC) to further this effort. The working group will compare SysML V2 and AADL semantics and modeling styles, align semantics and define mapping rules between both languages, develop specific KerML/SysML library elements to support real-time-embedded and safety-critical system design and development, and define use and test cases to validate the library.

The working group's contributions will support the precise engineering of safety-critical, real-time, embedded systems and allow for code generation as well as verification and validation.

The DoD will benefit from having a single-language approach covering both high-level MBSE and low-level safety-critical embedded systems and precise semantics leveraging SysML V2 extensible semantics capabilities. Developing a common tool set built across the SysML V2 open application programming interface, the SEI-led working group will support the DoD's *Digital Engineering Strategy* to promote the use of digital representations of systems and components and the use of digital artifacts to design and sustain national defense systems.

Professional Organization Participation Boosts SEI Staff and Mission

The active participation of SEI staff members in professional societies and organizations helps the SEI serve its Department of Defense (DoD) sponsor, external customers, and the nation. In 2023, many SEI researchers held leadership positions or attained special honors in professional organizations.

Conference Committees

- **Bjorn Andersson** was a program chair for the IEEE Explainability of Real-Time Systems and their Analysis (ERSA) 2023 workshop. **Dionisio de Niz**, **Mark Klein**, **Carol Smith**, and **Violet Turri** were on the conference program committee.
- **Linda Parker Gates** was the co-chair of the International Association for Strategic Planning (IASP) Global Conference 2023.

Society Leadership and Memberships

- **Grace Lewis** was elected first vice president of the IEEE Computer Society for 2024. She was also second vice president and vice president for technical and conference activities in 2023.
- **Rick Kazman** was elected to the IEEE Computer Society board of governors for 2023–2025.
- **Andrew Kotov**, **Brigid O'Hearn**, and **Scott Sinclair** were promoted to senior members of IEEE.
- **Eric Ferguson** and **Sam Procter** were promoted to senior members of the Association for Computing Machinery (ACM).
- **Jonathan Frederick** was the vice president of the (ISC)² Pittsburgh chapter.

Editorships

- **Leigh Metcalf** was co-editor-in-chief of the ACM journal *Digital Threats: Research and Practice*.
- **Ipek Ozkaya** completed five years as *IEEE Software* magazine editor-in-chief.
- **Anita Carleton** became the chair of the *IEEE Software* magazine advisory board.
- **Bill Claycomb** was an associate editor for the journal *Counter-Insider Threat Research and Practice*.

Standards Development

- **Jérôme Hugues** was the technical lead for the SAE AS-2C Architecture Analysis and Design Language (AADL) committee working on the definition of the language and co-chair of the Systems Modeling Community on Real-Time Safety-Critical Systems (AADL V2 for SysML V2).
- **Carol Smith** was a member of the IEEE P7008 Standard for Ethically Driven Nudging for Robotic, Intelligent, and Autonomous Systems working group. Smith also joined the National Institute of Standards and Technology Generative Artificial Intelligence Public working group.
- **Hasan Yasar** chaired or co-chaired the committees for the IEEE DevOps, IEEE Software Configuration Management, and Open Group Zero Trust Implementation standards development.
- **Robin Yeman** is the vice chair of the National Defense Industrial Association (NDIA) Agile Delivery for Agencies, Programs and Team (ADAPT) and co-chair of the Agile and systems engineering working group of the International Council on Systems Engineering (INCOSE).
- **Roman Danyliw** was a security area director in the Internet Engineering Task Force (IETF).
- **Chris Inacio** co-chaired the IETF network file systems (NFS) working group.

Other Honors

- **Carol Smith** was a member of the ACM Distinguished Speaker Program.
- **Alexis Presti-Simpson** was appointed a University Fellow on the Norwich University Board of Fellows on the College of Sciences and Mathematics committee to advise on academic programs in coordination with the Dean of the Faculty of the Office of the President.


Bjorn Andersson


Dionisio de Niz

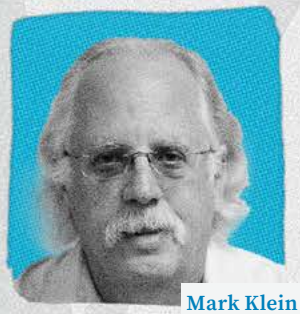
Mark Klein


Carol Smith


Violet Turri


Linda Parker Gates


Grace Lewis


Rick Kazman


Andrew Kotov


Brigid O'Hearn


Scott Sinclair

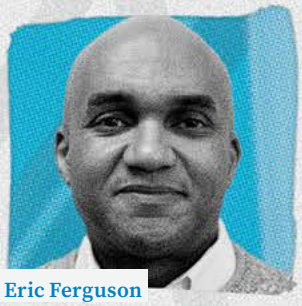
Eric Ferguson


Sam Procter


Jonathan Frederick


Leigh Metcalf


Ipek Ozkaya


Anita Carleton


Bill Claycomb


Jérôme Hugues


Hasan Yasar


Robin Yeman


Roman Danyliw


Chris Inacio


Alexis Presti-Simpson

Leadership

CMU Leadership

Farnam Jahanian
President

James H. Garrett, Jr.
Provost and Chief Academic Officer

Theresa Mayer
Vice President for Research

SEI Executive Leadership

Paul Nielsen
Director and Chief Executive Officer

David Thompson
Deputy Director and Chief Operating Officer

Tom Longstaff
Chief Technology Officer

Anita Carleton
Director, Software Solutions Division

Gregory J. Touhill
Director, CERT Division

Matt Gaston
Director, Artificial Intelligence Division

Heidi Magnelia
Chief Financial Officer

Mary Catherine Ward
Chief Strategy Officer

Sandra Noonan
General Counsel

Board of Visitors

The SEI Board of Visitors advises the Carnegie Mellon University president, university provost, and SEI director on SEI plans and operations. The board monitors SEI activities, provides reports to the president and provost, and makes recommendations for improvement.

Russell Crockett

Cofounder of Aeon Blue Technologies; Principal and Owner of RTC Energy LLC; Trustee, Carnegie Mellon University

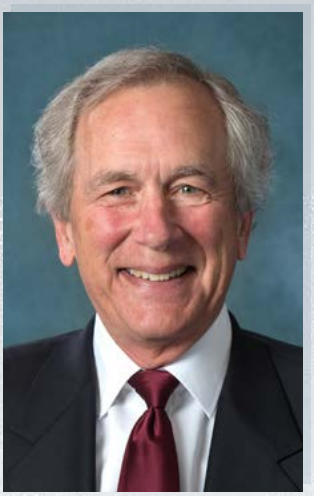

Philip Dowd

Private investor; former Senior Vice President, SunGard Data Systems; Emeritus Trustee, Carnegie Mellon University

Tom Love

Chief Executive Officer, ShouldersCorp; Founder of Object Technology Group within IBM Consulting

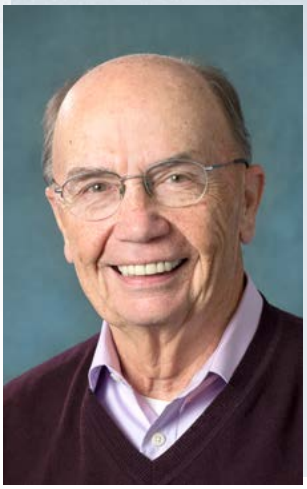

Donald Stitzenberg

Consultant and founder, CBA Associates; Emeritus Trustee, Carnegie Mellon University; former Executive Director of Clinical Biostatistics at Merck; retired member, New Jersey State Bar Association

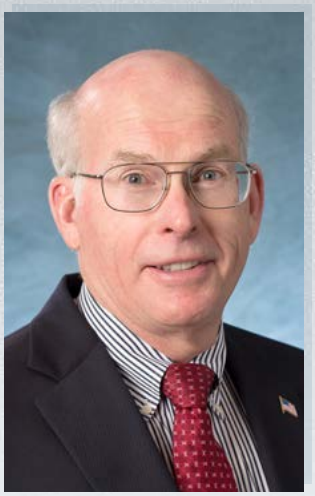

John M. Gilligan

Chair, SEI Board of Visitors; President and CEO, Center for Internet Security (CIS); former President and COO, Schafer Corporation; former President, Gilligan Group; former Senior Vice President and Director, Defense Sector of SRA International; former CIO for the Department of Energy and the U.S. Air Force

Elizabeth A. Hight

Former Vice President of the Cybersecurity Solutions Group, U.S. Public Sector, Hewlett Packard Enterprise Services; Rear Admiral, retired, U.S. Navy; former Vice Director of the Defense Information Systems Agency

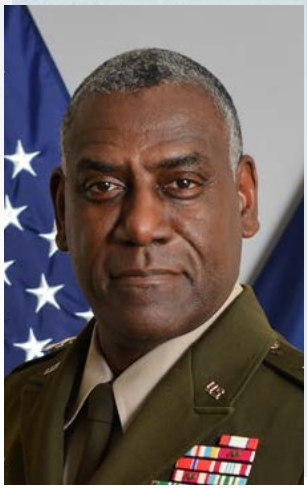

Cedric T. Wins

Superintendent, Virginia Military Institute; Major General, retired, U.S. Army; former Commanding General of the Army Combat Capabilities Development Command (CCDC); former Commander, Army Research, Development, and Engineering Command (RDECOM); former Director of Force Development in the Army Office of the Deputy Chief of Staff

Key Publications

Articles

Casey, W. A. & Cai, Y. Editorial: Renaissance of Biomimicry Computing. *Mobile Networks and Applications*. Volume 28. December 14, 2022. Pages 486–489. **https://doi.org/10.1007/s11036-022-02066-7**

Dramko, Luke; Lacomis, Jeremy; Yin, Pengcheng; Schwartz, Ed; Allamanis, Miltiadis; Neubig, Graham; Vasilescu, Bogdan; & Le Goues, Claire. DIRE and Its Data: Neural Decompiled Variable Renamings with Respect to Software Class. *ACM Transactions on Software Engineering and Methodology*. Volume 32. Issue 2. March 29, 2023. **https://doi.org/10.1145/3546946**

Ernst, Neil A.; Klein, John; Bartolini, Marco; Coles, Jeremy; & Rees, Nick. Architecting Complex, Long-Lived Scientific Software. *Journal of Systems and Software*. Volume 204. October 2023. **https://doi.org/10.1016/j.jss.2023.111732**

Ferreira, T.; Ivers, J.; Yackley, J. J.; Kessentini, M.; Ozkaya, I.; & Gaaloul, K. Dependent or Not: Detecting and Understanding Collections of Refactorings. *IEEE Transactions on Software Engineering*. Volume 49. Issue 6. June 1, 2023. Pages 3344–3358. **https://doi.org/10.1109/TSE.2023.3244123**

Hugues, Jérôme. Special Issue on Reliable Software Technologies (AEiC2022). *Journal of Systems Architecture*. Volume 135. February 2023. **https://doi.org/10.1016/j.sysarc.2022.102809**

Jin, W.; Zhong, D.; Cai, Y.; Kazman, R.; & Liu, T. Evaluating the Impact of Possible Dependencies on Architecture-Level Maintainability. *IEEE Transactions on Software Engineering*. Volume 49. Number 3. March 1, 2023. Pages 1064–1085. **https://doi.org/10.1109/TSE.2022.3171288**

Levine, Alan & Tucker, Brett Alan. Zero Trust Architecture: Risk Discussion. *Digital Threats: Research and Practice*. Volume 4. Issue 1. March 31, 2023. **https://doi.org/10.1145/3573892**

Márquez, Gastón; Astudillo, Hernán; & Kazman, Rick. Architectural Tactics in Software Architecture: A Systematic Mapping Study. *Journal of Systems and Software*. Volume 197. March 2023. **https://doi.org/10.1016/j.jss.2022.111558**

Procter, Sam. The OSATE Slicer: Graph-Based Reachability for Architectural Models. *Journal of Object Technology*. Volume 22. Number 2. July 2023. Pages 2:1–14. **https://samprocter.com/wp-content/uploads/2023/06/ecmfa23-slicer.pdf**

Romagnoli, Raffaele; Krogh, Bruce H.; de Niz, Dionisio; Hristozov, Anton D.; & Sinopoli, Bruno. Runtime System Support for CPS Software Rejuvenation. *IEEE Transactions on Emerging Topics in Computing*. Volume 11. Number 3. July–September, 2023. Pages 594–604. **https://doi.org/10.1109/TETC.2023.3267899**

Romagnoli, Raffaele; Krogh, Bruce H.; de Niz, Dionisio; Hristozov, Anton D.; & Sinopoli, Bruno. Software Rejuvenation for Safe Operation of Cyber-Physical Systems in the Presence of Run-Time Cyberattacks. *IEEE Transactions on Control Systems Technology*. Volume 31. Issue 4. July, 2023. Pages 1565–1580. **https://doi.org/10.1109/TCST.2023.3236470**

Shomo, Paul; Echeverría, Sebastián; & Sowell, Jesse. Introduction to the Special Issue on the Lifecycle of IoT (In)security. *Digital Threats: Research and Practice*. Volume 3. Issue 4. February 16, 2023. Pages 1–2. **https://doi.org/10.1145/3569901**

Spring, Jonathan M. Analysis of How Many Undiscovered Vulnerabilities Remain in Information Systems. *Computers & Security*. Volume 131. Article number 103191. August 2023. **https://doi.org/10.1016/j.cose.2023.103191**

Sridhar, Upasana; Tukanov, Nicholai; Binder, Elliott; Low, Tze Meng; McMillan, Scott; & Schatz, Martin D. SMaLL: Software for Rapidly Instantiating Machine Learning Libraries. *ACM Transactions on Embedded Computing Systems*. July 12, 2023. **https://doi.org/10.1145/3607870**

Takeshita, Jonathan; Karl, Ryan; Gong, Ting; & Jung, Taeho. SLAP: Simpler, Improved Private Stream Aggregation from Ring Learning with Errors. *Journal of Cryptology*. Volume 36. Article 8. March 3, 2023. **https://doi.org/10.1007/s00145-023-09450-w**

Tamburri, Damian A.; Kazman, Rick; & Fahimi, Hamed. On the Relationship Between Organizational Structure Patterns and Architecture in Agile Teams. *IEEE Transactions on Software Engineering*. Volume 49. Issue 1. January 1, 2023. Pages 325–347. **https://doi.org/10.1109/TSE.2022.3150415**

Wetherell, Mark A.; Lau, Shing-Hon; & Maxion, Roy A. The Effect of Socially Evaluated Multitasking Stress on Typing Rhythms. *Psychophysiology.* Volume 60. Issue 8. August 2023. **https://doi.org/10.1111/psyp.14293**

Books

Johnson, Suzette & Yeman, Robin. *Industrial DevOps: Build Better Systems Faster.* IT Revolution. 2023.

Spafford, Eugene H.; Metcalf, Leigh; & Dykstra, Josiah. *Cybersecurity Myths and Misconceptions: Avoiding the Hazards and Pitfalls that Derail Us.* Addison-Wesley Professional. 2023.

Conference Presentations

Booz, Jarrett. *Developing Hands-On Cybersecurity "Challenges" to Improve Training & Assessment.* TechNet Indo-Pacific 2023. November 2023.

Chick, Timothy A. *Does Your DevSecOps Pipeline Only Function as Intended?* InfoSec World 2023. September 2023.

Costa, Dan. *Insider Threats in the Software Development Lifecycle.* 2022 DoD Weapon Systems Software Summit. December 2022. **https://www.sae.org/binaries/content/assets/cm/content/attend/2022/dod/presentations/17-costa-insider-threats-22-09-30-2.pdf**

Metcalf, Leigh. *Mitigating Common Cognitive Biases in Cybersecurity Practice.* Women in Cybersecurity (WiCyS) 2023. March 2023.

Nichols, William Richard. *Automated Continuous Program Estimation for Pipelines and Factories.* First International Boehm Forum on COCOMO and Systems and Software Cost Modeling. November 2022. **https://drive.google.com/file/d/1jx5woFZ5gaXNjf3pa26GznwduXOfEHuA/view?pli=1**

Prevost, Katherine & Shimeall, Timothy J. *IPFIX and DPI Information in a Big Data Environment.* FloCon 2023. January 2023. **https://insights.sei.cmu.edu/library/ipfix-and-dpi-information-in-a-big-data-environment/**

Scanlon, Thomas. *Are Deepfakes Really a Security Threat?* ISC2 Security Congress 2022. October 2022.

Sherman, Mark. *Should I Trust ChatGPT to Review My Program?* InfoSec World 2023. September 2023.

Shevchenko, Natasha. *Modeling Cyber Threats with MBSE.* Department of the Air Force Modeling and Simulation Summit 2023. May 2023.

Smith, Carol. *Implementing Responsible, Human Centered AI.* 4th Annual AI World Government Conference. October 2022.

Somlo, Gabriel L. *Self-Hosting (Almost) All the Way Down.* Free and Open Source Software Developers' European Meeting (FOSDEM) 2023. February 2023. **https://archive.fosdem.org/2023/schedule/event/rv_selfhosting_all_the_way_down/**

Trzeciak, Randall F. *Measuring Insider Risk Program Effectiveness.* InfoSec World 2023. September 2023.

Tucker, Brett. *Ransomware, Defense & Resilience Strategies.* 4th Annual AI World Government Conference. October 2022.

Tucker, Brett. *Striking the Balance: Measuring and Managing the Complexity of Cyber Environments.* FloCon 2023. January 2023. **https://insights.sei.cmu.edu/library/striking-the-balance-measuring-and-managing-the-complexity-of-cyber-environments/**

Tyzenhaus, Laurie. *The Four Pillars of Cybersecurity.* FIRST 2023 Conference. June 2023. **https://www.first.org/resources/papers/conf2023/FIRSTCON23-TLPCLEAR-Tyzenhaus-The-4-Pillars-of-Cyber-Security.pdf**

Yasar, Hasan. *Realities of SBOM: What Is Under the Hood of SBOM?* InfoSec World 2023. September 2023.

Conference Papers

Chu, Simon; Shedden, Emma; Zhang, Changjian; Meira-Góes, Rômulo; Moreno, Gabriel A.; Garlan, David; & Kang, Eunsuk. Runtime Resolution of Feature Interactions through Adaptive Requirement Weakening. Pages 115–125. In *2023 IEEE/ACM 18th Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*. May 2023. **https://doi.org/10.1109/SEAMS59076.2023.00025**

Fang, Hongzhou; Cai, Yuanfang; Kazman, Rick; & Lefever, Jason. Identifying Anti-Patterns in Distributed Systems with Heterogeneous Dependencies. Pages 116–120. In *Proceedings: IEEE 20th International Conference on Software Architecture Companion (ICSA-C)*. March 2023. **https://doi.org/10.1109/ICSA-C57050.2023.00035**

Gallagher, Shannon K.; Whisnant, Austin; Hristozov, Anton D.; & Vasudevan, Amit. Reviewing the Role of Machine Learning and Artificial Intelligence for Remote Attestation in 5G+ Networks. Pages 602–607. In *Proceedings: 2022 IEEE Future Networks World Forum (FNWF)*. October 2022. **https://doi.org/10.1109/FNWF55208.2022.00111**

Hatcliff, J.; Hugues, J.; Stewart, D.; & Wrage, L. Formalization of the AADL Run-Time Services. Pages 105–134. In *Leveraging Applications of Formal Methods, Verification and Validation: Software Engineering*. October 2022. **https://doi.org/10.1007/978-3-031-19756-7_7**

Ivers, James; Nord, Robert L.; Ozkaya, Ipek; Seifried, Chris; Timperley, Christopher S.; & Kessentini, Marouane. Industry Experiences with Large-Scale Refactoring. Pages 1544–1554. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE 2022)*. November 2022. **https://doi.org/10.1145/3540250.3558954**

Kazman, Rick & Chen, Hong-Mei. Architecture of Complexity Revisited: Design Primitives for Ultra-Large-Scale Systems. Pages 6956–6965. In *Proceedings of the 56th Annual Hawaii International Conference on System Sciences*. Volume 2023. January 2023. **https://scholarspace.manoa.hawaii.edu/items/23698e0d-ac1c-4562-b794-d77388bd5c55**

Kumara, Indika; Pecorelli, Fabiano; Catolino, Gemma; Kazman, Rick; Tamburri, Damian Andrew; & Van Den Heuvel, Willem-Jan. Architecting MLOps in the Cloud: From Theory to Practice. Pages 333–335. In *Proceedings: IEEE 20th International Conference on Software Architecture Companion (ICSA-C)*. March 2023. **http://dx.doi.org/10.1109/ICSA-C57050.2023.00076**

Lefever, Jason; Cai, Yuanfang; Kazman, Rick; & Fang, Hongzhou. Towards the Assisted Decomposition of Large-Active Files. Pages 126–130. In *Proceedings: IEEE 20th International Conference on Software Architecture Companion (ICSA-C)*. March 2023. **http://dx.doi.org/10.1109/ICSA-C57050.2023.00037**

Maffey, Katherine R.; Dotterrer, Kyle; Niemann, Jennifer; Cruickshank, Iain; Lewis, Grace A.; & Kästner, Christian. MLTEing Models: Negotiating, Evaluating, and Documenting Model and System Qualities. Pages 31–36. In *2023 IEEE/ACM 45th International Conference on Software Engineering: New Ideas and Emerging Results (ICSE-NIER)*. May 2023. **https://doi.org/10.1109/ICSE-NIER58687.2023.00012**

Morales, Jose Andre; Hamed, Jeffrey; Reynolds, Douglas; Shepard, David; Antunes, Luiz; Yankel, Joseph; & Yasar, Hasan. Experiences with Secure Pipelines in Highly Regulated Environments. In *Proceedings of the 18th International Conference on Availability, Reliability and Security*. August–September 2023. **https://doi.org/10.1145/3600160.3605466**

Nahar, N.; Zhang, H.; Lewis, G.; Zhou, S.; & Kästner, C. A Meta-Summary of Challenges in Building Products with ML Components – Collecting Experiences from 4758+ Practitioners. Pages 171–183. In *2023 IEEE/ACM 2nd International Conference on AI Engineering – Software Engineering for AI (CAIN)*. May 2023. **https://doi.org/10.1109/CAIN58948.2023.00034**

Savell, Thomas C.; Kam, Ambrose; Shevchenko, Nataliya; & Tucker, Brett. A Cyber Attack Forecasting System. In *Proceedings of the Interservice Industry, Training, Simulation, and Education Conference (I/ITSEC)*. November–December 2022. **https://www.xcdsystem.com/iitsec/proceedings/index.cfm?Year=2022&AbID=111901&CID=944#View**

Schenker, Fred & Hugues, Jérôme. You Can't Wait for ROI to Justify Model-Based Design and Analysis for Cyber Physical Systems' Embedded Computing Resources. In *Annual Acquisition Research Symposium Proceedings & Presentations*. June 1, 2023. **https://insights.sei.cmu.edu/library/you-cant-wait-for-roi-to-justify-model-based-design-and-analysis-for-cyber-physical-systems-embedded-computing-resources/**

Turri, Violet & Dzombak, Rachel. Why We Need to Know More: Exploring the State of AI Incident Documentation Practices. Pages 576–583. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. August 2023. **https://doi.org/10.1145/3600211.3604700**

Wilder, Gregory; Miyamoto, Riley; Watson, Samuel; Kazman, Rick; & Peruma, Anthony. An Exploratory Study on the Occurrence of Self-Admitted Technical Debt in Android Apps. Pages 1–10. In *Proceedings: 2023 ACM/IEEE International Conference on Technical Debt (TechDebt)*. May 2023. **https://doi.org/10.1109/TechDebt59074.2023.00007**

Yang, Chenyang; Brower-Sinning, Rachel A.; Lewis, Grace; & Kästner, Christian. Data Leakage in Notebooks: Static Detection and Better Processes. Pages 1–12. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering.* October 2022. **https://doi.org/10.1145/3551349.3556918**

Yang, Chenyang; Brower-Sinning, Rachel A.; Lewis, Grace; Kästner, Christian; & Wu, Tongshuang. Capabilities for Better ML Engineering. In *CEUR Workshop Proceedings, 2023 Workshop on Artificial Intelligence Safety*. February 2023. **https://ceur-ws.org/Vol-3381/41.pdf**

Zhai, Lijing; Kanellopoulos, Aris; Fotiadis, Filippos; Vamvoudakis, Kyriakos G.; & Hugues, Jérôme. A Modular Approach to Verification of Learning Components in Cyber-Physical Systems. In *AIAA SCITECH 2023 Forum*. January 2023. **https://doi.org/10.2514/6.2023-0131**

Keynotes

Longstaff, Tom. *Generative AI in Cyber: Change Is Coming in a Big Way*. Generative Artificial Intelligence in Cyber, International Information Integrity Institute (i-4) Global Forum. June 28, 2023.

Longstaff, Tom. *SEI Thoughts on AI T&E and Related Topics.* AI T&E Defense Science Board (DSB). April 13, 2023.

Nielsen, Paul. *Cybersecurity and Artificial Intelligence: How They Help and Challenge Each Other.* 2023 IEEE International Conference on Cyber Security and Resilience. August 1, 2023.

Nielsen, Paul. *Hope and Hype: AI and Autonomy.* Process Insights USA. September 25, 2023.

Nielsen, Paul. *National Security Readiness: Advances in Software, Cyber, and AI.* AFCEA New Horizons Symposium 2023. March 27, 2023.

Ozkaya, Ipek. *AI-Augmented Software Engineering: Opportunities and Implications.* 26th Ibero-American Conference on Software Engineering (CIbSE 2023). April 26, 2023. **https://conf.researchr.org/info/cibse-2023/keynotes**

Ozkaya, Ipek. *Are You Ready to Engineer and Sustain AI Systems.* 5th Software Engineering for Machine Learning Applications Symposium. October 29, 2022. **https://semla-polymtl.github.io/semla2022/program.html**

SEI Technical Reports

Alberts, Christopher J.; Bandor, Michael S.; Wallen, Charles M.; & Woody, Carol. *Acquisition Security Framework (ASF): Managing Systems Cybersecurity Risk*. CMU/SEI-2022-TN-003. Software Engineering Institute, Carnegie Mellon University. 2022. **https://insights.sei.cmu.edu/library/acquisition-security-framework-asf-managing-systems-cybersecurity-risk/**

Booz, Jarrett; Arora, Leena; Vessella, Joseph; Kaar, Matt; Allen, Dennis M.; & Hammerstein, Josh. *Challenge Development Guidelines for Cybersecurity Competitions*. CMU/SEI-2022-TR-005. Software Engineering Institute, Carnegie Mellon University. 2022. **https://insights.sei.cmu.edu/library/challenge-development-guidelines-for-cybersecurity-competitions/**

Chick, Timothy A.; Pavetti, Scott; & Shevchenko, Nataliya. *Using Model-Based Systems Engineering (MBSE) to Assure a DevSecOps Pipeline is Sufficiently Secure*. CMU/SEI-2023-TR-001. Software Engineering Institute, Carnegie Mellon University. 2023. **https://insights.sei.cmu.edu/library/using-model-based-systems-engineering-mbse-to-assure-a-devsecops-pipeline-is-sufficiently-secure/**

Kazman, Rick; Echeverría, Sebastián; & Ivers, James. *Holistic View of Architecture Definition, Evolution, and Analysis*. CMU/SEI-2023-TR-004. Software Engineering Institute, Carnegie Mellon University. 2023. **https://insights.sei.cmu.edu/library/a-holistic-view-of-architecture-definition-evolution-and-analysis/**

2023 Featured Research Teams

SEI Team Leads First Independent Study on Technical Debt in Software-Intensive DoD Systems
Ipek Ozkaya and Brigid O'Hearn (project leads), Julie Cohen, Forrest Shull

p. 4

President's Cup Competition Expands Access to SEI Cybersecurity Simulations
Josh Hammerstein (project lead), Jarrett Booz, Rotem Guttman, Dominic Ross

p. 4

SEI Establishes First AI Security Incident Response Team
Lauren McIlvenny and Mike Mattarock (project leads), Eric Heim, Shing-hon Lau, Nathan VanHoudnos

p. 5

SEI Quantum Experts Join Pittsburgh Computing Research Organizations
Daniel Justice, Jason Larkin

p. 5

CISA Adapts Innovative SEI Approach to Transform Vulnerability Management Landscape
Allen Householder (technical lead), Eric Hatleback, Vijay Sarvepalli, Jonathan Spring, Laurie Tyzenhaus, Chuck Yarbrough

p. 8

CMU Collaborations Enhance Outcomes for U.S. Government
Will Kleiber, Kris Rush, Hasan Yasar, Scott McMillan, Josh Hammerstein, Grace Lewis

p. 9

Reducing the Risk of UEFI's Hidden Security Challenges
Vijay Sarvepalli (project lead)

p. 10

Workshop Identifies Critical Needs for U.S. Leadership in Software Engineering, AI Engineering
Anita Carleton (project lead), Erin Harper, Ipek Ozkaya, John E. Robert, Douglas Schmidt, Forrest Shull

p. 13

Bringing Zero Trust Practices to Army Tactical Networks
Tim Morrow (project lead), John Yager, Tom Scanlon, Dan Costa, Alfred Schenker, Mary Catherine Ward, Nadine Bodnar, Alexander Curtis, Chad Hershberger, Ryan Lehman, Nicholas O'Connor, Andrew Schlackman, Gregory Seroka, Mary Warren, Andrew Wilkey, David Schulker, Jeff Mellon, Nicole Pavetti, Austin Whisnant

p. 14

SEI Support for Long Range Standoff Program Spurs New Engagements
David Walbeck (project lead), Stephen Beck

p. 16

Assuring Trustworthiness of AI for Warfighters
Christopher Fairfax (program manager), Robert Beveridge, Cole Frank, Jonathan Frederick, Matt Gaston, Derek Gobin, Matt Hale, Eric Heim, Daniel Justice, Rick Labiak, Mike Mattarock, Andrew Mellinger, William Nichols, Carol Smith, John Stogoski, Oren Wright

p. 19

Pathfinding Project Explores Large Language Models for the Intelligence Mission
Shannon Gallagher (project lead), Hollen Barmer, Robert Beveridge, Tyler Brooks, Bryan Brown, Eric Heim, Angel McDowell, Andrew Mellinger, Will Nichols, Swati Rallapalli, Jasmine Ratchford, Nathan VanHoudnos, Nick Winski

p. 20

Supporting the Human and Technical Elements of Responsible AI for National Defense
Robert Beveridge, Matthew Hale, Sumanyu Gupta, Katie Robinson, Carol Smith, Alex Steiner

p. 22

Extending SysML V2 with AADL Concepts to Support Engineering and Certification of Safety-Critical Systems
Jérôme Hugues, Joseph Seibel, Lutz Wrage

p. 24

Copyright

Credits

Manager, Communication Services

Janet Rex

Manager, Public Relations

Richard Lynch

Manager, Communication Design

Cat Zaccardi

Manager, Technical Communication

Tamara Marshall-Keim

Editor-in-Chief

Paul Ruggiero

Editorial

Jenna Bodnar

Ed Desautels

Claire Dixon

Felicia Evans

Lope Lopez

Sheela Nath

Sandy Shrum

Pennie Walters

Barbara White

Design

Christopher Baum

Illustration

Christopher Baum

David Biber

Kurt Hess

Photography

Carnegie Mellon University
Communications & Marketing
Photography

David Biber

Digital Production

Mike Duda

SEI Pittsburgh, PA
4500 Fifth Avenue
Pittsburgh, PA 15213-2612

SEI Arlington, VA
4301 Wilson Boulevard
Suite 200
Arlington, VA 22203