



Cyber Security Division Transition to Practice Technology Guide

Volume 2



**Homeland
Security**

Science and Technology



Thank you for your interest in the U.S. Department of Homeland Security (DHS) Science and Technology Directorate's (S&T) Transition to Practice (TTP) Technology Guide. This technology guide is the result of extensive foraging efforts to identify cybersecurity technologies developed at Department of Energy and Department of Defense affiliated laboratories. We're excited to share these promising cybersecurity technologies with you.

Through the TTP Program, S&T is identifying innovative, federally funded cybersecurity research that addresses cybersecurity needs and is helping to transition this research into the Homeland Security Enterprise through partnerships and commercialization. This guide represents an important step in that process as all of the technologies included in this guide are ready to be piloted in an operational environment or to be transitioned into a commercially available product. If you're interested in piloting, licensing, or commercializing one of the technologies, please note that the DHS S&T TTP program is funding test and evaluation activities to validate technology performance, capability claims, and interoperability; and red teaming to find, reduce, and eliminate potential vulnerabilities.

This technology guide, which is updated and published annually, is the second volume and it features nine new technologies, along with eight technologies from the first volume. To help direct future publications please reflect on the cybersecurity capability gaps in your own organization, and share your thoughts with the TTP Program Manager (ST.TTP@hq.dhs.gov). Your input will help us identify timely solutions and inform future research efforts. Again, it's our pleasure to introduce you to the TTP program and these newly developed cybersecurity tools resulting from government funding.

Sincerely,

A handwritten signature in black ink that reads "Douglas Maughan".

Douglas Maughan
DHS S&T Cyber Security Division
Director

A handwritten signature in black ink that reads "Michael Pozmantier".

Michael Pozmantier
DHS S&T Cyber Security Division
TTP Program Manager



DHS S&T Mission

Strengthen America's security and resiliency by providing knowledge products and innovative technology solutions for the Homeland Security Enterprise

Goal 1:

Rapidly develop and deliver knowledge, analyses, and innovative solutions that advance the mission of the Department

Objectives:

- Provide knowledge, technologies, and science-based solutions that are integrated into homeland security operations, employing 24-36 month innovation cycles from project inception through operational testing
- Strengthen relationships with DHS components to better understand and address their requirements
- Focus on high-priority needs, through rigorous project selection and regular review of the entire R&D portfolio
- Implement processes that strengthen project management, evaluation, and accountability within the Directorate

Goal 2:

Leverage technical expertise to assist DHS components' efforts to establish operational requirements, and select and acquire needed technologies

Objectives:

- Provide scientific and engineering advice and services to strengthen DHS acquisition processes
- Encourage the private sector (with a focus on small business engagement) to develop technologies relevant to the HSE
- Incent owners of critical infrastructure and key resources to adopt technologies that reduce vulnerabilities and increase resilience

Goal 3:

Strengthen the Homeland Security Enterprise and First Responders' capabilities to protect the homeland and respond to disasters

Objectives:

- Better understand the needs and requirements of First Responder communities, including those on the front line of border protection and transportation security
- Create high-impact technologies and knowledge products – such as standards and protocols – that facilitate the safety, effectiveness, and ease with which First Responders do their work
- Advance the interoperability of communications equipment for First Responders
- Increase First Responders' access to information on best practices and product performance standards

Goal 4:

Conduct, catalyze, and survey scientific discoveries and inventions relevant to existing and emerging homeland security challenges

Objectives:

- Ensure effective construction and utilization of S&T laboratories in support of homeland security missions
- Improve S&T's knowledge and use of relevant national and international research and facilities, with a focus on DOE National Labs and DoD efforts
- Leverage academia to address Homeland Security needs and nurture the future technical workforce of the HSE
- Collaborate with OSTP and other government agencies to develop the national policy and strategic plan for homeland security research and development

Goal 5:

Foster a culture of innovation and learning, in S&T and across DHS, that addresses challenges with scientific, analytic, and technical rigor

Objectives:

- Increase S&T and the Department's awareness of cutting edge research and technology developments pertinent to DHS missions
- Promote a culture of openness, continual learning, innovation, and collaboration within S&T Directorate and across DHS
- Internally promote synergies and eliminate programmatic redundancies by creating mechanisms and processes to increase information sharing
- Support the development of a high-performing technical workforce at DHS
- Streamline business processes to increase organizational efficiency and effectiveness



The Cyber Security Division (CSD) is a Key Component in the President's National Strategy

Threats on the Internet change fast and cyber security is one of the most challenging areas in which the Federal government must keep pace. Next-generation cyber security technologies are needed to enhance the security and resilience of the nation's current and future critical infrastructure and the Internet.

In the Department of Homeland Security (DHS) Science & Technology Directorate (S&T), the CSD enables and supports research, development, testing, evaluation, and transition for advanced technologies in cyber security and information assurance. This full lifecycle of activities evolved in response to the President's National Strategy to Secure Cyberspace and the Comprehensive National Cybersecurity Initiative (CNCI).



The CNCI establishes a multi-pronged approach the Federal government will take in identifying current and emerging cyber threats, shoring up current and future vulnerabilities in telecommunications and cyberspace,

and responding to or proactively stopping entities that wish to steal or manipulate protected data on secure Federal systems.

The S&T Cyber Security Division addresses these objectives by:

- Discovering new solutions for emerging cyber security threats to the nation's critical infrastructure;
- Driving security improvements to close critical weaknesses in today's technologies and emerging systems; and
- Delivering new, tested technologies to defend against cyber security threats and making them available to all sectors through technology transfer and other methods.

CSD Focuses on Critical Vulnerabilities in the Cyber Security Landscape

Internet Infrastructure Security—Developing security protocols for the existing Internet infrastructure (browsers and routers, essential to daily Internet operation) so that users are not redirected to unsafe websites or pathways by malicious actors.

Critical Infrastructure/Key Resources—Securing the information systems that control the country's energy infrastructure including the electrical grid, oil and gas refineries, and pipelines, to reduce vulnerabilities as legacy, standalone systems are networked and brought online.

National Research Infrastructure—Providing the infrastructure that enables development and testing of technologies to address cyber security issues including botnets, worm propagation and defense, and denial-of-service defenses that protect Internet websites against attack; providing a data repository to support the cyber security research community.

Leap-Ahead Technologies—Develop “leap-ahead” technologies that will achieve orders-of-magnitude improvements in cyber security. One of CNCI’s goals is to achieve a reliable, resilient, and trustworthy digital infrastructure.

Our vision is a cyberspace that supports a secure and resilient infrastructure, that enables innovation and prosperity, and that protects privacy and other civil liberties by design. It is one in which we can use cyberspace with confidence to advance our economic interests and maintain national security under all conditions.

— *Quadrennial Homeland Security Review, 2010*

Cyber Security Education—Helping to foster adequate training and education programs critical to the nation’s cyber security needs by providing opportunities for high school and college students to develop their skills and by giving them access to advanced education and exercises through team competitions.

Identity Management—Evaluating and developing proof-of-concept solutions, and conducting pilot experiments of identity and access control architectures and technologies, as well as data privacy protection technologies for the homeland security community.

Cyber Forensics—Developing new cyber forensic analysis tools and investigative techniques to help law enforcement officers and forensic examiners address cyber-related crimes.

Software Assurance—Developing tools, techniques, and environments to analyze software, address the presence of internal flaws and vulnerabilities in software, and improve software security associated with critical infrastructure (energy, transportation, telecommunications, banking and finance, and other sectors).

S&T: Preparing for Next-Generation Cyber Threats

In the coming years, several cyber security challenges must be addressed. The most critical of these include Enterprise-Level Metrics, Combating Insider Threats, Combating Malware and Botnets, Digital Provenance, Situational Understanding and Attack Attribution, and Usable Security.

Transition to Practice: Accelerating the Pace of Technology Transition



Homeland Security

Science and Technology

Michael Pozmantier

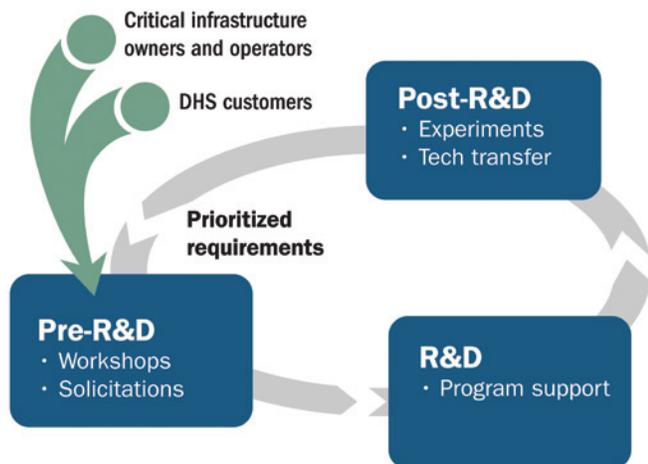
michael.pozmantier@hq.dhs.gov

In 2011, the Networking and Information Technology Research and Development (NITRD) Program of the White House named ways the Federal Government can rapidly improve the security of the Nation's cyber infrastructure. From that list, one of the NITRD's top priorities is to accelerate the transition of cybersecurity research into widespread deployment and use via the marketplace. As one of the agencies designated to address this priority, the Department of Homeland Security (DHS) tasked the Science and Technology Directorate (S&T) Cyber Security Division (CSD) with creating the Transition to Practice (TTP) Program. This newly initiated Program aligns with objectives four (Coordinate and redirect research and development efforts) and nine (Define and develop enduring "leap-ahead" technology, strategies, and programs) of the Comprehensive National Cybersecurity Initiative (CNCI). The TTP Program builds on the S&T Directorate's process of funding projects through the full research and development life cycle: research, development, test and evaluation, pilots, and transition.

In accordance with NITRD's recommendations for accelerating technology transition, the TTP Program's goals are to: (1) identify mature technologies that address an existing or imminent cybersecurity gap in public or private systems that impact national security, (2) identify and fund necessary improvements identified during pilot programs and test and evaluation activities, and (3) introduce new cybersecurity technology throughout the entire Homeland Security Enterprise through partnerships and commercialization.

The TTP Program is targeting technologies that are most likely to successfully transition to the commercial market within two years, and that will have a notable impact on the cybersecurity of our Nation's networks or systems; this is an ambitious endeavor with enormous potential for positive impact. Additionally, the TTP Program will provide a connection point for cybersecurity researchers, the Federal Government, and the private sector to ensure technology transitions from the research lab to the Homeland Security Enterprise.

For further information about the TTP Program, please send an email to ST.TTP@hq.dhs.gov.



Cyber Security Division R&D Lifecycle

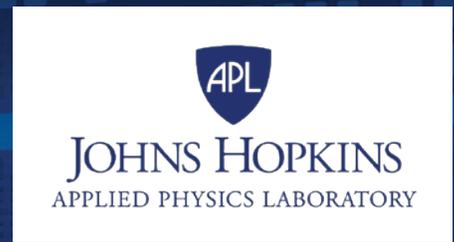
Year Two Technologies:

- **CodeDNA: Scalable, High-Speed, High-Volume, Shareable Malware Detection**
- **Quantum Security**
 - **Velocirandor: Quantum Random Number Generator**
 - **Quantum Secured Communications: Security for the Nation's Infrastructure**
- **CryptAC: Securing Data for Public Clouds**
- **LOCKMA: Lincoln Open Cryptographic Key Management Architecture**
- **Digital Ants: Dynamic & Resilient Infrastructure Protection**
- **PACRAT: The Blended Physical and Cyber Risk Analysis Tool**
- **SerialTap: Enabling Complete Situational Awareness in Control Systems**
- **SecuritySeal: Critical Protection for Your Supply Chain**
- **WeaselBoard: Zero-Day Exploit Protection for Programmable Logic Controllers (PLCs)**

CodeDNA: Scalable, High-Speed, High-Volume, Shareable Malware Detection

Margaret Lospinuso
margaret.lospinuso@jhuapl.edu

Laura Glendenning
laura.glendenning@jhuapl.edu



Overview

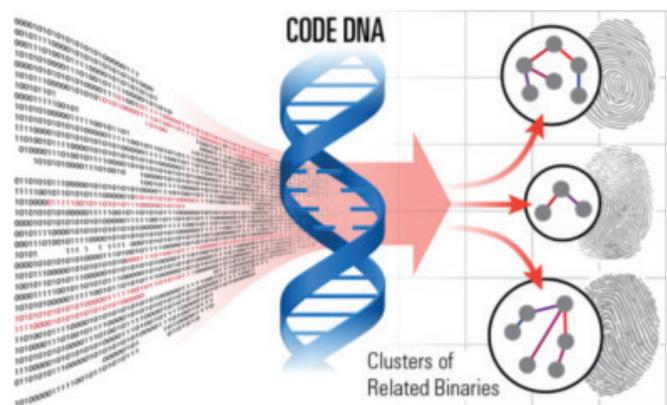
Malware attacks by external agents pose a continuing threat to government and commerce; information security costs are significant, and rising. CodeDNA is a scalable, shareable technology that facilitates community-based defense against malware attacks. CodeDNA has a significantly higher malware variant detection accuracy compared to other industry benchmarks. Attackers generally base new attacks on previously developed code; CodeDNA exploits this efficiency by reporting the codebase relationships between malware binaries. CodeDNA detects families of attacks and supports a navigable means of exploring attack family development, leading to rich insights and useful predictions about what a broad range of future zero-day attacks may look like, so that the defenders can detect them instantly.

Customer Need

Defense against malware is expensive; the economics of information security currently favor the attackers. Defenders bear the added costs of each attack individually, with little ability to achieve economies of scale, whereas attacker costs rise very little with each added attack. Current malware detection technology using checksums and similar signatures is not robust enough to support network malware ingest rates with the fully automated, inexpensive, accurate, and efficient linking of malware variants in a form suitable for secure sharing between communities of interest that is needed to achieve economies of scale.

Our Approach

CodeDNA provides a reliable, fully automated, fast means for identifying related malware binaries and linking variants. By creating highly compressed, shareable fingerprints of malware instances, CodeDNA facilitates sharing the burden of recognizing new malware variants and analyzing relationships and attribution. Defenders with access to a common repository of CodeDNA fingerprints can quickly learn what is already known, identify variants, and readily share information about newly arrived malware, thus reducing the economic burden on individual defenders. This shifts the advantage to the defenders and becomes a platform for understanding attacker plans. Incoming binaries are compared with a stored fingerprint database by a fast matching process that lends itself to low-cost open-source cloud processing. CodeDNA comparisons provide a distance metric (i.e., measure of similarity) between multiple fingerprints and support immediate drill-down into selected regions of the malware, without requiring manual intervention, cloud or code expertise, or malware reverse-engineering expertise.



CodeDNA: Scalable, High-Speed, High-Volume, Shareable Malware Detection

Benefits

CodeDNA supports fast crowd-sourcing of information by providing a robust malware identifier (fingerprint) that is deterministic and repeatable for correlating reports, analyses, and other information about attackers, yet cannot be used to re-create the original malware. CodeDNA users do not need expertise in reverse engineering, malware analysis, or code-matching algorithms, and can share fingerprints without sharing malware binaries. CodeDNA high-volume fingerprint matching has a cloud-based implementation based on open-source Hadoop, providing reliable malware detection technology whose cost grows only linearly with ingest rates. CodeDNA fingerprints are robust against common malware polymorphism using code padding and rearrangement. CodeDNA relationship data support predictions on the nature of a broad range of future zero-day attacks.

Competitive Advantage

The lack of an automated, repeatable, robust alternative to signature-based malware detection for fast clustering of malware into families has stymied attribution and crippled attempts by defenders to collaborate and join forces.

CodeDNA effectively identifies clusters of related malware in very large datasets and reports the degree of similarity. For example, CodeDNA recently found 1.8 million clusters (i.e., groups of related malware binaries) in a sample of 3.6 million binaries with unique checksum identifiers provided by Offensive Computing (*offensivecomputing.net*), thus demonstrating the ability to match 90,000 malware samples per hour on inexpensive cloud nodes running Hadoop. We believe our algorithm will scale linearly in time and cost for handling higher volumes of malware ingest. CodeDNA recognized 1,000 of 4,800 malware samples provided by the Georgia Institute of Technology as malware variants

that were not identified by 10 leading anti-virus vendors, demonstrating correlation of variants not achievable with checksum-based technologies. In a recent cyber espionage data set 76% of 293 binaries proved to be strongly related to one another when evaluated with CodeDNA. A sample of 32,000 malware binaries matched against Windows 7 using CodeDNA showed that malware authors use Windows 7 code, but did not report false positives.

Next Steps

The prototype CodeDNA is ready to be piloted and tested within a malware processing laboratory environment, followed by a move to enterprise-level testing. We envision embedding the fast cloud implementation of CodeDNA into an existing malware processing system that would provide unpacking, decryption, and de-obfuscation. Rules for processing CodeDNA matches would then lead to automatic blocking of known malware and its variants while also updating records of attempted attacks. We are searching for a transition partner to pilot CodeDNA in an enterprise environment. We are also seeking a government sponsor to fund continued research focused on mining the malware relationship data provided by CodeDNA to investigate predictive analyses of malware development. These analytical methods could expand the fast recognition of never-before-seen variants that is vital to anticipating malware developers' next moves.

Velocirandor: Quantum Random Number Generator

Jane Nordholt
jnordholt@lanl.gov

Richard Hughes
rxh@lanl.gov



Overview

Velocirandor is a small, low-cost, deployable solution to one of the most difficult problems in modern secure communications: the generation of secret random numbers (keys) at high rates.

Customer Need

Secure communication requires secret keys for use as cryptographic parameters in applications ranging from cloud computing, to secure online sessions (e.g., SSL), to hand-held device security. Keys are random numbers that adversaries must not be able to predict, influence or monitor. These requirements have consistently proven to be very difficult to achieve, and poor randomness is a common weakness in cryptography. There is a widespread need for a low-cost, compact, deployable source of high-rate cryptographic random numbers.

Our Approach

Velocirandor captures the randomness arising from properties of light that reflect its composition as a beam of elementary particles called photons. Due to the fundamental Laws of Quantum Physics, the results of certain measurements on light are intrinsically unpredictable. Velocirandor extracts this quantum randomness from a compact light source via an optical detection system, and provides true random bits at high rates (multi-Gbps) through standard interfaces. The random bit outputs pass all of the available statistical randomness test suites used to evaluate cryptographic random number generators.

Benefits

The random numbers produced by Velocirandor come with the ultimate security guarantee of an inviolable law of nature. No adversary could ever predict or influence the output. Velocirandor is affordable with component costs of about \$100 per unit. It has a small form factor (the prototype is approximately the size of a pager), and could be further miniaturized for incorporation into a handheld device. It provides the very high rate randomness (up to and beyond the 6 Gbps of the prototype) needed for modern applications such as secure cloud computing. In its current form, Velocirandor is amenable to manufacturing/automated assembly, and to integrated-photonics mass production with further development.

Competitive Advantage

Unlike conventional, true, random number generators that capture electrical or thermal noise, Velocirandor's quantum randomness cannot be influenced or monitored without detection. Deterministic random number generators use the output of known cryptographic algorithms with a short, secret seed value as input. Compromise of the seed enables an adversary to reproduce the entire output bit stream. But owing to the laws of quantum physics, no adversary can predict or reproduce the output of Velocirandor. Velocirandor is 1,000x faster and one-tenth the cost of other quantum random number generators that are commercially available.

Next Steps

There are currently two U.S. Patent applications for this technology. LANL seeks partnering or licensing with an original equipment manufacturer (OEM) or a major vendor of secure communications systems.



Quantum Secured Communications: Security for the Nation's Infrastructure



Jane Nordholt
jnordholt@lanl.gov

Richard Hughes
rxh@lanl.gov

Kevin P. McCabe
kmccabe@lanl.gov

Overview

Quantum Secured Communications (QSC) leverages Quantum Key Distribution (QKD) to replace all of the key management services provided by a public key infrastructure (PKI). QSC can authenticate and encrypt commands and data from one networked device to another over optical fiber. Devices can be anything from infrastructure control equipment, to financial trading systems, to tablet computers that are no longer connected to the optical fiber providing unprecedented speed and low maintenance costs for secure communications.

Customer Need

The cost effectiveness of networked devices is dependent on strong, long-term system security but today's cryptographic software needs constant updates and has an unknown secure lifetime. Adversaries have access to exponentially more computing and networking power each year to defeat present-day cryptography, but countering this threat with increased key lengths causes unacceptable communications latency. At the same time the risks of cryptographic failures such as those that allow intrusions into financial trading systems or false command injections into infrastructure devices are severe and a successful attack could cripple a major part of the US economy.

Our Approach

QSC uses QKD and Los Alamos National Lab (LANL)-developed techniques based on it to provide all of the cryptographic utilities required to replace key management in services such as TLS/SSL. It adds quantum user authentication with lightweight, low-latency built-in or retrofit protection for any networked device. QSC's security is based on the laws of quantum

mechanics and provides fast, reliable services with much shorter yet more secure keys providing long-term security guarantees without upgrade or maintenance costs.

Benefits

QSC replaces conventional cryptographic key and user management which has many vulnerabilities as well as maintenance and operational costs. QSC provides faster, cheaper cryptographic services with long-term system security. A central Trusted Authority securely manages the keys among users and can authorize users or devices on the fly. These techniques plus small, inexpensive, manufacturable components from LANL make it affordable.



Competitive Advantage

Moore's Law and human ingenuity are working against public-key cryptography key management systems, which also need upgrades that are difficult and expensive to perform on deployed hardware. LANL's team has been working to advance QSC for 20 years and has achieved many firsts. They have now turned to making QSC cheap and reliable for broad applicability and have more than 25 related US and foreign patent filings.



Next Steps

LANL seeks to partner with or license QSC to an Original Equipment Manufacturer (OEM) or other vendor for use in networked devices.

CryptAC: Securing Data for Public Clouds

Gene Itkis
itkis@ll.mit.edu

Darby Mitchell
mitchelljd@ll.mit.edu



MIT Lincoln Laboratory

Overview

CryptAC provides cryptographic access control, enabling secure storage of data in public clouds. CryptAC presents a seamless view of fine-grained access control and data organization, returning control of data security to the data owners. Furthermore, it separates data security from storage management, enabling seamless interoperability with multiple cloud providers.

Customer Need

Public cloud storage offers at least an order of magnitude reduction in cost for many government and commercial organizations, while enhancing data availability, ubiquity, and redundancy.

However, these advantages are currently only achievable by outsourcing data management to a third party, which requires surrendering control over the data and its security. Typically, a cloud service provider (CSP) guarantees, as part of a service level agreement (SLA), that data will be protected. But is this sufficient? Can the Federal Government afford to give up control over data security and rely on SLAs? Security breaches at CSPs and traditional websites highlight the danger of this approach. As a result, many government organizations have yet to take advantage of public clouds. A similar case can be made for commercial organizations concerned about proprietary information or even individual consumers concerned about privacy of their data.

Our Approach

Cryptographic access control relies on rigorous mathematical principles, rather than the threat of

litigation, to protect data. MIT Lincoln Laboratory has developed a framework for seamless cryptography and key management that returns control of data security and sharing to its owners by providing flexible, cryptographically enforced access control policies ensuring data confidentiality, integrity, and authenticity.

In our approach, data is cryptographically protected on the client device before being uploaded to public cloud storage, so only encrypted data is stored on the public cloud. Each data item is protected by a randomly generated content key that is itself protected (i.e. encrypted) by a public key accessible to the user. This encrypted (or *wrapped*) key is then embedded in the stored object as a cryptographic *permission*. Only the owner of the associated private key can exercise the cryptographic permission and access the content. In the simplest example, the public-private key pair is owned exclusively by an individual user. More complex schemes will support dynamic group keying.

To access the data, an authorized user would retrieve the protected content from the cloud, and exercise the embedded cryptographic permission to remove the protection (i.e. decrypting the data) on the client device. Only authorized parties with the appropriate cryptographic permissions can access the content. Unauthorized parties cannot extract the content key from the permission, because they do not possess the necessary private key. The CSP therefore never has access to the unprotected content. Using this and similar cryptographic methods, a set of access control policies are defined. These policies ensure data confidentiality, integrity and authenticity, while enabling secure sharing with individual users and/or dynamic groups.



Future enhancements to the technology include support for attribute-based access control policies and error correction coding across cloud providers.

This approach restricts administrators to managing data storage without requiring access to data contents, and empowers users to maintain total control over their data security.

Benefits

- Advantages of public cloud storage including reduced cost, improved availability, archiving and versioning, and ubiquitous access to data.
- Effective and secure sharing of data.
- Promotes incremental feature evolution and adaptability to changing security threats.
- Protection from insiders, including local and CSP administrators.
- “Plug-and-play” flexibility of selecting CSPs.

Competitive Advantage

Traditional data protection and access control tools rely heavily on local operating system permissions, since access controls and file system storage are inextricably coupled in traditional operating systems. This approach is quickly becoming obsolete in the new cloud-computing paradigm, where the data is typically replicated across many geographically dispersed systems, none of which are under the data owner’s control.

Cloud computing has implicitly introduced a new abstraction layer for data storage. This abstraction has many benefits; for example, different cloud storage providers can be added in a “plug-and-play” fashion. However, it necessitates new access controls decoupled from the underlying storage platform, which is exactly what CryptAC provides.

The effectiveness of cryptographic protection depends crucially on key management. Compared to existing protection offered by cloud services (e.g. Dropbox), CryptAC does not require data owners to trust the CSP for key management. Unlike CSP provided tools (e.g., client-side encryption in the Amazon SDK for Java) which can lead to vendor lock-in, CryptAC supports seamless integration across multiple cloud providers.

Compared to traditional client side encryption tools, CryptAC provides seamless key management and integrated support for secure sharing, as well as superior integrity protection.

MIT Lincoln Laboratory is a nonprofit, federally funded research and development center (FFRDC) whose mission is to conduct research to address problems critical to national security. MIT Lincoln Laboratory has a long and distinguished history as an impartial, independent and trusted advisor to the Federal government.

Next Steps

CryptAC provides a solid foundation on which future access control systems can be built and customized to satisfy the needs of individual government and corporate enterprise customers.

Pilot testing in an operational environment would provide an opportunity to optimize reusability, customer experience and performance while offering sponsoring organizations an opportunity to influence the future direction of the technology.

LOCKMA: Lincoln Open Cryptographic Key Management Architecture

Roger Khazan
rkh@ll.mit.edu

Dan Utin
danu@ll.mit.edu



Overview

LOCKMA is a software component designed to significantly simplify the task of adding cryptographic protections and underlying key management to software applications and embedded devices, such as mobile devices, unmanned vehicles, and sensors, as well as larger systems. LOCKMA stands for Lincoln Open Cryptographic Key Management Architecture.

Customer Need

There is a strong market need for cryptographic technology that is “seamless”, i.e., easy to integrate and use, efficient, secure, and comprehensive. While modern cryptography offers strong, proven, efficient ways to secure applications and devices, it is rarely used outside of a few established use-cases. The fundamental reason is the lack of generic, easy-to-deploy, and easy-to-use solutions for key management. Just as conventional locks require physical keys, cryptographic algorithms require digital keys to function. Managing these keys and making them available to authorized remote devices when needed, while protecting these keys in storage and in-transit, is a complicated problem.

Our Approach

LOCKMA provides just such a “seamless crypto” solution by combining the following three sets of functions into a self-contained, easy-to-use, rigorously architected and verified component:

1. Powerful, modern, NSA-approved cryptography to enable applications to protect their data at-rest and in-transit over communication channels.

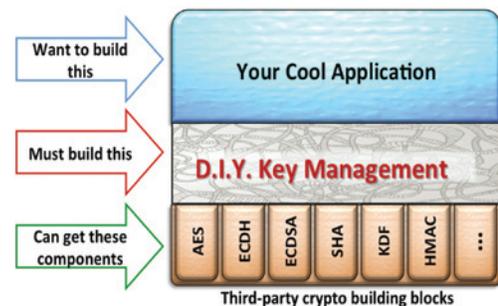
2. Standards-based identity management to help applications create, establish, and verify cryptographically-strong identity credentials.
3. Advanced, standards-based key management functions for generating, protecting, and securely distributing cryptographic keys to authorized recipients, based on their crypto identities, thereby enabling the use of LOCKMA’s crypto primitives for data protection.

LOCKMA is architected as a next-generation “seamless crypto” solution, based on several highly successful high-assurance realizations of this concept in advanced military applications.

LOCKMA is highly portable, has virtually no dependencies, is extremely resource efficient, and is decoupled from specific types of communication channels. It is beneficial to a wide variety of applications and is straightforward to integrate.

Benefits

The following illustration represents what an application developer interested in securing his/her application has to do today, versus in the future with help from LOCKMA.

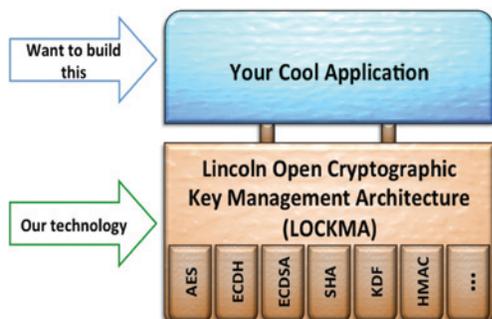


D.I.Y. key management is expensive, and often results in flawed security and hampered usability

This work is sponsored by the Assistant Secretary of Defense for Research & Engineering under Air Force Contract FA8721-05-C-0002.

LOCKMA: Lincoln Open Cryptographic Key Management Architecture

Without LOCKMA, the developer has to figure out how to combine low-level cryptographic functions into a secure design that supports all the high-level security functions required by the application: data protection, cryptographic identity management, and key management.



LOCKMA enables strong, reliable, usable crypto protections at low cost

In contrast, LOCKMA handles all of these functions “under the hood”, in a holistically architected and verified design, and provides a simple, intuitive interface to the application for invoking these functions. Using LOCKMA’s interface, an application developer can create cryptographic identities, use these identities for secure key distribution, and then use the distributed keys for protection of the application’s data.

By using LOCKMA, the effort and expense of securing an application can be reduced by at least an order of magnitude, from several man-years to several man-weeks (based on two recent uses of LOCKMA). Perhaps even more importantly than significant cost savings, the benefit for application developers in using LOCKMA is in being able to offer their users security that is both highly-dependable and easy-to-use.

Competitive Advantage

LOCKMA provides a self-contained solution that allows developers to easily integrate cryptographic protections into their applications. In contrast, existing cryptographic software libraries provide only a partial solution, lacking built-in support for key management and identity management.

Existing libraries often go for breadth, supporting many types of cryptographic algorithms, modes, and key lengths. The presence of so many options complicates the interface and makes the application developer’s job harder, not easier. LOCKMA focuses on making the addition of strong, usable cryptographic protections to applications as easy and inexpensive as possible. As such, LOCKMA implements only those algorithms approved by NIST and the NSA that are necessary for the job.

Furthermore, unlike existing key management enterprise solutions, LOCKMA enables device and applications to secure their data end-to-end, without having to trust any centralized key servers.

In 2012, LOCKMA was recognized by the prestigious R&D 100 award; a realization of LOCKMA as an FPGA core resulted in two USPTO patent applications and won the MIT Lincoln Laboratory Best Invention Award.

Next Steps

Seamless cryptography is a high-impact area with a possibility of making crypto protections ubiquitous in future products. We welcome opportunities to discuss how LOCKMA can help stakeholders secure their applications of interest.

Digital Ants: Dynamic & Resilient Infrastructure Protection

A. David McKinnon
david.mckinnon@pnnl.gov

Glenn Fink
glenn.fink@pnnl.gov



Overview

Digital Ants uses dynamic, decentralized mechanisms inspired by nature to provide mobile, resilient cyber security for protecting large enterprise IT networks and critical infrastructures. Individual ant-like sensor programs swarm to the location of anomalies and enable human operators to focus on areas and issues of concern.

Customer Need

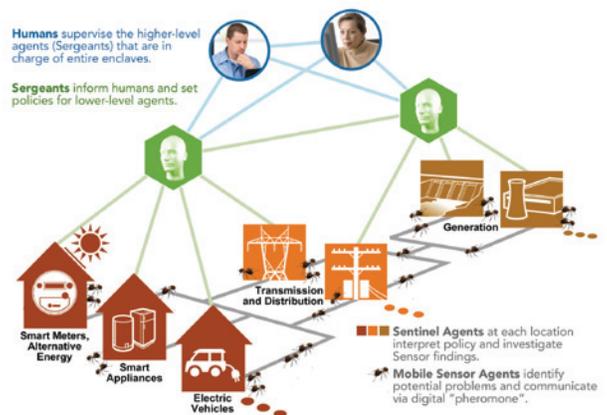
Existing cyber security mechanisms, most of which are static and centrally controlled, are straining to protect our networks. Today's enterprise networks are larger, more dynamic and more enmeshed than ever. For example, "bring your own device" is forcing traditional IT enterprises into unplanned growth and uncontrolled permeability. Even traditionally static infrastructures such as utility networks are becoming significantly more dynamic as smart devices enable two-way communication and increase customer involvement. Going forward, our cyber security frameworks must become equally dynamic and even more resilient.

A lightweight, extensible framework is needed that can address the ever-changing cyber security landscape. For example, protecting legacy devices and ensuring that protections for modern devices are future-proof is a significant concern. It is not cost-effective for utility networks to replace devices designed for 20-year deployments with new devices every 3-5 years to mitigate the computational burden of new cyber security mechanisms. Extensibility is needed to mitigate the risk of future, unknown (zero-day) malware.

Funded through the Department of Energy's (DOE) Office of Electricity Delivery and Energy Reliability (OE), Cybersecurity for Energy Delivery Systems (CEDS) R&D Program.

Our Approach

Systems in nature routinely solve adversarial problems that are even thornier than cyber security. This inspired us to use ant-colony behaviors like foraging and swarming as foundations for a cyber security paradigm shift. Digital Ants' sensors are lightweight, interpreted programs that are always on the move. These sensors roam from machine to machine within an infrastructure, via hosted software, constantly gathering and evaluating metrics (e.g., CPU usage, network bandwidth, memory use, etc.). When a sensor identifies an anomalous value, it leaves behind a digital pheromone trail. The trail attracts other sensors to the machine where the anomaly was found, similar to the way that real ants use chemical pheromones to mark a path to a food source. As more sensors arrive, they review additional metrics for anomalous values and report them to the resident Digital Ants sentinel program. If a large number of sensors have joined together, forming a sensor swarm, the sentinel informs a Digital Ants sergeant. The sergeant, located on another machine, analyzes the strength of the swarm and the severity of the reported data before informing a human operator of a potential issue.



The Digital Ants Framework within the Smart Grid

Digital Ants: Dynamic & Resilient Infrastructure Protection

A unique strength of the Digital Ants approach is the framework's ability to leverage swarm intelligence to identify previously unknown cyber security concerns. Individually, each sensor provides only a partial answer. However, a swarm of sensors can identify potential cyber security issues. Digital Ants require no centralized control, enabling the framework to scale to very large infrastructures. Furthermore, because any combination of sensors can cause a swarm, even unknown (zero-day) malware can be found.

Benefits

The Digital Ants framework is a lightweight, cyber defense that will protect very large infrastructures, even millions of devices. The decentralized approach means no single point of failure exists. Data analysis occurs at the edge devices rather than centrally, so infrastructure owners can avoid overbuilding their networks and over-investing in high performance storage and analysis servers.

Costly and complex supervised training is not required to field Digital Ants, nor do operators have to constantly update sensor profiles or malware signatures. Instead, the system learns in a simple, unsupervised manner giving personnel more time to focus on operations and security trends. Digital Ants can protect an infrastructure without overburdening CPUs because sensors do not run all of the time on all of the nodes.

The ant-inspired paradigms used by Digital Ants enables a high degree of resilience. Multiple copies of each sensor type constantly patrol the infrastructure ensuring that each node is routinely visited, even if one or more sensors fail. False positives are reduced because the Digital Ants can automatically analyze many kinds of anomalies. Even if adversaries are aware of the Digital Ants, they cannot predict the sensor movements, which provides another level of resilience to attacks.

Competitive Advantage

Digital Ants' lightweight framework enables sensor deployment on devices with the most modest resources. Traditional host-based intrusion detection and intrusion prevention systems consume too much processing power and memory for resource-constrained devices found in critical infrastructure.

Digital Ants rapidly swarm around anomalous behaviors using only simple rules and local communication rather than complex machine-learning algorithms. This enables Digital Ants to rapidly identify new malware infestations, even before signatures can be generated for them. In fact, the collection of sensors that form a swarm can help identify new pathologies and create new signatures.

Next Steps

The Digital Ants Framework has been extensively tested in laboratory and workstation environments. Going forward, this new cyber security paradigm needs to be integrated into embedded device firmware and deployed to the field. We are seeking partners that are willing to pilot Digital Ants technology in their infrastructure devices and large networks. We are also seeking systems integrators interested in harnessing the power of swarm intelligence to create new resilient protection products. Finally, we seek research sponsors to help us continue to tune this revolutionary new approach to cyber security.

PACRAT: The Blended Physical and Cyber Risk Analysis Tool

Doug MacDonald

douglas.macdonald@pnnl.gov



Proudly Operated by **Battelle** Since 1965

Overview

How secure are your assets and infrastructure? Without the right tools to properly assess the vulnerabilities of your most important assets, how can you answer that question? The Physical And Cyber Risk Analysis Tool (PACRAT), is a vulnerability and risk analysis software package that blends the methodology and assessment processes used in the physical and cyber security domains. This blended approach provides an accurate and comprehensive assessment of your overall security strategy, taking into account system level interactions and interdependencies.

Customer Need

Every industry, market sector, or business has valuable assets they need to protect. This could be a product, proprietary process, intellectual property, national security asset, or critical infrastructure element. Most organizations use a combination of physical protection and cyber security measures in their overall protection strategy to thwart attackers from attacking assets in each domain.

The cyber and physical domains must be analyzed together to thoroughly understand how each can affect the other. This holistic approach is critical to determining the resources needed to properly protect each asset. The approach and methodology used in these types of assessments ultimately determines the accuracy of the results, and directly affects the final risk determination.

Being wrong can have catastrophic consequences.

Our Approach

PACRAT uses a blended approach developed by an integrated team of physical protection specialists and cyber security experts with decades of experience in real-world, boots-on-the-ground assessments. These experts were cross-trained in the process and methodology each domain currently uses. PNNL combined and modified these approaches to provide a comprehensive modeling and simulation capability that can evaluate every avenue of approach, using both electronic and physical pathways.

PACRAT builds upon the industry standard Adversarial Timeline Analysis System and incorporates usability features of many of the most widely used analysis tools, but adds functionalities like capturing system level interactions and interdependencies and a “backtracking” capability. These elements are critical to properly assessing the true risk to an asset, operation, or facility with modern integrated security systems.

PACRAT has been provisioned for a Value-Added Module to assist in prioritizing investment upgrades. This automated process will recommend improvements to the cyber-physical systems based on increased performance parameters selected by the analyst. The result is an automated “what if” analysis.

Benefits

PACRAT’s ability to blend the physical and cyber domains into a single vulnerability and risk assessment capability provides a more accurate and comprehensive analysis than can be achieved by looking at these domains independently.



PACRAT: The Blended Physical and Cyber Risk Analysis Tool

Many organizations consider the cost of performing a comprehensive physical or cyber security assessment (ranging from \$200,000 to \$300,000 for a medium-sized facility or campus) to be frivolous with no immediate benefit. This expenditure can be difficult to justify when risk consequences are not immediately realized. However, the cost of not performing assessments, or even worse, not properly performing them, can have grave consequences.

The Great Blackout of 2003 in the Northeast is an example of underestimating the consequence of failure. It resulted in 11 deaths and cost the United States economy an estimated \$10 billion. Similar results may be possible if critical infrastructure elements are not properly protected against a coordinated cyber-physical attack.

The Stuxnet worm (which can be classified as a physically enabled cyber attack) wreaked havoc on the uranium enrichment programs in Iran, setting the program back several years, and forcing the replacement of thousands of extremely expensive centrifuges.

Compromised protection systems in the banking industry can deny customers vital access to company websites and delay financial transactions. One report estimated the cost to the financial services industry to be \$32,000 for every minute of downtime. New terrorist organizations have vowed to attack the U.S. by going after critical elements of our way of life.

The cost of not adequately protecting vital assets could be astronomical in terms of lost revenue, rebuilding and reconstruction costs, environmental restoration, damage to a company's reputation, even loss of life.

Competitive Advantage

PACRAT blends the physical and cyber domains and allows for backtracking attack pathways; no other assessment techniques or tools can do this.

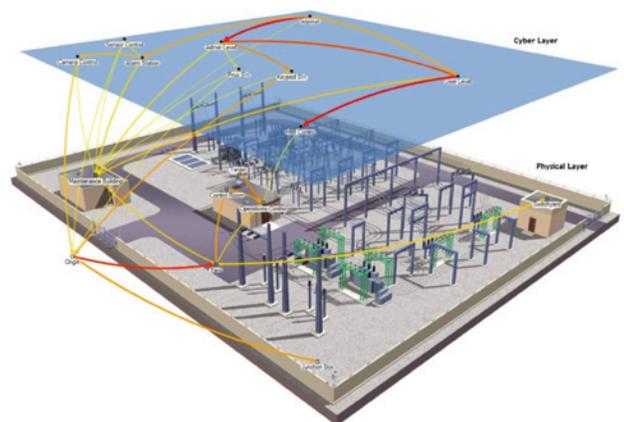
Currently used vulnerability analysis and risk assessment software packages were developed decades ago and have not kept up with the advancements in technology and the increased system level interactions introduced through automation and today's integrated security systems. Additionally, practitioners in the cyber security and physical protection domains have fundamental differences in how they apply their craft.

Our subject matter experts have been cross trained to understand the intricacies of each process and methodology, and have been able to articulate that in the PACRAT tool.

Next Steps

Two prototype PACRAT assessments have been performed to date. The software tool and subject matter expertise is ready to be taken to the next step. This could be tailoring it to a specific industry or entity to perform detailed analysis in a strategic area, or licensing the technology to be used internally or repurposed for other industries.

In either case, PACRAT is well poised to make sure adequate protection is in place to protect your most valuable assets.



SerialTap: Enabling Complete Situational Awareness in Control Systems

Thomas W. Edgar
thomas.edgar@pnnl.gov

Eric Choi
eric.choi@pnnl.gov



Overview

The SerialTap is a low cost embedded device for passively tapping serial line communication and transmitting it over an Ethernet network for comprehensive control system situational awareness.

Customer Need

Industrial control systems (ICS) and IT networks are converging. Historically, due to physical separation, ICS has had limited exposure to the vast number of Internet cyber security threats and vulnerabilities but with the merger they now become a problem. ICS now require cyber security solutions to defend against these threats.

Large portions of ICS are still operated with legacy serial communications which have largely been ignored by the cyber security community. This has led to one of the biggest challenges for control system operators: retrofitting cyber security solutions to legacy systems. IT cyber solutions focus on routable networks and are unable to work with legacy serial communications which prevents ICS owners from monitoring their entire infrastructure. As demonstrated by new malicious attacks like StuxNet, operators cannot trust the self-reported behaviors of field devices. The ability to monitor traffic in these legacy communication environments is necessary to provide complete situational awareness of ICS security state.

Our Approach

Legacy communication in ICS is often RS-232/485 serial communication. The SerialTap is a small, embedded device that is placed passively in-line on the legacy links between process control devices. It collects the data sent

between the devices, determines message boundaries, and transmits those messages via a secure UDP packet.



Figure 1: SerialTap Prototype

The SerialTap is designed specifically to fail without affecting the communications between the process control equipment (*fail-safe*). The cost of fixing a communication failure can be very high and additional equipment in the communication path increases the failure risk. The SerialTap is designed as a passive tap to remove the risk of SerialTap failure to the process control system. Loss of power or a processor failure *will not impact* process control communication.

The SerialTap is designed to encapsulate data and transmit it to a centralized location to *leverage current enterprise analysis solutions*, such as cyber security incident and event management systems. Serial communication is collected by the SerialTap, processed to determine message framing, and transmitted via secured UDP to a user configurable IP address. Centralizing analysis enables detection of system-wide anomalous patterns.

SerialTap: Enabling Complete Situational Awareness in Control Systems

Benefits

The SerialTap provides data monitoring capabilities to enable complete situational awareness in process control systems. The SerialTap is a cost effective, non-intrusive **add-on** to monitor and verify the activity in legacy serial communications systems providing the following benefits:

- **Passive failure** to reduce additional risk to ICS
- **Inexpensive** design for system wide coverage
- **Integrates easily** with common IT enterprise security solutions

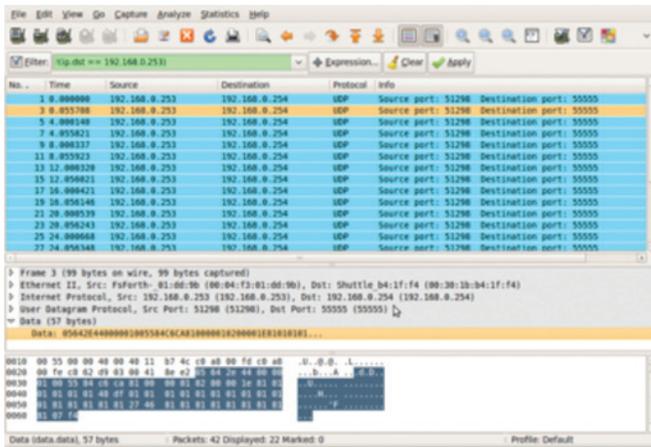


Figure 2: SerialTap Transmitted Data Monitored in Wireshark

Competitive Advantage

There are no known direct competitors to the SerialTap at this time. There are embedded devices available today that perform protocol translation, such as serial to IP, but none of them perform passive tapping. There are two categories of products that compete in the application domains that are attempting to achieve

similar end results to the SerialTap, however, these solutions fall short in two categories. The first category is designed for troubleshooting applications and requires physically connecting a computer to an adapter located in the field site. This prevents remote collection and analytics across the control system. The second category is designed to be active bump-in-the-wire solutions, similar to IT firewalls or application proxies. While these enable active protection, such as blocking known malicious traffic, the most common attacks leverage legitimate commands. In addition, it increases per unit cost and risk due to potential communication failure. The SerialTap is a cost effective method to centralize legacy communication for analysis to detect anomalous behavior in the context of the entire system and not just a single device.

Next Steps

The SerialTap prototypes have been developed and tested in laboratory environments. We would like to partner with an industry asset owner to pilot SerialTap in an operational environment and to develop and demonstrate integration with an enterprise situational awareness tool. Furthermore, it is not within our scope or capability to manufacture this technology. We, therefore, are looking for a manufacturing partner to license this technology to provide its benefits to industry.

SecuritySeal: Critical Protection for Your Supply Chain

Todd Bauer, tmbaue@sandia.gov

Robert Brocato, Jason Hamlet, Brian Wroblewski



Sandia
National
Laboratories

Overview

SecuritySeal is a combined hardware and software solution that enables cryptographically secure authentication of a seal and any object it is affixed to, providing anti-counterfeiting protection, tamper detection, and supply chain risk management for high value assets. SecuritySeal is remotely readable and the level of security is scalable to the application.

Customer Need

Global trade in counterfeit goods will top \$1.7 trillion per year by 2015. Counterfeit products pose health, safety, and security risks and create performance deficiencies that have widespread negative consequences. Microelectronics, pharmaceuticals, and chemicals are common targets for counterfeit traders. Counterfeit goods hit the manufacturing and retail industries causing loss of reputation, legal exposure and loss of sales. Legitimate manufacturers investment in R&D, materials, and human capital are diminished by counterfeits that capitalize on a brand's reputation without making the same investments.

Our Approach

SecuritySeal leverages Physical Unclonable Functions (PUFs) to create physical seals that can be used to verify that a system is authentic. PUFs are derived from the inherently random, physical characteristics of the system from which they are sourced, which makes their outputs physically and computationally impossible to predict or reproduce. The PUF output is used as a fingerprint to authenticate a system. SecuritySeal implants PUFs in both an integrated circuit (IC) that is responsible for

data processing, and in a tamper-detecting seal that is applied with adhesives to the object to be protected. The IC PUF is based on well-characterized circuit designs and the seal PUF is based on screen-printed resistors on a flexible film. The screen-printed resistors have unique values that depend on the characteristics of the surface to which the seal is adhesively attached. We simultaneously measure the PUFs from the seal and from the IC and combine them to create a system-level signature that is unique to any particular IC-seal combination. To mitigate man-in-the-middle and playback attacks that can exploit authentication using raw PUF signatures, we use a cryptographic challenge-response protocol using public/private key pairs seeded from the PUF signature. For authentication, the verifier and SecuritySeal exchange an encrypted symmetric key using Diffie-Hellman key exchange. After this process, the verifier and SecuritySeal are in possession of a shared key. The verifier then chooses a random value, encrypts it with this key, and challenges SecuritySeal with this encrypted value. Only the original, unmodified SecuritySeal will be able to generate the key needed to correctly decrypt the challenge. If SecuritySeal returns the correct result to the verifier, then the verifier is assured that the correct SecuritySeal is in place and that the item it protects has not been tampered with. After authentication, the keys and PUF measurements are erased.

Benefits

SecuritySeal enables cryptographically secure authentication of physical seals. It is widely applicable in scenarios ranging from safeguarding nuclear material to warranty fraud prevention. SecuritySeal can help combat counterfeiting of high-value consumer goods



SecuritySeal: Critical Protection for Your Supply Chain

and can satisfy ePedigree Track and Trace requirements for the pharmaceutical industry. The security level can be tailored to the application through selection of cryptography algorithms, bit generation requirements, and communication protocols. SecuritySeal can be configured for hard-wired or wireless interrogation. Each instance of SecuritySeal is unique and cannot be replicated. Unlike most cryptographic systems, SecuritySeal generates secret keys as they are needed, rather than storing them, making the system less vulnerable to attack.

Leveraging its PUF values, SecuritySeal can provide on-board encrypted memory to store application specific data like product lot information. Because the key used to protect this memory is generated from the unique PUF values, the key does not need to be stored in memory in the SecuritySeal.

The PUF-based authentication employed by SecuritySeal can be used independently of the seal to permit authentication of integrated circuits (ICs).

This has anti-counterfeiting value to manufacturers and users of ICs. This authentication capability has national security implications as it permits authentication of ICs in deployed systems, which allows detection and deterrence of modification or substitution to critical systems.

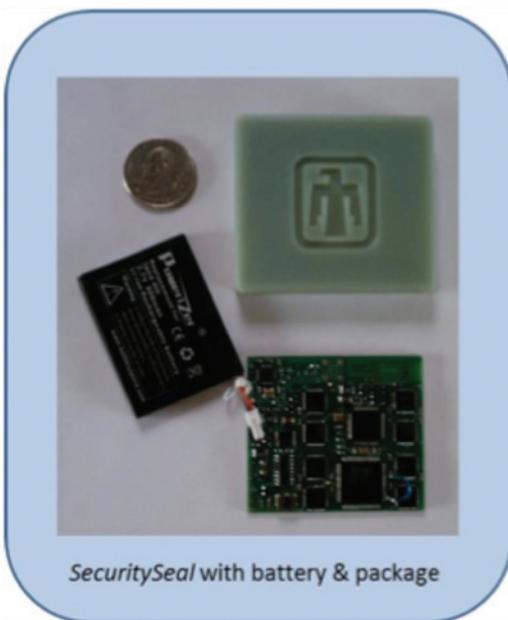
Competitive Advantage

Existing solutions fall short in either of two critical areas: 1) seals that are bound to the object that they protect are easy to counterfeit or 2) seals that are difficult to counterfeit are not robustly bound to the system they protect. SecuritySeal effectively and efficiently overcomes both vulnerabilities. SecuritySeal's unclonability and highly adaptable security level render it valuable in a wide range of applications that require the verification of the integrity of a seal, from protecting nuclear material to detecting warranty fraud.

The SecuritySeal technology is protected by US patent number 8,516,269.

Next Steps

Currently, SecuritySeal has successfully completed prototype demonstrations. It is ready to be piloted and tested within operational environments to secure high value assets. We are actively seeking a partner to bring SecuritySeal to market.



SecuritySeal with battery & package

WeaselBoard: Zero-Day Exploit Protection for PLCs

John Mulder
jmulder@sandia.gov



Overview

WeaselBoard provides zero-day exploit protection for programmable logic controllers (PLCs). By capturing and analyzing backplane traffic among PLC modules, WeaselBoard detects changes to process control settings, sensor values, module configuration information, firmware updates, and process control program (logic) updates. WeaselBoard detects zero-day attacks with minimum intrusion and footprint.

Customer Need

Critical infrastructures, such as electrical power plants and oil refineries, rely on PLCs to control essential processes. State of the art security cannot detect attacks on PLCs at the hardware or firmware level. This renders critical infrastructure control systems vulnerable to costly and dangerous attacks.

Most attacks on control systems focus on network communications, Windows PCs, and PLC logic, but not on PLCs at the hardware or firmware level. PLCs are currently not monitored for security compromise.

There is a critical need to inspect and monitor PLC hardware and firmware, and create an assurance platform for responding to attacks as these systems scale up in the future. Millions of dollars in equipment damage, lost uptime, and ultimately, casualties among operating personnel can be prevented by early detection.

These industrial control system (ICS) components receive little attention as an asset requiring security monitoring. Recent high-profile events like the Stuxnet attack (2010) and Digital Bond's Basecamp (2012) have highlighted this critical vulnerability.

Our Approach

WeaselBoard captures and analyzes backplane communications between PLC modules. WeaselBoard connects directly to the PLC backplane either in a chassis or an ICS and forwards inter-module traffic to an external analysis system.

Analysis software displays the backplane traffic, which is similar to network traffic, but is based on proprietary physical layer protocols. WeaselBoard takes the signals from the backplane and extracts fields at each protocol layer.

The analysis software uses two mechanisms to identify malicious behavior: a rule set and a machine-learning algorithm. The rules-based mechanism causes an alert when predetermined behavior is seen, and can be customized to process-specific limits. The machine-learning algorithm is a Bayesian classifier trained to alert on traffic classified into known bad states.

Operators can detect any compromise that affects the process because WeaselBoard alerts on the effects of the attack in progress, not on signatures of previously catalogued attacks. This allows zero-day exploits to be detected, unlike systems using signature-based detection methods.

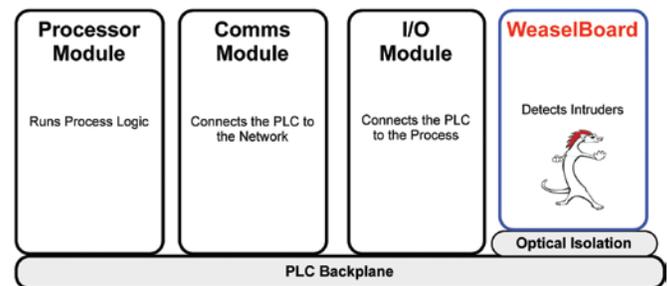


Figure 1: WeaselBoard in a Chassis

WeaselBoard: Zero-Day Exploit Protection for PLCs

The system reports unusual PLC behavior using a standard network reporting tool (syslog) and therefore works with common industry collection and correlation tools.

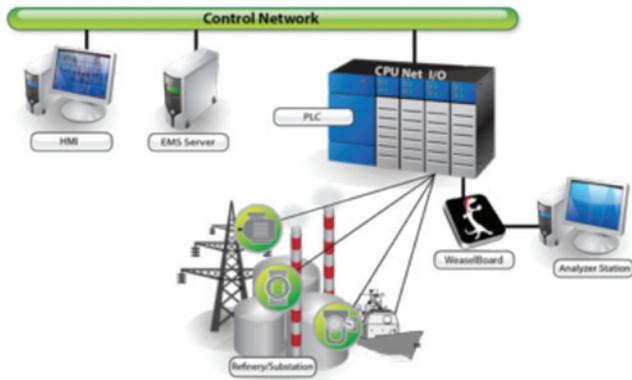


Figure 2: WeaselBoard in an Industrial Control System

Benefits

WeaselBoard detects zero-day exploits against PLCs as soon as the state of the PLC changes instead of after serious damage has occurred.

WeaselBoard addresses the problem of low-frequency, high-impact attacks from sophisticated adversaries that use zero-day attacks against PLCs. Backplane analysis provides defenders with low-level PLC behavior in real time, enabling early detection. By detecting attacks in the early stages, asset owners can mitigate or stop malicious attacks before damage occurs.

PLC devices control billions of dollars worth of production, manufacturing and utility equipment in the United States. These processes require high availability and any cyber attack could result in casualties among operating personnel, lost uptime and costly equipment damage.

WeaselBoard is anticipated to sell for less than \$500 per unit when mass-manufactured. Interoperability with existing network monitoring will facilitate integration and minimize the training needed for WeaselBoard users.

Competitive Advantage

Many security systems monitor Windows PC activity and network communications. No other security system monitors and protects PLCs. The benefit of looking at PLCs directly is that they are simpler and more consistent, so malicious activity is easier to detect.

Control system security products provide network firewalls, network intrusion detection, and assessment scanning. These tools can detect known attacks on PCs and networks, but leave the systems vulnerable to zero-day exploits that are aimed at the PLCs. There is no tool that provides direct, real-time monitoring of PLC integrity.

Industry practice forces critical infrastructure owners to react to malicious attacks after the damage has occurred, without the ability to detect PLC exploits at the firmware or hardware level.

WeaselBoard detects changes in the PLC and the process. This revolutionary capability in PLC monitoring is a novel and unique approach, protected under a 2013 US patent application. WeaselBoard fills the gap that currently exists for protection of Industrial Control Systems.

Next Steps

WeaselBoard has been tested in a variety of systems at Sandia and government laboratories, it has been validated using control system physical processes to provide realistic environments. Sandia National Laboratories is continuing to develop this exciting breakthrough technology.

WeaselBoard is seeking a pilot partner to test the system within an operational environment.



Year One Technologies:

- **NeMS (Network Mapping System): Network Characterization and Discovery Tool**
- **PathScan: Finding the Attacker Within**
- **Choreographer: A Moving Target System to Thwart Automated Network Attackers**
- **Hyperion: Detecting Vulnerabilities and Sleeper Code, Analyzing Malware, and Assuring Software**
- **USB-ARM: Architecture for USB-based Removable Media Protection**
- **Hone Technology: Producing Insight by Correlating Machine and Network Activities**
- **MLSTONES: The DNA of Cyber Security - An Organic Model for Identifying Cyber Events**
- **CodeSeal: Tamper-proof Trust Anchors**

ount Number

SSN

NeMS (Network Mapping System): Network Characterization and Discovery Tool



Celeste Matarazzo
matarazzo1@llnl.gov

Domingo Colon
colon3@llnl.gov

Overview

The Lawrence Livermore National Laboratory Network Mapping System (NeMS) is a software-based network characterization and discovery tool. NeMS produces a comprehensive representation of IP-based computer network environments constructing visual representations of the targeted network based on observed behavior. The tool provides an iterative analysis platform from which network security managers and information technology (IT) personnel can explore the findings of each mapping operation.

Customer Need

Understanding the components, structure, and activities of a computer network is the first step in many cyber defense and cyber mission assurance operations. Mapping operations are needed to discover the network topology, including routers, switches, and end hosts as well as services running on these devices. The data is processed and stored to produce a map of the target network environment that may be viewed and analyzed by the appropriate IT and security personnel of the organization.

Our Approach

NeMS applies a combination of active mapping, passive network traffic analysis, and host discovery techniques to accomplish the characterization of the network environment. Dedicated computer hardware is used to maintain performance and to provide a platform for follow-on analysis. NeMS is also implemented as a “mapping appliance,” a virtual machine containing an in-memory database-backed application for active mapping that can be placed behind firewalls, on disconnected networks, or

on other geographically or logically separated networks. The data from these mapping appliances can be combined out of band into a main mapping database for a complete network map without requiring special access. Data from all components is merged into a single data store for analysis and visualization.

NeMS identifies and uncovers the network environment through a combination of:

- Discovering active devices
- Identifying communication paths
- Discovering open ports and associated services and applications
- Identifying active routing directives
- Discovering previously unknown devices
- Discovering unknown routing behaviors
- Identifying and processing transactions between hosts and users of the network
- Labeling content and resequencing network traffic

Benefits

NeMS creates a queryable graph of any IP network with details of network entities, attributes, roles, and logical relationships. It can be utilized from outside a firewall or from any vantage point within a network, including multiple vantage points. The tool addresses the need for mapping IP networks to achieve network situational awareness without requiring extensive network preparation or prior knowledge and without compromising the security posture of the mapped environment. NeMS has many controls to enable the mapping operations to meet performance and security requirements. The system can be applied to



NeMS (Network Mapping System): Network Characterization and Discovery Tool

government networks and commercial networks and can be used to generate a new map, corroborate or update existing maps, or fuse with additional types of data and information that may be available in the enterprise.

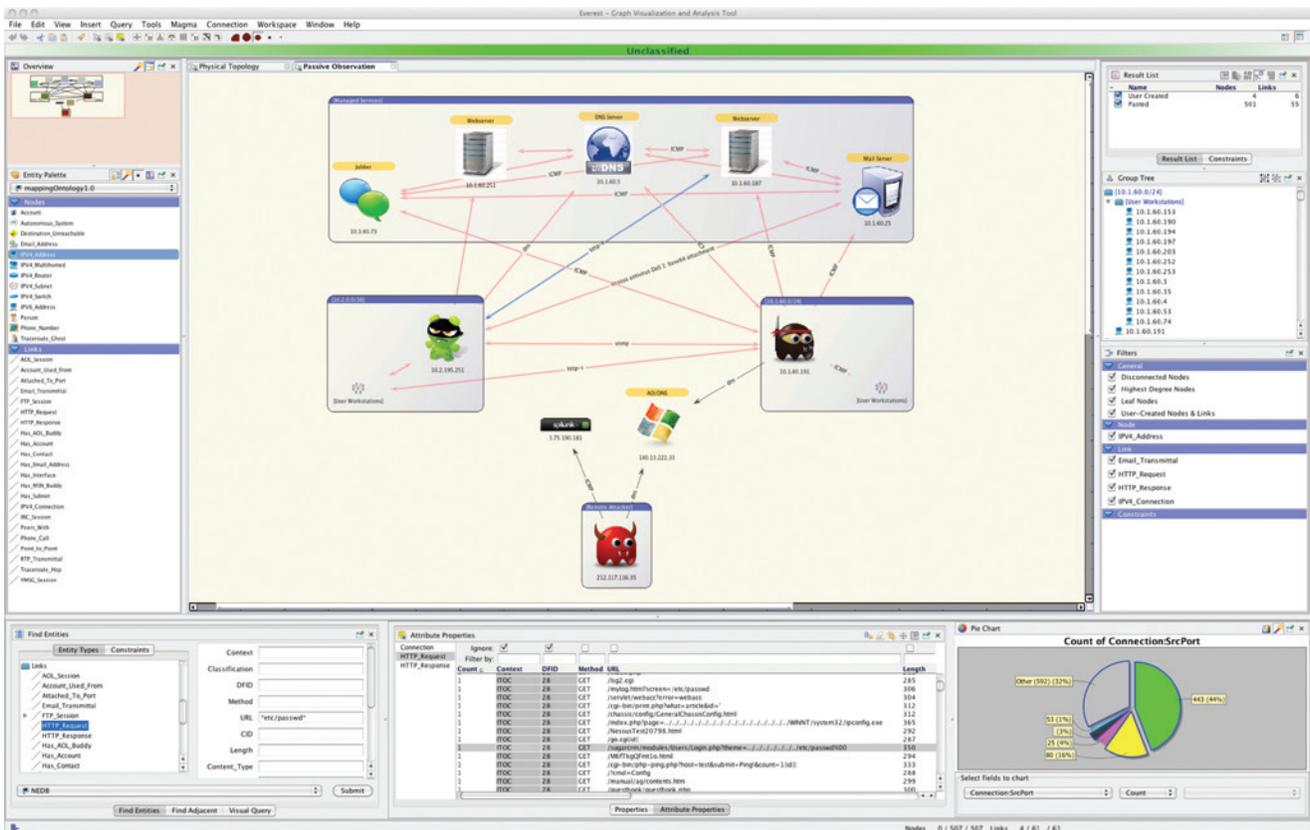
Competitive Advantage

Current tools for mapping networks are often slow and intrusive on network operations; they also require special exceptions to network security. The NeMS tool is designed and configurable to minimize disruptions and impacts on the target operational network and to require minimal intervention by network security staff. The tool uses well-established and tested network query tools and mechanisms while mapping. The system has a modular structure that allows the easy addition of new

capabilities. For example, host-based sensors or asset information can be integrated into the system by adding the interface to the persistence layer with appropriate analysis and visualization primitives.

Next Steps

The NeMS tool is currently available for pilot deployment for your network mapping and analysis needs.



PathScan: Finding the Attacker Within



Joshua Neil
jneil@lanl.gov

Curtis Hash
chash@lanl.gov

Overview

PathScan quickly detects the movement of hackers once they are inside a computer network.

Customer Need

Hackers can and do penetrate perimeter defenses. For example, users clicking on phishing e-mails allow hackers to bypass firewalls and intrusion detection systems, providing a foothold in the network. Testing indicates click rates on phishing e-mails are as high as 10%. To get to the core network assets, hackers must traverse the network after this initial penetration. There is a need to quickly identify hackers once they have penetrated perimeter defenses, but before they can access core network assets.

Our Approach

PathScan targets the traversal behavior of hackers by building behavioral models to reflect normal activity, followed by passively monitoring network traffic and comparing it with the behavioral models. Our approach proceeds as follows:

- Build statistical models to characterize the network traffic between each pair of communicating computers.
- Break the network into millions of small paths.
- Passively monitor each path and test whether the data observed is likely to be normal according to the models built in Step 1 or, alternatively, it appears to be caused by a hacker moving along this path.

PathScan has two modes of operation.

- *Online*: Currently, PathScan is operational on

LANL's unclassified network, analyzing millions of communications every minute.

- *Forensic*: PathScan can also be run in a forensic mode; it has proven effective in fleshing out attacks initially identified by security incident responders, discovering additional compromised machines that were undetected by the original investigators.

With a single commodity Symmetric Multi-Processing (SMP) machine, we are able to rapidly analyze LANL's 20,000 node unclassified network, examining the network in near real-time. We require network connectivity information in the form of DNS or NetFlow data. The output is a ranked list of the most anomalous hosts along with a heat map, as depicted in Figure 1.

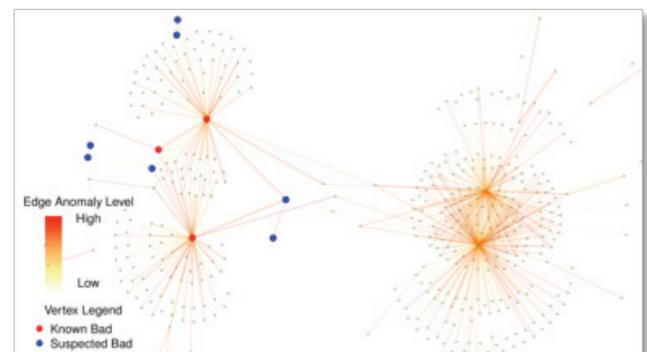


Figure 1: Heat map of the detection of an actual Advanced Persistent Threat attack on LANL's network.

In Figure 1, dots are machines and lines are communications between machines. The color indicates the anomaly level. Trillions of objects were examined, the most anomalous of which, pictured here, contain the truly compromised machines as confirmed by forensic investigation. In addition, several machines identified

PathScan: Finding the Attacker Within

in this plot were later determined to be compromised, yet the initial investigation missed these compromises, indicating the value of PathScan to aid forensic discovery of the attack.

Benefits

Detection is needed before attackers get to core assets. Early detection allows network operators the ability to shut down only those machines that are determined to be compromised, avoiding the shutdown of the entire network. This prevents the exfiltration of important data, but even more importantly, the difference between detecting an attack within the first few minutes and detecting the attack after several hours can be the difference between a minor security incident and an extremely costly attack. It has been shown that allowing an attacker to exist within a network for more than a few hours allows that attacker to penetrate the core machines, such as Active Directory servers. Compromise of these servers forces network operators to shut down the entire network, possibly for several weeks, in order to ensure effective removal of the compromise.

Competitive Advantage

Commercial products mostly look for exact signatures of previous attacks, whereas our methods are statistical in nature, allowing us to detect both known and zero-day behavior. Many government solutions only monitor data at the perimeter, while ours examines internal data, finding the attacker once they are inside. Finally, academic approaches generally do not scale well or are only applied to synthetic, non-realistic networks. PathScan, on the other hand, has scaled to very large networks, has scaled up to millions of computers, and has been validated on large operational networks.

Next Steps

We are actively developing PathScan. We seek partners to provide:

- *Pilots:* Live networks are needed for continued validation. Access to these networks will aid us in ironing out operational issues. In addition, any attacks identified on these networks will aid PathScan researchers in tuning the approach for future attack identification. These pilot network operators will be provided with analysis results and real-time protection using PathScan.
- *Commercialization:* The PathScan team at LANL is a research and development (R&D) organization, not a commercial software shop. We seek commercialization partners to provide software solutions that will encapsulate the PathScan technology in a commercial-grade software environment.
- *Research Funding:* While PathScan has proven effective in identifying Advanced Persistent Threat (APT) activity, research on improvements are ongoing. More advanced statistical models and better data collection mechanisms are important to stay ahead of the threat. We seek R&D funding to support the further improvement of PathScan.

Choreographer: A Moving Target System to Thwart Automated Network Attackers

Craig A. Shue
cshue@ornl.gov



Overview

Attackers regularly compromise the public-facing servers that organizations use to fulfill their missions. We regularly change the locations of these servers, disorienting attackers while still providing reliable connectivity for legitimate users.

Customer Need

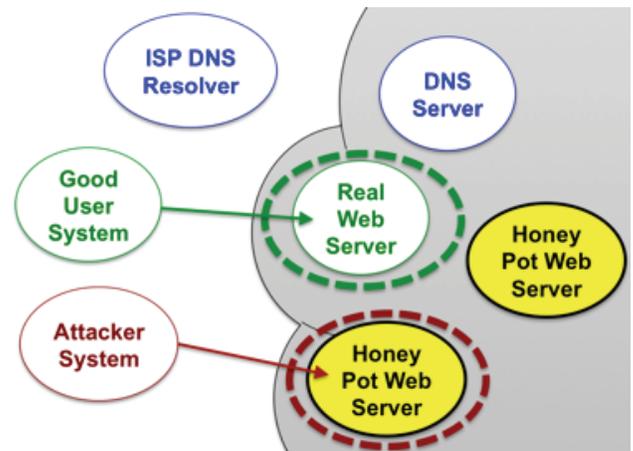
Most organizations face ongoing and damaging attacks on their public-facing servers. At the same time, such servers are critical to the mission objectives of these organizations. Over time, attackers can scan organization servers and learn about the infrastructure and defenses, allowing the attackers to tailor their assault on our infrastructure. Each public server can become an attack vector and a foothold into an organization's network for adversaries. The cost of a security failure can be high. In 2011, the average cost of a data breach was estimated at \$5.5 million. Other costs may be less quantifiable, including damage to the customer's reputation or even the customer's ability to complete its mission.

Our Approach

We frequently change the public addresses of protected servers, which 1) makes it challenging for attackers to guess the server's address, and 2) allows us to seamlessly redirect an attacker to monitoring infrastructure (called a "honey pot").

When contacted by a legitimate user—one without a prior history of attacks—the DNS server provides the correct address for the server and creates a network mapping to maintain the link.

This approach allows the DNS server to grant or deny access to legitimate users and seamlessly transition malicious users to honey pots upon detection. Organizations can use prior history to make decisions, protecting themselves based on past actions by a network.



Benefits

- Our approach reduces attacker scanning effectiveness from around 100% to less than 1% for most network deployments.
- We can limit access to authorized requestors.
- We can study the diverted users, 95% of which are likely to be malicious users and parole the legitimate users.
- We enable policy decisions based on the source network, incentivizing ISPs to remove malicious clients from their networks.
- Deployment is straightforward and requires only minor changes in infrastructure.
- Performance overheads are minimal for smaller DNS zones and organizations can select which systems to place in protected zones.

Competitive Advantage

While traditional firewalls can thwart access, they are based on signatures, and a single misconfiguration allows arbitrary attackers into the network. Dynamic and adaptive network approaches do not support migrating ongoing connections from the old address to the new one, causing connections to break. Our supports established connections even when addresses change. Our approach makes an explicit decision before the connection starts and during the connection, if needed. Unlike intrusion detection systems that rely on anomaly detection or attack signatures, our approach can detect and thwart zero-day attacks.

Next Steps

We are ready to begin early rollout in non-critical production environments. We are looking for partners interested in transitioning the technology.

Hyperion: Detecting Vulnerabilities and Sleeper Code, Analyzing Malware, and Assuring Software

Stacy Prowell
prowellsj@ornl.gov

Rick Linger
lingerr@ornl.gov



Overview

Hyperion is a new technology that computes the behavior of software, including malware, in all circumstances of use, without the need for source code. Hyperion operates on compiled binaries, rather than source code, to approach the ground truth of processor operations.

Customer Need

The first day a vulnerability is announced, a vendor loses, on average, \$860 million in market value (Telang and Wattal 2007), and software security incidents cost an average of \$300,000 (Aberdeen 2010).

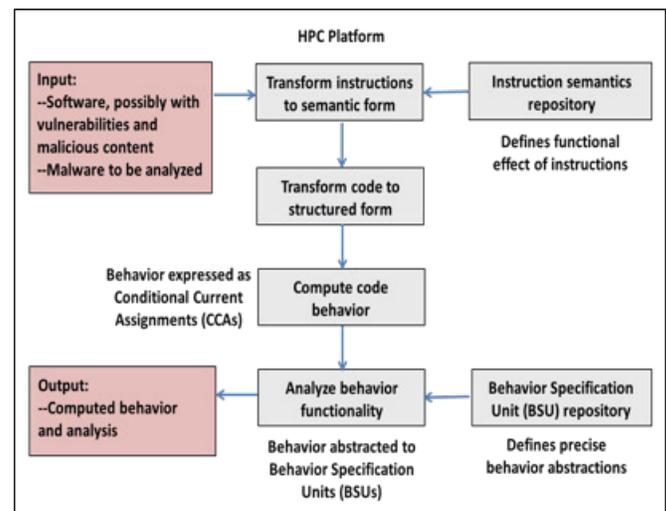
There is a growing need for more complete analysis of software to account for all possible behavior, whether legitimate or malicious, without the uncertainties of approximations and heuristics. Such a capability will help assure newly developed and acquired software, reduce damage from vulnerabilities, and assist in analyzing malware.

Our Approach

For any critical software functionality, the Hyperion system generates associated program behaviors and the complete set of conditions under which they occur. These behaviors can be automatically checked for known malicious signatures and inspected by domain experts to assure correct operation and the absence of malicious content.

The key is Oak Ridge National Laboratory's (ORNL) Function Extraction (FX) technology that directly computes the behavior of software binaries, no matter

how they were originally coded. FX transforms programs into procedure-free, "as-built" specifications based on deep semantic analysis that enables new kinds of reasoning. ORNL is the sole provider of FX, leveraging its institutional expertise in big data and High Performance Computing to address scale up and performance. Hyperion applies the mathematics of denotational semantics to transform input code into a functional representation, transform it into a structured form, compute its behavior, and abstract that behavior according to Behavior Specification Units provided by domain experts.



Benefits

Hyperion provides a repeatable, cost-effective means to achieve assured software. It permits validation with high confidence of the security of software for the deployment environment. It also permits discovery of the functionality of malware, even malware that is obfuscated and hidden. Because Hyperion coalesces and aggregates related behavior, malware that is distributed throughout legitimate code is revealed as just more

Hyperion: Detecting Vulnerabilities and Sleeper Code, Analyzing Malware, and Assuring Software

cases of behavior. The technology has also been applied to polymorphic and metamorphic malware.

Competitive Advantage

Existing approaches to high-assurance software include testing, inspection, and scanning. Even the best testing can exercise only a small subset of possible executions, and inspections are time-consuming and subject to human fallibility. Scanning methods are largely syntax-based, depend on a priori signatures, and can be subverted by simple variations. Behavior computation used by Hyperion does not look for specific artifacts in code; rather, it computes all behavior, legitimate and malicious, to permit complete analysis and assurance.

Next Steps

ORNL is ready to customize, deploy, and support the Hyperion system as required by sponsors. Hyperion algorithms have been developed to execute in computing clusters that are readily available on the market for others to use. If necessary, these algorithms can be adapted for specific sponsor requirements.

The system is driven by definitions of the functional semantics of instructions. These semantics have been developed for a comprehensive subset of the Intel x86 instruction set and for the MSP430 processor. In addition, the system can apply semantic subject-matter abstractions of behavior provided by domain experts.

Beyond x86 and MSP430 code, the system is ready to be tailored for sponsor requirements for other instruction sets. Custom user interfaces can also be provided to integrate the technology with specific operational environments.

R. Telang and S. Wattal, "An Empirical Analysis of the Impact of Software Vulnerability Announcements on Firm Stock Price," *IEEE Transactions on Software Engineering*, 33(8): August 2007, pp: 544-557.

Aberdeen Group, "Securing Your Applications: Three Ways to Play," August 31, 2010, quoted in

<http://blogs.aberdeen.com/it-security/quantifying-business-value-of-application-security-cost-avoidance-cost-savings/> (retrieved on 10/2/2012)

USB-ARM: Architecture for USB-based Removable Media Protection

Logan Lamb
lamblm@ornl.gov



Overview

USB-ARM provides a simple, efficient, and customizable layer of security that brokers all communication between removable media and the operating system. USB-ARM guarantees that a set of user-defined criteria are met prior to allowing access to the removable media.

Customer Need

While the convenience of USB devices and removable media increase productivity, they also provide an effective attack vector for malicious software. Currently organizations have to compromise on their solutions for handling removable media and protecting against malware. Generally, organizations will either ban removable media use, resulting in lost productivity, or rely on a single anti-virus solution to eliminate infection. Unfortunately, no anti-virus tool has a 100% detection rate and many policies can be circumvented. The cost of such realities can be tremendous. The average cost of a cybersecurity incident totaled \$214,000 in 2012, ballooning to \$5.5 million if involving a data breach; it is only expected to increase as threats continue to evolve. Organizations require extensible tools to handle the evolving malware threat. With USB-ARM, an organization does not have to choose between technologies and can easily incorporate all enforcement tools into their removable media policy.

Our Approach

USB-ARM installs a driver that brokers all communication between the removable media and the operating system. Upon recognition from the operating system, USB-ARM blocks all communication to the device until a set of user-defined criteria are met. For example, a configuration might employ McAfee anti-virus, followed by AVG anti-virus, and finally an executable

detection engine. Access to the media is granted only if McAfee and AVG found nothing suspicious. Access to a given file is granted transparently if not flagged as an executable by the detection engine.



Benefits

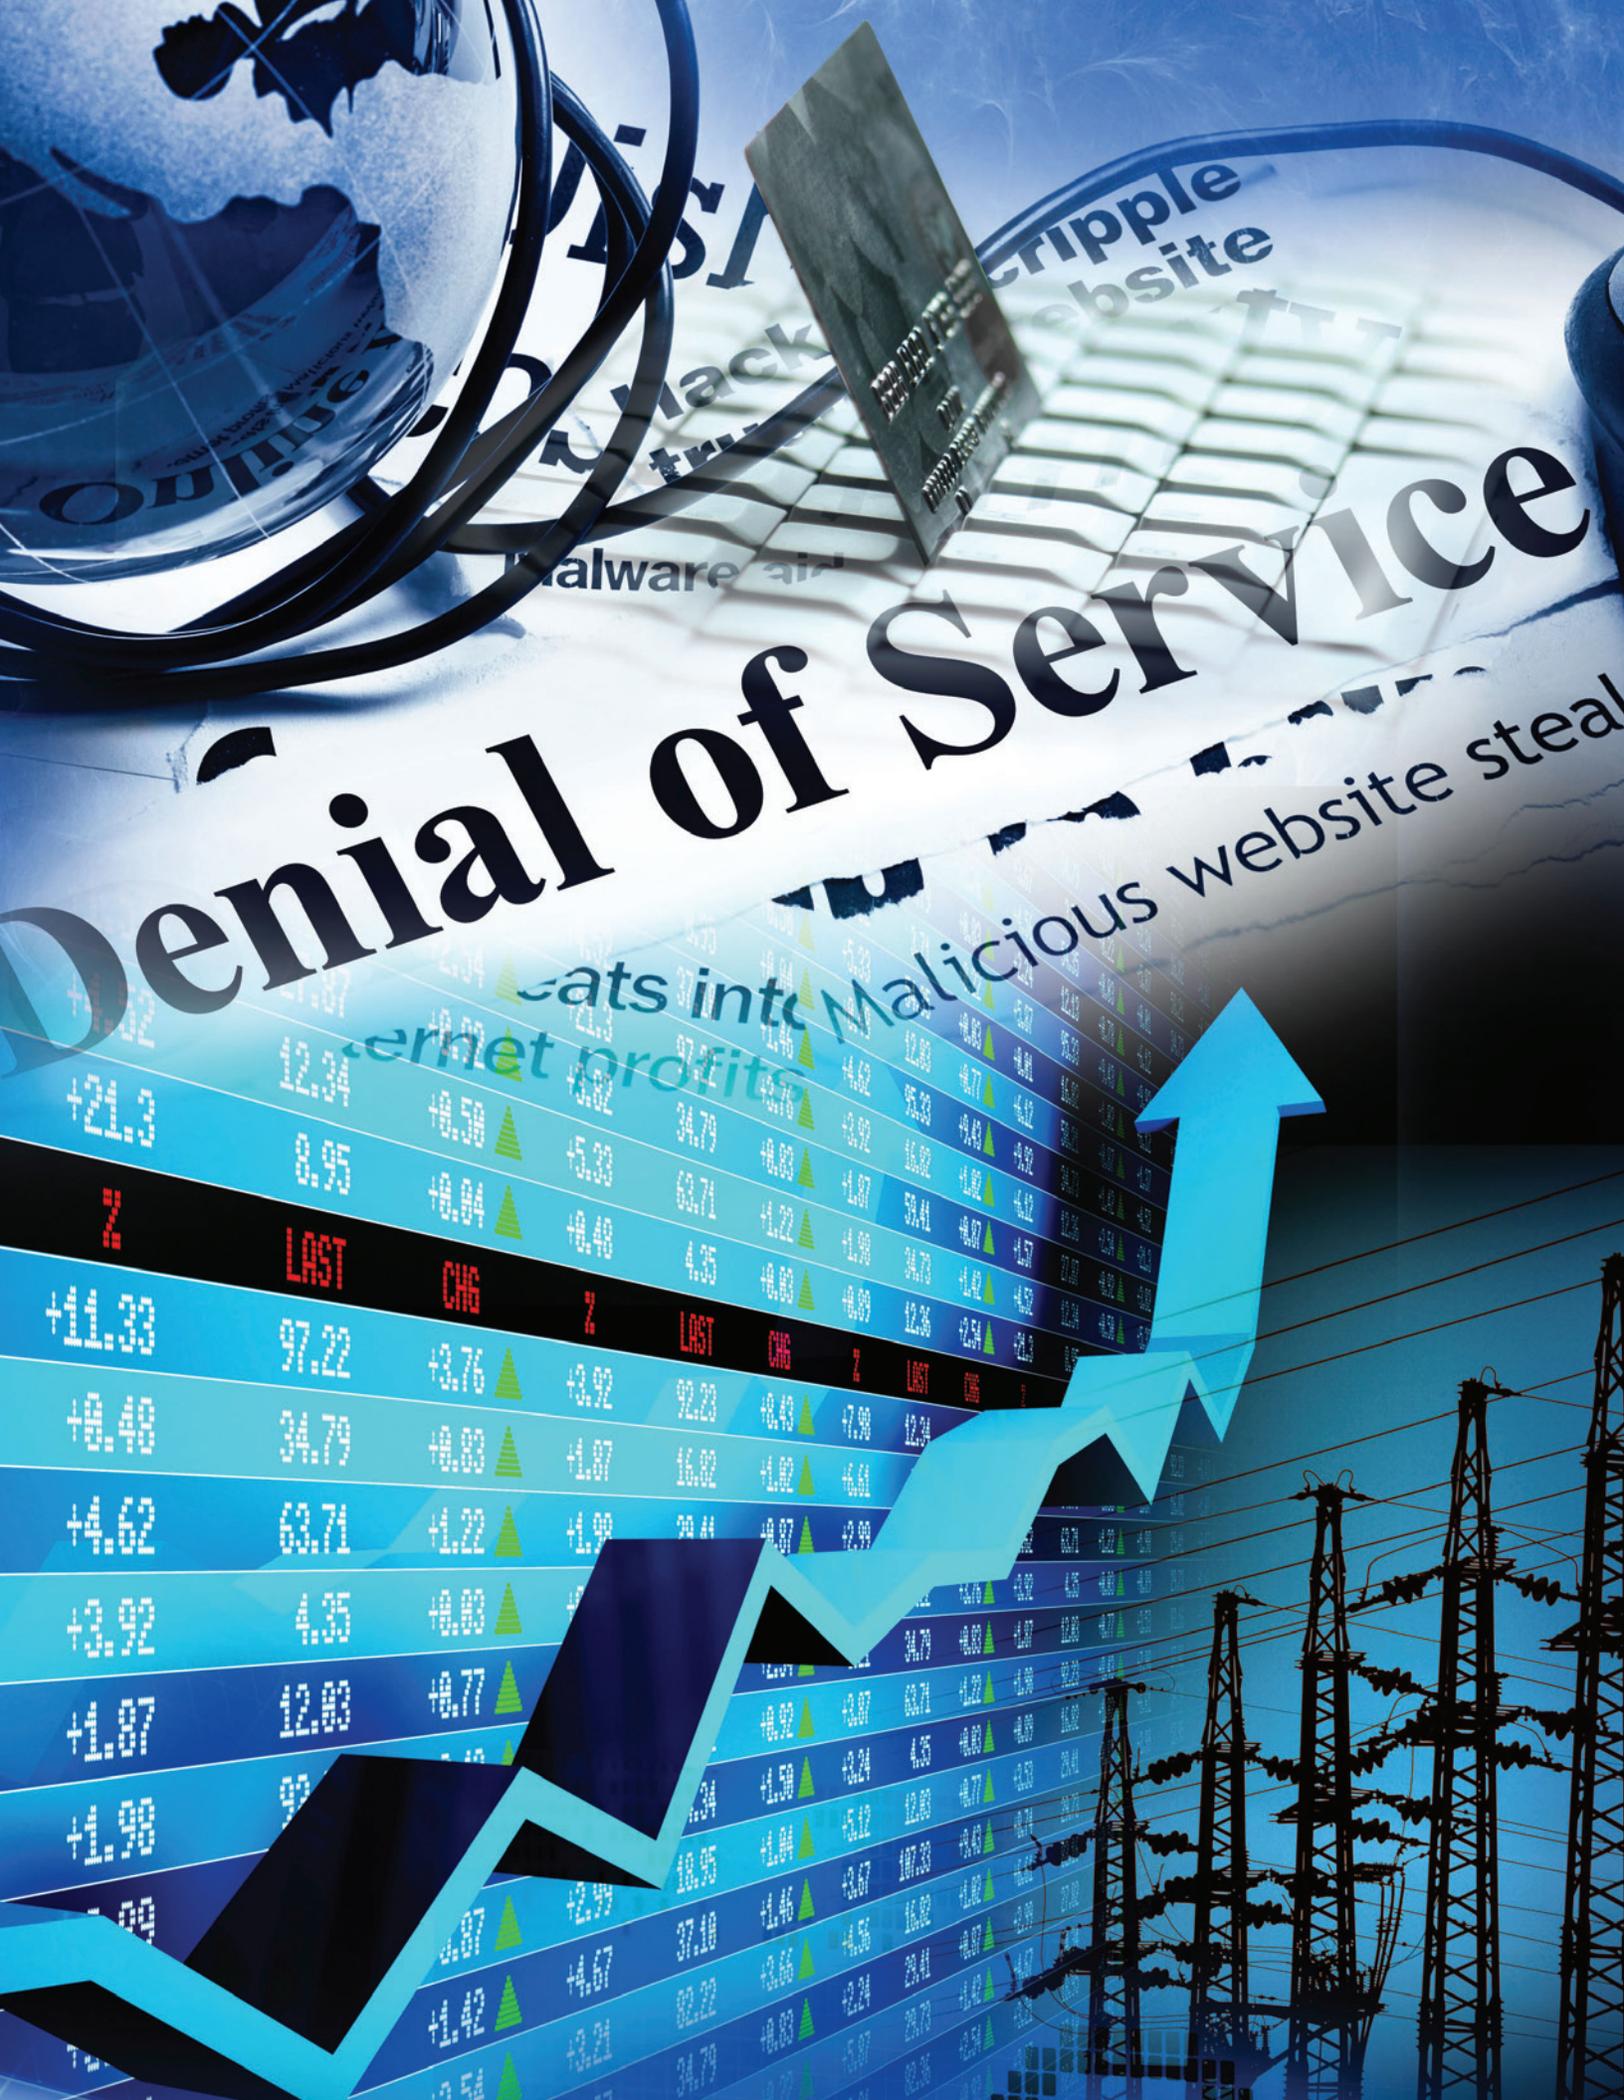
USB-ARM guarantees that a set of user-defined criteria are met prior to allowing access to the removable media. USB-ARM eliminates any possible race conditions between security software and the execution of malware on removable media. Unlike current mechanisms, USB-ARM facilitates sequential use of multiple anti-virus engines, ensuring maximum protection. An organization can decide for itself what security properties are used to identify “clean” media. USB-ARM is simple, efficient, and transparent to the user. It is as effective as the sum of the user-defined stages, allowing customization to an organization’s needs.

Competitive Advantage

USB-ARM prohibits access to removable media until all user-defined stages have been completed successfully. This capability allows organizations to easily tailor and extend their removable media policy. Currently, no other tool has this capability.

Next Steps

We are seeking partners interested in piloting and commercializing USB-ARM.



Denial of Service

Symbol	LAST	CHG	%	LAST	CHG	%	LAST	CHG	%
+	12.34	+0.58	▲	34.79	+0.83	▲	63.71	+1.22	▲
+	8.95	+0.04	▲	63.71	+1.22	▲	4.35	+0.03	▲
+	21.3	+0.09	▲	97.22	+3.76	▲	92.23	+0.43	▲
+	12.34	+0.58	▲	34.79	+0.83	▲	16.02	+1.02	▲
+	8.95	+0.04	▲	63.71	+1.22	▲	29.41	+0.97	▲
+	21.3	+0.09	▲	97.22	+3.76	▲	92.23	+0.43	▲
+	12.34	+0.58	▲	34.79	+0.83	▲	16.02	+1.02	▲
+	8.95	+0.04	▲	63.71	+1.22	▲	29.41	+0.97	▲
+	21.3	+0.09	▲	97.22	+3.76	▲	92.23	+0.43	▲
+	12.34	+0.58	▲	34.79	+0.83	▲	16.02	+1.02	▲
+	8.95	+0.04	▲	63.71	+1.22	▲	29.41	+0.97	▲
+	21.3	+0.09	▲	97.22	+3.76	▲	92.23	+0.43	▲
+	12.34	+0.58	▲	34.79	+0.83	▲	16.02	+1.02	▲
+	8.95	+0.04	▲	63.71	+1.22	▲	29.41	+0.97	▲
+	21.3	+0.09	▲	97.22	+3.76	▲	92.23	+0.43	▲
+	12.34	+0.58	▲	34.79	+0.83	▲	16.02	+1.02	▲
+	8.95	+0.04	▲	63.71	+1.22	▲	29.41	+0.97	▲
+	21.3	+0.09	▲	97.22	+3.76	▲	92.23	+0.43	▲
+	12.34	+0.58	▲	34.79	+0.83	▲	16.02	+1.02	▲
+	8.95	+0.04	▲	63.71	+1.22	▲	29.41	+0.97	▲
+	21.3	+0.09	▲	97.22	+3.76	▲	92.23	+0.43	▲

Hone Technology: Producing Insight by Correlating Machine and Network Activities

Glenn Fink

glenn.fink@pnnl.gov



Proudly Operated by **Battelle** Since 1965

Overview

Hone is a host-based cyber security tool that provides a new kind of data: correlated Host and Network data. Hone bridges a fundamental gap in Internet protocol design to enable powerful insight for cyber defenders.

Customer Need

In 2012 cyber attacks cost on the average US company \$8.9 million, and these companies had to fend off an average of 102 successful cyber penetrations every week.¹ Computer systems require continuous monitoring, not just regulatory compliance. But continuous monitoring produces enormous amounts of host and network data, and defenders only have time to concentrate on the most important information to make the best use of their analysis dollars.

Amazingly, the design of Internet protocols prevents analysts from correlating network activities to the responsible processes on communicating machines. Network routing is kept separate from the routing of those communications among machine processes. This separation makes it impossible for analysts to determine with certainty which processes are responsible for which communications, so they cannot easily isolate root causes of break-ins.

Our Approach

The unique contribution of the Host-Network (Hone) sensor is that it bridges the networking and processing parts of monitored machines, enabling analysts to know which program is responsible for each network activity. With Hone, analysts can quickly and accurately find the root cause of suspicious activities seen on the network.

Rather than collecting all the available host and network data as we currently do, Hone collects only the network data that has an effect on the monitored machine, and it tags it with the responsible process's identifier. This data is the most important for homing in on problems.

Hone is installed in the kernel, the deepest part of the operating system, via a system patch and a small host-based agent that runs on each monitored machine. Wireshark, a free network analysis tool, has been modified to view the correlated data from Hone.

Benefits

While Hone relies on a host-based agent, it can be easily rolled out and managed in an enterprise via the Windows Installer. Hone's performance impact is so small that it cannot be easily distinguished from normal operational variability. Hone enables unprecedented visibility inside the monitored machine, greatly reduces analyst workload, and provides 100% accurate process attribution. Knowing the responsible program enables us to make analytical use of detailed process information including open files, registry entries, libraries, and user information.



1. Ponemon Institute, LLC, 2012 Cost of Cyber Crime Study: United States. October 2012.



Hone Technology: Producing Insight by Correlating Machine and Network Activities

Competitive Advantage

Without Hone it is impossible to determine with certainty the process responsible for each communication. Analysts may rely on deep-packet inspection technologies to approximate the correlation, but these technologies can cost hundreds of thousands of dollars apiece, they can violate user privacy, and they still can provide only a guess about the process-packet correlation.

Connection-filtering, host-based firewalls are another alternative, but they only operate on the connection level, not every packet. This means that once you grant blanket permission for an application to access the network, you have no further control. Only Hone provides the precision to control communications at the packet level coupled with 100% certain attribution of responsible processes.

Another kind of solution, multi-host management, uses host-based agents and provides monitoring services that can rival Hone's in some ways, although none can give the absolute certainty that each packet goes to a specific process as Hone can. These solutions are often quite expensive in every way, and they require elaborate infrastructure.

In contrast, Hone is simply a data source that can be taken advantage of by other security information management tools agencies already own. By integrating Hone data into existing intrusion-detection systems, firewalls, and other defensive systems, these defenses would gain an additional layer of insight into the meaning of network and host activities.

Next Steps

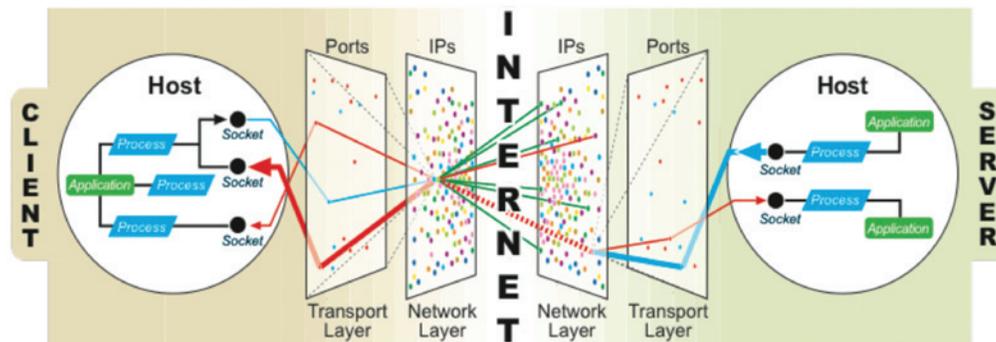
We are seeking clients to sponsor further development, for pilot testing, and to license Hone for use in new and existing technologies. Hone provides a new kind of correlated data. Partners could help us build upon this the ability to collect the logs centrally, to provide further analysis of the log data, and to enact network controls based on the analysis. For instance, by embedding Hone data in network flows, the network infrastructure could make access control decisions based on what process caused the activity.



Hone currently supports Linux and Microsoft™ Windows™ 7 and 8. We also have prototype Mac OS X and Android sensors that could be completed with sponsorship. The Linux version of the sensor is open-source, and we welcome collaborative development effort.

Because every Internet-capable device uses the same protocols, Hone will create a revolution in cyber defense. The correlation of communications to processes that once had to be done in analysts' heads just became automatic and 100% accurate. So ask yourself,

What would you be able to do if you could collect only the most important cyber data and rapidly isolate root causes of cyber break-ins?



Hone provides correlated data that can enable an unprecedented end-to-end view of network communications

MLSTONES: The DNA of Cyber Security - An Organic Model for Identifying Cyber Events

Elena Peterson
elena@pnnl.gov



Overview

MLSTONES is a set of tools that support a methodology that can help you quickly find the needle (cyber event) in a haystack of data, even if you don't know which needle is there and the haystack is full of other types of needles you aren't interested in. MLSTONES can also help you identify new cyber events that are not already known.

Customer Need

Our reliance on cyber systems permeates virtually every aspect of national infrastructure. From banking, finance, and industry to agriculture and distribution, from national defense to power generation and delivery, software and the data it produces are the lifeblood for maintaining critical infrastructure, information, and the U.S. strategic advantage over our adversaries. The volume of data generated has outpaced our ability to effectively analyze it fast enough to prevent many forms of cyber attacks. In most cases, new forms of attacks cannot be detected with current methods. We need a method to drastically reduce the amount of data to be analyzed, to quickly characterize a cyber event, and to identify previously unseen types of attacks before they are executed.

Our Approach

We've translated several biology and bioinformatics concepts onto cyber defense data. Specifically, we've created a methodology that uses the concepts of protein identification and families, inheritance, and function to apply to a number of cyber-based data types. The MLSTONES process creates cyber "proteins" and then uses protein alignment techniques to generate families of proteins; it does so very quickly. With this method, we can then create a single representation of an entire

family of entities, thus reducing the amount of data to analyze by several orders of magnitudes.

We can also infer the function of a "cyber protein" by its relationship to other similar proteins. This is the same process used in biology to discover similar proteins. This helps to identify completely new (zero-day) cyber threats.

```
>SERVER1
QLMAQMLQQANNNNNNNQLMAQIIMQALQLMQATMGNIQINAGQQQMLQL
MALAWRWRWVWVQTAGMMLLQAAQLMAMLQQAQLMLAAMLAMLATMAGQ
MLQMALATMAGQQQTAGMQIQMALILQMQLALAWVWVWVWVWVWVWVWVWV
MAGQQVWVWVWVWVWVWVWVWVWVWVWVWVWVWVWVWVWVWVWVWVWVWV
MLAMLQQANNNNNNNQMLQIAQMLIQAMQLAMQLAQTMG
>SERVER2
QMLQQQLAMQMLQQQLAMIIQQMLQAAQMLAMQQLAQTMGAMLAMLAM
QQLAQLAQVQLAMMQLATMAQTMAGQQTMAGQAMQLQIMLQANNNNNNNI
LALALALALALALANNNNNNNNNNNNNNNNNNNNNLALALALALALALALA
NNNNNNNNNNNNLALALALALALALALANNNNNNNNNNNNNNNNNNNNN
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNLLLLLLLLLLLLLLLLAAAAA
AAAAAAAAAAAAAAAAAALMLQAQTMAGQLAM
>SERVER3
MLAMLANNNNNNQMLAMLAQMLAMLAQQLAMLATGMAMLAAMLAQML
AAMLATGMALALLAALLAALLAALLAALLAALLAALLAALLAALLAALLA
QLLAAMLAMLATGMAQLLAQQTGMQQQLLAAMLATGMAQQTGMQQLALLA
ALLAALLAAMLATGMAMLALLAALLAAMLALLAAMLANNNNNNNNVTGMANN
LLAAMLAMLAQLLAAMLAMLAAMAMLATGMA
```

Figure 1: An example of "cyber proteins"

For example, to analyze a very large catalog of software, the MLSTONES team has created a mapping of machine codes to the amino acids that comprise a "protein." We use this mapping and some scalable, parallel protein alignment tools to generate families of similar binaries and, finally, create a single representation (motif) for each family. We've now reduced the data to analyze a new binary by several orders of magnitude and can very quickly place a new binary into its family membership. Even a previously unseen piece of software can be characterized by its behavior at the machine level without analyzing or executing the code. We've found that our approach

also works with many types of cyber data. We are currently researching new mappings for understanding malicious network activity and have analyzed text-based data, such as error logs and server behavior.

Benefits

The MLSTONES process can reduce extremely large data sets to much smaller sets of family motifs that enable identification or classification in near real-time. We can identify new objects of interest that are similar to known items and also identify completely new classes of objects. Our tools are customizable to the specifics of the data being characterized. With some research, completely new types of cyber data can be classified just by designing a new transformation function. While methods and tools similar to the MLSTONES approach have long been used in computational biology, none can match the speed of MLSTONES. Because MLSTONES can handle and process data in near real-time, we can apply it to the volumes and velocities of data found in cybersecurity applications.

Competitive Advantage

No other known technology uses this approach and obtains the same results as MLSTONES. There are specialized methods for analysis of other types of data, but none can support analytics on the scale required for cyber data. Generating signatures for cyber applications typically occurs in one of two ways—reactively or with expert knowledge. Reactive signatures are generated by reverse engineering their details and then building new rule sets or exact patterns for finding the same event in the future. Antivirus and network intrusion detection, for example, primarily operate in this mode and are plagued by the fact that they often cannot recognize new events that are highly related to prior events. On the other hand, expert knowledge signatures are obtained by asking subject matter experts to intuit what they believe are the most relevant attributes

to look for. While these are not constrained by the same limitations as reactive approaches, they can be heavily biased by the subject matter expert and may still be defeated by outside-the-box mutations on prior strategies. MLSTONES offers a third option that is guided, but not limited by prior events. MLSTONES can recognize similarities that are distantly related but still statistically significant. MLSTONES is also not biased by experts and can be used to **discover** relevant attributes from a large data set.

Next Steps

MLSTONES is currently under development for government clients. The general tools are being fully developed and tested at the Pacific Northwest National Laboratory. They're also being incorporated into a workflow engine and into a workbench environment. We are also in the process of researching the capability to identify malicious network activity in near real-time. We would like to put the tools into a pilot operational environment in order to fully test their capabilities and their scalability.

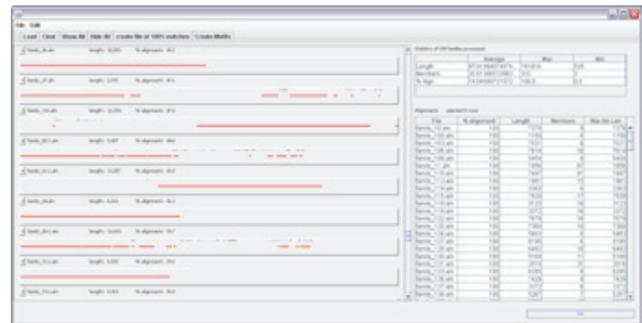


Figure 2: Visualization of family motif generation

CodeSeal: Tamper-proof Trust Anchors



Sandia
National
Laboratories

Adrian Chavez
adrchav@sandia.gov

John Solis
jhsolis@sandia.gov

Overview

CodeSeal is a cryptographically secure code obfuscation technology developed at Sandia National Labs that provides tamper-proof trust anchors to protect hardware and software running in compromised systems from malicious tampering.

Customer Need

The proliferation of counterfeit information technology products is a constant security threat faced by government departments and infrastructure operators. High-end counterfeit products are providing backdoor access to secure and sensitive systems due to compromised government supply chains. A recent study by KPMG and the Alliance for Gray Market and Counterfeit Abatement (AGMA) estimates that one in 10 IT products sold globally are counterfeit. How many IT products are currently operating in your organization?

Critical software must execute securely and with high fidelity in robust environments and with increasing pressing demands. Critical infrastructure systems require absolute assurance. Since many components are vulnerable throughout a product's lifecycle, we must assume that these systems are compromised before we receive them. IT organizations need cost-effective solutions that can be retrofitted into these systems to provide essential security properties of authenticity, confidentiality and integrity.

Our Approach

Trust anchors are functional elements that can be introduced into information systems to provide

Funded through the Department of Energy's (DOE) Office of Electricity Delivery and Energy Reliability (OE), Cybersecurity for Energy Delivery Systems (CEDS) R&D Program.

unbiased measurement and unimpeded control capabilities. These elements provide verification that systems are functioning correctly and can serve as a foundation for additional, independent security services. CodeSeal provides trusted execution in untrusted environments. CodeSeal is a cryptographically secure obfuscation technology that ensures trust anchors are tamper-proof and that an adversary cannot derive their function. CodeSeal's trust anchors serve to greatly reduce the risk of an adversary inserting malicious functionality into a protected system.

CodeSeal technology uses a customized compiler to obfuscate a software program, hiding the program's functionality from analysts and reverse engineers. The obfuscated code is executed with the aid of a trust anchor, which interprets the obfuscated code and ensures its integrity. Protecting this trust anchor is critical, as it is the key to deriving the functionality of the obfuscated code. Obfuscated code can only execute when in communication with the trust anchor, and remains obfuscated when executing and at rest.

Benefits

CodeSeal's obfuscation routines are based on established and widely accepted cryptographic standards that are provably secure. This allows our technology to benefit users by satisfying important security properties:

- Obfuscated code behaves as a true black box when the trust anchor is properly protected.
- The original algorithm experiences at most a polynomial time slowdown. Lab testing has demonstrated a linear slowdown with a coefficient of two.



CodeSeal: Tamper-proof Trust Anchors

- An adversary cannot detect what the device is measuring.
- An adversary cannot understand or modify program functionality.
- An adversary cannot subvert the system, as any modification will be immediately evident.

Competitive Advantage

Traditional software obfuscation tools typically operate on a source program by manipulating function and variable names or a program's control flow. Several public tools exist for a variety of languages, e.g., Proguard for Java. None are capable of preventing a dedicated, patient, and well-funded adversary from decompiling and reverse-engineering the obfuscated code.

The CodeSeal solution addresses the shortcomings of traditional obfuscation techniques:

- CodeSeal correctly assumes that an adversary is capable of analyzing complex systems.
- CodeSeal is provably secure and cannot be reversed engineered.
- CodeSeal can be configured to quickly assess software for vulnerabilities at a fraction of the time and cost of traditional tools.

Next Steps

Currently, CodeSeal is at technology readiness level of 5 with several demonstrations prototyping the technology. It is ready to be piloted and tested within an operational environment to secure critical software from malicious tampering on potentially compromised systems.

Through laboratory testing, performance metrics will be gathered and the algorithm will be further refined and customized for specific industry applications. The

algorithms within the CodeSeal technology have been designed to be highly optimized through pipelining in hardware. A software implementation is currently available for demonstration, and performance will drastically improve with a hardware implementation.

