2016 YEAR IN REVIEW

Carnegie Mellon University
Software Engineering Institute

The Software Engineering Institute (SEI) is a federally funded research and development center (FFRDC) sponsored by the U.S. Department of Defense and operated by Carnegie Mellon University.

The SEI's mission is to advance the technologies and practices needed to acquire, develop, operate, and sustain software systems that are innovative, affordable, trustworthy, and enduring.

The 2016 SEI Year in Review highlights the work of the Institute undertaken during the fiscal year spanning October 1, 2015, to September 30, 2016.

MESSAGE FROM THE DIRECTOR

Our military and civilian organizations increasingly look to software for ways to reach better decisions and assure their missions. This heavy reliance on software continually creates new challenges for the men and women of the Carnegie Mellon University Software Engineering Institute (SEI). In 2016, SEI technical staff tackled a range of tough, emerging problems confronting the Department of Defense (DoD), civilian government agencies, and industry. Among these are challenges related to increased software complexity, security concerns arising from hyperconnectivity, the effective adoption of autonomous systems, and the best ways for humans and machines to interact. The SEI's R&D in these and other areas addresses immediate needs facing the organizations we serve and helps the software engineering and cybersecurity communities understand challenges that lie just over the horizon.

One such challenge is establishing trust between humans and autonomous systems. The Defense Science Board's study on autonomy, which I had the honor to co-chair, published a report in 2016 that underscores the importance of this trust. Human-robot partnerships can maintain the nation's edge in mission settings, and assuring trust is crucial to accelerate DoD adoption of autonomous systems. SEI research on "Why Did the Robot Do That?" aims to build trust by creating a means for autonomous systems to explain their actions using natural language. When humans understand why autonomous systems behave as they do, their trust in these systems grows.

In other work related to autonomy, the SEI's participation in the Robotic Operating System for military robots (ROS-M) is helping to foster innovation and security in unmanned systems while reducing system development time and costs. Our experts made key contributions to the ROS-M Cybersecurity and Software Process working groups and worked specifically to support the U. S. Army Tank Automotive Research, Development, and Engineering Center (TARDEC) on this effort.

Automation can also provide the DoD a critical edge in software development. To this end, the SEI is researching the use of automated code repair to reduce software vulnerabilities. Our researchers are working with the DoD Software Assurance Community of Practice Working Group and others to produce tools they hope will reduce secure coding rule violations requiring manual inspection by two orders of magnitude.

All SEI technical work demonstrates our ongoing commitment to fulfilling our mission as a DoD research and development center focused on software and cybersecurity concerns.

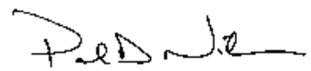

Paul D. Nielsen
Director and CEO

TABLE OF CONTENTS

STRATEGY

The SEI achieves its goals through technology innovation and transition. The SEI creates usable technologies, applies them to real problems, and amplifies their impact by accelerating broad adoption.

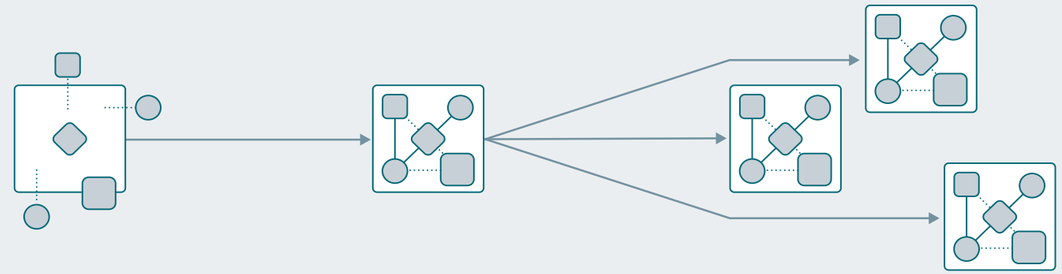

CREATE

The SEI addresses significant and pervasive software engineering problems by

- motivating research
- innovating new technologies
- creating prototypes and open-source software
- identifying and adding value to emerging or underused technologies
- improving and adapting existing solutions

SEI technologies and solutions are suitable for application and transition to the software engineering community and to organizations that commission, build, use, or evolve systems that are dependent on software. The SEI partners with innovators and researchers to implement these activities.

APPLY

The SEI applies and validates new and improved technologies and solutions in real-world government and commercial contexts. Application and validation are required to prove effectiveness, applicability, and transition potential. Solutions and technologies are refined and extended as an intrinsic part of the application activities.

Government and commercial organizations directly benefit from these engagements. In addition, the experience gained by the SEI informs

- the “Create” activities about real-world problems and needed adjustments, technologies, and solutions
- the “Amplify” activities about needed transition artifacts and strategies

The SEI works with early adopters to implement the “Apply” activities.

AMPLIFY

The SEI works through the software engineering community and organizations dependent on software to encourage and support the widespread adoption of new and improved technologies and solutions through

- advocacy
- web-based communication and dissemination
- books and publications
- certifications
- courses
- leadership in professional organizations
- licenses for use and delivery

The SEI accelerates the adoption and impact of software engineering improvements.

The SEI engages directly with the community and through its partners to amplify its work.

GENERATING IMPACT

The mission of SEI technical work is to improve efforts in the U.S. Department of Defense (DoD) and other organizations to obtain the benefits from software while controlling cost and risk associated with their software-enabled systems. Improvements result from capabilities we deliver that provide confidence in the behavior of large software-based systems.

Our approach to generating impact involves a strategic technical framework and a customer engagement strategy.

STRATEGIC TECHNICAL FRAMEWORK

Our work delivers measurable impact through capabilities such as tools and models for decision analytics, risk-reducing virtual system integration, software cost management, and automating infrastructure cybersecurity. Research and development activities in seven core areas underpin the capabilities we deliver:

- **Autonomy and Counter-Autonomy**. Develop and apply methods and technologies for autonomous and semi-autonomous systems, including the development and understanding of evidence that indicates the trustworthiness, dependencies, and vulnerabilities of autonomous systems.
- **C4ISR Mission Assurance**. Develop methods for C4ISR systems to effectively adapt or predictably degrade while continuing to effectively achieve their missions.
- **Cybersecurity**. Develop improved systems, repeatable practices, and capable personnel to enable cyber missions.
- **Data Modeling and Analytics**. Develop and apply mathematically rigorous data collection, analysis, and visualization techniques for system acquisition, development, adaptation, feedback, and algorithms in support of national defense missions.
- **Human-Machine Interactions**. Invent, assess, and improve comprehensible, safe, and trustworthy techniques and technologies for humans to use and team with machines.
- **Software Engineering and Information Assurance**. Develop and apply practices and tools that enable the acquisition, development, and fielding of high-quality, secure software-based systems in a predictable and affordable manner.
- **System Verification and Validation**. Build and apply practical, mathematically grounded, and evidence-based methods and tools to increase confidence in the entire systems engineering lifecycle and the quality of the resulting systems.

ENGAGEMENT STRATEGY

The SEI conducts applied research and development with funding from the Office of the Secretary of Defense, project work plans with the DoD and other federal agencies, and collaborative research and development agreements with non-federal organizations, including industry.

To coordinate this portfolio of work in a way that facilitates transition from the lab to the field, the SEI relies on an engagement strategy. This strategy ensures that the SEI is aligned with the needs of the DoD, can innovate to solve DoD challenges, and establishes the needed organizational relationships.

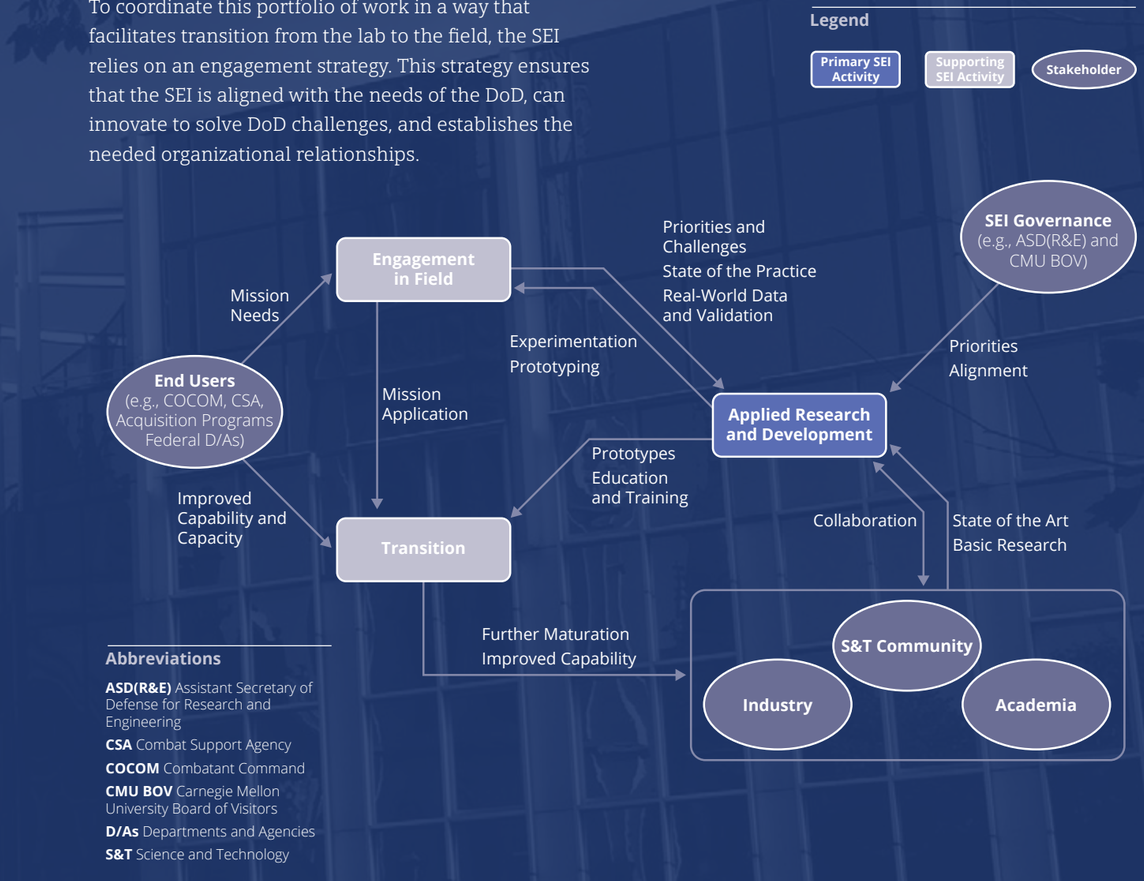

FUNDING SOURCES

In FY 2016, the SEI received funding from a variety of sources in the Department of Defense, civil agencies, and industry.

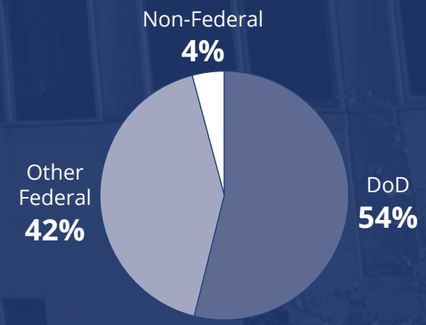

IN THE NEWS

Peter Feiler Named SEI Fellow

In 2016, the SEI honored principal research scientist Peter Feiler by naming him an SEI Fellow. Feiler, whose career at the SEI spans 31 years, became the eighth SEI Fellow, a designation bestowed on staff who have made an outstanding contribution to the work of the SEI and to whom SEI leadership looks for valuable advice on advancing the Institute's mission.

"Peter joins a select group of SEI legends who have contributed so much, not only to SEI, but to the nation and the global software engineering community," said SEI Director and CEO Paul Nielsen.

Feiler is the technical lead and author of the SAE AS-2C Architecture Analysis & Design Language (AADL) standard. AADL is a framework that allows analysis of system (and system of systems) designs prior to development and supports an architecture-centric, model-based development approach throughout the system lifecycle. AADL lowers development and maintenance costs while improving reliability and safety. Feiler has lent his expertise in this area to several Department of Defense projects, including the Joint Multi-Role Technology Demonstrator, Future Vertical Lift, and the DARPA High-Assurance Cyber Military Systems program.

"I appreciate the recognition for my contributions as an SEI member, in particular in the last 16 years as technical lead of the SAE AADL standard," said Feiler. "This would not have been possible without contributions by the SAE AADL committee members and by the SEI team working with me."

Before joining the SEI, Feiler led research on software technology at the Siemens Corporation, where he also served as system architect for the software development environment in a large-scale product development. Feiler holds a PhD in computer science from Carnegie Mellon University and is a senior member and member of ACM, IEEE, and SAE International.

Hayes Testifies on Use of Agile Approaches in Social Security Systems Modernization

On July 14, 2016, the House Subcommittee on Social Security convened a hearing on the proposed Social Security Administration (SSA) information technology modernization plan, which incorporates Agile approaches.

To lend insight on these issues, the subcommittee called on the SEI's Agile in Government (AIG) team, represented by Will Hayes, to testify. Hayes is principal engineer on the AIG team. His oral testimony reinforced points expressed in written testimony prepared by Hayes, Suzanne Miller, Eileen Wrubel, and Alyssa Le Sage.

The SSA, despite recent updates to its physical assets, such as computers and data centers, continues to struggle with applications created decades ago in the COBOL programming language. While Agile practices can help the SSA develop new and more effective capabilities, it can also pose challenges to traditional government oversight and management practices.

In his testimony, Hayes reminded the subcommittee that the planning and development cadence of Agile approaches "place a premium on consistent use of short iterations with stable staffing dedicated to a single stream of technical work. This new cadence offers more oversight opportunity, but with different measures of success."

Hayes also noted that Agile relies on uncovering user needs through collaborative interaction. This approach can be difficult to achieve in a government setting. "It is not yet clear how we will build the capacity for government personnel to interact more frequently with developers," said Hayes. The SSA workforce, he noted, is already being asked to accomplish more with limited resources.

SEI Security Services Team Earns Defense Security Service Honor

On June 8, 2016, the Defense Security Service (DSS) honored the SEI Security Services team with the James S. Cogswell Outstanding Industrial Security Achievement Award. SEI Security Services was one of just 42 honorees chosen from approximately 13,500 cleared facilities, all of which undergo recurring assessment by DSS through the National Industrial Security Program. DSS bestows this honor to security services teams who maintain the highest standards for security, exceed National Industrial Security Program Requirements, and demonstrate leadership in establishing best practices. DSS presented the award at the annual NCMS training seminar in Nashville, Tenn.

Representing the SEI were Chief of Staff John Bramer, Chief Information Officer David Thompson, Facility Security Officer Kara Branby, Industrial Security Specialist Bryan Stake, Industrial Security Administrator Allison Zust, Industrial Security Specialist Angela Raible, Security Manager Jason Hawk, and Information Assurance Coordinator Ryan Gindhart.

"The SEI is fortunate to have the team of dedicated security and information assurance professionals that we do," said Thompson. "They consistently ensure that we as an organization fulfill our role in properly handling and protecting the sensitive information entrusted to us by our sponsor. It was a true pleasure to see their work recognized by the government."

To learn more about the James S. Cogswell Award, visit **www.dss.mil/isp/partnership.html**.

CHIEF TECHNOLOGY OFFICER (ACTING)
JEFFREY BOLENG

Assuring the Software That Enables Autonomy

In 2014, the U.S. Department of Defense (DoD) unveiled its Third Offset Strategy. Twice before, the military has pursued offset strategies: in the 1950s with nuclear deterrence and in the 1970s–1980s with advanced intelligence, surveillance, and reconnaissance (ISR) platforms, stealth technology, and other thrusts. The goal has been to use technology to offset adversaries' numerical advantages.

The Third Offset targets system autonomy and other technologies. Autonomy involves systems taking action on behalf of humans. The DoD goal is to use autonomy to keep humans out of harm's way, ensure faster and more accurate decision making, automate cyber defense, and process massive amounts of sensor data on a scale humans cannot approach.

The essence of autonomous systems or automation is software. Autonomy has entered our consciousness because of machines such as unmanned aerial systems (UAS) and driverless cars. Software enables the machine to understand its environment and react accordingly, sensing when it is operating beyond its design specifications and modifying its behaviors based on what it learns about its operating environment.

The 2016 Defense Science Board study on autonomy, co-chaired by SEI Director and CEO Paul Nielsen, draws attention to the need to address human and societal trust in autonomous systems in order to accelerate DoD adoption of autonomous capabilities.

At the SEI, we established autonomy and counter-autonomy as one of our technical work areas with the aim of building and understanding the evidence that autonomous systems are trustworthy and operating to their intended behaviors. We are pursuing four concerns:

- **Explainability**: SEI research that spans FY16 and FY17 explores trust in robot behavior. We are developing algorithms that enable robots to explain their actions in plain language that humans can understand and that can predict behavior.
- **Repetition and training**: New research in FY17 will lead toward the creation of a novel set of battlefield capabilities that integrate cyber effects with tactical operations for the front-line soldier.
- **Building trustworthy systems from untrustworthy components**: Our ongoing work in verifying distributed, adaptive, real-time (DART) systems (such as autonomous multi-UAS missions) is producing validated techniques to assure systems that may be composed of components with unknown provenance.
- **Continuous runtime verification**: Many autonomous systems rely on components that use machine learning to achieve mission success and cannot be fully verified prior to deployment. We are addressing continuous runtime assurance challenges in new research.

In executing our autonomy and counter-autonomy research portfolio, the SEI benefits from its collaborations with CMU. We conduct robotic research, for example, at the National Robotics Engineering Center. Our researchers also collaborate with world-leading CMU faculty in areas such as biometrics and human–machine interaction.

By virtue of its institutional knowledge about software and cybersecurity, the SEI is in a prime position to contribute to solving trust issues in autonomy that have been identified as a critical obstacle to the adoption of autonomous capabilities by the DoD.

At the SEI, we established Autonomy and Counter-Autonomy as one of our technical work areas with the aim of building and understanding the evidence that autonomous systems are trustworthy.

Photo: CMU NREC

Collaborations with Carnegie Mellon University Drive Several Key Projects

The SEI has a history of driving results by collaborating with government, industry, and academia. As a federally funded research and development center located at Carnegie Mellon University (CMU), the SEI has ready access to collaborative opportunities with leading researchers in fields essential to our work and that help advance the state of the art in software engineering and cybersecurity.

In 2016, the SEI worked with experts from CMU on a number of projects highlighted in the Year in Review. For instance, Will Klieber and other SEI staff working to advance the field of automated code repair (see page 36) collaborated with CMU professor Claire Le Goues, a leading researcher in the use of genetic programming for automated code repair. In their approach, genetic programming uses computational analogs of biological mutation and crossover to generate new program variants and to search for a variant that produces the desired result for all test cases. The SEI team also worked with CMU's Christian Kästner, who has pioneered work on symbolically analyzing code under all possible build configurations. The goal of this collaborative effort is to develop repairs that work for all possible build configurations.

SEI researcher Stephanie Rosenthal is working on ways to develop trust in robots and to understand, through straightforward verbal communication, why a robot behaved the way it did in certain situations. (See "Why Did the Robot Do That" on page 12.) Rosenthal collaborated on the natural language component of this project with Siddhartha Srinivasa of CMU's Robotics Institute and Manuela Veloso of CMU's Machine Learning Department. The three are investigating how robots can communicate in plain English about the actions they take and the decisions they make.

Another SEI team is researching solutions for verifying safety- and mission-critical systems (see page 18). This work, undertaken by the SEI's Sagar Chaki, Scott Hissam, and Dionisio De Niz, centers on two projects: Verifying Distributed, Adaptive, Real-Time (DART) Systems and Auto-Verification of Software with Timers and Clocks (STAC). The work aims to head off problems in complex Department of Defense (DoD) systems, such as missile defense, or safety features, such as automatic braking in automobiles. The SEI team is working with CMU's National Robotics Engineering Center, which is applying the SEI's techniques in Husky, an all-terrain robotic development platform.

These collaborative projects represent just a few of the ways in which the SEI engages in mutually beneficial collaborations with colleagues on the CMU campus. By teaming with experts at CMU, a global research university annually rated among the best for its programs in computer science and engineering, the SEI advances the state of the art and tackles some of the toughest challenges facing the DoD and industry.

RESEARCHER
STEPHANIE ROSENTHAL

Why Did the Robot Do That?

A first responder is looking for disaster survivors with a search-and-rescue robot. Suddenly, the robot swerves. Why did the robot do that? Did it spot a victim? Avoid danger? Malfunction?

"It's not always clear why a robot acts the way it does," said Stephanie Rosenthal of the SEI's Emerging Technology Center. Autonomous robots sense their environment and use this information to decide what actions to take. Bystanders can only guess how these robots make decisions by observing their behavior. "If humans don't understand a robot's reasoning, how can they trust it to do its job?"

Rosenthal wants to build this trust through verbal communication, or natural language. "A robot needs to explain what it's doing in a way that's easy for people to understand," she said. With Siddhartha Srinivasa of Carnegie Mellon University's (CMU) Robotics Institute and Manuela Veloso of CMU's Machine Learning Department, she's investigating how robots can communicate in plain English about the actions they take and the decisions they make. This helps the robot's users to understand its behavior.

A pilot project with Joshua Peschel of Iowa State University is putting Rosenthal's technology into action. It's being deployed on robotic boats developed by Peschel's company, Senformatics. The pilot will study whether the boat's explanations affect user trust among water rescuers, environmental monitors, and other users.

Trust isn't just an academic issue. Autonomous robots deliver medications in hospitals and move goods in warehouses, and self-driving cars are taking to the streets. As robots grow more sophisticated, they'll interact more closely with people, who will need to know if they work properly. Will humans feel the need to constantly supervise robots? Will they perform a dangerous task themselves instead of letting a robot do it? This lack of trust undermines the very idea of human–robot partnerships.

Rosenthal chose natural language because it expresses more information than flashing lights and other non-verbal signals. However, natural language also makes communication more complicated. Poor language choices cause misunderstandings even among humans, let alone humans and machines. To identify key words and phrases that describe a robot's actions, Rosenthal turned to crowd sourcing. She created tasks for crowd members to perform, collected their explanations, and extracted the vocabulary and language patterns used most often.

To find out which words and phrases were easiest to understand, each crowd member read a sentence and described how a robot would act. The most accurate words and phrases became the building blocks of new explanations. The result? Clear, understandable explanations of the robot's activities.

Rosenthal is also investigating what kinds of explanations people prefer to hear. Are detailed explanations always necessary? Could a physical demonstration of the robot behavior help instead of an explanation? Rosenthal is working on algorithms to tailor explanations to each user's preferences. Everyone should get exactly the information needed to understand what the robot did.

Rosenthal's plans for future research include creating explanations and demonstrations that help people to predict a robot's future behavior.

For more about this work, visit insights.sei.cmu.edu/sei_blog/2016/12/why-did-the-robot-do-that.html.

Photo: U.S. Army

"If humans don't understand a robot's reasoning, how can they trust it to do its job?"

— STEPHANIE ROSENTHAL, SEI EMERGING TECHNOLOGY CENTER




RESEARCHERS
SARAH SHEARD, MIKE KONRAD, WILLIAM NICHOLS, CHARLES B. WEINSTOCK

SEI Determines the Effects of System Complexity on Aircraft Safety for the FAA

In the realm of aerospace, software error can be catastrophic. The SEI has been working on the challenge of software complexity in aerospace systems to understand and prevent such catastrophes. Our long history of work with the System Architecture Virtual Integration Program (SAVI) represents one thread of our research in this area.

A number of aerospace stakeholders participate in SAVI, including the Federal Aviation Administration (FAA), whose goal is to lower development costs of complex aerospace systems. In 2014, because of the SEI's history of collaboration in SAVI, the FAA awarded the SEI a two-year project to research the effect of complexity on aircraft safety.

The SEI research team, led by Sarah Sheard, included Mike Konrad, William Nichols, and Charles B. Weinstock. The team investigated how complexity manifests in software-reliant systems of the avionics domain, how to measure that complexity early in the development lifecycle with virtual models, and how to tell when too much complexity might lead to safety problems.

This work culminated in a formula for calculating how many ways a failure can propagate from one system component to another. This information can be used as a basis for estimating the size of a safety argument.

"This result will strengthen the case for aircraft and parts manufacturers to address complexity by using safety assurance cases," said Sheard. "It will also help manufacturers understand the reasons for creating and maintaining safety cases, and it will help manufacturers estimate the effort required to demonstrate safety."

The FAA invited the SEI team to present its results to the FAA and members of the aircraft industry at the 2016 FAA Streamlining Assurance Processes Workshop.

To learn more about this research effort, see the FAA project's series of papers in our digital library: resources.sei.cmu.edu/library/asset-view.cfm?assetID=483758.

RESEARCHER
ART MANION

CERT Division Works with DoD and DC3 to Shape Vulnerability Disclosure Policy

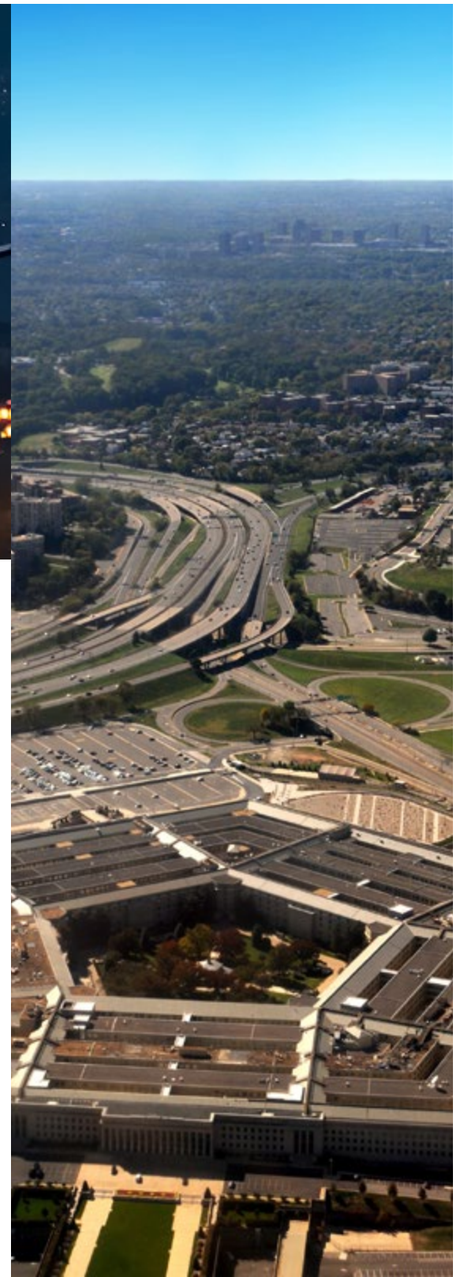

For years, when security researchers in the field discovered a vulnerability in an organization's software or systems, they could safely report their findings under the organization's vulnerability reporting policy. This was not the case, however, when it came to government and Department of Defense (DoD) systems. Why? Because in the absence of clear vulnerability reporting and disclosure policies, researchers feared legal consequences. In such an environment, organizations might find out about a vulnerability only after it has been publicly disclosed or used in attacks. Sometimes the organization never learns of the vulnerability, and you can't defend against something you don't know about.

In 2016, however, the DoD began to put in place policies to foster a closer relationship with the security research community. In particular, the DoD announced a vulnerability disclosure policy to provide clear guidelines to researchers conducting vulnerability research on DoD web properties. The aim is to foster good-faith research that can inform DoD security efforts that help ensure the DoD accomplishes its mission in defense of the United States.

"The Vulnerability Disclosure Policy is a 'see something, say something' policy for the digital domain," said former Secretary of Defense Ash Carter in a DoD press release. In the press release, Carter noted the DoD's interest in encouraging the legitimate work of computer security researchers. "This policy gives them a legal pathway to bolster the department's cybersecurity and ultimately the nation's security," he said.

In creating this new policy, the DoD worked with the DoD Cyber Crime Center (DC3) and the SEI's CERT Division. Both the DoD and DC3 drew upon CERT's nearly 30 years of coordinated vulnerability disclosure experience to help shape the policy. CERT experts provided advice on policy and helped create processes that are flexible enough to handle the many exceptions that arise during coordinated vulnerability disclosure.

"The DoD program sets an example for other organizations," said Art Manion, technical manager of the CERT Vulnerability Analysis team. "All software, systems, and sites have vulnerabilities. Mature organizations recognize this and focus on coordinated disclosure policy and practices, which is exactly what the DoD is doing with this new policy."

To review the DoD Vulnerability Disclosure Policy, visit hackerone.com/deptofdefense.





RESEARCHERS
TIMOTHY CHICK, CHRIS INACIO, ANGELA MOSQUEDA, KEN NIDIFFER, TOM SCANLON

SEI Plays Key Role in JFAC Stand Up

To truly protect DoD systems, the security focus needs to shift from a perimeter-defense-focused approach to an engineered-in approach."

— TIMOTHY CHICK, SEI CERT DIVISION

Photo: U.S. Navy

In 2015, the Department of Defense (DoD) launched the Joint Federated Assurance Center (JFAC), a federation of DoD organizations that promotes and enables software and hardware assurance within defense acquisition programs, systems, and supporting activities. JFAC member organizations and their technical service providers interact with program offices and others to provide software and hardware assurance expertise and support, including vulnerability assessment, detection, analysis, and remediation services. JFAC also provides information about emerging threats and capabilities, software and hardware assessment tools and services, and best practices.

Drawing on the SEI's long experience in the field of software assurance, its expertise in establishing computer security incident response teams, and other fields related to JFAC's mission, the Deputy Assistant Secretary of Defense for System Engineering (DASD[SE]) engaged the SEI to support JFAC's mission. The SEI's primary focus was on standing up the JFAC Coordination Center (JFAC-CC) to establish initial operating capability. Contributing to this effort were the SEI's Tim Chick, Chris Inacio, Angela Mosqueda, Ken Nidiffer, and Tom Scanlon. Among the tasks the SEI team was charged with were the following:

- administering the JFAC-CC, a support operation that evaluates and analyzes user issues and coordinates resolution
- increasing software assurance awareness by providing training and software assurance tool demonstrations to the DoD community
- conducting a comprehensive analysis of COTS software assurance tools—and their licenses—used by JFAC
- contributing to a gap analysis of software assurance technology and a user experience report on JFAC tools

"To truly protect DoD systems, the security focus needs to shift from a perimeter-defense-focused approach to an engineered-in approach. A key part of that for JFAC is shifting software assurance from a step in a development process to an integrated part of the process from end to end," said Chick. He also noted that the JFAC initiatives have made great strides toward improving the security of DoD applications by increasing the awareness of the tools and techniques used to achieve software assurance and by connecting programs throughout the DoD with expert resources.

The SEI also supported the launch of the JFAC enterprise software-licensing pilot, which put software assurance tools in the hands of over 60 different DoD groups. The pilot provided immediate impact: millions of lines of code were scanned and thousands of potential security issues were detected and addressed. "I expect the long-term impact of the various JFAC initiatives to have an even greater impact going forward," said Chick.



RESEARCHERS
SAGAR CHAKI, SCOTT HISSAM, DIONISIO DE NIZ

New Solutions for Verifying Safety- and Mission-Critical Systems

An incident report for the MIM-104 Patriot system noted a system clock that was off by only 0.3 seconds prevented the system from detecting an incoming missile. Carmaker Acura issued a recall to repair automated safety systems that incorrectly braked for a non-existent obstacle when traveling next to a guard rail. Software errors such as these demonstrate the challenges that the Department of Defense (DoD) shares with industry when developing complex systems. Verification of these systems is especially problematic when they have safety-critical, autonomous, distributed, adaptive, and real-time components. The systems must meet the challenge of dual verification: logical correctness of instructions and execution at the right time.

Traditional verification techniques are inadequate for the scale and complexity of today's software-reliant systems. Needed capabilities take too long to field, largely because testing is a lengthy process. The SEI's Dionisio de Niz noted, "With testing, you find only the kinds of errors that you test for." His team member Scott Hissam added, "When cyber-physical systems interact with the environment, there are infinite possibilities. You can't test for them all."

For example, sensing and actuation must occur in sync with events in the environment. The time between sensing a car crash and inflating the airbag should not exceed 20 milliseconds. No amount of testing can include all possible forms of this interaction with the environment. Verification has broader coverage for potential errors and can be applied earlier in the development lifecycle. Consequently, it can reduce cost by revealing errors earlier.

The SEI has researched verification problems for more than 20 years and made noted advances in software model checking and static analysis. In two recent projects—Verifying Distributed, Adaptive, Real-Time (DART) Systems and Auto-Active Verification of Software with Timers and Clocks (STAC)—team lead Sagar Chaki and team members de Niz and Hissam continue this work to improve verification techniques. To help ensure that the research is targeted to DoD-relevant problems, Chaki's team works with Stanley Bak of the Air Force Research Laboratory at Wright-Patterson Air Force Base. And the Carnegie Mellon University National Robotics Engineering Center is applying the SEI team's techniques in the Husky system, an all-terrain robotic development platform.

For DART, the SEI team developed a method to produce high-assurance software for cyber-physical systems composed of multiple agents, such as a team of robots that communicate, coordinate, and adapt to uncertain environments to achieve safety-critical and mission-critical goals. They created an architecture that isolates the safety-critical parts of the system from the mission-critical ones. Then they used automated analyses that allow DART systems to self-adapt in changing environments.

For STAC, the team investigated formal verification of safety properties in software that accesses system clocks and uses their values to set timers and perform computations. An important contribution of this research is its model of time in network behavior. The approach includes verifying STAC systems at the source code level, thereby reducing the differences between the verified system and the executed system.

"The goals of these two projects are tightly connected," Chaki explains; "together, they address logical verification and timing verification, so that software does the right thing at the right time."

Both projects also use automation to scale verification methods to complex, distributed, real-time systems. These new techniques for verification build on the SEI's substantial body of work to improve verification methods and reduce the cost of assurance.

For more about STAC, visit insights.sei.cmu.edu/sei_blog/2016/12/verifying-software-with-timers-and-clocks-stacs.html.

For more about DART, visit insights.sei.cmu.edu/sei_blog/2016/10/verifying-distributed-adaptive-real-time-systems.html.

Photo: U.S. Air Force

> *When cyber-physical systems interact with the environment, there are infinite possibilities. You can't test for them all."*
>
> **— SCOTT HISSAM, SEI SOFTWARE SOLUTIONS DIVISION**

Photo: U.S. Air Force

> *Our scientific approach is novel because of its use of multiple static analysis tools, the large variety of features used to develop classifiers, and the competing classification techniques."*
>
> **— LORI FLYNN, SEI CERT DIVISION**

RESEARCHER
LORI FLYNN

Using Machine Learning to Improve Static Analysis of Source Code

Software engineers sometimes use multiple static analysis tools to detect code flaws because different tools generally find different sets of code flaws. The challenge is choosing tools that strike a balance between selectivity (a fraction of flagged problems [alerts] that are true code flaws) and sensitivity (finding code flaws). Most engineers select tools with high sensitivity for many types of potential code flaws, which produce long lists of potential coding errors, including many false positives. What's more, when engineers use multiple static analysis tools, they find more flaws, which only compounds the problem of generating too many alerts to analyze.

CERT Secure Coding team researchers developed a novel technique to address this problem that introduces machine learning to static analysis. The solution uses alert archives from multiple static analysis tools and produces sets of classifiers that accurately predict whether a static analysis alert is true or false. The classifiers are built using audit archives containing output alerts from multiple tools and other metadata for each codebase, along with analyst determinations (e.g., true or false positives) for alerts. The eventual goal is a fully automated and accurate statistical classifier integrated with an alert auditing framework that efficiently uses analyst effort and facilitates removal of code flaws. The FY16 goal was to create accurate alert classifiers for the data sets.

Three Department of Defense (DoD) collaborators participated in this research project. They used the enhanced-SCALe auditing framework tool developed as part of this project to audit their own code. Enhanced-SCALe is based on the CERT SCALe system and includes added data collection, an archive sanitizer, offline installs, and a virtual machine.

One significant finding of this project is that using the tool name as a classifier feature increased classifier accuracy. In other words, data from running multiple static analysis tools on the same codebases was helpful. The team used data from years of CERT code audits and the DoD collaborators. Many classifier variants were developed (all but one filtering total data to use a subset for the classifier). We created each classifer using a randomly selected 70 percent of that data set, then tested them on the remaining 30 percent of the data set. Classifier accuracies ranged from 88 percent to 91 percent using the largest data set.

Though the method needs to address a wider range of possible code flaws (which the team is addressing in FY17), it has been shown to decrease the amount of data engineers must manually examine. It also more accurately identifies legitimate errors. This new machine learning approach will help engineers focus their effort on fixing legitimate errors when integrated with an auditing system that uses the classifiers to order the alerts. One of the SEI's DoD collaborators plans to do this integration in FY17.

CERT researcher Lori Flynn explained, "Our scientific approach is novel because of its use of multiple static analysis tools, the large variety of features used to develop classifiers, and the competing classification techniques we compare (Random Forest, Lasso Logistic Regression, CART, and XGBoost)." Flynn added that her team's goal is to automatically classify 90 percent of flagged anomalies as true and false positives with 95 percent accuracy. If successful, the new method and subsequent software tools will significantly reduce the effort needed to inspect static analysis results and prioritize confirmed defects for repair.

For more information about the SEI's research on using machine learning to improve static source code analysis, visit http://resources.sei.cmu.edu/library/asset-view.cfm?assetid=474252.



RESEARCHERS
ROBERT BEVERIDGE, JOHNATHAN FREDERICK, GEOFF DOBSON

Air National Guard and Air Force Reserve Units Develop and Test New Skills in Cyber Lightning Challenge

In June 2016, the SEI hosted "Cyber Lightning," a three-day joint training exercise involving Air National Guard and Air Force Reserve units from western Pennsylvania and eastern Ohio. Participating in the exercise were members of the 911th Airlift Wing, operating out of the Pittsburgh International Airport Air Reserve Station; the 171st Air Refueling Wing, operating out of the Pittsburgh International Airport; and the 910th Airlift Wing, operating out of the Youngstown-Warren Air Reserve Station in Ohio.

"All the participants work in traditional base communication squadrons," said the SEI's Robert Beveridge. "Their workload in maintaining computer systems does not provide the opportunities to gain hands-on cybersecurity skills in protecting the organizational networks. The Cyber Lightning exercise provided these men and women a chance to learn and test new cybersecurity skills in an environment that mimics real Department of Defense networks, and it aligns with the desire of senior leaders in the Air Force Reserve and Air National Guard to help develop the cyber cadre."

On the first day of Cyber Lightning, SEI staff trained participants on techniques such as log analysis, firewall management, vulnerability scanning, traffic analysis, and intrusion detection systems. SEI staff also provided the participants a threat brief. The second day was devoted to mission planning for the competition phase of the exercise, and participants used what they learned on the first day to scan their networks and perform a vulnerability analysis.

"The teams found the vulnerability analysis portion challenging," said Beveridge. He noted that this part of the exercise introduced concepts such as identifying key cyber terrain, performing a qualitative risk assessment of those critical systems, and prioritizing the vulnerabilities to mitigate in a limited time frame.

On the third day, all three teams engaged in a competition in which they applied the skills and techniques they learned on Day 1 and the clues obtained during pre-planning and the network scan conducted on Day 2. Their objective was to find malicious traffic and activity on their networks.

"The teams did a good job identifying authentic malware that has been developed and used by attackers to infiltrate and steal secrets from large corporate networks over the past few years," said Jonathan Frederick, cybersecurity exercise developer and trainer at the SEI.

"This is a great effort for the squadron," said Maj. Kelly Quigley, commander of the 910th Airlift Wing communications squadron. "This is an opportunity for our men and women to learn about how cyber teams do their business and learn new skills."

Lt. Col. Joseph Sullivan of the 171st Communications Flight of the Pennsylvania Air National Guard also found value in Cyber Lightning. "The training received was relevant to our daily mission," noted Sullivan. "The additional training and exercises on intrusions and malware detection provided our base communications personnel training they haven't received to date. Even though this training doesn't make them experts, they now have a true understanding of the importance in remaining vigilant in protecting Air Force systems."

The success of Cyber Lightning could pave the way for similar events. "We hope there are future opportunities to conduct this type of exercise again with other services and other units," said Beveridge.

For more on the SEI's efforts in cyber workforce development, visit www.cert.org/cyber-workforce-development.

> *"This is an opportunity for our men and women to learn about how cyber teams do their business and learn new skills."*
>
> **— MAJ. KELLY QUIGLEY, COMMANDER OF THE 910TH AIRLIFT WING COMMUNICATIONS SQUADRON**



RESEARCHERS
ANITA CARLETON, FORREST SHULL, DAVE ZUBROW

SEI Analysis Spurs SMARTer Air Force Data System

To help Air Force acquisitions specialists make better decisions, the Air Force developed the System Metric and Reporting Tool (SMART). SMART is a key component of the Air Force's Acquisition Workbench, which is moving to an app store approach. The data in SMART is typically collected monthly and used to generate reports across multiple weapon system capability portfolios. After years of experience with the tool, the Air Force Life Cycle Management Center (AFLCMC) believed it could get more out of its investment using the SMART data more effectively to observe trends and correlations.

To help meet the challenge of getting more out of SMART to better manage this vast enterprise, Tim Rudolph, chief technical officer for the AFLCMC, turned to experts on the SEI's Measurement and Analysis team. The SEI team included Anita Carleton, deputy director, SEI Software Solutions Division; Forrest Shull, assistant director, Office of Empirical Research; and Dave Zubrow, associate director of empirical research, SEI Software Solutions Division.

In support of the overall Secretary of the Air Force (Acquisition) (SAF/AQ) functional mission, Rudolph engaged the SEI team to determine whether SMART data could identify programs benefiting from corrective action earlier in the acquisition lifecycle. By catching and fixing problems earlier, the Air Force could ensure better results at lower cost. The Air Force also wanted to use SMART data to enable ongoing analysis, improve the understanding of program health, and better predict the future health of its programs.

In support of these goals, the SEI conducted an analysis of the data reported to the SMART system by program managers and program executive offices. The team focused on data from Major Defense Acquisition Programs (MDAP) and Major Automated Information Systems (MAIS) from the top two acquisition categories (ACAT).

"We applied data mining techniques to study leading indicators in acquisition data in the SMART system," said Shull. "One key finding was that there was a bias in data reporting toward healthy assessments. More objective data, more consistent data, and more finely grained data were needed to achieve its program management goals."

"The depth and breadth of the SEI's knowledge is critical to leveraging information innovatively for the Air Force," said Rudolph. AFLCMC is pursuing the SEI's recommendations at SAF/AQ.

"Beyond the specifics of our analysis, I think our work helped communicate some broader messages about data," said Zubrow. "One is the untapped potential of longitudinal analyses and visualizations to discover patterns in program performance and health. We also noted the importance of better guidance and training as well as automated data quality checks to improve SMART data quality to make it a more useful tool for program management."

For more about SMART, visit acc.dau.mil/CommunityBrowser.aspx?id=631173.

"The depth and breadth of the SEI's knowledge is critical to the Air Force programs and innovation in general."

—TIM RUDOLPH, CHIEF TECHNICAL OFFICER, AFLCMC

Photo: U.S. Air Force

RESEARCHERS
DAN PLAKOSH, JAY MARCHETTI

Converting a Major U.S. Navy System from 32- to 64-Bit Architecture

A centralized, automated, command-and-control (C2) and weapons control system deployed by the U.S. Navy has played a key role in the United States' ability to project naval power around the globe since the 1980s. This key Navy asset was designed as a total weapon system, from detection to kill. An official Navy description notes, "The computer-based command and decision element is the core of the combat system. This interface makes the system capable of simultaneous operations against multi-mission threats: anti-air, anti-surface and anti-submarine warfare."

But in the spring of 2015, the Navy faced a difficult task: it needed to update this weapons system by converting its basic software architecture from a 32-bit foundation to a 64-bit foundation. "It was a major undertaking, one that could potentially affect millions of lines of computer code," noted Jay Marchetti, a senior member of the SEI's technical staff.

The Navy asked its contractor for the system to assess the risks and schedule for the conversion. And, aware of the SEI's reputation as an unbiased, independent expert in software engineering, the Navy also asked the SEI for a second opinion regarding the scope, costs, portability, and risks associated with the migration of such an important system from one architecture to another.

"The resulting engagement was good for the Navy and good for the SEI," said Dan Plakosh, a senior engineer at the SEI who worked on the project with Marchetti. "We were able to help them, while at the same time demonstrating that recent advances in code analysis are applicable in large projects."

Through the SEI's analysis, the Navy gained a clearer picture of the 64-bit migration, including the amount of effort it would likely take, how it could be undertaken incrementally, and the technology trends driving the required completion time frame. "We were able to deploy tools across a much wider swath of the code than prior analyses, providing higher confidence in the migration effort estimates as opposed to manually reviewing just a fraction of the total codebase and then extrapolating those figures for the full project," Marchetti explained.

The SEI's review—using automated tools and the latest static analysis techniques—looked at a substantial portion of the Navy weapons system's code. By so doing, the SEI team demonstrated that static analysis tools are essential for accurately identifying 32- to 64-bit conversion risk areas, particularly in very large codebases. "The SEI approach was faster and much more accurate in finding conversion risks, reducing the overall risk and cost for the program," Plakosh said.

The engagement also uncovered research opportunities in static analysis tools that are funded in the SEI's FY17 research portfolio, and it spurred the Navy to request the SEI to propose a 64-bit migration prototype effort for one of the weapon system's elements to develop and document the tools and processes utilized.

> *"It was a major undertaking, one that could potentially affect millions of lines of computer code."*
>
> **— JAY MARCHETTI, SEI SOFTWARE SOLUTIONS DIVISION**

Photo: U.S. Navy

Photo: U.S. Army

The enhancements in ROS-M will help the military to develop secure, reliable robots that can carry out a wide variety of missions, from transporting goods in self-driving convoy trucks to disposing of explosives.




RESEARCHERS
JONATHAN CHU, ANDREW MELLINGER, DAN KLINEDINST, NEIL ERNST

SEI Lays the Groundwork for Open-Source Operation of Military Robots

Creating an environment for developing unmanned systems that spurs innovation and reduces development time: That's the goal of the ROS-M project, an add-on to the open-source Robotic Operating System (ROS) that's designed to meet the unique needs of military robots. In 2016, SEI researchers provided key input to the technical scope of ROS-M.

"The SEI made significant contributions to the ROS-M Cybersecurity and Software Process working groups," said Jonathan Chu of the SEI's Emerging Technology Center (ETC). "Our job was to fill strategic gaps in the working groups in which we had the most involvement." Other SEI participants in these working groups included Andrew Mellinger (ETC), Dan Klinedinst (CERT Division), and Neil Ernst (Software Solutions Division).

ROS-M builds software and hardware simulation tools, cyber assurance checking, a code repository, and a training environment for warfighters into the commercial version of ROS 2.0. These enhancements will help the military to develop secure, reliable robots that can carry out a wide variety of missions, from transporting goods in self-driving convoy trucks to disposing of explosives.

"We want to build upon the success of ROS in academia, industry, and events like the DARPA Robotics Challenge," said Chu. ROS-M promotes code sharing and reuse, especially for components whose distribution is restricted due to national security or export control concerns. This will reduce risks associated with testing, lower development costs, and foster cooperation between researchers and developers.

The SEI's work on ROS-M is part of its support for the U.S. Army Tank Automotive Research, Development and Engineering Center (TARDEC), which researches and develops advanced technologies for ground systems.

SEI participants in the Cybersecurity working group found that requirements for Department of Defense (DoD) systems do not always provide guidance for implementing prescribed controls for unmanned ground vehicles and mobile systems. Many of these requirements are verified by the organization responsible for the systems, and it is not clear who owns them. To help unmanned systems that use ROS-M to achieve their cyber maturity goals, they strongly recommended identifying which organizations own requirements and drive their enforcement.

SEI participants in the Software Process working group were particularly concerned with sustaining system adoptability and facilitating the transition of systems to the DoD. They assumed that, while system components would be ready for ROS integration, they would not contain many commercial off-the-shelf components. Integrating these components could involve significant effort.

Another assumption was that technologies would be transferred from original incubators all the way through deployment in the hands of warfighters in theater. ROS-M modules would have a correspondingly broad range of maturity levels, from very immature technologies to those with formal evaluation by the U.S. Army Test and Evaluation Command. To raise the baseline of ROS-M capabilities and improve compliance with DoD requirements, they recommended articulating this spectrum of technological maturity. Inclusivity in the ROS-M requirements needs to be balanced with stringency; otherwise, they will lose their impact.

For more on ROS-M and TARDEC, visit dtic.mil/ndia/2016GRCCE/Saowski.pdf.

RESEARCHER
SCOTT McMILLAN

Setting the Standard for Big Learning Evaluation

In 2016, the SEI launched the Big Learning Benchmarks project, a big data project that seeks to establish the standard for evaluating large-scale machine learning ("big learning") platforms. These platforms are crucial to a variety of government tasks that employ huge data sets.

As big data grows ever bigger and collection speeds become faster, large-scale machine learning is necessary for analyzing and deriving meaningful information from that data. However, in spite of a great deal of research and development of scalable machine learning platforms, there is little consistency on how these platforms are evaluated. Data sets, applications, and metrics are often chosen on a case-by-case basis, so it is difficult to standardize or replicate results. So, every time the government needs to evaluate one of these platforms, it must start from scratch. By developing big learning benchmarks, the SEI intends to bring a standard approach to how the performance of big learning platforms is measured and reported.

The first step for the Big Learning Benchmarks project was to design, acquire, and configure a capable compute cluster for researching big learning. To get the cluster up and running, the SEI's Scott McMillan and his project team collaborated with Garth Gibson and Eric Xing of Carnegie Mellon University's (CMU) Parallel Data Lab (PDL). The cluster went online on July 31, 2016.

Housed in the Data Center Observatory on the CMU campus, the distributed cluster has massive storage and computing power. "We had four things on our list for this cluster: lots of compute nodes including GPU capability, large amounts of memory, high network bandwidth, and massive amounts of storage," said McMillan. The cluster delivers more than 400 TB of storage and 42 compute nodes, each with a 16-core processor, its own graphics processing unit (GPU), and 64 GB of RAM—and a 40 GB Ethernet connection between the nodes themselves and between the compute nodes and the storage. To put these numbers in perspective, consider that 400 TB of storage could hold approximately 124 million photos or 12 million minutes of video.

The cluster is currently being used by SEI researchers and Carnegie Mellon's Big Learning Group faculty members and graduate students to perform research into large-scale machine learning algorithms like deep neural nets for image and video classification, capabilities McMillan points out have important government applications. "The government has huge data sets—images they need to classify and graphs they need to analyze," says McMillan. "Optimizing tasks like these—whether for speed or for how many resources they use—can enable government organizations to get meaningful information from their data in a timely and cost-effective way."

Next steps for the project are to use the cluster to run multiple machine learning applications on public or artificial data sets, including video data sets, that represent challenges experienced by government stakeholders. The data set could be as large as hundreds of terabytes and could have millions of features.

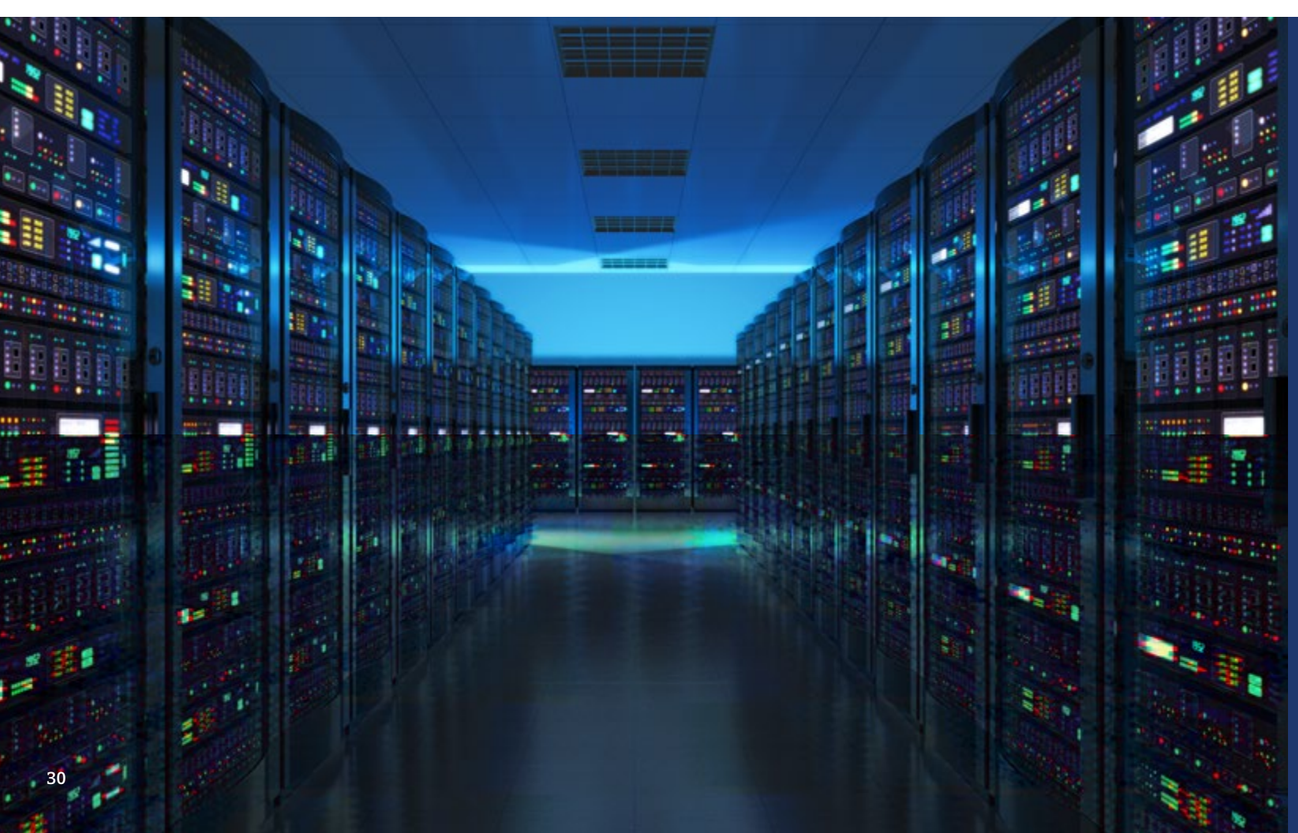

Optimizing tasks … can enable government organizations to get meaningful information from their data in a timely and cost-effective way."

— SCOTT McMILLAN, SEI EMERGING TECHNOLOGY CENTER


RESEARCHERS
DAVID TOBAR, DAVID ULICNE

Improving Cybersecurity and Resilience at the United States Postal Service (USPS)

In 2014, the United States Postal Service (USPS) experienced a cyber attack that compromised the personally identifiable information of more than 800,000 employees and over 2 million customers. The USPS recognized the need to improve, and it reached out to the SEI to help bolster its cybersecurity posture and operational resilience. Its ultimate goal was to protect the critical capabilities and assets of USPS and enhance its ability to continue business operations under degraded conditions.

The SEI collaborated with the USPS Corporate Information Security Office (CISO) to help develop and implement a cybersecurity strategy that integrated numerous recommendations into strategic improvement initiatives. The SEI's Risk and Resilience research team developed metrics, based on the CERT Resilience Management Model (CERT-RMM), to track USPS progress. "The measurement activities of the CERT-RMM are an essential element in executing the strategy," said David Tobar of the CERT Division's Cyber Risk Management team. "SEI researchers used the CERT-RMM to assess USPS performance in important areas—including risk management, system assessment and authorization, software development, incident management, and policy development—and made recommendations based on those assessments."

"To better train USPS CISO employees on cybersecurity, we teamed with USPS CISO management to create a new training program called the CISO Academy," said the SEI's David Ulicne, technical lead on the project. "This groundbreaking approach to cybersecurity workforce development offers a 12-week curriculum, including tracks for program managers and technical staff as well as courses delivered through the SEI CERT Division's Simulation, Training, and Exercise Platform (STEPfwd) and the Federal Virtual Training Environment (FedVTE)."

The SEI is seeking to continue research into cyber workforce development by assisting and measuring the impact of how other organizations strengthen their cybersecurity teams and culture.

To learn more about CERT-RMM, visit cert.org/resilience/products-services/cert-rmm.

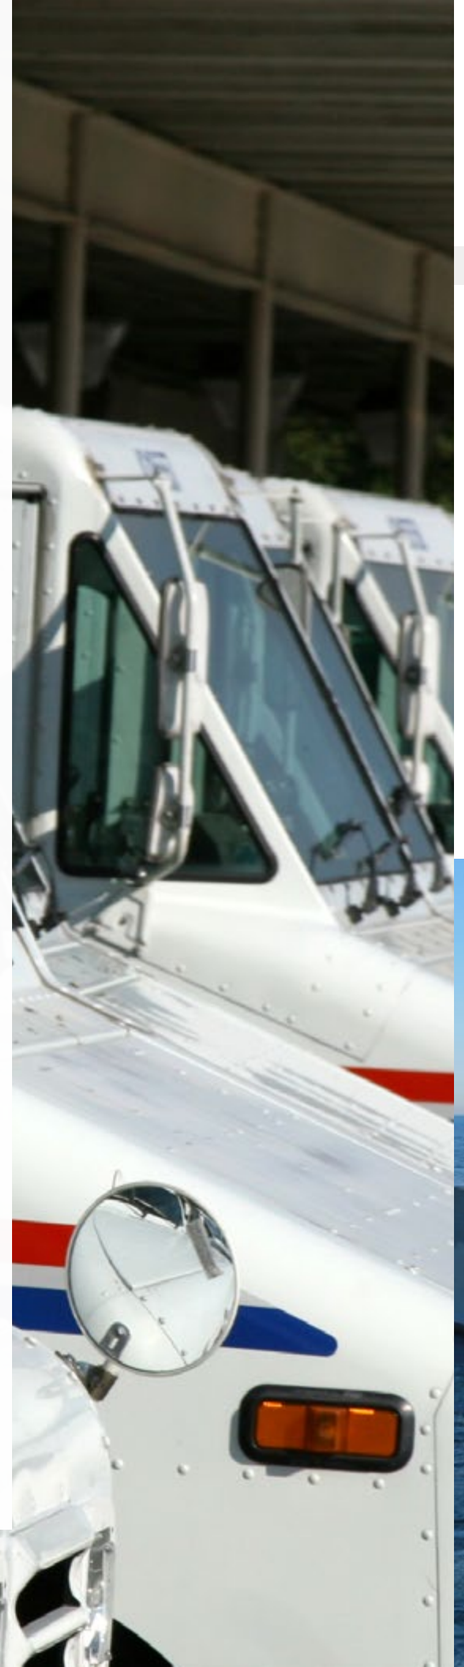



RESEARCHERS
STEVE BECK, TED MARZ, JOHN ROBERT

Providing Time-Critical Software Analysis

In 2016, the SEI continued to expand its capabilities and offerings in what the Institute calls "independent software analysis," or ISA. John Robert, a technical director in the SEI's Software Solutions Division, explained how ISA works.

"The SEI can provide time-critical independent software analysis to urgently help programs gain insight into and address software issues," Robert said. "We're able to apply our experience and software expertise to help Department of Defense programs identify and resolve immediate problems."

Recent examples of ISA in use include

- providing the Office of Naval Intelligence (ONI) assistance in assuring that its acquisition satisfies architectural goals and best practices. According to the SEI's Ted Marz, this involved analyzing the software, documentation, and architecture of the proposed system and providing recommendations to assure its sustainability.
- applying a "wider objective lens" at the enterprise level for the addition of a new warehousing system by the U.S. Marine Corps to help ensure that it would integrate within a larger system of interoperable systems. According to the SEI's Steve Beck, the SEI "focused the Marines on envisioning the desired end state for the entire enterprise," noting that "no single system that belongs to an enterprise should be examined solely as an island."

Photo: U.S. Navy

RESEARCHERS
CHRIS MAY, JONATHAN FREDERICK

SEI STEM Initiative: High School Students Get Crash Course in Cyber-Kinetic Tactical Operations

Forward Operating Base Kyle buzzed with activity. Under a canopy of camouflage netting, Combat Mission team 227 from the U.S. National Cyber Mission Force worked to coordinate the efforts of Task Force 44, a Navy SEAL unit deployed to the small island of Paraiso in the Indian Ocean. Loudspeakers squawked with the urgent chatter of the task force. Monitors displayed real-time views from SEAL body cams and the surveillance drone hovering above the island, which had succumbed to a well-coordinated band of pirates. Task Force 44 had been deployed to the island to rescue a prominent journalist taken hostage by the pirates.

If some members of team 227 sometimes whooped with youthful abandon, they could be forgiven. They were, after all, local high school students from the Pittsburgh region having a great time. The 75 students had gathered at the SEI CERT Division's Distributed Learning Center for a three-day program on cyber techniques used in mission support operations. The program culminated in a rescue mission executed in a sophisticated training environment created by the SEI to support Department of Defense training initiatives.

"This is a first for us," said Chris May, technical director of the CERT Division's Cyber Workforce Development team. "In this exercise, we connected cyber and kinetic missions in real time." May explained that his team created the cyber component of the exercise using the SEI's STEPfwd training environment. May's team then integrated a virtual kinetic battle simulator produced by a third-party vendor. This integration resulted in a rich training environment that extended from the cyber realm into the realm of events taking place in a simulated combat environment.

"Our goal for these kinds of events is to address a gap in teen education and help develop and inspire the next generation of elite cybersecurity professionals," said the SEI's Jonathan Frederick, who helped design the exercise.

For more information on the work of the SEI CERT Division's Cyber Workforce Development team and the STEPfwd training environment, visit cert.org/cyber-workforce-development.

"Our goal for these kinds of events is to address a gap in teen education and help develop and inspire the next generation of elite cybersecurity professionals."

— JONATHAN FREDERICK, SEI CERT DIVISION

RESEARCHER
WILL KLIEBER

Using Automated Code Repair to Reduce DoD Software Vulnerabilities

Department of Defense (DoD) codebases contain billions of lines of C code containing an unknown number of errors. These errors can lead to security vulnerabilities. Static code analysis tools can help find errors, but these tools are typically used late in the development process and generate a huge number of error warnings. Even after excluding false positives, the volume of actual coding errors can overwhelm developers. Consequently, only a small percentage of the vulnerabilities identified are eliminated. But help is on the way: recent work by the SEI aims to make code repair at the DoD much more manageable.

Research by the SEI's Will Klieber and the Secure Coding team has revealed that many security-related software bugs follow common patterns that can be used to automatically repair code and eliminate security vulnerabilities. Building on this research, the SEI CERT Division's Secure Coding team is developing source code transformation tools that automatically fix vulnerabilities caused by violations of rules defined in the CERT Secure Coding Standards. These tools have the promise to identify and repair errors much faster than the manual review of thousands of alerts, and do so at much lower cost.

The SEI's work on automated repair is based on three premises:

1. Many security bugs follow common patterns.
2. By recognizing a pattern, the developer's intention can be inferred (the Secure Coding team calls this *inferred specification*).
3. The code can then be automatically repaired to satisfy the inferred specification.

Take the case of memory allocation. Security bugs often result when the allocation exhibits a pattern such as "p = malloc(*n* * sizeof(T))," where *n* is attacker controlled. If *n* is very large, integer overflow occurs, and too little memory gets allocated. This condition sets the stage for a buffer overflow. In this example, the inferred specification is "Try to allocate enough memory to hold *n* objects of type T," and the repair is to insert code to check if overflow occurs and, when it does, to simulate malloc returning NULL.

The goal of the automated code repair project is to enable development teams to mitigate all unhandled violations by reducing the number of rule violations that require manual inspection by two orders of magnitude—from thousands to tens.

Secure Coding team members are engaging DoD Software Assurance Community of Practice members on the project, and the SEI has engaged with the U. S. Army Communications-Electronics Research, Development and Engineering Center (CERDEC) to provide feedback and technology transition. Specifically, in FY17, CERDEC will evaluate the integer-overflow repair tool on DoD codebases.

Project members also collaborated with Carnegie Mellon University professors Claire Le Goues and Christian Kästner. Le Goues is a leading researcher in the use of genetic programming for automated code repair. In this approach, there are three inputs: a defective program, test cases that exercise a fault in the program, and test cases that exercise normal program behavior. Genetic programming uses computational analogs of biological mutation and crossover to generate new program variants and to search for a variant that produces the desired result for all test cases.

Kästner has pioneered work on symbolically analyzing code under all possible build configurations (combinations of compile-time options). Large projects often have tens or even hundreds of compile-time options, and in the worst case, the number of possible build configurations grows exponentially with the number of options. For example, a project with 20 compile-time options might have a million possible build configurations—way too many to analyze individually. This situation requires a symbolic approach. Most existing work on static analysis considers only one build configuration at a time; the Automated Code Repair team hopes to eventually develop a repair that works for all possible build configurations.

Automated code repair promises to greatly reduce the number of vulnerabilities in a codebase, freeing the organization to focus on fixing the remaining coding errors, developing secure code, and achieving the organization's software assurance goals.

To learn more about this work, visit resources.sei.cmu.edu/library/asset-view.cfm?assetID=474244.

The goal of the automated code repair project is to enable development teams to mitigate all unhandled violations by reducing the number of rule violations that require manual inspection ***by two orders of magnitude.***

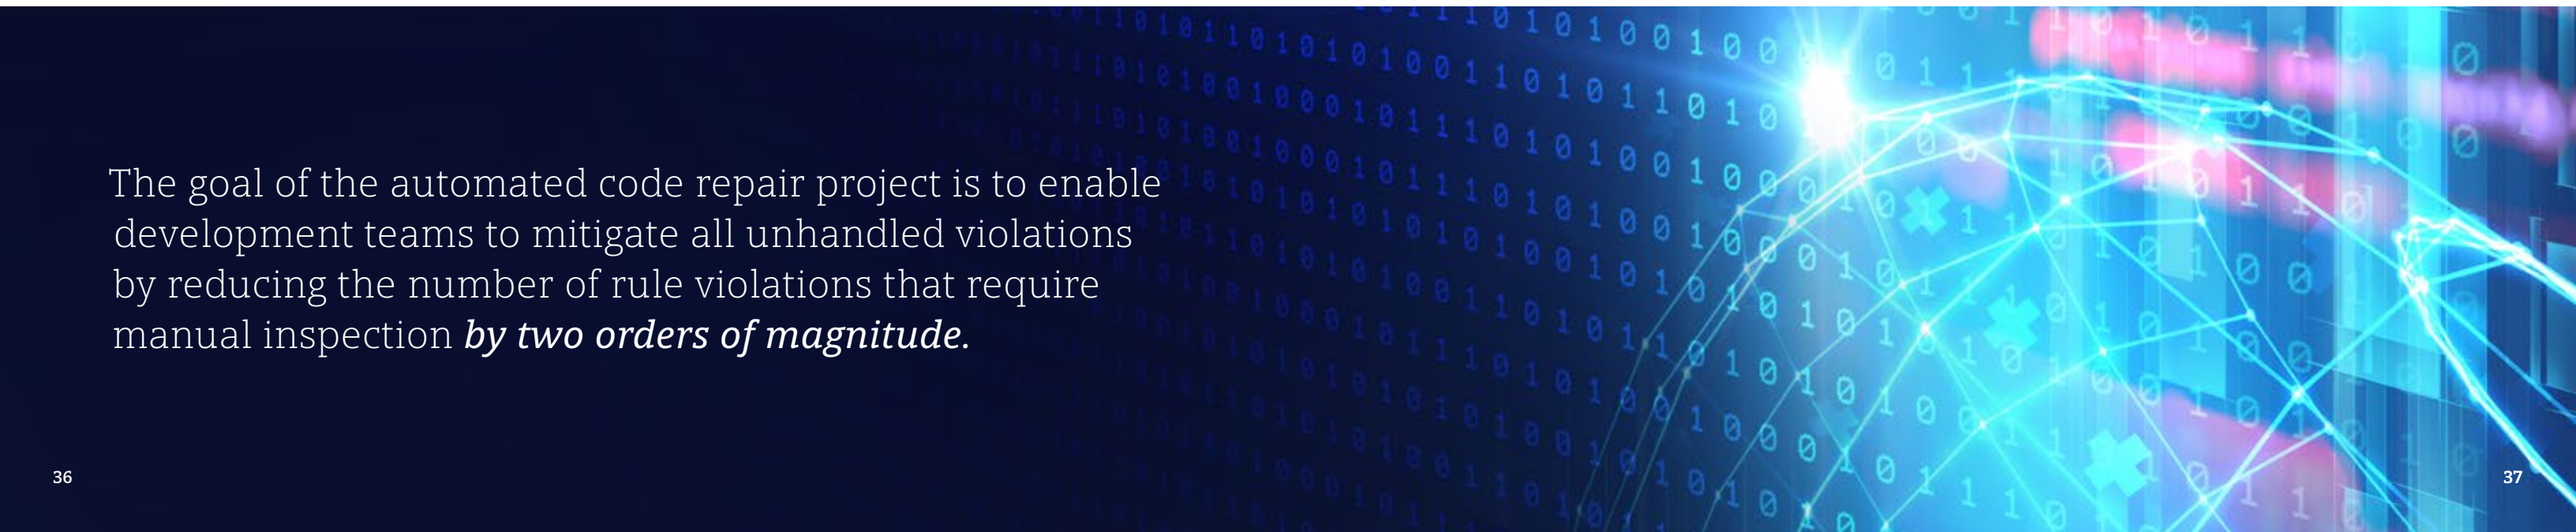

TRANSITION

The SEI accelerates the impact of software and cybersecurity improvements by working to promote adoption of improved capabilities by the defense industrial base and the wider software and cybersecurity communities. The SEI does this by creating standards, prototypes and tools, technical guidance, and platforms for knowledge and skill acquisition.

STANDARDS

The SEI develops standards that improve the software ecosystem on which the Department of Defense (DoD) relies. For instance, the CERT Secure Coding Initiative has been leading the community development of secure coding standards for common programming languages. Many of these proposed practices are in use by major participants in the supply chain for DoD software-reliant systems, including Cisco Systems and Oracle. The SEI has also worked to integrate several research technologies into the Architecture Analysis and Design Language standard, making it extensible and semantically well defined. Application of the standard promotes the virtual integration of system building and testing activities—an approach that supports DoD objectives of achieving integrated warfighting capabilities and delivering solutions sooner to warfighters.

PROTOTYPES

SEI researchers develop software prototypes that test proposed solutions, like the smartphone app developed in collaboration with the Carnegie Mellon University Human–Computer Interaction Institute. Called the Edge Mission-Oriented Tactical App Generator (eMONTAGE), this software program for mobile devices enables warfighters to mash data from multiple sources and view the results on a unified display—all without writing code. SEI researchers have demonstrated an eMONTAGE prototype at the U.S. Special Operations Command/Naval Postgraduate School (NPS) Tactical Network Testbed and at NPS's Joint Interagency Field Exploration (JIFX).

TOOLS

The SEI systematically builds software tools, especially those that address acute cybersecurity needs. Fuzz testers and debuggers developed by the SEI's CERT Division, for example, can position military software engineers to meet requirements outlined in the 2013 National Defense Authorization Act for software assurance testing. Other SEI tools facilitate security analysis in large networks, enable analysts to rapidly query large sets of data traffic volumes, process packet data into bidirectional flow records, and simplify the building of analysis environments.

TECHNICAL GUIDANCE, WORKFORCE DEVELOPMENT, AND KNOWLEDGE SHARING

The SEI shares the progress and results of its research through a host of media avenues, including

- technical reports, blog entries, webinars, and podcasts available on its websites
- articles in prestigious professional journals and publications geared to practitioners
- books in the SEI Series in Software Engineering published by Addison-Wesley

The books in the SEI Series often form the basis for education materials and training courses offered by the SEI and others. The SEI offers classroom and eLearning courses in software acquisition, network security, insider threat, software architecture, software product lines, software management, and other areas.

In 2012, the SEI introduced the CERT STEPfwd (Simulation, Training, and Exercise Platform) to help cybersecurity practitioners and their teams continually build knowledge, skills, and experience. In addition, SEI researchers collaborated with educators from around the United States to develop the first curriculum for software assurance, the Master of Software Assurance (MSwA).

The IEEE Computer Society and Association for Computing Machinery, as well as community leaders in curriculum development, formally recognized the MSwA Reference Curriculum as suitable for creating graduate programs or tracks in software assurance.

CMU LEADERSHIP

Subra Suresh
President, Carnegie Mellon University

Farnam Jahanian
Provost, Carnegie Mellon University

SEI EXECUTIVE LEADERSHIP

Left to Right:

Edward Deets, Director, Software Solutions Division; John Bramer, Chief of Staff; Mary Catherine Ward, Chief Strategy Officer; Robert Behler, Deputy Director and Chief Operating Officer; Paul D. Nielsen, Director and Chief Executive Officer; Peter Menniti, Chief Financial Officer; Jeff Boleng, Chief Technology Officer (Acting); David Thompson, Chief Information Officer; Matthew Gaston, Director, Emerging Technology Center

ORGANIZATION

SEI DIRECTOR'S OFFICE

Paul D. Nielsen
Director and Chief Executive Officer

Robert Behler
Deputy Director, Chief Operating Officer

Jeff Boleng
Chief Technology Officer (Acting)

SOFTWARE SOLUTIONS DIVISION

Edward Deets
Director

Anita Carleton
Deputy Director

David Zubrow
Chief Scientist

CERT DIVISION

Bill Wilson
Director (Acting)

Greg Shannon
Chief Scientist

Roman Danyliw
Chief Engineer

EMERGING TECHNOLOGY CENTER

Matthew Gaston
Director

Eric Werner
Deputy Director

OFFICES OF THE CHIEF OF STAFF/ CHIEF INFORMATION OFFICER

John Bramer
Chief of Staff

David Thompson
Chief Information Officer

FINANCIAL & BUSINESS SERVICES

Peter Menniti
Chief Financial Officer

STRATEGIC INITIATIVES

Mary Catherine Ward
Chief Strategy Officer

SEI LEGAL

Sandra Brown
SEI General Counsel

BOARD OF VISITORS

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

Barry W. Boehm
TRW Professor of Software Engineering, University of Southern California; Director, University of Southern California Center for Software Engineering

John M. Gilligan
President, Gilligan Group; former Senior Vice President and Director, Defense Sector of SRA International; former CIO for the Department of Energy

Alan J. McLaughlin
Chair, Board of Visitors; Consultant; Former Assistant Director, MIT Lincoln Laboratory

Gilbert F. Decker
Consultant; former President and CEO, Penn Central Federal Systems Company; former President and CEO of Acurex Corporation; former Assistant Secretary of the Army/Research, Development, and Acquisition

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

Donald Stitzenberg
President, CBA Associates; Trustee, Carnegie Mellon University; former Executive Director of Clinical Biostatistics at Merck; Member, New Jersey Bar Association

Philip Dowd
Private Investor; former Senior Vice President, SunGard Data Systems; Trustee, Carnegie Mellon University

Tom Love
Chief Executive Officer, Shoulders Corp; Founder of Object Technology Group within IBM Consulting

SEI STAFF

FULL-TIME & PART-TIME STAFF AND OTHER CONTRIBUTORS

Michael Abad-Santos
Lisa Abel
Joshua Acklin
Steve Ader
Laura Aguera
Cecilia Albert
Christopher Alberts
Michael Albrethsen
William Aldrich-Thorpe
Dennis Allen
Noelle Allon
Amanda Alvarez
Laura Anderson
William Anderson
Bjorn Andersson
Eileen Angulo
Christopher Antalek
John Antonucci
Luiz Antunes
Jeffrey Apolis
Michael Appel
Leena Arora
Eric Azebu
Felix Bachmann
David Bahm
Marie Baker
Karen Balistreri
Vincent Balistreri
Aaron Ballman
Jeffrey Balmert
Ronald Bandes
Michael Bandor
Hollen Barmer
Peter Barrett
Jeffrey Basista
Jason Batchelor
Mary Batchelor
Barbora Batokova
Daniel Bauer
Christopher Baum
Dwight Beaver
Stephen Beck
Robert Behler
David Belasco
Stephany Bellomo
Jonathan Bender
Brian Benestelli
Kristi Berneburg
James Besterci
Robert Beveridge
Philip Bianco
David Biber
Daniel Bidwa
Darlene Bigos
Tracy Bills
Robert Binder
Marla Blake
Stacie Blakley
Stephen Blanchette
Jeffrey Boleng
Elaine Bolster
Randall Bowser
Andrew Boyd
Diane Bradley
Ben Bradshaw
Eric Bram
John Bramer
Kara Branby
Erica Brandon
Pamela Brandon
Heidi Brayer
Scotty Brennan
Rex Brinker
Rita Briston
Nanette Brown
Rhonda Brown
Lisa Brownsword
Andrew Bunker
John Bush
Matthew Butkovic
Palma Buttles-Valdez
Anthony Calabrese
Rachel Callison
Sara Cammarata
Kimberley Campbell
Linda Campbell
Linda Canon
Peter Capell
Anita Carleton
Cassandra Carricato
William Casey
Harry Caskey
Kelly Cassidy
Thomas Castrodale
Anthony Cebzanov
Adam Cerini
Sagar Chaki
Mary Jo Chelosky
Timothy Chick
Leslie Chovan
Mary Beth Chrissis
Natalie Chronister
Jonathan Chu
Matthew Churilla
Jason Clark
Kathleen Clarke
William Claycomb
Matthew Coates
Cory Cohen
Julie Cohen
Sanford (Sholom) Cohen
Mary Lou Cole
Rebekah Coley
Matthew Collins
Anne Marie Connell
Carol Connelly
John Connelly
Robert Conway
Michael Cook
Stephen Cooney
Rebecca Cooper
Alexander Corn
Daniel Costa
Jennifer Cowley
Susan Cox
Randy Crawford
Rita Creel
Lucy Crocker
Larry Crowe
Stephanie Crowe
Natalie Cruz
Jerome Czerwinski
Rebecca D 'Acunto
Roman Danyliw
Rosemary Darr
Jeff Davenport
John Dayton
Dionisio de Niz
Daniel Decapria
Edward Deets
Grant Deffenbaugh
Julien Delange
Kareem Demian
Matthew Desantis
Edward Desautels
Aaron Detwiler
Jill Diorio
John Diricco
Robert Ditmore
Claire Dixon
Geoff Dobson
Patrick Donohoe
William Dormann
Audrey Dorofee
Margie Drazba
Elke Drennan
Michael Duda
Carly Duffy
Michael Duggan
Madelaine Dusseau
Ladonna Dutton
Karin Dwyer
Patrick Dwyer
Sean Easton
James Edmondson
Danielle Edwards
Eileen Eicheldinger
Robin Eisenhart
Linda Elmer
Harold Ennulat
Lover Epps
Neil Ernst
Jared Ettinger
Felicia Evans
Heather Evans
Michele Falce
Kimberly Farrah
Mariane Fazekas
Jeffrey Federoff
Peter Feiler
Eric Ferguson
Luis Figueroa
Donald Firesmith
Kodiak Firesmith
Lori Flynn
Justin Forbes
John Foreman
Kunta Fossett
Summer Fowler
Tracey Fox
Jonathan Frederick
David French
Jason Fricke
Michelle Fried
Michael Fritz
Brent Frye
Michael Gagliardi
Suzette Gambone
Douglas Gardner
Matthew Gaston
Linda Parker Gates
David Gearhart
Jeffrey Gennari
Martin Geshev
Cara Giannandrea
Travis Gibbons
Stephen Gifford
Lisa Gillenwater
Ryan Gindhart
Jamie Glenn
Walter Goss
Wassie Goushe
Bruce Grant
Michael Greenwood
Russell Griffin
Phillip Groce
Jacqueline Grubbs
Rajasekhar Gudapati
Rotem Guttman
David Guzik
Shannon Haas
Bart Hackemack
Nancy Hags
William Halpin
Jeffrey Hamed
Josh Hammerstein
Jeffery Hansen
Eric Hanson

Stephen Hardesty
Erin Harper
Jeffrey Havrilla
Jason Hawk
William Hayes
Matthew Heckathorn
Jessica Hedges
Stephanie Hedges
Steven Henderson
Sharon Henley
Christopher Herr
Kurt Hess
Charles Hines
Scott Hissam Barbara
J. Hoerr Bryon Holdt
Charles Holland
Andrew Hoover
Angela Horneman
Allen Householder
Joshua Howell Ryan Howley
John Huber
John Hudak Clifford Huff
Lyndsi Hughes
Jennifer Hykes Chris Inacio
Terry Ireland
James Ivers
Jerry Jackson
Vanessa Jackson
Michael Jacobs
Michael Jehn William Jones Jacob Joseph
Alejandro Jove Gavin
Jurecko Matthew Kaar

Stephen Kalinowski
Eliezer Kanal
Jennifer Kane
Rachel Kartch
Mark Kasunic
Harry Kaye
Tracey Kelly
Robert Kemerer
Brent Kennedy
Jennifer Kent
Carolyn Kernan
Kimberly King-Cortazzo
John Klein
Mark Klein
Stacy Klein
William Klieber
Dan Klinedinst
Georgeann Knorr
Andrew Kompanek
Michael Konrad
Kiriakos Kontostathis
Keith Korzec
Andrew Kotov
Paul Krystosek
Robert Kubiak
Amy Kunkle
Zachary Kurtz
David Kyle
Michael Lambert
Joel Land
Debra Lange
Mark Langston
Vincent Lapiana
Frank Latino
David Law
Alyssa Le Sage
Bernadette Ledwich
Ryan Lehman
Harry Levinson
Darrell Lewis

Grace Lewis
Alena Leybovich
Amy Leyland
Braden Licastro
Joshua Lindauer
Reed Little
Constance Loffreda-Mancinelli
Todd Loizes
James Lord
Melissa Ludwick
Richard Lynch
Rudolph Maceyko
Harold Major
Lisa Makowski
Arthur Manion
Jay Marchetti
Attilio Marini
Tamara Marshall-Keim
Theodore Marz
Lisa Masciantonio
Laura Mashione
Michael Massa
Roxanne Matthews
Jeffrey Mattson
Christopher May
Joseph Mayes
Michael McCord
Patricia McDonald
Roy McDonald
Michelle McGee
Shane McGraw
James McHale
David McIntire
Donna McIntyre
Donald McKeon
Janis McKinney
Bernadette McLaughlin
Michael McLendon
Joseph McLeod
Scott McMillan

Jason McNatt
Deborah McPherson
Ryan Meeuf
Andrew Mellinger
Peter Menniti
Thomas Merendino
Jennifer Mersich
Leigh Metcalf
Bryce Meyer
Toby Meyer
Bertram Meyers
Chase Midler
Matthew Milazzo
Amy Miller
Cassandra Miller
Christopher Miller
Gerald Miller
Suzanne Miller
Samantha Misurda
Soumyo Moitra
Elizabeth Monaco
Justen Monroe
Austin Montgomery
Andrew Moore
Jose Morales
Maria Morales
Damon Morda
Gabriel Moreno
John Morley
Edwin Morris
Mervyn Morris
Timothy Morrow
Anna Mosesso
Angela Mosqueda
Jamie Moyes
David Murphy
Paul Murray
Mark Musolino
Lynne Marie Naelitz
Cynthia Nesta

Gail Newton
William Nichols
William Nichols
Paul Nielsen
Crisanne Nolan
Robert Nord
William Novak
Marc Novakouski
Simon Novelly
Kevin Nowicki
Jasmine Oates
Nicholas O'Connor
Sharon Oliver
Kyle O'Meara
Luke Osterritter
Nancy Ott
James Over
Ipek Ozkaya
Mari Ann Palestra
Timothy Palko
Mark Palmquist
Amanda Parente
Allison Parshall
Carmal Payne
David Pekular
Kelwyn Pender
Brenda Penderville
Matthew Penna
Samuel Perl
Sharon Perry
Alexander Petrilli
Thomas Petrus
David Phillips
Kevin Pitstick
Patrick Place
Daniel Plakosh
Thomas Podnar
Shauna Policicchio
Mary Popeck
Jason Popowski

Douglass Post
Jerome Pottmeyer
Katherine Prevost
Andrea Prilla
Samuel Procter
Sean Provident
Michael Rattigan
Traci Radzyniak
Angela Raible
James Ralston
Donald Ranta
Frank Redner
Marc Reed
Sonia Reed
William Reed
Robert Reeder
Aaron Reffett
Colleen Regan
Nicholas Reimer
David Reinoehl
Janet Rex
Clifford Rhoades
Louis Richards
Nathaniel Richmond
Michael Riley
Stacey Rizzo
John Robert
Lawrence Rogers
James Root
Robert Rosenstein
Sheila Rosenthal
Stephanie Rosenthal
Dominic Ross
Adam Rousseau
Bradley Rubbo
Daniel Ruef
Robin Ruefle
Brad Runyon
Kristopher Rush
Mary Lou Russo

Mary Lynn Russo
Charles Ryan
Samuel Salinas
Miranda Salva
Venkatavijaya Samanthapudi
Thomas Sammons
Geoffrey Sanders
Concetta Sapienza
Emily Sarneso
Vijay Sarvepalli
Jeff Savinda
Thomas Scanlon
Alfred Schenker
David Scherb
Robert Schiela
Andrew Schlackman
Steve Scholnick
Patricia Schreiber
James Schubert
Carol Schultz
Kenneth Schultz
Edward Schwartz
Giuseppe Sciulli
Tina Sciullo-Schade
Philip Scolieri
David Scott
Shirley Scott
William Scully
Johnathan Seaburn
Joseph Seibel
James Semler
Stephen Serafin
Gregory Seroka
Gregory Shannon
Sarah Sheard
David Shepard
Mark Sherman
Nataliya Shevchenko
Deana Shick
Timothy Shimeall

Linda Shooer
Sandra Shrum
Forrest Shull
George Silowash
Matthew Sisk
Michelle Slusser
James Smith
Holly Smith
Lenny Smith
William Snavely
Timur Snoke
Gabriel Somlo
Tara Sparacino
Debra Spear
James Spencer
Derrick Spooner
Bryan Stake
Lauren Stanko
Jonathan Steele
Lizann Stelmach
Katie Stewart
Robert Stoddard
John Stogoski
Edward Stoner
Kirk Striebich
Jeremy Strozer
Gregory Such
Siobhan Sullivan
David Svoboda
Seth Swinton
Rebecca Sylvester
Michael Szegedy
Lucille Tambellini
Joe Tammariello
Michael Theis
Marcia Theoret
Jeffrey Thieret
Kimberly Thiers
Alisa Thomas
Mark Thomas

William Thomas
David Thompson
David Tileston
David Tobar
Michele Tomasic
Barbara Tomchik
Carolyn Tomko
Helen Trautman
Matthew Trevors
Peter Troxell
Donovan Truitt
Randall Trzeciak
Laurie Tyzenhaus
David Ulicne
Jeanette Urbanek
Vijay Sai Vadlamudi
Justin Valdengo
Karen Van Buren
Christine VanTol
Nathan Vanhoudnos
Satya Venneti
Joseph Vessella
Aaron Volkmann
Alexander Volynkin
Robert Vrtis
Todd Waits
Kurt Wallnau
Cynthia Walpole
Rand Waltzman
Mary Catherine Ward
David Warren
Mary Warren
Trina Washington
Garret Wassermann
Charles Weinstock
Adam Welle
Eric Werner
James Wessel
Austin Whisnant
Barbara White

Amanda Wiehagen
Akia Williams
Keegan Williams
Pamela Williams
Mary Wilson
William Wilson
Nicholas Winski
Robert Wojcik
Brandon Wolfe
Carol Woody
Jonathan Woytek
Lutz Wrage
Michael Wright
Eileen Wrubel
Joseph Yankel
Charles Yarbrough
John Yarger
Hasan Yasar
Lisa Renee Young
Cat Zaccardi
Mark Zajicek
Gene Zambrano
Marianne Zebrowski
John Zekany
Christine Zobel
David Zubrow
Allison Zust

OTHER CONTRIBUTORS

Brian Averi
Travis Breaux
Sandra Brown
Grady Campbell
David Carney
Anne Carrie
Peter Chen
Larry Druffel
Robert Ellison
Robert Ferguson
David Garlan
Charles Garvey
Thomas Glazier
David Gluch
John Goodenough
Raghav Goyal
Charles Hammons
Daniel Jack
Frederick Kazman
Mary Ann Lapham
Claire le Goues
Carrie Lee
Sung Lee
Shen Li
Michael Maass
John McGregor
Joseph McManus
Nancy Mead
Julia Mullaney
Kenneth Nidiffer
Linda Northrop
Cathy O'domes
Jeffrey Pinckard
Raghvinder Sangwan
Rosario Scalise
William Scherlis
Emilie Schlauch
Bradley Schmerl
Doug Schmidt

Lui Sha
Douglas Sicker
Lynda Silipo
Charles Wallen

AFFILIATES

Yoshihiro Akiyama
Diego Vallespir

Copyright

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM17-0079

Credits

MANAGER, COMMUNICATION SERVICES

William Thomas

MANAGER, CORPORATE & TECHNICAL COMMUNICATIONS

Janet Rex

MANAGER, PUBLIC RELATIONS

Richard Lynch

EDITOR-IN-CHIEF

Ed Desautels

EDITORIAL

Hollen Barmer
Heidi Brayer
Ed Desautels
Claire Dixon
Tamara Marshall-Keim
Gerald Miller
Nancy Ott
Sandra Shrum
Barbara White

DESIGN

Christopher Baum

ILLUSTRATION

Kurt Hess

DIGITAL PRODUCTION

Mike Duda

PHOTOGRAPHY

David Biber

Tim Kaulen, Photography and Graphic Services, Mellon Institute

Nancy Ott

WEB DESIGN

Barbora Batokova

Christopher Baum

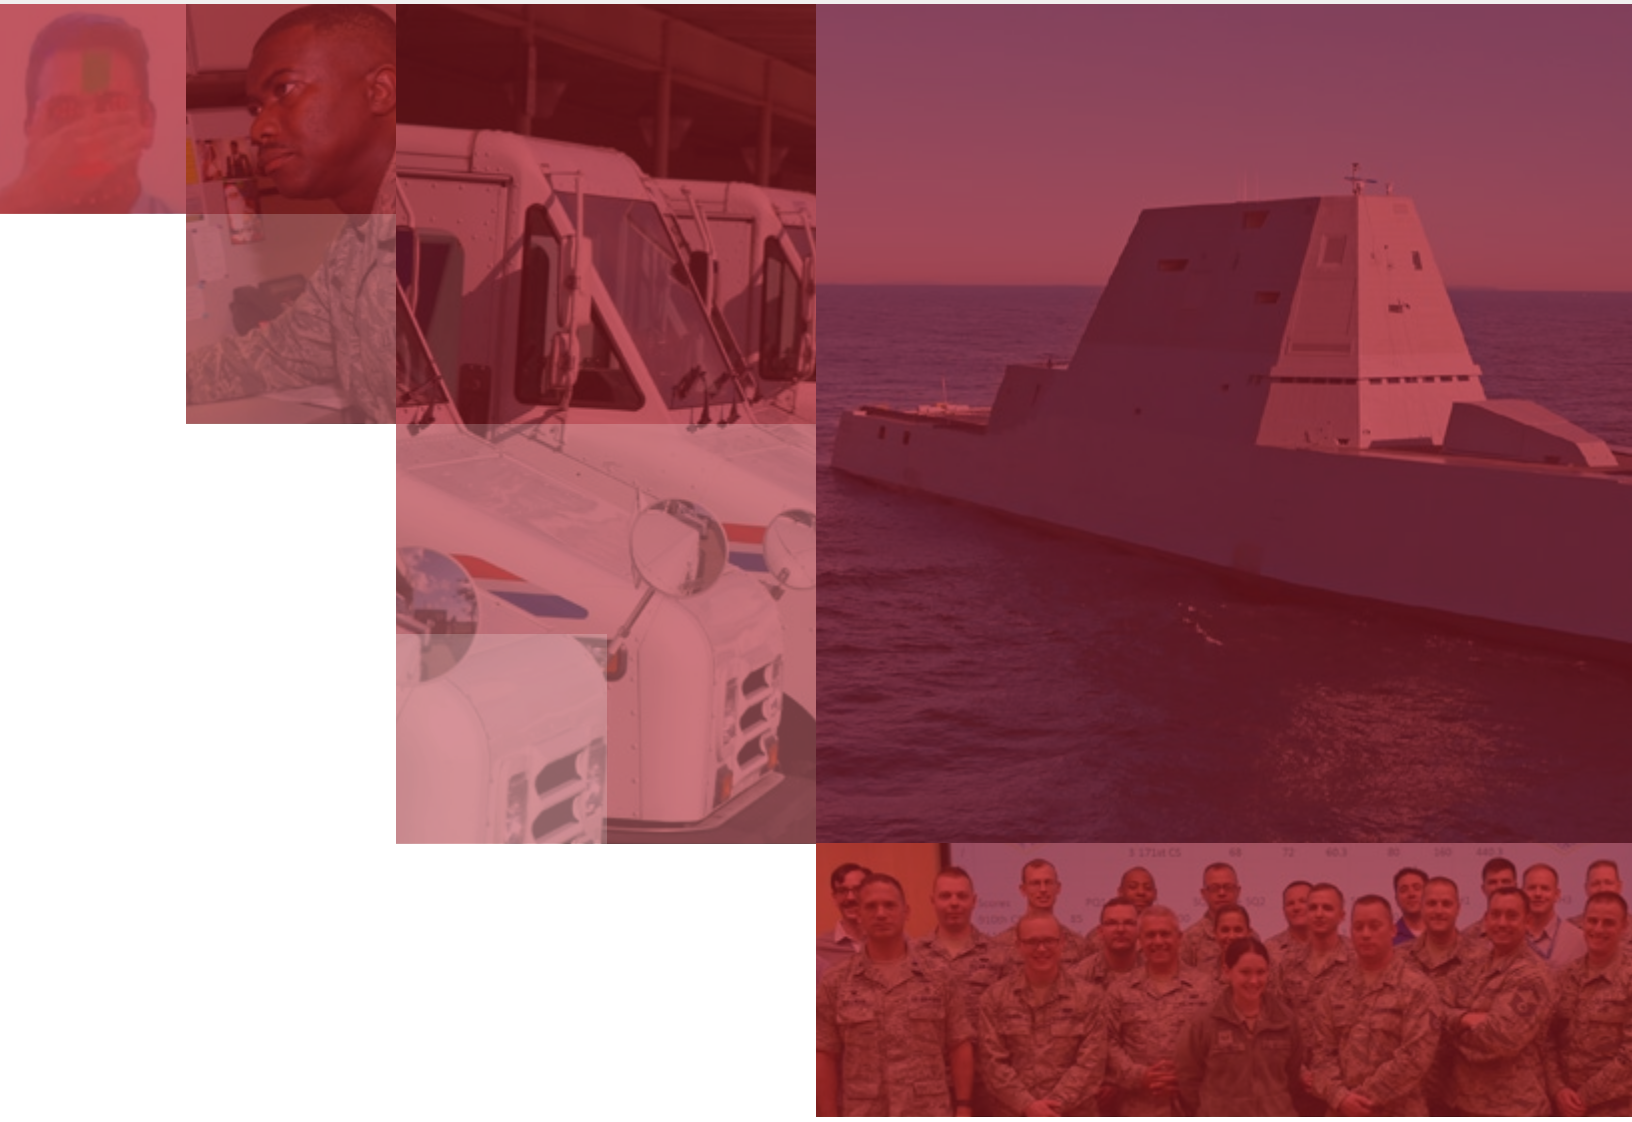

SEI PITTSBURGH, PA

4500 Fifth Avenue
Pittsburgh, PA 15213-2612

SEI WASHINGTON, DC

Suite 200
4301 Wilson Boulevard
Arlington, VA 22203

SEI LOS ANGELES, CA

2401 East El Segundo Boulevard
El Segundo, CA 90245