Cyber Security Division Transition to Practice Technology Guide

Fiscal Year 2016

Homeland Security

Science and Technology

This page is left blank intentionally.

Introduction

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 culmination of extensive foraging effort to identify cybersecurity technologies developed at Department of Energy National Laboratories, Department of Defense affiliated laboratories and National Science Foundation funded academic institutions. 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 ready 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 fourth volume and it features eight new technologies, along with sixteen technologies from the last two volumes. As of January 2016, four of 24 technologies from the first three years have licensed, one open sourced while in the program and numerous others are in various stages of the licensing process. We're excited for the research teams and their licensing partners and wish them success on their journey to the marketplace. Ultimately, their success will result in better cybersecurity for the nation, the global internet community and you.

To help direct future publications please reflect on the cybersecurity capability gaps in your own organizations, 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 from federal government R&D community.

Sincerely,

Douglas Maughan

Douglas Maughan
DHS S&T Cyber Security Division
Director

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

CONTENTS

Department of Homeland Security (DHS) Science and Technology Directorate (S&T) Cyber Security Division (CSD)

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

Threats on the Internet change fast and cybersecurity is one of the most challenging areas in which the Federal government must keep pace. Next-generation cybersecurity 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 cybersecurity 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 cybersecurity 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 cybersecurity 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 cybersecurity 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 cybersecurity research community.

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

Cyber Security Education—Helping to foster adequate training and education programs critical to the nation's cybersecurity needs by providing opportunities for high

> *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*

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 cybersecurity 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.

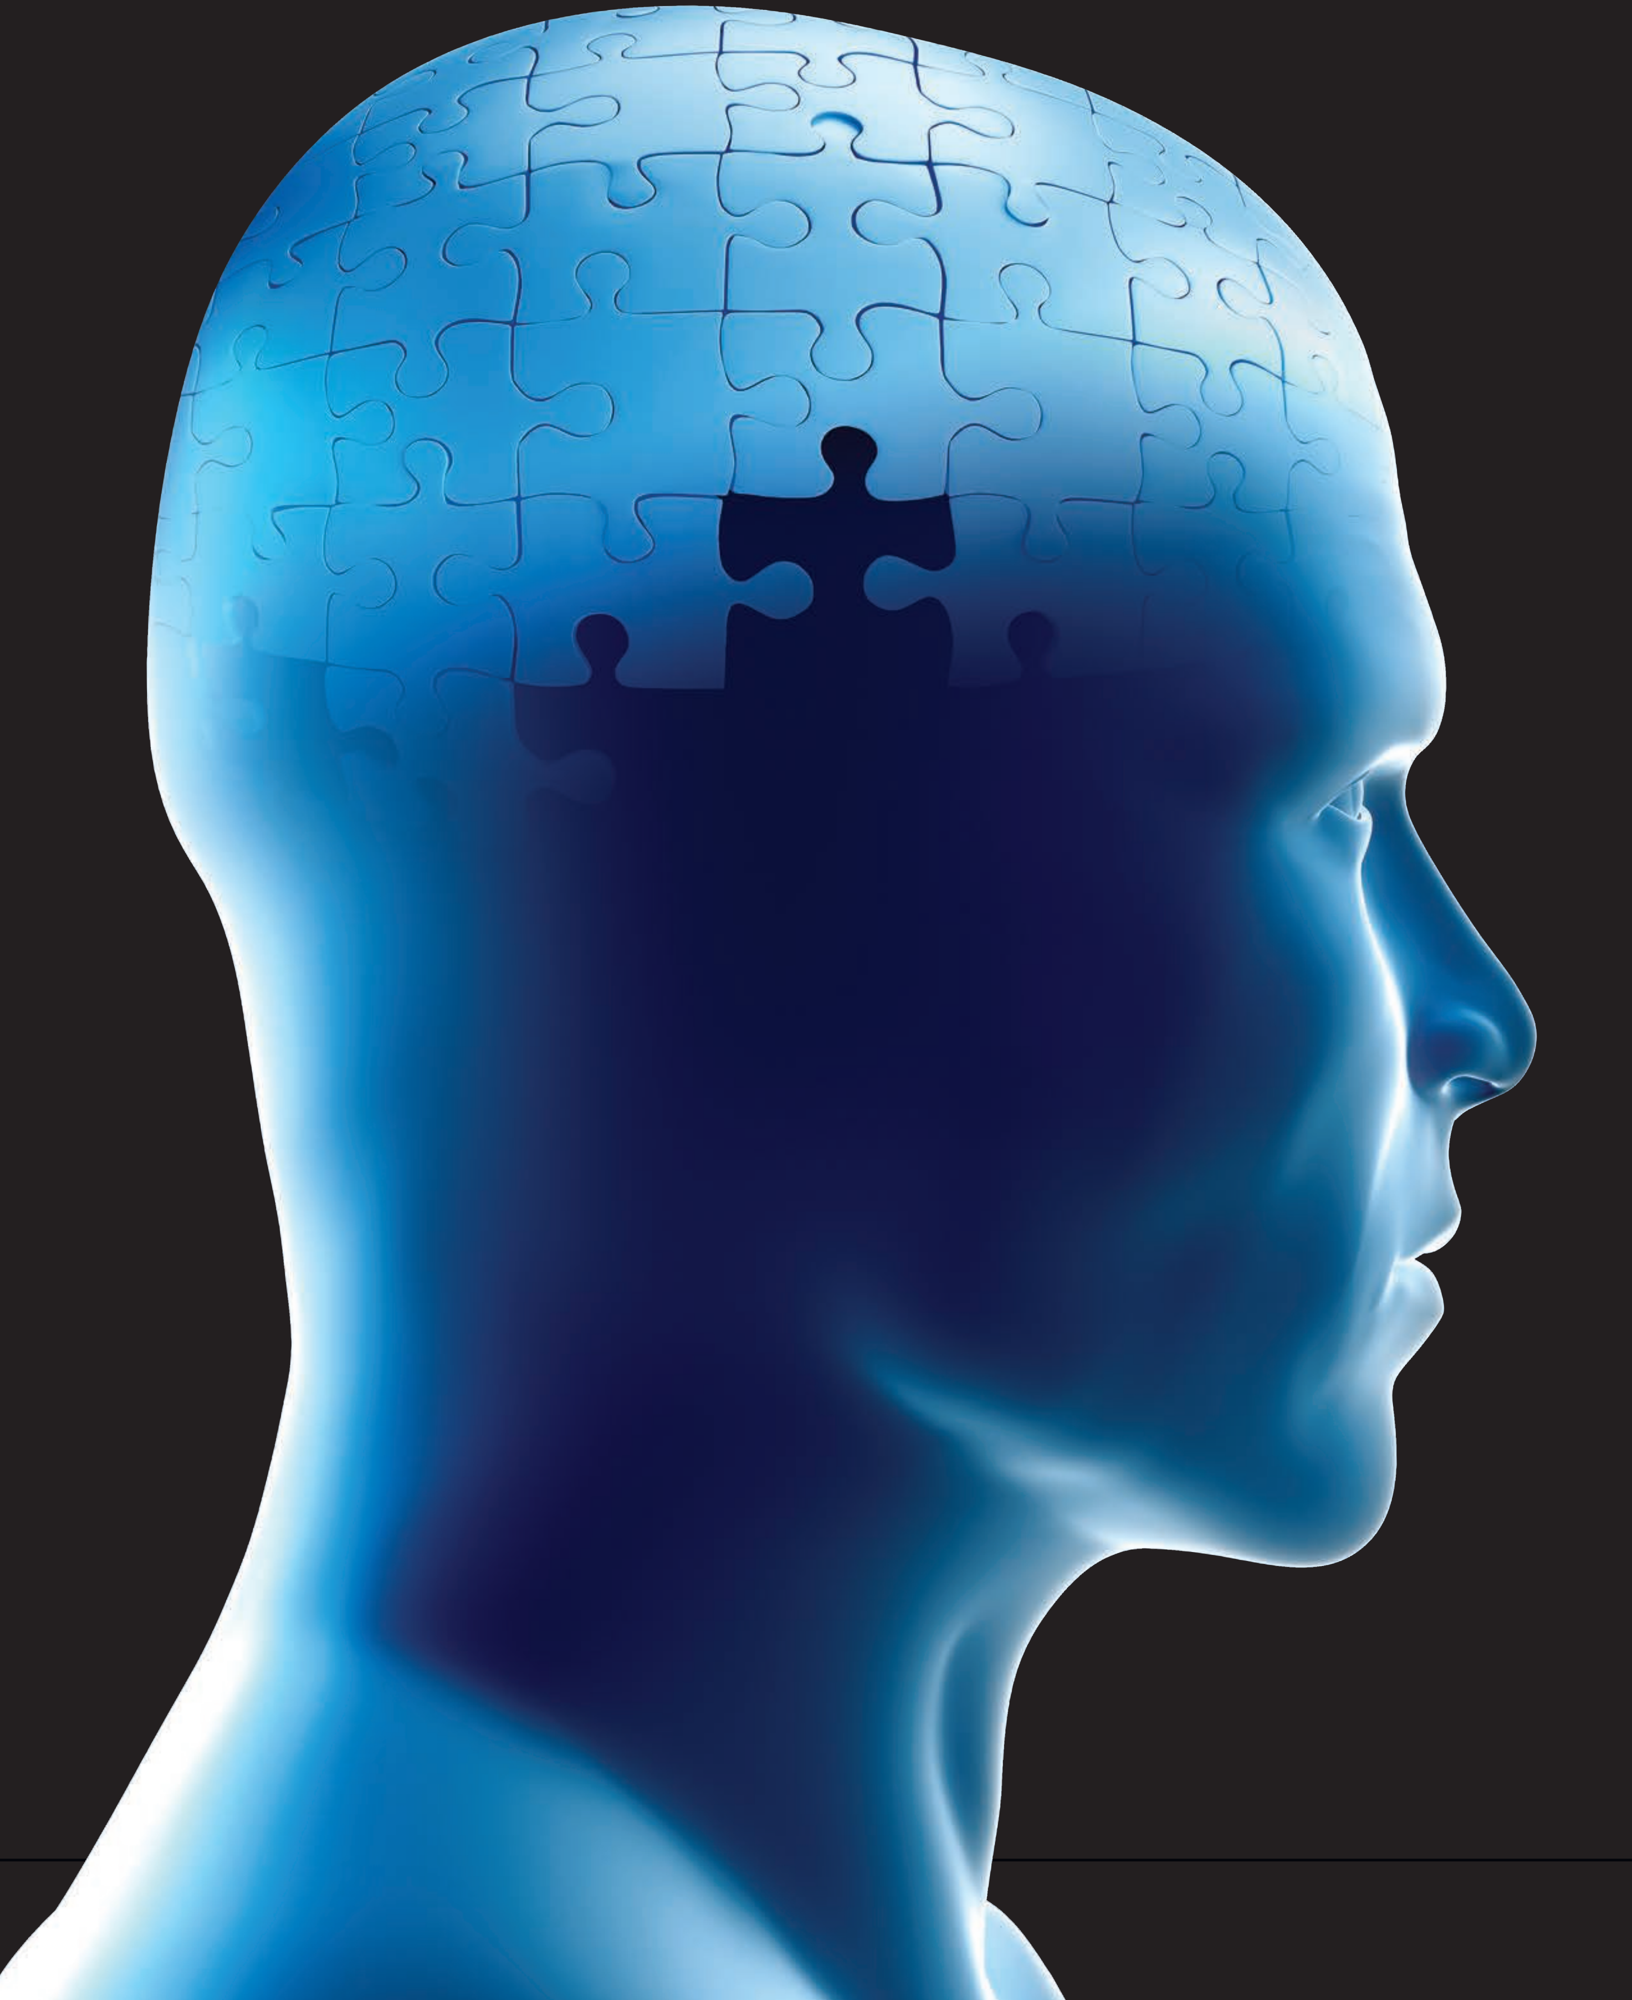

Transition to Practice: Accelerating the Pace of Technology Transition

Michael Pozmantier
michael.pozmantier@hq.dhs.gov

Turning Cybersecurity Research Into A Reality

Cybersecurity research is seldom commercialized, a fact all too familiar to researchers. The White House's Networking and Information Technology Research and Development (NITRD) program recognizes the potential gains that could be achieved through better commercialization practices. Since 2011, the federal government has made accelerating the transition of cybersecurity technology into widespread deploy-ment and use via the marketplace a priority in order to im-prove our nation's cybersecurity infrastructure. In response, the Department of Homeland Security tasked the Science and Technology Directorate (S&T) Cyber Security Division (CSD) with creating the Transition to Practice (TTP) program.

How It Works

The TTP program, which was initiated in the spring of 2012, builds on S&T's process of funding projects through the full research and development life cycle: research, development, test and evaluation, pilots, and transition.

The federal government spends more than $1 billon on un-classified cybersecurity research every year. However, very little of that research is ever integrated into the marketplace. The divide between the research phase and commercialization phase is commonly referred to as the "Valley of Death." Research is often stranded in that divide because researchers do not have the necessary resources to take the research to the marketplace, limited communication between researchers and the private sector commercialization community, and the lack of a clear understanding of the transition process among researchers, the private sector, and end users.

TTP Goals

The TTP program's goals are to: (1) identify mature technol-ogies that address an existing or imminent cybersecurity gap in public or private systems that impact national security, (2) increase utilization through partnerships, product develop-ment efforts, and marketing strategies, and (3) improve the long term ability for federal government research labs to more efficiently transition technology. The TTP program targets federally funded cybersecurity research that demon-strates a high probability of successful transition to the com-mercial market within two years and is expected to have a notable impact on the cybersecurity of our nation's networks or systems.

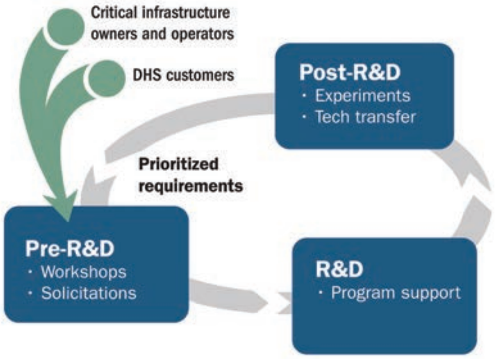

Cyber Security Division R&D Lifecycle

The Value

The TTP program is developing better lines of communica-tion between researchers and the investment community and funding activities that will improve the likelihood that tech-nologies will transition. For example, S&T conducts opera-tional test and evaluation to ensure stability, functionality, and refinement through technology pilots. Research teams will also be active participants in the commercialization pro-cess, thereby gaining valuable and lasting hands-on experi-ence with the commercialization process.

The goal of the TTP program is not only to accelerate the transition of cybersecurity research, but also to build lasting connections and processes that can be adopted by others and become self-sustaining—in essence, to build a lasting bridge over the "Valley of Death".

For more information about the TTP Program, email ST.TTP@hq.dhs.gov.

FISCAL YEAR 2016 TECHNOLOGIES:

REnigma: A Tool to Reverse Engineer Malware

Socrates: Graph Analytics for Discovering Patterns and Relationships in Large Data Sets

PcapDB: Optimized Full Network Packet Capture for Fast and Efficient Retrieval

REDUCE: Collaborative, Statistically Guided Exploration of Malware Similarities

Dynamic Flow Isolation: Adaptive Access Control to Protect Networks

TRACER: Timely Randomization Applied to Commodity Executables at Runtime

FLOWER: Network FLOW AnalyzER – Deep Insight Into Network Traffic

SilentAlarm: Detecting Abnormal Network Traffic

REnigma: A Tool to Reverse Engineer Malware

Julian Grizzard
julian.grizzard@jhuapl.edu

James Stevens
james.stevens@jhuapl.edu

Johns Hopkins
Applied Physics Laboratory

Overview

When an organization is under cyber attack, there are numerous questions that need to be answered in a timely fashion. How did the attacker get in? How bad is the damage? Who is behind the attack? How can further damage be prevented? To maximize the impact of an attack, the adversary's goal is to increase the difficulty of answering these questions. Obfuscation of executable code prevents static analysis, encrypted communication prevents network analysis, and anti-analysis techniques prevent dynamic analysis. REnigma helps malware analysts regain the upper hand against advanced malware techniques by transparently and precisely recording the execution of malware, and it enables analysis that can extract the level of detail necessary to answer the vital questions needed to understand and recover from a cyber attack quickly and accurately.

Customer Need

The analysis of malware used in a cyber attack is a very manual, time consuming, low throughput, and costly process requiring days to weeks to give the answers needed to clean up the attack and prevent further damage. Existing approaches include static analysis that cannot cope with packed malware and dynamic analysis that is either easy for malware to detect or is extremely time-consuming. It is critical that defenders utilize state-of-the-art techniques that are transparent and provide quick, scalable, and in-depth analytic capabilities.

Our Approach

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) REnigma technology uses Virtual Machine Record and Replay (RnR) to precisely and transparently record execution of malicious code so that an analyst can then replay and analyze the execution in detail. RnR provides many powerful techniques for malware analysis that are not possible today because it enables the ability for the analyst to "rewind" to any previous state in the system without affecting the execution of the code under test. This approach enables instruction-level analysis algorithms such as exploit detection and data flow analysis to operate without being detected by anti-analysis techniques used by malware. For example, if a malicious code sample outputs encrypted data on the network, an analyst can use REnigma to backtrack to the plaintext data in memory or recover the encryption key used for exfiltration.

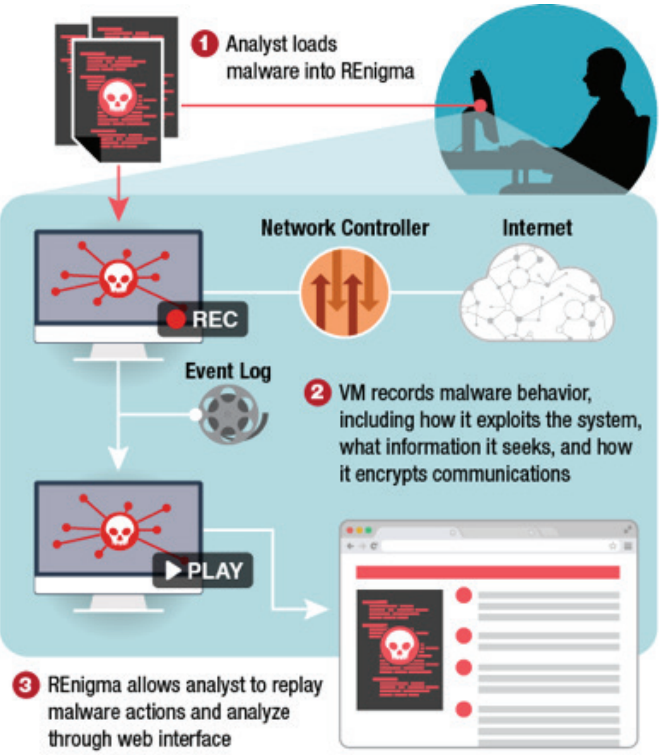
1 Analyst loads malware into REnigma
Network Controller
Internet
REC
Event Log
2 VM records malware behavior, including how it exploits the system, what information it seeks, and how it encrypts communications
PLAY
3 REnigma allows analyst to replay malware actions and analyze through web interface

REnigma analysis consists of three stages.

First, the analyst loads suspect malware into REnigma. Second, REnigma launches a virtual machine, copies the malware into the virtual machine, and begins recording execution. During this stage, the malware executes inside the virtual machine exactly as it would in a normal system so that its behavior is captured. Additionally, the analyst can configure network access to the virtual machine to either expose the malware to an untethered "live" Internet connection, capturing remote command-and-control communication, or a controlled "fake" Internet connection that responds with false data. In the final stage, REnigma performs automated analysis and generates a report summarizing the malware's behavior, such as exploitation methods, indicators of compromise, and decrypted command-and-control communication. The analyst can also use an indexed timeline to quickly jump to points of interest captured during recoding and manually examine the behavior in detail.

Benefits

REnigma provides an analyst with a new capability to analyze malicious malware samples to understand their functionality at a level of detail not previously possible. Additionally, REnigma is designed to integrate standard tools, allowing the analyst to retain and leverage existing skills. For example, a user can replay execution and stop at various points during replay and dump system memory. This memory image can be fed into Volatility, which is an industry standard tool for extracting key artifacts from raw memory dumps. Furthermore, REnigma incorporates a framework to create new modules that can extract arbitrary information during replay. Advanced analysts can employ REnigma's modules as well as create their own custom modules.

Competitive Advantage

The key technology behind REnigma is JHU/APL's Virtual Machine Record and Replay capability. Record and replay research prototypes over the past 20 years were never fully developed, were not robust, did not have high performance, or are not available to others. JHU/APL developed RnR by modifying the open source Linux kernel and QEMU software systems. The RnR capability can record operating systems and applications running at speed with a modest 5% slowdown compared to a virtual machine that is not recorded. During replay, the precise execution of a malware sample can be recreated with instruction-level precision.

REnigma's ability to perform in-depth instruction-level analysis without disturbing the code sample reduces the need for expert reverse engineers to load the code in tools like IDA Pro and manually edit it to remove anti-analysis checks and force the malicious code to execute. REnigma allows security conscious organizations to avoid immediately resorting to manual reverse engineering as anti-analysis techniques become increasingly common, potentially saving tens of thousands of dollars per code sample analyzed.

Next Steps

JHU/APL researchers have used REnigma to examine newly discovered malware found on the JHU/APL network. We seek operational partners looking for cutting edge malware analysis tools, and plan to work closely with the partners to improve REnigma's capabilities. We are also seeking additional sponsorship to further develop these capabilities and scale up our deployments.

Socrates: Graph Analytics for Discovering Patterns and Relationships in Large Data Sets

JOHNS HOPKINS
APPLIED PHYSICS LABORATORY

Cetin Savkli
Cetin.Savkli@jhuapl.edu

Ryan Carr
Ryan.Carr@jhuapl.edu

Mike Lieberman
Mike.Lieberman@jhuapl.edu

Overview

SOCRATES is a flexible, easy to use graph analytics tool designed to discover patterns and relationships in large scale and complex data sets. Such data sets can be found in cyber, social, financial, energy, and biological domains. SOCRATES has been successfully used to discover previously unknown patterns in real world big data sets. Examples include detecting illegal international trade, discovering unknown associates of persons of interest from travel patterns, and detecting anomalous flight behaviors. SOCRATES can be readily applied to cyber and cybersecurity data.

Need

Scope and complexity of data sets as well as evolving analytic challenges make rapid development of analytics a critical need. A key challenge in big data analysis is the human skillset required to properly store and analyze immense quantities of data. People who excel at analysis may not have the necessary computer science knowledge relevant for big data analysis and vice versa; drawing conclusions from big data requires both skill sets.

Another critical challenge is development of unsupervised methods for analysis of data. Analytic approaches based on classification or rule based techniques are unsuitable for large scale and complex data sets as data is typically high dimensional and shows great variability. Relying on past examples of bad behavior is not sufficient for detecting future threats. There is a need for development of scalable unsupervised analytics to learn patterns directly from data.

Many big data questions are domain independent. For example: determining correlations in data; discovering patterns of behavior and associated anomalies; discovering links and networks; identifying critical nodes for network resiliency, and spread of virus/information; determining central nodes, leaders, and power brokers. All of these analytic questions must be addressed in a manner that fusion of data and implementation of analytic ideas are both simple and scalable.

Approach

SOCRATES is flexible, easy to use graph analytics software tool designed to discover patterns and relationships in large scale and complex data sets. It features several advances in parallel computing and scalable distributed storage and uses a flexible graph model to represent complex data sets and knowledge. Every attribute of data is automatically indexed for fast random access and analytics processing.

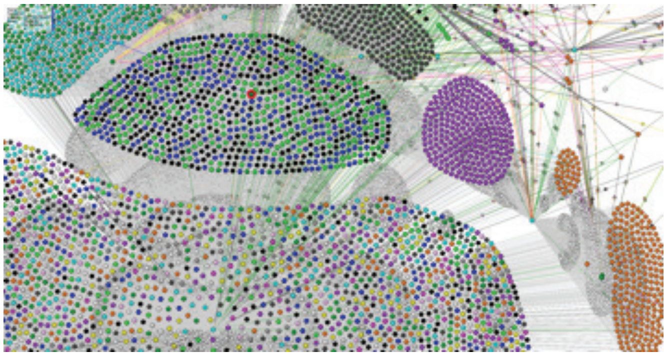

Detection of anomalous activity from netflow data in a network with 2,000,000 links.

SOCRATES' analytic capabilities are based on a probabilistic representation of data that captures a concise expression of knowledge. It uses this approach to provide anomaly detection and classification capabilities for high dimensional data including temporal behaviors. Most of the analytics and supporting correlation libraries are parallelized to take advantage of the computing power of a distributed hardware cluster.

SOCRATES also provides a library of link inference and network clustering algorithms. These algorithms work together to facilitate community based behavior analysis.

Benefits

SOCRATES provides a simple and scalable software platform and a library of unsupervised machine learning algorithms for big data analytics. Simplicity of analytics at scale allows developers and sponsors to rapidly explore ideas and leads to increased productivity both in terms of results and cost. Implementation of analytics has been done with a focus uncovering probabilistic knowledge and patterns in large scale data into without an assumption on the availability of ground truth or categorization. SOCRATES provides a robust set of parallelized algorithms for anomaly detection in high dimensional spaces, temporal analysis, and correlation analysis.

SOCRATES' flexibility of graph representation facilitates fusion of diverse sources of data and simplifies management of data complexity. Automated indexing of attributes and advanced query capability facilitates rapid implementation of analytic ideas on complex data sets.

The combination of these benefits has led to development of analytic capabilities that have been successfully used to discover previously unknown patterns on real world data sets that can readily apply to cyber and cybersecurity data.

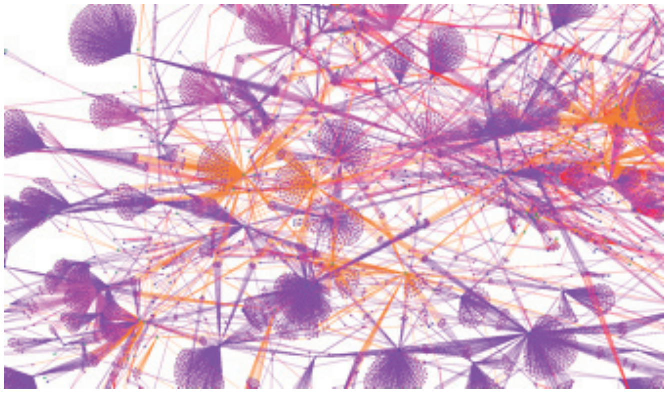

Graph of global trade transactions with more than 1 billion links analyzed using Socrates to find anomalous transactions.

Competition

In addition to providing a robust set of analytic capabilities for behavior analysis, anomaly detection, and graph analytics, SOCRATES overcomes key issues in automated analysis of large data sets. NoSQL systems such as Accumulo & HBase face challenges that make analyzing big data difficult. SOCRATES provides secondary indexing for improved query performance, locality control to avoid unnecessary movement of data, and a schema that overcomes database maintenance challenges.

Traditional relational database management systems (RDBMS) also face challenges when dealing with big data. SOCRATES provides table structures that are flexible enough to easily support new kinds of data and better parallelization to increase scalability.

SOCRATES offers key advantages over the alternatives: a) The locality of graph elements can be controlled, a feature essential for not moving data in large scale graph analytics; b) All of the attributes of graph elements are indexed for fast query processing; c) Provides a parallelization paradigm that is close to standalone programming; and d) Cluster is not centrally managed.

The biggest competitive advantage SOCRATES provides is to make big data analysis as simple as possible and that has been the key to its success.

Next Steps

SOCRATES is a flexible, easy to use, large scale data analytics tool for use by technical users in a controlled environment. The success of analytic results using SOCRATES has sparked sponsor interest and it is being prepared for deployment at various sponsor sites. We seek additional partners who can deploy and apply SOCRATES data analytics to their large cyber and cybersecurity datasets.

PcapDB: Optimized Full Network Packet Capture for Fast and Efficient Retrieval

Paul Ferrell
pferrell@lanl.gov

Shannon Steinfadt
shannon@lanl.gov

Overview

Capturing and analyzing network traffic is key to protecting and detecting intrusions on networks and systems. PcapDB is a database system designed from the ground up. It optimizes full network packet capture for fast and efficient retrieval. Packets are reorganized and indexed by flow before they are ever written to disk. PcapDB gives cyber analysts and incident responders fast search and retrieval capabilities while limiting disk access, unlike any other open source tool. It is a software solution designed for easy deployment on commodity hardware. PcapDB allows for large-scale installations at a significantly lower cost than existing commercial solutions.

Customer Need

Cyber security incidents are often discovered hours, weeks, or even months after they happen. On average, advanced persistent threat (APT) actors are inside networks and systems for one year, with multi-year durations up to five years of entrenchment reported, before being discovered. To fully assess the threat and scope of cyber incidents, analysts need access to full network packet capture (pcap) files for as much of the incident time frame as possible. Full packet capture gives an analyst a complete picture of network traffic during the time of the incident, similar to a black box flight recorder. Pcap allows reconstruction of malware transfers, downloads, command and control messages, and exfiltrated data.

To be effective, the captured packets should be easily searchable and have a long a history (the largest amount of pcap) as possible. Response time is also important; while adversaries may have months to prepare an attack, analysts have comparably little time to create a comprehensive and effective response once detected. Full packet captures enable analysts to assess and respond to the current attack, and prepare for the next one.

Our Approach

Full packet capture, while highly dependent upon the speed at which a system can move packets from a network interface onto disk, is not otherwise a very hardware intensive process. PcapDB uses the abundance of CPU cycles and memory in modern servers to optimize the captured packets for indexing and storage before they are ever written to disk. It organizes the data by the entire transfer (flow) rather than individual packets. The flow-ordered pcap (fcap) provides data in units the users expect rather than how it happens to come in over the wire.

The indexing structures and algorithms of PcapDB are designed to take full advantage of this flow ordering and other unique properties of the packet capture problem. In doing so, we minimize disk interaction, thus greatly improving system performance. Additionally, the indexing techniques consume significantly less disk space than other solutions, leaving 99% of the disk for captured packets. This allows users to make the most of the disk investment required by packet capture systems.

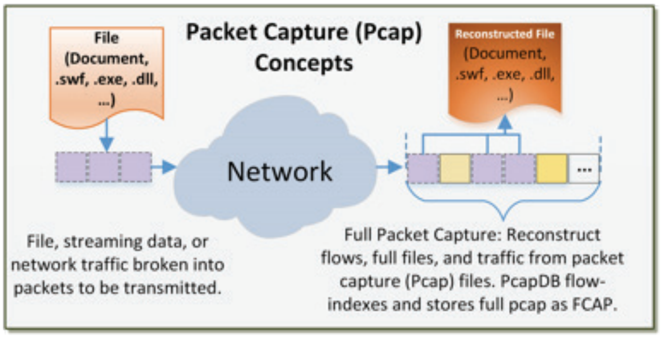
Packet Capture (Pcap) Concepts
File (Document, .swf, .exe, .dll, ...)
Network
Reconstructed File (Document, .swf, .exe, .dll, ...)
...
File, streaming data, or network traffic broken into packets to be transmitted.
Full Packet Capture: Reconstruct flows, full files, and traffic from packet capture (Pcap) files. PcapDB flow-indexes and stores full pcap as FCAP.

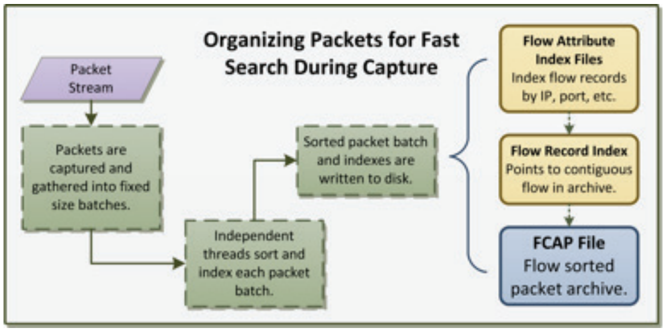
Organizing Packets for Fast Search During Capture
Packet Stream
Packets are captured and gathered into fixed size batches.
Independent threads sort and index each packet batch.
Sorted packet batch and indexes are written to disk.
Flow Attribute Index Files
Index flow records by IP, port, etc.
Flow Record Index
Points to contiguous flow in archive.
FCAP File
Flow sorted packet archive.

PcapDB is inherently parallel, designed to scale across multiple capture nodes to handle higher capture rates. Queries are resolved using a Map-Reduce-like model, first distributed amongst the various capture nodes and then combined into the final result on a central search system. This allows for distributed capture that is scalable to meet different capture needs and requirements.

Benefits

PcapDB pares down the packet capture features to exactly what cyber analysts and incident responders need to do their job effectively: swift searching across the longest, logically ordered, history achievable. The result is a system where users can rapidly investigate intrusions and potential threats from an institution with a greater chance for success.

PcapDB's optimizations greatly reduce the amount of time needed to retrieve stored packet captures. Individual network flows can be retrieved with a single, contiguous disk read with little to no reading of unnecessary data. Indices can similarly be searched with only a few disk seeks each. Since reading from disk is by far the slowest part of data retrieval, the elimination of wasteful disk access allows PcapDB to retrieve results almost as quickly as physically possible on modern hardware.

PcapDB eliminates the need to pay for costly commercial packet capture hardware. We provide a guide for building affordable capture systems using commodity hardware. The savings can be used to expand the usefulness of the resulting system. Rather than purchasing a capture system that can only hold a few days' worth of capture data, PcapDB can store weeks or months of capture data on low-cost commodity Serial Attached SCSI (SAS) disks and "just a bunch of disks" (JBOD) enclosures. The longest history possible is key when investigating a cyber incident.

Competitive Advantage

PcapDB index files are designed specifically around packet data and the write-once, read-many nature of captured data to allow for extremely efficient usage of precious storage space.

Existing commercial packet capture technologies, while capable of high speed capture, prioritize increasing their feature sets in order to sell the capture hardware rather than improve their basic packet storage and retrieval functionality. These additional "features" typically result in an index structure that consumes a significant fraction of the available disk space. This disk space is better utilized storing captured packets, giving a longer history. PcapDB can retain over 90% more capture data on disk over some commercial systems with huge indexes and unnecessary software.

Commercial systems are typically built upon commodity hardware sold at a 10x markup for $200K or more with just a few days of retention. PcapDB capture nodes, with 150 TB of SAS storage (~14 days of capture at 1 Gb/s), can be built for as little as $20K and only require minor assembly. Lowering hardware costs by an order of magnitude allows users to affordably scale their deployments both to capture at higher data rates and to have a longer data retention period.

Next Steps

The core functionality of PcapDB is complete. The user interface is rapidly nearing a deployable state for a production environment. PcapDB is currently being tested on networks with live data, up to 100 Gb/s of bandwidth. We seek partners to deploy and utilize this technology at their sites.

REDUCE: Collaborative, Statistically Guided Exploration of Malware Similarities

Juston Moore
jmoore01@lanl.gov

Overview

REDUCE is a software toolset enabling cyber security analysts to rapidly discover relationships between malware samples, to extract temporal threat intelligence, and to develop actionable signatures for known and emerging threats. REDUCE performs automated static code analysis and identifies similarities between malware samples in order to support knowledge sharing about related pieces of malware. By integrating with well-established reverse engineering tools, REDUCE speeds up both deep-dive reverse engineering efforts and custom signature generation.

Customer Need

Large enterprise networks are constantly under attack by a huge number of threat actors. Each threat actor commonly deploys many variants of the same malware in order to defeat anti-virus defenses. Responding to the constant emergence of new and evolving malware threats, cyber defenders require the ability to easily compare new malware with previously seen malware samples, and to efficiently develop and deploy custom signatures to guard against targeted attacks.

Manual analysis of malware yields invaluable information, but analyzing many similar malware samples one at a time is very inefficient. Cyber security analysts currently lack tools to perform in-depth analyses of more than one or two pieces of malware at once. Analyzing similarities between sets of malware, rather than characteristics of individual samples, allows an analyst to work more efficiently and to easily transfer knowledge from previous analyses on similar pieces of malware.

Our Approach

REDUCE provides a centralized software toolset for automatically performing static analysis on a collection of malware samples. Once samples have been uploaded and processed, REDUCE uses statistical techniques to identify similar code sections across many malware samples. Identified similarities enable an analyst to leverage existing knowledge about a small set of samples in order to rapidly make inferences about the authorship and technical characteristics of new samples.

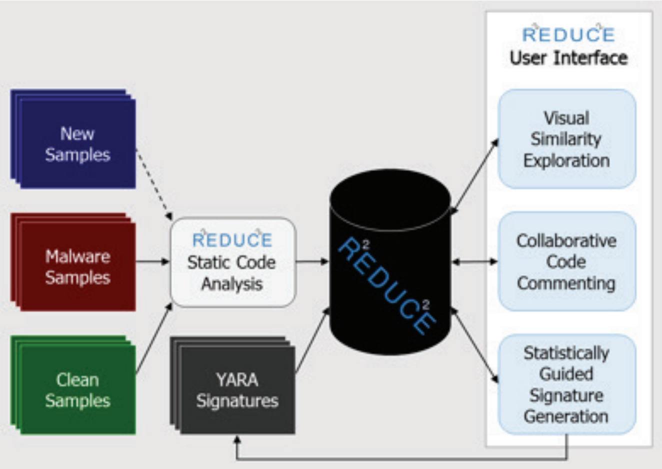

In addition to providing guidance and insight during an in-depth examination of malware, REDUCE also allows an analyst to quickly construct robust signatures for known and emerging threats. The REDUCE signature generation process is guided by information theoretic principles, which help to identify features that are distinctive to the set of interest, while eliminating features common across software samples in general. Signatures are generated in the open-source YARA format. YARA signatures can be deployed on a wide variety of security appliances and easily translated into other formats to be deployed across an enterprise.

Benefits

REDUCE is a practical toolset developed by analysts in an operational incident response environment. REDUCE helps analysts to focus on specific code patterns and threat actors rather than peculiarities of individual malware samples.

REDUCE enables a reverse engineering workflow and complements tools commonly used by security analysts. The REDUCE user interface gives analysts a big-picture view over collections of malware, while also facilitating deep-dive investigations. Using statistically guided reverse engineering, analysts can uncover similarities among a potentially large set of related malware samples.

For both experienced reverse engineers and junior analysts, REDUCE shortens the time required to construct effective signatures for new and emerging threats. REDUCE rapidly uncovers shared code between new malware samples and known samples, even if existing signatures do not detect the new samples. REDUCE also serves as an expanding knowledge resource and enhances collaboration between analysts by propagating comments between similar functions in different samples.

Competitive Advantage

The REDUCE toolset improves upon the capabilities of popular commercial and open-source binary similarity tools, such as Zynamics BinDiff, diaphora, and radare2, all of which compare only two malware samples. REDUCE identifies similarities across multiple malware samples.

Many machine learning systems, including commercially deployed solutions such as the VirusTotal, perform similarity analysis for sets of malware samples. However, these systems act as black boxes, and do not allow an analyst to directly interact in the decision-making process. In contrast to most of these systems, REDUCE displays specific similarities identified along with existing knowledge about those particular patterns. By involving analysts in critical decision-making processes, REDUCE produces information and deployable signatures that capture operational awareness criteria.

Manually constructed YARA signatures are the current de facto format for government and industry intelligence reporting. By utilizing the REDUCE signature generator, analysts at Los Alamos National Laboratory have reduced the time needed to produce quality YARA signatures from hours or days to minutes. Signatures created with REDUCE have consistently identified more malware samples of interest than commercial anti-virus or YARA signatures obtained from open source industry threat reports. REDUCE signatures also have a very low false positive rate in testing.

Next Steps

The REDUCE software is in the final phases of development and testing by security analysts within Los Alamos National Laboratory's Computer Security Incident Response Team. We seek to establish partnerships for pilot testing at other sites with in-depth reverse engineering capabilities. REDUCE can be used for malware analysis and signature generation by security practitioners with little knowledge of reverse engineering. We welcome feedback on the interpretability of the tool and its ability to work with existing reverse engineering processes.

Dynamic Flow Isolation: Adaptive Access Control to Protect Networks

Thomas Hobson
thomas.hobson@ll.mit.edu

Richard Skowyra
richard.skowyra@ll.mit.edu

Veer Dedhia
veer.dedhia@ll.mit.edu

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

This material is based upon work supported by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract No. FA8721-05-C-0002. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Assistant Secretary of Defense for Research and Engineering.

Overview

Dynamic Flow Isolation (DFI) improves network security by dynamically changing access control in response to the current operational state or business need. DFI leverages Software-Defined Networking (SDN) to apply security policies on-demand to all systems on an enterprise network. Communications between individual users and services can be enabled, disabled, or rate-limited based on both automatic and human-in-the-loop decision systems.

Customer Need

Today's networks suffer from an over-connectivity problem that results in nearly all end-hosts being mutually reachable at all times. Paths that may need to exist under a particular set of circumstances are typically available under all circumstances. Leaving these unused paths in place allows adversaries to use them to move laterally within a network, often initially entering through easily compromised hosts (e.g. via phishing emails) and moving towards higher-value targets.

Improved mechanisms are required for selectively enabling and disabling paths in order to achieve the minimal connectivity demanded by the operational state or business need. For instance, at any given time a device should only be able to reach those destinations required for the current task of the actively logged in user.

Our Approach

DFI is a software-based solution that integrates with SDN controllers to align network connectivity with the operational state by changing access control in response to both automatically and manually generated events. It integrates with third-party network services such as authentication servers and intrusion detection systems (IDS) to be informed of events indicative of changes to operational state. Network connections are enabled, disabled, or rate-limited according to these events and in conjunction with a specified policy.

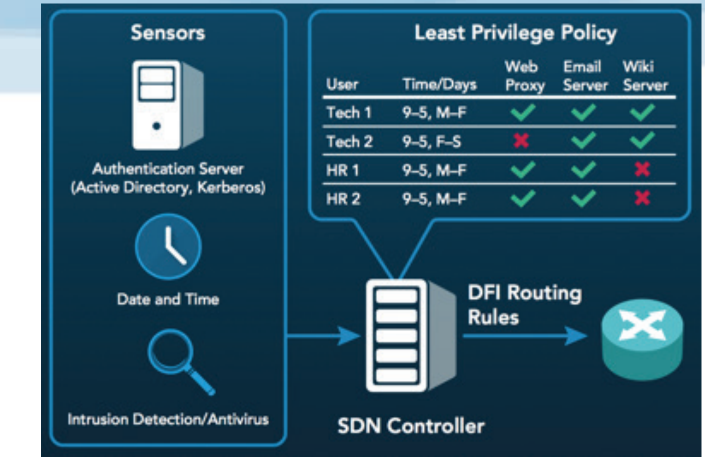
Sensors
Authentication Server (Active Directory, Kerberos)
Date and Time
Intrusion Detection/Antivirus
Least Privilege Policy
DFI Routing Rules
SDN Controller

User	Time/Days	Web Proxy	Email Server	Wiki Server
Tech 1	9–5, M–F	✓	✓	✓
Tech 2	9–5, F–S	✗	✓	✓
HR 1	9–5, M–F	✓	✓	✗
HR 2	9–5, M–F	✓	✓	✗

DFI enforces adaptive access control in accordance with a specified policy and sensor input about the current state.

DFI leverages recent advances in SDN to dynamically control network access in a manner that's scalable to large enterprise networks. A small policy enforcement kernel is implemented within SDN controllers that updates access rules for all switches on the network. DFI is designed to work with existing SDN hardware, be portable across SDN controllers, composable with other SDN applications, and extensible with new third-party services.

Benefits

DFI provides a means to enforce the principle of least privilege on networks and enables numerous capabilities that are impractical on existing networks.

- DFI can enforce granular connectivity changes in response to user login and logoff activity: the network can be configured to grant a device limited access to network resources until a user successfully logs in, at which point connections to specific resources are enabled, as required by that user.
- DFI can be utilized to implement contingency plans and quickly enable alternative network configurations.

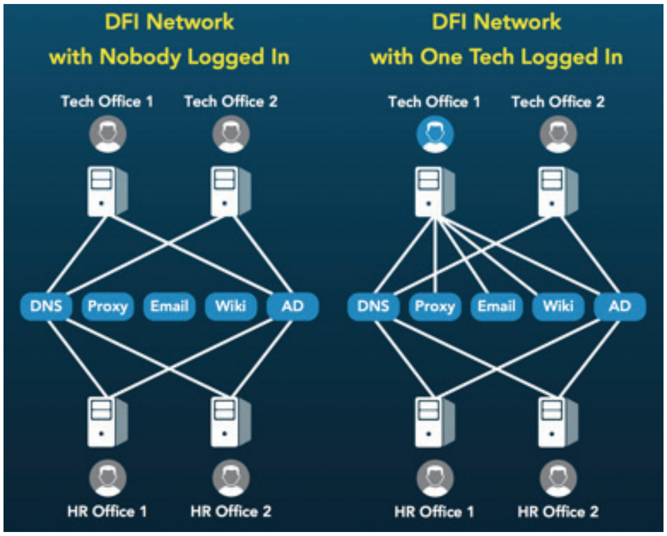

Authentication-triggered network access control

- Devices can be quarantined during sensitive operations or when an intrusion has been detected via a human analyst or IDS.
- Incident response teams can take advantage of DFI to quickly alter the network in order to prevent active adversaries from accessing the most critical components while maintaining the ability to observe adversary behavior.

Competitive Advantage

Existing firewalls are effective at protecting against certain threats from external networks but do little to prevent attacks that involve lateral movement within an internal network. Additionally, firewalls are statically configured, with changes or updates requiring manual effort, on human timescales. In contrast, DFI allows for firewall-like functionality for every port on every switch, with updates that can be applied automatically at computer timescales.

Network Access Control (NAC) solutions allow enforcement of pre- and post-admission policies, primarily designed to restrict the access of non-compliant devices (e.g. unauthorized or missing patches). DFI, in addition to its broader capabilities for dynamically shaping network access, enables similar compliancy mitigations but via an entirely software-based approach that eliminates the need for proprietary hardware, is designed to be extensible and scalable, and allows for visibility into and control over policy enforcement mechanisms.

Next Steps

We're continuing to expand our integration with third-party services (e.g. IDS's) and to pilot DFI on government networks.

We're looking for opportunities to broaden the utilization of the technology and are seeking organizations interested in deploying DFI within their environment, networking or security-centric companies interested in integrating DFI into their product suites, and incident response teams seeking proactive ways to both prepare for and to shape a network during an attack.

TRACER: Timely Randomization Applied to Commodity Executables at Runtime

Hamed Okhravi
hamed.okhravi@ll.mit.edu

This material is based upon work supported by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract No. FA8721-05-C-0002. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Assistant Secretary of Defense for Research and Engineering.

Overview

Timely Randomization Applied to Commodity Executables at Runtime (TRACER) protects closed-source Windows applications against sophisticated, modern attacks by automatically and transparently re-randomizing their sensitive internal data and layout.

Customer Need

Sophisticated, large-scale attacks against popular closed-source applications such as Adobe Reader, Internet Explorer, Java, and Flash have become widespread in recent years. With such attacks, adversaries can take control of a computer remotely to exfiltrate sensitive information or steal user data. These attacks often compromise millions of machines at once.

A significant problem contributing to large-scale attacks is the homogeneity of the targets. When the attackers develop an attack against an application, since all installations of that application look alike, it will be easy for them to compromise millions of computers at once.

Another factor contributing to this problem is the closed-source nature of the applications running on the proprietary Windows operating system. According to a report, more than 90% of desktop computers run Microsoft Windows with closed-source applications. Many cyber protections and defenses rely on having the source code available which makes them non-applicable to such environments.

To properly protect against large-scale attacks on closed-source applications, a diversification technique is needed that changes the sensitive internal data and layout of an application.

Our Approach

TRACER automatically and transparently randomizes key internal data and layout of an application at runtime to prevent modern control hijacking attacks. Figure 1 illustrates this concept at a high level.

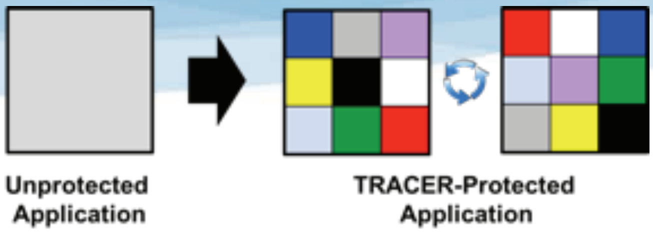

Figure 1: TRACER frequently randomizes the sensitive internal data and layout of the application after every possible leakage point.

The internal data and layout include stack cookies and heap metadata that protect static and dynamic memory, as well as the addresses of dynamically linked libraries (DLLs). Attackers frequently target these sensitive regions because they allow the attacker to take control of the application remotely. Figure 2 illustrates the concept of TRACER protecting DLLs.

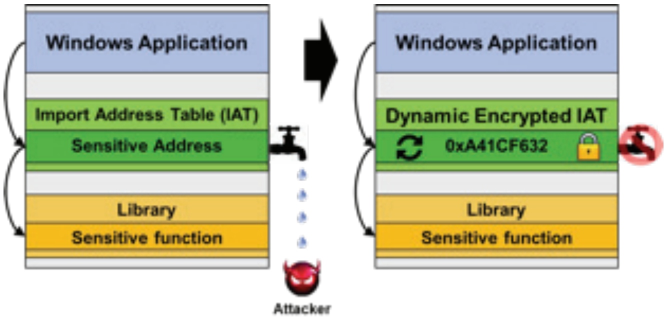

Figure 2: TRACER prevents leakage of application's sensitive internal data or layout to attackers. The addresses of linked libraries are dynamically re-encrypted after every possible leakage point.

TRACER re-randomizes the sensitive internal data and layout at every output from the application. Since vulnerable applications can leak how their internals have been randomized, it is crucial to continuously re-randomize these values. A time-based re-randomization would still be vulnerable because the leakage and the attack can happen within the short period of time. As such, TRACER implements an output-based re-randomization strategy to thwart a potential attacker. With this re-randomization strategy, any information leaked by the application will be stale when attackers attempt to exploit it.

Benefits

TRACER prevents the most common and sophisticated control hijacking attacks against Windows applications. According to a survey, these attacks constitute more than 50% of attacks and are commonly used by Advanced Persistent Threat (APT) actors.

TRACER is implemented as a single DLL, and unlike other defenses in this domain does not require access to the source code or modification of the Windows operating system. TRACER only takes minutes to install on each machine and is seamless to operate after the initial installation. TRACER does not interfere with normal maintenance, patching, software inventory, or debugging facilities of an enterprise network.

TRACER incurs an average 12% increase in execution time with common Windows applications. The overhead is often masked by the normal application execution delays and is not noticeable by the users. Since most computer systems under-utilize resources such as the CPU, the incurred overhead is likely to be acceptable in most enterprise environments.

Competitive Advantage

The main competitors for TRACER are randomization techniques such as memory layout randomization (Address Space Layout Randomization), compiler-based code randomization, and instruction set randomization techniques. All these techniques employ a “one-time” randomization strategy which makes them vulnerable to information leakage attacks. Using information leakage, attackers can analyze how the application has been randomized and undo the impact of randomization. In fact, information leakage attacks are used frequently to bypass one-time randomization defenses in the wild.

By re-randomizing the sensitive internal data and layout of an application every time any output is generated, TRACER renders leaked information stale and resists attacks that can otherwise bypass randomization defenses. TRACER provides security guarantees that are stronger than all of the previously mentioned techniques.

Unlike many other defenses for closed-source applications, TRACER does not rely on emulation techniques that incur unacceptably high overhead.

Next Steps

TRACER is ready to be deployed in an enterprise to protect all commonly used Windows applications. TRACER has been tested in a laboratory environment. We are seeking partners to deploy and test TRACER in operational environments to help us improve the technology and mitigate any unknown, large-scale deployment challenges.

TRACER can alternatively be implemented within the operating system itself. We are currently exploring such deployment opportunities with operating system vendors.

FLOWER: Network FLOW AnalyzER – Deep Insight Into Network Traffic

Darren Curtis
Darren.Curtis@pnnl.gov

Overview

FLOWER (Network FLOW AnalyzER) is a software application that performs deep IPv4/IPv6 packet header inspection in real-time to collect bi-directional network conversations between computers. It automatically combines unidirectional Internet Protocol (IP) packets into bi-directional network flows. FLOWER can be deployed anywhere in an enterprise using a passive network tap so it cannot be detected.

Customer Need

Enterprise networks, including virtual cloud networks, are under constant attack. Cyber attack tools and hacker communities give cyber-criminals an asymmetric advantage over network and system administrators acting as cyber defenders. This allows the attackers to readily breach networks, create backdoors, and infect systems, leading to costly data loss or theft of intellectual property. Some US government networks have experienced up to 25,000 attempted attacks every day.

There exists an urgent need for a novel detection system that can collect data about network activity throughout the enterprise, providing cyber-defenders insight into traffic patterns, abnormal data flows, and forensic data to study and learn about potential breaches and identify insider threats.

Our Approach

FLOWER utilizes a simple but powerful algorithm based on the IP specifications to parse and aggregate up to 1 million IPv4/IPv6 packet headers per second. FLOWER passively monitors all network traffic. Data can be collected over the same enterprise network being monitored or on a private data collection network.

FLOWER can be deployed on small inexpensive data collection appliances throughout an enterprise and at the perimeter, allowing cyber defenders to selectively target resources to monitor and incrementally deploy appliances to scale to their needs.

FLOWER produces simple CSV-formatted output files that are compatible with existing conventional data analysis tools and emerging cyber analytic research efforts. One example is Trelliscope, a public domain data visualization tool used to analyze data from the IEEE VAST 2013 challenge.[1]

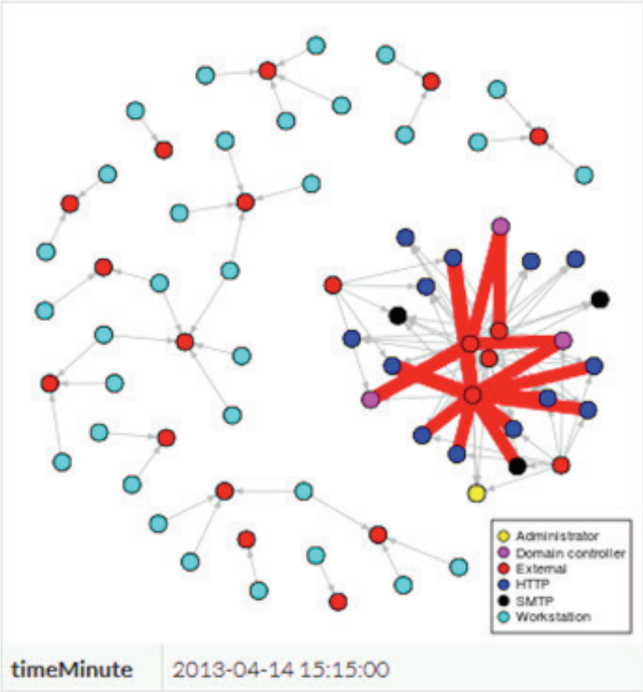

Trelliscope display used to identify a breach using an Administrator account to take control of the Network Domain Controllers and modify web server configurations to redirect users to the adversary's external web server.

Benefits

FLOWER can continue to capture network data even when the traffic exceeds the capacity of network routers and switches. It automatically handles partial or fragmented packets until all fragments have been received. In addition, FLOWER can be configured to specify the maximum number of simultaneous conversations, the number of minutes to wait to mark a conversation

[1] *http://ieeevis.org/year/2013/info/call-participation/vast-challenge*

complete if no packets have been seen, and the maximum number of minutes before forcing a conversation record to be written to a log file.

FLOWER can process packet capture (pcap) files generated by tcpdump, Wireshark, and other pcap generating tools.

FLOWER is a turnkey solution that can scale to meet needs of small to enterprise size networks.

Competitive Advantage

FLOWER has been deployed at over 100 US government sites and many private corporations collecting data of trillions of network flows since 2010.

Network management products from companies such as SolarWinds analyze Cisco NetFlow data to identify traffic patterns and network bottlenecks. Data collected by FLOWER data can be used for those purposes but also provide cyber security analysts more insight in the content of the network traffic to help identify unusual access patterns.

FLOWER provides the ability to recursively parse packet headers to identify the outermost network flow as well as the innermost network flow data encapsulated in a tunnel such as IPv6-in-IPv4, IPv4-in-IPv6, GRE, IPv6, and IPv6-Teredo tunnels.

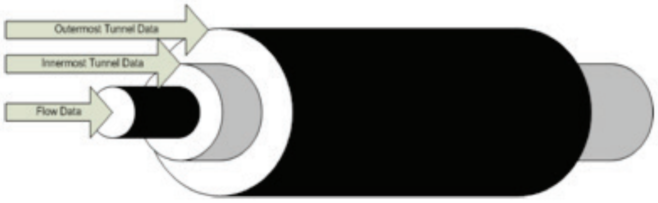

A representation of a tunneled network flow where the outer layer could be IPv4 and the inner layer could be IPv6 with IPv6 TCP data.

Next Steps

FLOWER has been successfully used to detect and mitigate coordinated attacks. Cyber analysts have been able to process data and develop techniques for identifying signatures of potential attacks and modify network configurations to deter future attacks.

We are seeking partners interested in using FLOWER to learn more about their enterprise network traffic and identify potential breaches and insider attackers. We are also seeking partners willing to support a pilot of a customized FLOWER application using micro appliances and a Mobile Adhoc NETwork (MANET) in their enterprise.

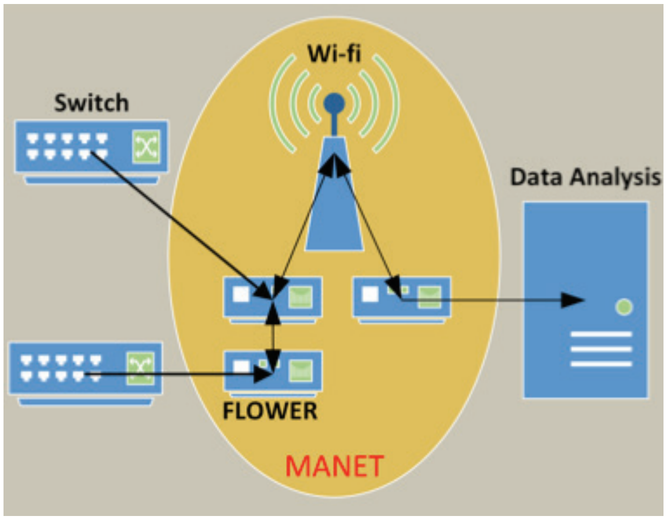

Using FLOWER appliances to collect data using a secured MANET for offline data analysis.

SilentAlarm: Detecting Abnormal Network Traffic

Joel Doehle
joel.doehle@pnnl.gov

Overview

SilentAlarm is an inference-based technology for detecting abnormal network traffic that depends on dynamic network behavior knowledge rather than static signatures and thresholds. It characterizes network behavior as likely malicious and enables the detection of zero-day attacks and polymorphic malware without needing prior knowledge of their specific characteristics.

Customer Need

To enact effective cyber security, organizations need to be able to detect unknown malware targeting unidentified vulnerabilities. Many existing security solutions are signature-based, utilizing a detailed description of a given malware's characteristics to detect its presence on their systems and networks. These technologies require specific foreknowledge of a malware's form or function. As a result, they are incapable of detecting unknown malware or attacks exploiting as-yet undeclared vulnerabilities. This leaves organizations vulnerable to adversaries using tactics such as "zero-day" attacks and polymorphic malware. Better solutions are needed to enable organizations to identify and address these and other types of attacks without the detailed foreknowledge required to develop a signature.

Our Approach

SilentAlarm utilizes dynamic behavioral analysis to detect abnormal network traffic resulting from malware and malicious intrusions and takes action to address the malicious software behind them.

SilentAlarm provides this capability by employing a Bayesian inference model that analyzes and correlates network traffic using dynamic network behavior knowledge in order to construct hypotheses regarding likely malicious activity.

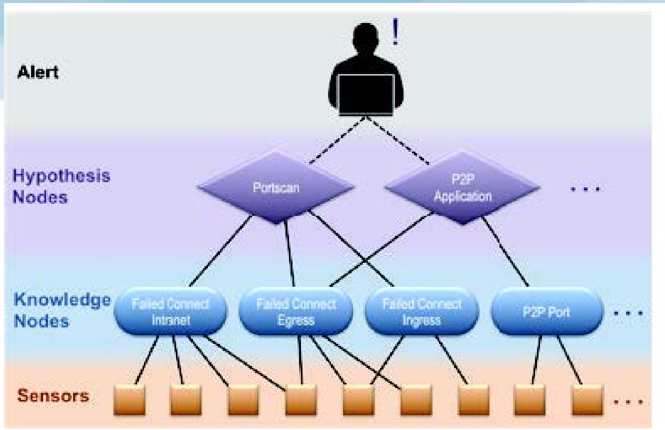

Figure 1. SilentAlarm Inference Network Architecture

Network events are fed into SilentAlarm through various types of sensors (e.g., network anomaly, protocol anomaly) that are already in place on a network. The traffic collected at these sensors is ingested by knowledge nodes that are associated with a particular network behavior (e.g., failed or successful SMTP, failed intranet connection). These knowledge nodes characterize the ingested traffic based on the metrics of prior network behavior.

These characterizations are pushed up to hypothesis nodes that construct hypotheses regarding the likely malicious nature of the observed traffic. Each hypothesis node conducts reasoning about one type of malicious action (e.g., port scanning). In order to do so, it subscribes to the characterizations of one or more knowledge nodes and assigns weighting values to the results of each. Collectively, these characterizations enable the hypothesis node to deduce whether the observed traffic is indicative of a particular malicious action.

When such a determination is made and the associated confidence value is greater than or equal to an alert value, a security action can be performed (e.g., sending an alert to a system administrator or disabling or restricting network access to a particular resource).

Network behavior data is continuously collected through the deployed sensors. This data – along with administrator input and feedback into the knowledge and hypothesis nodes regarding characterization and confidence values – enables SilentAlarm to "learn" and refine its understanding regarding abnormal network behavior.

Benefits

SilentAlarm is able to characterize network traffic as likely malicious based on knowledge of prior network behavior and an inferential understanding of what constitutes abnormal network behavior.

In this way, SilentAlarm is not wholly dependent on static thresholds and signatures, as many traditional malware detection methodologies are. As a result, SilentAlarm can detect previously unknown malware (including polymorphic malware and attacks against zero-day vulnerabilities) based on its behavior in a network.

This technology is device and network agnostic. It can be adapted to integrate with presently deployed sensors on an enterprise environment. It is also highly scalable across various network sizes.

SilentAlarm has been proven effective in an active enterprise environment, serving as an integral component of the security of the PNNL network for several years. Upon initial deployment, SilentAlarm correctly identified 200 machines that were infected with "zero-day" type malware, out of a network of 10,000 computers. In continued operation, SilentAlarm identified zero-day type malware on the network within three minutes of machine compromise.

Competitive Advantage

SilentAlarm is highly adaptable and extensible across varying network environments with myriad sensors. It offers proven behavioral based anomaly detection that provides an enhanced complement to signature-based solutions.

Additionally, we possess a patent[1] on the technology that prevents others from developing the same type of product.

Next Steps

Several years have passed since SilentAlarm's initial development and operational use. The technology needs to be updated to reflect current network behaviors and to develop sensors for today's network event logging technology.

We would like to partner with a sponsor who is interested in renewing this technology and pursuing potential commercial opportunities for it.

[1] *Goranson, Craig A., and John R. Burnette. "Methods and systems for detecting abnormal digital traffic." U.S. Patent 7,908,357, issued March 15, 2011.*

This page is left blank intentionally

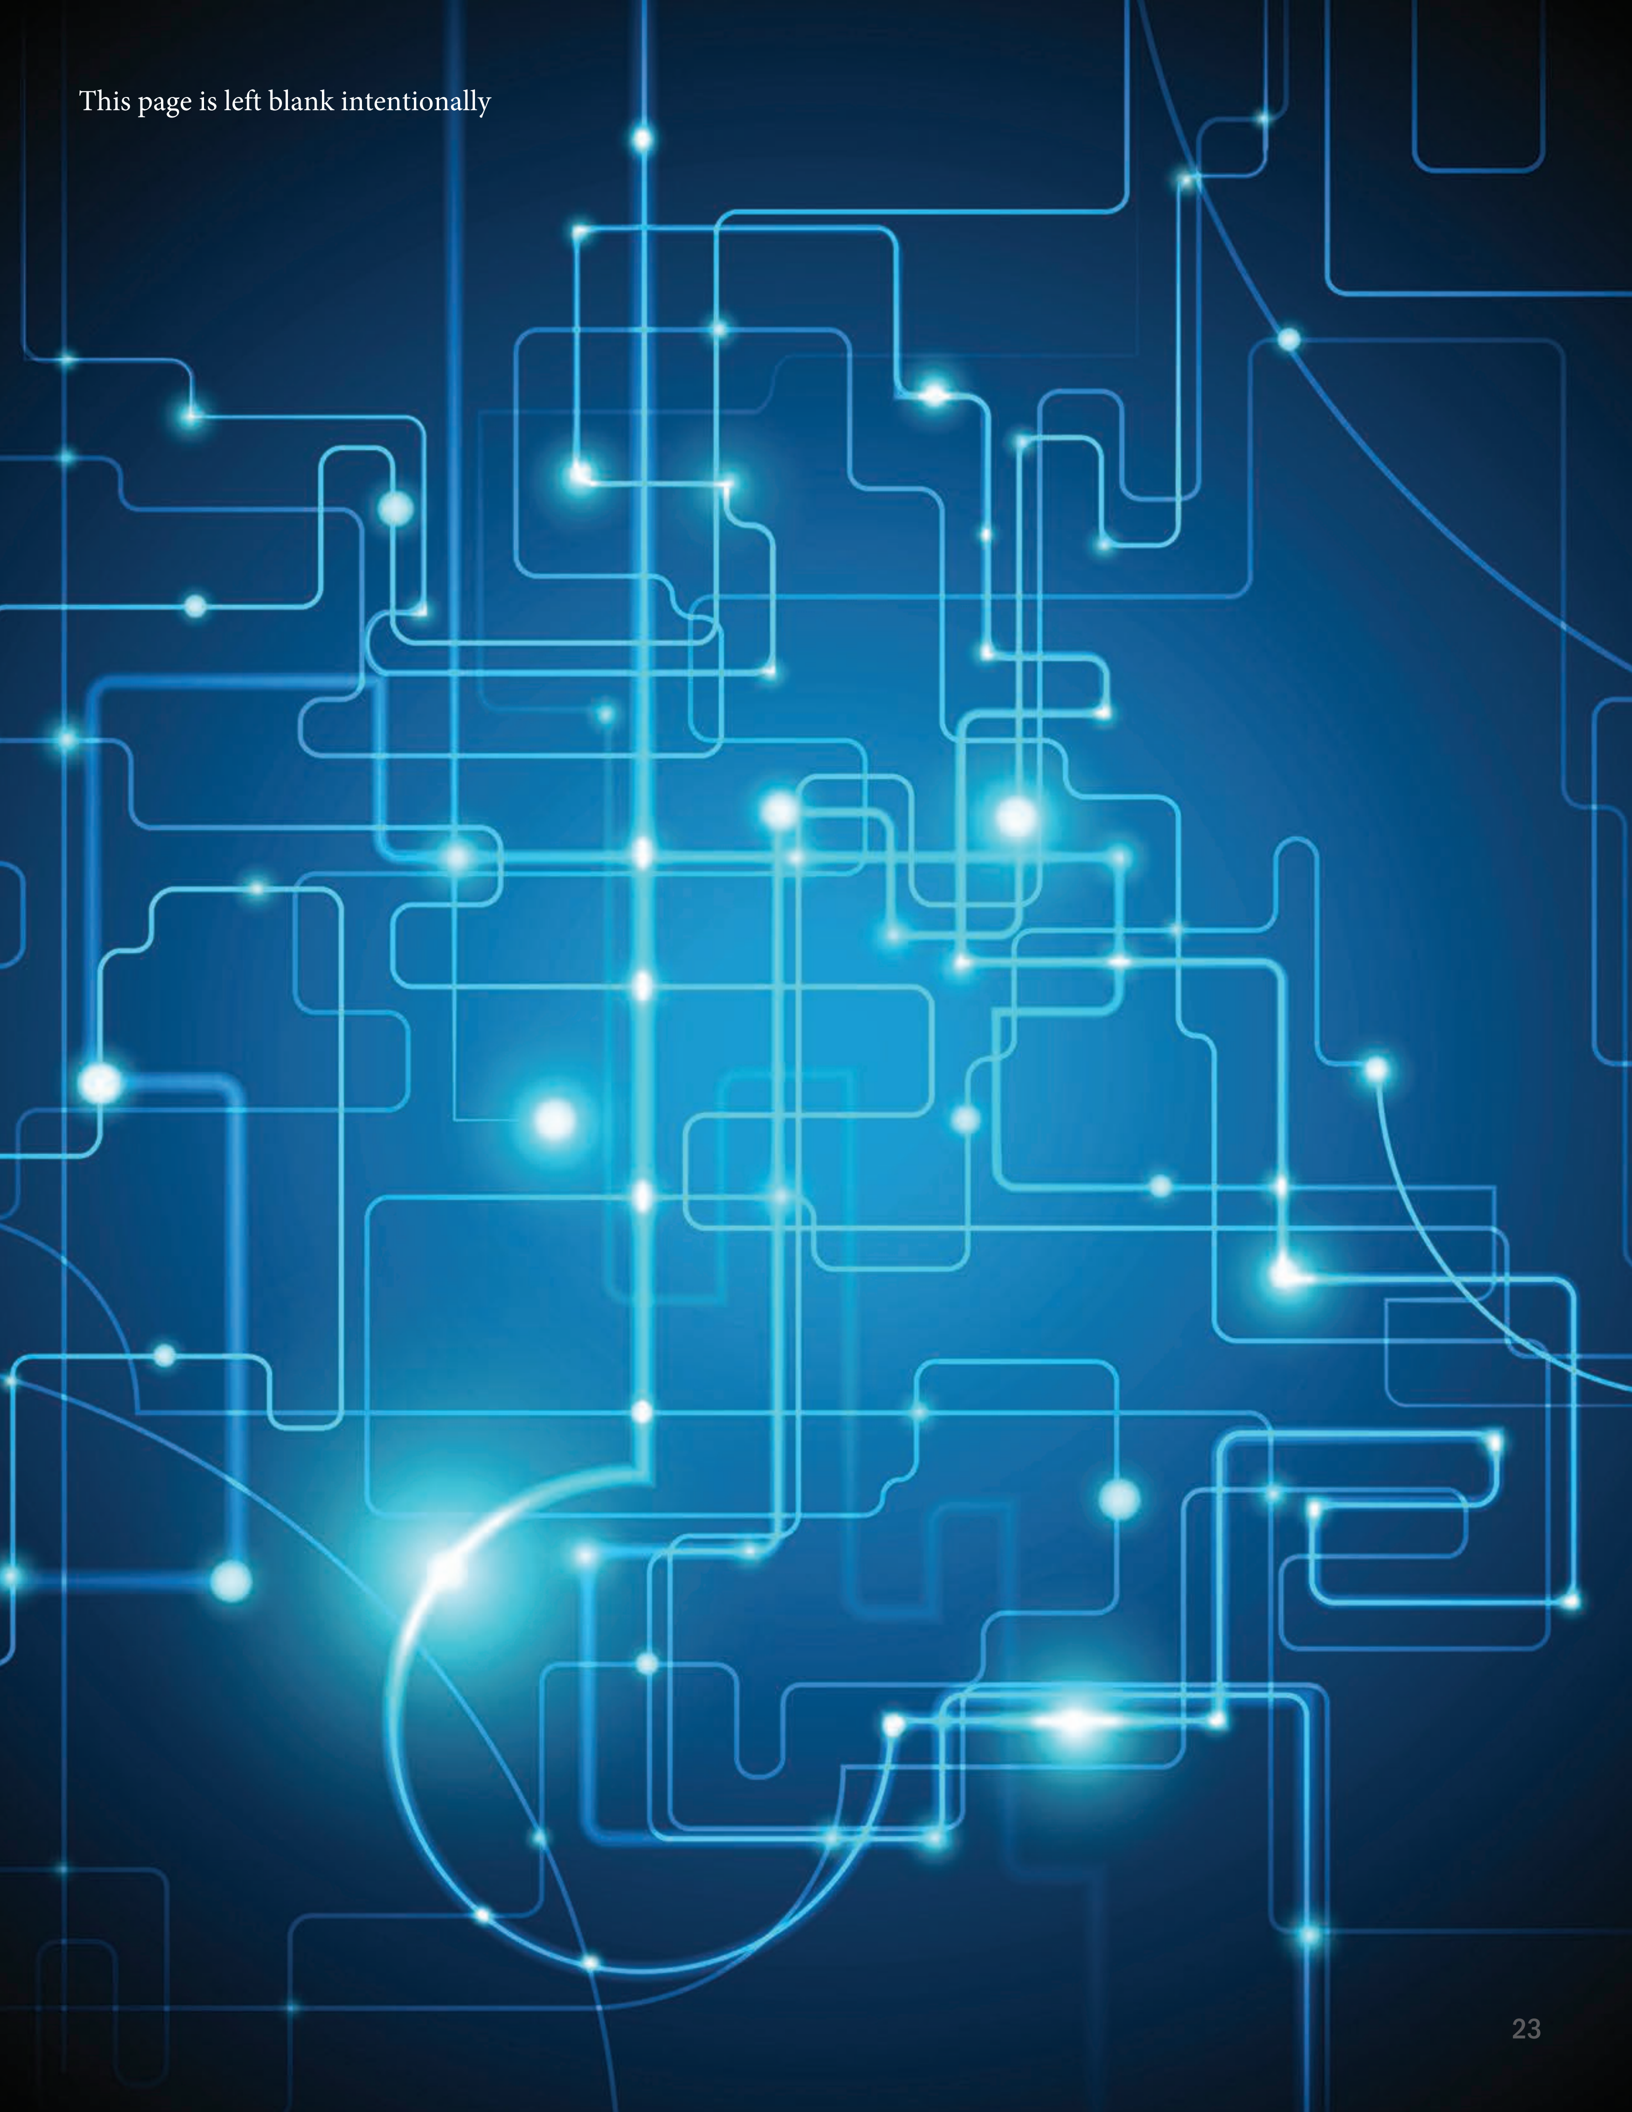

FISCAL YEAR 2015 TECHNOLOGIES:

- Autonomic Intelligent Cyber Sensor (AICS): Cyber Security and Network State Awareness for Ethernet-based Industrial Control Networks
- Situ: Discovering and Explaining Suspicious Behavior
- Scalable Reasoning System (SRS): Threat Landscape Analysis for the Cyber Defender
- Dynamic Defense & Network Randomization
    - Dynamic Defense: Proactively Defending Control Systems Against Emerging Threats
    - Network Randomization: Moving Target Defense for Computer Systems
- SCOT: Turning Cyber Data into Incident Response Threat Intel
- AMICO: Accurate Behavior-Based Detection of Malware Downloads
- ZeroPoint: Advanced Weaponized Document Detection and Analytics

Autonomic Intelligent Cyber Sensor (AICS): Cyber Security and Network State Awareness for Ethernet-based Industrial Control Networks

D. Todd Vollmer
denis.vollmer@inl.gov

Craig Miles
craig.miles@inl.gov

Overview

The Autonomic Intelligent Cyber Sensor (AICS) provides autonomous cybersecurity and state awareness for Ethernet-based industrial control networks. It employs Autonomic Computing techniques and a Service Oriented Architecture to: 1) automatically discover network entity information, 2) automatically deploy deceptive virtual hosts, and 3) automatically identify anomalous network traffic with very high accuracy.

Customer Need

Industrial Control System (ICS) networks facilitate communication among critical infrastructure and form an attack surface that must be secured. Maintaining state awareness and detecting anomalies are notoriously difficult tasks in traditional IT networks due to their inherent complexities, such as the presence of heterogeneous hardware and software, dynamic network composition and usage patterns, and decentralized control. ICS networks can have similar complexities, however the control system traffic tends to be more observable and amenable to predictive modeling.

Ensuring ICS network cybersecurity in the face of these complexities entails both real-time monitoring of network host composition and agile response to changing network conditions. Neither of these capabilities are well met by manual actions alone. A cyber sensor is needed that automatically reacts to changing network compositions and conditions, while simultaneously attaining the highest possible accuracy and lowest false positive rates in detecting anomalous traffic. Such a sensor will obviate much of the human intervention presently required to effectively monitor evolving industrial networks for anomalies.

Our Approach

AICS employs three major analysis components plus standards based communication channels to monitor and protect ICS networks:

Network Identity Identification (NEI): The NEI performs asset discovery by passively monitoring ICS network traffic. For each host discovered on the network, the NEI catalogs its IP and MAC addresses, and attempts to identify its operating system. The NEI continually updates this network model to reflect the present composition of hosts on the network, thereby providing network state awareness.

Dynamic Honeypot (DHP): The DHP utilizes the NEI's constantly evolving network model to automatically configure and deploy deceptive virtual network hosts, otherwise known as honeypots, which imitate the real hosts on the network. These honeypots serve to draw the focus of malicious intent, and thereby provide a decoy attack surface that is easily monitored for anomalous activity.

Intelligent Anomaly Assessment (IAA): The IAA selectively monitors a prescribed list of host network traffic for anomalous activity while adjusting its own sensitivity based on observed global network trends. Statistical features are extracted from the traffic of each network host into feature vectors. A fuzzy logic based anomaly detection algorithm is then used to compute an anomaly score for each vector that expresses the belief that the current window of packets contains anomalies. The anomaly score is compared against the dynamic sensitivity threshold to determine whether to raise an alert.

Communications: AICS captures control traffic by listening on the ICS network switch's SPAN ports. Network host and alert information is delivered externally over the open-standard IF-MAP protocol and syslog. IF-MAP anomaly alerts are raised through a publish/subscribe style messaging system, enabling network stakeholders to selectively receive only those types of alerts which interest them. The AICS communications approach supports flexible deployment options including the ability to deploy multiple sensors with potentially overlapping host monitoring duties.

Autonomic Intelligent Cyber Sensor (AICS): Cyber Security and Network State Awareness for Ethernet-based Industrial Control Networks

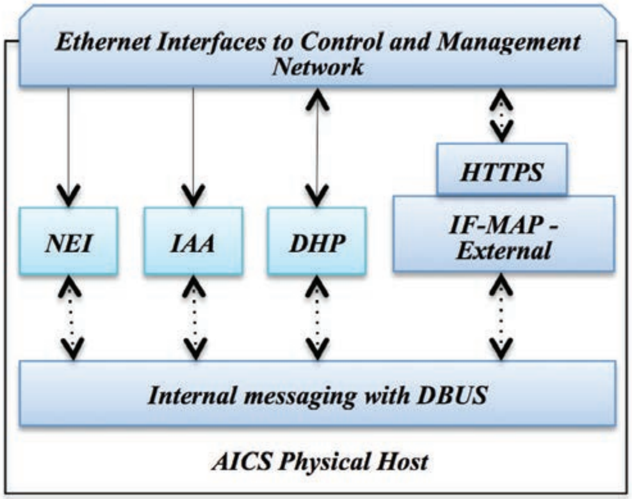
Ethernet Interfaces to Control and Management Network
HTTPS
NEI
IAA
DHP
IF-MAP - External
Internal messaging with DBUS
AICS Physical Host

Figure 1. AICS Architecture

Benefits

AICS employs a modular framework that is deployable on commonly available hardware, and provides for automatic gathering of network host information, automatic deployment of dynamic virtual honeypots, and automatic identification of anomalous network traffic.

AICS reduces the need for human intervention in maintaining network state awareness and anomaly detection. Dynamic honeypots are automatically configured and deployed based on passive network observations, reducing dependence on human network expertise and configuration effort. AICS anomaly detection does not rely on human created rules. Instead it automatically learns normal traffic patterns directly from observation of the network. Additionally the anomaly detection algorithm is designed to minimize false alerts.

AICS has been shown to be effective in its ability to automatically configure itself and detect network anomalies within a controlled laboratory setting. For instance, while anomalous traffic was injected into a test ICS network, AICS was able to correctly label packets to specific hosts as either normal or anomalous with greater than 99.8% accuracy [1].

The modular nature and common communications infrastructure of AICS provides a flexible base for evolving its functionality in the future. This modular nature and common communications interface allows deployment of multiple AICS devices to achieve scalability. Further, AICS delivers alerts and other information via a common interface, which provides for easy integration with products such as system information and event managers or other data correlation solutions.

Competitive Advantage

AICS is an autonomous, intelligent cyber sensor that learns about its environment in order to maximize its own situational awareness and thereby maximize the efficacy with which it detects anomalies. This is in contrast to other state-of-the-art network awareness frameworks that often require intense intervention by skilled humans. Further, the modular design, extensibility, and standards based communication of AICS provides for quick and reliable integration with other systems.

AICS was developed by Idaho National Laboratory (INL), a Federally Funded Research and Development Center (FFRDC) whose mission includes protecting the cybersecurity of critical infrastructure. INL is internationally recognized for its expertise in providing cybersecurity for critical infrastructure, including industrial networks towards which AICS is targeted.

Next Steps

Given the acumen AICS has exhibited in experimental settings, it is ready for phased transition into real ICS networks. Thus, INL is seeking partners for Beta evaluation and commercialization of AICS for broad application to Ethernet-based ICS networks.

[1] Vollmer, T.; Manic, M.; Linda, O., “Autonomic Intelligent Cyber-Sensor to Support Industrial Control Network Awareness,” *Industrial Informatics, IEEE Transactions on* , vol.10, no.2, pp.1647,1658, May 2014.

Situ: Discovering and Explaining Suspicious Behavior

John Goodall
jgoodall@ornl.gov

Joel Reed
reedjw@ornl.gov

Oak Ridge National Laboratory

Overview

Situ is a scalable, real-time platform for discovering and explaining suspicious behavior that current technologies cannot detect.

Customer Need

Despite the best efforts of cybersecurity analysts, networked computing assets are regularly compromised, resulting in the loss of intellectual property, the disclosure of state secrets, and financial damages in the billions. A 2014 report from the Center for Strategic and International Studies estimated the global cost of cyber crime at $400 billion annually. There has also been a rise of sophisticated attack groups that continually develop novel methods of penetrating networks that current technologies are typically unable to detect.

Signature-based security systems are effective at detecting known attacks, but are unable to detect novel or sophisticated attacks. Indeed, automated security systems will never be capable of detecting all malicious activity.

Network operators need tools to help identify suspicious behavior that bypasses automated security systems.

In the deluge of data in today's networks, operators cannot be expected to discover suspicious activity without better tools. Further, operators need to understand what makes an event suspicious to determine the importance and impact of the event. Highlighting such suspicious behavior helps operators focus their limited time on the most suspicious events within vast amounts of data.

Our Approach

Situ combines anomaly detection and data visualization to provide a distributed, streaming platform for discovery and explanation of suspicious behavior to enhance situation awareness.

Our novel approach to anomaly detection is based on unsupervised, probabilistic modeling. Key to our approach is modeling events in different contexts or at multiple scales; each event is modeled and scored by multiple anomaly detectors to identify different kinds of anomalous behavior. For example, a context may group events by day of the week or hour of the day to build a model of temporal behavior for each computing asset.

The anomaly detectors update the behavior models online as new data is streamed into the system. The detectors score each event for each context based on the likelihood of new events occurring given the probability model of prior behavior. Scoring the anomalousness of events for multiple contexts provides analysts with an understanding of *why* an event is anomalous. By examining these contexts, operators can understand how different event features contribute to the overall anomaly score.

The architecture of Situ is designed to scale to very high data rates on commodity hardware—hundreds of thousands of events per second. The system stores data on compute nodes for very fast updates and queries. Scored events are published to a data store for archival review and historical analysis. Scored events are also pushed immediately to a web-based visualization to allow operators to monitor the most suspicious events in real-time.

Benefits

Situ helps network operators discover and understand suspicious events that would otherwise go undetected. It reduces the huge volumes of raw network data to a smaller, manageable number of events that should be examined by human domain experts. By highlighting suspicious activity operators can find novel attacks, but can also be made aware of insider threats, policy violations, misconfigurations, and new kinds of behavior that may require some investigation. Through the application of multiple contexts, Situ can look for a wide range of activity. Different contexts perform better for

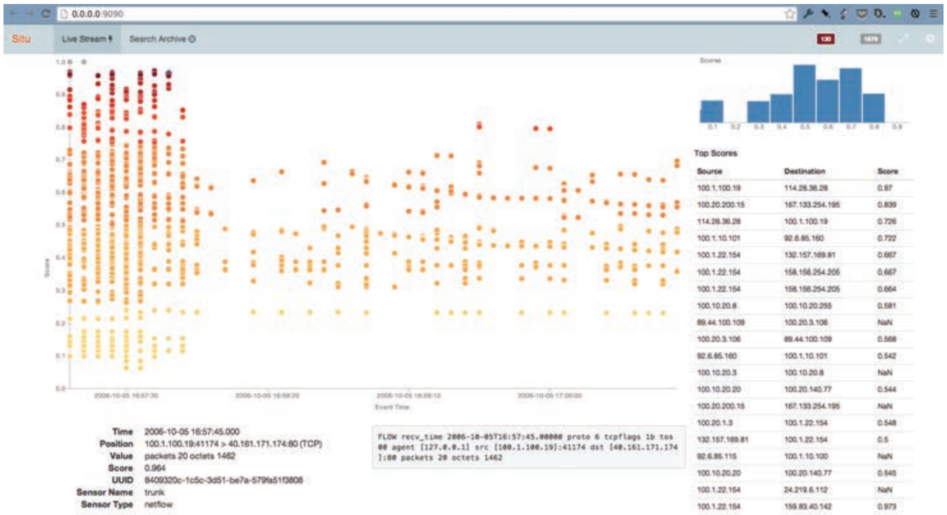

Top Scores

Source	Destination	Score
100.1.100.19	114.28.36.28	0.97
100.20.200.15	167.133.254.195	0.839
114.28.36.28	100.1.100.19	0.726
100.1.10.101	92.6.85.160	0.722
100.1.22.154	132.157.169.81	0.667
100.1.22.154	158.156.254.205	0.667
100.1.22.154	158.156.254.205	0.664
100.10.20.8	100.10.20.255	0.581
89.44.100.109	100.20.3.106	NaN
100.20.3.106	89.44.100.109	0.568
92.6.85.160	100.1.10.101	0.542
100.10.20.3	100.10.20.8	NaN
100.10.20.20	100.20.140.77	0.544
100.20.200.15	167.133.254.195	NaN
100.20.1.3	100.1.22.154	0.548
132.157.169.81	100.1.22.154	0.5
92.6.85.115	100.1.10.100	NaN
100.10.20.20	100.20.140.77	0.545
100.1.22.154	24.219.6.112	NaN
100.1.22.154	159.83.40.142	0.973

Time	2006-10-05 16:57:45.000
Position	100.1.100.19:41174 > 40.161.171.174:80 (TCP)
Value	packets 20 octets 1462
Score	0.964
UUID	8409320c-1c5c-3d51-be7a-579fa51f3808
Sensor Name	trunk
Sensor Type	netflow

```
FLOW recv_time 2006-10-05T16:57:45.00000 proto 6 tcpflags 1b tos
00 agent [127.0.0.1] src [100.1.100.19]:41174 dst [40.161.171.174
]:80 packets 20 octets 1462
```

Figure 1. Situ's streaming user interface shows the most critical events

different kinds of attacks. Multiple contexts can also help explain why an event is suspicious since the varying scores will point operators at certain kinds of behaviors.

Situ is generally applicable to other domains, such as intelligence analysis and cyber-physical infrastructure protection, that require real-time behavioral monitoring.

Competitive Advantage

Situ's probabilistic approach to anomaly detection has several advantages over other methods. Signature-based discovery systems can only identify *known patterns* of malicious behavior. Situ complements such systems by highlighting suspicious behavior that existing systems cannot detect.

Machine learning offers a more robust approach, but typically requires labeled training data, which is rarely available and usually out of date. Situ requires no labeled training data, making it easier to deploy in operational environments. Machine learning approaches typically train periodically offline. Situ trains online so that data models are always up to date.

Other approaches to anomaly detection in cybersecurity commonly help identify atypical events or time windows where an anomaly occurred. Situ goes further and helps operators understand *why* something is anomalous through the scoring and reporting on multiple contexts.

Many approaches to anomaly detection and attack discovery operate in batch mode (e.g. map-reduce jobs in a Hadoop store), which ignores the reality of the speed of cyber attacks. By the time detection takes place, the attacker may have come and gone. Situ operates on real-time streaming data, minimizing the time from the observation of an event by a sensor to the reporting of the event to the operator.

Finally, other approaches to attack and anomaly detection typically have large numbers of false positives, which leads operators to mistrust or ignore alerts. Situ has an adjustable false positive rate that allows an operator to define the acceptable percentage of false positives to set the threshold for discriminating anomalous from normal behavior.

Our visualization approach is unique in that it focuses on streaming data, reducing the time it takes to be notified of important events.

Next Steps

We are currently improving the user interface by creating multiple visualizations that allow analysts to seamlessly move back and forth between a view of the streaming data and a visual query interface to search through archival data.

We are looking for potential pilot and test sites, as well as commercialization and transition partners to put Situ into the hands of the operators who need it.

SRS: Threat Landscape Analysis for the Cyber Defender

Scott Dowson
scott.dowson@pnnl.gov

Oriana Love
oriana.love@pnnl.gov

Rick Riensche
rmr@pnnl.gov

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

Overview

Cyber defenders need to stay abreast of patterns and emerging trends in the threat landscape to effectively protect their networks. The Scalable Reasoning System (SRS) is a solution that automates data collection from various sources, analyzes the data to identify trends and hot topics, and provides a visual interface to explore the information.

Customer Need

To effectively prepare for and counter cyber threats, cyber defenders must actively survey many sources of information. Only by monitoring a broad spectrum of information resources (social media, threat reports, open source media, etc.) can the full threat landscape be pieced together. Manually discovering, harvesting, and reading data from these sources is time consuming. Tracking emerging trends against historic patterns or correlating reports across multiple sources is a taxing process that carries the risk of missing critical pieces of information. Cyber defenders need a single, consistent, and reliable collection and analysis strategy for information—a system that automatically extracts topics, themes, and trends in the data and visually presents the relevant and emerging threats.

Our Approach

SRS is a flexible framework that encompasses the 1) harvesting, processing, and management of data; 2) analytics to extract, correlate and summarize; and 3) interactive visualizations to explore and interpret the information. SRS was designed from the beginning as an extensible component-based system, so end users are empowered to customize the application to suit their needs.

Drawing from a library of data harvesting components, the system monitors and automatically retrieves data from sites using a variety of data exchange technologies. This retrieval includes pulling data from file systems, data warehouses, and web interfaces. The system can be easily adapted to new data sources as they emerge using the published software development kit.

Analytic components process, extract, and correlate categorical and topical features from the data. For unstructured text, keywords are automatically extracted, correlated, and visualized over time. This capability is used to both identify long trending patterns and detect emerging new patterns. For structured data—including temporal data— distributions and facets are calculated, providing the means to filter and pivot within the data collection.

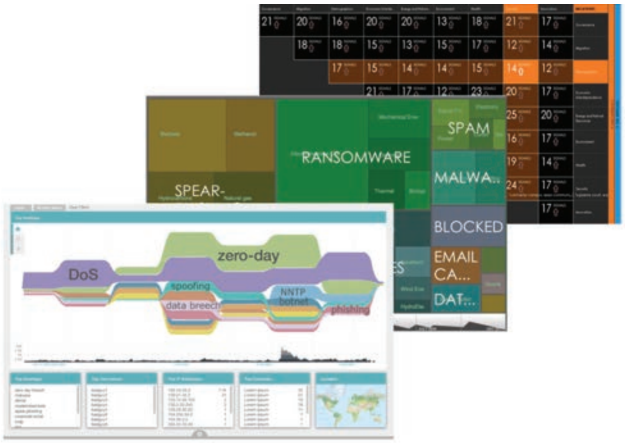

Figure 1. Interactive visualizations of thematic trends, ontologies, and alerts detected.

As requirements change or as new algorithms emerge, the extensible plug-and-play nature of the SRS framework allows new components to be developed and integrated to expand the collection and analysis capabilities of the system, keeping the system current and relevant.

SRS is designed to provide data and analytic products through web services and to present the information in an interactive web-based interface. This feature allows the defender to explore and interact with the data using a variety of visual widgets. Users can visually explore the breadth of information; monitor the reported trends; or drill in to focus on newly discovered information.

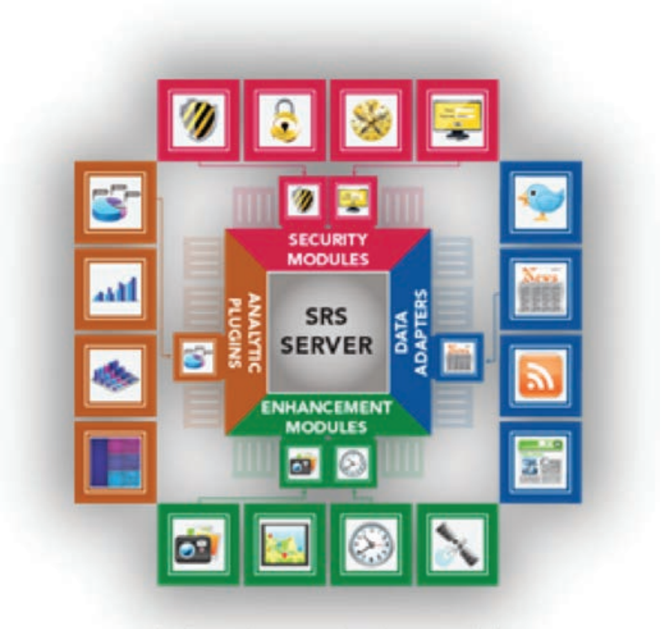

Figure 2. Modular plug-and-play architecture

Benefits

SRS provides situational awareness and alerting for both emerging threats and countermeasures as reported by the selected sources. The system is data agnostic and can be adapted to ingest most data sources. These sources are continuously monitored through a web harvesting engine that performs the arduous task of parsing and processing the data for their salient features. This frees up more time for the cyber defender to analyze the data through the interactive dashboard, which provides the visual means to explore and identify patterns in the data. The dashboard is platform independent and can be customized to meet enterprise and user needs.

Competitive Advantage

Existing tools, such as news aggregators, are useful to cast a wide net and collect information from a specific set of sources. However, these still require the user to manually read and assimilate all the data. SRS can automatically and continually ingest data from a customized set of data sources and extract the data's key features, which are then provided to the user via an interactive visualization. When appropriate, predefined analytics can be applied and presented to draw user attention to particular features.

Other services provide threat intelligence products based on meta-analysis of cyber threat data. Although these products are a very rich source of data, the threat landscape can be further broadened by incorporating other data sources. By combining threat intelligence with other data sources, SRS provides the means for cyber defenders to visually explore, discover, and monitor the full, dynamic landscape.

Next Steps

We are seeking partners interested in participating in a user study to help us learn and understand their specific needs and use cases, and in supporting a pilot of a customized SRS application in their enterprise.

Dynamic Defense: Proactively Defending Control Systems against Emerging Threats

Adrian Chavez
adrchav@sandia.gov

Jason Hamlet
jrhamle@sandia.gov

Overview

Sandia National Laboratories (SNL)* is investigating and developing dynamic defense techniques to better secure critical systems operating within the energy sector. Currently, it is extremely difficult to detect threats within control system networks until it is too late. Using dynamic defense techniques, SNL has developed a set of machine learning algorithms to detect system patterns that deviate from normal operation and respond in an appropriate manner depending on the scenario. Detection coupled with a set of appropriately chosen responses to mitigate malicious traffic patterns, our "chess master" engine in the diagram below, provides situational awareness to an operator and uncertainty to an adversary. We developed these security enhancements while meeting the unique time-critical constraints faced by control systems.

Customer Need

The ability to quickly recognize and appropriately respond to threats is critical to control system security. One can see from ICS-CERT alerts and advisories that critical infrastructure systems continue to be an active target for adversaries. ICS-CERT is reporting over a 20% increase in incidents from 2014-2015 alone**. Each incident is a security threat to these high-consequence, high-availability systems and deserves an appropriate response strategy that can be activated quickly.

Our Approach

SNL developed dynamic defense algorithms to detect and trigger responses that mitigate attacks on a host system. The algorithms apply an ensemble of machine learning algorithms to detect traffic that deviates from a trained baseline or resembles previously observed attacks. Once detected, a response to mitigate the specific threat is triggered or an alert is generated for operator intervention. A unique set of machine learning algorithms are employed within each host and the specifics of those sets periodically and randomly change, presenting a dynamic, difficult to predict defense posture to the adversary. Our solution works in both Windows and Linux operating systems.

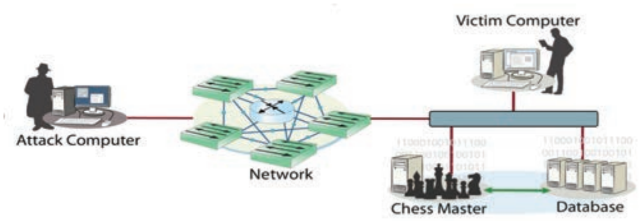

Benefits

Dynamically defending systems against threats launches appropriately chosen mitigations to counter attacks quickly. Our modular implementation provides a framework to integrate new protective measures that counter past, present and future threats. New responses can easily be integrated to mitigate new threats, which is essential for maintaining high availability systems.

Competitive Advantage

Our solution has yielded higher accuracy rates and lower false-positive rates than those in published literature when compared against the same datasets. Our accuracy rates continue to improve as we refine our algorithms and train on more datasets.

Next Steps

We are currently developing our dynamic defense framework to allow for additional response modules to easily be integrated into our existing solution. We seek pilot partners to validate our detection algorithms within a laboratory environment and to transition our technology into industry.

*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.
**www.securityweek.com/critical-infrastructure-incidents-increased-2015-ics-cert

Network Randomization: Moving Target Defense for Computer Systems

Adrian Chavez
adrchav@sandia.gov

William M.S. Stout
wmstout@sandia.gov

Overview

Computer systems continue to use predictable communication paths, static configurations, and unpatched software, all of which benefit an adversary. Sandia National Laboratories (SNL) has developed a prototype implementation of a moving target defense solution that efficiently randomizes IP addresses, application port numbers, and network communication paths while maintaining network connectivity, functionality, and performance. Introducing randomness, uncertainty, and unpredictability thwart attacks and shift the advantage back to the defender. Applying these protective measures converts computer systems into moving targets, adding an additional layer of defense in the early stages of an attack.

Customer Need

The first step an adversary takes is to gain reconnaissance information about a system of interest Cyber security incidents have risen dramatically, from just 5,503 in 2006 to 67,168 in 20141. In 2014 ICS-CERT found that 55% of incidents were APT related[1]; furthermore, 53.22% of all ICS-related incidents were network scanning/probing[2]. Many of these incidents are enabled by the broad availability of system information that is openly available to anyone upon request or observation. Randomization of such information is a promising solution that can protect a system against these early stages of an attack.

Our Approach

SNL's network randomization solution can be retrofitted into existing computer systems in a scalable and transparent manner. Software Defined Networking (SDN) technology allows our randomization schemes to be inserted directly into the network layer, so our solution is transparent to the end devices and scalable. We depend on an SDN controller within the network to manage the randomization of network configurations. Each of the SDN switches is responsible for communicating with the controller to learn the random IP address, port number, and path assignments for traffic traversing the network.

Benefits

Our solution can be rapidly introduced into an existing network using OpenFlow capable hardware switches. If adding new hardware is infeasible, software-based switches, such as Open vSwitch, can be used. The randomness of network configurations provides an environment that is continuously changing and difficult for an adversary to target.

Competitive Advantage

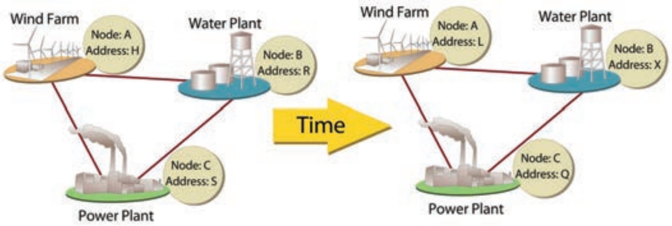

Moving target defense strategies often involve introducing agent software onto each node in the network to randomize network configurations. This approach is effective in small environments but does not scale to large networks such as critical infrastructure networks. We are taking the next step to put research to practice and have developed a prototype that is scalable, efficient, and effective in defending against adversaries in the early stages of an attack. Latency introduced is minimal (<20ms in our test environment) and continues to improve as our development progresses.

Next Steps

We seek pilot partners to deploy our randomization algorithms at a larger scale than our test environment (300 nodes). We ultimately seek to transition our technology into industry and integrate our solution with other management systems.

References in footer section of tech sheet:

1. http://www.gao.gov/assets/680/671253.pdf

2. https://ics-cert.us-cert.gov/sites/default/files/Annual_Reports/Year_in_Review_FY2014_Final.pdf

SCOT: Turning Cyber Data into Incident Response Threat Intel

Todd Bruner
tbruner@sandia.gov

Overview

The Sandia Cyber Omni Tracker (SCOT) is a cybersecurity incident response management system and knowledge base. Designed by cybersecurity incident responders, SCOT provides a new approach to manage security alerts, analyze data for deeper patterns, coordinate team efforts, and capture team knowledge. SCOT integrates with existing security applications to provide a consistent, easy to use interface that enhances analyst effectiveness.

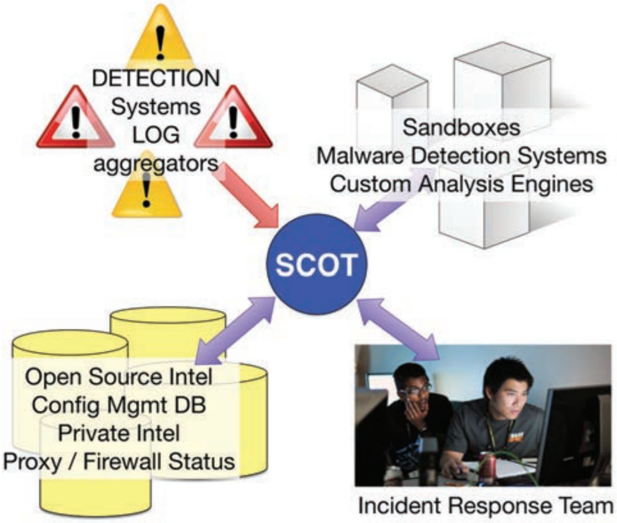

Customer Need

Incident response (IR) teams utilize many systems to detect, collect and analyze cybersecurity event data. These systems, while solving pieces of the puzzle, often fail to give the analyst a holistic view of what is happening and their team's response to those events. Many systems do not have the flexibility to work with the IR processes to research and document those activities. Research is not easily shared and searchable, so the team's effectiveness decreases, especially when key personnel are on vacation or take other positions. Without a ready corpus of examples of past events, training new team members becomes a lengthy process. Each additional tool adds cognitive load to the analyst and the tool's maintenance needs take the analyst away from the primary task of IR.

Our Approach

Focused on removing the friction between analysts and their tools, SCOT enables analysts to document and share their research and response efforts. As a software suite that integrates data from detectors, analysis, and other information sources, it provides real time updates of the team's work to keep the team informed and coordinated. SCOT automatically identifies indicators to help the analyst discover and respond to advanced threats. Centralization of the data reduces the contextual shifts necessary to access each detection system. Fusing detection data with the accumulated team knowledge allows the team to quickly discover that a new alert might be part of a larger campaign. In addition, SCOT automates and simplifies common analyst tasks to increase analyst's effectiveness by freeing them to concentrate on cybersecurity – not tool mastery.

Benefits

The number of alerts Sandia's IR team has seen has nearly doubled in the past several years. SCOT enabled the team to keep up with this increase without adding additional team members. As a training tool, new team members started contributing in weeks, instead of months. In just over 4 years SCOT has amassed a database of over 700K indicators from analyst and alert input. These indicators help the team spot an adversary's methods and tactics, as well as highlighting common

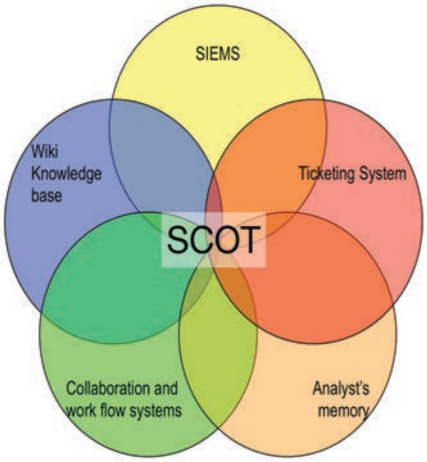

targets within the enterprise. SCOT processed over 1.6 million alerts since deployment, while maintaining 99.9% availability, and required minimal administration. SCOT is fully scalable to meet higher loads.

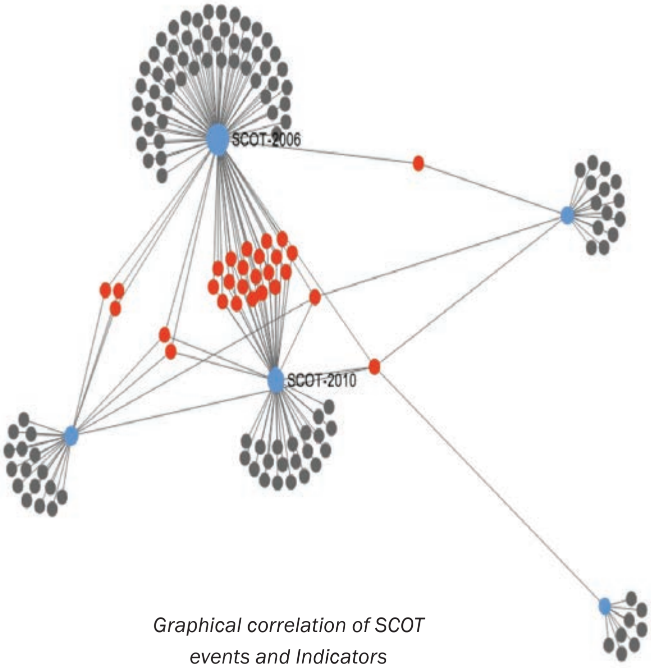

Graphical correlation of SCOT events and Indicators

Competitive Advantage

Sandia's incident response team realized several advantages using SCOT over other solutions. SCOT's ease of use eliminated the steep learning curve of traditional SIEMS and captured team knowledge much more effectively. Designed for cybersecurity, SCOT allows the IR team to enter data easily, instead of struggling to conform to a ticketing system designed for other purposes. While workflow systems handle linear workflows easily, SCOT is purpose built for the looping nature of cybersecurity investigations. SCOT also solves the challenges of keeping wikis, spreadsheets and documents up-to-date and accessible to an IR team. While top-notch analysts may be able to keep everything in their brains, SCOT will capture their knowledge for when they go on vacation or to other employment.

Combining the best attributes of these solutions, SCOT has been enthusiastically adopted by Sandia's IR team as an indispensible tool that enhances the productivity of the team and helps us keep an edge on our adversaries.

Next Steps

Start building your organizational memory and turn security data into intel your IR team can use. Please go to http://getscot.sandia.gov for more information on licensing and how to obtain SCOT. Sandia is actively developing SCOT and looking for ideas and contributors. We seek opportunities for collaboration and custom development. Please contact tbruner@sandia.gov for additional information.

AMICO: Accurate Behavior-Based Detection of Malware Downloads

The University of Georgia

Roberto Perdisci
perdisci@cs.uga.edu

Kang Li
kangli@cs.uga.edu

Overview

AMICO is a novel open source software system for accurate behavior-based detection of malware downloads in live web traffic. Once deployed at the edge of a network, AMICO automatically learns how to distinguish between malware and benign software downloads by observing the download behavior of the network users themselves. After the initial learning phase, AMICO is able to automatically detect new (including zero-day) malware downloads in the monitored web traffic, and can alert network security personnel with detailed incident reports about the detected events.

Customer Need

Sensitive computer networks are under constant attack. Cyber criminals can gain almost unrestricted access to a network by leveraging malicious websites to force users to download and run malicious software. This allows the attackers to implant malware into the network, and to create a backdoor that can lead to costly data breaches and loss of intellectual property.

Most networks rely on traditional antivirus software to protect themselves from malware downloads. Unfortunately, security researchers have repeatedly demonstrated that anti-virus defenses are only partially effective and may miss more than 65% of the latest malware threats.

Other existing malware download defenses make extensive use of URL blacklists, to prevent users from accessing known malware distribution sites. However, by nature these blacklists lag behind the new threats and fail to detect a significant number of new malware.

Therefore, there exists an urgent need for novel malware download detection systems that can better protect a network's perimeter by accurately detecting new, never-before-seen malware files and the related malware distribution sites.

Our Approach

AMICO's behavior-based approach to detecting malware downloads is based on the following main intuition: to evade existing defenses, malware distribution operations must be *agile*.

For example, to avoid antivirus detection, malware developers make heavy use of code obfuscation and polymorphism to frequently change their malware files. On the other hand, benign executable files usually change only when a new version is released.

Furthermore, to evade URL blacklists, malicious websites that distribute malware need to frequently relocate, for example by changing their domain name and IP addresses. On the other hand, benign executable files are typically hosted at professionally operated service providers with a stable domain name and network infrastructure.

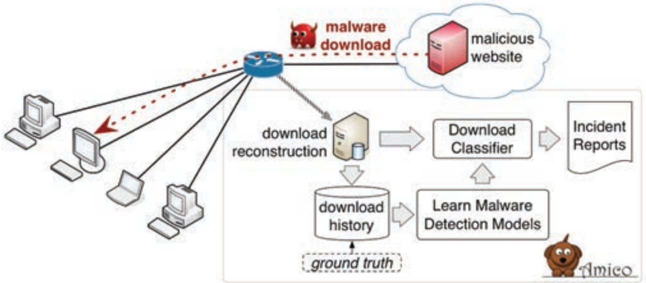

To leverage these intuitions, AMICO combines advanced network traffic monitoring with artificial intelligence and data mining methods.

AMICO passively monitors all web traffic at the edge of a network. Every time a user downloads an executable file, the system performs an on-the-fly reconstruction of the download from the network traffic, and stores the file into a download history database, along with provenance information regarding *who* (i.e., what machines) downloaded the file and *where* (i.e., what website) the download came from.

During an initial training period, some of these download events are first labeled as either *benign* or *malware*, using the partial ground truth provided by existing antivirus tools. Given these labeled events and statistics about the download behavior of the network users collected during the training phase, AMICO automatically learns a web traffic model that can be used to accurately classify future malicious file downloads based simply on their provenance characteristics.

Benefits

AMICO is able to efficiently reconstruct and accurately classify new malware file downloads by passively monitoring web traffic from the network edge. It explicitly leverages the fact that modern malware distribution operations are highly agile, and turns the attackers' strategy into an advantage for the defenders.

AMICO automatically learns how to distinguish between malware and benign software downloads by observing the download behavior of the network users, providing a defense that can self-adapt to the deployment network and further improve detection accuracy.

Competitive Advantage

AMICO provides a fully open source and easy to deploy solution for detecting malware downloads in live web traffic.

AMICO's download classifier does not rely on signatures, and therefore is not affected by malware code polymorphism and obfuscation. Instead, AMICO leverages malware polymorphism as a feature to enable a more accurate detection of malware download events. Furthermore, AMICO does not rely on URL or domain name blacklisting, and does not need to run malware files in a sandboxed environment.

Unlike existing defenses, AMICO is able to detect never-before-seen malware download events by leveraging their provenance characteristics, and by automatically learning from the download behavior of the network users themselves. Therefore, AMICO provides an effective complement to current antivirus and malware defense solutions.

Next Steps

AMICO has been tested via pilot deployment in a large academic network serving tens of thousands of users, where it was able to detect more than 95% of all new malware file downloads and about 80% of malware files missed by existing defenses.

Pilot testing in other operational environments would provide an important opportunity to improve performance, usability, and to compare AMICO to other existing defense solutions. In addition, we are seeking partners and sponsors who are interested in fostering the widespread adoption of AMICO.

Funded through the National Science Foundation's (NSF) Division of Advanced Cyberinfrastructure (ACI), Cybersecurity Program

ZeroPoint: Advanced Weaponized Document Detection and Analytics

THE UNIVERSITY *of* NORTH CAROLINA *at* CHAPEL HILL

Kevin Z. Snow
kzsnow@cs.unc.edu

Fabian Monrose
fabian@cs.unc.edu

Overview

The ZeroPoint Platform provides highly effective, high-throughput, next-generation detection and diagnostics of exploit payloads embedded in documents distributed via email and the web, content used in so-called drive-by downloads and attacks on network servers.

Customer Need

Today, the widespread proliferation of document-based exploits distributed via massive email and web-based attack campaigns is an all too familiar strategy. Attackers use this tactic to kickoff full-scale data breaches by weaponizing documents and web content to gain total access to the recipient's computer. In 2012 these data breaches cost an average of $5.5 million per incident, a figure on the rise as organizations increase their online presence and threats become more sophisticated. In August 2014, for example, several large financial institutions lost gigabytes of data to cyber criminals targeting the financial sector. Sadly, contemporary defenses have failed to keep pace with the relentless onslaught of evasive techniques that are readily available from off-the-shelf attack toolkits. In light of this ever-present threat, there is a need to empower organizations to allow end-users to safely use email and browse the web.

Our Approach

The ZeroPoint Platform is a network appliance that analyzes documents, email, web content, and server interactions collected from network border traffic and operator-submitted content. Potentially hazardous documents or web content are launched or replayed in their target application to dynamically unpack embedded exploit payloads, and then application memory is inspected to discover those payloads. The key to the ZeroPoint approach is a patented *"execution of data"* technology that uses an advanced micro-OS built into the analysis engine to enable fast, accurate inspection of data or memory to identify exploit payloads. This core technology takes advantage of hardware virtualization to inspect all data by directly *executing* it to discover what lurks within, without relying on any form of software emulation. There is no need to guess whether a resource is malicious based on trivially obfuscated file content, post-infection behavior that can easily be disguised, or out-of-date signatures. ZeroPoint hones in on the small portion of an attack the adversary cannot omit or quickly adapt – the exploit payload – by leveraging the fact that exploits operate under practical constraints that bound their operations in ways that make them detectable.

Benefits

The ZeroPoint Platform enables users to safely use email and browse the web with the confidence that attacks are promptly discovered at the first stage, before data is lost. The platform transparently provides complete network-wide protection with no downtime to deploy, inspects each document or web page in about one second, and produces virtually no false alarms. Our core technology has already been validated on the University of North Carolina Chapel Hill campus (29,000 students with 5 Gbps average load) and large-scale empirical analysis spanning 10,000 weaponized documents. ZeroPoint's diagnostic functionality also enables operators to preemptively block connections to malicious domains found in the inspected content.

Competitive Advantage

Contemporary approaches for detecting attacks have relied on antivirus *signatures* of previously observed attacks. Unfortunately, the delay between the first use of an attack and the deployment of its signature is too often measured in weeks and months. Meanwhile, the attackers continually compromise users. Moreover, signatures are widely known to produce many false alarms. A myriad of recent solutions and products–most based on sandboxing technology–claim to avoid the pitfalls of signatures and protect against zero-day attacks. However, these containment mitigations are complex and costly to deploy and manage on endpoints. Instead, ZeroPoint provides transparent and network-centric detection. Other detection solutions leverage sandboxes for

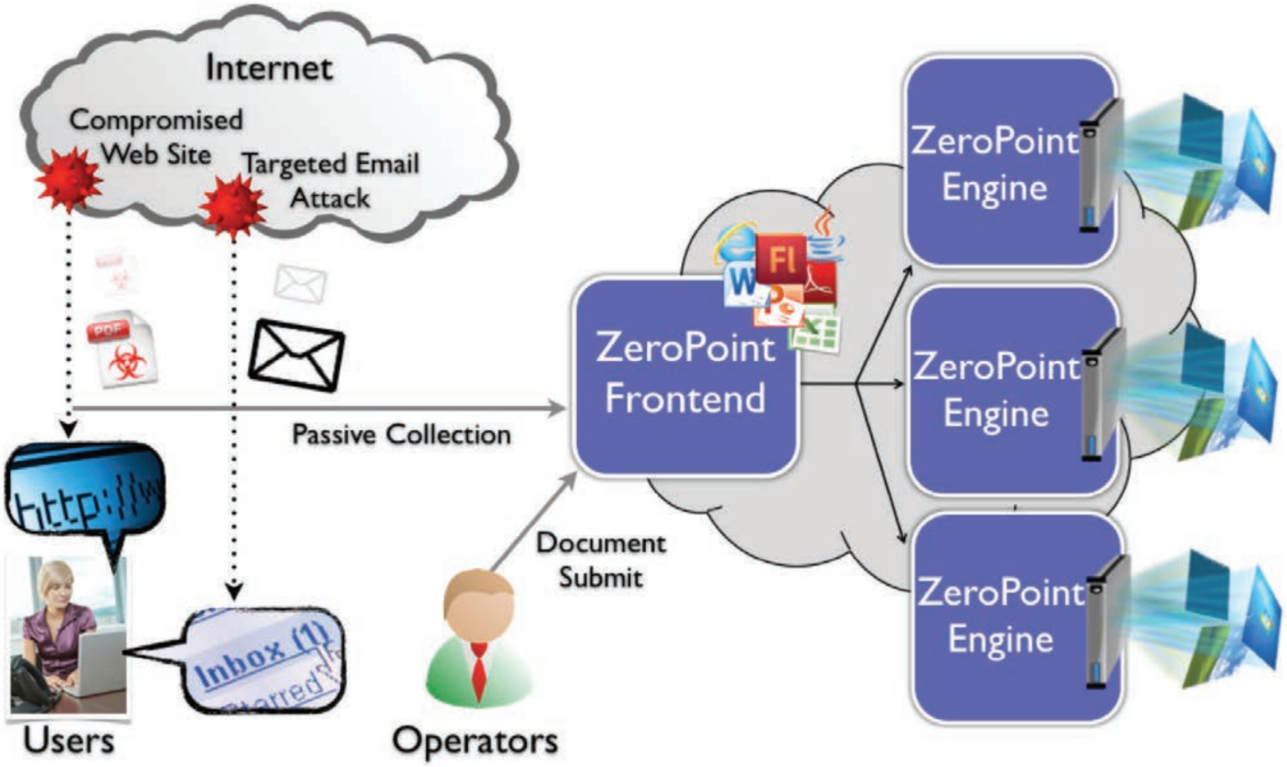

Figure 1 - The deployment scales-up with multiple analysis engines in a cloud-ready model, or rolls all components into one rack-mounted server in stand-alone deployments capable of tens of thousands of inspections a day.

behavioral analysis. Unfortunately, that behavior is easily camouflaged with benign activity, only revealing itself after an extended period of time, or is limited to the analysis of executable files. Rather than fruitlessly attempting to keep up with the fast pace of new attack signatures and easily disguised behaviors, our technology turns the tide by moving away from this status quo and avoiding signatures and observable post-infection behavior altogether. Our underlying technology has not required any signature, behavior, or heuristics updates over several years, and yet we continue to find weaponized documents where other solutions fail – a testament to the solid foundation on which ZeroPoint is built. In short, ZeroPoint takes a unique approach that is faster, more accurate, and more informative than other solutions.

Next Steps

Two U.S. Patents that protect the core technology are pending.   We seek commercialization of our technology through partnering or licensing with a major vendor of network security products. We also seek pilot deployments with large organizations for our stand-alone or cloud-ready prototypes.

Funded through the National Science Foundation's (NSF) Division of Advanced Cyberinfrastructure (ACI), Cybersecurity Program

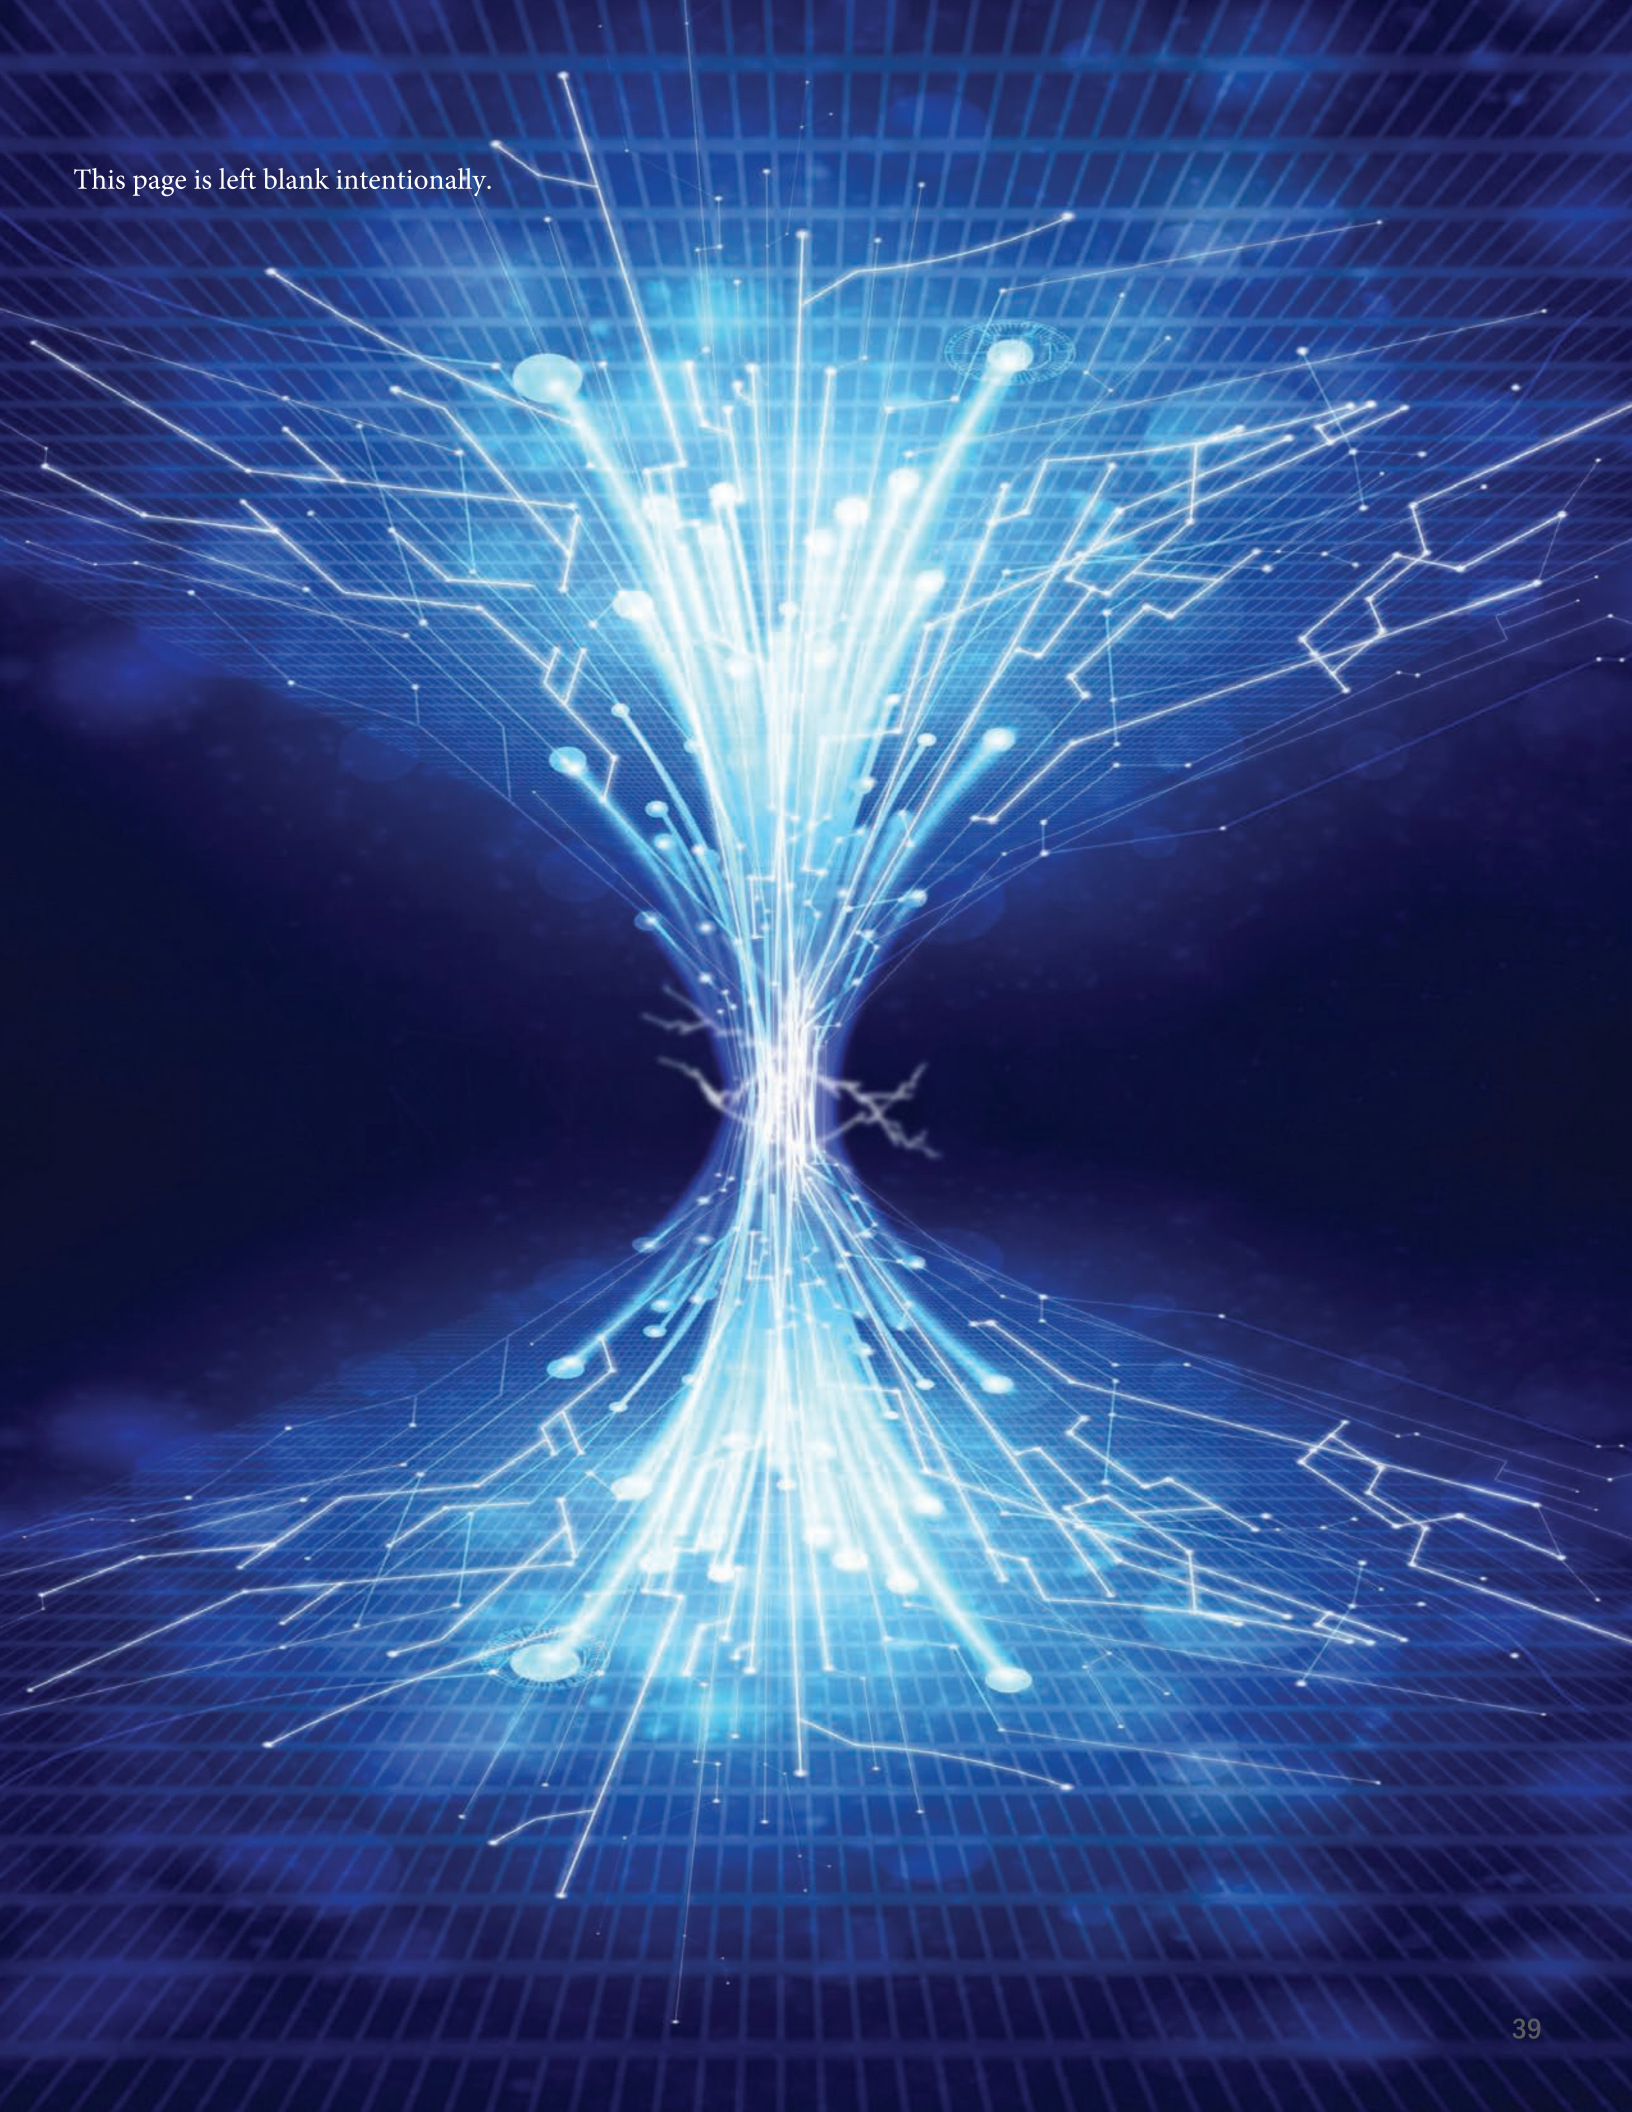

This page is left blank intentionally.

FISCAL YEAR 2014 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
- SENSECL: 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

Matthew Breiner
Matthew.Breiner@jhuapl.edu

Shaku Harshavardhana
Shaku.Harshavardhana@jhuapl.edu

JOHNS HOPKINS
APPLIED PHYSICS LABORATORY

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 very high malware variant detection accuracy when measured against 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; today the economics of information security 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 does not support the kind of malware variant identification needed to achieve economies of scale with information sharing between communities of interest.

Our Approach

CodeDNA provides a reliable, fully automated, fast means for identifying related malware binaries and linking variants. With automated creation of 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 to the entries in a stored fingerprint database utilizing a fast, scalable matching process that lends itself to low-cost open-source cloud processing. CodeDNA comparisons provide a similarity score 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.

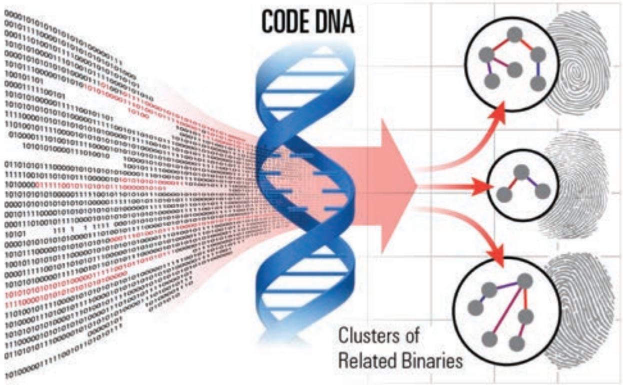

Benefits

CodeDNA supports 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's high-volume fingerprint matching is implemented as a parallel streaming process that runs on inexpensive hardware. 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 (possibly 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 current techniques. In a recent cyber espionage data set 90% of 1000 binaries proved to be strongly related to one another when evaluated with CodeDNA. Another 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. In addition, the recently added fingerprinting for PDF format files demonstrates that adversaries re-use their PDF exploits as well as their executable binary exploits.

Next Steps

The prototype CodeDNA is ready to be piloted and tested within a malware processing environment, followed by a move to enterprise-level testing. We are currently running CodeDNA as part of APL's perimeter defense, and are seeing interesting results. We envision embedding 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 believe that CodeDNA can be expanded beyond our currently supported file formats (x86, Mono/.Net, PDF and JavaScript) to include other executable and digital media formats. 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

Kevin P. McCabe
kmccabe@lanl.gov

Raymond Newell
raymond@lanl.gov

Los Alamos
NATIONAL LABORATORY

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-photonic 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

Velocirandor and Quantum Secured Communications have been licensed and are now part of Whitewood Encryption Systems, Inc. For more information visit Whitewood Encryption System's website (www.whitewoodencryption.com).

Quantum Secured Communications: Security for the Nation's Infrastructure

Kevin P. McCabe
kmccabe@lanl.gov

Raymond Newell
raymond@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.

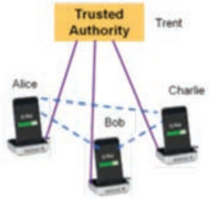

Competitive Advantage

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

Quantum Secured Communications and Velocirandor have been licensed and are now part of Whitewood Encryption Systems, Inc. For more information visit Whitewood Encryption System's website (www.whitewoodencryption.com).

SENSECL: Securing Data for Public Clouds

Dr. Gene Itkis
itkis@ll.mti.edu

Jorge Coll
jorge.coll@ll.mit.edu

Benjamin Kaiser
benjamin.kaiser@ll.mit.edu

This material is based upon work supported by Defense Information Systems Agency and Department of Homeland Security under Air Force Contract No. FA8721-05-C-0002. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of the Air Force.

Overview

MIT Lincoln Laboratory's Self-Enforcing Security for the Cloud (SENSECL) provides cryptographic access control, enabling secure storage of data in public clouds. SENSECL 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 service providers (CSPs).

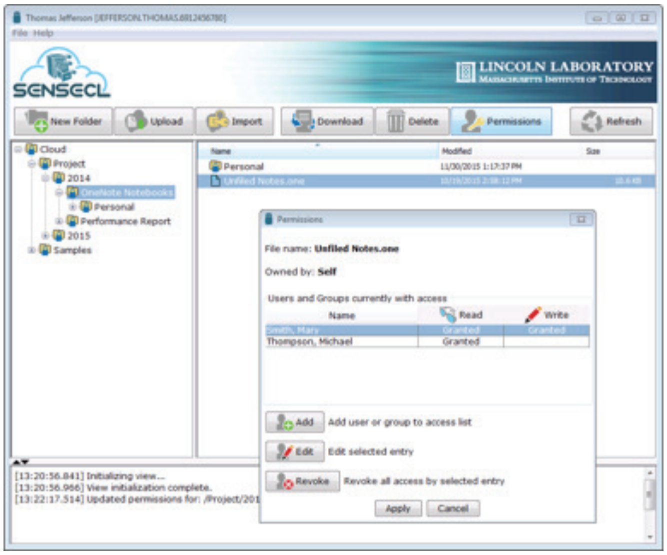

Name	Modified	Size
Personal	11/30/2015 1:17:37 PM	
Unfiled Notes.one	10/19/2015 2:38:12 PM	10.6 KB

File name: **Unfiled Notes.one**

Owned by: **Self**

Users and Groups currently with access

Name	Read	Write
Smith, Mary	Granted	Granted
Thompson, Michael	Granted	

Add user or group to access list

Edit selected entry

Revoke all access by selected entry

```
[13:20:56.841] Initializing view...
[13:20:56.966] View initialization complete.
[13:22:17.514] Updated permissions for: /Project/201
```

Customer Need

Commercial cloud storage offers a significant reduction in cost for many government and commercial organizations while enhancing data availability, ubiquity, and redundancy. These advantages are currently only achievable by outsourcing data management to a third party (i.e., to a CSP), which requires surrendering control over the data and its security. Typically, a 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, called CryptAC, for seamless cryptography and key management providing flexible, cryptographically enforced access control policies ensuring data confidentiality, integrity, and authenticity. SENSECL is built on top of this framework.

In CryptAC, data is cryptographically protected before leaving the client device; therefore, it is protected both at rest and in transit. Each data item is protected using a unique, independent, randomly generated content key that is itself encrypted by the public key of the authorized user and embedded in a cryptographic permission. Only the owner of the associated private key can exercise the permission and access the content. In the simplest example, an individual user owns the public-private key pair. 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 on the client device. Only authorized parties with the appropriate permissions can access the content. Unauthorized parties cannot extract the content key from the permission because they do not possess the necessary private key. Therefore, the CSP and other unauthorized parties never have access to the unprotected content. Using this and other cryptographic methods, access control policies are defined and enforced without having to rely on other parties for the enforcement. These policies ensure data confidentiality, integrity and authenticity while enabling secure sharing with users and groups that have a need-to-know.

This approach restricts administrators to managing data storage without requiring access to data content 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
- Provides an extensible framework enabling 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 on local operating system permissions, since access controls and filesystem storage are typically inextricably coupled. In the cloud computing paradigm, this approach is obsolete, as data is typically replicated across 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 has many benefits, but it necessitates new access controls decoupled from the underlying storage platform, which is exactly what SENSECL provides.

The effectiveness of cryptographic protection depends crucially on key management. Compared to existing protection offered by cloud services (e.g. Dropbox), SENSECL 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, SENSECL supports seamless integration across multiple cloud providers.

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

MIT Lincoln Laboratory is a non-profit, 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

SENSECL provides unique self-enforcing access control, which affords many benefits over traditional access control mechanisms. Even more importantly, SENSECL offers a solid foundation upon which future access control systems can be customized to meet the needs of individual government and corporate customers. SENSECL has already been successfully deployed as part of a government pilot.

LOCKMA: Lincoln Open Cryptographic Key Management Architecture

Roger Khazan
rkh@ll.mit.edu

Dan Utin
danu@ll.mit.edu

MIT Lincoln Laboratory

This material is based upon work supported by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract No. FA8721-05-C-0002. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Assistant Secretary of Defense for Research and Engineering.

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.

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.

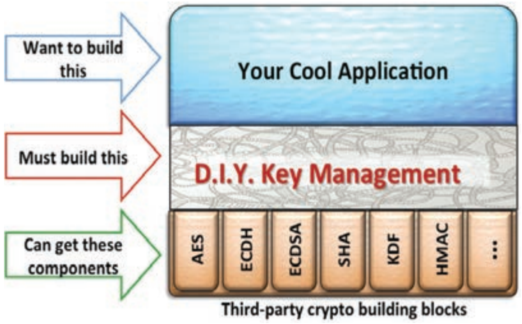

D.I.Y. key management is expensive, and often results in flawed security and hampered usability

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.

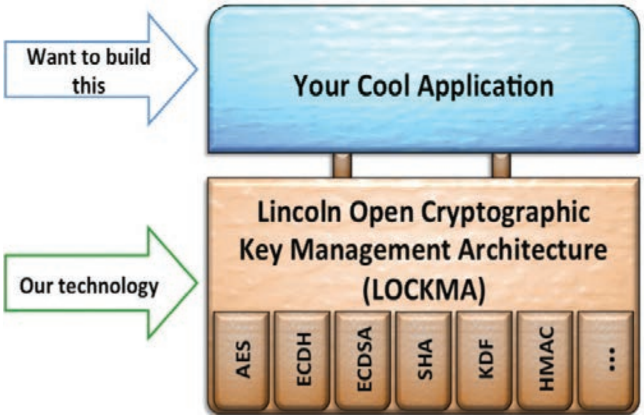

LOCKMA enables strong, reliable, usable crypto protections at low cost

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

Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

Overview

Digital Ants* is a nature-inspired resilient cybersecurity technology designed to protect large enterprise networks and next-generation critical infrastructures. Individual ant-like software agents swarm to the location of anomalies and enable human operators to focus on areas and issues of concern.

Customer Need

Today's enterprise networks are larger, more dynamic and more enmeshed than ever. For example, "bring your own device" is forcing IT enterprises into unplanned growth and uncontrolled permeability. Even static infrastructures such as traditional utility networks are becoming more dynamic as smart devices enable two-way communication and increase customer involvement. Cybersecurity frameworks must become increasingly dynamic and resilient.

Existing cybersecurity monitoring solutions, most of which store data centrally or in a "cloud," are straining to keep up with this runaway growth. The "big data" produced by these tools is seldom actionable. Situational awareness and the ability to take immediate action suffer as analysts labor to sift through low-value data.

A lightweight, extensible cybersecurity framework is needed that can address the ever-changing landscape. The framework must be lightweight to protect legacy devices with limited processing and networking capabilities. Leaving these devices "defended" only by obscurity and poor connectivity is not a viable option. Neither is it cost-effective to replace devices designed for multiple decades of service with new devices every 3-5 years to accommodate the computational needs of new cybersecurity mechanisms. The framework must also be extensible because no one knows what challenges future, unknown (zero-day) malware will present.

Our Approach

Natural systems, like ant colonies, routinely solve adversarial problems that are often harder than today's cybersecurity challenges. The intricate design of ants and their ability to communicate and survive in adversarial environments inspired us to base Digital Ants on ant-colony behaviors of foraging and swarming. Digital Ants' *Sensors* are lightweight, interpreted programs that are always on the move. These sensors roam from machine to machine within an infrastructure, via an overlay network of host-based software agents called *Sentinels*. Sensors constantly gather machine behavioral metrics (e.g., CPU usage, network bandwidth, memory usage) and search for anomalies that cannot be explained by the Sentinel's prior observations. The Sentinels also look for artifacts observed by security monitors.

When a sensor and sentinel agree on an anomaly the sensor leaves behind a digital pheromone trail that attracts other sensors to the anomalous machine, similar to how real ants use chemical pheromones to mark a path to a food source. As more sensors observe unexplained anomalies the resulting pheromone concentration will enable the formation of a sensor swarm. Then the sentinel will inform a human-interface agent called the *Sergeant*. The sergeant, located on another machine, analyzes the

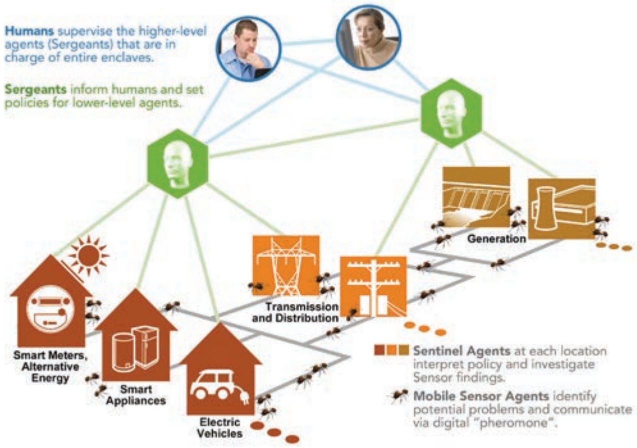
Humans supervise the higher-level agents (Sergeants) that are in charge of entire enclaves.
Sergeants inform humans and set policies for lower-level agents.
Generation
Transmission and Distribution
Smart Meters, Alternative Energy
Smart Appliances
Electric Vehicles
■■■ Sentinel Agents at each location interpret policy and investigate Sensor findings.
Mobile Sensor Agents identify potential problems and communicate via digital "pheromone".

The Digital Ants Framework within the Smart Grid.

*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.

strength of the swarm and the severity of the reported data before informing human operators of a potential issue.

A unique strength of the Digital Ants approach is the framework's ability to identify previously unknown cybersecurity concerns via swarm intelligence. Individually, each sensor provides only a partial indication of anomalous conditions. However, a sensor swarm will form where there are numerous unusual issues. Swarms focus operators on large, complex anomalies rather than individual issue alerts. Digital Ants' sensors require no centralized control, enabling the framework to scale to very large infrastructures. Furthermore, because any combination of sensors can trigger a swarm, unknown malware can be detected and the sensor composition of the swarm can be used to help classify the malware's signature.

Benefits

The Digital Ants framework is a lightweight, cyber defense designed to protect very large infrastructures, even millions of devices. Digital Ants minimizes network communication by performing the analysis at the edge. The framework does not overburden edge CPUs because sensors do not run all the time on all the nodes. Thus, infrastructure owners avoid the cost of high-end centralized storage and analysis servers.

Automatic learning reduces the human cost of configuration and supervision. As the sentinels observe their systems they will silently and efficiently learn to differentiate between normal behaviors and previously unknown or potentially serious situations. Digital Ants handles the low-level anomalies automatically, giving human operators more time to focus on operations and security trends.

Digital Ants' simple, nature-inspired design means there is no single point of failure, making the system naturally resilient to attack. Even if many sensors fail, our approach ensures each node is routinely visited. Although adversaries may be aware of the Digital Ants, they cannot predict the sensor movements providing yet another level of resilience to attacks.

Competitive Advantage

Digital Ants' lightweight framework enables sensor deployment on devices with modest computational and networking resources. Current host-based intrusion detection and intrusion prevention systems consume significant processing power and memory and often require constant connection to external cloud providers. Unlike antivirus, Digital Ants is always learning, so operators do not need to constantly update sensors or malware signatures. Eventually, our goal is to replace heavyweight antivirus and other host-based agents that consume too many host resources with the flexible, lightweight framework provided by Digital Ants.

Next Steps

The Digital Ants Framework has been extensively tested in laboratory environments and on workstation-class machines. Our current implementation can be adapted to a number of potential applications. Going forward, this new cybersecurity paradigm must be field-tested and integrated into existing tools such as Security Information and Event Management (SIEM) products. We are seeking integrators interested in harnessing the power of swarm intelligence to create new resilient protection products. We also seek partners who will integrate Digital Ants technology into existing, domain-specific products that protect infrastructure and IT networks. Finally, we seek research sponsors to help us continue to tune and adapt this revolutionary new approach to cybersecurity.

PACRAT: The Blended Physical and Cyber Risk Analysis Tool

Doug MacDonald
douglas.macdonald@pnnl.gov

Pacific Northwest
NATIONAL LABORATORY
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 cybersecurity 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 cybersecurity 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 cybersecurity 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.

Many organizations consider the cost of performing a comprehensive physical or cybersecurity 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 cybersecurity 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.

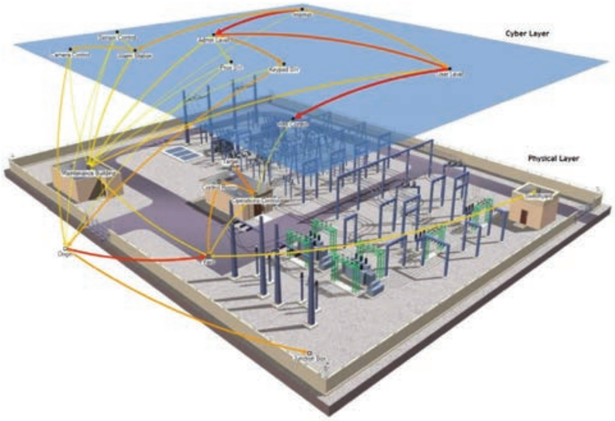

SerialTap: Enabling Complete Situational Awareness in Control Systems

Thomas W. Edgar
thomas.edgar@pnnl.gov

Eric Choi
eric.choi@pnnl.gov

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

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 cybersecurity threats and vulnerabilities but with the merger they now become a problem. ICS now require cybersecurity solutions to defend against these threats.

Large portions of ICS are still operated with legacy serial communications which have largely been ignored by the cybersecurity community. This has led to one of the biggest challenges for control system operators: retrofitting cybersecurity 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.

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.

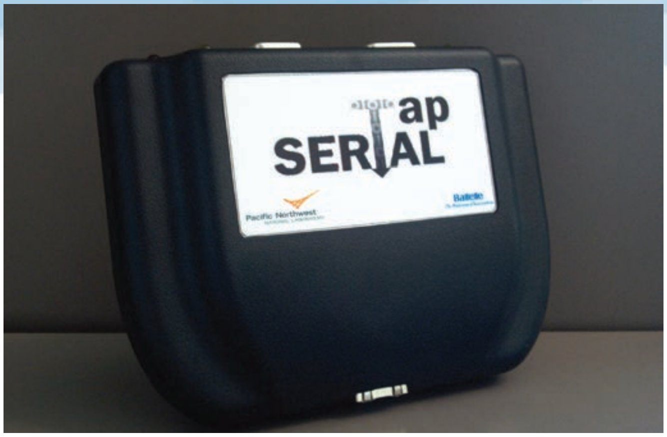

Figure 1: SerialTap Prototype

The SerialTap is designed to encapsulate data and transmit it to a centralized location to ***leverage current enterprise analysis solutions***, such as cybersecurity 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.

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

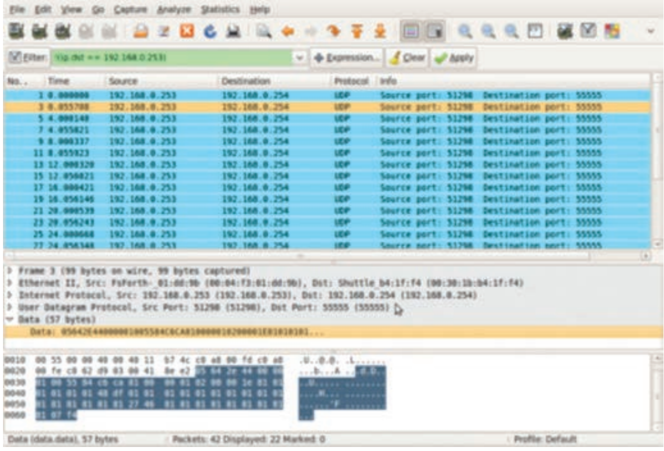
File Edit View Go Capture Analyze Statistics Help
Filter: !(ip.dst == 192.168.0.253)
Expression...
Clear
Apply

```
No.  Time        Source         Destination    Protocol  Info
  1  0.000000    192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
  3  0.055700    192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
  5  4.000140    192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
  7  4.055821    192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
  9  8.000337    192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
 11  8.055823    192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
 13  12.000320   192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
 15  12.050821   192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
 17  16.000421   192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
 19  16.056146   192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
 21  20.000539   192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
 23  20.056243   192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
 25  24.000668   192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555
 27  24.056348   192.168.0.253  192.168.0.254  UDP       Source port: 51298  Destination port: 55555

▷ Frame 3 (99 bytes on wire, 99 bytes captured)
▷ Ethernet II, Src: FsForth-_01:dd:9b (00:04:f3:01:dd:9b), Dst: Shuttle_b4:1f:f4 (00:30:1b:b4:1f:f4)
▷ Internet Protocol, Src: 192.168.0.253 (192.168.0.253), Dst: 192.168.0.254 (192.168.0.254)
▷ User Datagram Protocol, Src Port: 51298 (51298), Dst Port: 55555 (55555)
▽ Data (57 bytes)
    Data: 05642E440000001005504C6CA81000000102000001E01010101...

0010  00 55 00 00 40 00 40 11  b7 4c c0 a8 00 fd c0 a8   .U..@.@. .L......
0020  00 fe c8 62 d9 03 00 41  8e e2 05 64 2e 44 00 00   ...b...A ...d.D..
0030  01 00 55 04 c6 ca 81 00  00 01 02 00 00 1e 01 01   ..U..... ........
0040  01 01 01 01 48 df 01 01  01 01 01 01 01 01 01 01   ....H... ........
0050  01 01 01 01 01 01 27 46  01 01 01 01 01 01 01 01   ......'F ........
0060  01 07 f4                                           ...

Data (data.data), 57 bytes    Packets: 42 Displayed: 22 Marked: 0    Profile: Default
```

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

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 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.

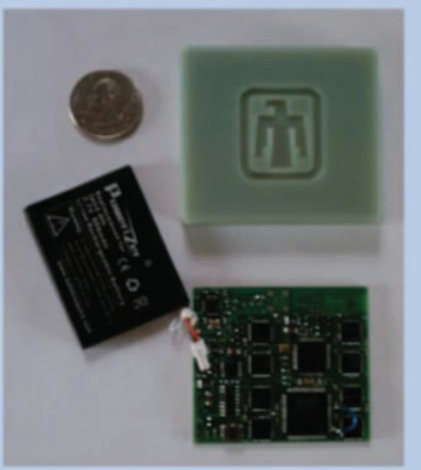

SecuritySeal with battery & package

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.

WeaselBoard: Zero-Day Exploit Protection for PLCs

John Mulder
jmulder@sandia.gov

Sandia National Laboratories

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.

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

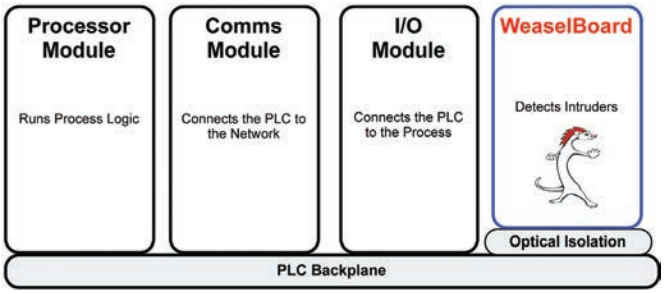

Figure 1: WeaselBoard in a Chassis

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.

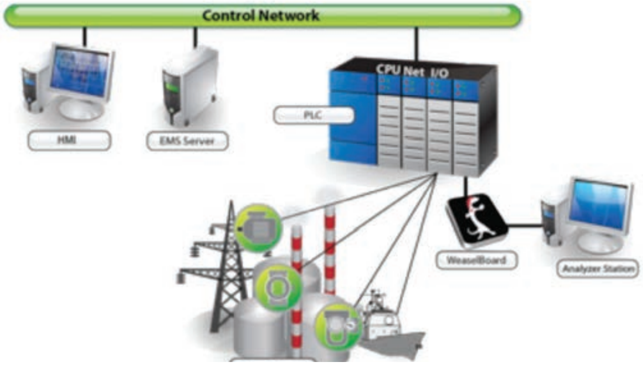

Figure 2: WeaselBoard in an Industrial Control System

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.

```
//v3.0
r(i=0;a&&i<a.length&&(x=a[i])&&x.oSrc;i++) x.src=
//v3.0
{ if(!d.MM_p) d.MM_p=new Array();
_preloadImages.arguments; for(i=0; i<a.length; i++)
 d.MM_p[j]=new Image; d.MM_p[j++].src=a[i];}}
//v4.01
 if((p=n.indexOf("?"))>0&&parent.frames.length
(p+1)].document; n=n.substring(0,p);}
[n]; for (i=0;!x&&i<d.forms.length;i++) x=d.forms
ayers.length;i++) x=MM_findObj(n,d.layers[i].document
 d.getElementById(n); return x;}

arguments; document.MM_sr=new Array; for(i=0;i<
null){document.MM_sr[j++]=x; if(!x.oSrc) x.oSrc=x.

//v3.0
r(i=0;a&&i<a.length&&(x=a[i])&&x.oSrc;i++) x.src=
//v3.0
{ if(!d.MM_p) d.MM_p=new Array();
arguments; for(i=0; i<a.length; i++)
```

FISCAL YEAR 2013 TECHNOLOGIES SUMMARY:

- 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

Fiscal Year 2013 Technologies Summary

Network Mapping System (NeMS)

Lawrence Livermore National Laboratory: NeMS, a software-based network characterization and discovery tool creates queryable graphs of any IP network with details of network entities, attributes, roles, and logical relationships. NeMS was licensed to Cambridge Global Advisors (CGA), which created a startup company, Quellum (http://www.quellum.com), to commercialize the technology.

For more information, contact Celeste Matarazzo, matarazzo1@llnl.gov or Domingo Colon, colon3@llnl.gov

PathScan

Los Alamos National Laboratory: PathScan, a network anomaly-detection tool utilizes statistical models to identify network behavior. Through behavioral models, the technology detects the movement of hackers once they breach the network and allows operational teams to triage and respond to security events in real time. PathScan was licensed to Ernst & Young LLP.

For more information,
contact Josh Neil, Joshua.Neil@ey.com

Choreographer

Oak Ridge National Laboratory: Choreographer, a moving target system thwarts automated network attackers by constantly changing the public addresses of protected servers. This makes it challenging for attackers to guess the server's address and allows a seamless redirection of an attack to a honey pot.

For more information, contact Craig Shue, cshue@ornl.gov

Hyperion

Oak Ridge National Laboratory: Hyperion, a malware forensics detection and software assurance technology computes the behavior of software, including malware, in all circumstances of use, without the need for source code. Hyperion was licensed to R&K Cyber Solutions LLC (http://www.rkcybersolutions.com), an application development and cyber solution company.

For more information,
contact Stacy Prowell, prowellsj@ornl.gov
2015 R&D 100 Award Winner

USB-ARM

Oak Ridge National Laboratory: USB-ARM protects the host against threats from removable media by installing an efficient and customizable layer of security that brokers device communication with the operating system. This tool blocks all communication to the device until a set of user-defined criteria are met.

For more information,
contact Stacy Prowell, prowellsj@ornl.gov

Hone Technology

Pacific Northwest National Laboratory: Hone, a host-based cyber sensor provides a new data source of correlated host and network data by forming a bridge between the networking and processing parts of monitored machines that enables the sensor to know which programs are responsible for malicious network activities. Hone has been made open source and can be found at https://github.com/HoneProject/, it is being actively used by Google, PNNL and others as a cyber security data collection tool.

For more information,
contact Glenn Fink, glenn.fink@pnnl.gov

MLSTONES

Pacific Northwest National Laboratory: MLSTONES, a set of tools that quickly categorizes data and compares attributes of the data to determine if it poses a threat. The methodology uses the concepts of protein identification and families, inheritance and function to apply to a number of cyber-based data types.

For more information,
contact Elena Peterson, elena@pnnl.gov

CodeSeal

Sandia National Laboratories: CodeSeal, a cryptographically secure code obfuscation technology that provides tamper-proof trust anchors; functional elements that are introduced into information systems to provide unbiased measurement and unimpeded control capabilities. The trust anchors protect critical hardware and software components from malicious tampering even when operating in a compromised environment.

For more information,
contact Adrian Chavez, adrchav@sandia.gov

Website
www.dhs.gov/cyber-research

Email
ST.TTP@HQ.DHS.GOV