2010
CERT® RESEARCH
REPORT

CERT

Software Engineering Institute
Carnegie Mellon

The CERT® Program's research mission is to leverage our operational experience and access to state-of-the-art research to fill gaps for our customers, collaborate with the cyber security research community, and provide longer-range perspectives on cyber-security research, technology, and policy.

Table of Contents

CERT Research Vision

Today we live in a world in which the threat of cyber attacks is ever-growing, and where threats from unknown sources are dynamic and constantly changing. It is seldom that a week goes by when articles on cyber security are not prominent in technical publications and popular media. The mission of CERT®, part of the Software Engineering Institute (SEI), a federally funded research and development center at Carnegie Mellon University (CMU), is to enable the survival of critical networked systems against contemporary threats and attacks by removing technical, maturity, information, and capacity barriers in cyber security and incident response.

Our stakeholders include the U.S. Department of Defense, the Department of Homeland Security, law enforcement, intelligence community, other U.S. federal agencies, state and local governments, and other operators of infrastructures critical to the national defense, cyber security, and the national economy; the providers of information communications technologies (ICTs) and services that support these system and network operators; the software development community; and computer security incident response teams with national responsibilities.

The overall goal of our program is improved practices and technologies that are widely understood and routinely used to protect, detect, and respond to attacks, accidents, and failures on networked systems. Better informed, trained, and equipped people will produce better systems that will be better managed to reduce operational risk and the impact of cyber attacks.

Our research strategy has been to build and maintain a technical center of excellence and innovation that uses its operational experience and expertise to look across the entire software life cycle (from requirements through development, deployment, operations, maintenance, and forensics) to

- identify new technologies, development practices, and management practices that would significantly improve networked systems security and enterprise resiliency
- mature these technologies and practices
- apply these technologies to meet the needs of the program's stakeholders
- transition these technologies into widespread use

This report attests to the successful execution of the CERT research mission in fiscal year 2010.

Greg Shannon
Chief Scientist, CERT
Software Engineering Institute, Carnegie Mellon University

2

Executive Summary

We are pleased to present the 2010 CERT Research Report. This year's report has some new features, including "greener" paper and a new look to parallel our exciting work environment. The CERT in the News section briefly summarizes publicity about CERT research in newspapers, magazines, and journals. We also describe sponsored workshops.

We strongly believe that security must be addressed at every phase of software and system development and operation, and in a variety of situations. This is reflected in the depth and breadth of research projects and other innovative activities at CERT, organized in this report by focus area:

- Malicious Code Research and Development – The malicious code research at CERT focuses on understanding malicious code itself and on leveraging the vast amount of data in the Malicious Code team's Artifact Catalog. This research supports efforts to mitigate the problem of malicious code and the development of tools and methods that streamline malicious code analysis while reducing associated costs.

- Secure Coding – The Secure Coding team works with software developers and software development organizations to reduce vulnerabilities resulting from coding errors before they are deployed. We strive to identify common programming errors that lead to software vulnerabilities, establish secure coding standards, educate software developers, and advance the state of the practice in secure coding.

- Software Security Assurance – The Software Security Assurance (SSA) team focuses on addressing security in the early life-cycle phases of acquisition and software development.

- Digital Intelligence and Investigation Directorate – Grounded in years of research and real-world experience, the Digital Intelligence and Investigation Directorate (DIID) focuses on "gap areas" not addressed by commercial tools or standard techniques.

- Incident Response – CERT supports the development of an international response community by helping organizations build incident response capability and by developing a common infrastructure of policies, practices, and technologies to facilitate rapid identification and resolution of threats.

- Insider Threat – Our insider threat research focuses on both technical and behavioral aspects of actual compromises. We produce models, reports, training, and tools to raise awareness of the risks of insider threat and to help identify the factors influencing an insider's decision to act, indicators and precursors of malicious acts, and countermeasures to improve organizational survivability and resilience.

- Network Situational Awareness – The Network Situational Awareness (NetSA) team's research focuses on monitoring large networks and analyzing bulk data collections, with the ultimate goal of detecting malicious activity. In NetSA's major research projects, analysts develop approaches to automated analysis, measure network phenomena, and determine the return on security investments in attack mitigation.

- Resilience Modeling and Analysis– Since 2001, CERT has been working in security process improvement and operational resilience management and engineering. Beginning with the OCTAVE® Method, CERT has been researching and developing tools, techniques, and methods that help organizations manage operational risk and improve operational resilience.

- Workforce Development – CERT makes training as realistic and accessible as possible through a web-based simulation, training, and evaluation platform called XNET.

Abstracts are provided here for the major research projects in each focus area, with more detailed project descriptions appearing in the report. Each focus area also highlights newer or smaller projects. An additional section provides a biography and short list of publications, presentations, and technical leadership activities for each author.

2010 Research Report Abstracts

Secure Coding Initiative

By analyzing thousands of vulnerability reports, CERT has observed that most vulnerabilities stem from a relatively small number of common programming errors. Software developers can take practical steps to eliminate known code-related vulnerabilities by identifying insecure coding practices and developing secure alternatives. In many cases, this can lead to language enhancements that are adopted by international standards development organizations. Compiler vendors can then implement these changes, and software developers can use the improved compilers in the supply chain and incorporate them into software-intensive systems. To enhance secure coding practices, CERT is working on the following projects: secure coding standards; standards development; automated analysis tools; secure compiler extensions; application conformance testing; SEI Team Software Process[SM] (TSP[SM]-Secure); and books, courses, training, and education.

Building Assured Systems Framework (BASF)

A framework is needed to organize research and practice areas focused on building assured systems. The Building Assured Systems Framework (BASF) addresses the customer and researcher challenges of selecting security methods and research approaches for building assured systems. After reviewing existing life-cycle process models, security models, and security research frameworks, the principal investigators used the Master of Software Assurance Reference Curriculum knowledge areas as the BASF. We mapped all major CERT research areas to the BASF, proving that the BASF is useful for organizing building assured systems research. We also performed a gap analysis to identify promising CERT research areas. The BASF is a useful structure for planning and communicating about CERT research. The BASF will also be useful to CERT sponsors to track current research and development efforts in building assured systems.

Supply Chain Assurance

Software has become a central element for all major organizations, and supply chains are an essential part of most deployed software systems, which often use commercial components or outsourced development and system integration. Reducing software vulnerabilities requires improving the assurance of the supply chain. Some software practices can reduce vulnerabilities, but they are not widely practiced. This research considers how to incrementally improve software supply chain assurance.

Measuring Software Security Assurance

For several years, the software engineering community has been working to identify practices aimed at developing more secure software. Although some foundational work has been performed, efforts to measure software security assurance have yet to materialize in any substantive fashion. As a result, decision makers (e.g., development program and project managers, acquisition program offices) lack confidence in the security characteristics of their software infrastructures. The CERT® Program at Carnegie Mellon University's Software Engineering Institute (SEI) has chartered the Security Measurement and Analysis (SMA) Project to advance the state-of-the-practice in security measurement and analysis. The SMA Project is researching and developing frameworks, methods, and tools for measuring and monitoring the security characteristics of large-scale, networked, software-reliant systems across the life cycle and supply chain. The project team has developed the SEI Integrated Measurement and Analysis Framework (IMAF) as well as several methods for implementing the framework. IMAF employs systemic analysis to integrate subjective and objective data from a variety of sources, including targeted analysis, status reporting, and measurement, to provide decision makers with a consolidated view of the security posture of large-scale, networked systems.

Security Requirements Engineering

Through the Security Quality Requirements Engineering (SQUARE) project, CERT researchers have developed an end-to-end process and associated tool for security requirements engineering, along with training courses to help organizations build security into the early stages of the production life cycle. The SQUARE methodology consists of nine steps that generate a final deliverable of categorized and prioritized security requirements. The process has been baselined and transitioned into practice. CERT researchers are currently leading the development of SQUARE for acquisition (A-SQUARE) and for privacy (P-SQUARE).

With SQUARE and other security requirements engineering processes as a base, we want to enable development and acquisition organizations to address security requirements engineering early in the life cycle. Ultimately we would like to see security requirements engineering reflected in standard development and acquisition processes, as well as in international standards. In our work with client organizations, we hope to capture additional data to support the business case for security requirements engineering.

Using Trusted Hardware as a Foundation for Cyber Security

The Software Engineering Institute's (SEI) stakeholders are increasingly operating in what we term extreme adversarial environments. Operating securely in such environments requires that we understand the range of characteristics of those environments and assess feasible approaches to maintain security and survivability properties under the severe conditions these environments impose. In FY 2010, this research evaluated the promise and limitations of using trusted hardware as a foundation for achieving demonstrably high assurance of end-to-end security properties of applications executing in extreme adversarial environments.

Automated support for trust is essential. Trust technology has the potential to overcome the limitations of existing approaches for achieving survivability, security, and dependability. The Trusted Platform Module and other hardware-based trust mechanisms are a step in the right direction but inadequate in current practice. Our FY 2010 research laid the groundwork for future work that will explore and exploit the concepts of trust and trustworthiness and provide a scientific basis for understanding the relationships among hardware, software, security, and trust. Our ultimate goal is to provide the capability to build and operate critical automated systems that will behave in a sufficiently trustworthy manner to consistently fulfill their missions, even when these systems are built and operated in extreme adversarial environments.

Malware Family Analysis: Correlating Static Features and Dynamic Characteristics on Large-Scale Projects

While static analysis and dynamic analysis are widely considered best practices for developing knowledge to characterize malware, approaches that combined the two modes of analysis have seen limited success in scaling to large problems involving thousands of malware programs.

We report on a study that merges the two modes of analysis in a scalable manner for a set of 61,889 *Zeus/Zbot* malware binaries. Specifically we are able to correlate functional behavior to static positions, thus producing a behavior map for static features in the malware corpus.

The result is a promising new development in techniques for knowledge and data discovery that improves the capability to identify and characterize malware families for large-scale projects.

Beyond Section Hashing

Assessing the similarity of executable files, especially malicious code (malware), remains a challenging problem for both human analysts and automated systems. Similarity estimates lower the cost of expensive human analysis by allowing analysts to quickly recognize files or data they have previously encountered. Our recent research into malware similarity has concentrated on decomposing formatted files into subcomponents and using hashing to observe identical matches of the subcomponents. While this enables human analysts to observe related files, using automated processes to capture these relationships has been stymied by confusing and sometimes incorrect results. This report analyzes the effectiveness of section hashing, specifically as used in two derived techniques: composite section hashing and section clustering. We analyze executable files collected in the CERT Artifact Catalog and observe that approximately 3 percent of all composite section hashes produce approximately 42 percent of all file MD5s. We observe connected components in the bipartite file-section graph and describe potential relationships in approximately 79 percent of the executable files in the Artifact Catalog. We describe problems with using these techniques to assert relationships, including irrelevant hash collisions and massively connected clusters of files. Finally, we describe future work, including content-derived weighting of section hashes.

Assessing the Benefits and Effectiveness of Network Sensors

This article reports on a model that evaluates how sensors on an organizational information network can benefit security-related decisions. Sensors are crucial for network security. They monitor traffic and help detect potential intrusions and attacks. Good decisions on how to acquire and deploy sensors are key to ensuring the best network security with constrained resources. Effective decisions require the knowledge of the benefits from a sensor if placed at a location to both justify the investment and efficiently allocate the sensors across competing locations. This project developed an integrated concept of the effectiveness of sensors from a managerial perspective and a model that will help sponsors of CERT make decisions regarding the acquisition and deployment of network sensors for security. No suitable model previously existed. The model was developed to address the concerns of the Department of Defense and other U.S. government organizations. The model included, among other factors, the value of sensitive information. It is a practical and actionable approach that can be integrated with strategic decisions about network security. The research also identified the data that should be collected for such decision making. Further collaborative work with CERT sponsors is planned.

How Fast Do Computers Move?

Watchdog organizations like Shadowserver track the growth and persistence of botnets around the globe. From this outside perspective, malware-infected computers are usually only visible through the IP addresses that they use to communicate over the internet. But the relationship between machines and internet protocol (IP) addresses is like the relationship between people and street addresses; an address can represent a single home, a high-rise apartment building, or a time-share. In this research we use geo-location of IP addresses from the Waledac botnet to examine the speed at which malware-infected computers appear to "travel" around IP space, with the ultimate goal of adjusting botnet population estimates for the inflation accrued by IP-mobile devices and network IP address allocation practices.

Measuring Operational Resilience: Moving from Uncertainty to Justified Confidence

Operational resilience, as defined by the CERT® Resilience Management Model (CERT®-RMM), addresses the ability of an organization to protect and sustain the resilience of mission-critical assets and services. An operationally resilient service is a service that can meet its mission under times of disruption or stress and can return to normal operations when the disruption or stress is eliminated. A service is not resilient if it cannot return to normal after being disrupted, even if it can temporarily withstand adverse circumstances. CERT-RMM incorporates the disciplines of security and business continuity as well as aspects of IT operations.

As organizations strive to improve their ability to effectively manage operational resilience, it is essential that they have an approach for determining which measures best inform the extent to which they are meeting their performance objectives. Resilience measurement and analysis (RMA) research builds upon foundational research methods in software and security measurement. Research results to date include deriving six high-level, business-driven objectives for operational resilience based on CERT-RMM, defining a candidate measurement template, and deriving a number of example measures using the template. Future work includes validating this approach and candidate measures based on CERT-RMM appraisals, surveys, reviews, and engagements with users of CERT-RMM.

Software Assurance Curriculum Project

Seeing the need for advanced education in software assurance and education for acquirers of assured software, the Department of Homeland Security (DHS) directed the SEI in 2009 to develop a curriculum for a Master of Software Assurance (MSwA) degree program. Substantial effort by a collaborative team of curriculum development and subject matter experts in CERT, Carnegie Mellon University, and other U.S. universities has resulted in a comprehensive body of foundational knowledge and course structure for graduate-level software assurance education—the MSwA 2010 Reference Curriculum, available at *http://www.cert.org/mswa/*. The curriculum is the first of its kind to be developed and has been recognized by the two leading computing professional societies: IEEE Computer Society and the Association for Computing Machinery.

In addition to the curriculum structure, CERT developed under the DHS directive course outlines, syllabi for nine MSwA core courses (eight lecture courses and a capstone project), a report detailing an appropriate undergraduate software assurance (SwA) concentration, a master list of references, and an initial set of course materials donated by educators that can be used in software assurance courses.

CERT in the News

CERT has far-reaching impact in the field of cyber security, and CERT researchers are often requested to provide expert opinions in the media. The following links to published articles are just a sample of media coverage that features CERT research and experts.

FBI Arrests, Charges Three Botnet Operators
PC Magazine, November 21, 2009

http://www.pcmag.com/article2/0,2817,2146395,00.asp

This article discusses how the government works with the CERT Coordination Center to identify individuals operating botnets. According to the article, "The government is working in conjunction with the CERT Coordination Center at Carnegie Mellon University as well as industry partners like Microsoft to uncover these [botnet operator] schemes. The effort has thus far uncovered more than 1 million computer IP addresses infected by botnets. The FBI and DOJ are working to contact those who have been affected by the scam."

Report: Dangers of Cyber Crime on the Rise
IT Business Edge, January 27, 2010

http://www.itbusinessedge.com/cm/blogs/poremba/report-dangers-of-cyber-crime-on-the-rise/?cs=39029

This blog post explores the results of the 2010 CSO Cyber Watch Survey. The article says that "…the results of the 2010 CSO Cyber Watch Survey, a cooperative effort between the U.S. Secret Service, Deloitte, the Carnegie Mellon Software Engineering Institute (CERT) and CSO Magazine, and a white paper from Deloitte's New Center for Security & Privacy Solutions, 'Cyber Crime: A Clear and Present Danger,' find that the cybercrime-fueled underground economy continues to breed a sophisticated arsenal of damaging tools and devices (malware, botnets, anonymizers) – and companies cannot keep pace or remain focused elsewhere."

Is Chasing Cybercrooks Worth It?
CNN.com, March 5, 2010

http://www.cnn.com/2010/TECH/03/05/cyberattack.prosecute/index.html?hpt=C2

CNN interviews Marty Lindner, CERT assurance manager working with Information Technology/ Security, about cybercriminals in this article that focuses on this concept: "This week's arrests of three men in connection with one of the world's largest computer-virus networks may seem like great news – perhaps even a sign authorities are starting to win the war against cyberthieves. But the real situation is more complicated." The article goes on to explain, "The people who actually write these programs – the keys to cybercrime – are almost impossible to catch and prosecute," said Marty Lindner, principal engineer with Carnegie Mellon University's Computer Emergency Response Team…Lindner said it's unclear if the authors of malicious code are doing anything illegal." Lindner shares this quote in the article: "The U.S. doesn't have jurisdiction on the [entire] planet Earth, so even if you can identify the author [of the malicious program], that doesn't give us the legal authority to get him, one, and two, it's not clear he's committing a crime…It's not illegal to write bad software. It's illegal to use it."

Detecting Malicious Insiders Before Data Breaches Damage Your Business
eWeek, April 6, 2010

http://www.eweek.com/c/a/Security/Detecting-Malicious-Insiders-Before-Data-Breaches-Damage-Your-Business-510011/

This article profiles the work of Dawn Cappelli, technical manager of Enterprise Threat and Vulnerability Management and the CERT Insider Threat Center. According to the article, "As intriguing as the idea of a mysterious cyber-criminal hacking his way into a corporate network sounds, the majority of data breaches are the work of insiders…Dawn Cappelli knows that well. As the technical lead of CERT's insider threat research at Carnegie Mellon's Software Engineering Institute, she has analyzed 450 cases of malicious insiders in search of common threads that businesses can use to develop security strategies." Cappelli is quoted in the article as saying, "If you look at these crimes, you can't detect it with technology alone because a system administrator is going to use his authorized access…and you can't tell if it's malicious or not unless you know when to look…Unless you put a strategy together that looks at the people, the process and the technology, it's going to be very hard to detect these things."

CERT Releases Fuzz Testing Framework
IEEE Computer Society, May 28, 2010

http://www.computer.org/portal/web/news/home/-/blogs/cert-releases-fuzz-testing-framework

This blog post, which announces the release of the fuzz testing framework, says that "…Will Dormann writes on the CERT blog that the Basic Fuzzing Framework (BFF) is a simplified approach to automated dumb fuzzing, a technique popularized for use in security research by hackers. With BFF anyone can easily test software application using the tool. Fuzzing finds vulnerabilities in software by sending random input to an application. This is the second such tool released by CERT."

The Barbarians Are Already Inside the Gates: Mitigating Insider Threats
TechRepublic, August 2, 2010

http://www.techrepublic.com/blog/security/the-barbarians-are-already-inside-the-gates-mitigating-insider-threats/4148

This IT security blog interviews Dawn Cappelli, technical manager of Enterprise Threat and Vulnerability Management and The CERT Insider Threat Center, about insider attacks on data: "According to Dawn Cappelli, technical manager for the threat and incident management division of the Software Engineering Institute CERT program, '…insider attacks continue to be seen as a bigger problem than anything that might come from the outside' (Brenner, 2010, p. 2). Dollars spent to prevent breaches and other information asset related incidents caused by employees may have a larger ROI than those spent on traditional controls."

CERT Team Examines Health-Care Security Risks
SEI, September 21, 2010

http://www.sei.cmu.edu/newsitems/healthcare_threats.cfm

This article outlines the different types of information security breaches that health-care organizations face. According to the article, "Far too often, the threats come from within an organization, according to Randy Trzeciak, a senior member of the technical staff at CERT and the insider threat team lead. Since its inception, the CERT insider threat team has studied internal malicious activity against organizations. The team has created a database of more than 400 insider threat cases that team members use to analyze potential indicators of malicious activity…With each passing year, medical facilities and hospitals rely more heavily on IT systems. This reliance makes them vulnerable to IT sabotage, which is often perpetrated at the hands of an employee. 'Employees who conduct IT sabotage are disgruntled. There is a perceived injustice on the part of the individual. Often, there has been a negative workplace event that caused the person to become disgruntled and want to exact revenge against the organization,' Trzeciak said."

Special Project: Recommending Cyber Security Research Topics

In the spring of 2010, the Director for Information Systems and Cyber Security in the Office of the Director, Defense Research and Engineering (ODDRE), in the Office of the Secretary of Defense (OSD) asked the SEI to provide input to ODDRE on cyber security research. While CERT continually seeks insight from its stakeholders and the research community, this request prompted a unique, in-depth study reaching beyond SEI and CMU on research challenges related to the development and operation of secured information systems.

CERT, in collaboration with CMU's CyLab and Institute for Software Research (ISR), identified key research areas that have the following benefits for the DoD:

- high payoff—will result in significant improvement in the mid to long term that will justify the DoD's investment
- doable—can be described with attainable and manageable goals and objectives
- substantive—is focused around hard problems
- disruptive—has potential for cutting-edge change in securing critical information systems

The team from CERT, CyLab, and ISR wanted to provide OSD with a broad view of potential areas meriting DoD attention. They reviewed the past 10 years of recommendations from security-related national studies. That review of relevant material provided perspective and prepared the team with background information to conduct interviews with cyber security thought leaders.

The team interviewed 28 cyber security professionals from industry, academia, and government. During the interviews, each of these thought leaders enthusiastically shared their visions on research needs and opportunities with the team.

After reviewing existing recommendations and conducting in-depth interviews, the team followed the suggestion of one of the interviewees to develop a vision of where we thought the state of practice should be in 10 years. The team developed that vision of desired attributes that include the following key concepts:

- Demonstrate safe and secure operations in a recognized malicious environment.
- Enable effective adoptability of secure and survivable systems.
- Develop the principles and practice supporting the science and engineering of secure systems. This effort must address and include architectures, design and development, software and systems construction, usability, policies, administration, and operations.
- Evaluate the effectiveness and results from security investments to enhance the value associated with security.

Based on this vision, the team's objective was to identify key research areas of focus particularly applicable to the DoD. Faced with the task of selecting key recommendations from a rich landscape of possible research areas, the team purposely kept the recommendations to a manageable set that could translate to actionable programs. After much debate and discussion, the team developed a set of six recommended areas of focus along with associated research projects, any one of which would contribute to the DoD. The recommended research areas of focus follow:

- Architect secure systems. Understanding how to encapsulate applications in protected, trusted environments is a challenge. Research is needed in how to compose trustworthy systems from secured components. Ideally architectures will be developed that will support secure operations in the presence of malware while at the same time providing effective countermeasures for insider threats.

- Protect the network fabric. Protecting the capabilities of the DoD and national networks will require a complete understanding of how to defend against large-scale attacks, gain and sustain situational awareness of the network(s), and restore networked services after a successful attack. Being able to correlate network behavior with network access may be a lucrative area of research with near-term payback. Providing reliable and secure mobile networks should also be an area of concern to the DoD.

- Develop secure software and systems. The software development life cycle must be adapted to deliver security and survivability. This will require improving programming practices and languages. It will also require improving static and dynamic analysis tools to identify potential vulnerabilities. Research is needed to design software structures for assuring security properties ranging from software architectures appropriate for secure and/or trusted systems to patterns for secure design. Research that supports composing trusted systems from diversely sourced components and encapsulating unvalidated components will contribute to supporting security in the DoD supply chain.

- Create usable security and resilient systems. The team identified that one of the common reasons that security is weak is the lack of human usable security interfaces. Research is needed to identify the criteria and evaluation mechanisms for effective and usable human interfaces. Developing new human interface concepts that support security administration is a worthy goal. Ideally, researchers can develop techniques to map security policy to system configuration to make the implementation of security safeguards more practical. Research in this area should include techniques for rapid recovery of system components to maintain or restore systems security as well as techniques for management and operations to detect and counter emergent threats.

- Enhance digital forensics. The team found that forensic analysis can be the key to mitigating the impact of repeat attacks. Ideally, future operating systems and networks would have embedded forensic support to provide real-time monitoring, reactive response based on rapid analysis, and forensic analysis. Research to support determining attack vectors, both the origin and identity of attackers, would increase possible responses to include mitigation, isolation, and retribution. In addition, with the wealth of evidence possible for collection, research is needed to support massive data retrieval and analysis.

- Support offensive operations. Lessons can be and should be learned from defensive operations that are applicable to offensive operations. Likewise, defensive countermeasures may be derived from offensive possibilities. Tools for active and adaptive offensive and defensive operations of attack, attribution, and retaliation are needed, along with an operational doctrine to support and guide the cyber war fighters.

The team also identified a crosscutting theme that influenced all the areas of focus. That theme was enhancing evaluation. The team recognized that it was imperative to address security metrics as well as effective measurement collection and use to support the need for continued investment in security. Those metrics and measurement techniques will rely on evidence-based analysis and will span across such critical areas as validation of components in a supply chain to expected and proper operation of executable code. Developing the measures will support assessment of security and survivability risks as well as evaluation of the quality of protection required.

The team presented these recommendations to ODDRE in April 2010. At that presentation, the team emphasized that the recommendations were not the research agenda for the SEI, CyLab, or ISR. Rather the recommendations were intended to be the unbiased view from a range of experts in cyber security on research areas of focus that should be of interest to the DoD. The set of recommendations were also presented to the U.S. Navy chief information officer, the U.S. Coast Guard, the SEI's Board of Visitors, and others in 2010. While the recommendations are not a set of actionable items, they have served to inform the research directions for CERT, CyLab, and ISR.

CRITICAL CODE

"Although it is not a principal focus of this report, cybersecurity is an unavoidable and critical dimension of software assurance. It is rarely possible to contemplate software assurance without also giving major attention to security considerations. This is particularly challenging because security, like assurance, must be addressed at every phase of development and the software lifecycle overall."

– Critical Code: Software Producibility for Defense

INSIDER tHREAt

SECURE CODING

SOFtWARE SECURItY ASSURANCE

A Time-Tested Leader in Cybersecurity

Critical Code

In 2010 the National Academy of Sciences published the extensive report *Critical Code: Software Producibility for Defense* that explores Department of Defense (DoD) needs and priorities for software research and suggests a research agenda. While the work of CERT is much broader than the scope of the Critical Code report, this report has meaningful overlap with CERT research and provides perspective into the research needs of the DoD. It is also significant as it is the cornerstone of the SEI research strategy.

Because software is essential to the U.S. government's military operations, there is significant value in research to sustain or enhance the DoD's software capability. The Critical Code report identifies seven technology areas where research would contribute to the DoD's software capability.

- Architecture modeling and architectural analysis. Research is needed to facilitate improvements in DoD's ability to manage system design, evaluation, development, and evolution at the architectural level to improve software producibility.

- Assurance: validation, verification, and analysis of design and code. Research on a number of assurance-related capabilities could greatly enhance the DoD's ability to develop highly capable, highly assured systems.

- Process support and economic models for assurance and adaptability. Research goals for this area include enhancing process support for assured software development and addressing supply chain challenges and opportunities.

- Requirements. Because requirements engineering is an ongoing activity throughout development, more expressive models and supporting tools for both functional and quality attributes and improved support for traceability and early validation are needed.

- Language, modeling, coding, and tools. Research on programming languages and associated capabilities would have a considerable influence on architecture, assurance, process, and measurement.

- Cyber-physical systems. There needs to be accelerated development of new conventional architectures for control systems and improved architectures for a wide range of embedded applications.

- Human-system integration. Research is needed to develop engineering and design practices that account for the ways in which humans integrate into systems as participants.

Several of these areas directly overlap with major research areas at CERT, confirming the importance and relevance of the role of CERT research in the security of software and systems that support the DoD's mission.

The Critical Code report notes that early engineering choices strongly influence feasibility of achieving high assurance and recommends expanding research focus on assurance-related software engineering technologies and practices. One CERT research project is Security Quality Requirements Engineering (SQUARE), a nine-step process to help organizations build security into the early stages of the production life cycle. Read more about SQUARE on page 45.

The CERT supply chain assurance research is relevant to the Critical Code discussion on complexity and supply chains. The changing character of the architecture and supply structure for software systems, enabled by advances in underlying software technologies, particularly related to languages, tools, and runtime architectures, introduce more complex architectures and supply chains. Read how CERT researchers are addressing these challenges on page 39.

The CERT secure coding research is another area of overlap. Critical Code discusses methods that prevent the introduction of defects or find them as soon as possible after they are introduced. The report recommends adopting secure coding practices as a preventive method. Read about secure coding on page 26.

Insider threat is a fascinating and vibrant area of CERT research and is related to the Critical Code human-system integration section. CERT experts are regularly tapped for expertise in the professional media; see the CERT in the News section for examples. Learn about insider threat research projects starting on page 15.

While CERT research extends beyond the research agenda proposed in Critical Code, the report is an important measure of the overall software research needs for the DoD and reinforces the significance of the research CERT is doing in key areas defined in the report.

ALL OF OUR WORK IN THE
CERT INSIDER THREAT CENTER
IS GROUNDED IN REALITY. WE ASK, WHAT IS
REALLY HAPPENING OUT THERE?
-Dawn Cappelli

insider THREAT

MOST ORGANIZATIONS ARE PREPARED
TO DEFEND AGAINST ATTACKS COMING
FROM OUTSIDE THEIR BUILDING OR THEIR
NETWORK. BUT MANY ARE NOT PREPARED
TO DEFEND AGAINST ATTACKS FROM
INSIDE THEIR ORGANIZATION.
-RANDY TRZECIAK

Insider Threat Overview

According to Randy Trzeciak, technical lead of the Insider Threat Outreach and Transition Team, "Most organizations are prepared to defend against attacks coming from outside their building or their network. But many are not prepared to defend against attacks from inside their organization." For the past decade, the CERT® Insider Threat Center has been gathering empirical data on malicious insider activity to guide organizations to prepare for, prevent, detect, and respond to malicious insider activity. Initial research focused on understanding insider IT sabotage, fraud, theft of intellectual property, and national security espionage. In 2010 the Insider Threat Center created an Insider Threat Lab to focus on finding solutions to these pressing problems.

Guidelines and Standards

By researching patterns in malicious insider behavior, the Insider Threat Center develops recommendations for organizational security. One key finding from 2010 research was that about 70 percent of theft of intellectual property by insiders occurs within 30 days of the insider's announcement of resignation. The team periodically captures these types of findings in the next version of The Common Sense Guide to Prevention and Detection of Insider Threats, now in its third edition. In 2010 the team began mapping these guidelines onto existing sets of best practices, such as standards by the National Institute of Standards and Technology (NIST). The Insider Threat Center is working to broaden the impact of its research by working with NIST, the Department of Defense (DoD), and another federal agency responsible for national security to mature the best practices into full NIST standards.

Insider Threat Database

It all starts with the data. "Our research addresses real-life problems that we hear from practitioners or leaders in government and industry," says Dawn Cappelli, technical manager of the Insider Threat Center. The Center bases all its work on cases of insider activity from the real world. All that data comes together in the Insider Threat database, which documents more than 550 insider threat cases. The database provides a rich source for empirical research on real cases of insider threat. Work began in 2010 on the Insider Threat Lab, where data can fuel experimental research using virtual systems to simulate insider attacks on networks.

Assessment

One of the goals of the Insider Threat Center is to help organizations assess their level of preparedness against insider threats. In 2010 the team enhanced their insider threat assessment methodology to measure insider threat preparedness. The work was sponsored by a federal agency responsible for national security, but assessments are available to any organization.

Insider threat is just one aspect of an enterprise-wide consideration of organizational resiliency. However, it involves a fundamental aspect of security often lost in the technology: people. The research of the Insider Threat Center integrates deep knowledge of information security with an objective, data-driven examination of the motivations and behavior patterns of malicious insiders, as well as organizational issues that may influence them.

Leveraging data to prevent, detect, and respond to insider threats can help small organizations, to be sure. But the team's engagement with the government also strengthens the protection of the nation's critical infrastructure. The Insider Threat Center also investigates the behavior patterns of malicious insiders in national security espionage cases. "We validate all of our findings against our database of actual cases," says Cappelli, "as well as first-hand experiences doing assessments of government and industry organizations."

Collaboration

To cover such broad domains, the Insider Threat Center collaborates with experts within the SEI and around the country. The Workforce Development team at the SEI has helped develop the XNET systems that are at the heart of the Insider Threat Lab simulations and exercises. The Software Engineering Measurement and Analysis initiative at the SEI collaborates on insider threat measurement. Outside of the SEI, the team works with behavioral scientists to paint the larger picture of the nontechnical with the technical. They also use the system dynamics methodology to develop comprehensive models of the insider threat problem alongside experts in the intelligence community, the DoD, and academia.

Getting the Word Out

Measuring how much the Insider Threat Center's research has mitigated the insider threat problem can be difficult. As the research increases the community's awareness of the problem, the community reports more cases of insider threats. Still, mitigation and awareness go hand in hand. So the Insider Threat Center leverages its empirical research into awareness-raising products like the annual Cybercrime Watch Survey, which the team has played a major role in developing along with the U.S. Secret Service and *CSO Magazine*.

The Insider Threat Center also conducts workshops that teach attendees how to develop an effective, comprehensive insider threat monitoring strategy. In 2010, the Center held 11 workshops in Arlington, Baltimore, St. Louis, and Dallas.

High-profile insider crimes heighten awareness of specific areas of concern. However, it is important that organizations take a proactive approach to insider threat mitigation and ensure their protection against all types of insider attacks, not just the ones that happen to be in the news. Cappelli says insider threat goes beyond the headlines. "Our research addresses real-life problems that we hear from practitioners or leaders in government and industry," she says. "In addition, we validate all of our findings against our database of actual cases, as well as first hand experiences doing assessments of government and industry organizations. All of our work in the CERT Insider Threat Center is grounded in reality. We ask, what is really happening out there?"

Insider Threat Vulnerability Assessment Measurement

Principal Investigators: Randy Trzeciak, Joji Montelibano, Andrew Moore, and Dave Zubrow

The CERT Insider Threat Vulnerability Assessment (ITA) enables organizations to gain a better understanding of insider threat and an enhanced ability to assess and manage associated risk. It merges technical, organization, personnel, business security, and process issues into a single, actionable framework. The instrument is structured to encompass all stakeholders in the fight against insider threat. The assessment's results

- enable organizations to gain a better understanding of their vulnerability to insider threat and an enhanced ability to assess and manage associated risks
- benefit all individuals involved in the insider threat vulnerability assessment process: information technology, human resources, physical security, data and business process "owners," legal and contracting personnel, and all levels of organization management.

To create the assessment instrument, hundreds of actual incidents of malicious insider activity were analyzed to determine what the key issues of concern were, both technical and non-technical. This work leverages the foundational Insider Threat Study that CERT conducted jointly with the U.S. Secret Service, e.g., [1, 2], and considers more recent modeling and analysis, e.g., [3, 4], in combination with the identification of insider threat best practices [5].

Thousands of issues of concern were used to create the assessment, including

- individual's actions
- organization's actions (or failure to take action in response to an individual's action, event, or condition)
- events and/or conditions that might have influenced an insider's decision to act maliciously

To structure the assessment and ensure all areas are consistently addressed, these issues were compiled into a series of workbooks: Information Techology, Software Engineering, Human Resources, Physical Security, Legal / Contracts, and Data Owners. These workbooks provide the foundation for the substantiation of each insider threat capability. They facilitate the interviews with subject matter experts, provide indicators to be validated when observing processes in operation, and direct the review of organizational practices, policies, and procedures in an attempt to determine institutionalization of directives.

While the assessment provides an organization a snapshot in time of their ability to address the threats posed by malicious insiders, we want the assessment results to be measurable so organizations can track changes in their capabilities related to insider threat over time.

The goal of any measurement activity is to be both accurate and consistent. The ITA methodology developed scoring criteria both at the capability level (superclass) and the indicator level (subclass) to reflect the true state of the organization being assessed. The scoring must be the same in the light of the same evidence regardless of who is doing the scoring or when. Two aspects of consistency are repeatability, the degree to which the same results are obtained and when repeated measurements are made by the same individual under the same conditions, and reproducibility, the degree to which different individuals or teams obtain the same results when they score the same evidence.

The ITA achieves repeatability, intra-rater reliability, by defining scoring criteria such that the same analyst inspecting the same object using the same procedure under the same environmental and time conditions should come up with the same result. Empirically, this is assessed by having individuals score the same evidence or evidence that is deemed equivalent multiple times. The ITA achieves reproducibility, inter-rater reliability, by defining scoring criteria so that when different analysts inspect the same object they come up with the same result.

A lack of consistency in repeatability or reproducibility means that the score is driven by differences in the measurement method (e.g., the teams and individuals) rather than the evidence and state of the organization. This might be thought of as a low signal to noise ratio. The goal of the guidance for scoring is to minimize the noise in the system and enhance its ability to accurately capture the signal.

The goal of ITA scoring is to measure an organization's preparedness level with respect to the ability to prevent, detect, and respond to the issues of concern included in the ITA workbooks. Each ITA capability and indicator are scored on a scale of 1 to 4, with the following qualifications:

- 4 – maximum countermeasures in place to address issue of concern; ability to prevent, detect, and respond to the issue of concern

17

- 3 – adequate countermeasures in place; ability to detect and respond

- 2 – minimal countermeasures in place; ability to detect and respond (to a lesser degree)

- 1 – no countermeasures in place; failure or fatal flaw in the ability to detect or respond

The scoring levels are equated to a sieve or filter. The evidence for a capability is initially compared to that for level 4. If the criteria are not satisfied, the capability evidence is compared to that for level 3. The evidence continues to be compared with the criteria for each level until it hits a level where it does satisfy all of the criteria.

The goal of the ITA methodology is to provide clearly written criteria, with objective, observable indicators, so that insider threat capabilities can be quantitatively assessed. This enables organizations to measure their preparedness level against insider threat vulnerabilities exploited in other organizations and provides them with the information they need to develop a plan of action to increase their ability to prevent, detect and respond to insider threats.

References

[1] Kowalski, E.F., M.M. Keeney, D.M. Cappelli, A.P. Moore, "Insider Threat Study: Illicit Cyber Activity in the Information Technology and Telecommunications Sector" Joint SEI and U.S. Secret Service Report, January 2008. http://www.cert.org/archive/pdf/insiderthreat_it2008.pdf

[2] Kowalski, E.F., T. Conway, S. Keverline, M. Williams, D. McCauley, D.M. Cappelli, B.W. Willke, A.P. Moore, "Insider Threat Study: Illicit Cyber Activity in the Government Sector," Joint SEI and U.S. Secret Service Report, January 2008. http://www.cert.org/archive/pdf/insiderthreat_gov2008.pdf

[3] Moore, A.P., D.M. Cappelli, R.F. Trzeciak, "The 'Big Picture' of Insider IT Sabotage Across U.S. Critical Infrastructures," in Insider Attack and Cyber Security: Beyond the Hacker, eds. Stolfo, S.J., et. al., Springer Science + Business Media, LLC, 2008. Also published in SEI Technical Report - CMU/SEI-2008-TR-009. http://www.cert.org/archive/pdf/08tr009.pdf

[4] Weiland, R.M., Moore, A.P., Cappelli, D.M., Trzeciak, R.F. Spooner, D., "Spotlight On: Insider Threat from Trusted Business Partners," Joint CyLab (CMU) and CERT (SEI), February 2010. http://www.cert.org/archive/pdf/TrustedBusinessPartners0210.pdf

[5] Cappelli, D.M., Moore, A.P., Trzeciak, R.F. and Shimeall, T.J., "Common Sense Guide to Prevention and Detection of Insider Threats," Joint CyLab (CMU) and CERT (SEI), 3rd Edition, September 2008 (updated from July 2006 and April 2005). http://www.cert.org/archive/pdf/CSG-V3.pdf
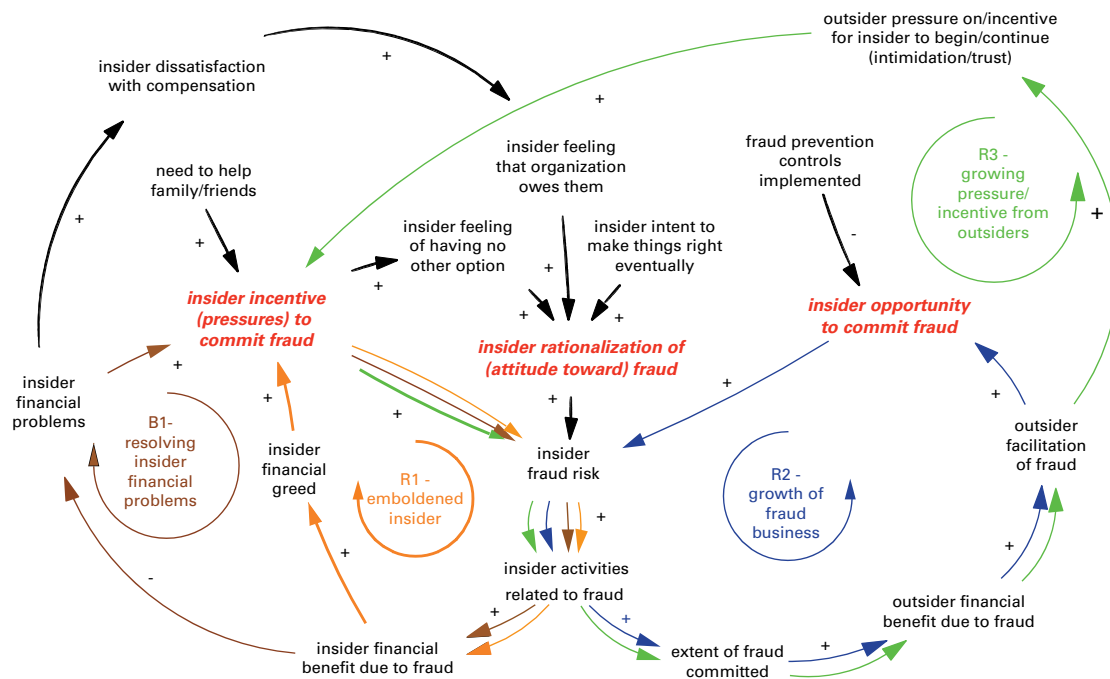
Modeling and Analysis of Insider Fraud

Principal Investigators:
Andrew Moore, Adam Cummings, and Derrick Spooner

A survey of 40 organizations in the Banking and Finance Sector showed that half were losing at least four to five percent of their total revenue to insider crime.[1] Eighty percent reported that the problem had grown worse since the economic downturn. Unfortunately, technical controls for these types of crimes are difficult to test in an operational environment because of the unpredictability of crime occurrence and the potential disruption to operations if the controls are ineffective. This project models crimes of insider fraud based on an extensive database of insider crimes in the banking and finance community. Model execution permits conducting experiments to validate the effectiveness of controls to mitigate insider fraud risk in a safe and controlled virtual environment. The CERT Program has previously developed models for insider IT sabotage [1, 2] and insider theft of intellectual property [3] and actively collaborates with the larger research community in this area [4, 5].

We define insider fraud as a crime perpetrated by a current or former employee, contractor, or other business partner who used the organization's information technology for the unauthorized modification of, addition to, or deletion of an organization's data for personal gain, or the removal of information leading to identity theft. This does not include corporate fraud where employees are acting on behalf of, or in the interests of, the organization.

One of the most famous fraud-specific models is the Fraud Triangle, developed by the criminologist Donald Cressey in the early 1950s. The model evolved through his interviews with imprisoned bank embezzlers to include three dimensions: pressure (what causes a person to commit fraud, often stemming from a significant financial need or problem), opportunity (the ability to commit fraud, perhaps as a result of weak internal controls or poor management oversight), and rationalization (the process of overcoming any personal ethical hesitations to commit the fraud). The Fraud Triangle has gained widespread support, most prominently by the AICPA's (American Institute for CPAs) Statement of Auditing Standards.

CERT adopted the Fraud Triangle as the basis for modeling the primary patterns of insider fraud. A portion of our preliminary model of insider fraud is depicted below. In system dynamics models, signed arrows represent the system interactions, where the sign indicates how the variable at the arrow's source influences the variable at the arrow's target. A positive (+) influence indicates that the values of the variables move in the same direction, whereas a negative (-) influence indicates that they move in the opposite direction.

---

[1] See "Bankers Gone Bad: Financial Crisis Making the Threat Worse," Security Dark Reading,
http://www.darkreading.com/insiderthreat/security/government/showArticle.jhtml?articleID=220301087
by K.J. Higgins.

Significant feedback loops are indicated in the model by a loop symbol, a loop label, and a loop name in italics. System dynamics models identify two types of feedback loops: balancing and reinforcing. Balancing loops—indicated as numbered "B" loops —describe aspects of the system that oppose change, seeking to drive variables to some goal state. Reinforcing loops—indicated as numbered "R" loops —describe system aspects that tend to drive variable values consistently upward or downward. Complex behaviors emerge as a result of combining balancing and reinforcing feedback loops.

The vertices of the Fraud Triangle are shown as red variables in the middle portion of the model. The model is composed of one balancing feedback loop and three reinforcing feedback loops:

- B1 (brown): resolving insider financial problems – This loop was identified as the original financial motivation for why insiders started engaging in fraud. The feedback is balancing since the fraud helps to mitigate any financial problems.

- R1(orange): emboldened insider – While initial motivations may have been due to a problematic financial situation, fraud may take on a life of its own, stimulating insider greed. This action has a reinforcing nature as the insider becomes accustomed to the extra income, thereby requiring the insider to continue the malicious activity.

- R2 (blue): growth of fraud business – This feedback loop relates to the expansion of the insider fraud due to the growth of the outsider-facilitated fraud business. This is reinforcing in nature, due to the outsider's continued influence, thus expanding the insider's opportunity to participate in the fraud scheme.

- R3 (green) – In addition to providing increased opportunity, the growth of the fraud business will likely result in outsiders pressuring the insider to increase their participation in the crime, e.g., by providing more and more information or modifications supporting the fraud. While there are limits to growth of the fraud business, there will likely be a period of escalation.

The patterns of insider fraud described above were evident in our preliminary analysis of cases. They form strong hypotheses as we move forward continuing to analyze cases of insider threat in the Financial Services Sector, as described in another article in the Insider Threat research focus area. Future work will refine and extend this preliminary model based on this analysis.

References

[1] Cappelli, D. M., Desai, A. G., Moore, A. P., Shimeall, T. J., Weaver, E. A., & Willke, B. J. (July 2006). Management and Education of the Risk of Insider Threat (MERIT): Mitigating the Risk of Sabotage to Employers' Information, Systems, or Networks. Proceedings of the 24th International System Dynamics Conference. Nijmegen, Netherlands.

[2] Moore, A. P., Cappelli, D. M., & Trzeciak, R. F. (2008). The "Big Picture" of Insider IT Sabotage Across U.S. Critical Infrastructures (Vol. Insider Attack and Cyber Security: Beyond the Hacker). (S. Stolfo, S. M. Bellovin, S. Hershkop, A. Keromytis, S. Sinclair, & S. W. Smith, Eds.) New York, NY: Springer Science+Business Media, LLC.

[3] Moore, A.P., D.M. Cappelli, T. Caron, E. Shaw, R.F. Trzeciak, "Insider Theft of Intellectual Property for Business Advantage: A Preliminary Model," in Proc. Of the 1st International Workshop on Managing Insider Security Threats (MIST2009), Purdue University, West Lafayette, USA, June 16, 2009. *http://www.cert.org/insider_threat/docs/Insider_Theft_of_IP_Model_MIST09.pdf*

[4] Anderson, D. F., Cappelli, D. M., Gonzalez, J. J., Mojahedzadeh, M., Moore, A. P., Rich, E., et al. (July 2004). Preliminary System Dynamics Maps of the Insider Cyber-Threat Problem. Proceedings of the 22nd International Conference of the System Dynamics Society.

[5] Rich, E., Martinez-Moyano, I. J., Conrad, S., Cappelli, D. M., Moore, A. P., Shimeall, T. J., et al. (July 2005). Simulating Insider Cyber-Threat Risks: A Model-Based Case and a Case-Based Model. Proceedings of the 16th International Conference of the System Dynamics Society. Quebec City, Canada.

Insider Threat Lab

Principal Investigators:
Michael Hanley, Christopher King,
George J. Silowash, and Derrick Spooner
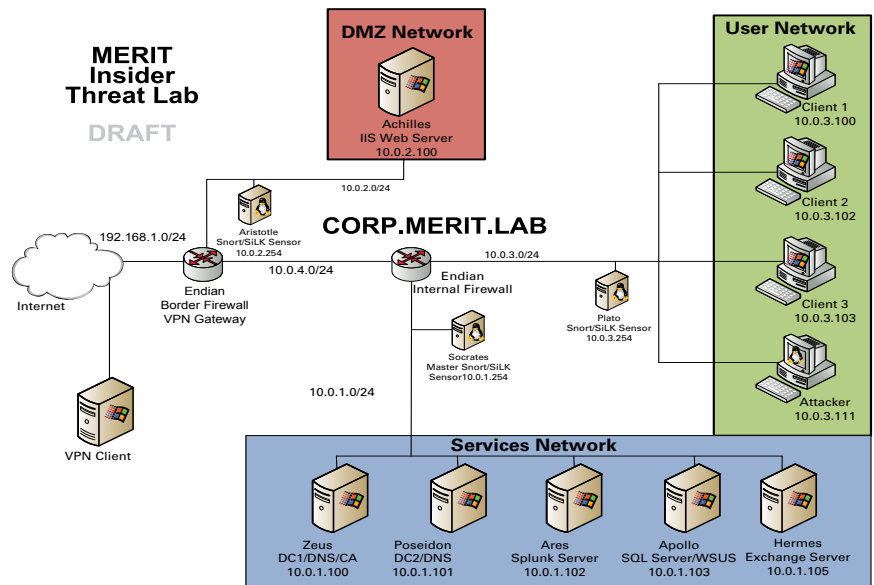
Since 2001, the Insider Threat Center at CERT has gathered over 550 cases of malicious insider activity. We have developed a coding and analysis methodology that allows the granular inspection of the various patterns of attack observed in each individual case. Understandably, analysis of these cases uncovers a large variety of organizational information architectures – however, we have identified commonalities in insider behavior and use of attack vectors that we have compiled into system dynamics models. These models simulate the events preceding and succeeding an insider act. They provide an organization with the means to detect early indicators of malicious activity, in the hopes of preventing any damage from occurring to information assets.

The CERT Insider Threat Lab provides a means to develop technical solutions to the attack vectors observed in these models. The Insider Threat Lab utilizes a virtual environment to provide a test bed so as to simulate an insider attack, and to evaluate the effectiveness of different technical solutions to prevent or detect the threat. The Insider Threat Lab simulates events from actual cases to develop scenarios and possible solutions. The goal is to identify potential areas where a particular technology may prove useful in preventing or detecting the insider's malicious activity [2].

For example, in theft of intellectual property cases, insiders have been known to steal proprietary information within 30 days of tendering their resignation [1]. With this knowledge, we crafted an open-source signature that would review email logs for abnormal activity at least 30 days prior to insiders submitting their resignation and at least 30 days after separating from the organization. Based on an actual case and this known vulnerability, the team was able to demonstrate an incident in which logs were reviewed for suspicious behavior that occurred within this 30-day timeframe. In this case, the insider emailed sensitive, proprietary information to someone outside of the organization. The organization depicted in the scenario had implemented centralized logging that recorded certain email attributes, such as sender, recipient, and

attachment size. The insider's email account was reviewed for past activity as part of an employee termination procedure. It was found that the insider emailed a large amount of documents to someone outside the organization.

In addition to simulating actual insider incidents, the Insider Threat Center is able to research configuration improvements for various hardware and software packages. Like the above example, further research into developing configuration guidance is also important for organizations. CERT's research into insider threats uniquely positions us to be able to offer guidance for properly configuring software and operating systems so that the organization will be more resistant to insider attack.

The Insider Threat Lab is also critical to developing training materials as well. The Insider Threat team has used the lab to simulate real life incidents to create demonstrations as well. These demonstrations are recorded and distributed to those wishing to learn more about ways to protect their organization. The demonstrations show the various technologies in use as well as attack scenarios and how they are detected.

As the threats evolve, technology will need to evolve with it. The lab environment allows the Insider Threat team the ability to test new technologies and strategies as they become available.

References

[1] Hanley, M., Dean, T., Schroeder, W., Houy, M., Trzeciak, R.F., and Montelibano, J. "An Analysis of Technical Observations in Insider Theft of Intellectual Property Cases" SEI Technical Note CMU/SEI-2011-TN-006, Carnegie Mellon University, February 2011.

[2] Hanley, M. "Deriving Candidate Technical Controls and Indicators of Insider Attack from Socio-Technical Models and Data" in Proc. Of the 2010 CAE Workshop on Insider Threat, November 2010.

Insider Threat in the Financial Services Sector

Principal Investigators:
Adam Cummings and Andrew Moore

Though external attackers compromising information systems is an established area of organization concern, malicious insiders commit equally damaging acts of fraud, sabotage, and theft of intellectual property across both public and private sectors. However, insider compromise does not usually receive the same level of attention from media, businesses, or government leaders. These attacks not only cost a great deal of time and money, but also pose a threat to critical infrastructure and national security. The Insider Threat Center at CERT has been funded by Department of Homeland Security, Science and Technology Directorate, to examine the threats faced by the U.S. financial services sector. DHS and CERT are working with the United States Secret Service (USSS) to build on an August 2004 project and accompanying report titled *Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector* [1]. The goals of this work are to

- develop a greater understanding of the behavioral, technical, and organizational factors that lead to insider malfeasance
- identify policies, procedures, and technologies that can mitigate the risk due to insider fraud

CERT technical staff began by identifying, selecting, and coding approximately 100 cases of attacks in the financial services sector committed by both insiders and outsiders. While the selection criteria for the outsider cases are still under review, the insider cases are closed cases and gathered as part of a close ongoing partnership with the USSS. In addition to case data provided by the USSS, sources of data include interviews with relevant parties, such as subjects released from prison, affected organization representatives, and law enforcement or legal representatives with knowledge of the case. Although the project does not ignore case types of sabotage or theft of intellectual property,

it will be primarily composed of fraud cases. This is partially the result of the prevalence of certain kinds of cases in the financial services sector: 82% of the financial services cases in our database are fraud related.

Though our previous research efforts have closely examined the role of insider attacks across various sectors and types of crime, roughly half of our cases in this project will be external attacks. With this new approach, we hope to gain an understanding of what differences, if any, exist between attacks that originate from a trusted insider versus an outsider. The research team believes this approach may highlight policies or procedures that would be effective at thwarting both types of malicious activity. For example, would technology or systems used to detect sensitive information being exfiltrated from a company be equally effective, no matter who originates the activity? If not, should they be configured or implemented in unique ways, based on the specific attack vector?

Though deriving possible mitigation strategies for this problem contains issues of minimizing both false positives and false negatives, Figure 1 provides an overview of our approach to ensuring the validity of our results. The lower left-hand corner of Figure 1 presents our preliminary findings, which are based on our analysis of representatives of the population of malicious insiders and outsiders in the banking and finance sector. These findings are analyzed through the lens of relevant existing theory and reviewed by individuals with expertise in malicious insider behavior in the sector. Where there is agreement, we identify a set of moderated findings with which we have a high degree of confidence. Findings in conflict with theory and/or expert
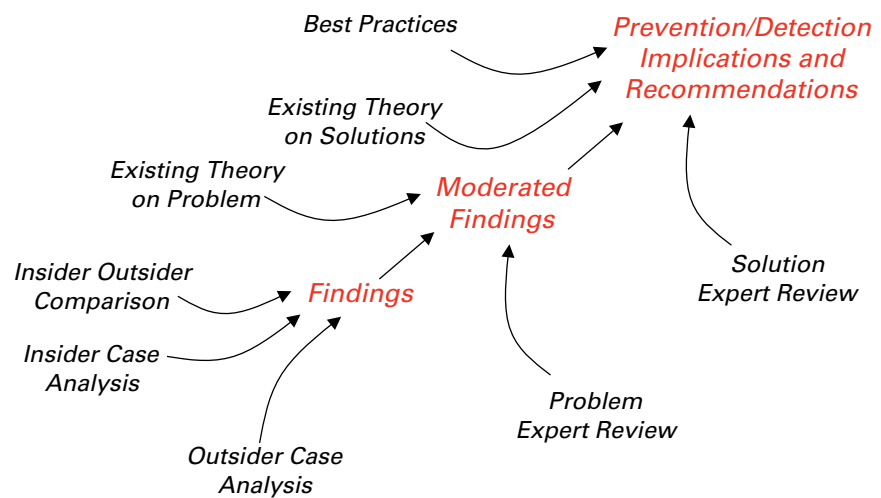
review are equally important as a source for additional analysis, case study, and experimentation. Further evidence may suggest either proposing new theory or debunking beliefs widely held by experts in the field. The findings from our study of the problem will suggest prevention and detection measures. Similar to our means of gaining

confidence in the findings of the study, we will gain confidence in the measures identified through mapping to best practices in the area and review by "solution" experts [2].

As a means of supplementing the case-based research, the Insider Threat Center at CERT is also teaming with the Department of the Treasury to enable direct interactions with the financial services sector. The team believes this will allow them to better understand and address the needs of the sector by meeting with high-level officials, such as chief security officers (CSOs), at large banking organizations. Rather than assuming that our case-based academic research will automatically help them protect themselves against these attacks, we are able to get to the heart of what the difficult aspects of this problem are for them to solve. It may also provide an opportunity to validate some of the team's preliminary findings and share any mitigation strategies they believe would have been effective to counter these costly incidents.

The final report, which should be released in late 2011, will analyze the actual incidents of insider crimes from inception to prosecution. Once the report has been released, the research team will travel to quarterly meetings of the Electronic Crimes Task Forces (ECTFs). ECTFs have been established across the nation to encourage communication and data sharing between federal, state, and local law enforcement, private industry, and academia. The findings will be presented to ECTF members, and feedback will be gathered about what additional strategies could be used to deal with a critical problem facing our financial services sector.

References

[1] Randazzo, M.R., Keeney, M.M., Kowalski, E.F., Cappelli, D.M., Moore, A.P. (2004, August) "Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector," Joint SEI and U.S. Secret Service Report. *http://www.cert.org/archive/ pdf/bankfin040820.pdf*

[2] Cappelli, D.M., Moore, A.P., Trzeciak, R.F. and Shimeall, T.J., "Common Sense Guide to Prevention and Detection of Insider Threats," Joint CyLab (CMU) and CERT (SEI), 3rd Edition, September 2008 (updated from July 2006 and April 2005). *http://www.cert.org/archive/pdf/CSG-V3.pdf*

Preventing the Federal Government from Being the Victim of Identity Theft

Principal Investigators: Alex Nicoll and Russ Griffin

In March 2005, the Internet Engineering Task Force finalized development on a significant enhancement in the way that the internet Domain Name System functions, called the Domain Name Services Security Extensions (DNSSEC). The changes were intended to significantly enhance users' ability to trust the responses they received from internet name servers and limit the ability of malicious actors to hijack their web (and other protocol) sessions. The change was slow to be adopted, and it wasn't until 2008 when security researcher Dan Kaminsky popularized a lesser known attack on the DNS system that the technology world was galvanized into action.

On August 22, 2008, the Executive Office of the President's Office of Management and Budget issued a new network security mandate (M08-23) for all government networks. A portion of this mandate requires that all government agencies that operate registered internet domains implement the security extensions (DNSSEC) to the Domain Name Services protocol. However, consistently assessing compliance with the mandate was problematic because the federal government alone had registered more than 1,800 domains. To address this problem a federal agency asked CERT to develop a tool that automatically diagnoses DNS security for a large number of domains at once.

The design team faced a significant challenge, simply because of the inherent complexity of DNSSEC. Every domain name (e.g., argylesocks.gov) is supported by at least two authoritative name servers. These name servers are the trusted repositories for all information about names and addresses (e.g., www.argylesocks.gov) contained within the domain, as well as all sub-domains (e.g., green.argylesocks.gov). The new DNSSEC extensions defined a relationship between a parent domain (e.g., .gov) and the child domain (e.g., arglyesocks. gov) that allowed the child to cryptographically "sign" any responses to a client's (user's) request for information from the child domain's name servers. The user could then verify that the child domain's response was correctly signed and also check with the parent domain to ensure that the child domain's response was valid and authorized.

The concept is simple, but actually verifying the signature is complicated. To ensure security, each authoritative name server for a domain must be able to send a correct signature and must possess the public key necessary to validate that signature. The server must also possess the public portion of a key used to create the signature, and the public portion of all keys contained on the server must itself be signed by a special key, which in turn is validated by each of the parent domain's authoritative name servers. In short, to do complete validation, each domain's authoritative name server must be checked against each parent domain's name server, validating all keys and signatures. Assuming that argylesocks.gov has three authoritative name servers, and the .gov domain has seven, there are 21 possible checks to perform for complete validation. Extending the validation to a third-level domain, such as green.argylesocks.gov, further increases the complexity and number of checks.

Unfortunately, current network security and diagnostic tools for DNS queries, such as Unbound and Dig, check only a single chain of authoritative servers because they assume that all servers for a given domain are in sync. However, this is often not the case. CERT designed its tool to perform these checks automatically for a large number of domains at one time and report on any faults in the process at a very granular level.

The end result of the work was a tool that the federal agency runs weekly to validate over 1,800 federal domains. It enables the agency's network administrators to correct any problems with its DNSSEC implementation.
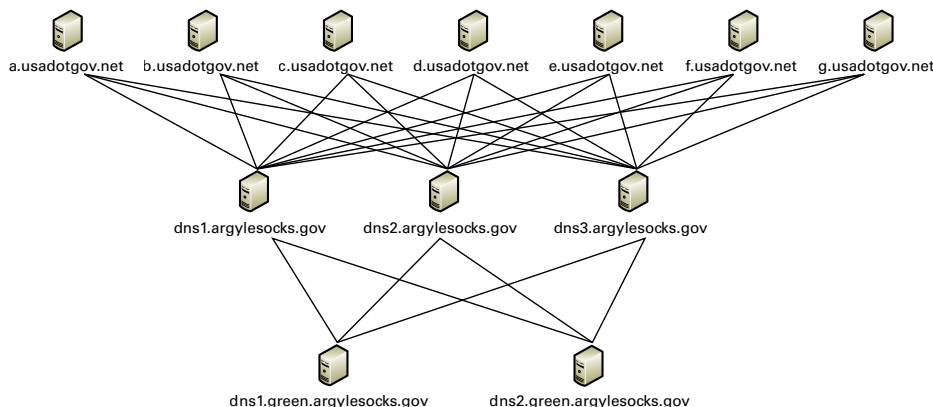

Figure 1: Sample set of verifications necessary during a DNSSEC configuration validation scan

SECURE CODING

NONCOMPLIANT CODE EXAMPLE

In this noncompliant code example, the macro CUBE() has undefined behavior when passed an expression that contains side effects.

```
#define CUBE(X) ((X) * (X) * (X))
/* ... */
int i = 2;
int a = 81 / CUBE(++i);
```

For this example, the initialization for a expands to

$$int\ a = 81 / ((++i) * (++i) * (++i));$$

which is undefined. (See rule EXP30-C. Do not depend on order of evaluation between sequence points.)

Compliant Solution

When the macro definition is replaced by an inline function, the side effect is executed only once before the function is called.

```
inline int cube(int i) {
  return i * i * i;
}
/* ... */
int i = 2;
int a = 81 / cube(++i);
```

NONCOMPLIANT CODE EXAMPLE

In this noncompliant code example the programmer has written a macro called EXEC_BUMP()

to call a specified function and increment a global counter [Dewhurst 2002]. When the expansion of a macro is used within the body of a function, as in this example, identifiers refer to the declarations in scope where the body occurs. As a result, when the macro is called in the aFunc() function it inadvertently increments a local counter with the same name as the global variable. Note that this example also violates recommendation DCL01-C. Do not reuse variable names in subscopes.

```
size_t count = 0;
#define EXEC_BUMP(func) (func(), ++count)
void g(void) {
  printf("Called g, count = %zu\n", count);
}

void aFunc(void) {
  size_t count = 0;
  while (count++ < 10) {
    EXEC_BUMP(g);
  }
}
```

WHEN IT COMES TO SOFTWARE SECURITY, SECURE CODING HAS ALWAYS BEEN THE ELEPHANT IN THE ROOM

ROBERT SEACORD

If you cannot distinguish between correct and incorrect code, the process you use to develop it is irrelevant.

ROBERT SEACORD

Secure Coding Overview

A relatively small number of software defects create many of the most commonly exploited software security vulnerabilities. CERT addresses the problem directly in the programming languages themselves. For the past seven years, the Secure Coding team has tackled some of the most critical vulnerabilities in C, C++, and Java. The Secure Coding team is working with the broader language development and programming communities to improve the security of common programming languages. The team's goal is to reduce the number of deployed vulnerabilities to a level that can be reasonably managed by existing vulnerability handling teams.

Recommendations and Standards

Industry partner Cisco helps maintain the Secure Coding wiki, where the programming community can comment on proposed recommendations and rules to secure coding standards. Robert Seacord, the Secure Coding team leader, says, "We try to target problems that can be addressed by revisions to the C language standard, where an improvement to the standard can be propagated out to a variety of compilers to drastically reduce the number of vulnerabilities." Seacord previously authored The CERT® C Secure Coding Standard, Version 1.0, and the Secure Coding team is at work on Version 2.0. In 2010, the Secure Coding team continued its work on a new standard for C++, as well as a standard for Java slated to be published in 2011.

One of the Secure Coding team's biggest successes early in fiscal year 2010 was the creation of the C Secure Coding Rules Study Group. The Study Group met regularly throughout the year to produce requirements for analyzers, mechanisms such as static analysis tools, tools within a compiler suite, and code reviewers that diagnose coding flaws in software programs. These recommendations require long and cautious research: the Study Group must examine code vulnerabilities, draft rules to fix them, build and run analysis checkers, and then iterate the process until the rules are robust and reliable. Over the past year, the Study Group has been maturing a base document in preparation for submission to the WG14-C, the International Organization for Standardization's (ISO) working group for the C programming language, in September 2011.

Analyzing Existing Software

Even with updated standards, software will inevitably contain code defects. The Secure Coding team created the Source Code Analysis Laboratory (SCALe) to assess conformance of source code to CERT secure coding standards. Trained analysts and automated analysis tools analyze source code submitted by software developers for conformance to CERT secure coding standards. In 2010, SCALe customers included the U.S. Department of Energy and the U.S. Department of Homeland Security.

Other 2010 research led to the as-if infinitely ranged (AIR) integer model. This largely automated mechanism eliminates integer overflow and integer truncation, two major causes of software vulnerabilities in the C and C++ programming languages. The Secure Coding team also prototyped the AIR integer model last year in the LLVM/Clang static analyzer. Collaborators on the AIR project included Plum Hall Inc., a publisher of standards and validation suites, and Carnegie Mellon University's School of Computer Science.

Secure Code in the Community

Much of the research of the Secure Coding team occurs at a very detailed level. But even minor changes to programming language standards can eliminate entire classes of software vulnerabilities. Nevertheless, the effort is not always visible. "A lot of our work has tremendous impact," says Seacord, "but the impact is spread across a very broad community." As dependent as we are on software to run everything from cell phones to power grids, that community includes us all.

Secure Coding Customer Spotlight

The Secure Coding team has many customer engagements with the U.S. Office of the Secretary of Defense, the U.S. Department of Homeland Security (DHS) National Cyber Security Division, and the U.S. Department of Energy (DOE). In 2010, the Secure Coding team worked with DHS to produce the Source Code Analysis Laboratory (SCALe). The Secure Coding team then used the SCALe to evaluate systems for DHS as well as DOE. Cisco and Qualcomm have both considered the secure coding standards developed by the Secure Coding team. The team also worked with the Office of Naval Intelligence and the Internal Revenue Service to analyze their legacy code. Developers, development organizations, and analyzer vendors are widely using the team's secure coding standards.

Secure Coding Initiative

Principal Investigators:
Robert C. Seacord, David Svoboda,
David Keaton, and Dean Sutherland

Problem Addressed

Software vulnerability reports continue to grow at an alarming rate, and a significant number of these reports produce technical security alerts. To address this growing threat to governments, corporations, educational institutions, and individuals, systems must be developed that are free of software vulnerabilities.

CERT takes a comprehensive approach to eliminating vulnerabilities and other software defects, starting with a detailed analysis of vulnerability reports originating from the U.S. Department of Defense (DoD) and other sources.

By analyzing thousands of vulnerability reports, CERT has observed that most vulnerabilities stem from a relatively small number of common programming errors. Software developers can take practical steps to eliminate known code-related vulnerabilities by identifying insecure coding practices and developing secure alternatives. In many cases, this can lead to language enhancements that are adopted by international standards development organizations. Compiler vendors can then implement these changes, and software developers can use the improved compilers in the supply chain and incorporate them into software-intensive systems.

In particular, CERT is working on the following projects:

1. secure coding standards. CERT provides a detailed enumeration of coding errors that have resulted in vulnerabilities and their mitigations through the development of secure coding standards for the most commonly used software development languages.

2. standards development. CERT participates in the development and evolution of international programming language standards to improve the safety and security of common programming languages.

3. automated analysis tools. CERT works with industry to develop tools that assist developers in building secure software.

4. secure compiler extensions. CERT produces safe and secure executables that are known to be free from several important classes of vulnerabilities, including buffer overflows.

5. application conformance testing. CERT offers testing of software for conformance to secure coding standards.

6. SEI Team Software Process^{SM} (TSP^{SM}-Secure). CERT integrates secure coding techniques into the TSP so that high-quality, secure software can be developed with predictable cost and schedule.

7. books, courses, training, and education: CERT creates books and courses that foster a security mindset and teach developers to code securely.
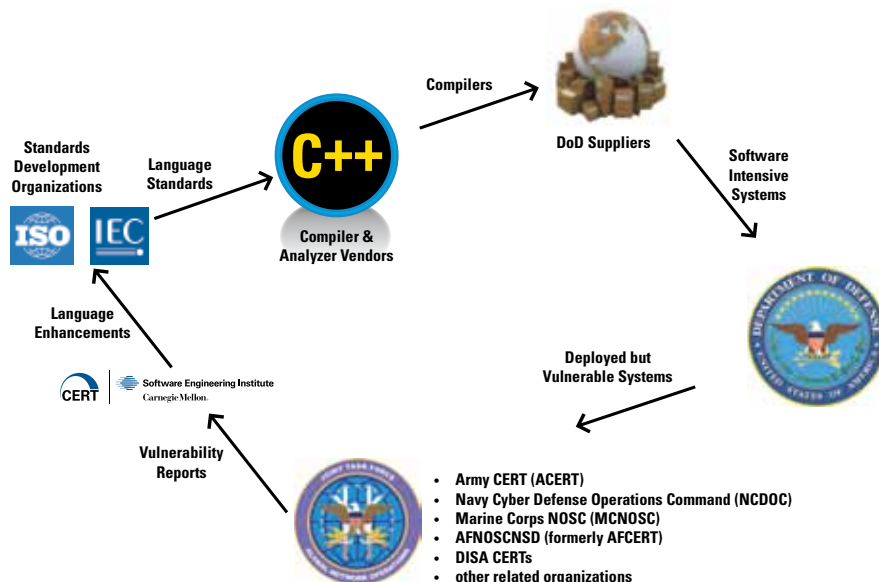
Research Approach

The foundations for secure coding work at CERT are secure coding standards for common programming languages such as C, C++, and Java. These coding standards define an enforceable set of guidelines against which the CERT® Source Code Analysis Laboratory (SCALe) can evaluate conformance.

Coding standards encourage programmers to follow a uniform set of rules and guidelines determined by the requirements of the project and organization, rather than by the programmer's familiarity or preference. Developers and software designers can apply these coding standards during software development to create secure systems.

The use of secure coding standards defines a set of rules and recommendations against which the source code can be evaluated for conformance. Secure coding standards provide a metric for evaluating and contrasting software security, safety, reliability, and related properties.

CERT coordinates development of secure coding standards by security researchers, language experts, and software developers using a wiki-based community process. More than 500 contributors and reviewers participated in the development of secure coding standards on the CERT® Secure Coding Standards wiki.

Standards Development

CERT participates in the development of international standards for programming languages to improve the safety and security of these languages. CERT chairs PL22.11 Programming Language C and is a voting member of INCITS PL22 Programming Languages, PL22.16 - Programming Language C++. In addition, CERT sends technical experts to International Organization for Standardization/International Electrotechnical Commission (ISO/IEC) working group meetings for SC22 Programming Languages and the C, C++, and programming language vulnerabilities working groups. CERT also chairs the WG14 C Secure Coding Rules Study Group.

Working with technical experts in these international standards bodies has led to the following advancements:

- the publication of TR 24731-1 [1] and its inclusion into a conditionally normative annex for C1X

- security improvements to C standard library functions

- deprecating the gets() function in C99 and removing it from C1X

- the inclusion of the Analyzability Annex into the conditionally normative annex for C1X [2]

- publication of ISO/IEC TR 24772, Guidance to Avoiding Vulnerabilities in Programming Languages through Language Selection and Use [3]

- formation of the C Secure Coding Guidelines Study Group within WG14 to study the problem of producing analyzable secure coding guidelines for C99 and C1X

CERT participation in international standards bodies improves the quality of our secure coding standards and processes and provides a channel for their adoption and publication as international standards.

Automated Analysis Tools

Secure coding standards alone are inadequate to ensure secure software development because they may not be consistently and correctly applied. Automated analysis tools, including static analysis tools, dynamic analysis tools, and tools within a compiler suite, can supplement manual security code audits. However, there are many problems and limitations in source code analysis. Static analysis techniques, while effective, are prone to both false positives and false negatives. For example, a recent study found that not one of five C and C++ source analysis tools was able to diagnose 41.5 percent of 210 test cases, while only 7.2 percent of test cases were successfully diagnosed by all five tools [4]. The same study showed that not one of six Java code analysis tools was able to diagnose 39.7 percent of 177 test cases, while 0 percent of the test cases were discovered by all six tools. Dynamic analysis tools produce lower false positives rates, but they are prone to false negatives along untested code paths. The National Institute of Standards and Technology (NIST) Static Analysis Tool Exposition (SATE) also demonstrated that developing comprehensive analysis criteria for static analysis tools is difficult because there are many different perspectives on what constitutes a true or false positive [5].

To address these problems, CERT is working with analyzer vendors and with the WG14 C Secure Coding Rules Study Group to precisely define a set of analyzable secure coding guidelines for C99 as well as for the emerging C1X major revision. Having such a set of guidelines and standardizing them through the ISO/IEC process should eliminate many of the problems encountered at the NIST SATE and also increase the percentage of defects found by more than one tool. In addition to developing a set of analyzable secure coding guidelines, CERT is coordinating a test suite under a Berkeley Software Distribution (BSD)-type license that will be freely available for any use. This test suite can then be used to determine which tools are capable of enforcing which guidelines and to establish false positive and false negative rates. Depending on the application, consumers of these tools may have different preferences for tools that can, for example, trade off a high false positive rate for a low false negative rate or vice versa.

In addition to working with commercial analyzer vendors, CERT has extended the Compass/ROSE tool (developed at Lawrence Livermore National Laboratory) to diagnose violations of the CERT secure coding standards in C and C++ language programs.

Security-Enhanced Open-Source C Compiler

Static analysis tools can be used during testing and maintenance to detect security flaws that can result in vulnerabilities. Solutions based only on static analysis place limitations on the language and cannot analyze legacy code without large numbers of false positives. Solutions based only on dynamic analysis have high overhead.

For any solution to make a significant difference in the reliability of the software infrastructure, the methods must be incorporated into tools that working programmers are using to build their applications.

Compiler producers constitute a segment of the software production supply chain, one that is quite different from the quality-tools producers. Each hardware company typically maintains some number of compiler groups, as do several of the large software producers. There are several specialized compiler producers. In addition, there is a significant community of individuals and companies that support the open-source GNU Compiler Collection (GCC). Adding these various groups together, we estimate that there are well over 100 compiler vendors.

Figure 1: Security-enhanced open-source C compiler

The CERT solution is to combine static and dynamic analysis to handle legacy code with low overhead. These methods can be used to eliminate several important classes of vulnerabilities, including writing outside the bounds of an object (for example, buffer overflow), reading outside the bounds of an object, and arbitrary reads/writes (for example, wild-pointer stores) [6]. The buffer overflow problem, for example, is solved by static analysis for issues that can be resolved at compile and link time and by dynamic analysis using highly optimized code sequences for issues that can be resolved only at runtime.

CERT is extending an open-source compiler to perform the Safe Secure C/C++ analysis methods as shown in Figure 1.
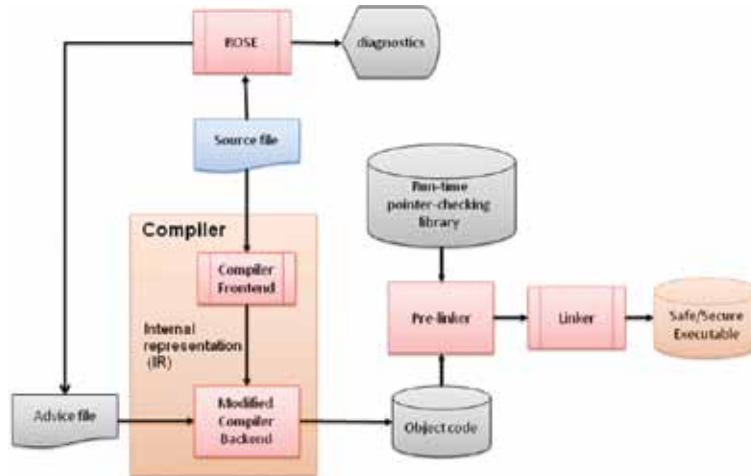
CERT has also completed a proof-of-concept implementation of the as-if infinitely ranged (AIR) integer model built upon Clang/LLVM. The AIR integer model produces either a value that is equivalent to a value that would have been obtained using infinitely ranged integers or a runtime-constraint violation. AIR integers can be used for dynamic analysis or as a runtime protection scheme. In either case, no changes are required to the source code. Consequently, the AIR integer model can be used with legacy systems by compiling C source code in analyzable mode.

At the 02 optimization level, our compiler prototype showed only a 5.58 percent slowdown when running the SPECINT2006 macro-benchmark. Although that percentage represents the worst-case performance for AIR integers (because no optimizations were performed in placing checks), it is still low enough for typical applications to enable this feature in deployed systems. AIR integers have also been proven effective in discovering vulnerabilities, crashes, and other defects in the JasPer image processing library and the FFmpeg audio/video processing library during testing with dumb (mutation) fuzzing.

Application Conformance Testing

The Source Code Analysis Laboratory (SCALe) is a research lab that tests software applications for conformance to one of the CERT secure coding standards. CERT secure coding standards provide a detailed enumeration of coding errors that have resulted in vulnerabilities for commonly used software development languages. The SCALe team analyzes a developer's source code and provides a detailed report of findings to guide the code's repair following the process shown in Figure 2. After the developer has addressed these findings and the SCALe team determines that the product version conforms to the standard, CERT issues the developer a certificate and lists the system in a registry of conforming systems [7].


Figure 2: Conformance testing process

When possible, the SCALe incorporates dynamic analysis and fuzz testing techniques in addition to the static analysis to identify coding defects and for true/false positive analysis.

For each secure coding standard, the source code is found to be provably nonconforming, conforming, or provably conforming against each guideline in the standard:

- provably nonconforming. The code is provably nonconforming if one or more violations of a rule are discovered for which no deviation has been allowed.
- conforming. The code is conforming if no violations of a rule can be identified.
- provably conforming. The code is provably conforming if the code has been verified to adhere to the rule in all possible cases.

Strict adherence to all rules is unlikely, and, consequently, deviations associated with specific rule violations are necessary. Deviations can be used in cases where a true positive finding is uncontested as a rule violation, but the code is nonetheless determined to be secure. This may be the result of a design or architecture feature of the software or because the particular violation occurs for a valid reason that was unanticipated by the secure coding standard. In this respect, the deviation procedure allows for the possibility that secure coding rules are overly strict. Deviations will not be approved for reasons of performance, usability, or to achieve other nonsecurity attributes in the system. A software system that successfully passes conformance testing must not present known vulnerabilities resulting from coding errors. Once the process is completed, a report detailing the conformance or nonconformance for each CERT C Secure Coding rule is provided to the customer.

TSP-Secure

The SEI Team Software Process (TSP) methodology, known for enabling dramatic improvement in productivity and product quality, is now being used for rapid, economic, and self-sustaining Capability Maturity Model Integration (CMMI®) implementation. TSP-Secure extends TSP to achieve the development of secure software systems by institutionalizing guidance offered by CERT, as illustrated in Figure 3. By implementing TSP-Secure, organizations can efficiently build high-quality, secure software while conforming to CMMI.

TSP-Secure incorporates the planning, process, quality, measurement, and tracking frameworks of TSP for secure software development and generates the practices and artifacts required to satisfy a Standard CMMI Appraisal Method for Process Improvement (SCAMPI^{SM}) Maturity Level 3 (ML3) appraisal. TSP-Secure requires selection of one or more secure coding standards during the requirements phase of the project. TSP-Secure teams use the application conformance testing processes as part of their own development processes to produce demonstrably conforming secure code.

Technical training for developers is delivered prior to project launch. A new team role, Security Manager, is defined. Additional launch meetings are specified and scripted. Some existing launch meetings are modified. These include modified scripts and forms. Process steps integrate the use of static analysis tools and other tools. At this time, development teams must be using C or C++ to take advantage of the security training, tools, and methods. We expect to extend them to Java development in the coming year. Finally, feedback loops put fresh information discovered in TSP-Secure projects back into our security and information repositories.
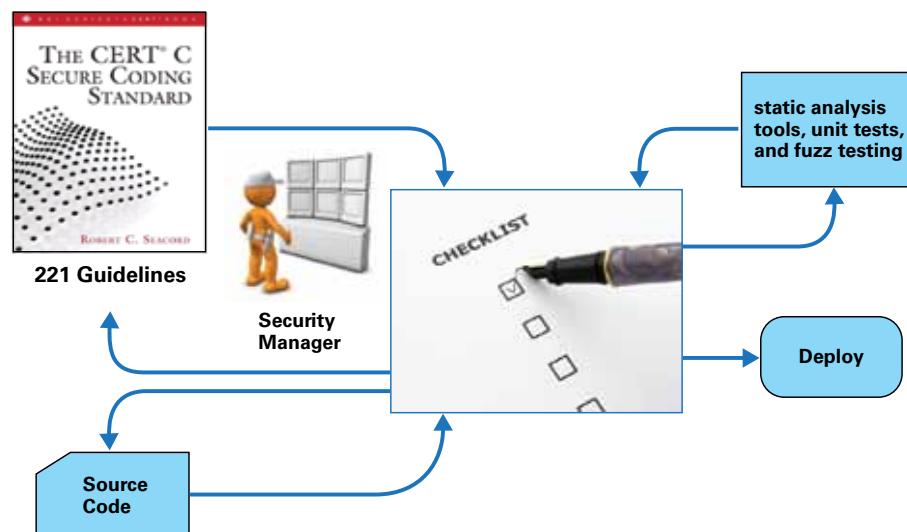

221 Guidelines

Security Manager

CHECKLIST

static analysis tools, unit tests, and fuzz testing

Deploy

Source Code

Figure 3: TSP-Secure

Books, Courses, Training, and Education

CERT has had two books published on secure coding: *Secure Coding in C and C++* [8] and the *CERT® C Secure Coding Standard* [9]. These books identify insecure coding practices, describe how insecure code can be exploited, and provide mitigation strategies.

CERT has developed a four-day Secure Coding in C and C++ course that identifies common programming errors in C and C++ and describes how these errors can lead to code that is vulnerable to exploitation. The course concentrates on security issues intrinsic to the C and C++ programming languages and associated libraries and is based on the CERT book by the same name. This course is currently being offered by the SEI and by SEI partner organizations.

CERT is also involved in teaching secure programming to undergraduates in the Computer Science department at CMU and secure software engineering to graduate students in CMU's Information Networking Institute, and it is working with other universities to improve their software security courses.

Expected Benefits

The goal of the CERT® SCALe is to reduce or eliminate vulnerabilities deployed in operational software by preventing coding errors or discovering and eliminating security flaws during implementation and testing. Organizations can benefit from this work by

- participating in the development of CERT secure coding standards and applying these standards in their software development process

- adopting, extending, and using static analysis tools (some of which are freely available) that have been enhanced to detect violations of CERT secure coding guidelines

- training their software development workforce through secure coding courses developed and offered by the SEI and SEI partner organizations

- using the resources of the CERT SCALe for conformance testing

- using TSP-Secure as their software development process

Recent Accomplishments

As-if Infinitely Ranged Integer Model

In 2010, CERT published a technical note on the as-if infinitely ranged (AIR) integer model, which provides a largely automated mechanism for eliminating integer overflow and integer truncation [10].

Java Concurrency Guidelines

The CERT Oracle Secure Coding Standard for Java provides guidelines for secure coding in the Java programming language. The goal of these guidelines is to eliminate insecure coding practices and undefined behaviors that can lead to exploitable vulnerabilities. Applying this standard will lead to higher-quality systems that are robust and more resistant to attack. In 2010, CERT published a technical report documenting the portion of those Java guidelines that are related to concurrency [11].

SCALe

CERT conducted software security assessments for the U.S. Department of Energy and the U.S. Department of Homeland Security. These assessments included systems developed in C, C++, Java, and Perl. CERT published a technical report describing the use of the SCALe for analyzing energy delivery systems [7].

References

[1] ISO/IEC. "Extensions to the C library, — Part I: Bounds-Checking Interfaces," ISO, Geneva, Switzerland, ISO/IEC TR 24731-1, Apr. 2006.

[2] T. Plum and R. C. Seacord, "ISO/IEC JTC 1/SC 22/ WG14/N1394 – Analyzability," ISO, Geneva, Switzerland, Aug. 2009. *http://www.open-std.org/jtc1/sc22/wg14/www/ docs/n1394.pdf*

[3] ISO/IEC. "Information Technology — Programming Languages — Guidance to Avoiding Vulnerabilities in Programming Languages through Language Selection and Use," ISO, Geneva, Switzerland, ISO/IEC TR 24772, 2010-10-01.

[4] C. Landwehr. (2008, July). IARPA STONESOUP Proposers Day. [Online]. Available: *http://www.iarpa.gov*

[5] V. Okun, R. Gaucher, and P. E. Black, "Static Analysis Tool Exposition (SATE) 2008," NIST, NIST Special Publication 500-279, June 2009.

[6] T. Plum and D. M. Keaton, "Eliminating Buffer Overflows Using the Compiler or a Standalone Tool," in Proc. of the Workshop on Software Security Assurance Tools, Techniques, and Metrics, Long Beach, CA, November 7-8, 2005. *https://samate.nist.gov/index.php/Past_Workshops*

[7] Robert C. Seacord, William Dormann, James McCurley, Philip Miller, Robert Stoddard, David Svoboda, and Jefferson Welch. "Source Code Analysis Laboratory (SCALe) for Energy Delivery Systems," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2010-TR-021, Dec. 2010.

[8] R. C. Seacord, Secure Coding in C and C++. Addison-Wesley Professional, 2005.

[9] R. C. Seacord, The CERT® C Secure Coding Standard. Addison-Wesley Professional, 2008.

[10] R. B. Dannenberg, W. Dormann, D. Keaton, R. C. Seacord, D. Svoboda, A. Volkovitsky, T. Wilson, and T. Plum, "As-If Infinitely Ranged Integer Model," in Proc. 2010 IEEE 21st Int. Symp. on Software Reliability Engineering (ISSRE 2010), San Jose, CA, Nov. 1-4, pp. 91-100. *http://www.computer.org/portal/web/csdl/doi/10.1109/ISSRE.2010.29*

[11] Fred Long, Dhruv Mohindra, Robert C. Seacord, and David Svoboda, "Java Concurrency Guidelines," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2010-TR-015, May 2010.

SOFTWARE SECURITY ASSURANCE

MEASUREMENT

SUPPLY CHAIN ASSURANCE

SECURE REQUIREMENTS ENGINEERING

| MATERIAL SOLUTION ANALYSIS | TECHNOLOGY DEVELOPMENT | ENGINEERING AND MANUFACTURING DEVELOPMENT | PRODUCTION AND DEPLOYMENT | OPERATIONS AND SUPPORT |

PRE-SYSTEMS ACQUISITION — SYSTEMS ACQUISITION — SUSTAINMENT
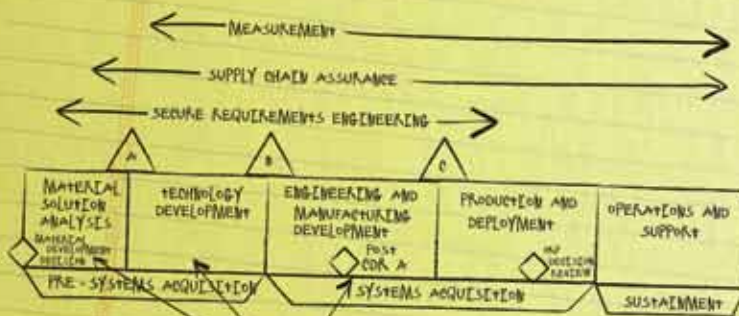
SOFTWARE SECURITY ASSURANCE NEEDS TO BE ADDRESSED IN THE EARLY STAGES OF THE DOD ACQUISITION LIFE CYCLE AND CARRIED THROUGH INTO OPERATIONS TO ALLOW REDUCTION IN SECURITY RISKS AND THEIR MISSION IMPACT.

Our company provided them with an opportunity to assess a many-faceted product and they responded graciously by sharing the different techniques they used to analyze the security aspects of our application. Their results gave us insight that has since influenced our application development and configuration.

— SQUARE CLIENT

We identify 23 activities that are essential to engineer complete and detailed security requirements. We use these 23 activities as a basis to compare five different requirements engineering processes. Our analysis shows that SQUARE incorporates more of these activities than other processes.

— Muhammad Umar Ahmed Khan and Mohammad Zulkernine, "On Selecting Appropriate Development Processes and Requirements Engineering Methods for Secure Software" presented at the 1st IEEE 32nd Annual IEEE International Computer Software and Applications Conference

Software Security Assurance Overview

The Software Security Assurance (SSA) team focuses on addressing security in the early life-cycle phases of acquisition and software development. Building security into software requires considerations beyond basic authentication/authorization and mandated operational compliance to identify and address the threat environment in which the resulting operational system must function. With greater security preparation, organizations have seen major reductions in operational vulnerabilities resulting in reductions in software patching. For example, Microsoft's own data shows that the patch levels for versions of Windows that were developed after the security "push" are half of what they were for earlier versions.[1]

Current approaches for software engineering apply a blend of training, frameworks, methods, tools, assessments, and best practices. Engineering software for effective security requires addressing all of these aspects to provide the ability to incorporate security as needed. The SSA team has developed frameworks, methods, assessments, and tools to support measurements and best practices identified to improve operational security and provide program management the ability to monitor software engineering to ensure effective consideration of security. A major gap in the security education of software engineers is being addressed through the development of curricula for colleges and universities. Transitioning the results of this research is a critical focus for SSA.

One unexpected finding of the team's research is that developing additional practices won't enable more organizations to implement software assurance into their life cycle. Instead, there's a critical need for better integration into the way software is designed and built. Wholesale change is difficult for organizations. So the SSA team has been developing practical guidelines and techniques and then piloting them to show results that are able to be replicated. If organizations can see it works, there's a better chance they'll implement it.

"It's like creating a cookbook," says Carol Woody, technical manager for SSA. "You build the recipe and then someone has to figure out how to cook it in their kitchen. We're developing customizable frameworks, methods, and techniques that organizations can tailor to their existing software acquisition and engineering practices."

The team worked on the following major research projects in 2010, collaborating with researchers in other SEI teams, at CMU, and at other universities and organizations world-wide.

Building Assured Systems Framework (BASF)

The SSA team developed the BASF, which provides a meaningful context and structure within which to describe, compare, and contrast research and development methods for building assured systems. It can also be used to identify gaps, prioritize new research projects, and stop or decommission current research projects that are not contributing useful results.

Supply Chain Assurance

Researchers developed an approach for assessing software supply chains and identifying the associated software assurance risks. SSA collaborated with members of the SEI's Acquisition team on this work.

Survivability Analysis Framework (SAF)

The SAF was a major area of research in fiscal year 2009 that informed Software Security Assurance research in fiscal year 2010. SSA researchers documented the SAF, an analysis technique for analyzing complexity and integration issues throughout the development life cycle for project management and stakeholders to ensure that development is proceeding toward an expected operational solution, for public release. The SAF was piloted for Joint Battle Mission Command and Control (JBMC2) in the analysis of a Time Sensitive Targeting mission thread for the Office of the Under Secretary of Defense (Acquisition, Technology, and Logistics) (OUSD [AT&L]). A second pilot analysis was completed for Time Sensitive Targeting information assurance for Electronic Systems Center, Cryptologic Systems Group, and Network Systems Division (ESC/CPSG NSD). The pilot results were documented in special reports for the U.S. Department of Defense.

---

[1] *http://www.microsoft.com/security/sdl/learn/measurable.aspx*

Software Security Measurement

This research focused on how to establish and specify a level of security and then how to measure, at each phase of the life cycle, whether that level of security has been achieved. The SSA team collaborated with members of the CERT Resilience Management team and the SEI Measurement and Analysis team.

Security Requirements Engineering

Several authoritative studies have shown that requirements engineering defects cost 10 to 200 times as much to correct once fielded than if they were detected during requirements development. The SSA team collaborated with other researchers and led several teams of CMU students in developing processes and tools to help organizations build security into the requirements engineering process.

Trusted Hardware for Cyber Security

This research evaluated the promise and limitations of using trusted hardware as a foundation for achieving demonstrably high assurance of end-to-end security properties of applications executing in extreme adversarial environments. It laid the groundwork for future work that will explore and exploit the concepts of trust and trustworthiness and provide a scientific basis for understanding the relationships among hardware, software, security, and trust.

Catastrophe Analysis

Along with researchers from the SEI Acquisition and System Design teams, SSA researchers analyzed key dynamics that take place and how they affect the country's technical infrastructure when catastrophes occur. The goal of this research is to understand complex failure in order to better build and operate technologies and address today's complex, software-dependent networked systems.

Complexity Modeling and Analysis

The SSA team partnered with SEI experts from Systems of Systems, Acquisition, and CERT Insider Threat teams to apply modeling techniques to analyze software assurance solutions for the increasingly complex, highly interconnected, rapidly changing state of software systems. The team created a modeling framework to examine the gaps, barriers, and incentives that affect the development and implementation of assurance solutions for complex systems.

The SSA team also supported the Department of Homeland Security (DHS) Processes and Practices, Measurement, and Workforce Education and Training Working Groups. This work informs the software engineering community about software assurance best practices and is available on the Build Security In (*https://buildsecurityin.us-cert.gov/bsi/home.html*) and Software Assurance Community Resources and Information Clearinghouse (*https://buildsecurityin.us-cert.gov/swa/*) websites.

Through their research and transition efforts, the SSA team has led the way for addressing security early in the software life cycle.

Building Assured Systems Framework (BASF)

Principal Investigators:
Nancy R. Mead and Julia H. Allen

Problem Addressed

There is no single, recognized framework to organize research and practice areas focused on building assured systems (BAS). Sponsors of the CERT Program's research could use such a framework to help address the following challenges, including customer "pain points" and general research problems:

- How do I decide which security methods fit into a specific life-cycle activity?

- How do I know if a specific security method is sufficiently mature for me to use on my projects?

- When should I take a chance on a security research approach that has not been widely used?

- What actions can I take when I have no approach or method for prioritizing and selecting new research or when promising research appears to be unrelated to other research in the field?

Such a framework could also help organize CERT research efforts.

Some organizations have already begun addressing BAS in research and development including

- organizations participating in the Building Security In Maturity Model [1]

- Microsoft's software development lifecycle (SDL) [2]

- Software Assurance Forum for Excellence in Code (SAFECode) consortium members [3]

- Oracle

- members of the Open Web Application Security Project (OWASP) using the Software Assurance Maturity Model (SAMM)

Efforts to incorporate BAS tend to be stronger in vendor organizations. However, they are weaker in large organizations developing systems for use in-house and integrating across multiple vendors. They are also weaker in small- to medium-sized companies developing products for licensed use. Furthermore, there are a variety of life-cycle models in practice—no single approach has emerged as

standard. Even in the larger organizations adopting secure software engineering practices, there is a tendency to select a subset of the total set of recommended or applicable practices. Such uneven adoption of BAS suggests the need for ways to measure results.

Research Approach

To understand previous and current work that could inform BASF development, we started by examining a number of existing software development and acquisition life-cycle process models, models for the development of more secure software, and research frameworks in software security and assurance. With this information, we formed a hypothesis that the recently developed Master of Software Assurance (MSwA2010) body of knowledge (BoK) [4] could serve as our starting point for the BASF. This makes sense given that the curriculum BoK draws extensively from more than 25 sources describing methods, practices, and technologies for software assurance and security (including the software security models considered in this report). Also, as the authors of this report, we led and contributed to the development of the MSwA2010 curriculum.

We tested this hypothesis by assigning "maturity levels" to each area of the MSwA2010 BoK. BoK areas include assurance across life cycles, risk management, assurance assessment, assurance management, system security assurance, system functionality assurance, and system operational assurance. We defined these levels as follows:

- L1—The area provides guidance for how to think about a topic for which there is no proven or widely accepted approach. The intent of the area is to raise awareness and aid the reader in thinking about the problem and candidate solutions. The area may also describe promising research results that may have been demonstrated in a constrained setting.

- L2—The area describes practices that are in early pilot use and are demonstrating some successful results.

- L3—The area describes practices that have been successfully deployed (mature) but are in limited use in industry or government organizations. They may be more broadly deployed in a particular market sector.

- L4—The area describes practices that have been successfully deployed and are in widespread use. Readers can start using these practices today with confidence. Experience reports and case studies are typically available.

To test this hypothesis further, we mapped existing CERT research work to the MSwA2010 BoK to see whether there were corresponding BoK areas for each research project. All major research projects did correspond to one or more BoK areas, either directly or indirectly. This gave us confidence that the BoK areas (and the research from which

they were derived) could be used as our initial framework. Once we mapped the current CERT research projects to the MSwA2010 BoK, we performed an initial gap analysis to identify some promising research areas for CERT.

The BASF helps to address some, but not all, of the four research questions stated previously. Since the BASF naturally covers the development life cycle, mapping a particular security method to the appropriate knowledge area(s) does help to answer the first question (relationship of security method to life-cycle phase). For the second question (security method maturity), considering knowledge area maturity levels in conjunction with examining a specific method provides information to help decide whether the method is sufficiently mature for use. The third question is a bit harder to answer and requires more work on the part of a BASF user. A cost/benefit analysis or risk assessment aids in answering the third question of whether it is worth taking a chance on a method that has not been widely used.

Expected Benefits

From a research perspective, researchers could consider periodically rating the maturity of their methods using the research approach described above. This would assist BASF users in deciding which methods to use. It would also be helpful if researchers and research methods users could begin to collect and provide cost/benefit data. All too often, researchers and research method users decide on a particular method but do not collect any information to determine whether the benefit justified the cost or to help inform future decisions.

We believe the BASF provides a context and structure for CERT's research work in building assured systems and that it can be used to show how various research efforts fit together. The gap analysis that we have done could be used to help in selecting new research and, to some extent, in prioritizing research projects. We anticipate that the BASF could be used in planning and justifying CERT's research program and communicating about it with others.

We expect that the U.S. Department of Defense (DoD) and other sponsors will find the BASF useful for tracking current research and development efforts in building assured systems and possibly in acquiring assured systems.

2010 Accomplishments

In 2010 we performed the research described above and documented the results in a technical report on the BASF [6].

Future Goals

To maximize its usefulness, the BASF needs to be more comprehensive. The BASF helps to address some, but not all, of the customer pain points. It is helpful in addressing the first and second questions, but is limited in its usefulness in addressing the third question. There are some areas of research that do not fit the BASF neatly. The BASF is not

intended to exclude these areas, but we recognize that some important research work does not fit the MSwA2010 topics directly. For example, our recent software assurance curriculum work is needed research, but it does not map directly to the MSwA2010 topics. As another example, some of our advanced work in intrusion detection and network analysis also does not map directly to these topics. This may suggest the need for follow-on work to broaden the BASF to provide a framework for a wider range of research activities.

References

[1] McGraw, Gary; Chess, Brian; & Migues, Sammy. Building Security In Maturity Model BSIMM v2.0. *http://www.bsimm2.com/ (Accessed March 2011)*

[2] Lipner, S. & Howard M. "The Trustworthy Computing Security Development Lifecycle." March 2005. *http://msdn.microsoft.com/en-us/library/ms995349.aspx*

[3] Software Assurance Forum for Excellence in Code (SAFECode). SAFECode. *http://www.safecode.org* (2011).

[4] Mead, Nancy R.; Allen, Julia H.; Ardis, Mark; Hilburn, Thomas B.; Kornecki, Andrew J.; Linger, Rick; & McDonald, James. Software Assurance Curriculum Project Volume I: Master of Software Assurance Reference Curriculum (CMU/SEI-2010-TR-005, ESC-TR-2010-005). Software Engineering Institute, Carnegie Mellon University, 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tr005.cfm*

[5] Allen, Julia H.; Barnum, Sean; Ellison, Robert J.; McGraw, Gary; & Mead, Nancy R. Software Security Engineering: A Guide for Project Managers. Addison-Wesley Professional, 2008.

[6] Mead, Nancy R. & Allen, Julia H., Building Assured Systems Framework (CMU/SEI-2010-TR-025). Software Engineering Institute, Carnegie Mellon University, September 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tr025.cfm*

Supply Chain Assurance

Principal Investigators:
Robert Ellison, Chris Alberts,
Rita Creel, Audrey Dorofee, and Carol Woody

Problem Addressed

The term "supply chain" has a long history in the business community and includes recent trends such as such just-in-time inventory. In the past, the business community considered supply chains as relevant only to the delivery of physical products. Now the business community uses the technology supply chain to develop most IT systems (hardware, software, public and classified networks, and connected devices), which together enable the uninterrupted operations of key government and industrial base actors, such as the Department of Defense, the Department of Homeland Security, and their major suppliers. While we have decades of physical supply chain data that have led to effective management practices, we have limited experience with software supply chains. While no perfect solution exists, much can be done to enable organizations to reduce risk effectively and efficiently while leveraging the significant opportunities afforded by supply chains.

On-time delivery and costs often get the most commercial attention, but some of the most serious risks are associated with system assurance, the confidence that the system behaves as expected. Software defects, such as design and implementation errors, can lead to unexpected behaviors or to system failure. Defects that enable an attacker to purposely change system behavior are often referred to as vulnerabilities. The source of such vulnerabilities is the supply chain, which includes commercial product vendors, custom development and integration contractors, and suppliers and subcontractors to those organizations. This research considers how to better manage the acquisition of software developed through a supply chain to reduce the likelihood of operational vulnerabilities.

Unfortunately, exploitable software defects are widespread. MITRE has analyzed successful attacks and identified more than 600 common software weaknesses, described in its Common Weakness Enumeration (CWE). Many of the CWE defects are widely known, as are the techniques that eliminate them. But those techniques are frequently not applied. For example, countermeasures for SQL injections are well established, yet SQL injections still rank second on the MITRE/SANS list of the top 25 most dangerous software errors. Veracode's State of Software Security Report released on September 22, 2010 warns that most software is very insecure. Regardless of software origin, 58 percent of all applications submitted to Veracode for testing did not achieve an acceptable security score upon first submission.

Software supply chain security issues do not vanish when an acquisition is completed. Product designers base their decisions on the data available and the threats known at the time of development. Product assessments performed as part of the initial acquisition for a commercial component are valid only at that time.

Some examples of sources of risks that may emerge during deployment include the following:

- New attack techniques and software weaknesses cannot be foreseen.

- Product upgrades that add features or change design can invalidate the results of prior risk assessments and may introduce vulnerabilities.

- Corporate mergers, new subcontractors, or changes in corporate policies, staff training, or software development processes may eliminate expected supply chain risk management (SCRM) practices.

- Product criticality may increase with new or expanded usage.

Research Approach

In an attempt to integrate development and acquisition practices with risk-based evaluation and mitigation of product vulnerabilities, the SEI has begun research that explores the complex dynamics of software supply chain risk and examines techniques, such as systematic risk assessment, based on key drivers [1], use of assurance cases [2], attack surface analysis and threat modeling [3, 4], and consideration of supply chains for systems as well as systems of systems [5].

Taking a systems perspective on software supply chain risks, this research considers current practices in software supply chain analysis and seeks some foundational practices. The role of an acquirer depends on the nature of an acquisition. Product development is completed in advance of an acquirer's product and supplier assessment. An acquirer seeks evidence that software developers have applied appropriate practices such as threat modeling and security testing. Acquirers need to understand the residual risks they will have to accept and accommodate in their operational implementation.

This research concentrates primarily on the role of the acquirer in software supply chain risk analysis for security. However, both suppliers and acquirers should perform such analysis, and it should consider the three components shown below and in Figure 1.

- attack analysis: factors that lead to successful attacks
- supplier: capability to limit product attributes that enable attacks
- acquirer: tradeoff decisions (desired usage and acceptable business risks)
- business risk assessment: identify attack enablers and possible business risks
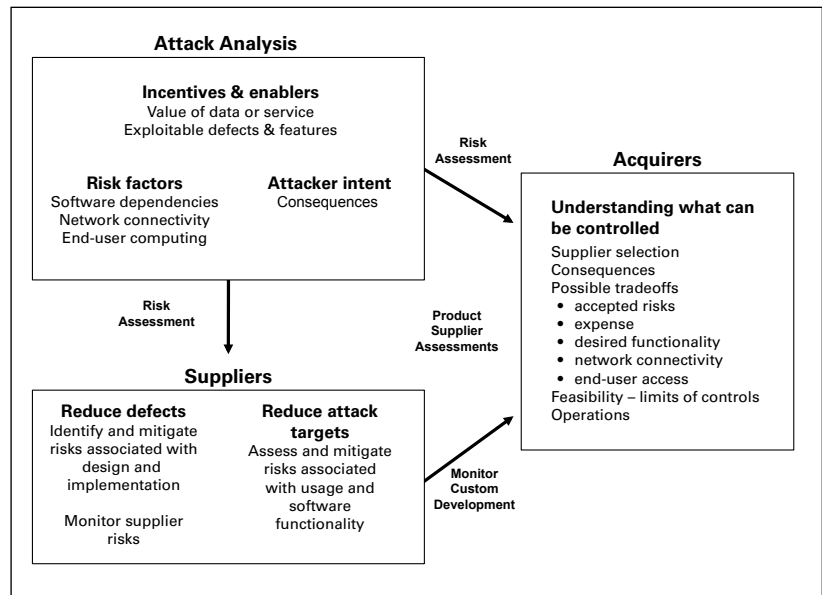- supplier/product assessment in terms of attack enablers and capability of supplier to manage them

Several factors, as shown in Figure 2, affect the occurrence of supply chain risks and the ability of an acquirer to manage them.

- Custom-developed software systems enable the acquirer to monitor and control risks during development. However, systems are increasingly constructed by integrating commercially available software, in which case the only controls might be to accept the risks or not to use a specific product.
- The owner of a system that participates in a system of systems has no control over or knowledge of the security risks of the other member systems.
- Expanded network connectivity and increased interoperability and dependencies among systems can increase the exposure of a system to adverse conditions. For example, a system for a large supplier has interfaces to their purchasers, manufacturers, and their transporters. Retailers, manufacturers, and suppliers are at risk when one of the other participating systems has been compromised.
- End-user software has always been a target for attackers. A large user community increases the likelihood of attack success. When the primary medium of data exchange was the floppy disk, an attacker might have used a Microsoft Word or Excel macro as malware. In 2010 the web is the dominant medium of data exchange, and web pages are used to install malware. Increased end-user connectivity, compromised mobile applications, and misconfigured end-user software increase the likelihood of end-user device compromise.
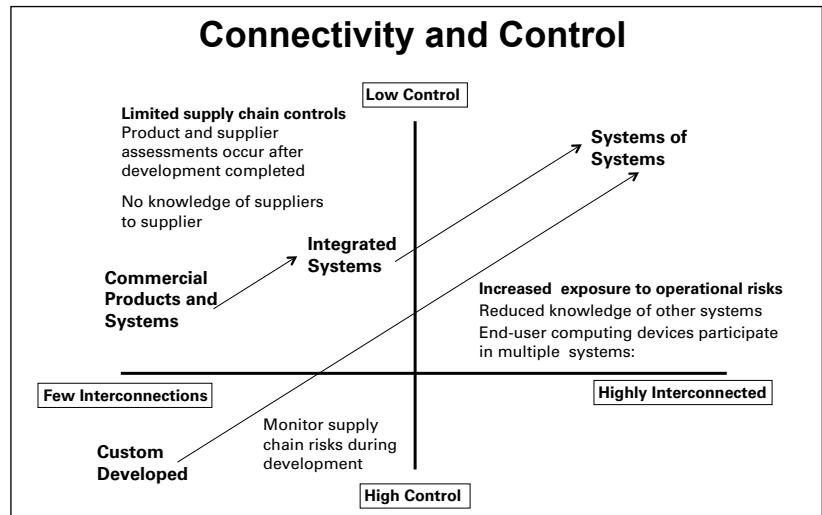

Figure 1

Figure 2

Expected Benefits

The expected acquirer benefits of this research include

- an understanding of the supply chain factors that can be effectively managed to reduce risks and the management of those factors during deployment

- for outsourced development and integration, acquirer practices to monitor and mitigate supply chain risks

- for commercial components, an identification of essential supplier and product attributes appropriate for an acquisition

2010 Accomplishments

In 2010, the SEI

- developed a supply chain risk model [1] and identified supply chain factors based on the type of acquisition [5] funded by the Department of Homeland Security (DHS)

- participated in DHS Software Assurance Working Groups and Forums

- held an internal SEI workshop with participation from members of the SEI Acquisition Support Program to identify supply chain issues that organizations supported by the SEI have encountered and to discuss how those concerns could be addressed

- presented the Supply Chain Risk Management Framework to the DHS Software Assurance Forum, March 2010

Future Goals

The SEI is proposing future work that will help acquirers build the capability to identify software supply chain risks, select mitigation solutions for key risks, and measure the effectiveness of solutions throughout the life cycle, as well as to obtain leading indicators related to software supply chain security.

As noted in the introduction, known software development practices exist that can reduce the occurrence of vulnerabilities. We are seeking organizations interested in helping us establish the risk reduction from incremental incorporation of such demonstrated practices into their acquisitions.

References

[1] Christopher Alberts, Rita Creel, Audrey Dorofee, Robert Ellison, and Carol Woody, "A systemic approach for assessing software supply-chain risk," in *Proc. 44th Hawaii Int. Conf. on System Sciences,* Kauai, HI, 2011. *https://buildsecurityin.us-cert.gov/bsi/articles/best-practices/acquisition/1230-BSI.html*

[2] Robert J. Ellison, John B. Goodenough, Charles B. Weinstock, and Carol Woody, "Evaluating and mitigating software supply chain security risks," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2010-TN-016, 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tn016.cfm*

[3] Robert Ellison and Carol Woody, "Supply-chain risk management: Incorporating security into software development," in *Proc. 43rd Hawaii Int. Conf. on System Sciences,* Poipu, Kauai, HI, 2010. *https://buildsecurityin.us-cert.gov/bsi/articles/best-practices/acquisition/1140-BSI.html*

[4] Robert Ellison and Carol Woody, "Considering software supply-chain risks," CrossTalk, vol. 23, no. 5, pp. 9-12, Sep.-Oct. 2010.

[5] Robert J. Ellison, Christopher J. Alberts, Rita C. Creel, Audrey J. Dorofee, and Carol Woody, "Software supply chain risk management: From products to systems of systems," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2010-TN-026, 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tn026.cfm*

Measuring Software Security Assurance

Principal Investigators:
Christopher Alberts, Julia Allen, and Robert Stoddard

Problem Addressed

Many organizations measure just for the sake of measuring, with little or no thought given to what purpose and business objectives are being satisfied or what questions each measure is intended to answer. However, meaningful measurement is about transforming strategic direction, policy, and other forms of management decision into action and measuring the performance of that action.

Effective measures express the extent to which objectives are being met, how well requirements are being satisfied, how well processes and controls are functioning, and the extent to which performance outcomes are being achieved. The basic goal of measurement and analysis is to provide decision makers with the information they need, when they need it, and in the right form. In recent years, researchers have begun to turn their attention to the topic of software security assurance and how to measure it.

Software security assurance is justified confidence that software-reliant systems are adequately planned, acquired, built, and fielded with sufficient security to meet operational needs, even in the presence of attacks, failures, accidents, and unexpected events. For several years, various groups within the software engineering community have been working diligently to identify practices aimed at developing more secure software. However, efforts to measure software security assurance have yet to materialize in any substantive fashion, although some foundational work has been performed [1].

As a result of the software engineering community's interest, the CERT® Program at Carnegie Mellon University's Software Engineering Institute (SEI) has chartered the Security Measurement and Analysis (SMA) Project to advance the state-of-the-practice in security measurement and analysis. The SMA Project builds on the CERT Program's core competence in software and information security as well as the SEI's work in software engineering measurement and analysis. The purpose of this new research project is to address the following three questions:

- How do we establish, specify, and measure justified confidence that a software-reliant product is sufficiently secure to meet operational needs?

- How do we measure at each phase of the development or acquisition life cycle that the required/desired level of security has been achieved?

- How do we scale measurement and analysis approaches to complex environments, such as large-scale, networked, software-reliant systems (e.g., systems of systems)?

In essence, the three research questions examine how decision makers (e.g., development program and project managers as well as acquisition program officers) can measure and monitor the security posture of large-scale, networked, software-reliant systems across the life cycle and supply chain.

Research Approach

Our research approach comprises the following activities:

- survey existing measurement and analysis approaches

- identify any limitations in existing approaches relevant to their application to large-scale, networked systems

- develop a framework for measuring the security characteristics of large-scale, networked systems

- develop a suite of methods and tools for implementing the framework

Our survey of traditional security measurement and analysis approaches indicated that they do not readily scale to today's large-scale, networked, software-reliant systems [1]. As a result, decision makers lack confidence in the security characteristics of their software infrastructures.

Traditional measurement and analysis approaches are based on the principle of system decomposition and component analysis, where the first step is to decompose a system into its constituent components. Next, the individual components are prioritized, and only the most critical components are analyzed in detail. Limitations of traditional approaches include the following:

- Only critical components are analyzed; non-critical components and interdependencies among components are not addressed.

- Causal relationships are presumed to be simple, direct, and linear. Non-linear relationships, such as feedback, are not analyzed.

- Confidence in the performance of critical components is not sufficient for establishing confidence in the performance of the parent system (or the parent system of systems).

Based on our research, we developed the SEI Integrated Measurement and Analysis Framework (IMAF), which is shown in Figure 1. IMAF employs systemic analysis to integrate subjective and objective data from a variety of sources, including targeted analysis, status reporting, and measurement, to provide decision makers with a consolidated view of the performance of large-scale, networked systems.
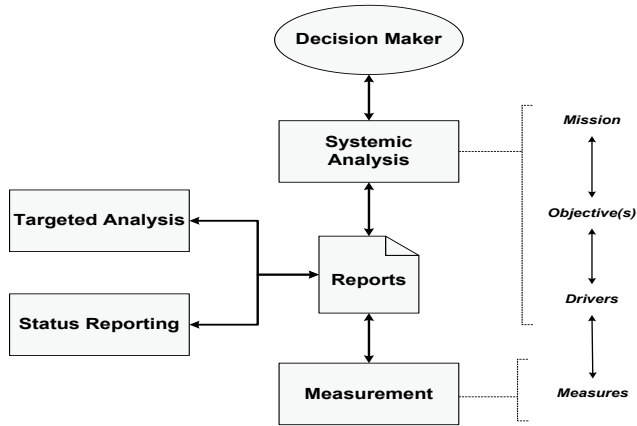

Figure 1: SEI Integrated Measurement and Analysis Framework (IMAF)

Systemic analysis is based on system theory. The underlying goal of system theory is to analyze a system as a whole rather than decomposing it into individual components and then analyzing each component separately [2]. In fact, some properties of a system are best analyzed by considering the entire system, including

- influences of environmental factors
- feedback and nonlinearity among causal factors
- systemic causes of failure
  (as opposed to proximate causes)
- emergent properties

The SEI approach for conducting systemic analysis requires identifying and analyzing a set of factors that have a strong influence on a system's mission and objectives. These factors are called drivers [3]. Figure 1 shows how drivers enable decision makers to link the security mission and objectives to measures that provide insight into a system's security characteristics. SEI experience shows that effective performance assessment requires approximately 15 to 25 drivers.

To assess secure development of software-reliant systems, we identified a total of 17 drivers. Nine drivers focus on programmatic issues: program security objectives, security plan, contracts, security process, security task execution, security coordination, external interfaces, organizational and external conditions, and event management. The remaining eight drivers examine product and operational attributes: security requirements, security architecture and design, code security, integrated system security, adoption barriers, operational security compliance, operational security preparedness, and product security risk management.

Finally, as illustrated in Figure 2, we have started to develop the following methods for implementing IMAF:

- The Software Security Review (SSR) is a method conducted by independent teams to assess the security characteristics of software-reliant systems. SSR is a driver-based approach that can be used to measure and monitor software security assurance across the life cycle and supply chain (including acquisition, development, and operations).

- Model-Based SSR incorporates predictive analytics, such as Bayesian Belief Networks (BBNs), into its analysis approach. Model-Based SSR enables quantitative analysis of software security assurance using a combination of subjective and objective data.

- Multi-View Decision Making (MVDM) is a coordinated approach for applying multiple security assessment methods. MVDM uses SSR to provide a broad view of software security assurance. An assessment team can use the findings of SSR to select and perform follow-on, "deep-dive" assessments. MVDM helps optimize security assessment activities by applying resources where and when they are most needed.
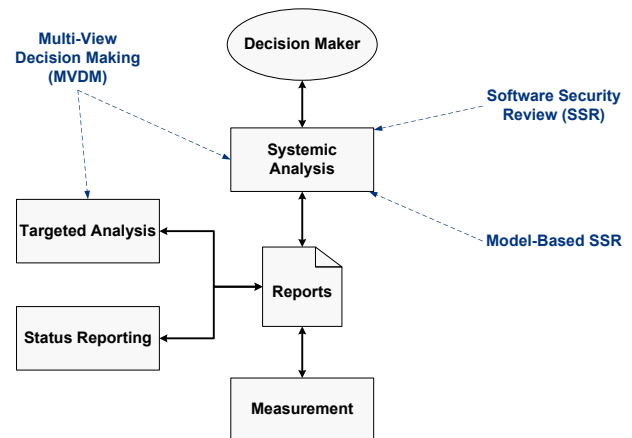

Figure 2: SEI measurement methods for software security assurance

IMAF and its associated methods provide a unique approach for software security measurement and analysis because they

- assess the behavior of large-scale, networked, software-reliant systems as a whole
- enable analysis of complex interrelationships and dependencies among a system's components
- establish justified confidence in the security characteristics of large-scale, networked systems across the life cycle and supply chain

We are currently beginning to pilot IMAF and its associated methods.

Expected Benefits

Expected benefits of this research include the following:

- Decision makers will have better tools for predicting and diagnosing security-related problems and for making well-informed decisions about security.

- IMAF and its associated methods will provide justified confidence in the security of software-reliant products that are acquired, developed, deployed, and sustained by acquisition and development programs.

- IMAF and its associated methods will provide a robust platform for conducting research in any security domain that requires measurement and analysis.

2010 Accomplishments

The 2010 accomplishments of the SMA Project include the following:

- developed the initial version of IMAF

- developed a prototype set of drivers for secure development of software-reliant systems

- initiated development of the SSR and MVDM assessment methods

- identified candidate security practices and measures related to selected drivers from the prototype set

- performed an initial mapping of security standards NIST 800-53 and ISO 27002 to the prototype set of drivers for secure development of software-reliant systems

- developed a notional Bayesian Belief Network using the prototype set of drivers

Future Goals

In 2011, we plan to make progress in the following areas:

- begin piloting the SSR and MVDM methods

- use the results of these pilots to refine IMAF, SSR, and MVDM as appropriate

- use the results of pilots to revise the prototype set of drivers for secure development of software-reliant systems

- begin development of driver sets focused on other parts of the life cycle and supply chain

- continue mapping security standards to driver sets

- develop Bayesian Belief Networks for selected driver sets

- begin development of Model-Based SSR

- mine data from SSR and MVDM pilots to identify a baseline set of software security measures

- explore applying IMAF to other security domains, such as incident management and operational security management

References

[1] Christopher Alberts, Julia Allen, and Robert Stoddard, "Integrated measurement and analysis framework for software security," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2010-TN-025, 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tn025.cfm*

[2] Nancy Leveson, "A new accident model for engineering safer systems," *Safety Science,* vol. 42, no. 4 pp. 237-270, April 2004.

[3] Christopher Alberts and Audrey Dorofee, "A framework for categorizing key drivers of risk," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2009-TR-007, 2009. *http://www.sei.cmu.edu/library/abstracts/reports/09tr007.cfm*

Security Requirements Engineering

Principal Investigator: Nancy R. Mead

Problem Addressed

When security requirements are considered at all during the system development life cycle, they tend to be general lists of security features such as password protection, firewalls, virus detection tools, and the like. These are, in fact, not security requirements at all but rather implementation mechanisms that are intended to satisfy unstated requirements, such as authenticated access. As a result, security requirements that are specific to the system and that provide protection of essential services and assets are often neglected. In addition, the attacker perspective is not considered, with the result that security requirements, when they exist, are likely to be incomplete. We believe that a systematic approach to security requirements engineering will help to avoid the problem of generic lists of features and to take into account the attacker perspective.

In reviewing requirements documents, we typically find that security requirements, when they exist, are in a section by themselves and have been copied from a generic set of security requirements. The requirements elicitation and analysis that is needed to get a better set of security requirements seldom takes place.

Much requirements engineering research and practice has addressed the capabilities that the system will provide. So while significant attention is given to the functionality of the system from the user's perspective, little attention is given to what the system should not do. In one discussion on requirements prioritization for a specific large system, as part of an earlier project that illustrated the need for attention to security requirements engineering, ease of use was assigned a higher priority than security requirements. Security requirements were in the lower half of the prioritized requirements. This occurred in part because the only security requirements that were considered had to do with access control.

Research Approach

CERT researchers have developed a methodology to help organizations build security into the early stages of the production life cycle. The Security Quality Requirements Engineering (SQUARE) methodology consists of nine steps that generate a final deliverable of categorized and prioritized security requirements. Although SQUARE could likely be generalized to any large-scale design project, it was designed for use in software systems.
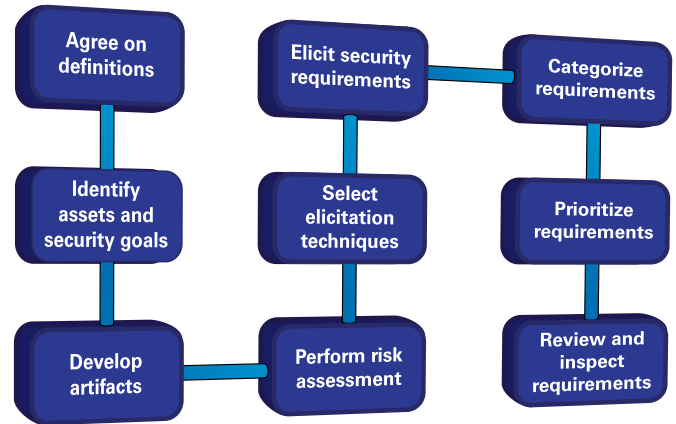

Figure 1

Subsequent to the development of SQUARE, pilot projects were conducted to validate the method, and associated tools, academic lectures, and industry workshops were developed and delivered. CERT collaborated with researchers in the U.S. and elsewhere in the research of the SQUARE method and its transition into practice. The method was recently extended to address privacy (P-SQUARE) and acquisition (A-SQUARE).

In our work with SQUARE we have found that many forward-thinking organizations already have documented processes and are not ready to embark on a whole new process. For those organizations, there is more benefit to enhancing their existing processes to ensure that security requirements are adequately addressed. Likewise, they may have existing requirements engineering tools in use and need to understand how to address security requirements in the context of those tools. Our current research is focused on both piloting existing methods such as SQUARE and using the results of those pilot efforts to enhance existing organizational processes. This approach can be applied to organizations that are concerned with security requirements, privacy requirements, or acquisition of secure developed products or commercial, off-the-shelf (COTS) products.

Expected Benefits

Industry studies have shown that addressing potential vulnerabilities during requirements engineering is 20 to 100 times less expensive than patching the vulnerability in a fielded system. Major software vendors are now addressing security requirements early, as are other leading companies. Vulnerabilities in mission-critical systems can lead to compromised systems, failures, and possible loss of life. While security needs to be addressed at every life-cycle phase, there are clear benefits to addressing it at the earliest possible stage, before architectural decisions have been made that may preclude optimal security solutions. Acquirers of developed software need to feel confident that security has been addressed early. Likewise, acquirers of COTS software need to factor security requirements into their purchasing decisions.

2010 Accomplishments

CERT was the client for a team of CMU Master of Software Engineering students in completing a robust SQUARE tool, which was made available for trial use and download at http://www.cert.org/sse/square. The team also created a set of five academic lectures on security requirements engineering available for download. Additional available materials include a tutorial and workshop/case study materials. SQUARE is included in books [1, 2] and papers, and is routinely cited in security requirements engineering research papers in the research literature. Extensions to SQUARE for privacy and acquisition have been documented in reports [3] and papers [4]. Prototype tools to support A-SQUARE and P-SQUARE are currently under development by two CMU student teams advised by CERT researchers.

Future Goals

Using our work on SQUARE and research by other organizations such as the Comprehensive, Lightweight Application Security Process (CLASP) and the Common Criteria for Information Technology Security Evaluation, we plan to extend our work in security requirements engineering, with a special focus on security and privacy, and especially on acquisition. Our goal is for security requirements engineering to become part of standard processes [5] as well as international standards. We are currently collaborating with researchers at CMU and at other universities and in industry. We are also looking for collaborators and sponsors so that we can pilot and extend security requirements engineering approaches, especially for acquisition, building on A-SQUARE. Ultimately we would like to develop a library of reusable security requirements that organizations could use as a starting point in their projects.

References

[1] Allen, J.; Barnum, S.; Ellison, R.; McGraw, G.; & Mead, N. R. *Software Security Engineering: A Guide for Project Managers*. Addison-Wesley Professional, 2008 (ISBN-13: 978-0-321-50917-8).

[2] Mead, N. R. Ch. 3, "Identifying Security Requirements Using the SQUARE Method," 44–69. *Integrating Security and Software Engineering: Advances and Future Visions* H. Mouratidis & P. Giorgini. Idea Group, 2006 (ISBN: 1-59904-147-2).

[3] Bijwe, A., Mead, N.R. Faculty Advisor, *Adapting the Square Process for Privacy Requirements Engineering*, CMU/SEI-2010-TN-022 Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2010.

[4] Abu-Nimeh, S. & Mead, N.R. "Combining Privacy and Security Risk Assessment in Security Quality Requirements Engineering.*"* Presented at the *Intelligent Privacy Management Symposium*. Stanford University, March 2010.

[5] (Invited Paper) Ingalsbe, J.A.; Shoemaker, D.; Mead, N.R. & Drommi, A. "Threat Modeling the Enterprise." *Journal of Information Systems Security (JISSec) Volume 5*, Number 3 (2009): 42-57.

Using Trusted Hardware as a Foundation for Cyber Security

Principal Investigators:
Archie Andrews, David Fisher, and Howard Lipson
CMU CyLab Collaborators: Anupam Datta and Jonathan McCune

Problem Addressed

The Software Engineering Institute's (SEI) stakeholders are increasingly operating in what we term extreme adversarial environments. These environments contain highly resourced adversaries such as nation-states, well-funded terrorist organizations, and large criminal organizations operating under the cover of a continuous barrage of amateur attacks against potential vulnerabilities. Operating securely in such environments requires that we understand the range of characteristics of those environments and assess feasible approaches to maintaining security and survivability properties [1] under the severe condition these environments impose. We believe that a successful approach to operating securely in extreme adversarial environments will ultimately depend upon obtaining and leveraging a better understanding of the relationships among hardware, software, security, and trust.

In FY 2010, this research evaluated the capabilities and limitations of using trusted hardware as a foundation for achieving demonstrably high assurance of end-to-end security properties of applications executing in extreme adversarial environments.

Research Approach

Hardware-based assurance of trust is an emerging area of research and open standards development. The Trusted Computing Group (TCG) is an international industry standards group that has developed and defined open standards for hardware-enabled trusted computing. It produced and continues to evolve the standard (TPM v1.2, ISO/IEC standard 11889, dated August 2009) for the Trusted Platform Module (TPM) chip, which provides for secure storage and generation of cryptographic keys, platform authentication (based on each TPM having a unique cryptographic key), and remote attestation (the ability to assure a third party of some aspects of its trustworthy status, such as its TPM identity and software configuration). The TPM's hashing function creates a cryptographic digest (i.e., a nearly unforgeable summary) of code or documents which can enable the detection of tampering by comparing a current digest to a previous (known good) instance of the digest to see if anything has changed.

There is an emerging area of research on using TPMs (or other hardware support) as anchors of trust upon which to build attestations about the security properties of systems [2, 3]. TPM-enabled computing devices provide a set of initial building blocks for leveraging hardware to improve security. In FY 2010, our research focused on evaluating the capabilities and limitations of applying the hardware features of the TPM to improve security as an important first step toward building trustworthy infrastructures and applications for secure and survivable operation in extreme adversarial environments.

Expected Benefits

The benefits of this research will be an improved understanding and approach to establishing and maintaining trust in extreme adversarial environments so that mission-critical applications operating in such environments can do so with higher assurance of achieving their desired security and survivability properties. Trust is of critical importance in all human activity. With unjustified trust, nothing can be adequately assured, safe, or efficient. Automated systems and networks impose additional problems of trust, but they do not traditionally provide adequate support for trust. Automated support for trust is essential for effective use of automated systems and networks. The TPM and other hardware-based trust mechanisms are a step in the right direction but inadequate in current practice. While they provide automated support for certain security aspects of trust, issues of trust go far beyond security. There is a need for investigation and understanding of the potential role of technology in supporting all aspects of trust in mission-critical software and systems.

2010 Accomplishments

This research in trusted computing accomplished the following in FY 2010:

- characterized the properties of extreme adversarial environments

- determined, characterized, and described the capabilities and limitations of a hardware-based trusted computing platform (specifically, the TPM) for improving security and survivability in extreme adversarial environments

- laid the groundwork for future research that will explore and exploit the concepts of trust and trustworthiness and provide a scientific basis for understanding the relationships among hardware, software, security, and trust

Extreme Adversarial Environment

We found that it was impractical to partition the potential activities of an extreme adversary into distinct classes of adversarial environments, such that specific security approaches would be effective for specific classes of these environments. We were compelled to conclude that in the case of a determined adversary with extreme skills

and resources, no reasonable application of the traditional security model (i.e., the "fortress model") could provide an impregnable defense. The fortress model assumes that the defender can successfully construct a defense that completely separates critical operations from the persistent attacker and successfully defend against every attack. A more extreme but more realistic assumption is that operation in a malicious environment is inevitable. Given that assumption, the question becomes whether it is feasible to isolate critical operations from that malicious environment to the extent necessary to enable trusted operations. The team therefore turned its research to the consideration of what degree of trust in the isolation of critical operations is enabled by a hardware-based trusted computing approach that relies on a trusted security platform.

Trusted Security Platform

A trusted security platform is one in which confidence in the secured operations, applications integrity, and isolation from malicious actions is enabled by an embedded Trusted Platform Module (TPM). The TPM is a simple, passive, integrated circuit chip (see Figure 1) intended to serve as a hardware root of trust for trusted infrastructure leading to trusted software applications. It is an inexpensive enabler for credential storage, device identity, and trusted applications and provides security capabilities that cannot be provided by software alone with the same degree of confidence. The trusted security platform and the TPM in particular offer a level of security capability unachievable in software alone. The TPM derives both trust and trustworthiness from its simple passive character. A more complex device, an active one with general-purpose computing capabilities, or one that shares registers with a CPU could not engender similar levels of trust.
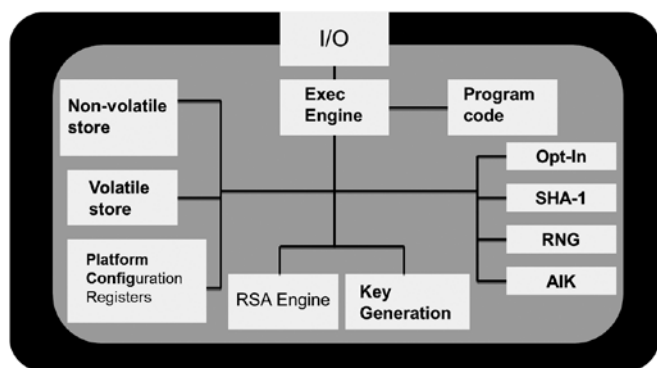

Figure 1: Components of a Trusted Platform Module

Despite these desirable attributes, in practice the TPM is rarely used even when the chip is present.[1] It is rarely used in the development of computational or network infrastructure and has had limited market penetration in end-user applications. To understand why, we examined the methods recommended and used for extending trust from the TPM to higher levels of application software. Building provably correct chains of trust starting from the BIOS and extending through several levels

of operating system and application software is theoretically sound [4], but in practice it is a tedious and brittle process that must be redone whenever any software along the way changes [5]. More practical methods are available (such as a Dynamic Root of Trust [6]) but with increased vulnerabilities and often lower levels of trust. The measurement technique used in conjunction with the TPM is vulnerable because it cannot detect changes made and restored between measurements. Neither does measurement provide protection for mutable storage.

Implications for Trust

A viable business model for exploiting trusted platform technology in a general-purpose platform has not yet emerged. However, as the number of deployed TPMs increases, so do the opportunities for commercial products dependent on the existence of deployed TPMs. A key appears to be integration and support by operating systems and application use of the resulting application programming interfaces (APIs). As a practical matter, applications developers and end users will not leverage the TPM unless its functionality is easily accessible. End users cannot be expected to develop the chains of trusted software required at the operating system level. More encouraging is the increased use of TPM chips in certain dedicated security applications such as Microsoft's BitLocker, storage of PKI private keys and other credentials (e.g., biometric identifiers), and potentially secure machine identities on sensitive networks. Part of the problem may be that without a large installed base of TPMs, there is little incentive for application developers, and without developers' demand, there is little incentive for the operating system (OS) to support them or for more systems to install them. The implication to the Department of Defense and others with a currently large installed base of TPMs is that the majority of those already deployed will unlikely be used without a critical mass of installations that triggers a broader market of applications and OS support.

As a foundation for improving the security of software-intensive systems, there is a need for hardware support beyond that provided by currently available trusted platform devices. The potential and realized benefits of the TPM derive from the ability to ensure integrity of information, processes, or identity either by physical isolation (e.g., key storage in the TPM) or logical isolation (e.g., encrypted communication). Hardware support for isolation of storage and communication is needed at many levels, including CPU register sharing through task switches, shared caches, uninitialized portions of memory pages, and direct memory access. Hardware could assist operating systems by isolating applications, eliminating security vulnerabilities inherent in current CPU architectures, and providing user-level security functions that are easily accessible through APIs. Hardware mechanisms to support dynamic roots of trust, to eliminate software intervention at the BIOS and the lowest levels of the OS, to provide immutable

storage as an alternative to measurement, and to facilitate logically atomic sequences of operations have the promise to make chains of trust less brittle and more trustworthy.

Our effort also looked at the character of trust and trustworthiness. The business model that has been successful for engaging the TPM is to provide end-to-end security products that happen to leverage TPMs. Other models have not been widely adopted and if fully realized would result only in security products and "feature sets," not necessarily in more secure or trustworthy applications. A business model of potentially greater effectiveness is the use of the TPM internal to applications, products, or systems to maintain and preserve trustworthiness that could otherwise be undermined by security flaws in lower-level infrastructure. The TPM could also be used to develop underlying infrastructures that are themselves more secure and trustworthy.

Future Goals

Support for trust will require more rigorous understanding of trust and trustworthiness, effective strategies for achieving trustworthiness and assessing trust, and development of automated tools for trust. Research in these areas will likely borrow from other domains with overlapping concerns. These domains include, most conspicuously, security, survivability, dependability, emergent behavior, and modeling and simulation.

In FY 2011, our research will focus on clarifying the principles of trust and then building on those principles to determine architectural characteristics and engineering methods necessary for building hardware-enabled, trustworthy software applications for networked, embedded systems.

Trust technology has the potential to overcome the limitations of existing approaches for achieving mission assurance. An effective trust technology will focus on mission fulfillment, survivability, and evolution of automated and networked systems. It will employ security methods when and where they are cost-effectively needed. It will emulate and adapt proven trust methods from everyday life. Our ultimate goal is to provide the capability to build and operate critical automated systems that will behave in a sufficiently trustworthy manner to consistently fulfill their missions, even when these systems are built and operated in extreme adversarial environments.

References

[1] H. Lipson and D. Fisher, "Survivability—A new technical and business perspective on security," in *Proc. 1999 Workshop New Security Paradigms,* Caledon Hills, Ontario, Canada, 1999, pp. 33-39. *http://www.cert.org/archive/pdf/busperspec.pdf*

[2] A. Datta, J. Franklin, D. Garg, and D. Kaynar, "A Logic of secure systems and its application to trusted computing," in *Proc. 30th IEEE Symposium Security and Privacy,* Oakland, CA, 2009, pp. 221-236.

[3] J. M. McCune, N. Qu, Y. Li, A. Datta, V. Gligor, and A. Perrig, "Efficient TCB reduction and attestation," CMU, Pittsburgh, PA, CMU-CyLab-09-003, 2009. *http://www.cylab. cmu.edu/research/techreports/2009/tr_cylab09003.html*

[4] W. A. Arbaugh, D. J. Farber, and J. M. Smith, "A secure and reliable bootstrap architecture," *IEEE Symposium on Security and Privacy,* 1997, pp. 65-71.

[5] B. Parno, J. M. McCune, and A. Perrig, "Bootstrapping trust in commodity computers," *IEEE Symposium on Security and Privacy,* 2010, pp. 414-429.

[6] J. M. McCune et al., "Flicker: An execution infrastructure for TCB minimization," in *Proc. 3rd ACM SIGOPS/EuroSys European Conference Computer Systems,* Glasgow, Scotland, U.K., 2008, pp. 315-328.

Trusted Computing in Embedded Systems Workshop

Sponsored by the Office of the Secretary of Defense, Director of Defense Research and Engineering (ODDR&E)
Nov. 2010 | Pittsburgh, PA

Co-Chairs: Archie Andrews, CERT; Jon McCune, CMU, CyLab. Committee: David Fisher, CERT; Virgil Gligor, CMU, CyLab; Howard Lipson, CERT; Adrian Perrig, CMU, CyLab

The workshop addressed the capabilities and limitations of employing trusted hardware-enabled components in embedded systems and discussed the research necessary to advance the field. It drew government, academia, and industry attendees; focused on embedded systems; and incorporated related disciplines. The workshop addressed new R&D and methods for enabling trust in embedded systems, lessons from R&D projects on embedded systems security, and gaps in current research. Findings include the following:

- Research is needed on isolation and memory management methods to improve the level of trust feasible in 8-bit microcontroller systems.
- Coherent trust models based on sound criteria and principles for trust are necessary to support both acquirers and researchers.
- A reference implementation of end-to-end use of trusted computing in an embedded system would be valuable for exploring the characteristics of such a system and understanding necessary limitations.
- Tool support is necessary to help embedded development and design teams incorporate security and trust into their requirements, specifications, designs, and implementations.
- Building a community that spans the safety, security, dependability, and trust communities requires research to identify the points of intersection and diversion, and it encourages sponsorship and leadership.
- Research is needed to appreciate the relationships inherent in cyber-physical systems and their impacts on trust.

Analysis of Catastrophic Failures

Principal Investigators:
Carol Woody and Linda Levine

As technology increases in pervasiveness and capability, we are increasing our dependency on it and reducing our ability to function in its absence. When catastrophes occur, what key dynamics take place and how do these dynamics affect our technical infrastructure? How can we understand complex failure in order to better build and operate future technologies and address today's complex, software-dependent networked systems? Addressing these questions requires analysis of multiple catastrophes, identification of patterns of value to future technology solutions, and exposure to opportunities for improvement. As software plays an increasing role in the functioning of normal operations, it is also becoming a growing component of catastrophes.

Understanding how software fails is a necessary element for establishing assurance.[1] Analysis of failures has been successfully used to identify and address vulnerabilities and weaknesses in code. We propose that the focus should be expanded to understand how highly interconnected software and systems fail. This expanded focus will help us understand the problems of

tomorrow as organizations adopt new infrastructures such as cloud computing, extended use of mobile devices, and total dependency on supply chains in the development and support of software. The best way to learn is to have things go wrong, and in catastrophes many things go wrong. With the growing complexity of highly networked systems and software, it is insufficient to consider them individually even though that is the way they are constructed and tested. The best examples of failures of these highly complex systems can be found in recent catastrophes.

Software engineering for safety has benefited extensively from studies of safety failures evidenced in disasters. Leveson writes, "The high frequency of accidents having complex causes probably result [sic] from the fact that competent engineering and organizational structures eliminate the simpler causes. On the positive side, the very complexity of accident processes means that there may be many opportunities to intervene or interrupt them. Therefore, thorough consideration of the conditions leading to accidents will be more useful than simplistic explanations" [1, p. 48].

Today, each individual catastrophe is studied independently to seek accountability and to identify successful short-term responses. The result is that systemic problems, within and across disasters, continue to worsen and feed future catastrophes. The absence of effective analysis has been reported in the many documents we have read for 9/11, Hurricane Katrina, Lockerbie, and other disasters. We have not found evidence of systematic mining of catastrophic failure examples to improve how we build systems and software.

In a preliminary analysis [2], we studied two cases—Hurricane Katrina [3] and 9/11 [4]—representing threats from natural

---

[1] Assurance is defined as the justified confidence that a system functions as intended and is free of exploitable vulnerabilities, either intentionally or unintentionally designed or inserted as part of the system at any time during the life cycle [5].

forces and terrorism. The lens used for this analysis was the Generic Error Modeling System (GEMS). The GEMS framework helps us to understand types of errors that occur in operational situations and distinguishes between skill, rule, and knowledge-based modes. Through this analysis we identified the following three issues.

1. Technology plays a problematic role given its fragility and dominance. We are dependent on technology and will become even more reliant on it in the future, yet the interconnectedness that it supports is very fragile. When technology fails us, we often do not have an alternative or substitute. For example, when the communication among first responders at the World Trade Center failed, no alternative actions were identified and implemented.

2. There is a coordination and centralization effect. Disaster response requires coordination among multiple organizations and roles. Decision making must be distributed and coordinated so that dependencies are understood and managed as a system of systems. For example, in Hurricane Katrina, the different organizations responsible for the levees created a patchwork quilt of ownership, and no entity (or shared standard) watched over the integrity of the whole system.

3. There is a failure to consider failure. Whether for social, political, or technical reasons, we have not adequately addressed failure modes and conditions. We have not distinguished, for example, between routine (or "normal") crises and unprecedented failure. Our plans assume that, as the saying goes, "failure is not an option."

Technology is too frequently built and validated to idealized operational conditions. The result is operational failure or extensive and expensive rework before implementation. Analysis of catastrophe helps us to see how complex, multi-system, multi-organizational environments fail as well as what works and what doesn't. Through the use of perspectives such as GEMS, our identification of patterns can be improved. And with improved pattern recognition, we can develop better ways to respond to catastrophes. We have just begun to apply this technique, and extensive work remains to be done with additional cases and perspectives.

References

[1] Leveson, Nancy G. *Safeware: System Safety and Computers.* Addison-Wesley, 1995.

[2] Levine, Linda & Woody, Carol. System of Systems Analysis of Catastrophic Events: A Preliminary Investigation of Unprecedented Scenarios. *IEEE International Conference on Technologies for Homeland Security.* Waltham, MA, November 8-11, 2010.*http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5654962&isnumber=5654927*

[3] Davis, Tom, chair. *A Failure of Initiative: Final Report of the Select Bipartisan Committee to Investigate the Preparation For and Response to Hurricane Katrina.* U.S. House of Representatives. US Government Printing Office, 2006.

[4] Kean, Thomas, chair. The Complete Investigation: *9/11 Report with Commentary by The New York Times.* National Commission on Terrorist Attacks upon the United States. St. Martin's Press, 2004.

[5] National Defense Industrial Association System Assurance Committee. *Engineering for System Assurance.* *http://www.acq.osd.mil/se/docs/SA-Guidebook-v1-Oct2008pdf (2008)*

Complexity Modeling and Analysis

Principal Investigators:
Christopher Alberts, Lisa Brownsword,
Andrew Moore, and Carol Woody

Our operational systems are increasingly interconnected within and across enterprise boundaries. In tandem, threats to our operational systems are more sophisticated, leveraging technical, organizational, and personal weaknesses. When addressing software assurance for today's complex ecosystems, policy makers, technology solution providers, and solution adopters should recognize the systemic nature of the operational and software assurance environments. Our assurance modeling framework provides a way to look across the assurance ecosystem to examine the gaps, barriers, and incentives that affect the formation, adoption, and use of assurance solutions. This framework lays important groundwork by (1) understanding the relationships between organizations and assurance solutions and how these relationships contribute to (and hinder) operational assurance and (2) identifying potential areas for improvement across a spectrum of technical and organizational areas.

Using Modeling to Understand a Software Assurance Ecosystem

Technology is too frequently the primary focus for operational assurance of highly interconnected, rapidly changing systems. Improving software assurance for these systems will require broad adoption of new types of assurance solutions.[1] However, forming those solutions requires a way to understand the complexities of the assurance ecosystem. The assurance ecosystem describes the range of interrelated elements that influence operational

assurance, including organizations, people, policies, practices, and technologies.

Modeling offers a means to structure, describe, analyze, and discuss those complexities. It provides a way to describe organizational behavior and social and technical elements that must work together to achieve results—a collaboration among solutions and participants.

The SEI Assurance Modeling Framework

The SEI has developed and piloted an Assurance Modeling Framework (see Figure 1) that provides a way to look across the assurance ecosystem and examine the gaps, barriers, and incentives that affect the formation, adoption, and usage of assurance solutions. The framework characterizes

- the current portfolio of organizations working in assurance

- assurance solutions (including those planned, funded, developed, and used)

- the interrelationships among organizations and assurance solutions

- the relative contributions of organizations and solutions to operational assurance

- future trends and their potential impacts on operational assurance

This modeling framework provides an approach for systematically assembling and analyzing the required information within an *assurance capability area*.[2] The general structure of the framework is shown in Figure 1. The modeling framework is composed of multiple *activity categories* (indicated by rounded rectangles). Each activity category provides insights on one or more than one of the framework information questions and produces one or more *views* (indicated by rounded capsules). Each view is a collection of models and data formed using one or more *methods* (indicated by rectangles). A *profile* is a set of views that collectively describe the relevant elements of the assurance ecosystem landscape for the selected assurance capability area.

Pilot use of the framework focused on the vulnerability management assurance capability area [1]. Vulnerability management is concerned with the prevention, discovery, and correction of vulnerabilities in systems, software, and networks. We selected two assurance solutions from the cluster of technologies related to vulnerability management: Common Vulnerabilities and Exposures (CVE) and Common Weakness Enumeration (CWE). From analysis of the assurance profile for CVE and CWE, we identified inefficiencies and candidate improvements for assurance adoption, including the following:

---

[1] An assurance solution is a policy, practice, or technology that contributes to system assurance. Assurance solutions have traditionally been developed for structured, tightly controlled operational environments with limited consideration of the complexity of their operational contexts.

[2] An assurance capability area is a set of related activities used to achieve an assurance property. Vulnerability management, incident management, and threat analysis are examples of assurance capability areas for the security assurance property.

- **The majority of the assurance solutions and participating organizations do not directly support operational assurance.** Their focus is on the products themselves; the connections to operational assurance are assumed. We identified limited support for the operational side where several roles rely primarily on manual or homegrown approaches.

- **Solution suppliers are motivated to identify and produce vulnerability patches quickly, but operational organizations are motivated to maintain system availability and responsiveness.** Operational organizations relate to the value of assurance solutions based on system availability, not timely patches.

- **There are important dynamics between the reactive and proactive responses to vulnerability management that affect the formation and adoption of assurance solutions.** The structure of the behavioral feedback suggests that a balancing point between proactive vulnerability prevention practices and reactive patch generation and release is needed to address immediate and longer term operational assurance.

- **Understanding the similarities and differences in user communities for seemingly similar assurance solutions can be critical to the successful adoption and usage of assurance solutions.** Each of the user communities for CVE (reactive) and CWE (proactive) applies different terminology, communication sources, and priorities.

Applying the framework to analyze the assurance needs for the prevention and management of malicious attacks is underway. The ability of successful attacks to hamper operational processes is increasing, and the potential for catastrophic impact is looming. Attackers use technology, people, and typical organizational implementation choices, such as password controls, trusted links, and standards, to set up successful attacks. As organizations change their response to thwart these, attackers identify new opportunities using many of the same tools and techniques applied by defending organizations. News sources[3] indicate that spending for security will continue to increase, but there is limited evidence that this investment addresses the real problems. There are numerous practices and technology tools for use, but where should an organization invest for effective operational assurance? Policy makers, program managers, and IT professionals need a way to analyze their problem space to explore the effectiveness of options prior to making choices.

The assurance modeling framework provides a way to look across the assurance ecosystem to tie the current environment to operational needs and identify ways in which policy, practices, and technology options can be applied to improve assurance. The assurance modeling framework focuses on describing the complexities of the interplay between technical and social elements of the assurance ecosystem. The range of methods used in the framework provides the data and perspectives required to balance crisis reaction with prevention to inform new or revised policies, practices, and organizational structures.

References

[1] Brownsword, L., Woody, C., Alberts, C. & Moore, A. *A Framework for Modeling the Software Assurance Ecosystem: Insights from the Software Assurance Landscape Project* (CMU/SEI-2010-TR-028). Software Engineering Institute, Carnegie Mellon University, 2010. http://www.sei.cmu.edu/library/abstracts/reports/10tr028.cfm
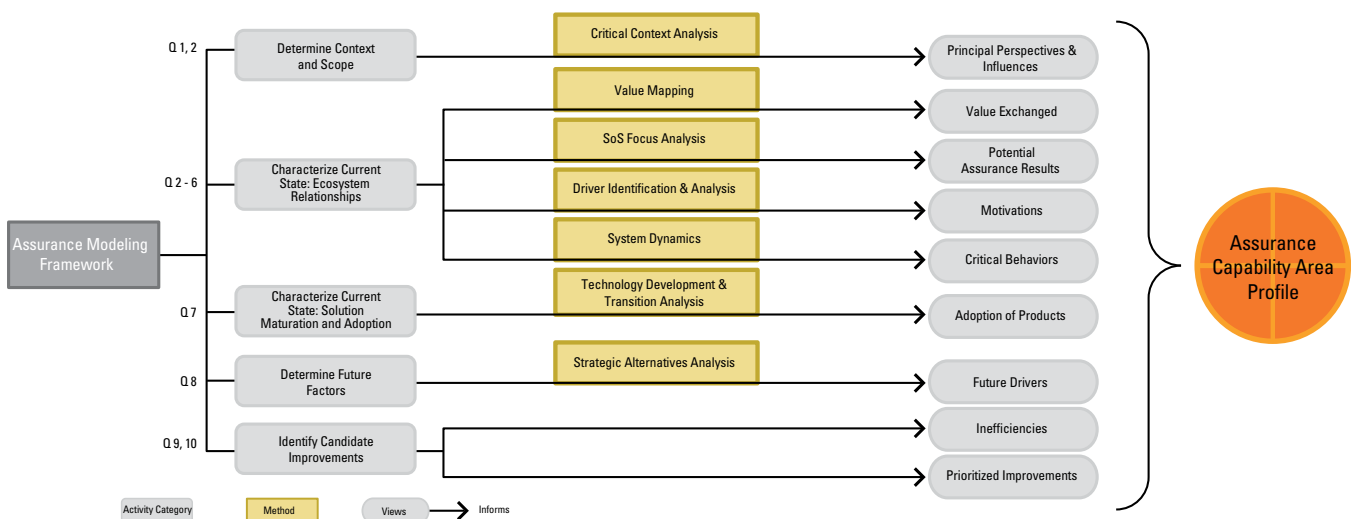
Figure 1: SEI Assurance Modeling Framework

---

[3] http://www.forrester.com/ER/Press/Release/0,1769,1320,00.html

Science of Cyber Security

"Research without operations is irrelevant. Operations without research are incompetent."

—Tim Shimeall

DIGITAL INTELLIGENCE AND INVESTIGATION DIRECTORATE

MALICIOUS CODE RESEARCH AND DEVELOPMENT

INCIDENT RESPONSE

NETWORK SITUATIONAL AWARENESS

RESILIENCE MODELING AND ANALYSIS

WORKFORCE DEVELOPMENT

Science of Cyber Security

Science of Cyber Security, published by JASON, an independent scientific advisory group that provides consulting services to the U.S. government on matters of defense science and technology, provides a fascinating look at how data is the cornerstone of cyber security research.

Cyber security threats are dynamic; attackers and their agendas are always changing, and the type of attacks evolves as attackers become familiar with defensive strategies. Consequently, there is currently no one science, including computer science, that covers all the issues related to cyber security.

The report concludes that the most important scientific attributes needed to help the field of cyber security progress are a common language and a set of basic concepts through which the security community can develop a shared understanding. The common language and basic concepts—indeed, all of science—are rooted in data.

This focus on data is the heart of the CERT research program. CERT was built from researching the data underlying security failures following the Morris worm incident, which brought ten percent of internet systems to a halt in November 1988. After this incident, the Defense Advanced Research Projects Agency (DARPA) charged the SEI with setting up a center to coordinate communication among experts during security emergencies and to help prevent future incidents, and CERT was created.

Data continues to influence how CERT looks at research problems and guides our perspective on research. CERT researchers have access to two decades of security data that help to solve security problems. In this way, CERT is contributing to a science of cyber security.

Another tie between CERT and the *Science of Cyber Security* report is the transition of CERT research into practice. A key recommendation from *Science of Cyber Security* is transitioning research results into tools available to developers.

CERT researchers focus on applied research. Researchers first develop and test new concepts and theories and then transition them into practice with both government and corporate customers. Read about specific examples of our applied research in the following sections.

Through our data-driven, applied research, CERT provides a practical perspective to creating a science of cyber security and, in doing so, helps to advance the field of cyber security.

CERT Research by the Numbers

- Insider Threat: 550+ cases in the Insider Threat database
- Secure Coding: 89 rules and 132 recommendations in the CERT C Secure Coding Standard
- Software Security Assurance: 2,439,889 unique visitors to the Build Security In website
- Digital Intelligence and Investigation Directorate: 140TB of active case data and 40TB of research data
- Malicious Code Research and Development: 17 million artifacts in the malware catalog
- Incident Response: 544 vulnerabilities catalogued
- Network Situational Awareness: 10 billion flow records collected per day on the two largest networks
- Resilience Modeling and Analysis: 50 resilience-focused assessments performed
- Workforce Development: 124 Government Forum of Incident Response and Security Teams (GFIRST) personnel trained through the XNET environment

"The three largest cases prosecuted by the Department of Justice were direct results of our technology."

– Rich Nolan

Digital Intelligence and Investigation Directorate Overview

Current tools and processes are inadequate for responding to increasingly sophisticated attackers and cyber crimes. Staff members within the CERT Program's Digital Intelligence and Investigation Directorate (DIID) are addressing that problem by developing technologies, capabilities, and practices that organizations can use to develop incident response capabilities and facilitate incident investigations. They are also developing advanced tools and techniques to address gaps that are not covered by existing resources.

DIID's work focuses on understanding digital information—developing methods for extracting information to support cyber crime investigations, gaining insights into attackers' techniques, and identifying trends. Traditional computer forensics focuses on analyzing digital evidence after an incident. While computer forensics is a component of our work, the DIID team is exploring broader implications and applications.

A primary focus of DIID's research is to explore methods for handling increasingly large data sets. Human-centric approaches cannot process large amounts of data quickly, if at all. To investigate alternatives, the DIID team is collaborating with researchers at Carnegie Mellon University in departments such as the School of Computer Science, the Robotics Institute, the department of Electrical and Computer Engineering, the CyLab Biometrics Center, and the Visual Intelligence Studio. Through these collaborations, DIID has been able to convert innovative, purely theoretical concepts into actual products and capabilities that can be used in digital investigations.

DIID's domain experience allows the team to identify specific, tactical problems and develop prototype solutions. They have established relationships with academia, industry, and law enforcement, so they are in a unique position to identify emerging challenges in the field of digital investigations and collaborate on solutions to meet those challenges. After removing any classified information to isolate the core issue, they consult with research resources available from the campus community to create solutions. By concentrating on the core issues, they are able to create solutions that not only apply to the specific cases but can be amplified into a much broader solution for similar cases.

The DIID team also takes an agile approach to development and relies on rapid prototyping. Taking time to test a solution under every possible condition before transitioning it to the community is impractical. They strive to improve an organization's capability as quickly as possible; a working solution that is still being honed may present a significant improvement to an organization with no capability.

Digital Intelligence and Investigation Directorate Customer Spotlight

In February 2010, Max Butler was sentenced to 13 years in prison for stealing millions of credit card numbers. In March 2010, Albert Gonzalez was sentenced to 20 years in prison for his involvement in the highly publicized case involving the TJX company. Both investigations and subsequent prosecutions relied on tools, techniques, and practices developed by DIID staff.

Representatives from the United States Secret Service requested DIID's help because they are familiar with our tools and capabilities. For years, DIID has worked closely with the law enforcement and intelligence communities to understand their challenges and develop solutions. The DIID team has developed a variety of resources for these communities, including tools such as LiveView, CryptHunter, and Forensic Appliance that authorized members of government and military organizations can access online. In 2010, there were 18,000 downloads of these tools.

MALICIOUS CODE

RESEARCH
AND
DEVELOPMENT

WE'RE WORKING TO IDENTIFY PATTERNS OF INTRUDER BEHAVIOR, THE TECHNIQUES USED BY INTRUDERS, WAYS INTRUDERS OBFUSCATE MALWARE, ETC.

—CORY COHEN

Malicious Code Research and Development Overview

Understanding Malware Data on a Vast Scale

For more than a decade, Cory Cohen and his colleagues in CERT Malicious Code have been collecting malware files from around the world. This has resulted in a database of enormous size: the Artifact Catalog. The catalog is a unique and powerful resource with which to study important aspects of the malware problem. The Artifact Catalog's massive scale, however, presents tough challenges to the CERT malware researchers.

"Our mission is to understand the malware in the Artifact Catalog," says Cohen, leader of the Research and Development (R&D) team. The R&D team supports the CERT mission by examining malware from the perspective of the entire catalog. Specifically, the team works to answer the following questions:

• What does it mean for files to be similar?

• What are the important features of malware?

• How to automatically extract those features?

• How to quickly and efficiently use what we learned for future analysis?

"We're working to identify patterns of intruder behavior, the techniques used by intruders, ways intruders obfuscate malware, etc.," adds Cohen. "This is a daunting job because the Artifact Catalog is so large. Today, we understand only a small percentage of the catalog. But we're making progress."

Cohen explains that tackling the issue of the Artifact Catalog's scale is essential for further progress. "Human malware analysis can provide deep insights, but it takes a long time to develop them," he notes. "We're working on ways to leverage that human analysis using automated methods, without completely replacing it. Automation means working at scale, and our scale is large." This forces the team to also work on many traditional research topics like data mining, algorithmic complexity, and large scale databases.

Challenges of Scale Spur Innovative Approaches

For small collections of files, it is possible to determine similarity by comparing files side by side. But with millions of files, this approach is not practical. "This would require an impractical amount of computing resources," notes Cohen. "So, we set about identifying which techniques could be automated in a way that would scale. For instance, an approach such as section hashing has greatly reduced the number of interesting files." Hash values of entire files can determine if two files are identical but cannot determine if they are merely similar because even a single byte difference will result in different hash values.

The R&D team is looking at decomposition approaches to hash smaller and smaller parts of the files. This works well for files with a structure, like Windows Portable Executable (PE) files which contain executable code that can be divided into functions or Adobe Acrobat documents containing streams of data. Cohen explains, "Function hashes show great promise because when two files contain the same functions, they might be similar or related, and that can be interesting."

Early Analysis Nets a Fresh Perspective

In terms of absolute numbers, the problem of malware has grown continually worse over time. However, Cohen observes that one of the most surprising and interesting findings of the Malicious Code R&D team's research is data that suggest not everything is gloom and doom. "If you simply count the number of malware attacks, and the number of individual instances of malware code used, we see a pattern in which the malware problem is growing worse over time," says Cohen. "But, when we apply our techniques to examine the data qualitatively and filter out the duplicate malware files, the picture is much different. What we've found is that the most prevalent families of malware are extremely prevalent. Setting those aside allows us to get a better picture." This picture shows a nearly flat, slightly upward trend for new malware families. The greatly reduced size of the data makes it easier to find correlations with malware files that pose the greatest threat to our national defense and critical infrastructure. "The fact that the trend, after filtering, is much flatter than expected is cause for hope," says Cohen. "The problem is still big, but it's moving in the right direction."

Current Research Efforts

The large set of malware data contained in the Artifact Catalog positions CERT to play a leading role in understanding the threats posed by malware. Complementing the rich data resources that CERT has are its operational malware analysis capability and its established working relationships with key sponsors and stakeholders. "These relationships help us understand the needs of the sponsors' analysts so we can tailor our research to meet the real-world problems confronting them," says Cohen. "We feel this makes our research much more relevant. We're working on real problems."

The articles that follow describe the research underway to address these real-world problems. "Malware Family Analysis" by William Casey, Cory Cohen, and Charles Hines describes work that centers on techniques for correlating runtime behavior with other static analysis techniques, such as section hashing and function hashing. This technique is enabling CERT researchers to learn new things about the malware and advances our goal for less expensive, more streamlined malware analysis that addresses the problems of the greatest interest to our sponsors. David French's article, "Beyond Section Hashing," describes efforts to assess the similarity of executable files. These efforts relate to reducing the vast amount of data in the Artifact Catalog to the data of interest, to reducing duplication of effort, to leveraging existing analysis, and to identifying malware by section. Finally, Jeffrey Havrilla is applying the lessons learned from analyzing malware code to the problem of malicious PDF files. His article, "Large-Scale Analysis of Malicious PDF Documents," describes the benefits of this approach and suggests its application to other non-executables.

Malware Family Analysis: Correlating Static Features and Dynamic Characteristics on Large-Scale Projects

Principal Investigators:
William Casey, Cory F. Cohen, and Charles Hines

Problem Addressed

Malicious software, or "malware," is software with malicious or criminal design goals. Examples include programs designed to pilfer personal or sensitive business information or to render computer equipment inoperable.

Identifying and characterizing malware is a dynamic and challenging problem, especially within the context of complex hardware and software systems.

Further compounding the difficulty is the adversarial use of obfuscation techniques that allow malware to evade detection and characterization.

These developments have necessitated novel and scalable techniques to identify and characterize malware families.

Research Approach

In general a malware family's structure is challenging to analyze. Robust techniques that involve multiple paths of data analysis and the fusion of these analysis components are more likely to yield useful knowledge. While each mode of analysis (i.e., static and dynamic) provides informative views into the data, we are developing techniques based on correlation between the separate modes to glean further insights into the family structure of malware. While these data fusion techniques are currently undergoing rapid development, they have already proven to be robust, dynamic, and capable of providing insight into the commonality and variability of large malware families.

This report demonstrates one recently developed method of deriving a behavioral map of data by linking dynamic characteristics to static features in the malware data. Further, this methodology is shown to scale to the *Zeus/Zbot* data set, which the cybersecurity community regards as a large and evolving malware family.

Static Features: Malware, like software in general, conforms to specifications to complete its process. Despite adversarial efforts to obfuscate code, malware's adherence to a process specification produces features that static analysis of binary code can reveal.

Our research group has organized static features into a hierarchy that represents scale and compositional relations.

For example, families of malware are composed of sets of files, files are further organized into sections, and code sections contain functions as a set of basic blocks, each of which are composed of sequences of code symbols that define runtime execution.
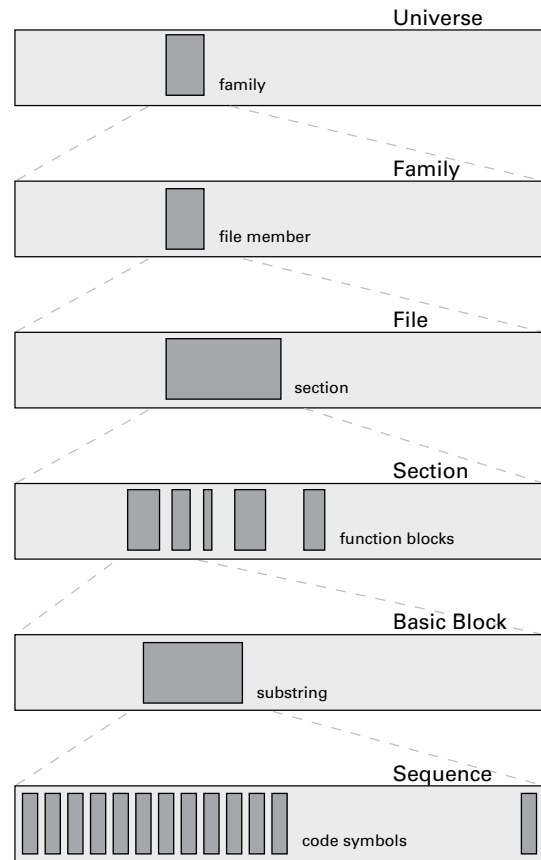

Figure 1: Organization and grouping of static features in software corpora

Dynamic Runtime Characteristics: Malware, when run as a process, affects a computer system by reading or writing files, opening communication pathways, modifying process images, and so on.

Dynamic analysis observes and records the effects of malware, whose distinct characteristics can each be associated to a binary program as a runtime characteristic.

Our runtime analysis focuses on selective characteristics that are highly specific to the operators of malware. These characteristics may differentiate one program from another within the same malware family and provide further knowledge into how this code is used and deployed.

For *Zeus/Zbot*, we focus on network communications observed during runtime, specifically on Internet Relay Chat (IRC) communication.

Problem Scale: The large size of contemporary malware families has motivated analysts to create methods that can scale to these large problems.

Table 1 shows the counts of various features of *Zeus/Zbot* that we obtained using our static and dynamic analysis modes.

Set	Set Description	Set Size
Z	Zeus binary files	61,889
$S(Z)$	Unique sections from Z	147,485
$F(Z)$	Unique functions from Z	686,899
$C(Z)$	Runtime characteristics from Z	81,840

Table 1: Scope and scale of Zeus/Zbot

Data Fusion: For our study of *Zeus/Zbot*, we sought to investigate how the static features could be linked to the dynamic characteristics and determine if these links could be used to identify locations in the family to focus analysis and reverse-engineering efforts. To explore this question further, we developed a quantity that indicates linkage between the static features and dynamic characteristics.

We began by specifying a runtime characteristic, c, which was fixed to a specific username in an IRC channel observed by our runtime analysis in multiple *Zeus/Zbot* files. The subset of *Zeus/Zbot* files in which this characteristic is expressed, called "the characteristic set," was denoted as Z_c.

We identified the set of all functions associated with the characteristic set as $F(Z_c)$. For each function f in $F(Z_c)$, we measured the specificity of this function to the characteristic set by forming a ratio comparing the frequency of occurrence of f in Z_c (the characteristic set) to the frequency of occurrence of f in Z (the *Zeus/Zbot* family at large). This comparison amounted to a weighing of the likelihoods, where the likelihood represents a probability that a function f occurs in any set S and is estimated as $|f(S)|/|S|$; the ratio of the number of files in S that contain f is normalized by the total number of files in S. We denoted the comparison as $r(c, f)$ and defined it in terms of the ratio weighing likelihoods:

$$r(c,f) = \frac{|f(Z_c)|/|\,Z_c\,|}{|f(Z)|/|Z|}$$

To understand how this ratio effectively identifies associations between static feature f (its position) and dynamic feature c, consider the following outcomes of the ratio:

- If the ratio $r(c, f)$ is less than one, then the function f occurs less frequently in the characteristic set than in the family at large.

- If the ratio $r(c, f)$ is equal to one, then the function f occurs with equal frequency in the characteristic set and in the family at large.

- If the ratio $r(c, f)$ is greater than one, then the function f occurs more frequently in the characteristic set than in the family at large.

The case where $r(c, f)$ is greater than one is operationally important because it provides leading indications of where to begin an analysis with limited resources.

Figure 2 shows the distribution of values for $r(c, f)$ plotted on a logarithmic scale. Here the characteristic c was the use of an IRC channel user name "VirUs," whose set Z_c had 13 files with 299 total unique functions in $F(Z_c)$, each of which produces a ratio plotted in log scale.

Expected Benefits

While the efficacy of this method for determining a cause-effect relationship between static features and dynamic characteristics has yet to be fully explored, the rank statistics


Figure 2: Significance ratios of functions of IRC channel use by user "VirUs." Data for 299 functions are plotted.

of $r(c, f)$ may prove highly relevant for analysis, possibly leading to a cause-effect finding. Further, the computation of $r(c, f)$ over the entire set of characteristics and functions observed in a malware family is within practical reach, which may lead to novel methods for exploring large malware data sets.

In general the results of this study are expected to benefit our overall ability to quickly identify binary program locations that are specific or significant to particular runtime behaviors. This ability impacts resource-constrained analysis projects that require developing new knowledge in large, uncharacterized sets of malware.

2010 Accomplishments

Studies on the *Aliser* and *Zeus/Zbot* malware families have improved our overall ability to characterize code in malware families. These studies explored several useful techniques for identifying code similarity in static features. Within the category of static features, we have primarily investigated association structure for files versus sections and files versus functions, leading to several findings [1, 2]. In addition we have expanded our techniques to include modes of analysis that blend static and dynamic techniques. We determined that the *Aliser* family was a "file infector" [1], and we reported several new findings on the *Zeus/Zbot* family [2].

These studies have also revealed more about the diversity in malware code. For example, we were able to identify characteristic features of the *Aliser* family in an edge-degree distribution of its file-versus-section association graph, which was later explained by *Aliser*'s file infection activity. These studies allow us to set expectations and understand the limits of what type of knowledge we may expect from each mode of our analysis. These studies have furthered our position that robust analysis techniques must include multiple analysis modes and that multimode correlation studies may prove to be important knowledge discovery tools.

Future Goals

We plan to further develop robust methodologies for the analysis of malware families, focusing on the following properties:

- *robust*—able to overcome missing data or failures in any one particular analysis pathway

- *dynamic*—able to adjust to evolving coding techniques of malware authors

- *high resolution*—provides meaningful summary of consensus and variation features across the malware family

- s*calable*—scales to large projects involving hundreds of thousands of files

We plan to further develop the use of rank statistics for behavior mapping of malware data. More generally we plan to explore the relations between analysis modes and how they may signal or produce useful knowledge to characterize malware families. In the coming year, we plan to further explore these issues in studies of additional malware families, including the *Poison Ivy* family.

We wish to go beyond analyses of particular malware families and understand just how varied and dynamic malware families can be as well as how various families relate and co-evolve in the Artifact Catalog data collection.

References

[1] W. Casey, C. Hines, D. French, C. Cohen, and J. Havrilla, "Application of Code Comparison Techniques Characterizing the Aliser Malware Family," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, 2010. Limited distribution report.

[2] W. Casey, C. Cohen, D. French, C. Hines, J. Havrilla, and R. Kinder, "Diversity Characteristics in the Zeus Family of Malware," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, 2010. Limited distribution report.

Beyond Section Hashing

Principal Investigator: David French

Problem Addressed

Human analysis of malicious code is a valuable but expensive process to which an analyst must bring an enormous amount of intuition, creativity, and experience. This process is complicated by the presence of human adversaries, who benefit directly from foiling efforts to analyze their malicious code. One way of reducing the time an analyst may spend on this task is by highlighting connections from a file currently being considered to one previously analyzed. Thus, the analyst may leverage past work and potentially reduce future effort. Therefore, methods by which similarity to existing files may be observed are of great interest in human analysis.

Over the past several years, the CERT Malicious Code team has developed a model for decomposing Portable Executable (PE) [1] files into well-defined, natural subcomponents, computing the hashes of these components, and then comparing those hashes across sets of files to observe similarity. Our previous research and decomposition model included entry-point clustering [2], section hashing [3], and function hashing [3]. These techniques have proved useful for guiding human analysts toward files that are potentially related to any particular investigation, primarily because analysts are able to reason about both the content of the decomposed data as well as its significance to their investigation. Given the success of these techniques in guiding human analysis, we consider whether these techniques are generally useful in analysis of large collections of malware.

Thus, the research problem for this report may be generally described as follows: given a particular feature decomposed from malicious files, can we reliably assert similarity or relation amongst all files that possess that feature? In this report, we concentrate specifically on section data. Sections are a unit of data storage defined within PE files, per Microsoft Corporation [1], and are used to aggregate data or code together.

Research Approach

We start with the assumption that any section hash collision between two files is considered significant. The rationale for this assumption is that for a section hash collision to occur the bytes from which the hashes were derived are almost always identical, representing the same byte values in the same order. Given this assumption, we have developed two different approaches for measuring the relationships between files based on the hashes of their section data. We selected these approaches because the data they produce may be created via automatic processes, rather than intense human analysis, and potentially offer a reduced analysis surface for observing relationships between files.

Composite Section Hashing

The first approach is called *composite section hashing*. In this approach, we treat the set of sections contained within a given file as a discrete set of quantized data, which may be extracted from and observed independently of the file from which the sections were derived. We consider the set of sections found in a particular file and attempt to normalize the observation of this set by applying a further transform. We use the *section hashes* (which are the MD5 checksums of individual sections in a PE file) as textual labels, apply a consistent ordering to these section hashes (which may be different from the order in which the original sections appeared in their PE file), and further compute the MD5 hash of this newly ordered set of section hashes. We describe the resulting MD5 hash as the *composite section hash* (or *CSH*) and assert that its value may be used as a label to describe a particular set of sections. We then assert that two files that share a composite section hash in fact share all the same section data. As sections are a reliable measure of the data actually loaded into memory and executed when a program is run, by observing that the sets of sections found in two files are identical, we may reasonably conclude that at least some portion of the runtime behavior of the two files is identical, and thus can establish a relationship between these two files based on static observations of their expected runtime behavior.

Section Clustering

The second approach is called *section clustering*. In this approach, we describe the connections between files and their sections as a bipartite graph of *file hashes* (the MD5 hash of an entire PE file) and *section hashes* (the MD5 hash of individual sections for each PE file). In this graph, an edge exists between a file hash and a section hash if a particular section is contained within a particular file. Given this graph, we may then observe relationships between sets of files based on the sections that the files have in common. We then observe *connected components* within this graph to derive sets of files that are related by their sections.

Expected Benefits

We expect that this research will demonstrate the extent to which section hashes may be used to automatically derive relationships between sets of files and will help us categorize malware by objective criteria such as section sharing. To the extent that this research is successful, we may automatically associate large numbers of files with each other. This represents concrete and defensible data for human analysts and allows them to rely upon automatic techniques to reduce their analysis surface. Achieving this goal has the direct benefit of saving expensive human analysis time in the cases for which the techniques we describe are successful.

2010 Accomplishments

Semantically Meaningless Sections

We have developed evidence that our initial assumption, which was that section hash collisions are always significant, is, in fact, false. We have identified two classes of section for which hash collision is most probably misleading or incorrect. Those two classes are the *empty hash* (which is observed by computing the MD5 of no input data) and *null hashes* (which are produced by computing the MD5 on sequences of bytes that are entirely value 0x00, or "null"). The empty hash has a well-known value, *d41d8cd98f00b204e9800998ecf8427e*, and is produced for sections with no on-disk content (for example, uninitialized data sections). Since there are no bytes to hash, asserting significance on collision of non-data is meaningless, and thus we exclude this case. Similarly, null hashes do not possess any significant semantic content and may accidentally arise from a number of methods (for example, data sections of null-initialized bytes padded out to the minimum section size). This lack of significant semantic content does not support assertions that collisions of null section hashes are significant.

Composite Section Hashing

Given our observations about empty and null hashes, we explicitly exclude such hashes from consideration when computing composite section hashes. When observing composite section hashes over all PE files in the Artifact Catalog (some 12,343,568 files as of this writing), we discover a total of 7,406,087 unique composite section hashes (CSHs). Of these 7.4 million unique CSHs, only 250,569 occur more than once, leaving 7,155,518 composite section hashes that derive from a unique file MD5. The remaining 5,188,050 unique file MD5s are thus created by the 250,569 duplicated composite section hashes. This means that approximately 3.38 percent of unique composite section hashes produce approximately 42 percent of all unique file MD5s in the Artifact Catalog.

We can further observe that a minority of CSHs produce the majority of unique file MD5s. The distribution for which each CSH produces a unique file MD5 seems to follow a power law distribution, as seen by the frequency graph in Figure 1. The x-axis represents distinct CSHs, and the y-axis represents how many unique file MD5s produce each CSH.
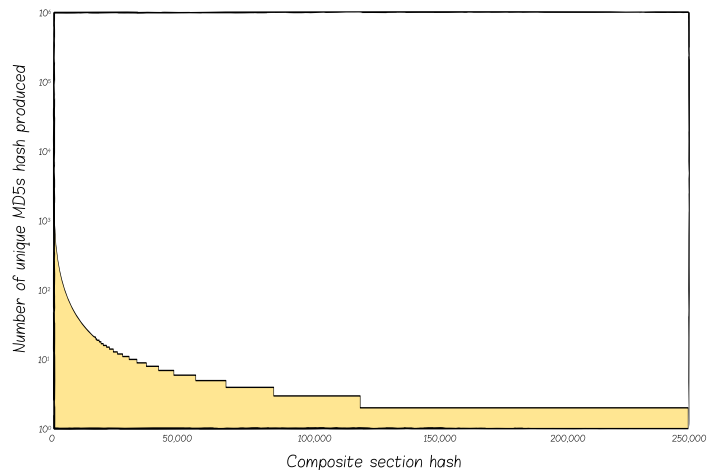

Figure 1: Composite section hash duplication rates

Understanding the extent to which a small number of files causes such a large increase in the potential analysis/comparison surface helps us infer motivation and methods of adversaries creating malicious files. We can reason in terms of economics and practical application about why malicious code authors may write their malware to support several tens of thousands of MD5-distinct copies. Files with different MD5s and identical sections can only vary in their header or in their slack. For example, such large numbers of files with identical section content might indicate an automatic process used to create minor variations of files, in order to foil antivirus detection. It may further indicate a large number of intended victims, where each file may contain unique target or unique command and control information.

Section Clustering

Having observed the degree to which certain types of files are strongly related, we may then expand our search to include all files connected by common sections. We apply a breadth-first search to the bipartite file-section graph and generate clusters of files (which are connected components in this graph) that share one or more sections. We again ignore empty and null hashes and exclude them from the bipartite graph. When observing section clusters for PE files in the Artifact Catalog (12,343,568 files as of this writing), we produce 373,522 distinct clusters of files, accounting for 9,766,425 PE files. These files represent 79.12 percent of all PE files in the catalog, leaving 2,577,143 files that do not share a section with any other files. Figure 2 shows an example of some of the clusters that are produced by this method (red shapes represent files, blue and black shapes represent sections, and green lines represent the edges).

We can observe additional properties of these clusters by examining their maximum depth. Of all the clusters, 356,986 connect their files by one or more section, with a maximum distance between any two files in a cluster of two edge traversals. These clusters account for 6,391,591 files (65.44 percent of all clustered files and 51.78 percent of all PE files in the Artifact Catalog). This is significant because it demonstrates that individual section hashes are

Figure 2: Visualization of several types of file-section clusters. Several types of subgraphs demonstrating types of relationships between file and section hashes. Red ovals indicate files, and blue/black squares indicate sections.

Future Goals

Several problems emerge when analyzing these file-section relationships. Sections that tie multiple files together are interesting for various reasons. Files connected by a section containing executable code are more significantly related than files connected by a section containing icons. Further, the presence of the mega-cluster of files defies analysis as to how arbitrary files within the cluster are related and represents our most significant challenge in clustering files based on sections. We can automatically observe these relationships, but we cannot automatically evaluate their significance. Thus, a future goal is to establish a weighting scheme by which the relationship established between files by a section may be interpreted as more or less significant, depending on the content of the section.

highly specific to sets of files. What is less clear is whether these files are actually the same malware. Since we cannot determine the content of a section based on its hash, we have no way of knowing whether the hash collision is meaningful to the behavior of the files in each cluster. These collisions could easily result from common sections that have low semantic value, such as imports or resources.

The remaining 3,374,834 files form clusters in which the distance between any two files (in terms of the number of edge traversals) may be greater than 2. There are 16,536 such clusters, including a single mega-cluster, comprising 1,457,902 files, whose maximum distance between any two files is 44 edge traversals. These clusters, in particular the mega-cluster, are much less significant because we cannot reliably understand how two arbitrary files in such clusters are related except by manually inspecting the files and hashes.

Using section clustering and composite section hashing, we can express a high-confidence relationship between files for more than half of the Artifact Catalog and can express a low-confidence relationship for a further quarter of the Catalog. Exploiting these relationships allows us to save analysts' time and reduce analysis surface. It also helps us place a more accurate upper bound on the number of truly distinct malware families represented by the 12 million files within the Artifact Catalog.

To achieve this, we must pursue a method of classifying sections as objectively as possible based on their actual content. This would allow us to identify additional classes of semantically meaningless sections (and avoid automatically relating files based on such section collisions) as well as express our confidence in file relationships based on semantics, rather than arbitrary byte sequences. Further, providing weights on the edges of the bipartite graph allows us to apply any of several well-known graph-cutting algorithms to the mega-cluster. In the case of malware, objectively identifying section content is stymied by the presence of obfuscation, malformation, and so on. Achieving this goal gives us additional insights into the specific modes of malware families and may recommend additional classifications of files based on their common characteristics.

References

[1] Microsoft Corporation, "*Microsoft portable executable and common object file format specification,*" *http://www. microsoft.com/whdc/system/platform/firmware/PECOFF.mspx*

[2] C. Cohen and J. Havrilla, "*Malware clustering based on entry points,*" *CERT Research Annual Report 2008*, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, 2009.

[3] C. Cohen and J. Havrilla, "*Function hashing for malicious code analysis,*" *CERT Research Annual Report 2009*, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, 2010.

Large-Scale Analysis of Malicious PDF Documents

Principal Investigator: Jeffrey S. Havrilla

Portable Document Format (PDF) files are not just static content. The versatility of the PDF format allows attackers to use a variety of features inherent in PDF documents. This exposes a large *attack surface* to many different digital weapons called *exploits*.

Antivirus names that label PDF exploits often inconsistently and insufficiently characterize most PDF malware in any meaningful way. To comprehend the malicious document phenomenon, we need to characterize the structure of the PDF artifacts in terms we can observe and measure. We gain more insight when we can identify natural file format boundaries to reason about relationships using standard analytical tools.

PDF documents can be broken down into a hierarchy of objects. Arbitrarily sized stream objects are particularly useful to attackers because they can be easily obfuscated and nested in ways that confuse many detection tools. We adapted a PDF parser tool to examine tens of thousands of malformed PDF files, without crashing, and reliably extracted their stream objects, uniquely identifying them using hashes. The result is a richer data set more suited to investigation using existing analysis tools and techniques.

Stream analysis shows a high degree of duplication of malicious, polymorphic PDF artifacts, likely caused by widespread use of exploit generator kits. Exploit generator kits are like cookie cutters that use slight variations in the decorations to lure victims, avoid antivirus detection, and exploit whatever vulnerabilities are useful in penetrating a target computer system. Kits are also useful in creating tables of pseudo-random, innocuous keywords to help PDF exploits avoid spam detectors and other malware detection tools. The challenge is separating the meaningful wheat (malicious code) from the distracting chaff (user content).

We extracted 90,000 PDF files from our malware corpus, yielding approximately 450,000 PDF stream objects. In the hash values generated, we found roughly 250,000 unique stream hash values, 4,400 of which (2 percent) connected over 95 percent of the existing PDF corpus (in general agreement with other published results [1]). We found the following types of data in the top 100 stream object hash values we manually inspected:

- payloads of binary *shellcode* often obfuscated to avoid easy detection [2]
- *JavaScript* used either to exploit a *vulnerability* or decode an obfuscated second stage used for *heap spraying*
- malformed images or font information used to exploit known PDF reader vulnerabilities
- apparently random words that may help confuse spam detection engines (English, Spanish, German, others)
- Adobe Extensible Metadata Platform (XMP) metadata
- PDF reader-specific text [e.g., *(Edited by Foxit Reader)*]
- dropped executable files and other PDF files
- other miscellaneous binary and text (e.g., page layout coordinates like "*0 0 595.28000 841.89000 re W n\r\n\r\n*")
- other degenerate objects (e.g., empty stream, white-space-only, etc.)

Correlation of the stream hashes was strongest in normalized de-obfuscated streams, but much work remains to understand the relationships. We gained one sample insight as a result of this work while recording an attack trend in an exploit generation kit called Neosploit in early 2010. Analysts noted the trend of PDF malware obfuscating itself by using PDF annotations in malicious documents to inhibit detection [3].

When looking at the ratio of incoming PDF stream objects correlated by hash value over time, we can see multiple instances of this exploit obfuscation technique being used (Figure 1). By isolating one particular instance, we can see evidence of the life cycle of a unique strain of PDF malware.


Figure 1: Ratio of collected getAnnots() stream obfuscation objects

Further analysis of malicious PDF documents will involve ongoing effort to peel away the latest exploit obfuscation techniques. Metadata embedded at both the document and object levels can be mined to tie individual attack trends to specific exploit generation kits. Correlating runtime behavior of PDF malware with such static attributes may result in new insight into how adversaries use this class of malicious code to perpetrate cyber attacks.

Some PDF software vendors such as Adobe have introduced new code execution sandboxes in their applications to thwart both known and future cyber attacks targeting their software [4]. We hope our effort to characterize our malware data will contribute to an emerging body of knowledge helpful in measuring their success and guide future improvements to more secure document standards [5, 6].

Glossary

Attack surface – The ratio of the subset of a computer system's resources reachable by an attacker versus the total number of functional subsystems.

Exploit – A hardware or software artifact or method designed to subvert the normal operation of a computer system by exploiting *vulnerabilities* or other security weaknesses.

Heap spraying – Filling the dynamically allocated area of an application's memory space with machine code (*shellcode*) at runtime, subsequently located and executed by a malicious program.

JavaScript – A dynamically interpreted scripting language with first-class functions used to manipulate document format and presentation in both network-oriented and file-oriented applications.

Shellcode – Machine code that can allow attackers to execute some functionality as if they were logged onto the local computer system. Shellcode is typically very small and designed to run in the same process as an exploited application.

Vulnerability – A design, implementation, or configuration error in a computer system that has the potential to be used by an attacker for some unintended security advantage (e.g., execute arbitrary commands or code, reveal secret information, or disrupt execution of a critical service or data operation).

References

[1] P. Baccas, "Malicious PDFs," presented at Virus Bulletin 2010, Vancouver, BC Canada, 2010. *http://www.virusbtn.com/pdf/conference_slides/2010/Baccas-VB2010.pdf*

[2] D. Stevens, "Malicious PDF documents explained," *IEEE Security & Privacy*, vol. 9, no. 1, pp. 80-82, Jan./Feb., 2011. *http://doi.ieeecomputersociety.org/10.1109/MSP.2011.14*

[3] J. Wolf, "PDF obfuscation using getAnnots()," *FireEye Malware Intelligence Lab Blog*, Jan. 14, 2010. *http://blog.fireeye.com/research/2010/01/pdf-obfuscation.html*

[4] K. Randolph, "Inside Adobe reader protected mode – Part 1 – Design," *Adobe Secure Software Engineering Team (ASSET) Blog*, Oct. 5, 2010. *http://blogs.adobe.com/asset/2010/10/inside-adobe-reader-protected-mode-part-1-design.html*

[5] J. Wolf, "Detailing the security risks in PDF standard," *Slashdot IT*, Jan. 2, 2011. *http://it.slashdot.org/story/11/01/02/0231242/Detailing-the-Security-Risks-In-PDF-Standard*

[6] A. Blonce, E. Filiol, and L. Frayssignes, "Portable Document Format (PDF) security analysis and malware threats," presented at Black Hat Europe: 2008. *https://www.blackhat.com/presentations/bh-europe-08/Filiol/Presentation/bh-eu-08-filiol.pdf*

We're helping national CSIRTs work together.

Jeffrey Carpenter

INCIDENT RESPONSE

CERT has been involved in incident response since its inception. We're addressing the big incidents affecting national and economic security.

Jeffrey Carpenter

Incident Response Overview

The CERT Program's work in the area of incident management currently focuses on two key areas: computer security incident response teams (CSIRTs) with national responsibility (National CSIRTs) and support for efforts to protect U.S. critical infrastructure. National CSIRTs play a key role in protecting the security of nations, economies, and critical infrastructures. They serve as central coordinating organizations within their countries for incident handling activities. Among other objectives, they seek to serve as a trusted point of contact and support incident reporting and mitigation across various sectors within a nation's borders.

Since its inception, CERT has supported critical infrastructure protection (CIP) and critical information infrastructure protection (CIIP). Homeland Security Presidential Directive 7 describes these infrastructures as "physical and cyber-based systems essential to the minimum operations of the economy and government. They include, but are not limited to, telecommunications, energy, banking and finance, transportation, water systems, and emergency services, both governmental and private."

Fostering Cooperation Among National CSIRTs

Jeffrey Carpenter, technical manager of the CERT Coordination Center, notes that National CSIRTs face the challenge of improving their capabilities in an environment of limited resources. "We're working on how to build more effective ways for National CSIRTs to collaborate, share information, and share solutions to common problems," says Carpenter. Better sharing mechanisms for National CSIRTS can reduce duplication of effort on the same problems, improving the capabilities of National CSIRTS while containing costs. "We're helping National CSIRTs work together," says Carpenter, "rather than individually, to solve problems common to all of them."

Currently, CERT is working with 75 functioning National CSIRTS. "We work with all of them to one degree or another. And, since 2006, we've sponsored an annual meeting that brings together technical staff from the National CSIRTs." The meeting serves as an opportunity for networking and collaboration in which National CSIRTs can present their work or research on issues, tools, and methods relevant to their community. Because of the CERT Coordination Center's long experience in this area, it is uniquely positioned to help.

"We've been at work in this area for more than 15 years. Having worked to help numerous countries establish a National CSIRT, we're now starting to achieve critical mass," says Carpenter. Many National CSIRTs have been developed using the advice of CERT and have adopted tools and processes originally developed by CERT. This has allowed CERT to turn its focus to helping these National CSIRTs become more effective.

Supporting Efforts to Protect Critical Infrastructure

CERT has long played a leading role in helping government and other organizations solve problems affecting the nation's critical infrastructure. "We're helping to inform and shape their efforts in the area of critical infrastructure protection," says Carpenter. In particular, CERT works with the Department of Homeland Security and the Department of Defense on improving their capabilities. This work focuses on

- research to identify new technologies and methodologies that support CIP
- research on critical infrastructure threats and vulnerabilities
- assist in the development of national critical infrastructure protection programs
- development of information- and tool-sharing capabilities

These efforts complement other projects in CERT aimed at developing information security risk assessments and methodologies, guidelines, and best practices centered on CIP. CERT is also collaborating with standards bodies to develop cyber security standards that support national CIP goals.

"CERT has been involved in incident response since its inception," says Carpenter. "Through our work in the areas of National CSIRTs and critical infrastructure protection, we're addressing the big incidents affecting national and economic security."

An Incident Management Body of Knowledge

Principal Investigators:
David A. Mundie, Robin Ruefle, and Sandi Behrens

In addition to those two key research areas, CERT studies the state of the practice of CSIRTs and other types of incident management capabilities. Based on analysis of that work, CERT produces best practice documentation and guidance, assessment tools, and educational products focused on process improvement and capacity building. One recent area of research is the identification of an Incident Management Body of Knowledge (IM BOK).

Despite considerable effort, incident management remains an ill-defined discipline. Disagreements persist on such fundamental issues as vocabulary, process models, competencies, and standards.

One technique which has recently gained currency as a way of solidifying and standardizing emerging disciplines is the Body of Knowledge (BOK), which has been used to capture essential knowledge and competencies in fields as diverse as massage therapy and process assessment. In this project we developed a body of knowledge for incident management.

We began by researching the body of knowledge methodology itself, which remains surprisingly nebulous. Based on the literature, we defined a ten-step systematic method for producing bodies of knowledge, starting with a controlled vocabulary of terms and progressing through taxonomies, static ontologies that capture the atemporal structures of the field being studied, dynamic ontologies that capture the processes and process attributes of that field, and intentional ontologies that capture the competencies and skills of the practitioner [1]. The last step in the method is to derive a meta-model which unifies the field with other related fields.

We then applied this ten-step method to incident management. For the controlled vocabulary we collected over 2,000 terms from a wide variety of sources, including five online information security dictionaries. For the dynamic ontology we examined 345 activities in 10 pre-existing process models from sources such as ISO 27002, ITIL, the U.S. Department of Homeland Security's (DHS) IT Security Essential Body of Knowledge, the CERT® Resilience Management Model, the Incident Management

Figure 1: Screenshot of facet map interface. The activities being examined are in the center, while the options for narrowing the activity set along four facets are on the right.
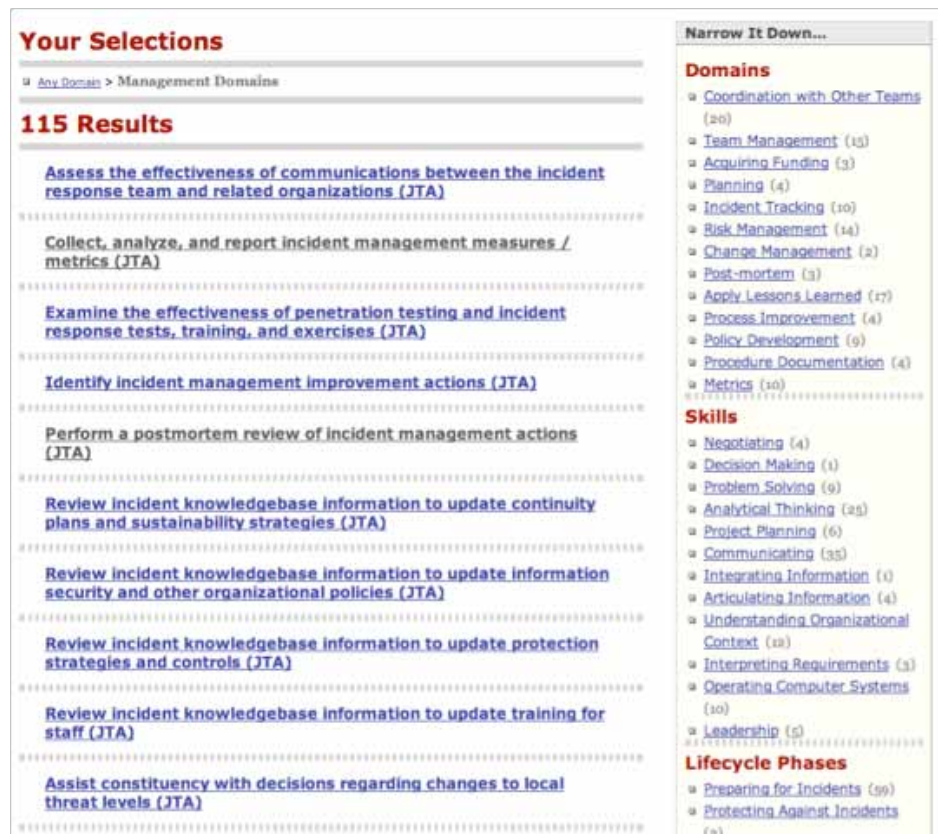
Capability Model, and the National Institute of Standards and Technology's (NIST's) 800-61. We showed that those activities can be synthesized into 23 clusters that represent the underlying activities of incident management. For the intentional ontology, we leveraged the SEI's previous experience in building competency-based bodies of knowledge such as the Competency Lifecycle Framework [2], the SCAMPI Lead Appraiser Body of Knowledge [3], and the Personal Software Process (PSP) Body of Knowledge [4] to incorporate a competency matrix into the model.
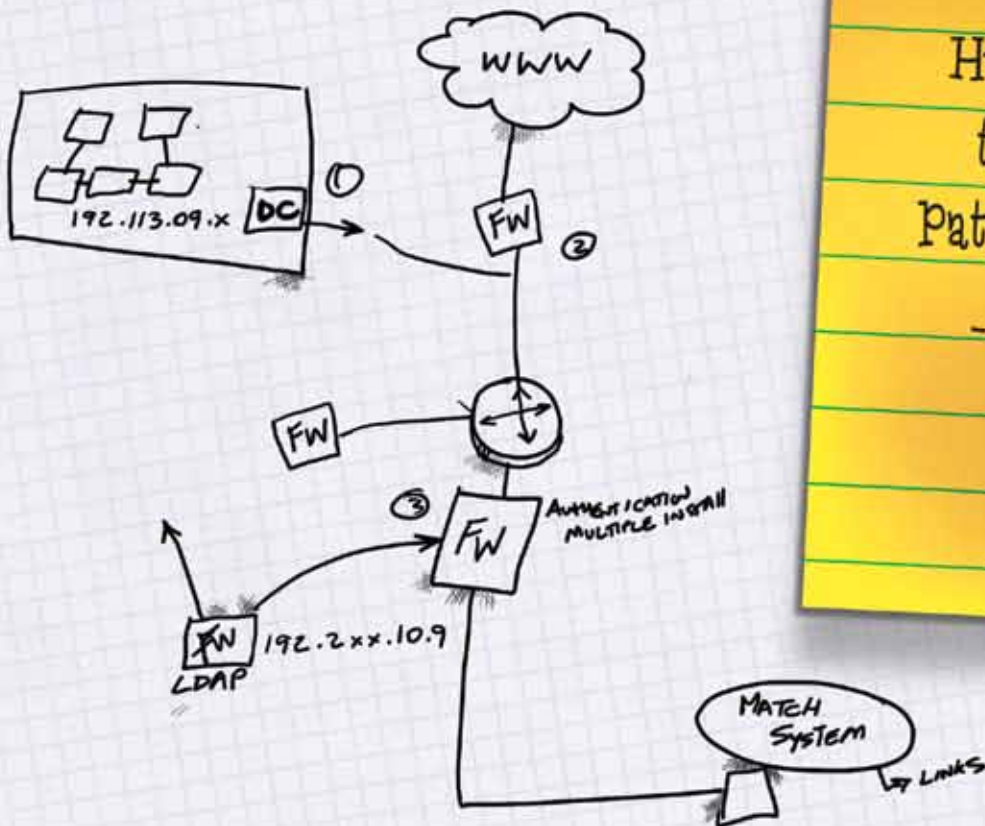
Navigating a body of knowledge this large can be difficult, so to provide a convenient interface we turned to facet maps, which have proven to be a simple yet powerful and flexible way to browse multi-dimensional spaces. See Figure 1 for a screenshot. Our facet map interface lets the user drill down to any area of the BOK along any of four dimensions: the knowledge domains, the skills, the incident management lifecycle phases, and the references for the activity.

The Incident Management BOK has six long-term goals: to improve benchmarking and gap analysis of incident management within organizations; to serve as the basis for creating certification programs which establish an individual's bona fides in incident management; to provide guidance for developing curricula, training requirements, and job competency descriptions; to enable the standardization of incident management at all levels; and to facilitate the creation of collective, expandable repositories for knowledge about incident management. Our current work is focused on exploring how the BOK can best be used to achieve those goals, on vetting the process taxonomy, and on finalizing the prototype implementation.

References

[1] I. Jurisica et al., "Using ontologies for knowledge management: An information systems perspective," in *Proc. ASIS Annual Meeting,* vol. 36, pp. 482-96, 1999.

[2] S. Behrens et al., "CMMI-based professional certifications: The competency lifecycle framework," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2004-SR-013, December 2004.

[3] S. Masters et al., "SCAMPI lead appraiser body of knowledge (SLA BOK)," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2007-TR-019, October 2007.

[4] M. Pomeroy-Huff et al., "Personal software process (PSP) body of knowledge, version 1.0," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2005-SR-003, August 2005.

Network
Situational
Awareness

www

192.113.09.x DC ①

FW ②

FW

③ FW AUTHENTICATION MULTIPLE INSTALL

FW 192.2xx.10.9
LDAP

MATCH
System
BY LINKS

Humans are
the best
Pattern matchers.

- Markus De Shon

Network Situational Awareness Overview

The Network Situational Awareness (NetSA) team focuses its research on monitoring large networks and analyzing bulk data collections, with the ultimate goal of detecting malicious activity. Bulk data include every transaction on large networks, either at connection summaries (network flow), specific application data (for example, DNS requests and responses and web proxy logs), or full packet capture. In NetSA's major research projects, analysts develop approaches to automated analysis, measure network phenomena, and determine the return on security investments due to attack mitigation.

Approaches to Automated Analysis

Large networks can generate billions of network transactions per day. Unassisted, people cannot possibly analyze such a huge amount of network transaction data. So the NetSA team does much of its work on approaches to automating that analysis. "Humans are the best pattern matchers," says Markus De Shon, team lead for the Trends and Metrics team in NetSA. "Analysts find security-relevant patterns first then figure out how to make machines see the same patterns."

NetSA analysts start with a hypothesis of how to detect patterns in bulk data. Then they develop new statistical models, algorithms, and methodologies, and apply advanced mathematical techniques, like machine learning, Bayesian techniques, and topology, to allow machines to reliably detect those patterns. Unreliable detection methods lead to many false positives and a lot of wasted effort, so analysts design the analytics carefully.

Similar research on network activity at a large scale usually only operates on snapshots of data (such as the annual Day in the Life of the Internet, or DITL, data sets collected by the Cooperative Association for Internet Data Analysis [CAIDA] and its collaborators). Through its work with the U.S. government, NetSA has access to ongoing large-scale data sources, which enable analysts to see the changes over time in the larger network context.

Measuring Phenomena of Networks

One of the biggest problems the NetSA team tackles is how to understand the phenomena behind data. "Some things are easy to measure, like simply counting data," says De Shon. "If you just count data you won't get the big picture." NetSA analysts develop complex statistical models that enable them to understand the big picture even in cases where they can see only a subset of the relevant network activity.

Botnets are a good example. Botnets are collections of automated computer programs or robots that can be controlled by one, or many, outside sources to conduct a range of malicious activities, from distributing spam and viruses to conducting denial-of-service attacks. When trying to measure the size of a botnet, often the most direct approach is taken, which is simply to count IP addresses. But anyone looking at the botnet's transactions on the network would see only a small subset of the botnet's communications, which may also be distorted by dynamic IP address assignment and network address translation devices on the network. NetSA's models quantify these distortions to enable CERT analysts to extrapolate these limited communications to an estimate of the true botnet size.

Measuring Return on Security Investment

Another significant project of the NetSA team is measuring the return on investment for attack mitigation. Measures of that return include the effectiveness of attack detection sensors, the extent of attack damage, and the value of the information being protected.

This work included developing a robust survey instrument with a level of sophistication usually found only in standardized testing. The value organizations place on something is subjective. Good surveys accurately collect what people believe. For example, military organizations may place different levels of sensitivity on troop movements at different times, which may influence how the organization values the security investment in protecting information about troop movements versus, for example, materiel transport.

Transitioning Research to the Real World

To apply its research, the NetSA team develops a proof of concept and then transitions it to customers' operations. Much of the NetSA work supports network defense operations in civilian government and the military, focusing on projects that deal with national security.

Another way NetSA transitions its research is through FloCon (*http://www.cert.org/flocon*). FloCon is an open workshop that provides a forum for operational analysts, tool developers, and other parties interested in the analysis of large volumes of traffic to showcase the next generation of flow-based analysis techniques. Flow is an abstraction of network traffic in which packets are grouped together by common attributes over time.

Network data can be overwhelming, but NetSA's research extracts the useful information from the data to enable useful interpretation and understanding.

About FloCon

The purpose of network flow data monitoring has changed significantly since FloCon began in 2004. Years ago, flow was the only way to understand activity on large networks as storage space was limited and deep-inspection technology simply could not keep up with traffic volume. Today, decreasing storage costs and increasing computing capabilities allow many products to generate huge volumes of deep packet data. However, flow analysis retains a vital role in the mission to understand network behavior.

The 2010 conference focused on flow data analysis within the context of other data sources. Presenters considered how flow is a piece of the puzzle, one data source among several which help you derive a holistic network view.

FloCon 2010 had an expanded demonstration session, and anyone with relevant tools or techniques was encouraged to register for the session. Similar to a poster session, the demonstration session provided an opportunity for informal interaction with the community to gain project feedback.

There were 86 attendees at FloCon 2010, including those from academia, industry, research, and government. International countries represented include Canada, Germany, Japan, Norway, and Switzerland.

January 11-14, 2010 | New Orleans, LA

Co-Chairs
Paul Krystosek, CERT
Sid Faber, CERT
Program Committee
Jim Downey, DISA PEO-MA
John Gerth, Stanford University
Joshua Goldfarb, IRS
Ray Kinstler, US-CERT
John McHugh, RedJack
Jeff Mauth and Troy Thompson,
    Pacific Northwest National Laboratory

Assessing the Benefits and Effectiveness of Network Sensors

Principal Investigator: Soumyo D. Moitra

Problem Addressed

Sponsors of CERT make significant investments in network sensors because they are a key component of security. The location of a sensor is a key determinant of that sensor's value. In view of constrained budgets, it is important for decision makers to know the value an organization gets from deploying a sensor at a particular location. Although sensors are and will continue to be deployed widely to defend networks against cyber attacks, there is currently no rigorous model in the public sector to guide decision making on sensor location. Thus there was a need to develop a model that would quantitatively assess the effectiveness and benefits of a sensor by location. Such a model would help decision makers with the acquisition and allocation process. It would also help prioritize the competing needs for sensors when resources are limited. This article reports on a model that has been developed for CERT sponsors to address this need.

The term "sensor" is used broadly to include not only the basic sensing device but also the hardware and software associated with monitoring, filtering, processing, and storing network traffic data for security purposes. Therefore the term includes intrusion detection and prevention systems (IDPSs) and security information and event managers (SIEMs).

Research Approach

The research project had four parts. In the first part, a comprehensive concept for the effectiveness of network sensors was developed. In the second part, a literature review of related topics was undertaken. In the third part, a model and methodology were developed to quantitatively assess the benefits of deploying a sensor at a given location. In the fourth part, a model for estimating the value of sensitive information was developed.

It is important to formalize the concept of sensor effectiveness in terms of an overall metric. In the first part, this was done from the managerial perspective, which is highly concerned with the issues of justifying investments in sensors and the optimal way to deploy them. From this conceptualization, methods were then developed to assess the return on investments in sensors. The overall metric is composed of one metric that considers the features and functionalities of the sensor and another metric that considers

the characteristics of the proposed location of the sensor. Both these factors have to be taken into account because the same sensor may provide different benefits at different locations, and different sensors may provide different benefits at the same location. This approach combines managerial and economic issues with the technical aspects of network security.

An extensive literature search was conducted in the second part of the project. Since this is an interdisciplinary issue, relevant literature on cost-benefit analysis, decision theory, public policy, effectiveness of preventive measures, network security metrics, and the extensive work in the areas of economics of information security and return on security investments was reviewed [1, 2, 3, 4, 5]. The key points arising from this review were that (1) this issue is of considerable interest and importance, (2) a number of models have been developed in the private sector, (3) many of the assumptions for the private sector do not hold for the public sector, and (4) it is difficult but possible to assess the benefits of security measures such as installing sensors [1, 3, 4].

There were, however, certain gaps in the current knowledge, most notably (1) a lack of a comprehensive set of attack categories and types of damages from cyber attacks; (2) a lack of a standard method to evaluate potential damages or mitigation effects; (3) a general lack of sensitivity analysis in spite of the acknowledged uncertainties; and (4) a lack of models for the network monitoring, detection, or incident handling processes.

These limitations suggested the need for a new model to assess benefits from sensors. This was the task of the third part of the project. The model that was developed consists of two modules: (a) an event tree that models the detection and response process and (b) an attack/damage matrix that utilizes data on potential damages as well as possible mitigation effects. The model is probabilistic in that it considers the probabilities of detection, prevention, and mitigation. The benefits were estimated as the reduction in expected damages as a result of having a new sensor. The analysis using the model involved extensive sensitivity analysis because the values of the different kinds of damages from different attack categories are uncertain. The possible attack rates were also varied to study their impacts. Finally, since several effects could be non-linear, such as the increase in damages as a function of the attack rate, non-linear effects were also considered.

The examination of the secondary data on estimates of damages from cyber attacks found large variations [2]. In particular, the variation in the estimates for the loss of sensitive information was extremely large, and it tended to dominate losses from other types of damages.

Therefore, in the fourth part, a new model and methodology were developed to estimate the value of sensitive information (VOSI). This is important because no standard method exists, and CERT sponsors need a comprehensive, consistent, and uniform method to estimate VOSI because of its

criticality. The method conceptualizes and estimates VOSI. It considers the loss to the owner when sensitive information is compromised and estimates the expected loss to an organization under different scenarios that include the threats to the organization and the potential losses. The estimation of the losses takes into account the kinds of compromises that can occur, the types of misuse, and the sensitivity of the information residing on the network.

Expected Benefits

The results from this research will benefit all organizations that need to plan for sensor deployment to defend their networks and the Global Information Grid. This includes almost all CERT sponsors (such as the Department of Defense (DoD), the Department of Homeland Security, and other U.S. government agencies). All of them deploy network sensors and depend heavily on them for network defense and information assurance. The model developed through this research provides a more rigorous and scientific basis for making decisions about sensors. The improved decisions will result in better security for a given budget [5]. The total budget for sensors is very large for U.S. government agencies, so improvements in the decision-making process will make a significant impact on the cost effectiveness of sensor deployment.

The research has resulted in quantitative, comprehensive, and practical metrics for sensor evaluation. While results of the model may not be the only input to final decisions, they will help with the managerial aspects of decisions. The model and the metrics specifically address concerns of CERT sponsors in ways that have not been done before, such as by including the attack categories of the DoD and considering the value of sensitive information. The results will help in deciding whether sensors are justified and, if so, where best to locate them. The research has identified the key data needs for making decisions in network security. The work has resulted in a method of assessing the value of sensitive information.

The model includes a novel module for the workflow at information assurance (IA) and network monitoring centers. The module can help managers analyze the effectiveness of their own workflows. The methodology developed can be used more generally by CERT sponsors for the allocation of security resources across all areas of network security (that is, not just for sensors, but for sensors, other security technologies, the Host Based Security System (HBSS), information assurance (IA) personnel, training, forensics, and other network defense measures).

2010 Accomplishments

The primary task of this project has been to investigate how network security decisions about sensors can be improved by making them more comprehensive and objective. The accomplishments of this research project can be summarized as follows:

- A conceptual model for a sensor effectiveness metric has been constructed. It will help guide managers and decision makers on how to think about the various aspects of sensors that are relevant to acquisition and deployment.

- A literature review that is comprehensive and up-to-date has been done. It has a large number of references that can serve as a source for further information for CERT sponsors.

- A model for estimating the benefits from sensors has been developed. Sponsors can use this model to estimate the benefits and the effectiveness of the sensors they plan to deploy. The model allows the user to estimate the incremental benefits of additional security measures. It is in the form of templates, and different scenarios can be analyzed depending on the needs of the user.

- The criticality of VOSI has been identified, and a method to address the issues of conceptualizing and estimating VOSI has been proposed. This is very important to many CERT sponsors.

The research has been documented in various reports that are available to CERT sponsors.

Future Goals

The future work that is planned includes

1. operationalizing the conceptual model for the effectiveness of sensors. That is, develop it in more detail and also develop a method of estimation. This will help CERT sponsors to easily apply the concepts.

2. refining and extending the model to incorporate some of the complexities that occur in reality, such as interaction effects among the variables and non-linearities. Further sensitivity analysis is also planned. These will be done according to the needs of CERT sponsors, which will be elicited.

The major challenges are data collection and developing case studies in collaboration with CERT sponsors. This is especially important for VOSI. The task of collecting the relevant data is complex and is rarely done at present. The research here has identified the required data and provides guidance on its collection.

Significant opportunities have arisen from this work. CERT can help its sponsors more actively with network security decisions and interact with them to integrate this approach into policy making. CERT can facilitate the application of the models and metrics widely among network security decision makers. Finally, CERT can help government agencies with data collection and estimation of sensor benefits through the application of these models.

References

[1] W. K. Brotby, Information Security Management Metrics, CRC Press, 2009.

[2] B. Cashell, "The economic impact of cyber-attacks," CRS Report for Congress, 2004.

[3] L. A. Gordon and M. P. Loeb, Managing Cybersecurity Resources, McGraw-Hill, 2006.

[4] K. J. Soo Hoo, "How much is enough? A risk-management approach to computer security," Working Paper, Consortium for Research on Information Security and Policy (CRISP), Stanford University, 2000.

[5] B. Schneier, Schneier on Security, Wiley Publishing, 2008.

How Fast Do Computers Move?

Principal Investigator: Rhiannon Weaver
External Collaborators: Chris Nunnery, Gautam Singaraju, and
Brent ByungHoon Kang at UNCC and George Mason University

Problem Addressed

A botnet is a collection of computers that have been compromised by a malicious software (malware) program, putting them all under the control of a single malicious operator or small group of operators. These "bot herders" use their armies of machines to gather intelligence about other networks through scanning, to send spam email on a large scale, or to cripple local servers or even national network infrastructures by distributed denial of service (DDoS) attacks. Tracking botnet size is important in order to understand the scope and spread of an infection. Security analysts rely on population estimates to prioritize threats and to measure the efficacy of clean-up strategies.

Often, watchdog groups like Shadowserver cannot identify infected machines directly, but they can identify IP addresses, protocols, and ports through which infected machines communicate.[1] But the relationship between machines and internet protocol (IP) addresses is like the relationship between people and street addresses; an address can represent a single home, a high-rise apartment building, or a time-share. A single infected mobile device such as a phone or laptop computer can cycle through IP addresses as it physically travels. Furthermore, two widely adopted network administration practices also complicate the relationship between IP addresses and even statically located machines:

- Network Address Translation (NAT, one-to-many): A network of many machines is configured to access the internet through a single machine with one external-facing IP address, often called a gateway or proxy. Gateway traffic is also often shuffled among two or more IP addresses over time to balance bandwidth across several assets, a technique known as load-balancing.

- Dynamic Host Configuration Protocol (DHCP, many-to-one): internet service providers (ISPs) often have a pool of IP addresses, any one of which can be provided dynamically to a machine via a temporary lease. Depending on the network configuration, DHCP leases can be valid for hours or days. One-day leases are common.

With NAT, a single IP address may represent hundreds of machines. With DHCP, even a stationary desktop computer can appear to travel, from an IP address perspective. One way to account for NAT-ted machines, mobile devices, and DHCP pools in population estimates is to model the typical range of movement of machines among IP addresses and to use this model to predict a likely range of addresses where a single machine may be observed.

Research Approach

The Waledac botnet, active during the latter half of 2009, represented one of the 10 largest networks of infected computers at the time [1]. From the period of December 4 through 22, 2009, researchers were able to collect log files from Waledac-infected machines as they checked in to the main command and control servers in order to receive instructions. Each check-in was associated not only with an IP address and timestamp, but with a unique hash ID for the infected machine. This identification system allowed researchers to track the IP address profiles of individual machines as they traversed IP address space.

Using the MaxMind GeoLite City database[2], we were able to associate each IP address with an approximate latitude and longitude. Using the timestamps observed in the logs, each machine in the botnet was assigned a mobility score equal to the average miles per hour traveled during its check-ins, from the first to last observance within the 18-day window. Distances were calculated using the Haversine (great circle) distance between points on the globe. A mixture of Gaussian models was fit to the set of non-zero mobility scores on the logarithmic scale, which was used to detect three anomalous outliers.

Expected Benefits

IP addresses are becoming more ephemeral as measures of individual infections, but there is little beyond heuristic rules of thumb to account for this inflation in population estimates. As IPv6 is more widely adopted, machines will have even more address space in which to travel. IP address mobility also affects the efficacy of countermeasures such as blacklisting. Studying the mobility of infected machines will help us not only to understand the underlying machine population, but also to estimate the effectiveness of implementing blacklists versus a more direct approach such as server take-downs.

---

[1] http://www.shadowserver.org

[2] http://www.maxmind.com/app/geolitecity

2010 Accomplishments

We conducted an extensive study of the IP address use and mobility for the Waledac botnet [2]. A total of 172,238 unique hashes were observed communicating through 548,997 IP addresses during the 18-day window. A majority of hashes in the botnet (63.6 percent) were associated with only a single IP address. A total of 55.9 percent of hashes were uniquely and disjointly paired with a single IP address, comprising 17.5 percent of all IP addresses, and another 7.7 percent of these static hashes appeared to share space behind a NAT. Thus nearly two-thirds of the botnet communicated through only 18.5 percent of all observed IP addresses. On the other hand, the top 1 percent of mobile hashes were very mobile, associated with 80 or more IP addresses, with a maximum of 428 IP addresses observed for a single hash. The mobile hashes comprised only 36.3 percent of all hashes but communicated through 81.4 percent of observed IP addresses.


Figure 1: Data (bar) and curve (blue/red lines) for a mixture model fit to mobility scores on the logarithmic scale

Figure 1 shows a histogram of the mobility scores for the 62,618 mobile hashes in the data set, shown on the logarithmic scale. The blue and red curves indicate the best mixture of Gaussian distributions that fit the data. Two groups are apparent. The large group (89 percent of mobile hashes) moved on average less than 10 miles per hour and likely represents machines with relatively long DHCP leases in geographically close pools of available addresses. The small group (11 percent of mobile hashes) appeared to move more aggressively, with some machines "travelling" on average 800 miles per hour or more as they changed location in IP address space. These machines may be associated with satellite links, very large address pools, or "bullet-proof" proxy networks.

The top 10 mobility scores are marked by solid black lines. Dashed lines mark the upper percentiles of 0.001, 0.0001, and 0.00001 in the model. An outlier analysis based on Monte Carlo simulations of this model flagged hashes with the top three mobility scores (23,272.11 MpH, 8,114.45 MpH, and 5,471.55 MpH) as statistical anomalies, and further analysis of the data set uncovered evidence to show that the hash IDs for these outliers were not tied to unique machines.

Future Goals

An observable host-to-IP relationship gives us a glimpse not only into a botnet, but into the networks it infects. Although the type of infection may change, the administration policies of infected networks—including DHCP regions, NAT policies, and throughput rates—remain relatively constant. In the future we hope to use mobility profiles from observable botnets such as Waledac to map out the boundaries of shared IP address pools and to profile network properties such as load-balancing thresholds and DHCP lease times among these networks. This information can be leveraged in the study of less visible botnets, such as Conficker, by, for example, using measured network profiles to adjust behavior-based population models. Not only can we transfer information directly between networks that house both kinds of infections, we can also use statistical models to extrapolate this information and infer the unobservable properties of a hidden botnet across the IP address space it infects.

References

[1] Thomas Claburn, "Microsoft decapitates Waledac botnet," in *Information Week online,* Feb. 25, 2010. *http://www. informationweek.com/news/hardware/desktop/showArticle. jhtml?articleID=223100747*

[2] R. Weaver, C. Nunnery, G. Singaraju, and B. Kang, "Entropy-based measurement of IP address inflation in the Waledac botnet," presented at FloCon 2011, Salt Lake City, UT, Jan. 2011. *http://www.cert.org/flocon/2011/presentations/ Weaver_Entropy.pdf*

Closed Networks: The Next Generation for Secure Network Design

Principal Investigators:
Sidney Faber and George Warnagiris

The Closed Computer Network model being developed by the Network Situational Awareness team at CERT seeks to make a significant improvement in securing critical information assets in a networked world. Although current best practices recommend network segmentation for improved security [1, 2], actual implementation details are sparse and lack standardization. Under the Closed Computer Network model, an organization can design a completely self-contained network yet still retain the ability to access information on public networks within well-defined constraints. Networked assets will only connect to the closed network when approved by a central authority, and this authority will dictate policy or disconnect assets that threaten network security.

In 2010 we conducted an initial study into the advantages of defending the closed network. The response actions used to defend the typical intranet were assessed in the closed environment, and the benefits of a Closed Computer Network became obvious. Not only is the closed network inherently more secure, but the defender can also pursue new directions in incident response and root cause analysis. Attribution of security incidents also becomes a realistic goal.

In 2011 we plan to significantly expand our work in Closed Computer Networks. This will include formally defining the need for closed networks, outlining scenarios where the closed network is applicable, and creating an architectural framework for a typical implementation. We anticipate this research will draw heavily on insider threat knowledge at CERT so that closed network defense is tuned to detect the malicious insider [3, 4]. Additional research will be done on defending the closed network, and recommendations will be made for implementing sensors on the closed network.

References

[1] G. Stoneburner, "Underlying technical models for information technology security," National Institute of Standards and Technologies, Information Technology Laboratory, SP800-33, Dec. 2001, pp. 16-17. *http://csrc.nist.gov/publications/nistpubs/800-33/sp800-33.pdf*

[2] K. Scarfone and P. Hoffman, "Guidelines on firewalls and firewall policy," NIST Special Publication 800-41 revision 1, Sep. 2009, pp. 2-5–2-7. *http://csrc.nist.gov/publications/nistpubs/800-41-Rev1/sp800-41-rev1.pdf*

[3] Stephen Band et al., "Comparing insider IT sabotage and espionage: A model-based analysis," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2006-TR-026, 2006. *http://www.sei.cmu.edu/library/abstracts/reports/06tr026.cfm*

[4] Michael Hanley, "Deriving candidate technical controls and indicators of insider attack from socio-technical models and data," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2011-TN-003, Jan. 2011. *http://www.sei.cmu.edu/library/abstracts/reports/11tn003.cfm*

Finding the Botnets
You Don't Know About

Principal Investigator: Evan Wright

Botnets—autonomous, network-enabled collections of malicious software agents—are increasingly using web traffic for communication amongst their members. They often need to communicate with each other or the command and control server. Historically, botnet members have communicated with one another or the command and control server via Internet Relay Chat (IRC) or Instant Messaging. Newer botnets use communication techniques such as communicating via HTTP to specific domain names or social networking media.

Law enforcement and defensive organizations may take specific domain names down when a domain has been determined to pose a threat to a large population. Botnet members use a technique called Domain Flux to avoid the botnet having a single point of failure. Domain Flux is a technique used by botnet members to map an Internet Protocol (IP) address to any number of computer-generated domain names. The IP address may be mapped to any subset of the domains and then use it to communicate to the rest of the botnet. If some of those domains are disrupted, the botnet will continue to function effectively. Domain flux affords higher fault-tolerance and resiliency in devices that are members of a botnet. Domain flux is resilient to intrusion detection system signatures, a common detection technique, particularly the detection of the signature behavior of botnet members communicating with explicitly blacklisted domains.

A human expert can distinguish between human- and computer-generated domains; consequently, a statistical classifier with the correct input data should be able to perform automatic classification. In 2010, work by the Network Situational Awareness team studied the effectiveness of using extracted lexicographical features of a domain name to predict if a domain name was human- or computer-generated. Some of these features include the name length; common letter frequency distributions; and use of numbers, vowels, consonants, and symbols. The computer-generated domains came from three botnets: Conficker, Kraken, and Srizbi. The team harvested human-generated domain names from web directories that humans manually review. The team extracted more than 100 features from domain names and classified them with a C4.5 decision-tree algorithm with 10-fold cross validation, which yielded 90 percent accuracy.

In the future, we plan to extract a greater number of features which will overall yield more predictive features from the domain names to increase accuracy while retaining scalability. In addition, other classifiers will be tested, including classifiers that apply feature selection algorithms. We also plan to be able to fully map out clusters of computer-generated IPs and domain names to infer populations on the internet as a whole. We designed our method to be scalable for large data sets. Passive Domain Name System (DNS) data repositories have been demonstrated to be scalable to internet-scale data with suitable hardware. Specifically, we plan to leverage large, passive DNS data sources on the internet at regular intervals to track the behavior of botnets that use domain fluxing techniques.

References

[1] Alper Caglayan, Mike Toothaker, Dan Drapaeau, Dustin Burke, and Gerry Eaton, "Behavioral analysis of fast flux service networks," in *Proc. 5th Annu. Workshop Cyber Security and Information Intelligence,* Knoxville, TN, Apr. 2009, pp. 48:1-48:4.

[2] Brett Stone-Gross, Marco Cova, Lorenzo Cavallaro, Bob Gilbert, Martin Szydlowski, Richard Kemmerer, Christopher Kruegel, and Giovanni Vigna, "Your botnet is my botnet: Analysis of a botnet takeover," in *Proc. 16th ACM Conf. Computer and Communications Security,* Chicago, IL, Nov. 2009, pp. 635-647.

RESILIENCE MODELING AND ANALYSIS

EC
ENVIRONMENTAL CONTROL

MA
MEASUREMENT AND ANALYSIS

COMM
COMMUNICATIONS

SC
SERVICE CONTINUITY

RECYCLED PAPER

Resilience Modeling and Analysis Overview

In 2010, the Resilience Modeling and Analysis team continued its research efforts in the resilience management research and development arena with work that included the first full-scale application of the CERT® Resilience Management Model (CERT®-RMM), the first Class A Appraisal, and the first Class C Appraisal, which focused on security improvements and internal business continuity and disaster recovery processes, respectively. These efforts demonstrated the ability of the model to improve operational resilience management from both a security and continuity standpoint. Throughout 2010, the team discovered the various applications for a resilience model. They were able to use the model in different ways to systematically address improvement challenges. In 2010, the team focused on commercial and industry use of the model and discovered that CERT-RMM could be used in improving the operational resilience of manufacturing processes that produce critical output, such as in the defense metals industry.

About CERT Resilience Modeling and Analysis

The CERT Resilience Modeling and Analysis team provides tools, techniques, and methods that help organizations characterize, diagnose, measure, and improve operational resilience. One of its main goals is to help organizations improve their security activities by framing them with resilience as the outcome. With CERT-RMM, security, IT operations, and continuity converge into a single model that can be used to catalyze an organization's improvements and improve its resilience posture. The ability to measure operational resilience and the attainment of critical infrastructure resilience are key factors in mission assurance in both the private sector and the Department of Defense.

Key Research

In 2010, the team kicked off the Resilience Measurement and Analysis Initiative. This work focuses on identifying and piloting metrics that can identify and measure an organization's level of operational resilience based on resilience processes.

The team also developed the CERT-RMM Compass, a lightweight assessment method to diagnose areas for improvement in managing operational resilience. This method can be used by organizations wishing to focus on improving security and resilience without investing in extensive appraisal activities. The Compass offers a quick way to apply the breadth of the model without attaining significant knowledge of the model's details.

In addition to refining existing work, the team began foundational research in several areas.

- Working to develop a capability model for incident resilience—incident resilience focuses on the capabilities that an organization needs to mature to ensure that it can continue to assure and achieve its mission whenever its critical assets and services are under stress—such as the stress that comes from an event or an incident

- Characterizing resilience and security postures—a reductionism-based approach that provides a framework for identifying and examining various measures of posture (such as the results of a penetration test) and using them to explain and characterize an emergent property such as posture

- Defining tools, techniques, and methods for incident management across disparate organizations or collaborative communities—for example, incident management across a power grid that is owned and operated by a number of different organizations in a particular region

Resilience Research Applied

To date, thousands of individuals in a variety of nations and across all types of industry, government, and academia have downloaded CERT-RMM. It's being used as the foundation for a number of assessment methods including the following:

- Cyber Resilience Review—Through DHS, this assessment method helps private organizations and local, state, and tribal governments to assess their ability to effectively manage critical infrastructure protection.
- National Cyber Security Review—The NCSR is a congressionally-mandated activity that will allow a baseline review of how the 50 states are performing with respect to cyber security initiatives. CERT-RMM Compass is the basis for this instrument.
- Incident Management Capability Model—CERT-RMM serves as the foundation for the development of this model that focuses on building incident management capabilities, particularly in developing nations.
- CERT-RMM is the foundation for the development of an assessment instrument that will help the U.S. Department of Health and Human Services to assess the readiness of health organizations to adopt and secure electronic health records.

In addition to their work on CERT-RMM, the CERT Resilience Modeling and Analysis team is currently the steward for the Smart Grid Maturity Model (SGMM), which is being used by many organizations to chart their smart grid transformation. More than 50 organizations have been appraised by the model to date. The SGMM and CERT-RMM can be used together to chart smart grid transformation and to manage operational resilience of the smart grid.

Moving Forward

The CERT Resilience Modeling and Analysis team will continue to research means for organizations to characterize and measure their resilience postures. The team's research will delve into the activities that organizations perform, looking for correlations between the maturation of these activities and improved resilience to events and hostile risk environments.

Measuring Operational Resilience: Moving from Uncertainty to Justified Confidence

Principal Investigator: Julia H. Allen

Problem Addressed

Organizations in every sector—industry, government, and academia—are facing increasingly complex business and operational environments. They are constantly bombarded with conditions and events that can introduce stress and uncertainty that may disrupt operations and critical services delivered to customers. Typically, the capability to withstand stress and disruption is measured by the way an organization has performed during an event or is described in vague terms that cannot be measured. For example, when organizations are asked to describe how well they are managing resilience, they typically characterize success in terms of what hasn't happened: "We haven't been attacked or had a significant service disruption; therefore we must be doing okay." Because there will always be new and emerging threats and unexpected, disruptive events, knowing how well an organization responded to one attack is necessary but not sufficient; it is more important to be able to predict how it will perform in the future when the risk environment changes.

Organizations lack the ability to assess and measure their capability for managing operational resilience, as they have no credible yardstick against which to measure. As organizations strive to improve their ability to effectively manage operational resilience, having an approach for determining what measures best inform the extent to which they are meeting their performance objectives is essential. The SEI has chartered the Resilience Measurement and Analysis (RMA) research project to advance the state-of-the-practice in operational resilience measurement and analysis.

Research Approach

Measurement is about transforming strategic direction, policy, and other forms of management decision into action and measuring the performance of such action. The RMA project addresses the following research questions, often asked by organizational leaders:

- How resilient is my organization? Are we resilient enough? How resilient do we need to be?

- Have our processes made us more resilient?

- Do we need to spend more on resilience? If so, on what? What are we getting for what we've already invested?

To inform these, this question is relevant:

- What should we measure to determine if performance objectives for operational resilience are being achieved? Do we know what our performance objectives are?

Most organizations today lack a reliable means for measuring either their operational resilience or their capability for managing operational resilience. The traditional disciplines of security, business continuity, and IT operations are typically compartmentalized, placing high-value services and their associated assets at risk. In addition, measuring the degree, state, or "posture" of an intangible quality attribute or emergent property is difficult even under normal operating conditions. The emergent property of operational resilience, however, can be most accurately observed and directly measured during times of stress and disruption. Unfortunately, this is often too late to be of benefit, and the organization is typically in too reactive a mode even to consider how to improve in anticipation of the next incident.

Looking to the fidelity and performance of the contributing processes may be a way to get more confidence and precision about an organization's state of operational resilience—it is, at least, one important indicator that is not typically being measured today.

Unlike other efforts to measure operational resilience, this research project uses as its foundation a process-based definition of resilience, as defined by the CERT® Resilience Management Model (CERT®-RMM) [1, 2]. CERT-RMM addresses the ability of an organization to protect and sustain the resilience of mission-critical assets and services.[1] The model defines an operational resilience management system (as shown in Figure 1) and provides a framework of goals and practices at four increasing levels of capability described in 26 process areas (PAs), each of which includes example measures.

---

[1] A service is a set of activities that the organization carries out in the performance of a duty or in the production of a product. A mission-critical or high-value service is one on which the success of the organization's mission depends. High-value assets (people, information, technology, facilities) are those upon which a high-value service depends.
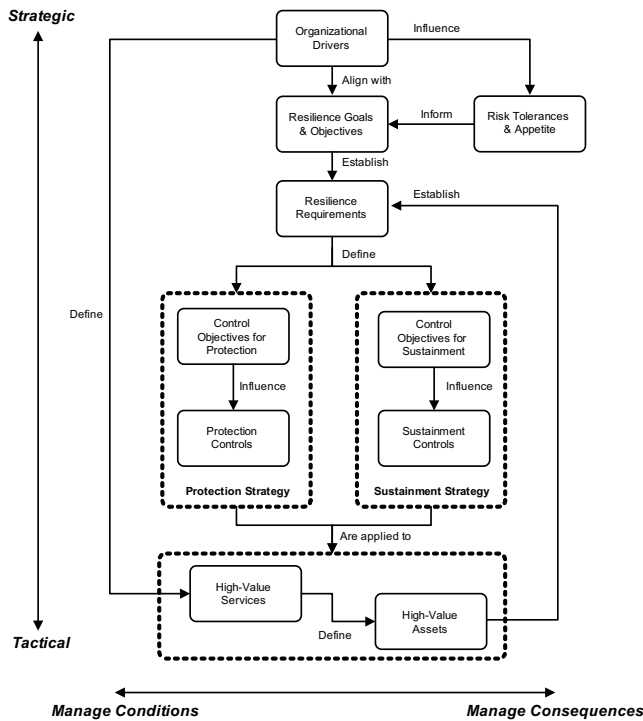
Figure 1: The operational resilience management system: The foundation for measuring resilience

As a process improvement and maturity model, CERT-RMM can guide an organization in using defined processes as a benchmark to identify its current level of organizational capability, setting an appropriate and attainable desired target for performance, measuring the gap between current performance and targeted performance, and developing action plans to close the gap. By using the model's defined processes as a foundation for performance and measurement, the organization can obtain an objective characterization of how it is doing, not only against a base set of functional practices but also against practices that indicate successively increasing levels of capability.

The first step in defining a meaningful measurement program for operational resilience and its effective management is to determine and express the required or desired level of operational resilience for an organization (organizational drivers in Figure 1, which may include strategic directives and critical success factors). An organization may be the enterprise, any business line or operating unit, or other form of business relationship, including partners, suppliers, and vendors. An organization can target a level of capability for one or more PAs, thus establishing a benchmark against which its operational resilience can be measured. Ideally, the targeted level for each process area is established during strategic and operational planning and when planning for continuity of operations, not as an afterthought during times of stress and service disruption. The targeted level should be no less and no more than that which is required to meet business mission objectives.

Meaningful measurement occurs in a context, so we further explore and derive example measures within the context of selected ecosystems (one example is shown in Figure 2), which are collections of process areas that are required to meet a specific objective. Example measures are derived using the approach shown in Figure 3 and defined using a measurement template.
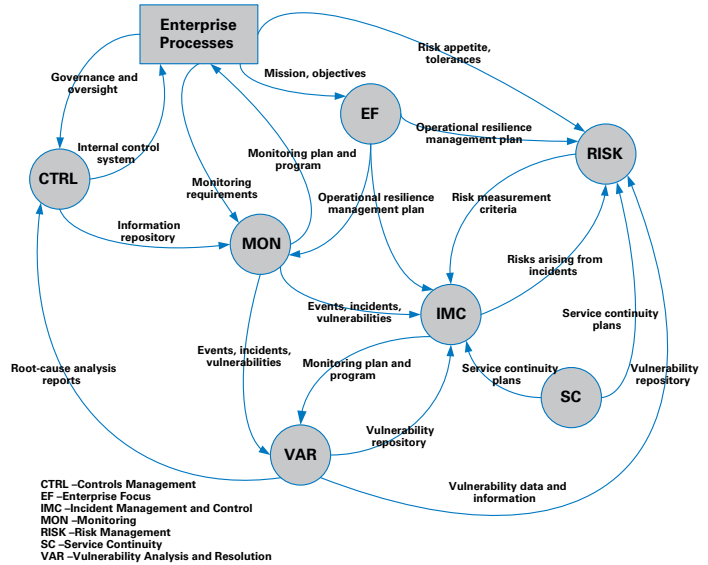

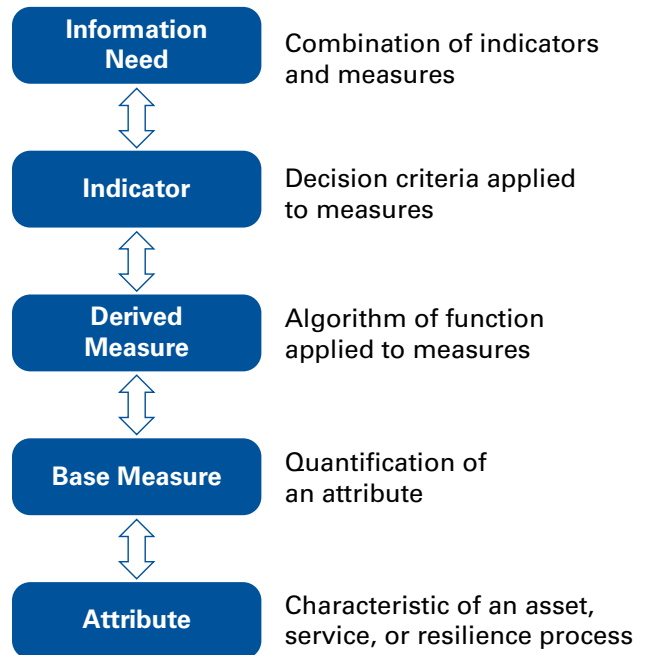Figure 2: An example CERT-RMM ecosystem for incident management

Figure 3: Resilience measurement relationships

The RMA research project provides a unique approach because

- it uses a process-based definition of resilience as its foundation (CERT-RMM v1.1)
- it builds upon the SEI's years of experience in developing and implementing capability maturity models and measuring organizational performance against them
- it is strongly influenced by a community of resilience practitioners
- its measures derive from a well-defined set of strategic and operational objectives (also shown as information need in Figure 3)

Expected Benefits

Potential benefits of implementing a measurement program based on this research include the following:

- Decision makers have better tools for predicting and diagnosing problems and for making better-informed decisions about their current state of operational resilience and where to invest.
- Measures provide justified confidence that high-value services and associated assets are meeting their performance objectives for operational resilience.
- This robust research method can be used for continuing exploration in any resilience domain that requires measurement and analysis.

2010 Accomplishments

The major accomplishment for 2010 was the publication of the first research report on this subject [3]. Research results included

- establishing six high-level objectives for operational resilience, including candidate measures that illustrate each of these
- defining the foundations for measuring operational resilience drawing from foundational measurement research methods such as GQ(I)M (Goal-Question-Indicator-Metric)
- defining a resilience measurement template and defining several examples measures using it
- applying the measurement derivation process to define three measures from three selected resilience ecosystems: managing risk, threat and incident management, and protecting information assets
- defining a candidate set of top 10 strategic measures for operational resilience (presented at the Computer Security Institute 2010 conference in November 2010)

Future Goals

Resilience measurement and analysis research in FY11 includes the following:

- Based on FY10 research results, revise and augment example measures presented in CERT-RMM v1.1. Issue as an addendum.
- Examine a range of CERT-RMM appraisal, survey, and review results together with results from the first CERT-RMM Users Group Workshop series for additional and revised candidate measures.
- Develop a detailed guide for CERT-RMM process implementation as the basis for defining measures to collect and the most effective tasks in the process to collect them.
- Seek pilot opportunities to validate measures for both process implementation and effectiveness. Effectiveness work will continue into FY12 and intends to provide evidence that supports or refutes the hypothesis that improving resilience processes contributes in some measurable manner to improving operational resilience.

References

[1] Richard A. Caralli, Julia H. Allen, Pamela D. Curtis, David W. White, and Lisa R. Young, "CERT® Resilience Management Model, v1.0," Software Engineering Institute, Carnegie Mellon University, Pittsburgh PA, CMU/SEI-2010-TR-012, 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tr012.cfm*

[2] Richard A. Caralli, Julia H. Allen, and David W. White, CERT® Resilience Management Model: A Maturity Model for Managing Operational Resilience. Addison-Wesley, 2010.

[3] Julia Allen and Noopur Davis, "Measuring operational resilience using the CERT® Resilience Management Model," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, CMU/SEI-2010-TN-030, September 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tn030.cfm*

Enhanced Methods
for Cyber Exercise

Principal Investigators:
Matthew Butkovic and James Stevens

Today's increasingly hostile cyber security environment includes rapidly evolving threats that many times outpace the ability of an organization to react. While it is hard to predict the best technical response to a threat, it is possible to develop plans for marshalling the necessary expertise and resources to coordinate a response, analyze options, and mitigate operational impacts. Performing cyber exercises has proven to be an effective tool to understand and improve the capability of organizations to prevent and respond to a broad range of threats.

The objectives of a cyber exercise focus on protecting, defending, and recovering critical assets and operations from a cyber attack or cyber incident. The exercise includes training or evaluating personnel on prevention, protection, response, and recovery procedures. It also includes practicing incident communication, command, and control processes with stakeholders and partner organizations.

Cyber exercises allow organizations to demonstrate critical capabilities, thus revealing how effectively they integrate their people, processes, and technology to protect their information assets and cyber-reliant services. Yet, despite the growing acknowledgement of cyber exercises as useful, they do not yet rival their traditional, all-hazards cousins (e.g., physical exercises) when it comes to innovation or utility—there is not a well-established body of knowledge and best practices to address the unique challenges of cyber exercises. The Homeland Security Exercise and Evaluation Program (HSEEP), codified by the U.S. Department of Homeland Security (DHS) and drawing extensively from the U.S. Department of Defense (DoD) exercise program, provides best practices for organizations preparing for disasters such as terror attacks, natural disasters, and pandemic outbreaks. However, the same level of richness does not exist for preparing for cyber attacks. This lack of detail in the HSEEP is one reason why cyber exercises often do not realize their full potential. Other reasons include

- Cyber exercises frequently lack the operational realism that significantly enhances their value.

- The lack of codified best practices leads organizations to use ad hoc formats and planning methodologies—neither of which promotes repeatability and familiarity across the cyber community, nor scalability to handle large exercises.

- Organizations struggle to deal with the unique complexities of cyber exercises arising from their inherent technical nature—a significant factor limiting their widespread use.

In 2010, the Software Engineering Institute's (SEI) CERT Program began working with the United States Department of Homeland Security to capture and codify cyber exercise methods. This compilation of cyber exercise best practices was presented to DHS as a report titled *Enhanced Methods for Cyber Exercise*. The work will also be released to a broader audience as an SEI technical report in 2011.

To enhance the state of the art of cyber exercises, methods must be developed for increasing operational realism, fostering repeatable processes, and managing complexity. The techniques presented in *Enhanced Methods for Cyber Exercise* specifically address these challenges, drawing upon subject matter experts and cyber exercise planners who have experience with them. Topics include

- cyber exercise planning team composition
- cyber scenario planning methodology
- live aggressor forces
- activities on operational networks
- simulators and ranges
- cyber-specific exercise documentation
- cyber-specific exercise control constructs
- cyber-exercise evaluation and improvement planning

During 2011, the CERT® Program will apply these advanced concepts to the creation of cyber exercise scenarios for specific critical infrastructure sectors. The energy and banking/finance sectors are anticipated to be areas of focus. CERT will also explore cyber exercise as a method to measure and improve an organization's response capabilities over time. The application of enhanced cyber exercise methods should allow organizations to demonstrate improvements to targeted capabilities.

Understanding Infrastructure Resilience Posture

Principal Investigators:
Matthew Butkovic, Samuel Merrell,
Philip Scolieri, and James Stevens

In the recent past, ensuring the protection of a nation's critical infrastructure was a more clearly defined undertaking.[1] A nation's critical infrastructure (i.e., its assets) was well understood and largely contained within its borders. Protection primarily involved the government ensuring the physical integrity of the assets that comprised the various infrastructures.

Today the task of protecting critical infrastructure is a significantly more complex undertaking due to multiple causes, including

- increased operational complexity in critical infrastructures
- less rigid boundaries for delineating critical infrastructure
- increased reliance on intangible assets (i.e., information)
- opening of borders and markets due to globalization
- increased reliance on information technology

The overall effect has been a significant increase in both the number and types of risks. This challenge is compounded by ownership and oparation of most critical infrastructure by the private sector, obscuring the issues of who can and who should manage infrastructure risks.

In 2009 CERT began working with the U.S. Department of Homeland Security (DHS) to develop a better understanding of the capabilities of owners and operators of critical infrastructures to protect and sustain their assets and services. Correspondingly, the CERT Program developed the Cyber Resilience Review (CRR), a method to measure the adoption and maturity of cyber resilience practices within an enterprise. The key goal of the CRR is to identify whether an enterprise exhibits capabilities for ensuring the resilience of a critical infrastructure service. Over 100 CRRs have been performed to date by DHS.

Version 1.0 of the CRR is an extensive interview-based process with questions derived from the CERT® Resilience Management Model (CERT®-RMM). CERT-RMM is a maturity model that describes organizational processes

necessary for ensuring the protection and sustainment (i.e., the resilience) of an organization's operations.

In 2011 the CERT Program is working to develop version 2.0 of the CRR method. The goal is to incorporate a number of enhancements to the CRR method designed to improve data collection and analysis, including development of

- a more robust question set and standard responses
- an algorithm to ensure objective scoring
- transition artifacts such as courseware and assessment guidance
- enhanced process improvement recommendations for sites participating in reviews

These changes will also facilitate the development of better recommendations for improvement for CRR participants. Many of these improvements will be achieved by utilizing the lightweight CERT-RMM Compass assessment as a foundation. The Compass consists of standard sets of multiple-choice questions and responses and a consistent scoring approach, easily incorporated into the CRR version 2. They provide insight into practices performed, incomplete practices, and institutionalizing factors (such as governance, training, policy, and measurement) that support the retention of practices under stressful conditions.

Version 2 of the CRR will also more precisely define an organization's service components for resilience of critical infrastructure services. This will provide a more granular view of an organization's capabilities and allow more meaningful comparisons between sector members. By gathering this type of data during the review, we hope to extend our research into critical infrastructure protection by beginning to answer the following questions:

- Are there differences in cyber security management profiles among the 18 critical infrastructure sectors? For example, do the various sectors exhibit different operational resilience patterns?
- Can we determine whether a given security management profile is sufficient for a given risk environment? For example, can you develop a security profile and determine whether it provides sufficient mitigation properties under various risk environments/operating environments?
- Are there general recommendations for improvement that would benefit the entire sector based on inspection of assessment data?
- Are there ways to measure critical infrastructure resilience and improve confidence in these measurements? What are effective measures of infrastructure resilience? What are the key dependencies linking critical infrastructure sectors?
- How well do critical infrastructure owners and operators understand their role in safeguarding the nation's critical infrastructure?

---

[1] Critical infrastructure is the physical and cyber based systems essential to the minimal defense and economic security of a nation, the minimal operation of its government, and the basic functioning of society. [Adapted from Executive Order 13010 (E.O. 13010) and Presidential Decision Directive 63 (PPD 63).]

The Smart Grid Maturity Model Updated

Principal Investigators:
James F. Stevens and David W. White

The 2009 CERT Research Annual Report introduced Version 1.0 of the Smart Grid Maturity Model (SGMM). The SGMM is a management tool that helps utilities to plan their smart grid journeys, prioritize their options, and measure their progress as they move toward the realization of a smart grid. The model describes eight domains containing logical groupings of incremental smart grid characteristics and capabilities that represent key elements of smart grid strategy, organization, implementation, and operation. Utilities use the SGMM to assess their current state of smart grid implementation, define their goals for a future state, and generate inputs into their road mapping, strategy, planning, and implementation processes.

In addition to introducing the model, the Report also announced plans for an update. This update to the model was to focus on architectural and usability improvements. The update was released on September 30, 2010, as Version 1.1 of the SGMM.[1] Prior to its release, the update was pilot tested with more than thirty utilities to ensure its quality and usability. Users of Version 1.1 benefit from a significantly improved model and supporting product suite that is built upon a refined version of the architecture created for the initial version of the model.

Because the architecture was retained but refined in Version 1.1, organizations can compare their Version 1.1 results against those obtained using earlier versions of the model.

Specific Version 1.1 improvements include

- an expanded SGMM Model Definition document

  - The model architecture has been codified and refined to ensure more consistent maturity progression within each domain.

  - Organizations still receive a maturity profile of their rating in each domain but no longer receive a single overall maturity rating.

- A consistent labeling scheme ensures easy mapping among model artifacts.

- New content better describes the SGMM levels and domains.

- New security and critical infrastructure characteristics have been incorporated.

- The characteristics now include more explanatory and educational text as well as more examples for clarification to support consistent understanding and application of the model.

- an updated and refined SGMM Compass survey

  - The new Compass survey includes demographic, scope, and performance questions.

  - Users can move easily between the Compass survey and the Model Definition with a one-to-one mapping between Model Definition characteristics and Compass questions.

  - 80 percent of Compass questions or answer options have been updated to elicit more accurate and consistent responses.

  - 29 new questions were added to support the new characteristics that were added to the model.

- a new SGMM Navigation process

  - The SGMM Navigation Process defines a five-step process for how an organization can use the model to help chart a technical, organizational, and operational path through its grid modernization effort.

  - SGMM Navigators are industry experts trained and certified by the Software Engineering Institute (SEI) to guide utilities through the process and use the outputs in their ongoing planning and implementation.

  - Users of the SGMM Navigation process report finding substantial value in the information sharing and consensus building that occur through the facilitated workshops.

  - This repeatable process also allows for consistent application of the model across markets, organizations, and time and increases the quality of SGMM community data.

---

[1] The lastest release of the SGMM is available at *http://www.sei.cmu.edu/goto/SGMM*.

The SGMM community continues to grow, with nearly 100 utilities having participated to date. The figure below shows their aggregate maturity profile.

In addition to expanding the size of the SGMM community, the SEI is making a concerted effort to increase its diversity. Among the steps taken to elicit broad-based input and participation was the creation of a Stakeholder Panel to represent the full range of SGMM stakeholders.

One question voiced by Panel members was to what extent the SGMM could be useful to all types—investor-owned, publicly owned, cooperative—and sizes of utilities. With the support of the Department of Energy and the American Public Power Association's Demonstration of Energy-Efficient Developments (DEED) research program, the SEI conducted a pilot study using the SGMM Navigation process with American Municipal Power (AMP) and 22 of its member utilities. The participating utilities found that the SGMM provided a common language and framework for discussing smart grid and recommended it for other public power utilities. At the same time, the SEI gained valuable insight into how the SGMM can be made accessible and useful to the public power sector, and it plans to continue to conduct this kind of outreach to the broad spectrum of U.S. utilities.

In 2011, the plans for the model focus on two objectives. The first is to significantly increase the number of utilities that are using the model and to continue to broaden the diversity of that community. As more and more utilities around the world participate and the SGMM experience base grows, it becomes an increasingly valuable resource for helping to inform the industry's smart grid transformation. The second objective is to develop a better understanding of connections between utility performance and domain maturity. This will help us improve and refine the model and lead to the development of better business cases and strategies for smart grid implementation.


Figure 1: Total SGMM community data—average and range

Evolving Incident Management to Incident Resilience

Principal Investigators:
John Haller and Robert Floodeen

Organizations often struggle to maintain operational resilience in the face of incidents that place high-value assets and services under stress. These types of incidents can disrupt the organization's mission and in some cases pose existential threats. The CERT Program's Incident Resilience research area explores how organizations can mature and measure their organic capability to withstand incidents and continue operations with minimal disruption.

Incident management is sometimes viewed as an activity executed by specific organizational units. These may be called the IT Security Department, Network Operations Center, or Computer Security Incident Response Team (CSIRT). This compartmentalized approach can be limiting if the focus is on handling and resolving incidents without being synchronized to the larger mission of the organization.

Incident resilience focuses on expanding the scope of activities that are required to ensure that the organization's mission is not impacted when critical assets are stressed. These activities include

- identifying high-value services and supporting assets to drive the detection, handling, and management of incidents

- analyzing and resolving critical vulnerabilities before they expose assets to stressors

- managing people, technology, and information with an eye toward continuity in the face of adversity

- directly involving a broad range of organizational capabilities—legal, public relations, and human resources, among others—in incident planning and recovery

- learning from incidents to improve the detection and prevention of organizational impact in the future

In an incident-resilient enterprise, all the right people work together. Effective planning has been done and can be executed. Operations continue no matter what.

In 2009 the Software Engineering Institute's (SEI) CERT Program began working with the Department of Homeland Security to develop methodologies and training material for National CSIRTs around the world. As part of this work, CERT is scheduled to release an initial technical note,

Best Practices for National Cybersecurity: Managing a National CSIRT with Critical Success Factors, in 2011. This handbook explores the use of critical success factors as a means to drive incident management—a key principle of an incident resilience approach—and describes best practices for managing and measuring the activity of organizations that typically conduct incident management.

The CERT Program will expand incident resilience research in 2011 by identifying the key activities in enterprises that compose this approach. CERT researchers will then apply this knowledge to areas such as critical infrastructure protection, protecting digital records containing sensitive personnel information, and collaborating with other countries around the world to manage cyber security. One of the principal deliverables in this work is the creation of an incident resilience capability model. Derived from the CERT® Resilience Management Model, this model will provide a means for measuring incident resilience capabilities in organizations that have existing processes for incident management or that are developing these processes.

The focus is on the maturity of organizations facing incidents that place stress on high-value assets and services. Organizations face factors that lead to ever greater uncertainty—from technological complexity to outsourced services, to reliance on intangible assets that drive business. Enterprises must develop the ability not just to handle incidents in the traditional sense, but to develop immunity to them. The Incident Resilience research area will help them.

93

XNET GIVES ORGANIZATIONS EASY AND CONTINUOUS ACCESS TO REALISTIC, HANDS-ON CYBER-TRAINING SCENARIOS AND ENABLES SYNCHRONOUS TEAM-BASED TRAINING THAT CAN SCALE OUT TO HUNDREDS OF PEOPLE THROUGHOUT THE WORLD.

— CHRIS MAY

WORKFORCE DEVELOPMENT

WHILE REFERENCE CURRICULA FOR SOFTWARE ENGINEERING EXIST, INCLUDING GRADUATE SOFTWARE ENGINEERING 2009 (GSWE 2009), AND THE SEI'S EARLIER SOFTWARE ENGINEERING CURRICULUM, THERE ARE FEW GRADUATE SOFTWARE ASSURANCE PROGRAMS OR TRACKS.

WITH THIS NEW MSWA CURRICULUM AVAILABLE, WE HOPE TO HELP EDUCATIONAL INSTITUTIONS CONTINUE TO PROVIDE CUTTING-EDGE PROGRAMS THAT PRODUCE SKILLED SOFTWARE ASSURANCE PROFESSIONALS.

NANCY MEAD

Workforce Development Overview

As cyber attacks become more sophisticated, commanders in the Department of Defense (DoD) are facing a new challenge: how can they evaluate and measure the mission readiness of their cyber warriors? How do they know the available training covers the right mission-essential tasks for their particular teams? Even with seemingly successful training, how do commanders know how staff will perform in the real world?

Chris May, technical manager of the CERT Workforce Development team, explains the problem this way: "There is plenty of training out there. But it usually does not translate well to operational missions where the primary goal is for service members to fight and win in cyberspace. If the troops are trained solely as individuals, without learning how to work as part of a team—like in the real world—how can they be expected to effectively support real environments that are much larger and more complex than the ones they were trained in? Our goal is to bridge that gap and make the training as realistic and accessible as possible."

The Workforce Development team is doing just that through their web-based simulation, training, and evaluation platform called XNET. XNET gives organizations easy and continuous access to realistic, hands-on cyber-training scenarios and enables synchronous team-based training that can scale out to hundreds of people throughout the world.

Moreover, XNET enables DoD organizations to effectively develop and evaluate mission readiness. The Workforce Development team helps DoD organizations identify their mission-essential tasks and then constructs training scenarios and simulation environments in XNET that focus on developing these operational capabilities. Consequently, XNET provides an unparalleled learning experience by providing staff with the same type of events, scenarios, and operating environments they will encounter while on the job and in the fight. "The exercises are customized so they address problems and situations cyber units will actually face, and the XNET platform allows groups to work together to solve those problems in real time," May says.

Another benefit of XNET is immediate assessment. May adds, "Evaluation is often conducted through written tests or observations in a classroom. However, to evaluate the troops' true capabilities outside of the classroom, we need to simulate real-world conditions. That's exactly what XNET allows us to do."

In 2010, the Workforce Development team focused on enhancing XNET functionality to allow people around the world to collaborate online in real time. For example, in June and November, the team participated in the International Cyber Defense Workshop (ICDW), which is sponsored by the Office of the Assistant Secretary of Defense (ASD) for Networks and Information Integration (NII). The goal of the workshop was to challenge the processes and technical response abilities of globally distributed, multi-national cyber defense teams. Participants disbursed across approximately 30 nations collaborated on live-fire incident response and computer forensics scenarios. ASD/NII and the Workforce Development team observed their activities and facilitated a friendly competition across time zones, languages, and cultures. "It was particularly fascinating to assess the progress of teams composed of members from four different countries on three separate continents, which in itself presented an absolutely unique research opportunity that is all but impossible to re-create organically," May says. "It was groundbreaking."

May explains that their observations can be evaluated and researched from a computer science and a psychological perspective. "We were excited to confirm that XNET could support that scale of dynamic collaboration. It's a tricky thing to pull off—to span languages, time zones, and continents, and yet still be able to create a scenario and a platform that enables that kind of seamless communication and collaboration over the internet."

Other educational projects at CERT, such as the first curriculum for a Master of Software Assurance (MSwA) degree program and recommendations for an undergraduate concentration in software assurance, complement the Workforce Development team's training efforts.

2010 has been a remarkable year for the CERT Workforce Development team. "We've raised the bar and the state of practice for cyber security training across the DoD," May says. "XNET makes it possible for cyber units and defense teams to train as they fight on a routine basis."

For more information on XNET, go to *http://xnet.cert.org*.

Software Assurance Curriculum Project

Principal Investigator: Nancy R. Mead

Problem Addressed

Complex software systems affect nearly every aspect of our lives, in areas such as defense, government, energy, communication, transportation, manufacturing, and finance. Protecting these systems against vulnerabilities and attacks is critical, so there is a growing demand for skilled professionals who can build security and correct functionality into software and systems under development. Yet there are few graduate software assurance programs or tracks that focus on developing assured software and, consequently, not enough professionals to meet the growing demand.

Research Approach

Recognizing the importance of software assurance education to meet this demand, CERT researchers collaborated on the software assurance curriculum with a team of educators from Embry-Riddle Aeronautical University, Monmouth University, and Stevens Institute of Technology. The focus of the software assurance curriculum project is to

- identify a core body of knowledge that educational institutions can use to develop Master of Software Assurance (MSwA) degree programs
- mentor universities in developing standalone MSwA degree programs and tracks within existing software engineering and computer science master's degree programs
- promote an undergraduate curriculum specialization for software assurance
- address community college needs

Architectural Structure of an MSwA2010 Degree Program

Preparatory Materials	Computing Foundations Software Engineering Security Engineering
MSwA Core	Assurance Across Life Cycles Risk Management Assurance Assessment Assurance Management System Security Assurance Assured Software Analytics System Operational Assurance
Electives	Courses Related to Assurance in Selected Domains
Capstone Experience	Project

Expected Benefits

The course structure for the MSwA 2010 Reference Curriculum supports the DHS objective of increasing the cyber security workforce by producing more educated graduates of software master's degree programs. This effort, in fact, directly contributes to accomplishing the goal of the National Institute of Standards and Technology (NIST) National Initiative for Cybersecurity Education (NICE) Formal Cybersecurity Education Track initiative—namely "to bolster formal cybersecurity education programs encompassing kindergarten through 12th grade, higher education and vocational programs" (source: *http://csrc.nist.gov/nice/aboutUs.htm*).

2010 Accomplishments

The MSwA Reference Curriculum is the first curriculum developed that focuses on assuring the functionality, dependability, and security of software and systems [1]. While reference curricula exist, including the SEI's groundbreaking software engineering curriculum, no reference curriculum existed that is focused solely on software assurance prior to the development of the MSwA.

The curriculum provides guidelines for a well-rounded education on key security and assurance topics, including assurance across life cycles, risk management, assurance assessment, assurance management, system security assurance, system functionality assurance, and system operational assurance.

Highlights of the curriculum include

- educational outcomes for students who graduate from a program based on the curriculum
- prerequisites expected of students entering an MSwA program
- curriculum architecture for both a standalone degree program and track
- a core body of knowledge that includes the fundamental topics to be taught in the curriculum
- implementation guidelines for educational institutions interested in establishing a program or track based on the curriculum

The MSwA Reference Curriculum has been formally recognized by the two leading computing professional societies, IEEE Computer Society and its partner, the Association for Computing Machinery (ACM) Education Board, as appropriate for a master's program in software assurance. This formal recognition signifies to the educational community that the MSwA Reference Curriculum is suitable for creating graduate programs or tracks in software assurance. The IEEE Computer Society and ACM have developed several computing curricula and are community leaders in curricula development.

In addition to the MSwA Reference Curriculum, the team developed undergraduate software assurance (SwA) course outlines [2]. These courses are intended to provide students

with fundamental skills for either entering the field directly or continuing with graduate-level education.

The team also created sample course outlines for the eight core courses in the MSwA Reference Curriculum (these do not include the capstone project) [3]. More detailed syllabi for the MSwA core courses are available and under review [4]. In addition the team has provided a master bibliography and selected lecture material and other materials to support faculty teaching software assurance. All are available on the CERT website at *http://www.cert.org/mswa/*. The MSwA team will discuss course offerings, review plans, and mentor colleges, universities, and governmental educational institutions at no charge.

To promote incorporation of software assurance information into formal degree programs, the MSwA project team offers flexible options. Educational institutions may choose from the following:

- implement the full reference curriculum to establish a standalone master's program in software assurance

- tailor the materials to offer a software assurance track within an existing graduate program in a related area, such as software engineering or information systems

- use the available undergraduate course outlines to prepare students for a career or additional graduate study in the field of software assurance

Additionally, managers or trainers within organizations may be able to use information from the curriculum to enhance the software assurance capabilities of their existing workforce.

MSwE with SwA Specialization

Preparatory Materials	Computing Foundations Software Engineering Security Engineering
GSwE Core	Ethics and Professional Conduct Systems Engineering Requirements Engineering Software Design Software Construction Software Testing Software Maintenance Configuration Management Software Engineering Management Software Engineering Processes Software Quality
MSwA Core	Assurance Across Life Cycles Risk Management Assurance Assessment Assurance Management System Security Assurance Assured Software Analytics System Operational Assurance
Capstone Experience	Project

Future Goals

Educational institutions have begun incorporating the curriculum into their offerings. Stevens Institute of Technology now offers a master's degree concentration in software assurance. A recent report [5] describes ways of incorporating software assurance content into Master of Science in Information Systems (MSIS) Programs.

The team is currently working on a project to help address community college software assurance needs, by providing a report that includes course outlines and supporting resources. Collaborators in this effort include Embry-Riddle Aeronautical University, Stevens Institute of Technology, and the ACM Two Year College Education Committee (TYCEC). In the future the team hopes to provide similar resources to address high school software assurance needs.

In order to fully transition the MSwA curriculum to educational institutions, there is a need to develop full course materials for the MSwA core courses, including slides, notes, homework assignments, exams, and readings. A corresponding one-semester certificate program should be developed to enhance the software assurance skills of government staff, especially acquisition personnel. Sponsorship for both of these efforts is needed.

References

[1] Mead, Nancy R.; Allen, Julia H.; Ardis, Mark; Hilburn, Thomas B.; Kornecki, Andrew J.; Linger, Rick; & McDonald, James. *Software Assurance Curriculum Project Volume I: Master of Software Assurance Reference Curriculum* (CMU/SEI-2010-TR-005, ESC-TR-2010-005). Software Engineering Institute, Carnegie Mellon University, 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tr005.cfm*

[2] Mead, Nancy R.; Hilburn, Thomas B.; & Linger, Rick. *Software Assurance Curriculum Project Volume II: Undergraduate Course Outlines (*CMU/SEI-2010-TR-019, ESC-TR-2010-019). Software Engineering Institute, Carnegie Mellon University, 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tr019.cfm*

[3] Mead, Nancy R.; Allen, Julia H.; Ardis, Mark; Hilburn, Thomas B.; Kornecki, Andrew J.; & Linger, Rick. *Master of Software Assurance Course Outlines*. Software Engineering Institute, Carnegie Mellon University, 2010. *http://www.cert.org/mswa/docs/MSwACourseOutlines.pdf*

[4] Mead, Nancy R.; Allen, Julia H.; Ardis, Mark; Hilburn, Thomas B.; Kornecki, Andrew J.; & Linger, Rick C. *Software Assurance Curriculum Project Volume III: Master of Software Assurance Course Syllabi* (CMU/SEI-2011-TR-013). Software Engineering Institute, Carnegie Mellon University, 2011 *http://www.sei.cmu.edu/library/abstracts/report/11tr013.cfm*

[5] Shoemaker, Dan; Mead, Nancy R., & Ingalsbe, Jeff. *Integrating the Master of Software Assurance Reference Curriculum into the Model Curriculum and Guidelines for Graduate Degree Programs in Information Systems* (CMU/SEI-2011-TN-004, ESC-TN-2011-004). Software Engineering Institute, Carnegie Mellon University, 2011. *http://www.sei.cmu.edu/library/abstracts/reports/11tn004.cfm*

Researcher Activities

Researcher Activities

Christopher Alberts

Christopher Alberts is a senior cyber security analyst in the CERT® Program at the Software Engineering Institute (SEI). He has 26 years of experience in software engineering and information technology, with particular emphasis on software assurance and cyber security. At the SEI, Alberts leads applied research and development projects in the field of measurement and analysis. He has developed practical and innovative methods, tools, and techniques that have been applied by people throughout government and industry organizations, both nationally and internationally. He has also co-authored two books, *Managing Information Security Risks: The OCTAVE^SM Approach* (Addison-Wesley 2002) and *Continuous Risk Management Guidebook* (Software Engineering Institute 1996). Prior to joining the SEI, Alberts worked at the Carnegie Mellon Research Institute, where he developed autonomous robots for hazardous environments, and at AT&T Bell Laboratories, where he helped automate AT&T's manufacturing processes. He has BS and ME degrees in engineering from Carnegie Mellon University.

Julia H. Allen

Julia Allen is a senior researcher within the CERT® Program at the Software Engineering Institute (SEI), a unit of Carnegie Mellon University in Pittsburgh, PA. Allen's areas of interest include operational resilience, software security and assurance, and measurement and analysis.

Prior to this technical assignment, Allen served as acting director of the SEI for an interim period of 6 months as well as deputy director/chief operating officer for 3 years. Her degrees include a B.Sci. in Computer Science (University of Michigan) and an MS in Electrical Engineering (University of Southern California).

Allen is the author of *The CERT Guide to System and Network Security Practices* (Addison-Wesley 2001) and moderator for the CERT Podcast Series: Security for Business Leaders. She is a co-author of *Software Security Engineering: A Guide for Project Managers* (Addison-Wesley 2008) and *CERT Resilience Management Model (CERT-RMM): A Maturity Model for Managing Operational Resilience* (Addison-Wesley 2010).

Selected Publications

(submitted to Addison-Wesley in July 2010; book published in December 2010) Caralli, Richard A., Allen, Julia H., & White, David W. *CERT® Resilience Management Model: A Maturity Model for Managing Operational Resilience*, Addison-Wesley, 2010.

Caralli, Richard, Allen, Julia, Curtis, Pamela, White, David & Young, Lisa. "Improving Operational Resilience Processes,"

International Workshop on Mission Assurance: Tools, Techniques, and Methodologies. The Second IEEE Internal Conference on Information Privacy, Security, Risk, and Trust (PASSAT 2010), Minneapolis, Minnesota. August 20-22, 2010.

Caralli, Richard A.; Allen, Julia H.; Curtis, Pamela D.; White, David W.; & Young, Lisa R. *CERT® Resilience Management Model, v1.0,* Software Engineering Institute, Carnegie Mellon University, CMU/SEI-2010-TR-012, May 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tr012.cfm*

Allen, Julia H. & Davis, Noopur. *Measuring Operational Resilience Using the CERT® Resilience Management Model,* Software Engineering Institute, Carnegie Mellon University, CMU/SEI-2010-TN-030, September 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tn030.cfm*

Alberts, Christopher, Allen, Julia, & Stoddard, Robert. *Integrated Measurement and Analysis Framework for Software Security,* Software Engineering Institute, Carnegie Mellon University, CMU/SEI-2010-TN-025, September 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tn025.cfm*

Technical Leadership

Moderated and posted 14 podcasts for the CERT Podcast Series: Security for Business Leaders. *http://www.cert.org/podcast*

Workshop organizer and program committee member for the International Workshop on Mission Assurance: Tools, Techniques, and Methodologies at the Second IEEE Internal Conference on Information Privacy, Security, Risk, and Trust (PASSAT 2010), Minneapolis, Minnesota. August 20-22, 2010.

Review board member for a special issue of the *International Journal of Secure Software Engineering* (IJSSE) Vol. 1, No. 4, on software security engineering education. Published October-December 2010.

Archie Andrews

Archie Andrews is a senior member of the technical staff in the Software Engineering Institute at Carnegie Mellon University. He leads the Secure Software and Systems directorate as part of the SEI Networked Systems Survivability Program. The Secure Software and Systems area is developing processes, tools, techniques, and standards for developing secure software and systems and assuring the security properties of these systems. Before joining the SEI, Mr. Andrews was the director of the Information Protection Technology business unit of the Advanced Technology

Institute, a private non-profit research and development management firm located in Charleston, South Carolina. He managed and led a number of projects that addressed information security and privacy protection issues for the Department of Defense, the Veteran's Administration and commercial firms primarily in the healthcare and insurance industries. He retired as a Colonel from the US Army in 1993. His last army assignment was director, US Army Computer Science School in Augusta, Georgia.

Lisa Brownsword

Lisa Brownsword is a senior member of the Acquisition Support Program at the Software Engineering Institute (SEI). She is currently supporting government programs in the application of system and software engineering practices for today's complex, software-reliant systems that interoperate with other critical systems in high threat environments. Lisa co-developed methods to analyze the organizational, governance, and management aspects for Systems of Systems (SoS) environments as a member of the SoS Practices initiative. She co-developed a framework for analyzing the software assurance landscape, with a focus on malicious software management, as part of the SoS Software Assurance (SoSSA) initiative.

Previously, Lisa was a member of the COTS-Based Systems (CBS) initiative where she developed the Evolutionary Process for Integrating COTS-based systems (EPIC). She founded and was the inaugural conference chair for the International Conference on Composition-Based Software Systems (ICCBSS). Lisa has over 20 years of experience in developing large, software-reliant systems along with training and consulting on a variety of software engineering practices. She has authored numerous articles and technical reports and delivered presentations at conferences and workshops worldwide. Recent publications include *A Framework for Modeling the Software Assurance Ecosystem: Insights from the Software Assurance Landscape Project,* an SEI technical report, and *Organizational Implications of Systems of Systems,* a tutorial presented at the NDIA Systems Engineering Conference.

Matthew Butkovic

Matthew Butkovic is an information and infrastructure analyst within the Resilient Enterprise Management Team of the CERT Program at the Software Engineering Institute (SEI). As a member of the team he performs information and critical infrastructure protection research and develops methods, tools, and techniques for resilient enterprise management. Butkovic has more than 15 years of managerial and technical experience in information technology (particularly information systems security, process design and audit) across the banking and manufacturing sectors. Prior to joining CERT in 2010, Butkovic was leading information security and business continuity efforts for a Fortune 500 manufacturing organization. He holds a BA from the University of Pittsburgh. Butkovic is a Certified Information Systems Security Professional (CISSP) and Certified Information Systems Auditor (CISA).

Dawn M. Cappelli

Dawn Cappelli, CISSP, is technical manager of CERT's Enterprise Threat and Vulnerability Management Team at Carnegie Mellon University's Software Engineering Institute. Her team's mission is to assist organizations in improving their security posture and incident response capability by researching technical threat areas; developing information security assessment methods and techniques; and providing information, solutions and training for preventing, detecting, and responding to illicit activity. Team members are domain experts in insider threat and incident response. Team capabilities include threat analysis and modeling; development of security metrics and assessment methodologies; and creation and delivery of training and workshops. Dawn has 30 years experience in software engineering, including programming, technical project management, information security, and research. She is often an invited speaker at national and international venues, is adjunct professor in Carnegie Mellon's Heinz College of Public Policy and Management and currently Vice-Chair for the CERT Computer Security Incident Handler Certification Advisory Board. Before joining CMU in 1988 she worked for Westinghouse as a software engineer developing nuclear power systems.

Selected Publications

Weiland, R.M., Moore, A.P., Cappelli, D.M., Trzeciak, R.F. Spooner, D., "Spotlight On: Insider Threat from Trusted Business Partners," Joint CyLab (CMU) and CERT (SEI), February 2010. *http://www.cert.org/archive/pdf/ TrustedBusinessPartners0210.pdf*

Moore, A.P., D.M. Cappelli, T. Caron, E. Shaw, R.F. Trzeciak, "Insider Theft of Intellectual Property for Business Advantage: A Preliminary Model," in Proc. of the 1st International Workshop on Managing Insider Security Threats (MIST2009), Purdue University, West Lafayette, USA, June 16, 2009. *http://www.cert.org/insider_threat/docs/ Insider_Theft_of_IP_Model_MIST09.pdf*

Cappelli, D.M., Moore, A.P., "Insider Threat Center at CERT Grows Solutions from Reality-Based Research," IA newsletter, Vol 13, No. 2, Spring 2010.

Talks, Panels, and Workshops

Cappelli, D.M., "Strategies for Mitigating Insider Threats; The 50,000 Foot View of Insider Threats," InfoSec World 2009 Conference, Orlando, FL, 7 March 2009.

Cappelli, D.M., "Securing the Weakest Link: Cyber Security Awareness & Education," National Association of State Chief Information Officers (NASCIO) 2009 Annual Conference, Austin, TX, 27 October 2009.

Cappelli, D.M., "The Keys to Successful Monitoring for Detection of Insider Attacks," RSA Conference 2010, San Francisco, CA, 4 March 2010.

Cappelli, D.M., Cummings, A.B., "Monitoring Strategies for Detection of Insider Attacks," GFIRST Conference, San Antonio, TX, 17 August 2010.

Cappelli, D.M., "Using Empirical Insider Threat Case Data to Design a Mitigation Strategy," ACM CCS 2010 Conference, Chicago, IL, 7/8 October 2010.

Technical Leadership

Vice Chair, SEI Computer Security Incident Handler Certification Advisory Board

Top Rated speaker at RSA for the past 3 years

Keynote address at Insider Threat Workshop at the ACM CCS 2010 Conference

Performed numerous reviews of technical papers for IEEE

Sponsored Workshops

Moore, A.P., Cappelli, D.M., "Accelerated Learning to Mitigate Insider Threat (ALtoMIT)," The ALtoMIT Workshop, Co-sponsored with the Air Force Research Laboratory, The Software Engineering Institute, Arlington, VA, 22-23 June 2010.

Cappelli, D.M.., Trzeciak, R.F. "Insider Threat Workshop," Baltimore, MD, 28-30 June 2010.

Cappelli, D.M., Trzeciak, R.F., Montelibano, J. "Insider Threat Workshop," Seattle, WA, 24-28 May 2010.

Cappelli, D.M., Cummings, A.B., "Insider Threat Workshop," Arlington, VA, 3-5 May 2010.

Cappelli, D.M., Trzeciak, R.F. "Insider Threat Workshop," DC3, St. Louis, MO, 21-24 January 2010.

William Casey

William Casey is a senior member of the technical staff of the Software Engineering Institute. He has worked to develop advances in malicious code detection and analysis by developing methods for practical applications. Casey has worked in the areas of threat analysis, code analysis, natural language processing, genomics, bio-informatics, and applied mathematics in academia, industry, and government. Casey received his PhD in applied mathematics from the Courant Institute at New York University. He also holds an MS in mathematics from Southern Illinois University Carbondale, a master's equivalency in computer science from the Courant Institute at New York University, and an MA in mathematics from the University of Missouri Columbia.

Cory F. Cohen

Cory F. Cohen is a senior member of the CERT technical staff, guiding the research and development work of the Malicious Code Analysis team. During his 14 years at CERT, he has worked as a security incident handler, a vulnerability analyst, and a malicious code analyst. His recent work has focused on large-scale automated analysis of malicious code samples collected by CERT.

Prior to joining CERT, Cohen worked for the University of Louisville as HP/UX system administrator in the engineering school where he managed the primary computing cluster.

He also worked for the university as an IDMS/R database administrator maintaining production payroll and student record systems.

Cohen holds a BS in information science and data processing from the University of Louisville.

Rita Creel

Rita Creel has been a principal engineer in CERT since 2010. She works with organizations to apply assurance methods and tools to the acquisition, development, operations, and sustainment of networked, software-reliant systems. Ms. Creel has over 25 years of experience spanning the software and systems life cycle. She has led software development teams, conducted research in software and systems measurement, developed and delivered training in software acquisition and measurement, and collaborated with government agencies to improve the performance and quality of software systems within cost and schedule constraints. Prior to joining CERT, she led a team in the Acquisition Support Program at the Software Engineering Institute providing expertise to the US Intelligence Community. Before the SEI, she worked for The Aerospace Corporation and TRW, Inc. (now Northrop Grumman) focusing on software considerations for space systems and associated ground equipment. She holds BS and MS degrees in Engineering and Computer Science.

Selected Publications

Alberts, C. J., Dorofee, A. J., Creel, R., Ellison, R. J., & Woody, C. (2011, January). A systemic approach for assessing software supply-chain risk, Proceedings of the Hawaii International Conference on System Sciences (HICSS).

Creel, R. (2007). Assuring software systems security: Life cycle considerations for government acquisitions. *https:// buildsecurityin.us-cert.gov/daisy/bsi/articles/best-practices/ acquisition.html*

Creel, R. & Ellison, R. (2008). Acquisition overview: The challenges. *https://buildsecurityin.us-cert.gov/daisy/bsi/ articles/best-practices/acquisition.html*

Creel, R. & Ellison, R. (2008). System-of-systems influences on acquisition strategy development. *https://buildsecurityin. us-cert.gov/daisy/bsi/articles/best-practices/acquisition.html*

Ellison, R. J., Alberts, C. J., Creel, R., Dorofee, A. J. & Woody, C. (2010). *Software supply chain risk management: From products to systems of system*s, CMU/SEI-2010-TN-026. Software Engineering Institute, Carnegie Mellon University.

Ferguson, R., Fowler, S. C., & Creel, R. C. (2009). *A method for assessing technical progress and quality throughout the system life cycle*, CMU/SEI-2009-TN-032. Software Engineering Institute, Carnegie Mellon University.

Adam Cummings

Adam Cummings is currently a member of the technical staff at CERT and a member of the Insider Threat center for three years. This team focuses on insider threat research, threat analysis and modeling, assessments, and training. Adam has over 10 years experience in information systems, information assurance, military communications, project management, and information technology education. He is a former officer in the United States Marine Corps, where he served as a communications officer, as well as a volunteer in the United States Peace Corps, where he served in West Africa. He holds an M.Sc. in Information Security Policy Management from Carnegie Mellon University and a BFA in Visual Journalism from Rochester Institute of Technology.

Talks, Panels, and Workshops

Cappelli, D.M., Cummings, A.B., "Monitoring Strategies for Detection of Insider Threats," GFIRST Conference, August 2010.

Cappelli, D.M., Cummings, A. "Insider Threats and Security Trends: Lessons Learned from Actual Attacks," GFIRST Conference, August 2010.

Technical Leadership

CERT Cybersecurity Compliance Validation (CCV) Team Lead, supporting five Department of Homeland Security (Federal Network Security) assessments of Federal civilian agencies.

Sponsored Workshops

Moore, A.P., Cummings, A. "Insider Threat Workshop": Dallas, TX, 19-20 May 2010.

Cappelli, D.M., Cummings, A.B., "Insider Threat Workshop": Arlington, VA, 4-5 May 2010.

Audrey Dorofee

Audrey Dorofee is a senior member of the technical staff in the Acquisition Support Program at the Software Engineering Institute, Carnegie Mellon. She has worked in the risk management, information security, and process improvement fields for nearly 20 years. Her work at the SEI has included development, training, and transition of advanced risk management methods, tools, and techniques. Her most recent work focuses on managing success and uncertainty in complex environments. Prior to the SEI, she worked for the MITRE Corporation and the National Aeronautics and Space Administration (NASA). She has co-authored two books, *Managing Information Security Risks: The OCTAVE^{SM} Approach* (Addison-Wesley 2002) and the *Continuous Risk Management Guidebook* (Software Engineering Institute 1996).

Robert J. Ellison

Robert J. Ellison is a senior member of the technical staff in the Networked Systems Survivability Program (NSS) at the Software Engineering Institute (SEI) at Carnegie Mellon University (CMU) and a founding member of the SEI. While at the SEI he has served in both technical and management roles. He participates in the evaluation of software architectures and contributes from the perspective of security and reliability measures. As a member of the NSS program, he is one of the developers of the Survivability Analysis Framework (SAF), has been the project leader for the Build Security In project, and currently leads the Supply-Chain Risk project. Dr. Ellison received his MS and PhD in mathematics from Purdue University. He is a member of the IEEE Computer Society and the Association of Computing Machinery (ACM).

Sidney Faber

Sid Faber is a member of the technical staff in the CERT Program at the Software Engineering Institute (SEI). As a member of the Network Situational Awareness (NetSA) analysis team, Faber supports sponsors by providing detailed reports of current and historical network activities. His current areas of interest include fusing massive network data sets, enabling analysts with tools and methods necessary to defend large networks, using large-scale DNS monitoring to detect malicious behavior, and designing closed networks for improved security.

Faber also serves as an adjunct faculty member at the Carnegie Mellon University Heinz College of Information Systems & Management and at the University of Pittsburgh School of Information Sciences.

Prior to joining the SEI, Faber worked as a security architect with Federated Investors, one of the largest investment managers in the United States. His experience includes more than fifteen years in software application security, development, and evaluation, and five years in the U.S. Navy Nuclear Power Officer program.

David A. Fisher

David A. Fisher is a senior research scientist in the Software Engineering Institute at Carnegie Mellon University where he conducts research on next generation information security. Dr. Fisher has held technical and executive positions in academia, industry, and government. His research interests include modeling and simulation, emergent behavior, and automated reasoning especially as they relate to security, HPC, and socio-technical systems. He has degrees in computer science (Ph.D. Carnegie Mellon 1970), electrical engineering (M.S.E. Univ. of Pennsylvania), and mathematics (B.S. Carnegie Mellon), and is a Senior Life Member of the IEEE.

Robert Floodeen

Robert Floodeen is a member of the technical staff, CERT® Resilient Enterprise Management Team in the CERT® Program at the Software Engineering Institute (SEI), a unit of Carnegie Mellon University in Pittsburgh, PA. Robert has 15 years of experience iniInformation security and computer network defense across federal and Department of Defense operations. Before transitioning from an SEI Visiting Scientist to full time technical staff in 2008, Robert led teams performing Intrusion Detection at the Pentagon, Army Research Lab, and for the Defense Research and Engineering Network (DREN). Additionally, he spent several years managing CSIRT operations for the Defense Threat Reduction Agency (DTRA). Robert holds degrees in Computer Science (BS, MS) and is adjunct at Carnegie Mellon University. Robert's entry with Brett Tjaden defining *Cyber Defense* has been accepted and will be printed in the McGraw-Hill Yearbook of Science and Technology, 2011 Edition.

Talks, Panels, and Workshops

Cappelli, D.M., Floodeen, R.W., Trzeciak, R.F., "The Key to Successful Monitoring for Detection of Insider Attacks," RSA Conference, San Francisco, CA, 3 March 2010.

Lead Instructor, "Cyber Readiness Exercise", pre-conference GFIRST, San Antonio, Texas, August 15-16, 2010.

Instructor "Common Cyber Defense Scenarios", International Cyber Defense Workshop, held virtually around the world - 220 participants, June 20–24, 2010.

Facilitator, 1-day workshop, "Exploring Common Cyber Attacks", FIRST Technical Colloquium, Hamburg, Germany, January 27, 2010

Instructor, "Incident Detection and Reporting Challenge", International Cyber Defense Workshop, held virtually around the world – 20 countries, November 8-12, 2009.

Technical Leadership

Education Committee Co-Chair, Forum of Incident Response and Security Teams (FIRST)

Program Committee, Forum of Incident Response and Security Teams (FIRST) 22nd Annual Conference, Miami, Florida, June 13-18, 2010

David French

David French is a senior member of the technical staff of the Malicious Code Advanced Research and Development team at CERT. He has been with CERT since 2010. He currently engages in research into malicious code, including large-scale analysis, fuzzy hashing, and data decomposition.

Prior to joining CERT, French worked for 15 years as part of the Defense Industrial Base as a software engineer and researcher, in a wide range of technical areas, including geospatial information systems, imagery archival and dissemination, visualization, steganography and digital data embedding, signal processing, and most recently, malicious code analysis. He holds a B.S. in Computer Science from Clarkson University.

Selected Publications

Casey, W., Cohen, C., French, D., Hines, C., Havrilla, J., Kinder, R., Diversity Characteristics in the Zeus Family of Malware, Carnegie Mellon University, CMU/SEI-2010-SR-030, Restricted Use

Casey, W., Hines, C., French, D., Cohen, C., Havrilla, J. Application of Code Comparison Techniques Characterizing the Aliser Malware Family, Carnegie Mellon University, CMU/SEI-2010-SR-031, Restricted Use

Talks, Panels, and Workshops

6th Annual US-CERT GFIRST, "Discovering Malicious Code Trends using Bulk Analysis", San Antonio, TX, August 2010

Sponsored Workshops

7th Annual CERT Malicious Code Collaboration Workshop, "Doctor StrangeHash: or, how I learned to stop worrying and love MD5", Arlington, VA, November 2010

John Haller

John Haller is an information and infrastructure security analyst at the CERT Resilient Enterprise Management Team in the CERT Program at the Software Engineering Institute (SEI). Haller is currently responsible for developing and fielding a capability model to improve national and community incident management, and for leading a research program focused on resilient operations. In other work at the SEI, Haller is supporting the development of expert systems to support law enforcement. Haller previously worked in the insurance industry and was a reserve U.S. Army officer. Prior to joining the SEI, Haller worked for the United States Postal Service, first as a cyber-crime analyst and then as a Special Agent for the Office of Inspector General. Haller holds a J.D. degree and a Master of Public and International Affairs degree from the University of Pittsburgh, and is a member of the Pennsylvania bar.

Michael Hanley

Michael Hanley is a member of the technical staff at CERT, part of the Software Engineering Institute. His research interests include insider threats, security metrics, digital forensics, and network security. Prior to joining the SEI, Michael worked for a Fortune 500 company on a large IT services contract in the manufacturing sector. During his tenure there, Michael played a key role in deploying and supporting systems across the globe.

He holds a M.Sc. in Information Security Policy and Management from Carnegie Mellon University and a B.A. in Economics from Michigan State University.

Talks, Panels, and Workshops

Hanley, M., Cappelli, D.M., "Insider Theft of Intellectual Property: A Profile of the Crime," InfoSec World, Orlando, FL, 21 April 2010.

Sponsored Workshops

Trzeciak, R.F., Hanley, M. "Insider Threat Workshop," Arlington, VA, 8/9 September 2010.

Jeffrey S. Havrilla

Jeffrey S. Havrilla has been a senior member of the technical staff at the Software Engineering Institute for over 10 years, primarily focused on software security engineering. His current area of work is analyzing malicious code and artifacts associated with computer security intrusions. Havrilla was previously the technical leader of the CERT/CC vulnerability analysis team, part of the CERT/CC focused on finding software vulnerabilities in deployed software using both static and dynamic analytical tools. Prior to working at the SEI, Havrilla worked at the University of Pittsburgh Medical Center and School of Medicine as a largescale database, network and research systems administrator and programmer. Havrilla has a Master's of Science in Telecommunications from the University of Pittsburgh School of Information Sciences, and is a member of the IEEE Computer Society and Internet Society (ISOC).

Selected Publications

Casey, W., Cohen, C., French, D., Hines, C., Havrilla, J., Kinder, R., Diversity Characteristics in the Zeus Family of Malware, Carnegie Mellon University, CMU/SEI-2010-SR-030, Restricted Use

Casey, W., Hines, C., French, D., Cohen, C., Havrilla, J. Application of Code Comparison Techniques Characterizing the Aliser Malware Family, Carnegie Mellon University, CMU/SEI-2010-SR-031, Restricted Use

Sponsored Workshops

7th Annual CERT Malicious Code Collaboration Workshop, "Bulk Analysis of Malicious PDF Objects", Arlington, VA, November 2010

Charles Hines

Hines is a senior malware researcher on the Malicious Code R&D team, which he joined early in 2010. He has been assisting in the code comparison efforts, primarily focusing on large-scale function-level comparisons, along with some side experiments in visualization of these (and similar) comparisons and experiments on alternate methods of storing and searching these large amounts of data.

Prior to joining CERT, Hines was involved in a variety of things, including malware-related R&D contracts for the Air Force Research Laboratory (AFRL), imagery-related work on contract for the National Geospatial-Intelligence Agency (NGA), and internal logic synthesis projects for IBM, as well as contributions to various open source projects in his spare time. He has a BS in computer engineering from Clarkson University.

David Keaton

David Keaton is a compiler writer with a background ranging from high-performance computing to embedded systems. He has six patents pending in computer architecture, and two patents in compiler-assisted security mechanisms. David is chairman of the U.S. committee standardizing the C programming language.

Christopher King

Christopher King is a member of the Insider Threat Center in the CERT Program, part of the Software Engineering Institute at Carnegie Mellon University. At CERT, Chris is researching insider threat technical controls, developing insider threat assessments, and analyzing cases of insider crime. Before coming to CERT, Chris worked at the National Security Agency in the Enterprise IA Systems Engineering Services division as an IT specialist. While at NSA, Chris researched new CND architectures utilizing non-persistent technologies.

Previously, Chris worked at the Defense Information Systems Agency as an information assurance manager and capability module team lead for the Net-Enabled Command Capability (NECC) program. Through his tenure there, Chris directed development of the first three capability modules through certification and accreditation, testing, and delivery. He also maintained the security accreditation of other C2 programs, assisted in the development of NECC's information security architecture, and contributed to the Global Command and Control System – Joint (GCCS-J) program.

Selected Publications

Manion, A., Togashi, K., Kousaka, F., Yamaguchi, M., McCaffrey, S., Kadane, J., King, C., Weiland, R., "Effectiveness of the Vulnerability Response Decision Assistance (VRDA) Framework," GFIRST Conference, 23-28 August 2009, Atlanta, GA.

Linda Levine

Linda Levine is a senior member of the technical staff at Carnegie Mellon University's Software Engineering Institute. Her research focuses on acquisition of software intensive systems; reasoning, systems thinking and patterns of failure; agile software development; diffusion of innovations; and knowledge integration and transfer. For a recent publication see Novak, W. and Levine, L. (2010). Success in acquisition: Using archetypes to beat the odds. (SEI Technical Report CMU/SEI-2010-TR 016). Pittsburgh, PA: Software Engineering Institute. *http://www.sei.cmu.edu/library/abstracts/reports/10tr016.cfm*

Levine holds a PhD from Carnegie Mellon University. She is a member of the IEEE Computer Society, Association for Information Systems, National Communication Association, and cofounder and Chair of IFIP Working Group 8.6 on Diffusion, Transfer and Implementation of Information Technology. Contact her at *ll@sei.cmu.edu*.

Howard Lipson

Howard Lipson is a senior member of the technical staff in the CERT Program at the SEI. Dr. Lipson has been a computer security researcher at CERT for 18 years. He is also an adjunct professor in CMU's Department of Engineering and Public Policy and an adjunct research faculty member at the Carnegie Mellon Electricity Industry Center. He has played a major role in developing the foundational concepts and methodologies necessary to extend security research into the new realm of survivability, and has been a chair of four IEEE research workshops on survivability. His research interests include the

analysis and design of survivable systems and architectures, software assurance, and critical infrastructure protection (primarily for the smart grid). Prior to joining CMU, Dr. Lipson was a systems design consultant. Earlier, he was a computer scientist at AT&T Bell Labs. He holds a Ph.D. in Computer Science from Columbia University.

Selected Publications

Highfill, D. (ed.), Bass, L., Brown, B., Brown, K., Carpenter M., Ivers, J., Kuruganti, T., Lipson, H., Nutaro, J., Searle, J., Shah, V., Smith, B., and Stevens, J., *Security Profile for Advanced Metering Infrastructure – Version 2.0*, Advanced Security Acceleration Project for the Smart Grid (ASAP-SG), UCA International Users Group, June 2010, 151 pp.

The SGMM Team, *Smart Grid Maturity Model – Model Definition (Version 1.1)*, CMU/SEI-2010-TR-009, Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, September 2010.

Talks, Panels, and Workshops

General Chair and Program Co-Chair, IEEE Smart Grid Survivability Workshop (CISW-SG 2010), October 2010.

Lipson, H. "Built for Survival – Defining and Achieving Survivability of Complex Systems," IEEE Smart Grid Survivability Workshop, October 2010.

Technical Leadership

General Chair and Program Co-Chair, IEEE Smart Grid Survivability Workshop, October 2010.

Program committee member, ACM Workshop on Digital Identity Management (DIM), ACM Conference on Computer and Communications Security, October 2009 & October 2010.

Program committee member, International Workshop on Security Measurements and Metrics (MetriSec), October 2009 & September 2010.

Program committee member, Workshop on Ethics in Computer Security Research (WECSR), January 2010.

Invited member, NERC-DOE Task Force on Coordinated (Cyber/Physical) Attack – High-Impact, Low Frequency Event Risk to the North American Bulk Power System. (Task force was active January to May 2010 to support preparation of a report.)

Nancy R. Mead

Nancy R. Mead is a senior member of the technical staff in the CERT Software Security Assurance group. Mead is also a faculty member in the Master of Software Engineering and Master of Information Systems Management programs at Carnegie Mellon University. She is currently involved in the study of security requirements engineering and the development of software assurance curricula.

Prior to joining the SEI, Mead was a senior technical staff member at IBM Federal Systems, where she developed and managed large real-time systems, worked in software engineering technology, and managed IBM Federal Systems' software engineering education department.

Mead has more than 150 publications and invited presentations. She is a Fellow of IEEE and the IEEE Computer Society and a Distinguished Member of the ACM. Mead received her PhD in mathematics from the Polytechnic Institute of New York, and a BA and an MS in mathematics from New York University.

Selected Publications

Mead, N.R., Shoemaker, D., Book Chapter "Novel Methods of Incorporating Security Requirements Engineering into Software Engineering Courses and Curricula," Chapter VI, *Software Engineering: Effective Teaching and Learning Approaches and Practices,* Eds, Ellis, Demurjian, & Naveda, IGI Global, pp. 98-113, 2008

Mead, Nancy R.; Allen, Julia H.; Ardis, Mark; Hilburn, Thomas B.; Kornecki, Andrew J.; Linger, Rick; & McDonald, James. *Software Assurance Curriculum Project Volume I: Master of Software Assurance Reference Curriculum* (CMU/SEI-2010-TR-005, ESC-TR-2010-005). Software Engineering Institute, Carnegie Mellon University, 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tr005.cfm*

Mead, Nancy R.; Hilburn, Thomas B.; & Linger, Rick. *Software Assurance Curriculum Project Volume II: Undergraduate Course Outlines* (CMU/SEI-2010-TR-019, ESC-TR-2010-019). Software Engineering Institute, Carnegie Mellon University, 2010. *http://www.sei.cmu.edu/library/abstracts/reports/10tr019.cfm*

Ashwini Bijwe, Nancy Mead, Faculty Advisor, *Adapting the Square Process for Privacy Requirements Engineering,* CMU/SEI-2010-TN-022 Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, July, 2010

Mead, N.R., Allen, J.H., *Building Assured Systems Framework,* CMU/SEI-2010-TR-025 Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, September, 2010

Talks, Panels, and Workshops

BrightTalk Webinar on April 8, 2010 titled "Advances in Privacy Requirements Engineering".

IACBP invited talk and SQUARE tool demo on July 22, 2010 "Squaring up your security requirements with SQUARE"

CyLab Seminar "SQUARE and Privacy Requirements Engineering", September 20 , 2010

How to Get Started in Software Assurance Education, Co-presenter, CSEET Workshop March 2010

Abu-Nimeh, S., Mead, N,R., Combining Privacy and Security Risk Assessment in Security Quality Requirements Engineering, Intelligent Privacy Management Symposium, Stanford University, March 2010

Technical Leadership

The first Nancy Mead Award for Excellence in Software Engineering Education and Training was presented to Mary Shaw on March 10, 2010 at the CSEET Conference Dinner.

REJ (Requirements Engineering Journal), Editorial Board

Editorial Review Board – International Journal on Secure Software Engineering (IJSSE), IGI Global

ICSE 2010 Tutorial Program Committee

CSEET 2010 Program Committee

Samuel Merrell

Samuel Merrell is a member of the technical staff on the Resilient Enterprise Management Team at CERT. Merrell works with organizations to improve their information security risk management capabilities. This work includes Critical Infrastructure Protection projects within the Department of Homeland Security and analysis of federal (DoD and civilian agency) information security programs, including Federal Information Security Management Act (FISMA) compliance efforts. Recent projects include assisting in the development of the CERT Resilient Enterprise Framework and evaluating Service Oriented Architecture initiatives within the U.S. Military. Prior to joining the SEI, Merrell spent seven years as the Information Technology Manager for a Pittsburgh-area community bank. Before that, he was an information technology consultant, primarily support the IBM AS/400. Merrell holds an undergraduate degree from the University of Pittsburgh, the Certified Information Systems Security Professional (CISSP) certification, and a number of SANS certificates, and is currently working towards a master's degree in Information Security at Carnegie Mellon University.

Soumyo D. Moitra

Soumyo Moitra is a senior member of the technical staff with the Network Situational Awareness Group at CERT/SEI. He has been involved with modeling and analyzing network traffic for security and monitoring. He is currently working on metrics for the cost-effectiveness of network sensors and modeling network security operations.

Joji Montelibano

Joji Montelibano is a member of the Insider Threat team at CERT. He has over 15 years experience in the fields of software development and network engineering. He began his career developing software for the petroleum and chemical industries, where he created customized simulation programs for companies such as Shell Oil, Sunoco, and Foster Wheeler. Prior to joining CERT, Joji was a senior information security analyst for the RAND Corporation, where his main projects focused on securing and ensuring the availability of military networks and communications. He holds an undergraduate degree in Chemical Engineering from Stanford University, and Master's degrees from Harvard University and the University of Southern California. His certifications include the CISSP, CSTE, CCNP, and ACSA.

Selected Publications

Hanley, M., Dean, T., Schroeder, W., Houy, M., Trzeciak, R.F., and Montelibano, J. "An Analysis of Technical Observations in Insider Theft of Intellectual Property Cases" SEI Technical Note CMU/SEI-2011-TN-006, Carnegie Mellon University, February 2011.

Porche, I.R., Montelibano, J., Comanor, K., Wilson, B., Rothenberg, J., and Schneider, M.J., "Navy Network Dependability: Models, Metrics, and Tools," RAND, 2010.

Talks, Panels, and Workshops

Montelibano, J., "Policy, Legal and Privacy Issues in the Fight against Insider Threats," Security 2010, 11th Annual Security Conference and Exhibition, Walter E. Washington Convention Center, Nov. 16-17, 2010.

Montelibano, J., "The Need for Robust Statistical Analysis of MANET Performance Data," 15th International Command and Control Research and Technology Symposium, Fairmont Miramar Hotel, Santa Monica, CA, June 24, 2010.

Sponsored Workshops

Cappelli, D.M., Trzeciak, R.F., Montelibano, J. "Insider Threat Workshop," Seattle, WA, 24-28 May 2010.

Andrew P. Moore

Andrew P. Moore is a senior member of the CERT technical staff. Moore explores ways to improve the security, survivability, and resiliency of enterprise systems through insider threat and defense modeling, incident processing and analysis, and architecture engineering and analysis. Before joining the SEI in 2000, he worked for the Naval Research Laboratory (NRL) investigating high-assurance system development methods for the Navy. He has over twenty years experience developing and applying mission-critical system analysis methods and tools, leading to the transfer of critical technology to both industry and the military. Moore received his BA in Mathematics from the College of Wooster and MA in Computer Science from Duke University.

Selected Publications

Moore, A.P., D.M. Cappelli, T. Caron, E. Shaw, R.F. Trzeciak, "Insider Theft of Intellectual Property for Business Advantage: A Preliminary Model," in Proc. of the 1st International Workshop on Managing Insider Security Threats (MIST2009), Purdue University, West Lafayette, USA, June 16, 2009. *http://www.cert.org/insider_threat/docs/Insider_Theft_of_IP_Model_MIST09.pdf*

Merrell, S., Moore, A. P., Stevens, J., "Goal-Based Assessment for the Cybersecurtiy of Critical Infrastructure," in Proc. of the 2010 IEEE International Conference on Technologies for Homeland Security (*http://ieee-hst.org/*), Waltham, MA, 8-10 November 2010.

Brownsword, L., Woody, C., Alberts, C.J., Moore, A.P., *A Framework for Modeling the Software Assurance Ecosystem: Insights from the Software Assurance Landscape Project,* Software Engineering Institute Technical Report (CMU/SEI-2010-TR-028), Carnegie Mellon University, August 2010. *http://www.sei.cmu.edu/reports/10tr028.pdf*

Weiland, R.M., Moore, A.P., Cappelli, D.M., Trzeciak, R.F. Spooner, D., "Spotlight On: Insider Threat from Trusted Business Partners", Joint CyLab (CMU) and CERT (SEI), February 2010. *http://www.cert.org/archive/pdf/ TrustedBusinessPartners0210.pdf*

Cappelli, D.M., Moore, A.P., "Insider Threat Center at CERT Grows Solutions from Reality-Based Research," IA newsletter, Vol 13, No. 2, Spring 2010.

Talks, Panels, and Workshops

Moore, A.P., "MERIT Interactive: Training Simulation for Mitigating Insider Threat," at the Accelerated Learning to Mitigate Insider Threat (ALtoMIT) Workshop, Co-sponsored with the Air Force Research Laboratory, The Software Engineering Institute, Arlington, VA, 22-23 June 2010.

Cappelli, D.M., Moore, A.P., "Lessons Learned from Actual Insider Attacks," at the Accelerated Learning to Mitigate Insider Threat (ALtoMIT) Workshop, Co-sponsored with the Air Force Research Laboratory, The Software Engineering Institute, Arlington, VA, 22-23 June 2010.

Moore, A.P., "Summary of Workshop on Accelerated Learning to Mitigate Insider Threat (ALtoMIT), Held 22-23 June 2010, Dagstuhl Seminar on Insider Threats, 22-26 August 2010.

Moore, A.P., "What Organizations Need to Know about Insider Cyber Fraud," BAI Combating Payments Fraud Conference, March 2010, Orlando, FL.

Moore, A.P., "DHS USSS Insider Threat Study: Research Methodology," DHS S&T Project Review Meeting, May 2010, Washington D.C.

Sponsored Workshops

Moore, A.P., Cummings, A. "Insider Threat Workshop": Dallas, TX, 19-20 May 2010.

Cappelli, D.M., Moore, A.P. "Insider Threat Workshop": Arlington, VA, 16-17 September 2009

Technical Leadership

Organized "Accelerated Learning to Mitigate Insider Threat (ALtoMIT) Workshop," Co-sponsored with the Air Force Research Laboratory, The Software Engineering Institute, Arlington, VA, 22-23 June 2010.

David Mundie

David Mundie is a member of the CSIRT Development Team in CERT. He has been at CERT since the year 2000, and has worked in a variety of projects including the EASEL simulation language, the K-BACEE knowledge-based automated component ensemble evaluation tool, a study of anomaly detection tool usability, function extraction, a study of insider threats from online social networks, and incident management capability metrics. As a member of the CDT, he develops and delivers workshops offered to CSIRT managers and incident handling staff. From 2006 to 2009 he was manager of outreach and training in the Q-CERT project which established a national information security team for Qatar. His current research interests include an incident management model for national teams, insider threat patterns, and an ethnographic study of malware analysis. Prior to joining CERT he worked at Texas Instruments and Western Digital on compiler development, test engineering, and process improvement.

Robin Ruefle

Robin Ruefle is the team lead for the CERT® CSIRT Development and Training (CDT) team. Her focus is on the development of management, procedural, and technical guidelines and practices for the establishment, maturation, operation, and evaluation of CSIRTs worldwide. As a member of the CDT, Ruefle develops and delivers sessions in the suite of courses offered to CSIRT managers and incident handling staff and has co-authored a variety of CSIRT publications include *Handbook for CSIRTs 2nd Edition, Organizational Models for CSIRTs Handbook, CSIRT Services, State of the Practice of CSIRTs, Defining Incident Management Processes for CSIRTs: A Work in Progress, The Role of Computer Security Incident Response Teams in the Software Development Life Cycle,* as well as numerous other articles and best practice guides. Current work includes development of a training and mentoring framework and BOK for incident management and the development of a process model for incident management.

Philip Scolieri

Philip Scolieri holds the position of information and infrastructure analyst within the Resilient Enterprise Management Team of the CERT Program at the Software Engineering Institute (SEI). As a member of the team he performs information and critical infrastructure protection research and develops methods, tools, and techniques for resilient enterprise management. Scolieri has over 25 years managerial and technical experience in both the engineering and information technology fields (particularly systems design and analysis, data and telecommunications, information systems infrastructure management and information systems security) across government and manufacturing sectors. Prior to joining CERT in 2010, Scolieri was the manager of security and compliance and leading disaster recovery efforts for a Fortune 500 manufacturing organization. He holds an MS and BS degree in Electrical Engineering from the University of Pittsburgh.

Robert C. Seacord

Robert C. Seacord leads the Secure Coding Initiative at CERT, located in Carnegie Mellon's Software Engineering Institute (SEI) in Pittsburgh, PA. CERT, among other security related activities, regularly analyzes software vulnerability reports and assesses the risk to the Internet and other critical infrastructure. Robert is an adjunct professor in the Carnegie Mellon University School of Computer Science and at the Information Networking Institute. He represents CMU at PL22.11 (ANSI "C") and is a technical expert for the JTC1/SC22/WG14 international standardization working group for the C programming language.

George J. Silowash

George J. Silowash is a cybersecurity threat and incident analyst in CERT at Carnegie Mellon University's Software Engineering Institute (SEI). He is part of the Threat Technical Solutions and Standards team. He has over nine years experience in the information technology field, including systems administration and information security. His current work includes insider threat research in the financial sector, researching technology for assisting in the detection of insider threats, and developing information security controls to enhance the security posture of government agencies. Other areas of interest include privacy and security, digital forensic investigations, and critical infrastructure security. Before joining CERT, Silowash was an information technology specialist focusing on information security for the United States Department of Justice, National Drug Intelligence Center. He was also a systems administrator for a healthcare company prior to working in the federal government. George is also an adjunct professor at Norwich University's Information Assurance Program. Mr. Silowash has a Master of Science in Information Assurance from Norwich University and is a Certified Information Systems Security Professional (CISSP).

Derrick Spooner

Derrick Spooner is currently an information security analyst at CERT. He is a critical member of the insider threat center which focuses on insider threat research, threat analysis and modeling, assessments, and training. He holds an MS in Information Security Policy Management from Carnegie Mellon University and a BA in Information Technology Leadership from Washington & Jefferson College.

Selected Publications

Weiland, R.M., Moore, A.P., Cappelli, D.M., Trzeciak, R.F., Spooner, D., "Spotlight On: Insider Threat from Trusted Business Partners", Joint CyLab (CMU) and CERT (SEI), February 2010. *http://www.cert.org/archive/pdf/TrustedBusinessPartners0210.pdf*

James Stevens

James Stevens is the team leader of the CERT Infrastructure Resilience Team in the CERT Program at the Software Engineering Institute (SEI). Stevens is responsible for developing, executing, and managing a research portfolio focused on improving the resilience of national infrastructures. Stevens is the technical lead for the SEI Smart Grid Maturity Model, a management tool that allows utilities to plan, quantifiably measure progress, and prioritize options as they move towards the realization of a smart grid. In other work at the SEI, Stevens has been active in developing and delivering various information security risk assessment, analysis, and management technologies for customers in the federal government and the private sector. Stevens has over 15 years experience in information technology (particularly information systems management and security). Stevens holds a BS degree in Electrical Engineering from the University of Notre Dame and an MBA from Carnegie Mellon University. Stevens is an IEEE Member and is a Certified Information Systems Security Professional (CISSP).

Robert Stoddard

Robert Stoddard joined the Software Engineering Institute (SEI) as a senior member of the technical staff in 2005 and is currently serving as a predictive analytics leader within CERT and SEPM CMMI High Maturity. Robert is a Motorola-certified Six Sigma Master Black Belt and holds American Society for Quality certifications in CSSBB, CRE, CQE, CSQE, and CQA. Robert earned a B.S. in Finance and Accounting, an M.S. in Systems Management, and completed most PhD course work in Reliability Engineering. Prior to the SEI, Robert served as a Motorola Quality Director and Distinguished Member of the Technical Staff (7 yrs), a Texas Instruments Software Quality Manager (14 yrs), and a US Army Finance Automation Officer.

Dean F. Sutherland

Dean F. Sutherland is a senior software security researcher at CERT. Dean spent 14 years working as a professional software engineer at Tartan, Inc. He spent the last 6 of those years as a senior member of the technical staff and a technical lead for compiler back-end technology. He was the primary active member of the corporate R&D group, was a key instigator of the design and deployment of a new software development process for Tartan, led R&D projects, and provided both technical and project leadership for the 12-person compiler back-end group. He received his Ph.D. in Software Engineering from Carnegie Mellon in 2008.

Selected Publications

Dissertation

Dean F. Sutherland. The Code of Many Colors: Semi-automated Reasoning about Multi-Thread Policy for Java. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA 15213, May 2008.

Refereed Papers

Dean F. Sutherland, William L. Scherlis, "Composable Thread Coloring", Proceedings of the 15th ACM Symposium on Principles and Practice of Parallel Programming, Jan. 2010. One of three nominees for the conference Best Paper award. (17% acceptance rate)

D. Sutherland, A. Greenhouse, W. Scherlis, "The Code of Many Colors: Relating Threads to Code and Shared State", Program Analysis for Software Tools and Engineering, Oct. 2002. (35% acceptance rate)

David Svoboda

David Svoboda is a software security engineer at CERT, at the Software Engineering Institute (SEI) in Pittsburgh, PA. David has been the primary developer on a diverse set of software development projects at Carnegie Mellon since 1991, ranging from hierarchical chip modelling and social organization simulation to Automated Machine Translation (AMT). His KANTOO AMT software, developed in 1996, is still (as of 2008) in production use at Caterpillar. David is also actively involved in several ISO standards groups: the JTC1/SC22/WG14 group for the C programming language, and the JTC1/SC22/WG21 group for C++.

Selected Publications

Mitamura, Baker, Nyberg, and Svoboda. "Diagnostics for Interactive Controlled Language Checking." Proceedings of EAMT/CLAW 2003.

Mitamura, Nyberg, Torrejon, Svoboda, Brunner, and Baker. "Pronominal Anaphora Resolution in the KANTOO Multilingual Machine Translation System." Proceedings of TMI 2002.

Many more KANT-related publications are available at *http://www.lti.cs.cmu.edu/Research/Kant/.*

Carley and Svoboda. "Modeling Organizational Adaptation as a Simulated Annealing Process." *Sociological Methods and Research*, August 1996, Vol. 25, No. 1: pp. 138-168.

Walker, Kellen, Svoboda, and Strojwas (1993). "The CDB/HCDB Semiconductor Wafer Representation Server," *Computer-Aided Design of Integrated Circuits and Systems*, February 1993, Vol. 12, No. 2, pp. 283-295.

Randy F. Trzeciak

Randy F. Trzeciak is currently a senior member of the technical staff at CERT. He is the technical team lead of the Insider Threat Outreach and Transition group in the Insider Threat Center at CERT; a team focusing on insider threat research; threat analysis and modeling; assessments; and training. Randy has over 20 years experience in software engineering, database design, development, and maintenance, project management, and information security. Before joining Carnegie Mellon University, Randy worked for Software Technology Incorporated, in Alexandria VA, as a consultant to the Naval Research Laboratory (NRL). He also is an adjunctprofessor at Carnegie Mellon's Heinz College, School of Information Systems and Management. Randy holds an MS in Management from the University of Maryland and a BS in Management Information Systems and a BA in Business Administration from Geneva College.

Selected Publications

Weiland, R.M., Moore, A.P., Cappelli, D.M., Trzeciak, R.F. Spooner, D., "Spotlight On: Insider Threat from Trusted Business Partners", Joint CyLab (CMU) and CERT (SEI), February 2010. *http://www.cert.org/archive/pdf/TrustedBusinessPartners0210.pdf*

Moore, A.P., D.M. Cappelli, T. Caron, E. Shaw, R.F. Trzeciak, "Insider Theft of Intellectual Property for Business Advantage: A Preliminary Model," in Proc. of the 1st International Workshop on Managing Insider Security Threats (MIST2009), Purdue University, West Lafayette, USA, June 16, 2009. *http://www.cert.org/insider_threat/docs/Insider_Theft_of_IP_Model_MIST09.pdf*

Talks, Panels, and Workshops

Trzeciak, R.F., "Defense Industrial Base: Technical Exchange", St. Louis, MO, 22 January 2010.

Trzeciak, R.F., "The Key to Successful Monitoring for Detection of Insider Attacks," RSA Conference, San Francisco, CA, 3 March 2010.

Trzeciak, R.F., "Risk Mitigation Strategies: Lessons Learned from Actual Insider Attacks," FS-ISAC Conference, Tampa, FL, 4/5 May 2010.

Trzeciak, R.F., "Understanding the Insider Threat: Lessons Learned from Actual Insider Attacks," FIRST Conference, Miami, FL, 14 June 2010.

Trzeciak, R.F., "Insider Threats and Security Trends: Lessons Learned from Actual Insider Attacks," GFIRST Conference, San Antonio, TX, 18 August 2010.

Sponsored Workshops

Cappelli, D.M., Trzeciak, R.F., "Insider Threat Workshop", DC3, St. Louis, MO, 23 January 2010.

Trzeciak, R.F., Longo, G. "Insider Threat Workshop," Arlington, VA, 23/24 March 2010.

Cappelli, D.M., Trzeciak, R.F., "Insider Threat Workshop," TMI, Arlington, VA, 13 April 2010.

Cappelli, D.M., Trzeciak, R.F., "Insider Threat Workshop", Baltimore, MD, 29/30 June 2010.

Trzeciak, R.F., Hanley, M. "Insider Threat Workshop," Arlington, VA, 8/9 September 2010.

George Warnagiris

George Warnagiris is a member of the technical staff in the CERT Program at the Software Engineering Institute (SEI). Warnagiris has done work in network monitoring, analysis, and training as part of CERT's Network Situational Awareness (NetSA) team. Prior to joining the SEI, he worked as a network engineer in the financial industry. Warnagiris holds a Bachelor of Arts degree in Computer Science from the City University of New York and is currently pursuing an advanced degree at Carnegie Mellon University.

Rhiannon Weaver

Rhiannon Weaver is a member of the technical staff for the Network Situational Awareness Group. She holds a BS in Mathematics and a BS in Computer Science from Penn State University, and a MS in Statistics from Carnegie Mellon University, where she is also pursuing her PhD in statistics. Weaver provides support for advanced modeling techniques for network anomaly detection and large-scale trending of Internet-wide phenomena. Her research interests include time series analysis of network data, data collection and inference in hierarchical and Bayesian models, as well as addressing the challenges of evaluating and applying advanced modeling and data mining techniques in operational environments.

David W. White

David W. White is a senior member of the technical staff at CERT. White is responsible for developing and implementing strategies that lead to the widespread dissemination and use of methods, techniques, and tools that help organizations manage information security risks. He is also a member of the development team for the CERT Resiliency Engineering Framework, a process improvement framework that provides guidelines for managing security and business continuity from an enterprise risk management perspective. White has a bachelor's degree in Civil Engineering and Public Policy from Carnegie Mellon University and a master's degree in Civil Engineering with a specialization in robotics from Carnegie Mellon University. He is currently based in New York City.

Carol Woody

Dr. Carol Woody has been a senior member of the technical staff since 2001. Currently she is the technical lead of the Survivability Analysis team, whose research focuses on cyber security engineering: building capabilities in defining, acquiring, developing, measuring, managing, and sustaining secure software for highly complex networked systems as well as systems of systems.

Dr. Woody has over 25 years of experience in software development and project management covering all aspects of software and systems planning, acquisition, design, development, and implementation in large complex organizations. Woody has a biographical citation in Who's Who in American Women. She is a senior member of IEEE and ACM, and a member of PMI and AIAA.

Dr. Woody holds a BS in mathematics from The College of William and Mary, an MBA with distinction from Wake Forest University, and a PhD in Information Systems from NOVA Southeastern University.

Selected Publications

Ellison, R. & Woody, C., "Supply-Chain Risk Management: Incorporating Security into Software Development", Proceedings of the 43rd Hawaii International Conference on System Sciences (HICSS) (CD-ROM), January 5-8, 2010, Computer Society Press, 2010

Woody, C. & Levine, L. "System of Systems Analysis of Catastrophic Events", Proceedings of the IEEE Homeland Security Technologies Conference (CD-ROM), November 2010

Ellison, R. & Woody, C., "Considering Software Supply-Chain Risks", CrossTalk, Vol. 23, No 5, September/October, 2010, pp 9-12

Brownsword, L, Woody, C., Alberts, C., & Moore, A., A Framework for Modeling the Software Assurance Ecosystem, Carnegie Mellon University CMU/SEI-2010-TR-028, August 2010

Ellison, R. & Woody, C., *Survivability Analysis Framework*, Carnegie Mellon University, CMU/SEI-2010-TN-013, June 2010

Talks, Panels, and Workshops

DHS SwA Forum, "The Way Forward for Mitigating Software Supply Chain Risk", Arlington VA, November 2009

DHS SwA, "SEI Measurement Project", McLean VA, December 2010

Annual Security Applications Conference (ACSAC), "Supply Chain Risk Management Framework", Austin TX, December 2010

Military Operations Research Society (MORS) Symposium, Quantico VA, "Analysis of Catastrophic Events: Preliminary Investigation", June 2010

Systems Engineering Research Center (SERC) Workshop, Software Security Engineering, Washington DC, April 2010

Technical Leadership

Elected IEEE Senior Member, August 2009

Elected ACM Senior Member, August 2009

Panel Facilitator, DHS SwA Forum, "Where is Academia Going and How Can the SwA Forum Help", Arlington VA, November 2009

DHS Software Assurance Working Group (SwA), "Software Assurance Principles", McLean VA, December 2010

Evan Wright

Evan is an analyst for the Network Situational Awareness Team (NetSA). Evan's research interests include next generation technologies, network design, routing protocols, and design of network attack tools. Prior to joining the SEI, Wright completed graduate school at Carnegie Mellon, where he obtained his MS in Information Security and Technology Management from the School of Engineering. He also holds a BS in Technology Systems from East Carolina University. Wright worked as a Network Administrator at ABC Phones in North Carolina and as a consultant for various other companies. Evan holds the Cisco Certified Networking Professional certificate and four other IT certifications.

Dave Zubrow

Dave manages the Software Engineering Measurement and Analysis initiative at the SEI, consults, and delivers professional education on measurement and analysis. He is an instructor for the Implement Goal Driven Measurement, the SEI's Six Sigma related courses, and Introduction to CMMI. He is also a certified SCAMPI High Maturity Lead Appraiser. His current research and development activities involve automated data anomaly detection as well as diagnostic assessments in the security area. Dave is a senior member of the American Society for Quality, a Certified Software Quality Engineer, and editorial board member and reviewer for several professional journals. Dave has helped numerous organizations establish and improve their use of measurement over the years. He is an avid scuba diver and rollerblader and enjoys live music.

New Researchers at CERT

Gregory Shannon

PhD, Purdue University, Computer Sciences

William Casey

PhD, Courant Institute New York University, Mathematics

Leigh Metcalf

PhD, Auburn University, Mathematics

Dean F. Sutherland

PhD, CMU, Software Engineering

Software Engineering Institute

Carnegie Mellon®