

ALternatives to Signatures (ALTS)

George M. Jones
John Stogoski

April 2014

Whitepaper

CERT-CC-2014-35

CERT® Coordination Center, Software Engineering Institute

<http://www.sei.cmu.edu>



Copyright 2014 Carnegie Mellon University

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

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 United States Department of Defense.

References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by Carnegie Mellon University or its Software Engineering Institute.

This report was prepared for the
SEI Administrative Agent
AFLCMC/PZM
20 Schilling Circle, Bldg 1305, 3rd floor
Hanscom AFB, MA 01731-2125

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This material has been approved for public release and unlimited distribution except as restricted below.

Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

* These restrictions do not apply to U.S. government entities.

Carnegie Mellon[®], CERT[®] and CERT Coordination Center[®] are registered marks of Carnegie Mellon University.

DM-0001213

Table of Contents

List of Figures	ii
Executive Summary	iii
1 Introduction	1
1.1 About This Report	1
1.2 Methodology	1
1.3 Intrusion Detection Background	1
2 Challenges	3
2.1 Changing User and IT Services	3
2.2 Changing Threat Landscape	3
2.3 Changing Marketplace	3
2.4 Volume and Variety of Data	4
2.5 Anomalous Behavior Is Not Malicious Behavior	4
2.6 The Need for Actionable Reporting	4
2.7 Supporting a Broad Community	4
3 Solutions Discussion	5
3.1 Hunting Operations	6
3.2 Security Information and Event Management	10
3.3 Content Inspection	12
3.4 DNS Analytics	15
3.5 List-Based Monitoring	17
3.6 Network Profiling	20
3.7 Deception	21
3.8 Anomaly Detection	23
4 Summary and Recommendations	26
5 Closing	29
6 Acknowledgments	30
Appendix A: Interview Questions	31
Appendix B: Taxonomies	33
Appendix C: Scalability of Staffing	35
Appendix D: Noise Reduction Techniques	36
Appendix E: Understanding What Should Be Happening	38
Bibliography	39

List of Figures

Figure 2: Hunting Operations Process	6
Figure 3: Collection and Analysis Architecture	7
Figure 4: Hunting Operations Example, Detailed Steps	9
Figure 6: General SIEM Architecture	11
Figure 7: Sandbox Evaluation of Email Attachments	13
Figure 8: Reputation List Solution	19
Figure 9: Intrusion Detection System Example	24
Figure 10: Solution Summary	26

Executive Summary

This report by the CERT® Coordination Center, part of Carnegie Mellon University's Software Engineering Institute, presents the results of a study of non-signature-based approaches to detecting malicious activity in computer network traffic. Our results are based on a survey of the academic literature on anomaly detection (AD) and interviews with personnel from security operations centers at organizations in key sectors.

One major theme emerged from discussions with operators: many organizations are reporting success with security operations that depend on “intuitive engineers” querying against customized back-end storage solutions, a process often referred to as “hunting operations.” These engineers have the knowledge required to understand the context surrounding the network data being analyzed. This includes an understanding of the operations of the organization, the structure of its networks, and sophisticated understanding of network protocols and the observable behaviors of normal and malicious activity on the network. As new types of attacks are identified, often using ad hoc analysis relying on non-signature-based approaches, analysts leverage opportunities to translate these analyses into repeatable, automatable processes for alerting on suspicious activity.

At the same time, there is a high degree of dissatisfaction with the flexibility, interoperability, scalability, and ultimately with the perceived value of commercial solutions for detecting malicious behavior. Cybersecurity operators are increasingly using open source software and custom storage solutions to address some of these problems. Commercial vendors are responding to the need by updating offerings to include more flexible collection, storage, and analysis platforms.

Hunting operations are still as much art as science, relying heavily on the availability of highly qualified staff with the time and opportunity to understand the environments they defend. Research and development of non-signature-based technologies should identify opportunities to effectively focus the use of valuable analyst time, support knowledge management to decrease the burden on individual analysts to retain context, and generally provide more powerful tools for human analysts.

* CERT® is a registered mark owned by Carnegie Mellon University.

1 Introduction

1.1 About This Report

The Office of the Under Secretary for Defense Policy (OUSDP) tasked the Software Engineering Institute (SEI) at Carnegie Mellon University (CMU) with independently surveying and assessing non-signature-based approaches, tools, and techniques to cybersecurity, with a specific focus on network-based detection of malicious activity. The goals of this study were to

- characterize current and emerging non-signature-based approaches, tools, and techniques
- identify and characterize industry best practices for non-signature-based approaches, tools, and techniques
- characterize the maturity of these approaches and the state of their adoption by industry
- identify the most promising current, emerging, and future non-signature-based approaches and technologies that merit consideration

The study had two phases. In the first phase, the SEI selectively surveyed the large body of anomaly detection (AD) literature to identify potential solutions.¹ In the second phase, the SEI interviewed participants from commercial companies. This report synthesizes the findings from both phases of the study.

1.2 Methodology

In the first phase of the project, the SEI selectively surveyed the large body of anomaly detection (AD) literature to identify potential solutions. In the second phase, the SEI interviewed participants from commercial companies. The participants were largely technical security managers and staff responsible for internal security engineering and operations. We used a standard set of questions (see Appendix A) to guide the discussion, but we encouraged participants to elaborate or address other relevant areas based on their own experiences. The results of the interviews provided a view of best practices and lessons learned from subject matter experts within the community, with the understanding that these organizations do not represent the whole community, but rather individual entities from critical infrastructure sectors. This report discusses the key challenges, potential solutions, and policy recommendations that came from the literature survey, interviews, and the experience of SEI researchers.

1.3 Intrusion Detection Background

What is malicious activity?

The term *unauthorized result* can be defined as increased access, disclosure of information, corruption of information, denial of service, or theft of resources caused by an attack. In this report, *malicious activity* means activity with the goal of obtaining an unauthorized result and that has a negative impact on Critical Success Factors [CAW10].

¹ Additional detail from the literature survey can be found in the paper *What You Know and What You Don't Know: A Survey of Anomaly Detection Techniques Applied to Network Intrusion Detection* [JM13].

What are the threats to mission success?

Organizations need to examine threats in terms of the organizational mission and operating model to understand the real-world possibilities of disrupting their business functions. Threats vary significantly across government and industry sectors. Some of the interviewed operators indicated that threats change due to corporate events such as high-profile customer engagements or international conferences. Such events add a timing element to data analysis. All organizational factors influence the specific technologies operators implement within their infrastructure and the processes analysts use to identify suspicious activity.

How do you measure success?

One of the survey's goals was to understand how the respondents judged their own effectiveness and how their organization's leadership evaluated it. Interviewees framed most of their answers in terms of something *not* happening, which is often hard to demonstrate. Cybersecurity personnel struggle to measure their success, and there does not appear to be a good solution. Interviewees' responses to this issue include the following:

- “Success is staying out of the *Washington Post*” and protecting the brand.
- Counting tickets produces metrics related to workload, but it is not at all clear that increased ticket counts correspond to improved performance.
- It seems that the important attacks are becoming more sophisticated and fewer in number. However, it is not clear how one would measure these variables together. One indication might be a decline in the number of tickets concurrent with increasing risk and loss.
- Compliance with standards and regulatory requirements is one way to measure success. But do these equate to improved security?
- Return-on-investment (ROI) calculations have not worked. Leadership now sees security as a cost of business.

2 Challenges

This section discusses the cyber community's key challenges to detecting malicious activity. These themes emerged from the survey of AD literature and interviews with the industry participants. Operational security groups are continuously trying to address these challenges, both internally and by working with vendors and researchers. The challenges impact the effectiveness of products and the feasibility of their wide adoption.

2.1 Changing User and IT Services

The rise of mobile device use, bring-your-own-device (BYOD) policies, cloud computing, and the share-everything culture of social networking is changing how networks are being used. The diminishing expectations of privacy and data protection, both at the personal and organizational levels, are influencing the community's definition of an incident. This impacts the role and capabilities of IT security groups in managing and controlling activities across the network. In addition, the idea of the network perimeter is vanishing, and enterprises have to become more focused on the end points. The result is that data is more distributed across many types of devices and harder to protect.

2.2 Changing Threat Landscape

The threat landscape is never static. Changing factors that affect the threat landscape include the motivation, capabilities, and resources of adversaries; vulnerabilities that are discovered as new devices and services are deployed or upgraded; and the value of assets, both to attackers and defenders.

2.3 Changing Marketplace

The cybersecurity marketplace has changed significantly over the last few years. Vendor products that provide sensors and management platforms that are proprietary, or silo-solutions, are seen as too restrictive and are falling out of favor. We have observed a change in emphasis from commercial products to open source solutions. Vendor-provided sensors are still being used, but they have become a source of data instead of the complete solution. Overall, commercial tools are seen to have insufficient flexibility, manageability, and interoperability. Several interview respondents reported dissatisfaction with the scalability of vendor solutions, particularly in storage and retrieval. One organization reported it largely builds its own tools due to the many problems with commercial tools. This organization found that there is a lack of application programming interfaces (APIs) for monitoring product performance and sharing data across platforms.

Open source products are providing alternatives for products such as antivirus (AV) and intrusion detection systems (IDS). Support structures for open source solutions have been enhanced significantly, providing the assurances enterprise customers require. In addition, some products are becoming part of larger platforms. For instance, Microsoft has begun building AV into their operating systems. This kind of combined functionality allows security organizations to shift where they spend their budgets.

2.4 Volume and Variety of Data

We are drowning in information, but starved for knowledge—John Naisbitt

The high volume and variety of data make it difficult to determine what is important, what should be collected, what analysts should look at, where it should be stored, and how to scale storage solutions [Gol12]. Some possible data sources in modern enterprises include proxy logs, email metadata, firewall logs, IDS logs, intrusion prevention system (IPS) logs, authentication data, dynamic host configuration protocol (DHCP) logs, domain name system (DNS) query logs, badge data, netflow, host-based collection, and crash-dump analysis.

2.5 Anomalous Behavior Is Not Malicious Behavior

For more than a decade, AD has been pursued as a way to increase detection rates. Gates and Taylor [GT06] list assumptions underlying the use of AD for intrusion detection. The most notable assumption is that attack behavior is somehow different from normal traffic, but attackers might hide their activities in normal traffic.

2.6 The Need for Actionable Reporting

Sommer and Paxson frame the need for actionable reporting this way: “Anomaly detection systems face a key challenge of transferring their results into actionable reports for the network operator. In many studies, we observe a lack of this crucial final step, which we term the *semantic gap*.... The next step then needs to interpret the results from an operator’s point of view — ‘*What does it mean?*’ Answering this question goes to the heart of the difference between finding *abnormal activity* and *attacks*” [SP10].

The high diversity of traffic on today’s networks increases the challenge. There is an increasing amount of variability in network traffic that operators need to manage. It is becoming more difficult to understand normal traffic and detect the important signal in the noise. Such variability and complexity leads to more difficulty in detecting attacks in large environments.

Proper interpretation of the data requires sufficient context. The same activity related to two different IP addresses could mean two very different things. The role of the server or device to which the IP address is assigned has a significant impact on the analysis of the event. Analysis done at large network consolidation points (e.g., enterprise internet gateways) is less effective at finding new or sophisticated threats due to reduced knowledge of the target environment. Further, analysts who do not have a strong understanding of acceptable or normal business processes have difficulty determining whether an anomaly is malicious. Simpler environments with smaller, well-defined processes make it easier to spot deviations and determine the proper mitigation and response measures.

2.7 Supporting a Broad Community

The cyber defense community comprises a wide range of members, from two-person shops to well-defined organizations with hundreds of analysts and operators. This reflects the range of missions and IT environments that are being supported. A single solution will not meet everyone’s needs, so the community must develop solutions that can address the existing range of operating environments.

3 Solutions Discussion

The interviews included discussions of several intrusion detection approaches, providing a good view into best practices from leading cyber organizations. This chapter provides brief descriptions of each approach along with information on costs, risks, benefits, and any policy recommendations. During the discussions, it became evident that many organizations have a strategy that integrates detection and protection, so some of these solutions perform both duties.

The study examined the following solutions:

- hunting operations
- security information and event management (SIEM)
- content inspection
- DNS analytics
- list-based monitoring
- network profiling
- deception
- anomaly detection

These solutions differ in nature and deployment location within the enterprise (see Figure 1). The hunting operations and SIEM solutions are predominately operating within the security enclave, with sensors sending data from locations across the enterprise networks. Content inspection services, such as web proxy and sandbox solutions, reside near the internet gateways or other egress points. DNS, access control, and network profiling infrastructures operate in the network by sending and receiving information to and from the security enclave. Deception solutions, such as honeypots, and intrusion detection sensors reside at internet gateways or other locations in the network based on specific needs. These differences affect their performance, cost structures, deployment complexity, and security risks. Most of the solutions examined use some form of AD.

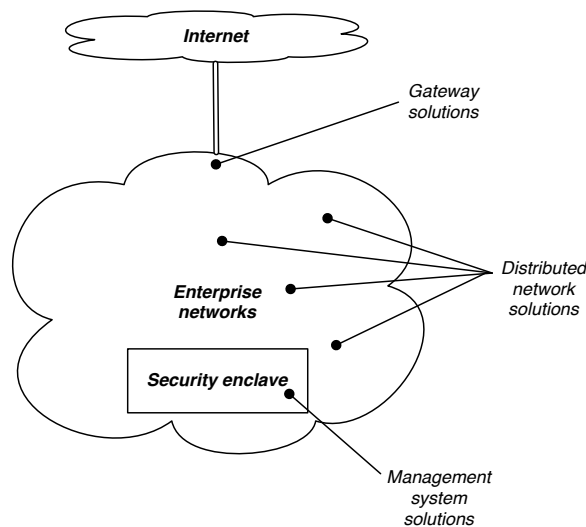


Figure 1: Security Deployment Illustration

3.1 Hunting Operations

Description

From our interviews, it seems that the discovery of new and unknown attacks is being accomplished by human analysts who have been given freedom to perform ad hoc analytic activities using all of their available resources. One respondent used the terms “intuitive engineer” and “hunting operations” to describe the process of analysts looking for malicious behavior. Another said, “Vendors will never know what an anomaly is in our environment; we do.” Trade publications provide some confirmation that hunting is emerging as an effective approach [Ros12], and government has extensively used the term as well as hunting processes.

The general process outlined in Figure 2 is to mine data, identify suspicious activity, conduct the investigation, and then codify an analytic around it, if appropriate. Analysts perform heuristic queries informed by expert knowledge of the allowed and expected behavior and the controls implemented. For instance, a small domestic financial firm does not expect any VPN connectivity from outside the United States. Such connections, if they occur, would be inherently suspicious. Once suspicious activities are identified, analysts investigate in more detail and determine whether it is benign or a security incident. This process is performed iteratively, sorting through the potentially huge amounts of information to identify incidents. It could be characterized as *throwing out the hay to find the needles*. Analysts will often focus on a few common protocols and apply their understanding of related business processes and their expected behaviors. Excluding the rest of the data makes the interesting things that require additional examination stand out. This process often results in the creation of a repeatable process that looks for suspicious IP addresses, domain names, certificates, and other artifacts. Automating the repeatable process often results in new signatures that can be deployed to existing infrastructure, minimizing analysts’ workload.

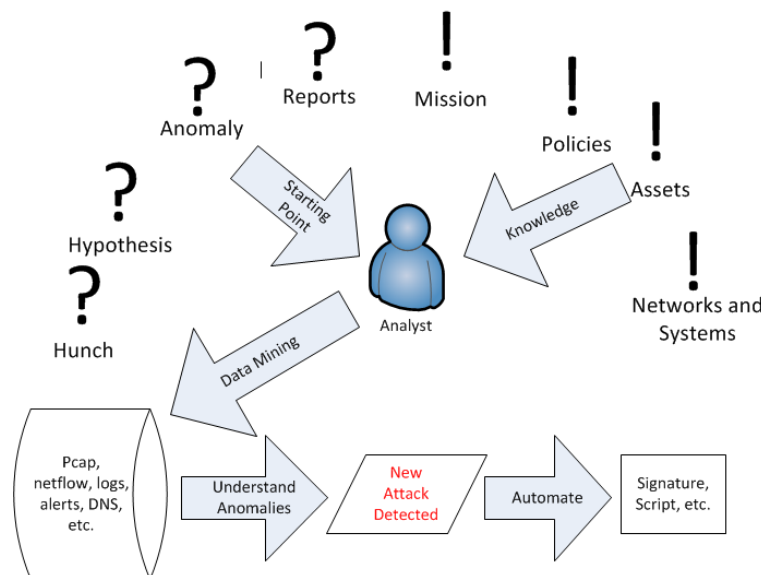


Figure 2: Hunting Operations Process

The discussions with the survey participants revealed that most shared the same view of architectures for detecting malicious activity. Nearly all organizations reported using custom, back-end data storage solutions, in many cases a Hadoop implementation, and open source software to support their detection and analysis functions. Technology developments have provided new “big data” capabilities that were not available a few years ago. This strategy is being used to address the dissatisfaction with the flexibility, interoperability, scalability, and poor perceived value of commercial solutions. Also, for situations where signatures are not available, deployed sensors are serving as data collection devices, and the intrusion detection activities are being performed in the centralized, analysis functions of the architecture. This approach is greatly improved when analysts understand the environment and have the required context to relate data to physical assets and business operations. Figure 3 illustrates the key components of the architecture required to support this style of analysis.

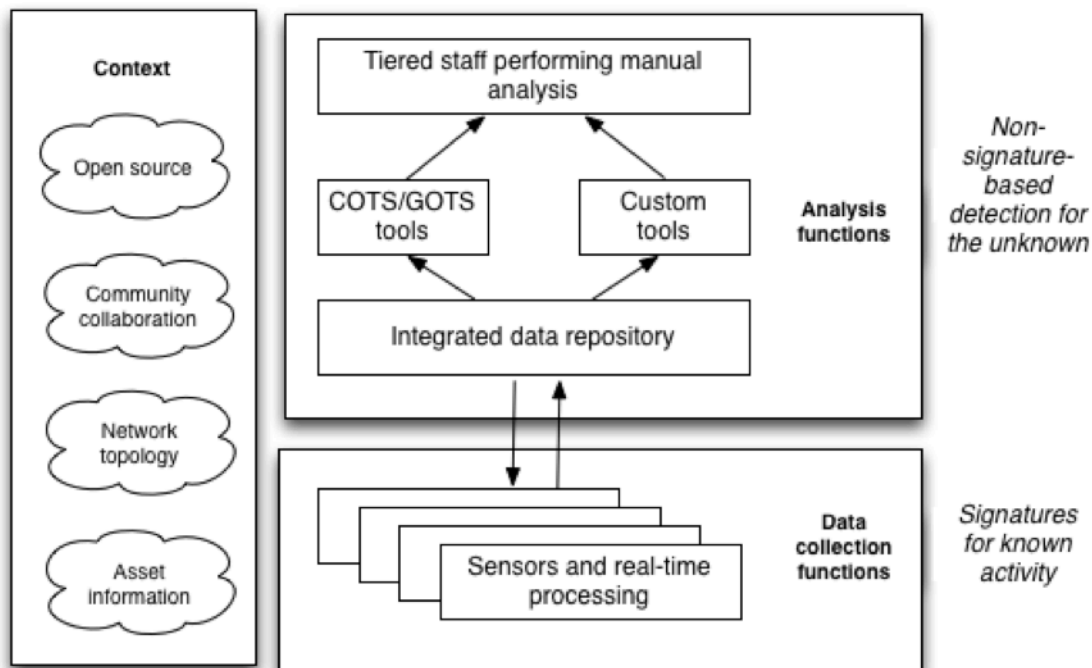


Figure 3: Collection and Analysis Architecture

Examples

Identifying suspicious VPN activity: Analysts develop an understanding of normal traffic patterns over time, such as the timing and source of VPN connections coming into the enterprise. For example, analysts can search log data, stored in a data repository, for connections that originate from international locations. This might suggest that a VPN connection has been established using stolen credentials. This information can be joined with access logs and travel records to validate the activity. Further investigation into what application activity occurred over the VPN tunnel and specific logs from those application servers can lead an analyst to the details needed to determine if an actual intrusion took place.

Using real-world context to identify spear phishing: One interviewee described a scenario in which the operator is notified when executives are going to external conferences or trips to important, high-

profile meetings. These occurrences change their threat profile and justify additional diligence. Analysts can examine email logs for suspicious sender and subject headings that could indicate a spear phishing attack related to recent meetings. Public activities can provide enough contextual information for an attacker to appear legitimate and fool a recipient into accessing a malicious website that infects the victim's computer. The analyst can then examine the logs to determine if this type of attack occurred and the nature of the specific attack.

Identifying beaconing malware: The example detailed in Figure 4 describes a process that was used successfully to detect machines infected with previously unknown malware. The malware was confirmed as unknown by a commercial threat feed provider and the process was transitioned to a SOC team for operational use.

Certain classes of malware “phone home” to command and control servers using beaconing behavior. Beaconing is traffic that repeats at regular intervals and/or with regular payload size. Detection of this communication could lead to the identification of infected systems and allow damage assessment and remediation.

This example is presented to illustrate the steps, types of data, activities, and knowledge required to perform hunting operations. It also lists, for each step, if the step can be automated. It assumes a level of familiarity with networking and security technology and terms such as network flow, DNS, blacklists, reputation lists, and beaconing. In the environment of this example, two of the major sources of data available were network flow data and a database of past passive DNS queries and responses. Available data will differ in different environments. Certain types of knowledge will also change in different environments: the internal network range, expected behavior of different systems, local policies, and assets. Others may change less from one environment to the next: DNS blacklists and IP reputation lists. Some types of activities are likely to remain constant: data selection, calculation of derived statistics, steps to understand what *is* happening, steps to understand what *should be* happening, indicator expansion, and manual investigation. Some steps of this analysis can be automated, mostly data acquisition and checking data against other sources. Later steps require manual analysis.

Step	Comments	Data					Activity Type					Knowledge Required	Can be automated?
		Netflow	Beacon Whitelist database	Reverse DNS	DNS blacklists	IP reputation lists	Data Selection	Calculate derived statistics	What is happening	What should be happening	Indicator Expansion	Manual Investigation	
Pull sample flow		✓					✓	✓				Target time-range, target networks	✓
Select likely beacon traffic		✓						✓				Heuristics network characteristics of beacon traffic	✓
Calculate entropy	This example used flow data (bytes per packet). Packet contents are another candidate for entropy calculation.	✓						✓	✓			Fields whose entropy is likely to indicate beaconing	✓
Flag low entropy flows	Low entropy means things (bytes per packet) are very similar. We assume beacons do this.	✓						✓				Choice of right entropy values.	✓
Check beacon whitelist, discard whitelisted flows	Check whitelists of beacons that are known to be OK.		✓							✓		Previously whitelisted beacons	✓
List IP addresses of remaining external beacon destination		✓					✓	✓					✓
Do reverse DNS lookups on IP addresses of potential external beacons	Find the names that the IP addresses mapped to. This is an example of "indicator expansion", finding more indicators of malicious activity given one. The location of the collection of the passive DNS data is important. Data collected locally (e.g. from the local bind server logs) is more likely to be accurate than data collected at the Internet peering points.	✓	✓						✓	✓		Timestamp from flows	✓
Manual inspection of domain names	Look for "weird" or "obviously malicious" domain names. Consult DNS blacklists. There has been work aimed at automating this step by automatically detecting domains generated by Domain Generation Algorithms (DGA).				✓					✓	✓	Intuition, knowledge of "normal" domain names	Partial via DGA detection
Consult IP reputation lists	Check external beacon addresses against IP reputation lists in an attempt to help determine intent.				✓					✓	✓	Understanding of reputation lists results	Partial. Lookup, not interpretation.
Investigate suspicious results	Manual investigation of suspicious results. Further examination of network and system data. Possible remediation such as reimaging compromised systems.										✓	Judgment, knowledge of policies, networks, systems, mission, normal behavior. Knowledge of investigative tools.	No
Update beaconing whitelist	Add obviously valid beacons to whitelist. See http://www.cert.org/flocon/2013/posters/allen-anand-behavioral-whitelists-of-beaconing-activity.pdf	✓									✓	Judgement, knowledge of protocols and hosts that "should" beacon	No

Figure 4: Hunting Operations Example, Detailed Steps

Risks and Vulnerabilities

The hunting operations approach is based on the concept of collecting large quantities of different types of data and storing it in a centralized repository to be analyzed. This raises a number of security concerns related to protecting the confidentiality and integrity of the data. The scope and level of detail can include very sensitive information pertaining to the business and its employees. The system must be protected against intruders gaining access to analyze the data themselves or possibly changing the information to cover their tracks. This also raises additional security concerns regarding the analysts and operators who have legitimate access to the system. Background investigations, training, auditing, and monitoring can help reduce the risks associated with security personnel with malicious intent.

Level of Adoption

Most of the interviewees had conducted some form of hunting operation, but the details of those implementations varied widely. The interviewees acknowledged that most small and medium-sized organizations do not have the funding and staffing resources to support a hunting operations component. These restrictions should lessen as more skilled staff become available and the technology continues to mature.

Costs and Benefit Considerations

The hunting operations solution is providing positive results and is considered by most to be an effective strategy. While it does require staffing and infrastructure resources, it can be scaled and customized depending on the organization's needs and budget constraints. Many of the software components are available via free, open source distributions, but adequately sized hardware will still be required. The solution also requires the collection of data from various sensors and services across the enterprise, which can create other infrastructure costs. Hunting operations have a major dependency on staff developing the required technical skills to support the database and perform the analytic operations. Finding qualified individuals is a challenge, and developing the skills internally takes time. As the associated tools and techniques continue to mature, it is believed that the solution's effectiveness in detecting more sophisticated attacks will continue to increase.

Policy Implications

Adoption: Given the effectiveness, low risk, and in some cases ease of deployment of hunting operations, further adoption should be encouraged.

Education and staffing: A human resources strategy will need to be developed for government and industry to increase the size of the analyst pool. Individuals will need to be trained in the universities and provided ongoing training. Initiatives will need to be developed to attract and retain quality staff.

3.2 Security Information and Event Management

Description

SIEM tools aggregate and correlate security data to identify and prioritize incidents. The tools attempt to relate data from multiple devices to a single event, providing the analyst with as much information as possible. SIEM products provide a number of functions, including log collection, real-time analysis, reporting, and operational dashboards. Figure 5 illustrates the general architecture of many of the

commercial solutions available. As in hunting operations, the approach requires the collection of data from various sources. However, SIEM solutions sometimes require specific collection devices that are proprietary, can automatically analyze aggregated data, and have scalability constraints. The interviews indicated a significant shift away from traditional SIEM products.

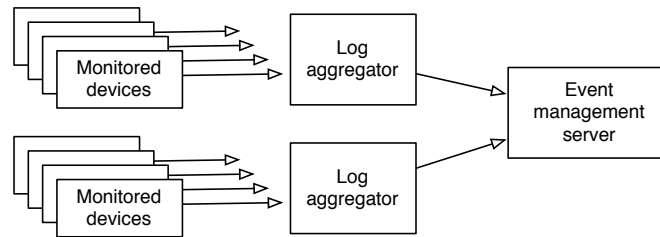


Figure 5: General SIEM Architecture

Examples

Vendors tend to focus on specific core functionality while including some capabilities in the other areas. For instance, HP's ArcSight product does real-time analysis of data and includes a robust graphical user interface for ops-center environments. Splunk is recognized for its log management capabilities that allow analysts to sift through large amounts of machine data. Symantec, McAfee, and IBM's Q1 Labs are other companies offering SIEM products. OSSIM provides one open source alternative.

Risks and Vulnerabilities

SIEM solutions have many of the same security concerns as hunting operations. One difference is that SIEM implementations often do not address the same wide scope of data that hunting operations collect. SIEM solutions require a significant amount of customized tuning for specific environments, thus increasing the chance that it will miss attacks. Handling the volume of data produced by security devices across large enterprises presents its own operational risks. The SIEM product's effectiveness depends on the capabilities of the sensors and other devices that are feeding it data. During deployment, integrating all these devices requires significant planning and coordination. Technical interfaces must be determined, potentially requiring configuration changes to operational networks. Many of the devices may be managed by other organizations, entailing the development of internal agreements and operational procedures. Lastly, as in other solutions, analysts must have sufficient context to assess SIEM alerts and determine the appropriate actions.

Level of Adoption

SIEM solutions are widely deployed. However, custom storage solutions and products that allow more free-form queries of large volumes of data and require minimal investment and set-up time are attracting users. SIEM vendors are tailoring products to emphasize areas that have proven successful and are now developing "big data" strategies such as integration with Hadoop.

Costs and Benefit Considerations

There seems to be a perception that SIEM systems have delivered minimal value compared to the significant capital and labor investments required for their deployment and ongoing operation. Many of the solutions are proprietary, requiring vendor support contracts and limiting the amount of

customization and integration that customers can perform. Open source solutions are providing a feasible alternative for many organizations.

The feedback we received about SIEM solutions came mostly from large enterprises. However, SIEM products may be better suited to the smaller data volume and lesser data diversity of small and medium-sized organizations. The workflow features of SIEM products provide a level of automation that may be sufficient to support their incident handling and analysis needs.

3.3 Content Inspection

Description

Several interview participants highlighted the value of sandboxing to inspect content that is coming into their networks (Figure 6). Email attachments and downloads from websites are common vectors for malware. From a kill chain perspective [HCA11], this would be the delivery phase of the attack. Sandboxing techniques do not depend on signatures. Rather, they examine the behavior of the file to assess its potential risk. Content inspection's behavior analysis approach and ability to process traffic in near-real time make it a defense against zero-day attacks.

Virtual machine technology provides an economical method of constructing closed environments where it is safe to open downloaded documents and code. This technique can help identify characteristics of the potential malware such as attempted connections back to a command and control (C2) system. Content inspection monitors the behavior of downloads in a protected environment, blocks malware before it gets to the user, and analyzes it to better understand emerging malicious code techniques and attack campaigns. This solution should be applied in conjunction with signature-based solutions such as virus protection, IDS, and firewalls.

A significant portion of the process is automatable using commercially available and open source solutions. Tuning allows administrators to incorporate specific policies, further reducing the need for manual intervention. For instance, applying a policy that bans email attachments from resulting in code execution or network activity may allow the sandbox device to decide whether to forward an attachment. In other situations, malware analysts will need to investigate the specifics characteristics of the file, possibly talking to the targeted recipient, and make the final decision on whether it can be forwarded. [FireEye](#) and [ThreatTrack](#) are two commercial sandbox execution solutions, and [Cuckoo](#) provides an open source solution. We are aware of two organizations that developed their own custom sandbox solutions.

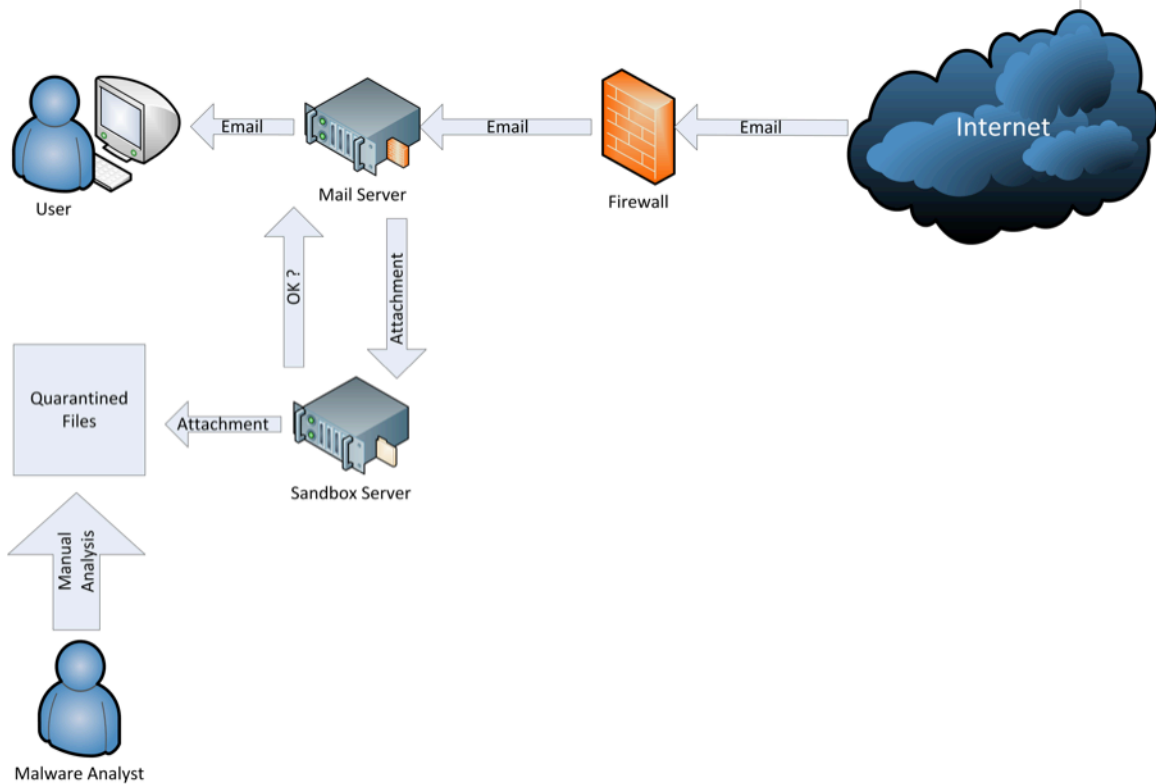


Figure 6: Sandbox Evaluation of Email Attachments

Examples

Sandboxing email attachments: Documents that are received via email attachments are a typical injection vector for malicious software. If the act of opening a document (e.g., a Word or PDF document) causes network traffic, changes to registry settings, or other system modifications, the document may be malicious. This is a case where intuitive logic (opening documents should not cause network activity) is applied to determine malicious intent.

Sandboxing web downloads: Downloads from websites can also result in code that executes in the user's environment. Sometimes this is expected (in the case of Java or executable files). Sometimes it is not. Web downloads can also be run and observed in virtualized environments. This is a case where behavioral analysis (e.g., attempted network connections) can be combined with other information (e.g., C2 blacklists, reputation lists, local policy) to determine intent.

Crawling web pages for executable content: Sandboxing allows for even more proactive solutions, such as crawling web pages, downloading executable content, executing it in a sandbox environment, and using that information to detect new malware or build reputation lists. See MITRE's [HoneyClient](#) or Microsoft's [HoneyMonkey](#) for two examples. Google also uses this technique in conjunction with its Safe Browsing project [[rajab2011trends](#)].

Risks and Vulnerabilities

- Direct attacks on the sandbox execution environment are possible, such as attempts to compromise the security of virtual machines running on the same physical hardware. These attacks are occurring and being discussed [[ormandy2007empirical](#)]. There have also been

attempts to evade honey clients [[kapravelos2011escape](#)]. This is just the latest example of the digital arms race: as effective defensive cybersecurity solutions are identified, adversaries target the solutions themselves and explore evasive techniques.

- Sandbox analysis may miss malicious behaviors (i.e., produce false negatives). For example, some malware attempts to detect if it is executing in a virtualized environment and, if so, behaves differently. Also, malicious behaviors may depend on time, user input, execution context, and other factors that cannot be duplicated in the sandbox environment.
- Sandbox execution may incorrectly flag content as malicious (i.e., produce false positives) resulting, for example, in a valid email attachment being quarantined.
- Delayed or blocked content might impact users.
- It is not realistic to replicate all possible computing environments; complexity increases as the number of applications increases.

Level of Adoption

Sandbox execution environments have been used in one form or another for a very long time, arguably since the invention of timesharing systems (including Unix and Windows). A sandbox security model has also been used at the heart of the Java programming language, the Google Android platform, and web browser and application security models.

The routine application of sandbox execution to web and email content to determine malicious intent is somewhat more recent and is growing in popularity as it proves effective. Mail servers, web proxies, application firewalls, and other products have long had the capability to pass content to antivirus products (signature-based solutions), and sandboxing extends that model to behavioral solutions (non-signature-based).

Costs and Benefit Considerations

The major benefit of content inspection is its ability to detect malicious behavior even in the absence of signatures; it can, in some cases, identify unknown attacks.

Cost considerations for commercial products include the software, hardware, and support costs. Open source solutions require knowledgeable support to deploy and maintain. These implementations process live traffic, increasing the care that must be taken to minimize impacts to users. Deployment times and costs are likely to be proportional to those for deploying antivirus scanners for incoming content. Once malware is detected, the content can be quarantined. Further decisions will have to be made about the level of manual inspection and analysis desired for quarantined content. This will have implications for staffing, for instance, the number of malware analysts required.

Sandboxing, instead of a complete block action, can be used in conjunction with reputation-based lists as a mitigation option for sites with low risk.

Policy Implications

Adoption: Encourage further deployment of sandbox solutions.

Research and development: Support research to identify inherent limitations of sandbox solutions, ways to enhance effectiveness, and the application of sandboxing beyond web and email (for example,

mobile, cloud, and virtual desktop environments). Ongoing development will be required to address countermeasures.

Staffing: Consider staffing levels required to support manual inspection (malware analysis) of quarantined content.

3.4 DNS Analytics

Description

DNS is a fundamental piece of internet architecture. Every client connecting to the internet, or enterprise networks using internet technology, makes requests to DNS servers, most often to resolve names to addresses before initiating connections to servers. The protocols are well understood and not encrypted. A DNS request often indicates an intended action or communication partner. These all combine to make analysis of DNS at every level a valuable source of security information.

When attempting to construct behavior anomaly detection systems, it is helpful to limit the scope of the applications in which anomalies need to be detected. A good target application will be common, simple, and widely used, but not obscured. DNS fits this profile perfectly. A DNS client is universally included in client and server operating systems, and most internet traffic starts with a DNS query. At the same time, the protocol is compact and simple, so data volumes are kept low and parsing is not prohibitively expensive. Finally, DNS traffic content is not encrypted, so analysts should always have access to the payload of the application data, unlike other common protocols such as HTTP or SSH, which are often or always encrypted.

Passive DNS collection has also cleared legal hurdles to adoption. A DHS privacy impact assessment for the Einstein system has approved passive DNS collection on U.S. federal government systems [[DHS12](#)].

Any DNS exchange contains information on at least three entities. Trivially, it indicates that a client asked a server for information. But the content of the messages is often far more useful. It provides the IP addresses on which domains reside, what name servers serve what domains, which domains have Domain Name System Security Extensions (DNSSEC) enabled, various information about e-mail handling capabilities (MX, SPF, DKIM), and domains that are related due to redirection (CNAME). By analyzing the content of DNS messages, the analyst can discover malicious resources. The DNS can also be used to deter internal clients from accessing such malicious resources by modifying DNS responses for such resources, such as by using BIND's request policy zone (RPZ) feature.

There are several monitoring point placement options, which give slightly different results [[PB11](#)]. In general, DNS messages can be passively observed from the network without loss. If the organization centralizes its DNS traffic by forcing all client traffic through a recursive resolver, the traffic can be monitored at this network link. But the better advantage of an organizational recursive resolver is that it is a single point at which to enforce DNS policy decisions. Monitoring can be done anywhere on the network that can see all the DNS traffic.

Examples

Detecting rapidly changing malicious resources: Malicious domains can use a network of domain names and IP addresses that change rapidly to avoid blocking attempts that use traditional signature-

based approaches. This technique is called *fast flux* when a single domain changes IP address rapidly. When a malicious resource changes both domain and IP rapidly, it is called *double flux*. Other features of a domain can also flux, such as name server, though this is less common.

Methods for detecting fast flux and double flux detect the behavior (see [Sto10] and [HGRF08]) and do not require a priori knowledge of suspicious domains, so any resource exhibiting this behavior can be blocked. The detection algorithms require a minimum amount of data before they can detect flux, usually about 10 domains in the fast flux network. Although these first 10 may still cause damage, the next 10,000 in the same flux network could be quickly detected and blocked.

Indicator expansion tool: Indicator expansion the “process of using one or more data sources to obtain more indicators of malicious activity by identifying those related to currently known indicators” [Spr12]. Passive DNS data provides a rich graph of connections among elements of the active internet. It is valuable because the DNS queries are passively observed, which provides a high degree of confidence the values are real and active, not provided spuriously by an attacker who has noticed the defender is scanning.

Analysts can use indicator expansion in various situational awareness activities to discover previously unknown indicators of malicious activity. DNS is not the only source of information for this process, but it is a central one.

Detect newly active domains: If a domain has never been seen before, it is more likely to be malicious. Some defenders choose to prevent access to a domain the first time it is requested on their network, an activity sometimes called *gray listing*. If the DNS messages are monitored, it is easy to determine if a domain has been seen before. If the enterprise has a central DNS recursive service, responses to such domains can be delayed or modified while the new domain is analyzed to determine if it is benign or malicious. This defense complements flux detection, for example.

Detecting parking: Domain names are occasionally put in storage for future use. This *parking* is achieved by pointing the domain name at a dummy IP address that the domain owner does not actually control. This prevents errors, and attackers can use it to circumvent defenses against newly active domains, discussed above. Algorithms can detect parking by its pattern of switching the domain to and from IP addresses that are not under the control of the same organization. Like detecting flux, detecting parking requires a minimum amount of data on a domain.

Risks and Vulnerabilities

Like other security solutions, a DNS repository contains information about the enterprise and may be a high-value target for an attacker. In the past, attackers have been able to inject spurious DNS messages into passive DNS monitoring solutions; however, these attacks have been mitigated, and current best practices seem to prevent any corruption of the passive DNS data.

As mentioned above, properly configured passive DNS collection does *not* have a significant impact on the privacy of the users within the organization. See [DHS12] and [HS12].

Level of Adoption

Many security and network defense organizations, including the United States Computer Emergency Readiness Team (US-CERT), currently make use of passive DNS data. Industry provides both self-

service and managed solutions. For example, Farsight, Inc. (formerly the Security Information Exchange at the Internet Systems Consortium) provides a large corpus of live passive DNS information. Damballa provides a managed DNS solution that makes use of multiple DNS analysis techniques. OpenDNS's services, which are available to organizations and individuals, block certain categories of domains, including known malware and phishing sites.

Costs and Benefit Considerations

DNS is a nearly ubiquitous application with low data volumes. Organizations should be able to analyze the application data without drastic changes to their IT infrastructure, though some dedicated DNS analysis hardware likely would need to be procured.

DNS analysis provides a rich data set for non-signature-based detection because the rules and expectations of a DNS exchange and the content of the messages are well defined. A reasonable baseline of expected activity already exists, and behavior outside that baseline is easier to detect than for other internet applications. The analysis methods are also comparatively mature, having had more than five years' worth of demonstrated, successful non-signature-based detection methods using DNS data.

Policy Implications

Adoption: Given the effectiveness, low risk, and in some cases ease of deployment of DNS analytics, further adoption should be encouraged. Staffing and cost will vary according to the specific DNS analytic.

Research and development: Defenses and countermeasures are not static. While DNS analytics are effective, there has been and will continue to be an evolution of measures (e.g., fast flux). The need for ongoing research and development efforts to address countermeasures will remain.

IT policy: In some cases, IT policy changes would support more effective deployment of DNS analytics, for example, requiring all DNS requests to a traverse particular set of DNS resolvers to enable logging and their effective use as points of control.

3.5 List-Based Monitoring

Description

Attempted connections to destinations that are deemed higher risk can signal suspicious traffic. Blacklists contain IP addresses or domain names that have been determined to be malicious and should be blocked. Analysis of the resulting log, event, and netflow data provide analysts with insights on potentially compromised machines or other malicious activity. Whitelists contain IP addresses and domain names that have been determined to be safe, enabling administrators to validate destinations and minimize operational impacts. These solutions are closely related to the DNS analytics solution. While list-based monitoring has provided valuable capabilities to operators, it has weaknesses due to the time delay of adding a destination to the list and determining its level of accuracy.²

Methods are being developed to dynamically calculate risk scores for IP addresses and domains. A scoring approach allows administrators to create ranges associated to different sets of policy actions,

² See [AA13] and [JS13] for processes to create two types of whitelists.

expanding the simple black-or-white approach. Several vendors as well as the SEI and others in the research community are now implementing these methods [APD]. Information such as passive DNS data, malware analysis, domain registration data, and border gateway protocol (BGP) data is fed into the process. Passive DNS data shows actual resolutions that are taking place on live networks, though this data is time dependent because the IP addresses for a given domain name can change over time. Malware analysis directly links malicious code and a domain name. BGP information can show a broader linkage between an IP address and other networks that are known to host malicious content. All the available information is processed to provide a numerical value or risk score. New risk scores can be calculated frequently based on the update intervals of the inputted data, significantly increasing the accuracy of the list. Before the data can be processed, individual organizations need to determine the relative importance of the various inputs and the organization's overall risk tolerance. For example, some organizations will decide to block traffic destined for locations having a risk score below 80. Others may decide to route traffic to a sandbox environment if it scores between 65 and 80 but block anything below 65. The unique nature of each organization will influence how conservative or aggressive it needs to be. Risk scores can change over time, so the list needs to be dynamic. An IP address may be re-allocated to a new customer and so become a legitimate site, or a host may be overtaken by an intruder and begin distributing malware. Unlike most block lists, reputation lists are managed dynamically.

This methodology commonly uses a proxy approach to assess the risk of individual websites. As previously mentioned, several vendors such as Blue Coat, Websense, Palo Alto Networks, and McAfee offer reputation-based appliances to their enterprise customers. These providers expend significant effort to categorize internet sites, allowing customers to select the categories they wish to accept and deny. A crowd-sourced solution that incorporates information from a broader community is also available [web13]. A related technique is implementing a so-called speed bump, which raises an alert to the end-user and asks the user whether or not to proceed. Speed bumps make the user think twice before connecting and have proven effective in blocking automated requests from malicious code. While these solutions could be considered to be signature-based, the methods used to develop them are based on behavioral analysis.

Examples

Blocking uncategorized domains: Another technique blocks domains that have not been categorized by the web proxy vendor. This effectively blocks many new and suspicious domains by default. This is a very low-cost, effective option if an organization has already deployed a web proxy with site categorization.

Mitigating an infected machine: Consider the scenario of a user's workstation becoming infected with malicious code. The code is installed and executed, and then it tries to connect to its command-and-control system via the HTTP protocol. The machine tries to resolve the DNS hostname for the malicious server. The enterprise's DNS server receives the request and checks the name against its lists to find any matches. Instead of resolving the DNS request, the DNS server replies with an IP address to an internal server where analysts can further examine the requested resource. The request from the malicious code is blocked, and the user's workstation is not further compromised. Figure 7 illustrates the reputation list solutions.

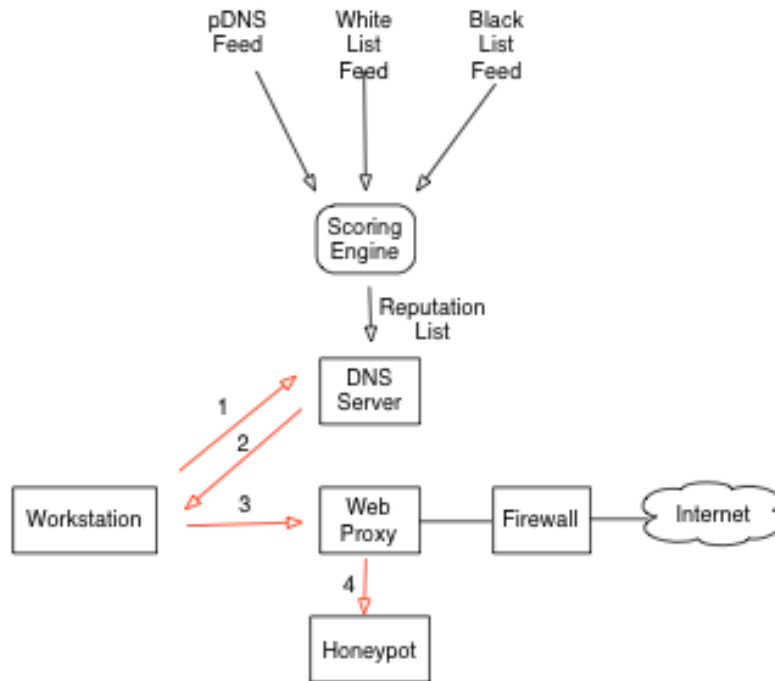


Figure 7: Reputation List Solution

Risks and Vulnerabilities

The operational risk associated with these lists includes an address or domain being incorrectly classified as malicious or benign. Lists can go stale if not updated and require localization. Attackers can compromise legitimate servers to avoid this protection until they are detected. The risk score data must be protected and received only from verified sources so an intruder cannot change values.

Costs and Benefit Considerations

The cost associated with these lists depends on the type of solution being implemented. Many lists are shared or obtained from external sources. Some are developed and maintained in-house. The reputation approach can be purchased or developed in-house with various levels of sophistication. Organizations can start with a small solution and grow it over time to lessen the risk, minimize costs, and demonstrate its value. This flexibility and ability to tailor its configuration to each organization's unique need makes list-based monitoring very attractive. The reputation approach, in particular, provides a dynamic capability to assess the risk of destinations using real data from operational networks. Also, the risks are calculated using a scale instead of a simple black-or-white determination of safety, providing enhanced fidelity and the ability to provide multiple sets of actions for mitigation based on risk score ranges.³ Because the foundation of the method is based on calculations, there is a substantial opportunity for automation that could minimize the organization's level of effort after deployment.

³ For instance, anything with a score above 75 is permitted, anything with a score between 55 and 75 requires additional authentication, and anything with a score below 55 is blocked.

Level of Adoption

Most enterprises and product vendors use list-based solutions. The mechanics for implementing lists are widely available. The data within the lists is inherently dynamic and will continue to change over time. Complementary techniques, such as the use of speed bumps, are growing in popularity.

Policy Implications

IT policy: DoD should begin monitoring and collecting passive DNS data at the enterprise level to provide increased data quality for reputation scoring and other related defense activities.

3.6 Network Profiling

Description

A network profile is an inventory of all assets on a network and the applicable data for each asset. For example, a profile of a small enterprise network may list a number of web servers, a web proxy, email servers, and a DNS server, along with the ports used by these assets and their average traffic volume. An up-to-date network profile provides a picture of a network's normal state, which can be compared to current network traffic to find assets that have recently been added or changed (broad AD). Analysts can also use profiling as part of their hunting operations to rule out expected traffic on a network, leaving a much smaller portion of unidentified traffic to analyze and to provide needed context for specific hosts. The network profile serves as a reality check for documentation and network maps that are often out of date and hard to manage across the multiple responsible organizations.

Network profiles can be developed in many different ways and to varying degrees of detail, but the basic process is straightforward. All active assets must be found, their traffic analyzed, and services listed. Additional profiling such as system fingerprinting can also be done. Scanning tools such as Nmap and Nessus can find assets and open ports, network management software can track assets, and traffic analysis tools such as SiLK or Wireshark can provide an accurate picture of assets, ports, and traffic volumes. Network profiling typically involves a mix of automatic and manual analysis and can use active or passive methods, depending on resources and policy. A basic profile can be completed and updated automatically.

Examples

The enterprise network profiled in the report *Network Profiling Using Flow* [WF12] benefitted in several ways from developing a profile. The process found two assets, one acting as an FTP server and one as a web mail server, that were previously undocumented and unknown to administrators. One web server hosting connections on port 443 (SSL) was found not to have the appropriate certificate installed and so was defaulting to unsecure connections. Analysis of traffic left over after profiling also revealed a workstation dual-homed to both the internet and the internal network, resulting in unacceptable risk to the network. Additionally, the new profile provided administrators with an accurate picture of assets for resource planning purposes.

Many network profiling training sessions [JW12], using small portions of business networks as case studies, have revealed that analysts gain a new, deeper understanding of the assets on the networks

they are assigned to protect. This new understanding leads to an enhanced ability to identify potential threats or misconfigurations on the network.

Risks and Vulnerabilities

As is true for all security methods, network profiling will not catch all malicious activity on a network. It is designed to provide snapshots of actual assets on a network. It is a useful tool for creating baselines for broad AD and separating normal traffic from junk and potentially malicious traffic. Some caveats are that a profile depends on coverage of sensors on the network, and services can run on unexpected, nonstandard ports and protocols.

Level of Adoption

Many administrators already use Nmap or similar tools to look for assets that should not exist on their network. However, Nmap can be slow on a large network, does not look at traffic volumes, does not look for changes from the previous scan, and in many cases is against policy. Most network administrators use a network monitoring program, which often comes as an add-on to purchased enterprise hardware and lists many active assets on the network. However, these programs are used for management purposes and are not typically designed for AD.

Profiling is currently being used in case studies and for training purposes.

Costs and Benefit Considerations

Many networks already have the tools in place to do network profiling. Automation and coordination of alerts may take time to set up. Analysis of unprofiled traffic will require expertise not held by typical network operations staff but well within the scope of a network defense analyst. As with any addition to a network security tool set, there is an added load on the network operations team, whose duty it is to make sure the tool set runs effectively. In this case, the added load should be minor.

Regular profiling has the added benefit of saving time during an incident. A security administrator with a firm understanding of the network and an up-to-date profile can complete an incident handling task more quickly and efficiently.

Policy Implications

Adoption: This capability should be integrated with asset management systems to provide security staff with more accurate and informative contextual information. Network profiling can be invaluable for those who need to understand networks to which they do not have access, for example, service providers who must have a picture of their customers' networks or those providing security for customer networks.

3.7 Deception

Description

Deception, as a strategy, implies misleading an adversary into acting in a way that suits the organization's overall goals, rather than the adversary's goals. The *misleading* approach systematically misrepresents (to the adversary, though likely not to benign users) the infrastructure and content in such a way that the adversary's actions are less productive and more visible than the

adversary intends. Deception methods implement such systematic approaches, with the approach varying depending on the method. Some deception methods, briefly outlined below, do not use signatures of adversary usage. Other deception approaches do use signatures, including tarpits (protocol response agents that obfuscate or conceal open communication channels on a computer network) [Bau10] and thin honeypots (virtual hosts that provide only enough functionality for initial response to common attacks) [Pro].

Deception may be applied at many levels of network architecture. At the highest level, diversification and load splitting among several service providers may obscure the network critical points of presence. This is commonly done to increase bandwidth, but by splitting traffic across multiple providers, network asymmetry may limit the ability of adversaries to interrupt or localize network services. Network services, particularly those most attacked and least integrated with protected information, may also be outsourced. The U.S. military has already shifted many of their public-facing web servers to cloud hosting, which makes their security a matter of contract rather than a matter for internal efforts. Attacks against those servers are not attacks against military infrastructures and are unlikely to yield increased access to military-operated services or computers. Other network services may also be outsourced, but coordination of outsourced services to obtain increased cost savings and security is a subject that requires further research, though some guidelines have already been published [AGM]. Even where specific servers are not outsourced, some network services may be moved to nonmilitary infrastructure via application hosting agreements, available commercially.

There has been some effort on generation of deceptive content [TJS12]. Creation of bodies of false information could be used to mislead persistent adversaries and reduce the confidence they might hold in information exfiltrated from organization networks. There are building blocks available in the open literature to systematically (and even automatically) construct such bodies of false information and to maintain the separation between false and true information, but further research and development would be required for an operational capability.

Examples

The most commonly cited network deception lies in the implementation of honeypots or honeynets. While the most minimal honeypots are often signature based, in the sense of providing responses expected by known attacks, higher-function honeypots [Pro] are full virtual servers that are implemented purely to waste adversary effort and to alert defenders to adversary activity. Honeynets are subnets of honeypots, designed to interact together but not link to organization resources, to more fully deceive adversaries as to their nature. Typically, outbound honeypot/honeynet-initiated contacts are blocked by surrounding infrastructure to limit the adversary's capability to exploit the honeypot or honeynet as a base for network attack.

Risks and Vulnerabilities

The liabilities of deception include adversary exploits of deceptive infrastructure (including entrapment issues and compromise of third-party information), separation of deceptive infrastructure and information from operational infrastructure and information without impacting either the effectiveness of the deception or the efficiency of operations, and fiduciary liabilities (including justification of the amount of effort expended on deception).

Level of Adoption

Deception solutions are not as common as many of the other solutions. The increased level of interaction with the attacker causes concern for many. Effectiveness measurement of deception is currently an unresolved research issue for network security. How to maintain network deception over time is also unknown currently. These liabilities and unknowns would need to be addressed prior to deployment of an operational capability.

Policy Implications

In general, specific policy considerations would need to be resolved prior to widespread deception as a defensive measure; specifically, how possible liabilities would be handled.

3.8 Anomaly Detection

Description

Over the last 25 years, the academic literature has explored anomaly detection (AD) as applied to intrusion detection. In Jones and Moitra [JM13], we provided a selective survey of the large body of AD literature and identified key terms and categories, important authors, and trends. We used the following definitions of AD, adapted from Chandola, Banerjee, and Kumar [CBK09]:

anomaly detection: “Anomaly detection [is] the problem of finding patterns in data that do not conform to expected behavior.”

anomaly: “These non-conforming patterns are often referred to as anomalies, outliers, discordant observations, exceptions, aberrations, surprises, peculiarities or contaminants.”

Many IDS/IPS vendors have added some form of AD within their devices to enhance signature-based detection. Examples include McAfee, Sourcefire, and Enterasys. The sensors monitor network traffic and develop a baseline of the normal traffic patterns. Using statistical calculations, they can identify traffic that diverges from the norm and notify the security administrator. Despite their automated capabilities, these products usually require significant tuning to reduce the number of false positives and focus their monitoring on desired activities. Because most enterprise networks experience frequent changes, this creates a recurring workload, a problem known as *concept drift*. Some products implement a very basic form of AD instead of a more sophisticated statistical solution, resulting in a wide range in effectiveness across the market. Successful deployments seem to apply AD to a narrow problem space and include domain-specific functionality to filter statistical anomalies, similar to an expert system. This allows the solution to focus on very specific things, reducing false positives.

There is also an increased use of wireless intrusion detection devices to identify unauthorized network access points. This varies across organizations based on their individual wireless policies. This area is still developing given the tremendous growth of handheld devices. Security management systems have not kept pace, and integration of data is a gap.

Examples

Data loss scenario: In a basic configuration, an AD IDS sensor would monitor user traffic going to the internet. The organization could be concerned about the loss of intellectual property and want to

monitor for suspicious file transfers to external sites. The sensor would detect the anomaly and alert analysts when a user's file transfer activity is significantly different than a reference group of users. This is in contrast to a DLP solution that might rely on hard-coded rules and thresholds that are not site- or user-specific. The SIEM would alert security analysts to possible incidents requiring investigation.

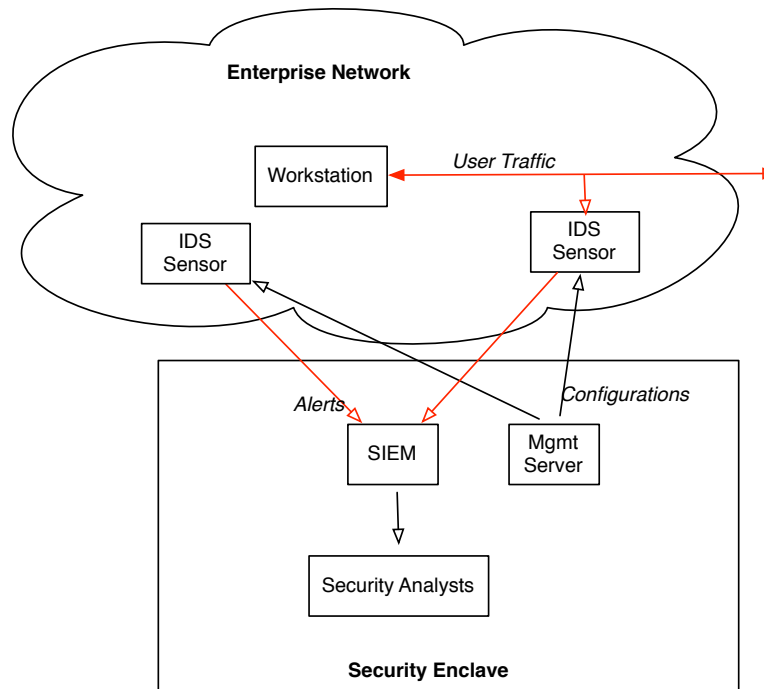


Figure 8: Intrusion Detection System Example

The Cuckoo's Egg: In 1986 Cliff Stoll [Sto89] began investigating what appeared to be a minor accounting error in a billing system and wound up catching a hacker selling secrets to the KGB. This fits the definition of AD: a small amount of data in the billing records did not conform to expectations.

Cross-domain collaborative anomaly detection: Boggs and colleagues [BHSS11] describe a system that performs payload inspection on web traffic, identifies potential attacks, shares the information across sites, and correlates anonymized results to reduce false positives. This is a far more complex example of AD, but it also fits the definition. The system had a training period of several weeks to determine normal or expected behavior, and it later alerted on items that did not fall within that range.

These examples illustrate the diversity of AD. It can be simple and manual or complex, involving multiple systems and complex algorithms. The bottom line is that AD systems must have some expectation and understanding of normalcy and a way to detect things that fall outside that definition.

Risks and Vulnerabilities

The sensors are deployed at various network locations within the enterprise and become targets themselves. They are not protected within an enclave with additional controls. Most solutions include

a centrally managed configuration that could be altered either intentionally or by accident, affecting the performance of the sensors.

Anomalous behavior is not always malicious: Gates and Taylor [GT06] list assumptions underlying the use of AD for intrusion detection. The most notable assumption is that attack behavior is somehow different from normal traffic, but attackers might hide their activities in normal traffic. The use of manual versus automated AD does not change this problem.

The need for actionable reporting: “Anomaly detection systems face a key challenge of transferring their results into actionable reports for the network operator. In many studies, we observe a lack of this crucial final step, which we term the *semantic gap*.... The next step then needs to interpret the results from an operator’s point of view — ‘*What does it mean?*’ Answering this question goes to the heart of the difference between finding *abnormal activity* and *attacks*.” [SP10]. Answering this question is both necessary and difficult.

Level of Adoption

One result of the industry survey confirmed that machine-learning-based AD sensors for intrusion detection are not widely deployed in operational environments. One organization reported that they “have machine learning and data mining experts, but we have largely abandoned it [for internal intrusion detection].... Statistical modeling has just not been effective. The problem is that internal security does not have labeled data.” Another respondent with expert knowledge of machine learning reported “trying out a number of anomaly detection techniques from the open literature [and being disappointed] every time.”

Our findings from the interviews confirm the observations, highlighted by the SEI’s Jones and Moitra [JM13], of Sommer and Paxson [SP10]. They examine “the surprising imbalance between the extensive amount of research on machine learning-based anomaly detection pursued in the academic intrusion detection community, versus the lack of operational deployments of such systems.” They identify five challenges to the successful application of machine learning and AD to the problem domain of intrusion detection: the difficulties of outlier detection in machine learning, the high cost of errors, the need for actionable reports, the diversity of network traffic, and difficulties with evaluation. These challenges help explain why these techniques are not widely deployed.

Costs and Benefit Considerations

A sensor-based AD solution requires a significant hardware deployment because the sensors are distributed across the network. Obviously, this depends on the specific enterprise and its tailored application. In addition, costs will include the labor required to configure and tune the solution before it can become operational. After deployment, operational costs will include the analysts’ time to respond to alerts, including any false positives as well as continual tuning as the enterprise changes. These operational costs have been larger than expected by many organizations. These costs are hampering the perceived value of the solution, and many organizations are not finding that the solution significantly increases their capabilities. A benefit of the sensor-based solution is that it monitors live traffic, so it takes less time to detect an issue than manual solutions that operate on resulting event and log data.

4 Summary and Recommendations

Discussions with the operational organizations provided a good view into the effectiveness of the various solutions. The cybersecurity field is still maturing, and these organizations continue to trial products and approaches to address gaps and improve capabilities. The ones that prove to be effective continue to be supported, and the non-effective solutions are retired. The technologies and techniques vary in their level of maturity and provide various opportunities for future innovation. It is important to note that the majority of the organizations we talked to were large enterprises from critical infrastructure sectors. These organizations seem to be early adopters, and there is probably a significant opportunity for additional deployments within other portions of the community.

Figure 9 summarizes each solution based on the evaluation factors that were considered. The factors are grouped into benefits, maturity, level of adoption, cost considerations, and risks. The values are qualitative assessments based on the research, interviews, and the SEI's experience. The matrix also contains information on policy options to consider. Policy types include adoption, education, research and development, staffing, and IT policy. Each column shows the relative values, providing a comparison of the solutions for each factor.

Solution	Benefits	Maturity	Adoption	Cost Considerations			Risks				Policy Recommendations				
	Effectiveness	Level of Maturity	Level of Adoption	Staffing Costs	Infrastructure Costs	Level of Automation	Direct Attack	High value target	Privacy Issues	Evasion	Adoption	Education	R&D	Staffing	IT Policy
Hunting Operations	High	low	medium	high	medium	medium	no	yes	yes	yes	yes	yes	yes	yes	yes
Security Information & Event Management	low	high	high	high	high	high	yes	yes	yes	yes					
Content Inspection	High	medium	medium	medium	low	high	yes	yes	yes	yes	yes		yes		
DNS Analytics	High	low	medium	low	low	high	no	no	no	yes	yes		yes		yes
List-based Monitoring	high	medium	high	low	low	high	no	no	no	yes					yes
Network Profiling	medium	low	low	low	medium	high	no	yes	no	yes	yes				
Deception	high	low	medium	medium	low	low	yes	no	yes	yes					yes
Anomaly Detection Systems	low	medium	low	high	high	high	no	no	yes	yes					

Figure 9: Solution Summary

Hunting operations is a major focus for many of the organizations that we interviewed. It is a relatively new approach that is being driven to address existing challenges with SIEM solutions as well as the technology developments associated with “big data” and data analytics. Though it is manual, this approach maximizes the ability of analysts and often results in discovery of indicators that make further automated detection possible. Advances in tool sets and development of best practices should be encouraged to expedite the enhancement of this solution.

DNS analytics, list-based solutions, and network profiling will continue to provide value. They do not require significant technology development or staff, and operational experience and sharing of techniques and specific lists will foster continual improvement.

Content inspection and deception techniques such as honeypots require ongoing research and development of new techniques to keep pace with the latest attacks. The infrastructure and processes will only remain effective if there is continued emphasis on analysis of threats and attack scenarios. Ongoing support is needed for the continued analysis of the changing threat landscape and development of new discovery techniques.

Cost continues to be an important consideration. Costs include the infrastructure costs associated with the hardware, software, and maintenance; the deployment labor; and the ongoing operational support labor. Most of the interview participants cited funding constraints that influence their strategy. Hunting operations, SIEM systems, and AD sensors have high costs, which is why their effectiveness has received high scrutiny. DNS analytics, list-based techniques, and deception configurations generally have low costs, enabling organizations to quickly test capabilities. The cost equation combined with the overall effectiveness drive whether solutions will be adopted and matured over time.

Recommendations

Promote the adoption of hunting operations, sandboxing, DNS analytics, and network profiling through policy guidance and funding allocations. Give engineering and operational groups the support they need to deploy the associated technologies and the staff they require for these solutions. Expedite the learning process by rotating staff to share knowledge. Consider establishing a framework to support distributed security operation centers using shared infrastructure and technologies to decrease deployment time and cost.

Develop a human resource and staffing strategy to support hunting operations. Scalability of staffing needs to be addressed by increasing funding, enhancing education, instituting training programs, attracting individuals to the field, and other steps to initiate an order-of-magnitude increase in the workforce. Security operation groups need the funding required to staff an effective analysis unit that can address salary and career path issues. University curriculums and position training plans need to be enacted to develop the skill base.

Fund research and development efforts supporting advancement of sandbox techniques, DNS analytics, and automation of hunting operations. Support the development of knowledge of specific attack scenarios to respond to the ever-changing security landscape. Given the trend toward application-specific attacks, work with the vendor community to understand application architectures and vulnerabilities. Provide environments to develop and test new techniques and analytics. Collaborate to share information on emerging threats and discovery techniques. Regarding hunting operations, monitor the scalability capabilities of hardware and software platforms to identify potential issues early, and continue to improve solution efficiencies. Research methods for sharing analytics across organizations, taking into account the uniqueness of organizations. Promote the use of pilot programs to encourage the development of analytic tools and expert systems that provide more value to the analyst.

Integrate researchers into operational environments to better inform efforts and increase likelihood of operational effectiveness. Promote two-way collaboration between R&D and operational staff, creating a stronger understanding of challenges, requirements, and innovation.

Augment IT policy on DNS designs to support the implementation of DNS analytics and the collection of passive DNS information at the enterprise level. Review the current architecture guidance and status of implementation. Support the deployment of the additional services required to support the monitoring, collection, and analysis of DNS-related information addressing performance and scalability needs. Develop a mechanism to allow various security operations groups supporting the enterprise networks to access the data.

Clarify policies regarding deception techniques to address concerns regarding potential liabilities. The liabilities include adversary exploits of deceptive infrastructure (including entrapment issues and compromise of third-party information), separation of deceptive and operational infrastructures, and fiduciary liabilities.

Assist in applying appropriate policies to address the security requirements for hunting operations and “big data” solutions that centralize extensive amounts of sensitive information from across the enterprise. Leverage guidance and solutions from the intelligence community as appropriate. Evaluate the use of cloud environments and the associated contract and service level agreement direction required to provide appropriate security of the information and services. Promote the implementation of granular access controls based on security policies and risk management considering the large data set. Work with the open source initiatives and the commercial marketplace to promote advancement of security features within “big data” offerings.

5 Closing

This study incorporated the views from academic literature and the operational experience of representatives from some of the critical infrastructures. We explored the use of non-signature-based solutions for intrusion detection, identifying techniques that are proving effective and potential trends for the future. The results of our survey, interviews, and analysis led to the following conclusions.

No single technical solution is enough.

We identified no single technical solution to address the problem of discovering new, unknown, and sophisticated attacks.

Hunting operations find new attacks.

Analyst teams, combined with “big data” technologies, are being formed to sift through the large quantities of available data to detect incidents. The general process is to mine data, understand anomalies, and then codify an analytic around the indicators identified. This also enables automation.

Signature solutions are still important.

Signature-based solutions are still important. They catch the older, known, less sophisticated attacks, reducing noise so analysts can focus on newer, unknown, and sophisticated attacks. Also, signature-based solutions can help automate detection of attacks identified via non-signature-based approaches.

Commercial solutions are dissatisfying.

The study participants expressed a high degree of dissatisfaction with commercial solutions in the areas of attack detection, integration, scalability, and perceived value. Open source software and custom storage solutions are being used to address some of these problems.

Additional skilled analysts are needed.

While automation may assist human understanding, it will not fully replace it. Analysts will remain an important part of the solution to discovering new, unknown, and sophisticated attacks.

6 Acknowledgments

We would like to thank the following people who made significant contributions to this report:

- the Office of the Undersecretary of Defense for Policy, for sponsoring the work
- the anonymous contributors who agreed to be interviewed and provide operational input
- anonymous contributors who provided input to our survey of academic AD work
- authors of works cited, with special thanks to Robin Sommer and Vern Paxson
- Soumyo Moitra, for his help researching and writing the summary of AD
- Rich Friedberg, Greg Shannon, and Brian Lindaur for providing important direction and feedback
- Paul Ruggiero for technical editing support
- the many others who contributed

Appendix A: Interview Questions

The SEI used the following questions to loosely guide the discussions in the interview phase of the project.

1 Current signature-based solutions

- What signature-based detection solutions are you using currently?
 - e.g., IDS, AV, etc.
- What gaps/deficiencies do you see in signature-based solutions?
- What is the level of maturity of this segment?

2 Current anomaly-detection-based solutions

- What anomaly detection solutions are you examining?
 - Vendors
 - Features
 - Analytics
- What is the level of maturity of this segment?

3 Other non-signature-based solutions?

- Besides anomaly detection, what other non-signature solutions are you considering?
 - Deception based (e.g., honeypots)
 - Behavioral analysis (scanning, quarantine, and analysis of malware)
 - Others?
- Are you emphasizing acquiring or developing solutions?
- What is the level of maturity of this segment?

4 Major challenges for non-signature-based solutions

- Describe some of the major challenges that you foresee in deploying non-signature-based solutions.
 - scalability?
 - normalization?
 - correlation?
 - staffing?
 - false positive rates?

5 The future

- What longer-term needs will you want the security marketplace to address?
- How do you see your solutions evolving over the next several years?

6 Measures of success

- What are your measures of success? In what order?
 - Prevention of successful attack?
 - Detection of attacks missed by signature-based solutions?
 - Cost reduction?
 - Risk reduction?
 - Return on investment? (How would you calculate it?)
 - Other?

Appendix B: Taxonomies

At the outset of this project, we believed that one of the major evaluation tasks would be to classify and compare various commercial and noncommercial tools. We believed there would be a need for a list of common terms and taxonomies to classify systems. While we still believe that these are needed for direct comparisons of specific systems, the major solutions identified in this report did not require such detailed comparisons.

For those wishing to perform such direct comparisons, we refer you to the following:

- “Glossary of Key Information Security Terms” [\[Kis10\]](#)
- “Guide to Intrusion Detection and Prevention Systems” [\[SM12\]](#)
- We adopted a taxonomy of anomaly detection features from Davis and Clark [\[DC11\]](#) and Estevez-Tapiador, Garcia-Teodoro, and Diaz-Verdejo [\[ETGTDV04\]](#).
- We adopted a taxonomy of anomaly detection techniques from Chandola, Banerjee, and Kumar [\[CBK09\]](#).

Finally, we developed the following taxonomy for classifying defensive cybersecurity solutions, including anomaly detection systems.

Cybersecurity solutions

- Defensive
 - Access control
 - Network based
 - Example: Checkpoint Firewall-1
 - Signature-based threat detection
 - Intrusion detection systems
 - Example: Snort
 - Antivirus system
 - Example: Symantec
 - Example: McAfee
 - [White|Black]listing
 - Non-signature-based threat detection
 - Intrusion detection systems
 - Specification based
 - Firewall
 - Example: Checkpoint Firewall-1
 - Example: Cisco ASA
 - Behavioral detection systems
 - Example: BotHunter
 - Misuse detection systems
 - Example: BotHunter
 - Anomaly detection systems
 - Anomaly detection features
 - Anomaly detection techniques
 - Small
 - Medium
 - Large

- Deception
- Visualization
- Reputation
- Sandbox execution
- Game-theoretic
- Hybrid system
- Components common to multiple approaches
 - Techniques
 - Noise reduction
 - Cross-site correlation
 - Single-source data analysis
 - Multisource data analysis
 - Actionable reporting
 - Understand context
 - Focus on assets
 - Reduce scope
 - Evaluation criteria
 - Data needs
 - Common data sets
 - Labeled data sets
 - Realistic data sets
 - Scale
 - Level of analysis
 - Time scale of analysis
 - Size of network, number of hosts
 - Performance considerations
 - Separate processing and analysis layers

Appendix C: Scalability of Staffing

Some of the solutions discussed have a significant dependency on the availability of trained, available analysts. Today, such staffers are scarce, and their distribution seems to be based on the size of the organization. For instance, larger, well-funded organizations are able to establish a security analytics unit, but they are experiencing challenges with hiring and retaining qualified staff. Smaller organizations are not able to support the same organizational model and need to rely on other mechanisms. They have a stronger dependency on deploying technology solutions to automate functionality. This broad range of environments creates the need for a variety of solutions. While today's enterprises have taken steps forward, their solutions are not the optimal, scalable solutions for the future.

Scalability of staffing needs to be addressed by increasing funding, enhancing education, instituting training programs, attracting individuals to the field, and taking other steps to initiate an order-of-magnitude increase in the workforce. Security operation groups need the funding required to staff an effective analysis unit. These analysts usually possess more skills and experience and so require higher salaries. Government should continue its coordination activities with universities to support the development of effective curricula and individual classes. Training programs need to be available for staff already in the workforce. The government can also be a key advocate for the field and attract people to it. Establishing and maintaining a positive culture around cyber defense functions can persuade smart, energetic students to pursue a related career. For instance, the idea of working at NASA has always interested a large group of young people, leading them to studies related to science and math. A similar atmosphere needs to be created for cyber defense. Lastly, individuals need to feel that they have the opportunity to make a difference and add value. They need to be positioned and equipped for success.

Despite increasing the workforce, additional automated capabilities will be required to keep pace with the exponential growth of attacks. Detection and response activities need to operate at network speed whenever possible. The current trend toward data mining and hunting operations by analysts is one of the best available options at present..

Appendix D: Noise Reduction Techniques

This appendix lists solutions that can be classified as *noise reduction* techniques. These help focus attention on relevant subsets of network activity, reducing the amount of information that must be analyzed.

Constrain User Behavior

Constraining user behavior is one way to reduce noise. It simultaneously reduces the actions that are possible, reducing the likelihood of successful attacks, and makes violations stand out.

An organization's ability to implement high levels of constraint depends on the organization's risk tolerance. For example, the financial sector places strong emphasis on integrity and confidentiality and can be more restrictive, whereas telecommunication carriers place more emphasis on availability and must be relatively unconstrained because their customer base expects open and unfiltered access to internet resources. The organization's relationship to its users is a key factor in determining the level of control that can be implemented. Employees in the financial sector can expect to tolerate more strict controls, whereas telecommunications customers expect fewer controls.

Constraining user behavior can include measures such as proxies, restricting allowable protocols, network segmentation and isolation, disallowing execution of unknown code (application white-listing), and disallowing personal and mobile devices. The key idea is to limit the number of services, servers, protocols, users, and other components that have to be watched, in order to increase understanding of expected behavior and make anomalies stand out.

NIST [[NA12](#)] and SANS [[SAN](#)] both provide examples of security controls.

Eliminate the Known Good

Several interviewees told the SEI that it is easier to find things if you reduce the area you have to search, or "throw out the hay to find the needles." This is similar to a log analysis process that Marcus Ranum described as "artificial ignorance" [[Ran](#)]. This report describes several solutions that support the idea of eliminating the known good.

Reduce Attack Surface

Reducing overall attack surface reduces noise. This includes methods such as exposing fewer systems to the public internet and exposing fewer services. Proxies, firewalls, and practices such as disabling unused services all reduce attack surface. See Manadhata and Wing [[MW](#)] for ideas on reducing attack surface.

Defense in Depth

Defense in depth reduces noise. For example, scanning activity detected by the firewall at an organization's internet gateway can probably be ignored, but indicators of successful compromise on the most sensitive internal networks and systems require attention. See May, Hammerstein, Mattson, and Rush [[MHMR06](#)] for a thorough discussion of defense in depth.

Implement Best Practices

While best practices are not a complete solution, they can reduce noise. For example, signature-based solutions such as AV and IDS can be used to identify known, common, or old attacks to allow analysts to focus on identifying new, unknown, and sophisticated attacks.

Threat Intelligence

Threat intelligence reduces noise by helping focus analysis activities. For example, information that a competitor or nation-state is interested in a particular unit of an organization's operation allows analysts to focus on that unit and possibly the identified adversary. But keep Donald Rumsfeld's "unknown unknowns" [[Rum02](#)] in mind: lack of threat intelligence does not mean lack of threat.

Indicator Sharing

Currently deployed indicator-sharing solutions reduce noise by helping focus analysts on known attacks. While they do not help analysts discover unknown attacks, they can shorten the process of identifying known malicious activity. See Collective Intelligence Framework (CIF) for one example of threat sharing [[You11](#)].

Whitelists, Blacklists, and Reputation Lists

Whitelists, blacklists and reputation lists can all be used to help focus analysis. See section 3.5, List-Based Monitoring.

Behavioral Analysis

Behavioral analysis can reduce noise by focusing attention on suspicious behaviors, both by internal staff as well as external entities. Several participants in the survey reported successful use of sandbox execution of suspect software as a way to evaluate its behavior and determine if it was malicious. See Section 3.3, Content Inspection, Section 3.4, DNS Analytics, and Section 3.7, Deception.

Appendix E: Understanding What Should Be Happening

This appendix lists solutions and techniques that increase understanding of *well-defined behavior*, which helps analysts understand what *should* be happening.

Use of Well-Defined Protocols

Most of the network protocols in use today have well-defined specifications that describe, at a protocol level, what *should* be happening. A large amount of work has been aimed at detecting protocol anomalies. See for example Sekar et al. [[SGF](#)] and Paxson [[Pax99](#)].

Network Profiling

Network profiling [[WF12](#)], described more fully in Section 3.6, provides a way to baseline networks, increase understanding, and detect changes.

Analysis in Context

Proper interpretation of the data requires sufficient context. Analysis done at large network consolidation points (e.g., enterprise internet gateways) is less effective at finding new or sophisticated threats due to reduced knowledge of the target environment. Further, analysts who do not have a strong understanding of acceptable or normal business processes have difficulty determining whether an anomaly is malicious. Simpler environments with smaller, well-defined processes make it easier to spot deviations and determine the proper mitigation and response measures.

Increase Security Controls

Implementing security controls such as those recommended by NIST [[NA12](#)] and SANS [[SAN](#)] has the double effect of increasing understanding of what should be happening and of what is happening. For example, a firewall with a default deny policy that allows inbound web traffic specifies what *should* (or is permitted) to happen. If, in addition, the firewall logs all other failed access attempts, the logs give an indication of what *is* happening. Both these may be useful in identifying unknown attacks.

Well-Defined Security Program

Having a well-defined and integrated security program that starts with clearly defined mission objectives is a foundational step toward improving an organization's ability to detect incidents. The types of policies and business processes influence the type of activity on a network. When analysts clearly understand the mission objectives and allowed and expected user behavior, they will be better able to detect malicious behavior. See the Critical Success Factors (CSF) [[CSWW04](#)], developed as part of the CERT® Resilience Management Model (CERT®-RMM) [[CER13](#), [CAW10](#)], which provide one way to categorize observable and measurable activity that is known to be consistent or inconsistent with an organization achieving its mission.

* CERT® is a registered mark owned by Carnegie Mellon University.

Bibliography

- [SAN] SANS. CSIS: 20 Critical Security Controls - Version 4.1. <http://www.sans.org/critical-security-controls/guidelines.php>, Accessed: 2013-08-01.
- [Pro] HoneyNet Project. The HoneyNet Project. <http://www.honeynet.org/about>, Accessed: 2013-04-20.
- [MW] Pratyusa K. Manadhata and Jeannette M. Wing. Attack Surface Measurement. <http://www.cs.cmu.edu/~pratyus/as.html>, Accessed: 2013-06-22.
- [Ran] Marcus J. Ranum. Artificial Ignorance: How-To Guide. http://www.ranum.com/security/computer_security/papers/ai/, Accessed: 2013-06-14.
- [Den87] Dorothy E. Denning. An intrusion-detection model. *Software Engineering, IEEE Transactions on*, (2):222-232, 1987. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1702202.
- [Sto89] Clifford Stoll. *The cuckoo's egg*, volume 1. Doubleday New York, 1989.
- [MF90] Roy A Maxion and Frank E Feather. A case study of ethernet anomalies in a distributed computing environment. *Reliability, IEEE Transactions on*, 39(4):433-443, 1990. <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=58721>.
- [Che92] Bill Cheswick. An Evening with Berferd in which a cracker is Lured, Endured, and Studied. In *Proc. Winter USENIX Conference, San Francisco*, 1992. <http://web.cheswick.com/ches/papers/berferd.pdf>.
- [FSM93] Frank Feather, Dan Siewiorek, and Roy Maxion. Fault detection in an ethernet network using anomaly signature matching. In *ACM SIGCOMM Computer Communication Review*, volume 23, pages 279-288. ACM, 1993.
- [KS94] Sandeep Kumar and Eugene H Spafford. A pattern matching model for misuse intrusion detection. 1994. <http://docs.lib.purdue.edu/cgi/viewcontent.cgi?article=2169&context=cstech>.
- [IKP95] Koral Ilgun, Richard A. Kemmerer, and Phillip A. Porras. State transition analysis: A rule-based intrusion detection approach. *Software Engineering, IEEE Transactions on*, 21(3):181-199, 1995. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=372146.
- [HL98] John D Howard and Thomas A Longstaff. A common language for computer security incidents. *Sandia Report: SAND98-8667, Sandia National Laboratories*, http://www.cert.org/research/taxonomy_988667.pdf, 1998. http://www.cert.org/research/taxonomy_988667.pdf, Accessed: 2013-06-14.
- [Pax99] Vern Paxson. Bro: a system for detecting network intruders in real-time. *Computer networks*, 31(23):2435-2463, 1999.
- [McH00] John McHugh. Testing intrusion detection systems: A critique of the 1998 and 1999 DARPA intrusion detection system evaluations as performed by Lincoln Laboratory. *ACM transactions on Information and system Security*, 3(4):262-294, 2000.
- [Axe00] Stefan Axelsson. The base-rate fallacy and the difficulty of intrusion detection. *ACM Transactions on Information and System Security (TISSEC)*, 3(3):186-205, 2000. <http://dl.acm.org/citation.cfm?id=357849>.
- [PES01] Leonid Portnoy, Eleazar Eskin, and Sal Stolfo. Intrusion Detection with Unlabeled Data Using Clustering. In *In Proceedings of ACM CSS Workshop on Data Mining Applied to Security (DMSA-2001)*, pages 5-8, 2001. <http://www.cs.ucsd.edu/~eeskin/papers/cluster-ccsdmsa01.ps>.
- [LX01] Wenke Lee and Dong Xiang. Information-theoretic measures for anomaly detection. In *Security and Privacy, 2001. S&P 2001. Proceedings. 2001 IEEE Symposium on*, pages 130-143. IEEE, 2001. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=924294.
- [SGF*02] R Sekar, Ajay Gupta, James Frullo, Tushar Shanbhag, Abhishek Tiwari, Henglin Yang, and Sheng Zhou. Specification-based anomaly detection: a new approach for detecting network intrusions. In *Proceedings of the 9th ACM conference on Computer and communications*

- security*, pages 265-274. ACM, 2002.
- [Rum02] Donald Rumsfeld. DoD News Briefing - Secretary Rumsfeld and Gen. Myers, February 2002. <http://www.defense.gov/Transcripts/Transcript.aspx?TranscriptID=2636>, Accessed: 2013-04-21.
- [AGM⁺03] Julia Allen, Derek Gabbard, Christopher May, Eric Hayes, and Carol Sledge. Outsourcing managed security services. Technical report, Software Engineering Institute, CERT Coordination Center, 2003. <http://www.cert.org/archive/pdf/omss.pdf>.
- [DDS⁺04] Niles Dalvi, Pedro Domingos, Sumit Sanghai, Deepak Verma, et al. Adversarial classification. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 99-108. ACM, 2004.
- [ETGTDV 04] Juan M Estevez-Tapiador, Pedro Garcia-Teodoro, and Jesus E Diaz-Verdejo. Anomaly detection methods in wired networks: a survey and taxonomy. *Computer Communications*, 27(16):1569-1584, 2004. <http://www.sciencedirect.com/science/article/pii/S0140366404002385>.
- [GC04] Fanglu Guo and Tzi-cker Chiueh. Traffic Analysis: From Stateful Firewall to Network Intrusion Detection System. *RPE Report*, 2004. <http://www.ecsl.cs.sunysb.edu/tr/TR164.pdf>.
- [Heo04] Heo. Abnormal traffic detection for network intrusion detection. 2004
- [CSWW04] Richard A Caralli, James F Stevens, Bradford J Willke, and William R Wilson. The critical success factor method: establishing a foundation for enterprise security management. Technical report, Software Engineering Institute, 2004. <http://www.cert.org/archive/pdf/criticalsuccessfactors0407.pdf>, Accessed: 2013-06-20.
- [WKL⁺05] Li Wei, Nitin Kumar, Venkata Lolla, Eamonn Keogh, Stefano Lonardi, and Chotirat Ratanamahatana. Assumption-free anomaly detection in time series. In *Proceedings of the 17th international conference on Scientific and statistical database management*, pages 237-240. Lawrence Berkeley Laboratory, 2005
- [BCJ⁺05] Michael Bailey, Evan Cooke, Farnam Jahanian, Niels Provos, Karl Rosaen, and David Watson. Data reduction for the scalable automated analysis of distributed darknet traffic. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, pages 21-21. USENIX Association, 2005.
- [MHMR06] Christopher J May, Josh Hammerstein, Jeff Mattson, and Kristopher Rush. Defense in Depth: 6] Foundation for Secure and Resilient IT Enterprises. Technical report, Software Engineering Institute, CERT Coordination Center, 2006.
- [GT06] Carrie Gates and Carol Taylor. Challenging the anomaly detection paradigm: a provocative discussion. In *Proceedings of the 2006 workshop on New security paradigms*, pages 21-29. ACM, 2006. <http://www.nspw.org/papers/2006/nspw2006-gates.pdf>.
- [Moo06b] David Moore. Anomaly Sampling (bringing diversity to network security). In *FloCon 2006 Proceedings*, Pittsburgh, PA, USA, 2006. CERT. http://www.cert.org/flocon/2006/presentations/anomoly_sampling1006_ppt.pdf.
- [Moo06a] David Moore. Anomaly Sampling. In *FloCon 2006 Proceedings*, Pittsburgh, PA, USA, 2006. CERT. http://www.cert.org/flocon/2006/presentations/anomoly_sampling0706_ppt.pdf.
- [MN06b] Ron McLeod and Vagishwari Nagaonkar. Anomaly Detection Through Blind Flow Analysis Inside a Local Network. In *FloCon 2006 Proceedings*, Pittsburgh, PA, USA, 2006. CERT. http://www.cert.org/flocon/2006/presentations/blind_flow_analysis2006_ppt.pdf.
- [MN06a] Ron McLeod and Vagishwari Nagaonkar. Anomaly Detection Through Blind Flow Analysis Inside a Local Network. In *FloCon 2006 Proceedings*, Pittsburgh, PA, USA, 2006. CERT. http://www.cert.org/flocon/2006/presentations/anomoly_detect2006.pdf.
- [Bin06b] James R. Binkley. Anomaly-based BotServer (and more!) Detection. In *FloCon 2006 Proceedings*, Pittsburgh, PA, USA, 2006. CERT. http://www.cert.org/flocon/2006/presentations/botserver2006_ppt.pdf.
- [Bin06a] James R. Binkley. Anomaly-based Botnet Server Detection. In *FloCon 2006 Proceedings*, Pittsburgh, PA, USA, 2006.

- CERT. <http://www.cert.org/flocon/2006/presentations/botnet0606.pdf>.
- [Gol06] Josh Goldfarb. Identifying Anomalous Network Traffic Through the Use of Client Port Distribution. In *FloCon 2006 Proceedings*, Pittsburgh, PA, USA, 2006.
CERT. http://www.cert.org/flocon/2006/presentations/clientport_dist1205.pdf.
- [WCS06] Ke Wang, Gabriela Cretu, and Salvatore Stolfo. Anomalous payload-based worm detection and signature generation. In *Recent Advances in Intrusion Detection*, pages 227-246. Springer, 2006. http://link.springer.com/chapter/10.1007/11663812_12.
- [SPM*06] Fritz Schneider, Niels Provos, Raphael Moll, Monica Chew, and Brian Rakowski. Phishing protection design documentation, 2006. http://wiki.mozilla.org/Phishing_Protection:_Design_Documentation.
- [EL06] E Earl Eiland and Lorie M Liebrock. An application of information theory to intrusion detection. In *Information Assurance, 2006. IWIA 2006. Fourth IEEE International Workshop on*, pages 16-pp. IEEE, 2006. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1610005.
- [Orm07] Tavis Ormandy. An empirical study into the security exposure to hosts of hostile virtualized environments. In *Proceedings of CanSecWest Applied Security Conference*. Citeseer, 2007.
- [GPY*07] Guofei Gu, Phillip Porras, Vinod Yegneswaran, Martin Fong, and Wenke Lee. Bothunter: Detecting malware infection through ids-driven dialog correlation. In *Proceedings of 16th USENIX Security Symposium on USENIX Security Symposium*, page 12. USENIX Association, 2007. http://static.usenix.org/event/sec07/tech/full_papers/gu/gu.pdf.
- [GPCR07] Sujata Garera, Niels Provos, Monica Chew, and Aviel D Rubin. A framework for detection and measurement of phishing attacks. In *Proceedings of the 2007 ACM workshop on Recurring malware*, pages 1-8. ACM, 2007.
- [PP07] Animesh Patcha and Jung-Min Park. An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Computer Networks*, 51(12):3448-3470, 2007. <http://www.sciencedirect.com/science/article/pii/S138912860700062X>.
- [BS08] Wim Biemolt and Werner Schram. Automatic anomaly detection using NfSen. In *FloCon 2008 Proceedings*, Pittsburgh, PA, USA, 2008.
CERT. http://www.cert.org/flocon/2008/presentations/biemolt_flocon2008.pdf.
- [Col08] Michael Collins. Attack Reduction and Anomaly Modeling in Popularly Targeted Protocols. In *FloCon 2008 Proceedings*, Pittsburgh, PA, USA, 2008.
CERT. http://www.cert.org/flocon/2008/presentations/collins_FloCon2008.pdf.
- [KI08] Tsuyoshi Kondoh and Keisuke Ishibashi. Identifying Anomalous Traffic Using Delta Traffic. In *FloCon 2008 Proceedings*, Pittsburgh, PA, USA, 2008.
CERT. http://www.cert.org/flocon/2008/presentations/kondoh_flocon2008.pdf.
- [Fan08] Robert L Fanelli. A hybrid model for immune inspired network intrusion detection. In *Artificial Immune Systems*, pages 107-118. Springer, 2008. Need copy. http://link.springer.com/chapter/10.1007/978-3-540-85072-4_10
- [HGRF08] Thorsten Holz, Christian Gorecki, Konrad Rieck, and Felix C Freiling. Measuring and Detecting Fast-Flux Service Networks. In *Proceedings of the 15th Annual Network and Distributed System Security Symposium*, February 2008.
- [Ish09] Keisuke Ishibashi. Detecting Anomalies in Inter-Host Communication Graphs. In *FloCon 2009 Proceedings*, Pittsburgh, PA, USA, 2009.
CERT. http://www.cert.org/flocon/2009/presentations/Ishibashi_GraphAnomalies.pdf.
- [SOC*09] Benjamin Sangster, T O'Connell, Thomas Cook, Robert Fanelli, Erik Dean, William J Adams, Chris Morrell, and Gregory Conti. Toward instrumenting network warfare competitions to generate labeled datasets. In *Proc. of the 2nd Workshop on Cyber Security Experimentation and Test (CSET'09)*, 2009. http://static.usenix.org/events/cset09/tech/full_papers/sangster.pdf
- [LMV09] Sungkornsarun Longchupole, Noppadol Maneerat, and Ruttikorn Varakulsiripunth. Anomaly detection through packet header data. In *Information, Communications and Signal Processing, 2009. ICICS 2009. 7th International Conference on*, pages 1-4. IEEE,

2009. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5397552.
- [CBK09] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Computing Surveys (CSUR)*, 41(3):15, 2009. <http://dl.acm.org/citation.cfm?id=1541882>.
- [Bau10] T Baults. Slow Down Internet Worms with Tarbits, 2010. <http://www.symantec.com/connect/articles/slow-down-internet-worms-tarbits>, Accessed: 2013-06-26.
- [APD⁺10] Manos Antonakakis, Roberto Perdisci, David Dagon, Wenke Lee, and Nick Feamster. Building a Dynamic Reputation System for DNS. In *USENIX Security Symposium*, pages 273-290, 2010.
- [Sto10] E. Stoner. DNS footprint of malware. In *DNS-OARC Workshop 2*, Denver, CO, 2010.
- [GAT10] Julie Greensmith, Uwe Aickelin, and Jamie Twycross. Detecting danger: Applying a novel immunological concept to intrusion detection systems. *arXiv preprint arXiv:1002.0696*, 2010. <http://arxiv.org/pdf/1002.0696.pdf>
- [Kis10] Richard Kissel. Draft Glossary of Key Information Security Terms, 2010. <http://csrc.nist.gov/publications/nistir/ir7298-rev1/nistir-7298-revision1.pdf>.
- [HSCZ10] Shun-Wen Hsiao, Yeali S Sun, Meng Chang Chen, and Hui Zhang. Cross-level behavioral analysis for robust early intrusion detection. In *Intelligence and Security Informatics (ISI), 2010 IEEE International Conference on*, pages 95-100. IEEE, 2010. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5484768.
- [SP10] Robin Sommer and Vern Paxson. Outside the closed world: On using machine learning for network intrusion detection. In *Security and Privacy (SP), 2010 IEEE Symposium on*, pages 305-316. IEEE, 2010. <http://www.icir.org/robin/papers/oakland10-ml.pdf>.
- [CAW10] Richard A Caralli, Julia H Allen, and David W White. *CERT Resilience Management Model (CERT-RMM): A Maturity Model for Managing Operational Resilience*. Addison-Wesley Professional, 2010.
- [KCKV11] Alexandros Kapravelos, Marco Cova, Christopher Kruegel, and Giovanni Vigna. Escape from] monkey island: Evading high-interaction honeyclients. In *Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 124-143. Springer, 2011.
- [AAKS11] Sajjad Arshad, Maghsoud Abbaspour, Mehdi Kharrazi, and Hooman Sanatkar. An anomaly-based botnet detection approach for identifying stealthy botnets. In *Computer Applications and Industrial Electronics (ICCAIE), 2011 IEEE International Conference on*, pages 564-569. IEEE, 2011. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6162198.
- [DC11] Jonathan J Davis and Andrew J Clark. Data preprocessing for anomaly based network intrusion detection: A review. *Computers & Security*, 30(6):353-375, 2011. <http://www.sciencedirect.com/science/article/pii/S0167404811000691>.
- [BHSS11] Nathaniel Boggs, Sharath Hiremagalore, Angelos Stavrou, and Salvatore J Stolfo. Cross-domain collaborative anomaly detection: so far yet so close. In *Recent Advances in Intrusion Detection*, pages 142-160. Springer, 2011. http://ids.cs.columbia.edu/sites/default/files/colaborative_RAID11_1.pdf.
- [MJH11] Peter Mullarkey, Mike Johns, and Ben Haley. Leveraging Other Data Sources with Flow to Identify Anomalous Network Behavior. In *FloCon 2011 Proceedings*, Pittsburgh, PA, USA, 2011. CERT. http://www.cert.org/flocon/2011/presentations/Mullarkey_Leveraging.pdf.
- [Fal11] Falliere, Nicolas and Murchu, Liam O and Chien, Eric. Symantec stuxnet dossier. Technical report, Symantec, 2011. http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf.
- [APL⁺11] Manos Antonakakis, Roberto Perdisci, Wenke Lee, Nikolaos Vasiloglou, and David Dagon. Detecting malware domains at the upper DNS hierarchy. In *Proceedings of the 20th USENIX Security Symposium, USENIX Security*, volume 11, pages 27-27, 2011.
- [SS11] Malek Ben Salem and Salvatore J Stolfo. Modeling user search behavior for masquerade detection. In *Recent Advances in Intrusion Detection*, pages 181-200. Springer, 2011. http://ids.cs.columbia.edu/sites/default/files/RAID_2011_0.pdf.

- [ATBT11] Hasty Atashzar, Atefeh Torkaman, Marjan Bahrolololum, and Mohammad H Tadayon. A survey on web application vulnerabilities and countermeasures. In *Computer Sciences and Convergence Information Technology (ICCIT), 2011 6th International Conference on*, pages 647-652. IEEE, 2011. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6316697.
- [HCA11] Eric M Hutchins, Michael J Cloppert, and Rohan M Amin. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare and Security Research*, 1:80, 2011.
- [PB11] D. Plonka and P. Barford. Flexible Traffic and Host Profiling via DNS Rendezvous. March 2011.
- [You11] W. Young. Security Message Standardization: the beginning of the end. In *5th Annual REN-ISAC Member Meeting*, April 9, 2011.
- [RBJ⁺11] M Rajab, Lucas Ballard, Nav Jagpal, Panayiotis Mavrommatis, Daisuke Nojiri, Niels Provos, and Ludwig Schmidt. Trends in circumventing web-malware detection. Technical report, Google, July 2011.
- [Jac12] Michael Jacobs. Indicator Expansion Techniques-Tracking Cyber Threats via DNS and Netflow Analysis. In *FloCon 2012 Proceedings*, Pittsburgh, PA, USA, 2012. CERT. <http://www.cert.org/flocon/2012/presentations/jacobs-indicator-expansion-techniques.pdf>.
- [JW12] George Jones and Austin Whisnant. Network Profiling with SiLK. In *FloCon 2012 Proceedings*, Pittsburgh, PA, USA, 2012. CERT. <http://www.cert.org/flocon/2012/presentations/jones-whisnant.pdf>, Accessed: 2013-04-06.
- [NA12] Nist and Emmanuel Aroms. *NIST Special Publication 800-53 Revision 3 Recommended Security Controls for Federal Information Systems and Organizations*. CreateSpace, Paramount, CA, 2012.
- [LCB⁺12] Wei Liu, Sanjay Chawla, James Bailey, Christopher Leckie, and Kotagiri Ramamohanarao. An efficient adversarial learning strategy for constructing robust classification boundaries. In *AI 2012: Advances in Artificial Intelligence*, pages 649-660. Springer, 2012.
- [VGS12] R Venkatesan, R Ganesan, and A Arul Lawrence Selvakumar. A Comprehensive Study in Data Mining Frameworks for Intrusion Detection. *International Journal*, 2, 2012. http://theaccents.org/ijacr/papers/december_2012/6.pdf.
- [Gol12] Josh Goldfarb. The UberData Source: Holy Grail or Final Fantasy? In *FloCon 2012 Proceedings*, Pittsburgh, PA, USA, 2012. CERT. <http://www.cert.org/flocon/2012/presentations/goldfarb-uber-data-source.pdf>.
- [SM12] Karen Scarfone and Peter Mell. Guide to intrusion detection and prevention systems (idps) (Draft). *NIST Special Publication*, 800(94-1), 2012. http://csrc.nist.gov/publications/drafts/800-94-rev1/draft_sp800-94-rev1.pdf.
- [KND12] Shubhalaxmi Kher, Victor Nutt, and Dipankar Dasgupta. A Prediction Model for Anomalies in Smart Grid with Sensor Network. 2012. http://ais.cs.memphis.edu/files/papers/csiirw8_submission_82.pdf.
- [Imp12] Imperva. Assessing the Effectiveness of Antivirus Solutions. Hacker Intelligence Initiative, Monthly Trend Report 14, 2012. http://www.imperva.com/docs/HII_Assessing_the_Effectiveness_of_Antivirus_Solutions.pdf.
- [ZYWZ12] Bin Zhang, Jia-Hai Yang, Jian-Ping Wu, and Ying-Wu Zhu. Diagnosing Traffic Anomalies] Using a Two-Phase Model. *Journal of Computer Science and Technology*, 27(2):313-327, 2012. Need copy. <http://link.springer.com/article/10.1007/s11390-012-1225-0>.
- [HMJ12] Fazirulhisyam Hashim, Kumudu S Munasinghe, and Abbas Jamalipour. On the negative selection and the danger theory inspired security for heterogeneous networks. *Wireless Communications, IEEE*, 19(3):74-84, 2012. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6231162.

- [MA12] David Mudzingwa and Rajeev Agrawal. A study of methodologies used in intrusion detection and prevention systems (IDPS). In *Southeastcon, 2012 Proceedings of IEEE*, pages 1-6. IEEE, 2012. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6197080.
- [HS12] C. Huth and J. Spring. The Impact of Passive DNS Collection on End-user Privacy. In *Securing and Trusting Internet Names 2012*. National Physical Laboratory, UK, March 2012.
- [DHS12] Privacy Impact Assessment for the National Cybersecurity Protection System. Technical report, United States Department of Homeland Security, July 30, 2012.
- [WF12] Austin Whisnant and Sid Faber. Network Profiling Using Flow. Technical Report CMU/SEI-2012-TR-006, Software Engineering Institute, CERT Coordination Center, August 2012.
- [TJS12] James L. Spencer Timothy J. Shimeall, Ronald M. Bandes. Scalable Cyber Deception (Revised). Technical report, Software Engineering Institute, August 2012. Publication NetSA-2012-12a, Available by request.
- [Ros12] Matthew Rosenquist. Cyber security Hunter teams are the next advancement in network defense. IT Peer Network, November 2012. <http://communities.intel.com/community/itpeernetwork/blog/2012/11/28/cyber-security-hunter-teams-are-the-next-advancement-in-network-defense>, Accessed: 2013-08-01.
- [web13] Web of Trust, 2013. <http://www.mywot.com/>, Accessed: 2013-08-13.
- [JS13] George Jones and Tim Shimeall. Behavioral Whitelists of High-Volume Web Traffic to Specific Domains. In *FloCon 2013 Proceedings*, Pittsburgh, PA, USA, 2013. CERT. <http://www.cert.org/flocon/2013/posters/jones-shimeall-behavioral-whitelists-high-volume-web-traffic.pdf>.
- [AA13] Brian Allen and Robert Annand. Behavioral Whitelists of Beaconing Activity. In *FloCon 2013 Proceedings*, Pittsburgh, PA, USA, 2013. CERT. <http://www.cert.org/flocon/2013/posters/allen-annand-behavioral-whitelists-of-beaconing-activity.pdf>.
- [RBL⁺13] Moheeb Abu Rajab, Lucas Ballard, Noé Lutz, Panayiotis Mavrommatis, and Niels Provos. Camp: Content-agnostic malware protection. 2013. <https://www.cs.jhu.edu/~moheeb/aburajab-ndss-13.pdf>.
- [Sic13] Svetlana Sicular. Gartner's Big Data Definition Consists of Three Parts, Not to Be Confused with Three "V"s. forbes.com website, 2013.
- [LSBHG13] Bingdong Li, Jeff Springer, George Bebis, and Mehmet Hadi Gunes. A survey of network flow applications. *Journal of Network and Computer Applications*, 36(2):567-581, 2013. <http://www.sciencedirect.com/science/article/pii/S1084804512002676>.
- [SJ13] Char Sample and George Jones. Anomaly Detection. In *FloCon 2013 Proceedings*, Pittsburgh, PA, USA, 2013. CERT. <http://www.cert.org/flocon/2013/presentations/sample-char-intro-anomaly-detection.pdf>.
- [CCN⁺13] Howard Chivers, John A Clark, Philip Nobles, Siraj A Shaikh, and Hao Chen. Knowing who to watch: Identifying attackers whose actions are hidden within false alarms and background noise. *Information Systems Frontiers*, pages 1-18, 2013. Need copy. <http://link.springer.com/article/10.1007/s10796-010-9268-7>.
- [CER13a] CERT. CERT Podcast Series: Security for Business Leaders, 2013. <http://www.cert.org/podcast/>, Accessed: 2013-06-21.
- [CER13b] CERT. CERT Resilience Management Model, 2013. <http://www.cert.org/resilience/rmm.html>, Accessed: 2013-06-21.
- [JM13] George M. Jones and Dr. Soumyo Moitra. What You Know And What You Don't: A Survey of Anomaly Detection Techniques Applied to Network Intrusion Detection. Technical report, Software Engineering Institute, 2013.