

450.1

Blue Team Tools and Operations

To: David Owerbach <0mamaloney0@gmail.com> April 28, 2020



PLEASE READ THE TERMS AND CONDITIONS OF THIS COURSEWARE LICENSE AGREEMENT ("CLA") CAREFULLY BEFORE USING ANY OF THE COURSEWARE ASSOCIATED WITH THE SANS COURSE. THIS IS A LEGAL AND ENFORCEABLE CONTRACT BETWEEN YOU (THE "USER") AND SANS INSTITUTE FOR THE COURSEWARE. YOU AGREE THAT THIS AGREEMENT IS ENFORCEABLE LIKE ANY WRITTEN NEGOTIATED AGREEMENT SIGNED BY YOU.

With the CLA, SANS Institute hereby grants User a personal, non-exclusive license to use the Courseware subject to the terms of this agreement. Courseware includes all printed materials, including course books and lab workbooks, as well as any digital or other media, virtual machines, and/or data sets distributed by SANS Institute to User for use in the SANS class associated with the Courseware. User agrees that the CLA is the complete and exclusive statement of agreement between SANS Institute and you and that this CLA supersedes any oral or written proposal, agreement or other communication relating to the subject matter of this CLA.

BY ACCEPTING THIS COURSEWARE, YOU AGREE TO BE BOUND BY THE TERMS OF THIS CLA. BY ACCEPTING THIS SOFTWARE, YOU AGREE THAT ANY BREACH OF THE TERMS OF THIS CLA MAY CAUSE IRREPARABLE HARM AND SIGNIFICANT INJURY TO SANS INSTITUTE, AND THAT SANS INSTITUTE MAY ENFORCE THESE PROVISIONS BY INJUNCTION (WITHOUT THE NECESSITY OF POSTING BOND) SPECIFIC PERFORMANCE, OR OTHER EQUITABLE RELIEF.

If you do not agree, you may return the Courseware to SANS Institute for a full refund, if applicable.

User may not copy, reproduce, re-publish, distribute, display, modify or create derivative works based upon all or any portion of the Courseware, in any medium whether printed, electronic or otherwise, for any purpose, without the express prior written consent of SANS Institute. Additionally, User may not sell, rent, lease, trade, or otherwise transfer the Courseware in any way, shape, or form without the express written consent of SANS Institute.

If any provision of this CLA is declared unenforceable in any jurisdiction, then such provision shall be deemed to be severable from this CLA and shall not affect the remainder thereof. An amendment or addendum to this CLA may accompany this Courseware.

SANS acknowledges that any and all software and/or tools, graphics, images, tables, charts or graphs presented in this Courseware are the sole property of their respective trademark/registered/copyright owners, including:

AirDrop, AirPort, AirPort Time Capsule, Apple, Apple Remote Desktop, Apple TV, App Nap, Back to My Mac, Boot Camp, Cocoa, FaceTime, FileVault, Finder, FireWire, FireWire logo, iCal, iChat, iLife, iMac, iMessage, iPad, iPad Air, iPad Mini, iPhone, iPhoto, iPod, iPod classic, iPod shuffle, iPod nano, iPod touch, iTunes, iTunes logo, iWork, Keychain, Keynote, Mac, Mac Logo, MacBook, MacBook Air, MacBook Pro, Macintosh, Mac OS, Mac Pro, Numbers, OS X, Pages, Passbook, Retina, Safari, Siri, Spaces, Spotlight, There's an app for that, Time Capsule, Time Machine, Touch ID, Xcode, Xserve, App Store, and iCloud are registered trademarks of Apple Inc.

PMP and PMBOK are registered marks of PMI.

SOF-ELK® is a registered trademark of Lewes Technology Consulting, LLC. Used with permission.

SIFT® is a registered trademark of Harbingers, LLC. Used with permission.

Governing Law: This Agreement shall be governed by the laws of the State of Maryland, USA.

SEC450.1

Blue Team Fundamentals: Security Operations and Analysis

SANS

Blue Team Tools and Operations

© 2020 John Hubbard | All Rights Reserved | F01_01

Welcome to SANS SEC450 – Blue Team Fundamentals: Security Operations and Analysis!

TABLE OF CONTENTS	PAGE
Course Outline	03
Welcome to the Blue Team	04
EXERCISE 1.0: Virtual Machine Setup	22
SOC Overview	25
Defensible Network Concepts	50
Events, Alerts, Anomalies, and Incidents	71
Incident Management Systems	90
EXERCISE 1.1: TheHive Incident Management System	113
Threat Intelligence Platforms	115
EXERCISE 1.2: MISP Threat Intelligence Platform	139
SIEM and Automation	141
Know Your Enemy	161
Day 1 Summary	176
EXERCISE 1.3: SIEM with the Elastic Stack	179

SANS | SEC450: Blue Team Fundamentals: Security Operations and Analysis 2

450.1 Table of Contents

This table of contents outlines the plan for 450.1.

Course Outline

Day 1: Blue Team Tools and Operations

Day 2: Understanding Your Network

Day 3: Understanding Endpoints, Logs, and Files

Day 4: Triage and Analysis

Day 5: Continuous Improvement, Analytics, and Automation

Course Outline

Welcome to Day 1 of SEC450. Here are the high-level topics that will be discussed throughout the next 5 books!

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. **Welcome to the Blue Team!**
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Welcome to the Blue Team!

- If you are new to the SOC/Blue Team, welcome!!
- One of the most challenging and in-demand careers in cybersecurity!
- If you think the offense side is hard... just wait
- LOTS of exciting development going on!
- This class will start you on your journey



Welcome to the Blue Team!

Welcome to SEC450: Blue Team Fundamentals: Security Operations and Analysis, and if you're new to the information security and Security Operations Centers, get ready for the ride of your life! Cyber defense is, without a doubt, one of the most exciting and challenging careers in information security. With attacks only becoming more frequent and more destructive, defense is a position that will be desperately needed for years to come. You may have looked at attacker's tools and wondered how such brilliant ideas were conceived. Trust me when I say that not only can equally amazing things be done on the defensive side, but if nothing else, there's even MORE opportunity for development. In the past, defense teams have lagged behind the newest and wildest attacks, left behind by the rapid pace of hacking tool development, but we are in a new and exciting time right now. Blue Team tools have caught up, and new defensive frameworks like MITRE's ATT&CK, for example, have helped bring the Blue Team up to speed and are continuing to be developed at a rapid pace. For the first time in 2018, we even saw a Blue Team Village at DEF CON and turnout was amazing!

There's an enormous number of exciting things going on in the defensive space and, in the author's opinion, there's never been a better time to be joining the Blue Team. This class is designed to jumpstart those who are new to defensive operations, whether you're a one-person team or a new member of a SOC, the goal is to give you the mindset, knowledge, and operational familiarity needed to be successful in modern cyber defense. Without further hesitation, let's get started...

About This Class

- This class is designed for defense team members and those who train, work with, and manage them
- We will cover a range Blue Team fundamentals:
 - **People:** Mindset, mental models, career progression, burnout
 - **Process:** Analysis, investigation theory, triage, and data flow
 - **Technology:** Network and host monitoring, understanding protocols, spotting attacks, scripting and automation
- Strategic, operational, and tactical level info
- Why was it written? **As a move toward standardized training, we need your help!**

About This Class

This class is designed primarily for those who are early in their careers or work closely with others who are. Whether you have a full SOC or not, the concepts covered will apply equally. Success in security operations is driven primarily by three factors: People, process, and technology. We will be covering all three.

Process-wise, we will cover the workflows and actions taken by the typical Blue Team member, and how to make them as efficient as possible, utilizing automation and eliminating manual work wherever possible. Technology-wise, we will cover the main tools and sources of data the typical security operations group have, and how to understand and wield them to their best potential. Finally, the people factor, that is why we're all here, right? In these 6 books, we will cover a broad range of topics with in-depth explanations and hands-on activities to help facilitate learning as rapidly as possible.

Why was this class created? Because the Blue Team desperately needs help! Our tools and process are rapidly maturing, and our technology is already outstanding when used correctly, but we have a dire shortage of people who are prepped and ready to jump into a security operations role. Therefore, the focus of this class will be developing you, the analyst, with skills that will help you immediately understand and jump in on alerts as they arrive.

Cybersecurity Skills Gap

2019 ISC² Global Workforce Study¹:

- Almost a **4.07M** worker shortage globally!
 - **561K** in NA, **291K** in Europe, **2.6M** in APAC, **600K** in LATAM
- **65%** of professionals believe they don't have enough people
- **56%** intended to pursue cybersecurity during education
- **42%** of professionals first job after education was not cybersecurity
- **Non-existent unemployment, shortage of workers, and high salaries mean **hiring/retention is difficult!****

Cybersecurity Skills Gap

How big is our people issue? According to the 2019 Global Workforce Study performed by ISC², we're looking at a shortage of over four million cybersecurity jobs! In our current state, 65% of teams already believe they don't have enough people! This is likely one of the reasons why cybersecurity is often stated to have "0% unemployment"—in some locations, this is close to true! Globally, ISC² found that the number is 2% on average, which is still extremely low.

So, what can be done about the issue? We need to bring in additional talent from other fields, and this is already happening. This year's ISC² surveys found that only 56% of professionals in cybersecurity intended to pursue cybersecurity during their education, 42% had a job outside of cybersecurity as their first job after education. To fill all the available demand, cybersecurity will need to extend its reach, and that has already begun. "Re-skilling" those from other technical and even non-technical roles has become a common source of new talent.

Even within security roles, the low unemployment rate can make it hard to keep those we have found. With so few people available, job-hopping is common since sometimes the only source of employees are those who are already employed elsewhere. Previous surveys also found that roughly 1 in 5 had changed jobs voluntarily in the past year, another high number showing that it is currently an employee's market within the industry.

[1] <https://www.isc2.org/Research/2019-Cybersecurity-Workforce-Study#>

The Difficulty of SOC Work

The problem is, SOC work can be tough...

- High barrier to entry
- Ticket/alert based, repetitive work
- Tiered structure may limit visibility and scope
- Over-prescribed workflow restricts freedom
- Repetitive clicking and information filling
- High turnover means "revolving door" of coworkers

We'll discuss how to fix this

- Let's make security operations fun and engaging!

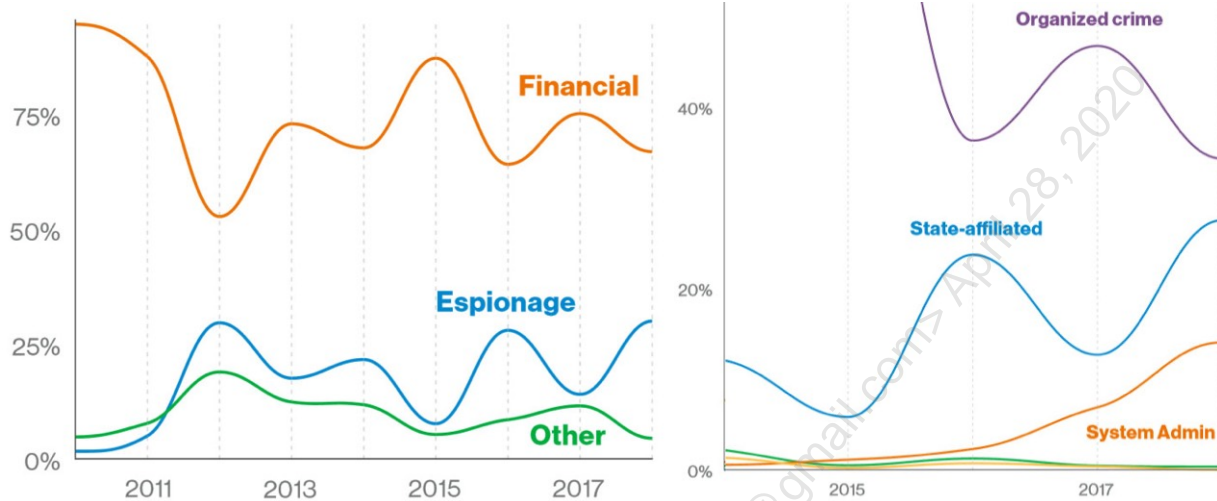
The Difficulty of SOC Work

One of the unfortunate sides of SOC work is that in some situations it can be repetitive and restricting. The author has heard SOCs described as everything from an awesome place to work where everyone is engaged and learning to "a revolving door" where "no one has any idea what they're doing," yikes! The processes used and the environment created within the group will have a huge impact on the effectiveness of the group, the engagement of its employees, and employee retention.

The goal of this class, aside from training in how to best understand and perform the duties of the job, is to help show ways to eliminate the misery associated with SOC work in some organizations—repetitive data entry, mindless form filling, overly restrictive workflow, and mountains of alerts. These issues can be tamed and should be aggressively tackled by any team facing them before they poison the environment of the group. These issues will be tackled head-on in Day 5. The author's wish is to truly help Blue Team jobs be looked at as an amazing and engaging place to work for the long term, and not "something you have to do a stint in" or as a steppingstone to something else. Whether or not this is the case in your organization will depend highly upon the SOCs attitude, processes, and past.

Why Are We Being Attacked?

From the 2019 Verizon DBIR



SANS

SEC450: Blue Team Fundamentals: Security Operations and Analysis

9

Why Are We Being Attacked?

Why are all these cyber attacks happening in the first place? According to the 2019 Verizon Data Breach Investigations Report, which is one of the most comprehensive reports compiled every year, most attacks are **financially motivated**. Attackers are after anything of value that can be sold on the black market or otherwise. The second most common reason is to gain a strategic advantage, otherwise known as **espionage**. Verizon notes that between financial gain and espionage, roughly 90% of breaches are accounted for. The remaining attacks are comprised of grudges, fun, or "other."

Who is perpetrating these attacks? The chart on the right (reduced in detail to fit on the slide) shows a breakdown of threat actors involved in breaches over the past few years. The message we can take from this is that **the most typical cybersecurity breach will consist of an external attacker looking to steal the data they can use to make a quick dollar**. Knowing this can be the first step in putting up a threat-aligned defense.

[1] Verizon 2018 DBIR: <https://www.verizonenterprise.com/verizon-insights-lab/dbir/>

How Long Does It Take To Realize It? Mandiant M-Trends 2019

GLOBAL MEDIAN DWELL TIME

Compromise Notification	2011	2012	2013	2014	2015	2016	2017	2018
All	416	243	229	205	146	99	101	78
External					320	107	186	184
Internal					56	80	57.5	50.5

DETECTION BY SOURCE

Compromise Notifications	2011	2012	2013	2014	2015	2016	2017	2018
External	94%	63%	67%	69%	53%	47%	38%	41%
Internal	6%	37%	33%	31%	47%	53%	62%	59%

How Long Does It Take To Realize It? Mandiant M-Trends 2018

How fast are we noticing that we're breached? Using the Mandiant M-Trends 2019 report, which covers all breaches Mandiant responded to in 2018, it seems to highly depend on your region and whether you were able to spot the breach yourself. While the global median dwell time was **71 days**, down from 101 days in the 2018 report, that doesn't tell nearly the whole story. Digging into the details by region and method of discovery, there appears to be an interesting bi-modal distribution. Those who identified the breach internally had a median dwell time of **50.5 days** in 2018 vs. **184 days** for breaches discovered by an external party.

What can we take from this? It's hard to make a perfect conclusion without knowing the underlying data, but it seems safe to conclude that those who discover their own breaches are likely paying significantly less as a result of incidents and seeing much less damage done than those who don't spot breaches themselves. How far are these breaches progressing in this time? Again, we don't have the raw data, but consider the fact that the Mandiant M-Trends report from 2016 stated that Mandiant's red team took an average of only 3 days to reach domain admin (the new number was not reported since then). It is likely every single one of these breaches had the potential to be a full-on disaster. The takeaway is that the number is still much larger than it should be, and that the Blue Team has a lot of work to do!

The Cure: Improving Cyber Security Operations

Cyber Security Operations:

*"Protecting the **confidentiality, integrity, and availability** of **information systems** of an organization through **proactive design and configuration, ongoing monitoring** of system state, **detection** of unintended actions or undesirable state, and **minimizing damage** from unwanted effects."¹*

The Cure: Improving Cyber Security Operations

On to the first piece of a defensive security mindset: What is cyber security operations? This definition is used by Chris Crowley, a SANS Sr. Instructor and security operations expert, as an all-encompassing mission statement. It's all about ensuring systems continue to perform as expected and facilitating doing that through good design and ongoing monitoring. When things do inevitably go wrong, it is Security Operations' job to ensure damage is minimized. Notice we did not say "there is no damage." This is an acknowledgement that no network is invincible, and that compromise will occur. The question will be: Is your team and technology properly configured to minimize that damage?

[1] <https://www.montance.com/mgt517/>

What is a Cyber Security Operations Center?

- **Cyber:** Related to information Systems
- **Security:** Intended use only
- **Operations:** Ongoing performance
- **Center:** A hub or nexus of activity
- The "**Blue Team**" for your organization:
 - A group of many acronyms – SOC, ISOC, SIRT, CSIRT, DART
 - Monitoring/hunting, incident response, threat intel, forensics
 - Responsible for security operations



What is a Cyber Security Operations Center?

If security operations is protecting and monitoring your data, and minimizing damage from occurring, the SOC is at the center of those activities, and houses (virtually or physically) the people, process, and technology that operationalizes those goals. This group of people is often synonymously called the "Blue Team." The terms Blue Team and SOC will be used interchangeably throughout this class. Depending on your organization and whether you have a defense team or an official "SOC", we are talking about the functions required to perform threat monitoring and detection, as well as potentially incident response and forensics. It may even include self-assessment functions such as vulnerability assessments and penetration testing. Ultimately, the specific groups that fall under the title are less important than getting these monitoring operations done fully and correctly. This is what truly helps us improve our time to detection and monitoring capabilities.

Blue Team Definition

Blue Team:


- The defensive security team
- Protects against both real and red team attacks

Purple Teaming:

- An exercise meant to bring together red and blue
- Ensures information sharing between groups

Red Teaming:

- The process of using TTPs to emulate a real-world threat with the goals of training and measuring the effectiveness of people, process, and technology used to defend an environment.¹



Blue Team Definition

There are multiple teams you may hear about in the cyber defense realm. This class is focused on the fundamentals for acting as an analyst and member of the blue or defensive team. The adversaries you will be facing as part of the Blue Team include real threat actors as well as a potential dedicated red team. The red team is an internal or external group whose job is to simulate real attacker Tactics, Techniques, and Procedures (TTPs) and challenge the Blue Team to improve their effectiveness. Joe Vest, a red teaming expert with SpectreOps, has a great definition for the red team: *They should measure the gap between "what is" and "what should be" to get the truth about security operations as a whole.*

You may also hear the term "Purple Team." Although there could be someone named to this position to facilitate communication between the red and Blue Teams, a purple team is often an *activity* and describes an exercise where the red and Blue Team sit together and simultaneously walk through a simulated attack. The goal is to learn real-time from each other how attacks can be launched and where their effects can be seen in the environment. It can be a highly effective way of training defensive teams consisting of both new and seasoned members. The output of such a test may surprise you. Defenses you thought would work may fail to produce evidence of the attack, leading to a high value finding and the ability to patch a dangerous and unknown oversight.

For those interested in diving into the details of red and purple teams, SANS provides SEC599: Defeating Advanced Adversaries - Purple Team Tactics & Kill Chain Defenses and SEC699: Purple Team Tactics - Adversary Emulation for Breach Prevention & Detection.

[1] <https://www.sans.org/course/defeating-advanced-adversaries-kill-chain-defenses>

Why Are We Here? (1)

Blue Team Truth #1: Compromise Will Happen

- **The question is...how it will affect you?**
 - **Outcome 1:** Adversary gets past the perimeter but is detected and fails to complete mission
 - **Outcome 2:** Adversary is not detected, runs free, causes huge impact!
- Not all adversaries will be blocked at the perimeter
- The acknowledgement of this fact and preparation for the consequences separate a good from bad SOC!
- **Goal: Detect and minimize damage** from compromise

Why Are We Here? (1)

One important truth to remember is that no matter how hard you try, compromise will happen to *some* extent. It's impossible to put up a defense that works perfectly forever. That's not to say it will be a full-on disaster, however. It's up to the Blue Team to determine how far an attack will progress and the damage it will be able to cause. The Blue Team affects these things by upwardly informing the business of the true risks in the cybersecurity landscape and driving good policy and configurations throughout the environment, as well as actively by being on point when things do inevitably start to go wrong.

Although preventing as many attacks as possible at the perimeter is ideal, it is not realistic; therefore, we must strive toward a goal of minimizing the damage of what does manage to get through the initial layers of defense. We do this by architecting our people, process, and technology for rapid detection of a compromise in progress, and work to rapidly stop its progress before severe damage is caused. Getting phished or getting a piece of malware installed on someone's laptop is not a big deal compared to what will happen if that attack is not noticed and the attacker uses this access over months to ultimately steal intellectual property or sensitive personal information and take that breach public. Therefore, it is the Blue Team's goal to catch the attacks as early as possible when it is easy to clean up and kick the attackers out of the environment!

Why Are We Here? (2)

Blue Team Truth #2: Your company does not solely exist to be secure

The team provides a "**loss prevention**" function:

- We **reduce cybersecurity risk to an acceptable level**
- Must strike **balance** between security and productivity
- Balance is defined by your organization/management
- Can be frustrating, but doesn't mean we can't try to influence
- Blue team must **inform those who make the risk decisions**
- Good information requires a deep understanding of your craft...

Why Are We Here? (2)

Given the previous slide, it's easy to want to jump in and lock everything down to the fullest. However, it's important to keep the higher-level perspective. Yes, the Blue Team can point out all the weak points the organization has and say everything must be locked down; but, remember, your company does not solely exist to be secure. Organizations exist to create value in one way or another, and cybersecurity can ultimately be viewed as a loss prevention function, similar to security guards at a physical store watching for shoplifters. As a store owner, you could make everyone go through a TSA-style search entering and leaving the store, but you would likely find such a store quickly out of business despite the 0% shoplifting loss.

There's a corollary to the above statement and that is "Organizations only get to stay in business if they effectively and correctly (enough) manage risk." Therefore, we can't have zero security either! There's a balance that must be struck between security and allowing the business to function without hindrance. This can be very frustrating at times, but every business will have its own take on what the major risks are (cyber and otherwise), and where the dial should be set on security vs. productivity. The goal of the Blue Team is ultimately to help the business understand the true risk they face and help them lock down things as much as possible, ideally causing minimal or zero productivity hit. This is where the art of designing a secure system comes in. At times, bargains and tradeoffs must be made, and clever solutions can be designed that will allow people and networks to operate without slowing down the pace of whatever it is that makes the company valuable.

The Blue Team Mission

Goal: Reduce effects of compromise to the minimum possible by

1. Staying up-to-date on current news and attack techniques
2. Utilizing threat intelligence for an advantage over attackers
3. Monitoring network and endpoints for signs of compromise
4. Triageing potential issues with speed and accuracy
5. Reacting quickly to scope and contain the incident
6. Remediating and recovering from the attack
7. Incorporating lessons to ensure it doesn't happen again

The Blue Team Mission

So, what role do you play in this picture? Your mission, at the end of the day, is to reduce the effects of compromise of your organization to the minimum possible. The way to do this is to be prepared to quickly identify the signs of compromise and put a stop to the progression of the incident. We can do this via a well-monitored network and set of endpoints, and the ability to alert on and triage potential issues quickly and accurately. This task is not easy. It requires lifelong learning and a drive to chase down adversaries. Attack techniques and exploits will change daily, and to add to the challenge, you will always have imperfect knowledge of the situation at hand. Learning to understand the protocols in use on the network, how to read and interpret log files, how to perform high quality analysis, and how data should and shouldn't flow goes a long way toward accomplishing this mission. Fortunately, that is what we will be focusing on in this class.

Modern Defense Mindset¹

A modern defense is our only hope, this means:

- Presumption of Compromise
- Detection Oriented Defense
- Proactive Detection: Hunt Teams
- Post-Exploitation Focus
- Response-Driven
- Risk Informed Strategy

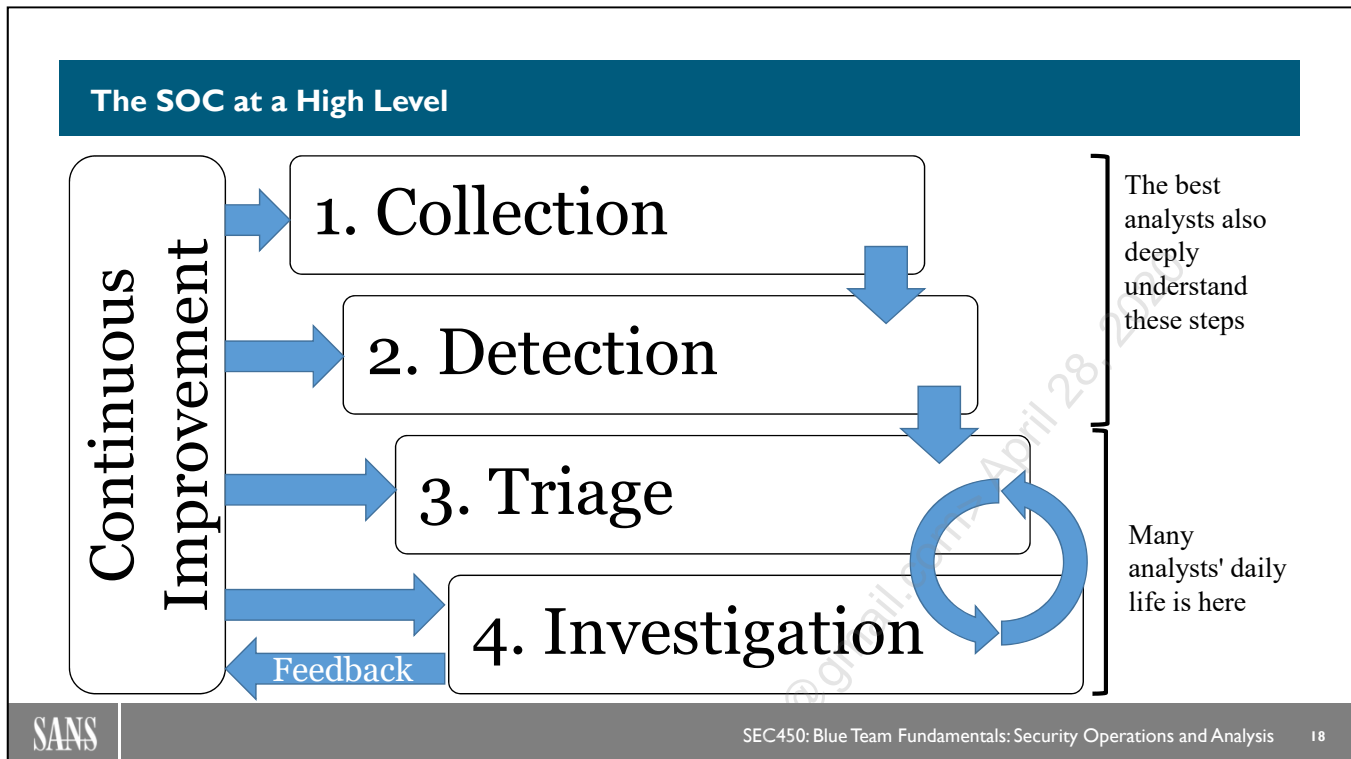
"Prevention is ideal, detection is a must"

Modern Defense Mindset

As a Blue Team, we need to make sure everyone is on the same page as to what constitutes a "modern defense mindset" and what it takes to defend the networks of today. SANS SEC511: "Continuous Monitoring and Security Operations" course (a great follow-on to this one) by Eric Conrad and Seth Misenar has a very well thought out rundown of some of the most important concepts, which are listed here. In short, we must acknowledge the likelihood of compromise at some point and put up a strong defense in terms of both prevention *and* detection technologies that can catch the things that do slip by. This proactive detection can be run by a hunt team or anyone else tasked with searching through the collected data looking for evidence of what has made it past the front gates and taken a foothold. This team will need to be post-exploitation focused in what they are looking for, as it is assumed to have already broken through the exploitation stage and is now present inside the environment.

Detection is great, but without the ability to do something about it, it's of no use, so detection must be followed by a rapid *response*. This means that once anything has been found, we need to be prepared both process and technology-wise to respond as fast as possible to prevent any further damage from being done. In most compromise situations, the cost of cleanup rises as the breach progresses, so catching it and stopping it as early as possible is of utmost importance. Finally, we must put up our defense in a risk-informed manner. Many teams will not have the budget or time to protect everything perfectly and evenly; therefore, knowing what your adversary might be after in your organization and how they might get there is vital for the optimization of resources. Placing more prevention and detection technology in front of the servers that hold the most important data and the people who have access to it is a rational way to maximize return on your security tool investment.

[1] SEC511: <https://www.sans.org/course/continuous-monitoring-security-operations>



The SOC at a High Level

If we were to break the average SOC's functions down into the most general set of steps, this would be it: Collect, detect, triage, investigate (and potentially incident response if you perform those duties as well). At a very high level, the SOC collects multiple types of data using various sensors and tools. Of all the data that is collected or seen, interesting items must be picked out of it at the detection stage. This process is fed by threat intelligence, attack technique knowledge, information about your environment and otherwise. Anything that is "interesting" becomes an alert and gets marked for investigation. There will almost always be more than one item to attend to at once, so triage will be a necessary skill in this process. Ideally, our SOC data organization tools help analysts pick the most important items first since that is the top priority of this step. How well we can pick out the most dangerous alerts depends on multiple factors: Knowledge of attack stages, methods of attack, risk scoring capability, data enrichment, experience with the environment, etc. all play a role in the quality possible in this step.

Once the most seemingly dangerous item is selected, we must dive in, triage, and investigate it to see if it truly is something bad going on, then, pass the case on to incident response if required. Analysts should be trained to do this in a rigorous way, free of cognitive bias and errors of analysis. Doing this step well, like all others, takes thorough training and improves with experience, and it is one of the steps we will focus on in this class. Once the investigation is complete, we must document the conclusion and look for any improvements that could be made to do it better next time. This leads to the overarching process that must happen at all stages and times—continuous improvement fed by the outcome of investigations. Each has an important role to play in the process. A weakness in any step will affect the overall detection capability while strengthening any step improve the cycle. Once the incident has been resolved and lessons learned have been recorded, we move back up to triage to move on to the next item.

When it comes to a SOC analyst's daily life, especially newer ones, we tend to live in steps 3 and 4. Learning how to select an alert and then triage it is a large topic unto itself and requires learning attack progression cycles and attack tactics. Investigation is a large topic as well. Both are necessary for an analyst to know and are covered in this course. The best analysts, however, understand the entire process and understand the collection and detection steps that are often tasks for other job roles such as detection, content, or data engineers.

Understanding the fundamental principles of why you see the data you see, how it is collected, how malicious things are identified by your detection tools, and how it is all aggregated for analysis makes you a more well-rounded analyst. The goal of this class is to explain all these processes and break down the required items to reason your way through each.

Licensed To: David Owerbach <0mamaloney0@gmail.com> April 28, 2020

SEC450 Goals

Goals of this class are to teach you:

- How the SOC collects data, and which types it collects
- Common attack tactics to look out for
- How to identify attacks and associated traffic and logs
- The SOC tools, information, and workflow for alert triage
- How to perform high-quality analysis and investigation
- How to make SOC life less painful and repetitive
- Set you up for success as a Blue Team member!

SEC450 Goals

Given these two important truths (that compromise will happen and that your company does not solely exist to be secure), be sure to keep that at the front of your mind throughout the week. The Blue Team will be at its best when it has a realistic view of what is possible and has a good working relationship with the business. Although most of this class will focus on inside the SOC data, processes and workflow, it's important to not lose sight of the big picture view of how the group fits within the organization at large. While that view is important and should not be forgotten, this class will focus heavily on the technical details, attack types, defensive measures, and in-SOC items toward the goal of creating a happy, well-functioning Blue Team set up to succeed at defending your organization!

Let's Keep In Touch

Author

- John Hubbard (@SecHubb)

Online

- **Slack:** See wiki for info
- #SEC450
- SANS (@SANSInstitute)
- CyberDefense (@SANSDefense)

Special Thanks to the SEC450 Reviewers/Instructors!

- Mark Orlando (@markaorlando)
- Mark Jeanmougin (@markjx01)
- Don Murdoch (@BlueTeamHB)
- John TerBush (@thegumshoo)
- Justin Henderson (@SecurityMapper)

Let's Keep In Touch

After class is over, be sure to stay in touch online via Twitter and in the SEC450 Slack channel! A link to the slack channel is provided in the home page of the class virtual machine.

A big thanks goes out to the course reviewers who have contributed their expertise and precious time to making SEC450 great!

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. **Exercise 1.0: Virtual Machine Setup**
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. **Exercise 1.1: TheHive Incident Management System**
8. Threat Intelligence Platforms
9. **Exercise 1.2: MISP Threat Intelligence Platform**
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. **Exercise 1.3: SIEM with the Elastic Stack**

This page intentionally left blank.

Lab Workbook and Digital Wiki

Workbook

- Permanent copy
- Usable for testing
- Doesn't use screen real estate

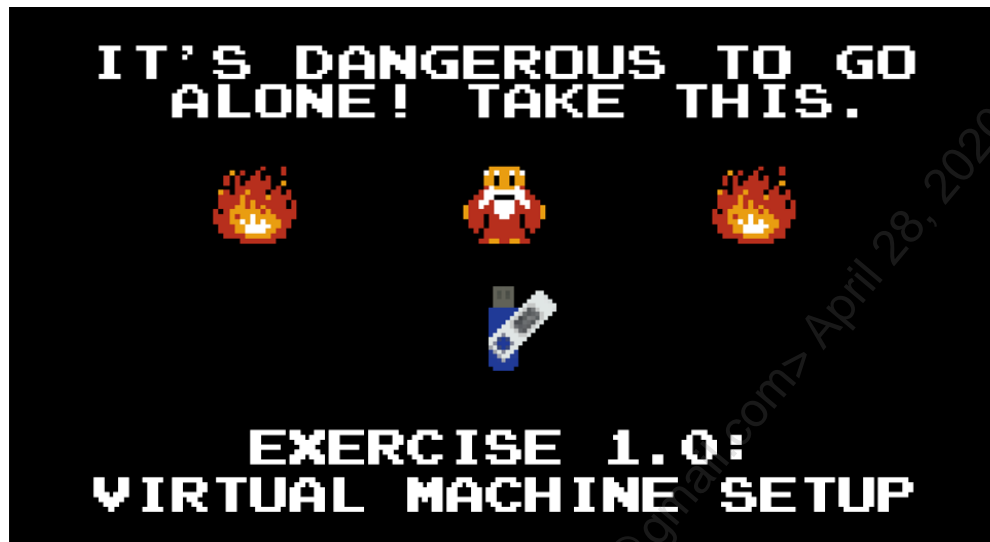
Digital Wiki

- Constantly updated
- Errata quickly fixed
- **Copy and paste**
- Clickable links
- Picture Zooming
- Lab Videos
- Reference material
- Instructor contact info

Lab Workbook and Digital Wiki

This class utilizes a digital wiki that is built into the virtual machine and contains all the information about the lab environment, cheat sheets, quick reference content, as well as digital versions of the labs themselves. Throughout the class, **we highly recommend you work from the digital wiki** instead of the workbook, as it allows the use of copy and paste for commands, which helps error-proof your experience.

Exercise 1.0: Virtual Machine Setup



Exercise 1.0: Virtual Machine Setup

Please go to Exercise 1.0 in the SEC450 Workbook. This lab consists of ensuring you have the lab environment up and running before diving into the rest of the course.

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. **SOC Overview**
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Licensed To: David Owerbach <0mamaloney@gmail.com> April 28, 2020

50,000 ft. Perspective

Before we dive into low-level, some important questions...

- What do we do and how do we do it?

In this section:

1. Core questions the SOC must answer
2. Identifying and meeting the business risk appetite
3. SOC org charts, roles, and functions
4. Critical information and documents to have on hand
5. Metrics introduction

50,000 ft. Perspective

Before we dive into the low-level details, it is worth having some high-level discussion on the mission of a defense team and how it should be derived, as well as how a SOC operates. In this section, we will discuss the "what we do" as well as the "how we do it" from a strategic perspective. It is important to understand things at this level in order to understand the unique goals and setup of your SOC. There are many valid ways to run and organize a cybersecurity operation, but each one will be (hopefully) optimized toward your own organization's specific goals. These goals should be defined and driven by your company as well as informed by where their SOC sits in the org chart, the functions it contains, and what the biggest perceived attack scenarios are.

Understanding the answers to these high-level questions ensures everyone on the team agrees on the top-level mission and the objectives and decisions that flow out of it. It will also ensure the business is happy with the team, sees the SOC's value, and continues to fund it. A team that ignores these critical lines of communication may find themselves pushed into obscurity, defunded, looked down on, or seen as the group to avoid talking to for fear of being seen as an impediment, which is a situation that will end poorly for all involved.

Four Core Questions¹

High-level direction setting questions:

1. What are we trying to protect?

Intellectual property, manufacturing process, personal data?

2. What are the threats?

Government-level attackers, organized crime, hackers, script kiddies?

3. How do we detect them?

Centralized network and endpoint data collection, continuous monitoring

4. How do we respond?

Quickly! With a variety of tools and a well defined (and automated) process

Four Core Questions

When it comes to looking at the mission of the SOC at the highest level, we can use the following four questions posed in "Crafting the Infosec Playbook"¹ to guide us in deriving the Blue Team's mission.

- What are we trying to protect?
- What are the threats?
- How do we detect them?
- How do we respond?

It is highly recommended you review and reconsider these questions periodically. The answers and capabilities of your team will change over time, as will the operations and needs of your organization. On the next slide, we will discuss how to codify these items into a document and make sure they stay up to date with inputs from a cross-functional committee.

[1] Crafting the Infosec Playbook: <http://shop.oreilly.com/product/0636920032991.do>

Aligning the SOC with the Organization

First, you need a **charter** approved by management describing:

- Constituency served
- Services to be delivered
- High-level mission statement
- SOC Scope
- Organizational structure

Steering committee

- Enumerates risk concerns from the business
- Aligns SOC capabilities and performance with business needs


Aligning the SOC with the Organization

One important item that a SOC must have early on is direction on what exactly it needs to focus on doing. This information comes from the SOC charter. The charter legitimizes and gives the Blue Team the authority to perform its duties within the environment. It is a list of the team's scope, duties, constituency to be served, and high-level mission that is blessed by management. The SOC charter is a document that everyone within the SOC should be familiar with to ensure everyone is on the same page about the purpose and scope of the group, how it is organized, and how it operates.

Another important aspect of answering the four questions on the previous slide is a steering committee. The steering committee helps continuously align the needs and risks identified by the business leaders to the capabilities and performance of the SOC and should help create the charter. The steering committee is a more formalized outward facing function of the SOC that ensures communication lines stay open, that the SOC is performing at the speed required of the business and is focusing resources on the appropriate risks. It is how we, as a Blue Team, speak with our constituency and make sure we're providing the services needed, and not just making up what we think is the appropriate defense. These two high-level items provide the initial and ongoing direction for what the SOC should be doing, who it protects, and how.

Finding the Organizational Risk Appetite

- Remember to consider the big picture
 - Organizations don't exist to be secure
- **Where is security on the priority list?**
 - Government/Military: Highest importance
 - New startup company: Low importance
 - Ask if your org. has a "**risk appetite statement**"
- **A mature security team understands the appetite**
 - Works within it (this doesn't mean you can't try to influence it)
 - It will change as management, company, and priorities change



Finding the Organizational Risk Appetite

One of the items that should be made clear through the steering committee is the organizational risk appetite. Everyone in the SOC should understand well what the business views as its biggest risks and how that translates to how it approaches its mission. Different organizations will have wildly different security priorities driven by the nature of the organization. Typically, government and defense-related organizations will have the lowest risk appetite and be willing to implement more intrusive controls to ensure breaches are minimized. Companies just getting off the ground may look at security as a function that should be done at a minimal level to ensure they can operate at the highest possible speed.

The Blue Team's job is to take this information as input and work within or inform management when the risk has been misjudged so that a proper adjustment can be made. Be aware that the risk appetite will change over time naturally as the organization matures, as well as when leadership changes. Keeping a finger on the pulse of the current organizational thinking is one of the important things the Blue Team must do. Some organizations keep a formalized "risk appetite statement." Reading this document (if it exists) should help make the thinking of upper management clear.

Meeting the Risk Appetite

- Consider the nature of your business
 - How damaging is it for a breach to occur?
 - What's the worst that could happen?
- **High security:** Whitelisting, highly segmented network, host isolation, strict email attachment policy
- **Low security:** AV, IDS, tools in detection mode
- **Your goal:** Find ways to crank security as high as you can *without hindering business process*
 - Find the "**Goldilocks zone**" – not too much or too little



Meeting the Risk Appetite

One of the driving items informing the SOC charter and the steering committee will be your organizational risk appetite. Is your company absolutely unwilling to tolerate a breach and will it take extreme measures to protect itself above all else? Or is it willing to go a little more relaxed and take a balanced approach to ensure that business gets done quickly?

Consider the questions above. Although ideally this information will be handed down from above, sometimes the answers for these questions can be somewhat self-evident. Stop for a minute and consider what is the worst possible thing that could happen to your organization: Is it a data breach or disruption of critical infrastructure you run? Consider how likely those things are to occur, and what it might take to get to that point in terms of access. Then consider the controls your company has in place to get in the way of that scenario. One goal of the steering committee is to help produce the answers and determine how much money and time the business is willing to invest to stop these things from happening. The SOC will then take these directives and convert them into controls that will need to be in place. Those who have a low tolerance for a breach will likely find themselves subject to strict controls such as whitelisting, highly segmented networks, isolated or air gapped system and strict lockdown of internet and email. Those with a high-risk tolerance may be OK with Antivirus, IDS, and basic email filtering and spam detection.

Your goal as a Blue Team member is to ensure these things stay in line. If the situation is out of alignment, you should feel empowered to make recommendations where there are gaps in coverage and help move toward the "goldilocks zone" of controls—just the right amount to get the job done at the risk level designated, without going overboard and ruining productive activity with overbearing security.

Risk Appetite Meets Reality

You work for a vaccines company:

- A vendor-built PC runs a critical production line
- Highly qualified build, **no extra software allowed**
- Requires outbound FTP transfer, inbound web status page
- Operating System is Windows XP, no updates coming
- **Consider: How do you secure this machine?**
 - Hint: The answer is not "don't allow it to be used" – you'll quickly be shown the door with this approach



Risk Appetite Meets Reality

An example situation you might run into where the level of risk desired may be at odds with the realities of the job you must do: Let's say you work for a vaccines manufacturing company and part of that manufacturing process takes a *highly* specific computer setup that has to be formally approved for all changes due to its sensitivity as part of the process. The build is based on Windows XP and cannot be modified beyond the stock software, so you cannot add monitoring or firewall software to the host, not even antivirus protection. In addition, you also can't update the operating system because the company went out of business and there is no alternative, so you're really stuck here. To make matters worse, this machine also needs to record data that it will send out via the insecure FTP protocol; and additionally, it hosts a status page via an old, likely insecure web server. This seems like a contradiction—you must be safe, but you also must run Windows XP on a machine performing a critical task. What do you do?

Some might be tempted to tell the business that this situation is so risky that it would be insane to continue to use the machine in any capacity at all! Be prepared to be rejected if this is the conclusion you propose. The answer you would likely receive to that response is "absolutely not, that machine is responsible for \$x/hr. of product, find another way!"

What else can we do? Our requirements are that we do not touch the inner workings of the machine; therefore, solutions would likely involve external security appliances applied as compensating controls. These would need to scan for viruses, block all unnecessary ports, and whitelist communications to and from the machine down to only sources that need to communicate with it. In addition, web application firewalls could be used to protect the web application. These are the types of complicated real-life situations you will encounter as a Blue Team member, and you will need to accept the situation, address the risk and requirements, and suggest a solution as well as interpret the data generated in order to detect a compromise under such a setup.

Accepting the Risk



SANS

SEC450: Blue Team Fundamentals: Security Operations and Analysis

32

Accepting the Risk

The concept gets joked about a lot in information security.¹ Do remember though, as badly as we want to make the perfect most impenetrable network, it's common that management wants to accept an identified risk. Cybersecurity is one piece of a complex business pie, and although security may be *your* entire world, to someone else it's just one factor in an equation of items to worry about to keep the business running.

Ultimately, the best we can do is give management accurate and complete information to do their job, which is making sure the business continues to run. If you disagree with a decision that is being made and think it may be overly risky and lead to issues down the road, the best you can do is try to communicate why in a more effective way or document the advice that was given and continue with your day.

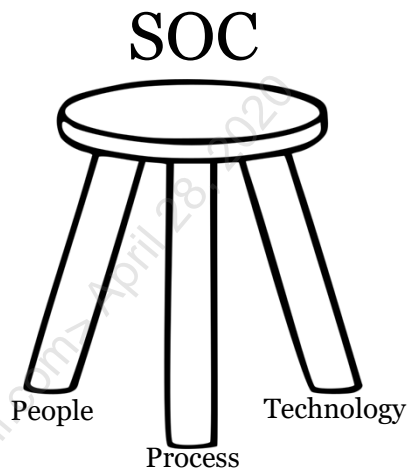
[1] "Host Unknown presents: Accepted the Risk": <https://www.youtube.com/watch?v=9IG3zqvUqJY>

The Components of a Blue Team

People: Performing analysis and investigation, design and run processes

Process: The defined sequence of events performed to achieve an end goal

Technology: Hardware and software used to accomplish the mission



The Components of a Blue Team

A fully functioning Blue Team requires three core components: People, process, and technology. First and foremost are the people. Without an engaged and well communicating team, no Blue Team can operate. The people are the heart and soul of the Blue Team operation. The selection of people on the team can single-handedly make or break it. No process or technology will be able to make up for a dysfunctional, untrained, or unhappy team.

Second is the process, which defines what exactly those people will be doing with their time and how they do it. Process definition involves identifying how tasks need to get done and how they should be performed in the Blue Teamer's day-to-day life.

Finally, technology is the enabler of efficient and well-defined processes. Technology allows us to monitor vast amounts of data at scale and is a force multiplier for our team members. If you're doing defense right, technology is NOT a replacement for people—it merely makes them better. If your analysts can be replaced by technology, they likely were doing the repetitive work that would've been better handled by automation in the first place, not doing the type of work they should be doing—analysis. These three components make up the "three-legged stool" of security operations. Without each one, the team will not be able to succeed.

People: The Most Important Component

- Interest, curiosity, self-motivation and passion
- Practical IT experience with diverse background
- "T" shaped people: Broad with at least one specialty
- Critical thinking: Analysis and synthesis skills
- Scripting and automation skills
- Soft-skills: Writing and presenting ability
- **Quality** over quantity
 - Remember: **Talent attracts talent**

People: The Most Important Component

What skills does a Blue Team look for when finding new members? One of the themes that will always appear is "quality over quantity"—the fact that one self-motivated, passionate, and capable individual dramatically affects the capabilities of a team through automation and a continuous improvement mindset.

The ideal analyst has a broad range of experience and the ability to go deep on at least one technical specialty. Ideally, they can also code and write well enough to make a convincing case about how or why something should be done. Another way of putting this is the "T-shaped" person as mentioned in Carson Zimmerman's excellent (free) book, *Ten Strategies of a World-Class Cybersecurity Operations Center*¹. This is someone who has a broad knowledge of IT and technology in general, as well as the ability to dive deep and possesses a unique talent that brings that capability to the team. One key thing to remember is "**talent attracts talent.**" The more unique skill and talent you can cultivate within the team, the easier it will be to draw additional skilled people in the future. People like working with coworkers who will challenge them and help them continue to grow, so do not overlook this important factor when building out your security operations group.

[1] <https://www.mitre.org/sites/default/files/publications/pr-13-1028-mitre-10-strategies-cyber-ops-center.pdf>

SOC Roles and Duties

Analyst: Investigate and triage alerts, incident response

Threat Intel: Collect info for tactical and strategic advantage over adversaries

Engineering and Infrastructure: Design and implement SOC infrastructure

System Administrator: Care and upkeep of SOC tools

Manager: Responsible for work prioritization and communication upwards

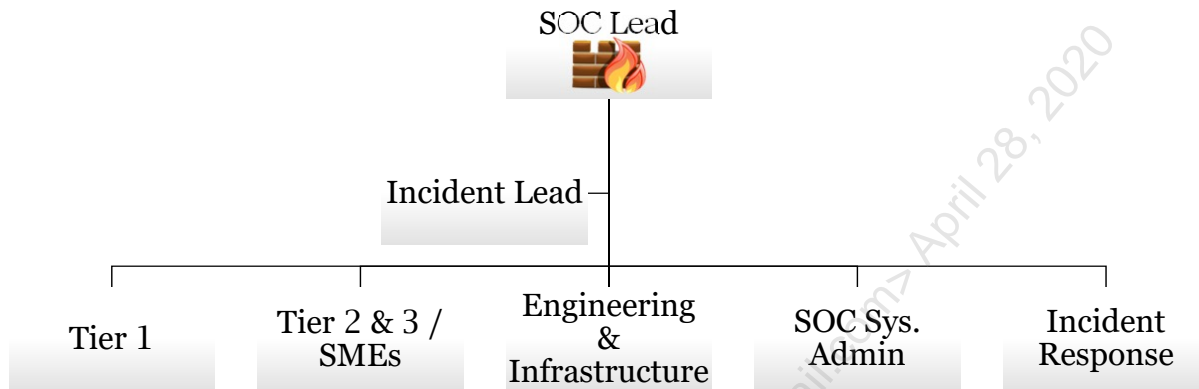
Incident Lead: Designated coordinator and communicator

SOC Roles and Duties

A SOC is made up of multiple vital roles, and they will not necessarily match up with each of the functions discussed in the previous slide. Whether or not you can staff dedicated roles for things like malware reverse engineering and threat intel will depend highly on the size of your defense team, but usually only in large environments are these roles separate. In many smaller organizations, people will perform multiple duties and must switch "hats" depending on the day and situation. Analyst jobs will typically cover jobs such as the command center, NSM, incident response, and forensics, but there is much more to it than that. Threat intelligence is a significantly different capability and skillset than the other alert triaging and incident handling type functions. Therefore, this is one of the first "systems" in the SOC to get its own dedicated role. In addition, we need people to design and architect the SOC-specific infrastructure. Engineering and infrastructure type roles bring this capability into the SOC. After systems are installed and configured, they also must be kept running, which is where system administrators also are brought in to help.

Beyond all the technical roles is also SOC management or even individual managers of any of the sub-groups within the SOC, depending on the organization's size. Managers help set direction and drive the operation of the SOC, but also importantly act as what I've heard described as a "bi-directional distraction firewall." Their job is to ensure those up the chain do not distract those in the SOC with outsider requests that can side-track them. They also distill the internal happenings of the SOC and communicate upwards the successes and issues the SOC has identified, eliminating the unneeded detail (we've all heard of the executive presentation where you are allowed one single slide). Doing this is a full-time job, and in some cases requires even more people during an incident. Incident leads can act as a dedicated information runner, keeping all concerned parties updated with current information in the middle of a large breach. Believe me when I say that in the heat of battle, someone doing this role can be a life-saver. Incident leads will take the distraction of communication and update writing off the analysts, which allows them to keep churning through the case at top speed.

How Are We Organized?¹



How Are We Organized?

Here is a chart showing a common organizational chart for a SOC. This is based on the outstanding free book *Ten Strategies of a World-Class Cybersecurity Operations Center* by Carson Zimmerman. This chart highlights the fact that there is much more to running a SOC than analysts doing triage. Operating such a highly technical capability in a large organization also likely involves the support of dedicated people doing content engineering—writing signatures and other detection capabilities for the tools, as well as the system administrators that keep it all up and running. Ensuring new appliances get architected, installed, and configured, as well as keeping the current technology up and running can require multiple full-time roles. As previously mentioned, threat intelligence and incident response are often also dedicated roles that work either within or very close with the SOC today.

[1] "Ten Strategies of a World-Class Cybersecurity Operations Center"

<https://www.mitre.org/sites/default/files/publications/pr-13-1028-mitre-10-strategies-cyber-ops-center.pdf>

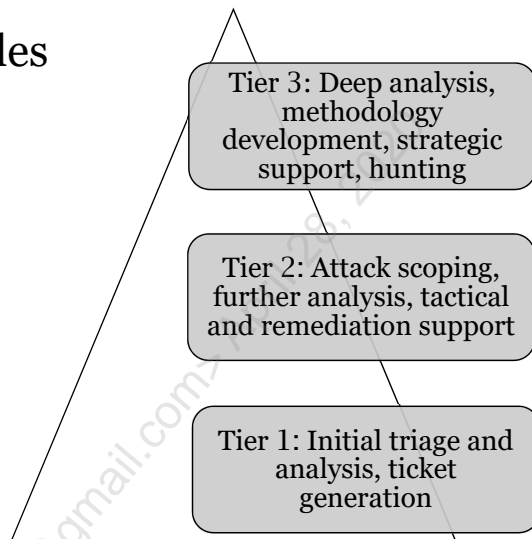
Tiered SOCs

Many SOCs have tiered analyst roles

- Tier 1: Learning the ropes
- Tier 2: Increasing capability
- Tier 3: Highly complex tasks



Typical duties for each tier...



Tiered SOCs

When it comes to analyst roles in the SOC, many teams break down into a tiered structure where tier 1 analysts represent an entry level role. As individuals gain experience, they may be promoted up the ranks into a higher tier. Tier 1, as an entry level role, typically involves a more highly defined process and tasks that will help newcomers understand the rules of the SOC and the data collected. Unfortunately, at least from the analyst's perspective, it may also involve restrictions on which tools an analyst is allowed to use, and which data can be viewed.

There are valid arguments both for and against tiers. Having a restricted tier 1 role helps focus learning and removes the temptation to try to use data that might further confuse issues. On the plus side, tightly controlling process and who can do what ensures the SOC is likely running at peak efficiency, and everyone knows exactly what is expected of them. The downside is that analysts who feel they can't see or use the "good" or "fun" tools and data may become quickly frustrated. We know that retention is a problem in many SOCs; therefore, exercise caution and careful consideration when assigning job tasks to tier 1. Overly restrictive tiers can lead to a situation where people may leave out of frustration if they don't get promoted as fast as desired. No one wants to operate the "revolving door" SOC.

As analysts gain familiarity with workflow, tiers 2 and 3 typically involve increasing amounts of freedom, less process and more complex tasks. These challenging tasks often top out in activities like malware reverse engineering and memory forensics—specific and niche activities that require high levels of expertise. The increased freedom is an acknowledgement that an analyst can be trusted with more dangerous files such as malware, or more sensitive incidents. This progression can be a great motivator and a highly efficient way to run the SOC when done correctly.

Tierless SOCs

Tierless:

- Everyone works together to get everything done
 - Must carefully manage alerts
- Even new analysts can use all available data and tools
 - Stay engaged, learn quickly
 - ...but must know limits
- Analysts more self-guided, teamwork crucial
- Senior and Lead titles for career progression

Tiered:

- Defined roles, clear path for promotion
- More structured processes, efficient handoffs and processing
- Often have less freedom to use all tools / data restrictions
- Less ability to explore and learn?
- Slow progression, repetitiveness may lead to retention issues

Neither is "right" – just optimized for different things

Tierless SOCs


Talking with many SOC analysts throughout the course of my career, I have found that many analysts become very frustrated with tiered SOCs. The reasons vary, but often it has to do with overly repetitive tasks or feeling held back despite their ability to take on more complex tasks. Therefore, the tierless operating model might be better for some teams. In tierless SOCs, analysts are generally given more freedom to learn and explore all the data and tools available without an artificially imposed ceiling. Anecdotally at least, it seems that analysts from these environments seem happier with their jobs and there is some research presented in Day 5 to back this up. Although this may seem less defined, is it worth the potential risk? Tierless SOCs, although perhaps less efficient, may boost retention over time, giving the SOC the ability to build and retain the talent that is required for highly complex tasks.

Tierless SOCs are not meant to be chaotic, and care must be taken when operating a SOC in this manner to ensure everything is still being done in a dependable and repeatable way. In this mode, it is expected that everyone will collectively be able to get everything done, and that everyone shares in the responsibility and is expected to contribute at the level they are capable of. To do this, everyone must know their own limits and be comfortable asking for help when they are reached. The benefits of this operating model are clear however, newer analysts are exposed to more techniques and complex analysis quickly and can take on more tasks as soon as they are ready, without having to wait for a promotion to be available.

Deconstructing the SOC Process and Technology

To more easily understand SOC functions, we will simplify them

- Abstract tools/functions into a simple “box” with **inputs, outputs**
- Deconstruction into inputs, outputs, and internal process shows how each item relates to each other



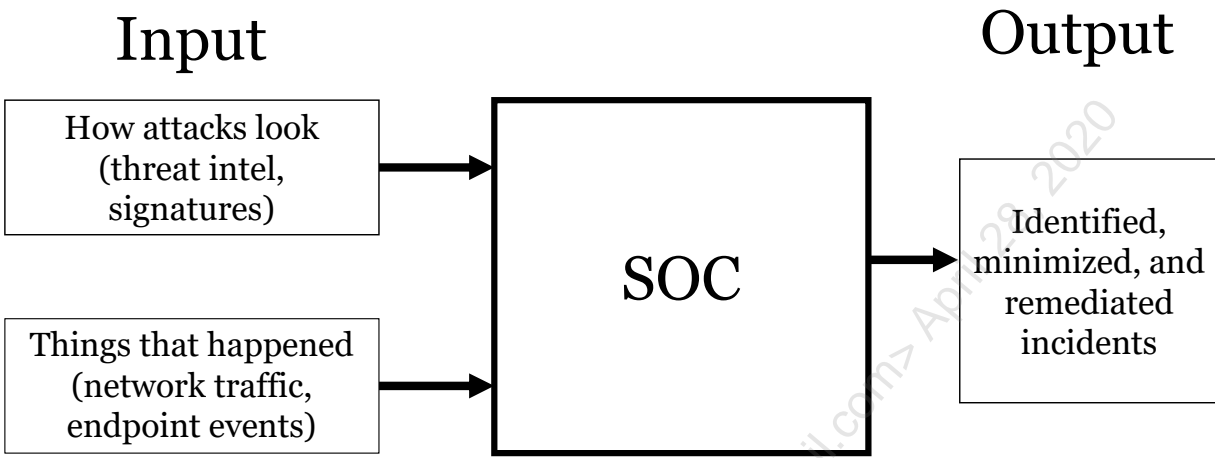
Deconstructing the SOC Process and Technology

One of the challenges for new SOC analysts is understanding how all the data they must deal with flows between all the tools in the SOC. While a SOC is a highly complex interconnection of many tools and processes, the operation can still be understood if you start of small and build upwards. To do this, throughout the course we will deconstruct and abstract complex systems into an individual set of inputs, internal processes, and outputs. Doing this makes it clear what type of input each system takes in, the internal processes that act on the input, and what it is expected as produce as output.

Take an intrusion detection system for example, while it is true this is a highly complex machine, we can generalize what it is doing to simplify it for quick understanding. At its essence, an IDS takes as input network traffic forwarded to it and a set of signatures for specific traffic contents it needs to look for. Internally it finds traffic that matches a signature from the input list in the given traffic and outputs alerts that will ultimately be forwarded to the SOC alert queue. We don't need to care too much about exactly how the IDS does its job to understand its role. We simply must understand that the output is a function of the inputs, which are traffic it can see, and the number and type of signatures that are active. Simple, right?

Once each individual function of the SOC is understood in this way it is much easier to recombine them and see the bigger picture of how they all work together. Through this method we will work to understand how the SOC can and should use each of the tools at its disposal.

The SOC Abstracted



Better output requires better input, “*garbage in, garbage out*”

The SOC Abstracted

Although building up from the bottom and understanding all the pieces is useful, we can also learn from the top down, abstracting away all the details to make the mission clear. What if we abstracted the entire SOC? At a very high level, what are the SOC's main inputs and outputs? As the slide above shows, we could generalize it by saying the SOC is a function that takes both the things that have happened in the environment and what attacks look like as input, and outputs identified, minimized, and remediated incidents. As with any system, if we want better output, we need better input (this is where the phrase “garbage in, garbage out” comes from). For a SOC, that means either more visibility or better knowledge of what an attack looks like. The “SOC” box, of course, is an abstraction of an extremely complicated set of interactions, but at its most basic, that's what we do, and those are the main variables that control it. It is through looking at our tools and systems in this highly abstracted view that, although imperfect, helps us understand the system and identify the levers that can be pulled when something isn't working as well as we'd like.

SOC Process and Technology Functions

Beyond **people**, we have **process** and **technology** to facilitate:

Core SOC Activities

- **Data Collection:** What's happening on the network / devices
- **Detection:** Identifying items of interest from data collected
- **Triage and Investigation:** Confirming and prioritizing detected issues
- **Incident Response:** Responding to and minimizing the impact of attacks

Specialty / Auxiliary Capabilities

- **Threat Intelligence:** Collecting information to improve attack detection
- **Forensics:** Supporting I.R. with deep research and reverse engineering
- **Self-Assessment:** Inventory, config monitoring, vuln. assessment, Red Team, etc.

Functional Components of a SOC

Here, the higher-level functions of a SOC are broken down into the items that would more likely be considered “core SOC” functions vs. the specialty fields. Core SOC items are more likely to be performed by SOC analysts or those that work very closely with them. Small teams may even have all these functions performed by the same group of people. The other specialty fields involve significantly different enough skillsets that they typically require another person to perform the duty. While large organizations may have dedicated forensics, threat intel, and penetration testing teams, smaller organizations often outsource these capabilities for cost efficiency.

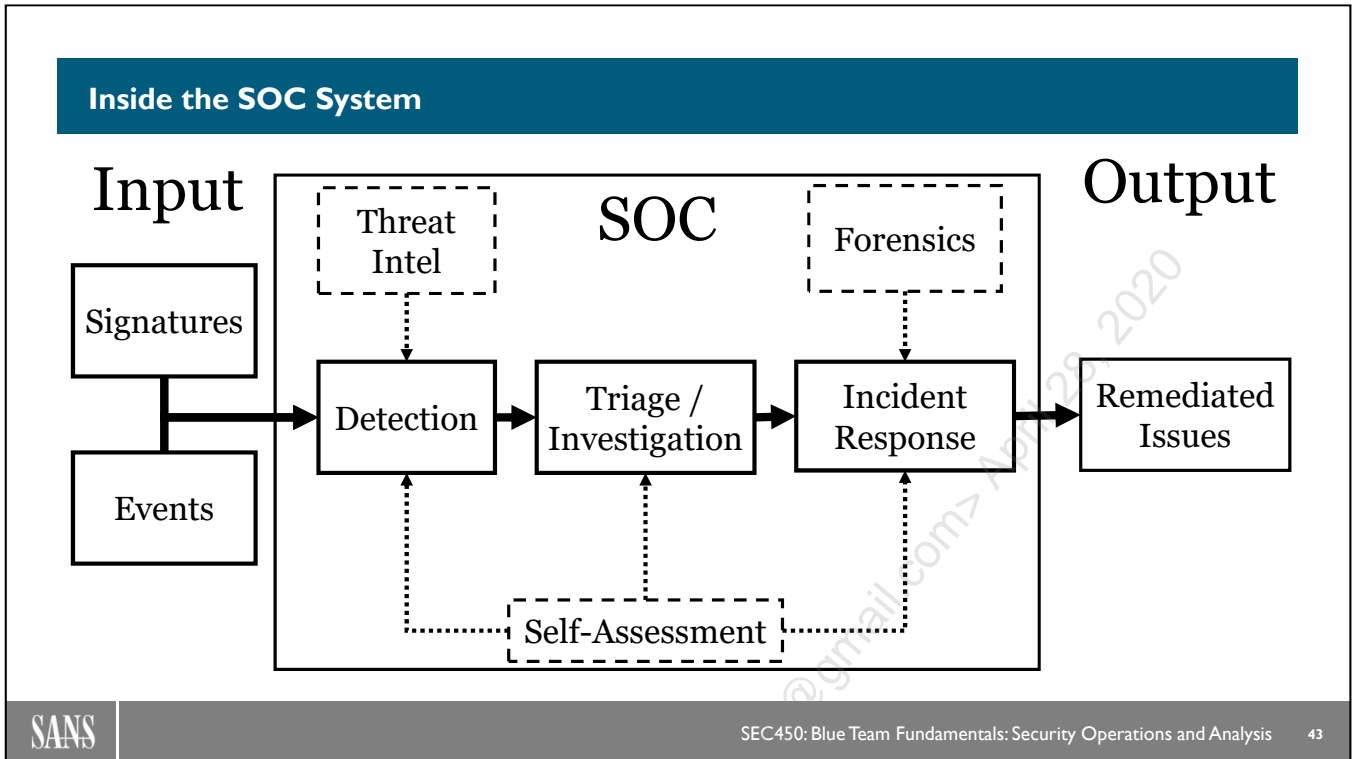
Core SOC:

- **Data Collection:** This is the technology and processes that enables us to understand what is occurring on both the network and the endpoints. As we will later discuss, this breaks down into the practices of Network Security Monitoring (NSM) and Continuous Security Monitoring (CSM).
- **Detection:** The goal of the detection function is watching the data collected from the network and endpoints and accurately identifying any potential compromises. This can be thought of as what your network and host IDS systems are doing, anti-virus, SIEM analytics, and anything else that watches everyday events and outputs alerts of possible compromise.
- **Triage and Investigation:** The triage and analysis functions are where all the identified alerts go to get prioritized and verified. Since in nearly every SOC there will be many potentially malicious events identified, it is the primary job of the SOC analyst to sort through them for criticality and verify whether an attack has indeed occurred.
- **Incident Response:** The incident response area is responsible for reacting to problems that are verified and ensuring the impact of the issue is minimized. In smaller SOCs, this falls under the scope of the analysts; in others, this may be a separate group called the CIRT or CSIRT. Regardless of the org structure, incident response is typically considered a core function of the blue team.

Specialty and Auxiliary Capabilities:

- **Threat Intelligence:** The mission of the threat intelligence group is to collect detailed high and low-level information on attack groups interested in the organization. The goal is to help give the Blue Team a tactical and strategic advantage over the attacker. If we can anticipate attacker goals, moves, and infrastructure ahead of time, it will be much harder for adversaries to accomplish their mission.
- **Forensics:** A specialized function focused on determining exactly what occurred during a breach. This may be traditional hard drive forensics, or something more specific such as memory analysis, malware reverse engineering, or even eDiscovery.
- **Self-Assessment:** This name is an umbrella term for multiple functions that may or may not be considered as directly within the SOC. This group contains things such as configuration monitoring, vulnerability assessment, penetration testing, and red teaming, and inventory. These activities are all similar in that they help the blue team perform their job effectively by either watching for potential issues (vulnerability management), or test the blue teams reaction to simulated threats (penetration testing and red teaming). Regardless of its position in the org chart, it is a critical piece of the security puzzle.

Licensed To: David Owerbach <0mamaloney0@gmail.com> April 28, 2020



Inside the System

If we were to take the simplified diagram of the previous slide and example the details of what’s going on inside the box for the SOC, what would we see? We could take the functional components previously mentioned and draw them all connected as well, each having their own inputs and outputs. Looking at this slightly more detailed version, the bold boxes show which components can be thought of as core to the SOC – collection detection, triage, investigation, and incident response. The dashed boxes are functions that help enhance the capabilities of the other functions and help us to get better results out of them. Self-assessment for example, is the function that tests our detection and response capabilities by simulating adversary activity and attack tactics. Forensics helps us understand our incidents in more depth by finding out the truth of what happened on an affected endpoint, what the capabilities of malware were, or what data has been stolen. Threat Intel helps sharpen our detection capabilities.

Throughout this class, we will inspect these individual functional components and the tools that help perform them, looking at what their inputs and outputs are, as well as what is inside them as well. Through this exercise, the goal is to ensure you are aware of now only how each individual system works, but how they interconnect into what you see here.

Critical SOC Information

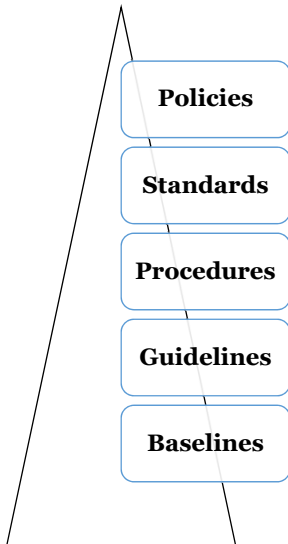
- **Network diagram:** Simplified version for easy reference
- **Points of visibility:** Taps and span ports, full PCAP
- **Data flow diagram:** How does traffic reach the internet?
- **Log flow diagram:** Where do logs come from/go?
- **Incident Response Plan:** What to do when things go wrong
- **Communication plan:** Who to inform, and when
- List of critical assets and points of contact
- Disaster recovery / business continuity plans
- Any other relevant policies, standards, procedures, guidelines

Critical SOC Information

In order to perform the main SOC functions, there are some important pieces of information and inputs for the SOC to have on hand. One item is an overall view of how data traverses the network, and where the SOC has points of visibility. Often, in an incident, the question "would we have seen that traffic?" must be answered. Without a firm grasp on how data flows over the network as well as where it can be seen, this will be a hard question to answer and can lead to confusion and slow down progress. Ideally, this information would include how both internet and internal network traffic is flowing, and where the SOC can see NetFlow, full PCAP, or logs from devices that would report any suspicious activity.

In addition to this info, an incident response plan, detailing the steps that should be taken in the case of a major incident, can keep things running smoothly during a disaster. A communication plan should also be in place, so time is not wasted looking up who to contact and how to contact them. A list of critical assets is of utmost importance as well. It's hard to defend your most important data or even recognize it's in jeopardy if you don't know where it is and the systems it sits on. Having a prescribed process for these items, at least in the case of major incidents, should be a priority.

Documents Analysts Must Be Familiar With



Policies: High level, broad, direction setting, mandatory

- "All systems plugged into the network must have antivirus installed"

Standards: Also mandatory, define how or how much

- "Configuration settings for antivirus agents must be..."

Procedures: Step-by-step instructions for a process

- "How to install and ensure antivirus is working"

Guidelines: Discretionary, suggested actions/recommended procedures where actual standards and procedures do not exist

- "Best practices for antivirus deployment"

Baselines: Highly specific settings list (CIS benchmarks)

Use Case / Playbook: SOC Specific prescriptive rules/procedures for detection

Documents Analysts Must Be Familiar With

One item that can be confusing for newcomers is the difference between all the various types of documentation involved in running a SOC. The types of documents you will likely run into may sound similar at first glance but do have content that can be distinguished from the others. These document types are listed here as well as their generally recognized purpose.

- **Policies:** Policies are high-level, broad, direction-setting documents that do not go into specifics, but lay out the general requirements for configuration items. They generally answer the "what" must be done. Items laid out in Policies should be viewed as mandatory.
- **Standards:** As opposed to Policies, standards give more specifics in that they specify "how" something gets accomplished or how much of something should be applied. Standards should also be considered mandatory compliance.
- **Procedures:** Procedures explain the step-by-step instructions for completing a specific task. These can be thought of as the lowest level documentation in terms of containing lots of specific detail.
- **Guidelines:** Of the previously discussed document types, guidelines are the ones that are *not* necessarily mandatory. They lay out suggestions and recommended actions or procedures of configuration or other best practices.
- **Baselines:** Highly detailed and itemized checklists. A perfect example is the security benchmarks provided by the Center for Internet Security for security operating systems and applications.¹
- **Playbooks / Use Cases:** Whereas you will likely see the other types of documentation in use across many parts of an organization, this is a type of documentation potentially unique to the SOC. Definitions vary org to org, but, in general, use cases are a fully documented reason and conditional logic you have and can implement on a particular security appliance, often a SIEM. SIEM use cases may be something like "detecting brute force login attempts" or "User attempting to access unauthorized resource." Sometimes, this term is used interchangeably with playbooks² and sometimes not. The author has seen playbooks

used as a term for something like procedures, but incident response or threat hunting specific in that they are "plays" to perform a certain action that will highlight potential evil. They also could be the logic and flow of information as implemented by an automation platform. Of the types of documentation listed here, these are the least well-defined, and it is likely worth following whatever convention is used in your environment.

[1] <https://www.cisecurity.org/cis-benchmarks/>

[2] Cisco's definition of a playbook: <https://blogs.cisco.com/security/using-a-playbook-model-to-organize-your-information-security-monitoring-strategy>

Licensed To: David Owerbach <0mamaloney0@gmail.com> April 28, 2020

Ensuring You Stay Funded

Outward communication is vital!

Metrics: Before creating, ask yourself

- Are terms well-defined? Meaningful?
- Is it **Actionable? Repeatable? Automated?**
- What "lever" can we pull to change it?



Consider some commonly requested metrics, are these good? Why?

- Incidents by delivery vector
- Unique hosts with viruses detected
- % systems with logs collected / % of network traffic monitored
- Number of attacks blocked?

Ensuring You Stay Funded

While we do not have space in this course to include a full discussion on metrics, analyst should be aware that metrics are what keep them funded. Consider the items listed above as core requirements for creating good metrics. Any metric you want to track should be actionable, repeatable, and ideally collected in an automated fashion. Metrics that do not follow these guidelines may end up being useless, inconsistent, inaccurate, or too burdensome to create.

This slide lists some example metrics that organizations might ask for—some of them are good, and some of them are not. Consider them in the light of the above advice.

- **Delivery vector:** If we know the answer to this question, is it something that we can act on? Absolutely! If we know that 90% of our attacks are coming through email, that can give us the justification to spend additional money on implementing and sandbox solution, or perhaps using more restrictive filters for what file attachments are allowed, or if URLs are rewritten.
- **Unique hosts with viruses detected:** This is an interesting metric; it is more specific and gives us an idea of the rate that attacks are passing the exploit stage. It tells us how many incidents have progressed to the point that a malicious file was actually dropped on the hard drive, implying prevention mechanisms at the network level were bypassed. From this metric, you could do a deeper dive into the data and try to understand how the malicious files that were detected are getting there. USB infections? Poor web filtering? Lack of policy stopping employees from downloading executables? Finding the cause and implementing controls to stop it is the lever that can be pulled to change this number, and the effect can be readily seen and understood as a clear improvement.
- **Percent of systems with log collection / traffic monitored:** This metric is useful for multiple reasons. For one, it gives you an idea of the coverage the Blue Team has in terms of the ability to detect attacks at a network and host level. In case the number is low, it also gives you a solid number to show management and express your concern for the lack of visibility. The action to take based on this metric is clear—more network and log collection, and the investment put into doing so will become readily apparent.

Number of attacks blocked: This metric is one that is extremely common to ask for, the problem is that what constitutes an "attack" is very poorly defined. Is a firewall block an attack? Is a visit to a malicious website that was blocked an attack? What if an event didn't progress to the point where we could tell if it would've been malicious or not? It would take a horrendous amount of time to gather the data for this number in most cases, and therefore, this is one of my favorite examples of a bad metric. In addition to the poor definition, what action would you take if you did collect this metric, and found it was up 20% vs. last month? It is not an actionable piece of data as it is too vague. The underlying data is likely to confuse people; it has too many edge cases and too much gray area to be meaningful.

There is a whole world of possible metrics out there. The question is, what story are you trying to tell, and which ones will help you tell it? A final item on metrics. Remember that not all metrics are interesting or meant to be reported "outward." Your board does not care about how many malicious domains you blocked or your average response time for tickets. What people at that level care about is risk reduction. When presenting metrics, especially the higher up you go, remember to tailor them to your audience and leave out unnecessary technical detail if it is not warranted.

Licensed To: David Owerbach <0mamaloney0@gmail.com> April 28, 2020

SOC Overview

- Your mission: Identify and reduce breaches
 - The **charter** gives you the power to do so
 - The **steering committee** guides you
 - Helps drive which controls are deployed
- The SOC is a complex set of systems that help accomplish this
 - People, technology, processes with input and output
 - Cmd. Center, NSM, Threat Intel, IR, Forensics, Self-Assessment
- Org charts vary for SOCs – there is no "best" setup
- Critical info must be gathered, monitored, and understood



SOC Overview

Throughout this section, we have gone over multiple high-level concepts for understanding the SOC, starting with how its mission is defined. We have stressed that although we live in a world of information security, we must stay connected to the bigger need—that of reducing risk to the business and keeping in mind throughout our work that this is the true goal. Planning and execution of this goal is key to success as a Blue Team, and metrics are one of the main ways this can be communicated. As the saying goes, "that which gets measured, gets done." Measuring your performance not only helps you improve it, but also crucially helps tell the story of how the investment in the Blue Team is paying back.

SOC roles and org charts may vary widely across organizations, but that's not to say there is a single right answer. The point of introducing the roles, tiered vs. tierless SOCs, and potential organization charts for the SOC is to make you aware that there are multiple valid configurations, and each one optimizes for a different thing. Where your team should personally fall is a decision that can be derived from the needs of the company, the SOC steering committee, and those who work in it.

We've also introduced the functional systems within the SOC and discussed how high-level systems thinking can help us break down each individual process and the interfaces between them. This simplifies the complex nature of all the Blue Team processes and makes it easy to understand and improve upon them. We will continue to present processes in this view throughout the class, analyzing input, output, and what internally drives the systems we will be using day to day.

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. **Defensible Network Concepts**
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Understanding Defensibility

In this module:

- What makes a network "defensible"?
- What is Network Security Monitoring (**NSM**)?
- What is Continuous Security Monitoring (**CSM**)?
- What tools perform network and endpoint event collection?
- Event collection content vs. format
- The importance of network Layer 7 and endpoint visibility
- The importance of data centralization

Understanding Defensibility

This module will discuss some of the fundamental questions about how a network can and should be monitored. It will cover Network Security Monitoring as well as Continuous Security Monitoring and explain the difference between the two. We will also dip into collection of logs, network architecture for monitoring, and the importance of data centralization.

Defensible Networks¹

According to Richard Bejtlich¹, a defensible network is:

- **Monitored:** Network and host data is captured and centralized
- **Inventoried:** Knowing your network
- **Controlled:** Traffic ingress/egress, network connection access...
- **Claimed:** Owners of services known, operation of assets planned
- **Minimized:** System attack surface is reduced
- **Assessed:** Weaknesses identified, defenses tested
- **Current:** Patched and known vulnerabilities addressed
- **Measured:** SOC and IT measure progress against previous steps

Difficulty



Gives you the chance to **resist** intrusion

Defensible Networks

A very important concept to keep in mind is the answer to the question, "What makes a network defensible?" In other words, what are the standards that a modern network must meet in order to be defensible against a modern adversary? According to Richard Bejtlich's well thought-out blog post from 2008 (yes, it still applies as much as ever), without an environment that meets these specifications, cyber defense is very difficult.¹

"A Defensible Network Architecture is an information architecture that is:

- **Monitored:** The easiest and cheapest way to begin developing DNA on an existing enterprise is to deploy Network Security Monitoring sensors capturing session data (at an absolute minimum), full content data (if you can get it), and statistical data. If you can access other data sources, like firewall/router/IPS/DNS/proxy/whatever logs, begin working that angle, too. Save the tougher data types (those that require reconfiguring assets and buying mammoth databases) until much later. This needs to be a quick win with the data in the hands of a small, centralized group. You should always start by monitoring first, as Bruce Schneier proclaimed so well in 2001.
- **Inventoried:** This means knowing what you host on your network. If you've started monitoring, you can acquire a lot of this information passively.
- **Controlled:** Now that you know how your network is operating and what is on it, you can start implementing network-based controls. Take this any way you wish—ingress filtering, egress filtering, network admission control, network access control, proxy connections, and so on. The idea is you transition from an "anything goes" network to one where the activity is authorized in advance, if possible. This step marks the first time where stakeholders might start complaining.
- **Claimed:** Now you are really going to reach out and touch a stakeholder. Claimed means identifying asset owners and developing policies, procedures, and plans for the operation of that asset. Feel free to swap this item with the previous. In my experience, it is usually easier to start introducing control before making people take ownership of systems. This step is a prerequisite for performing incident response. We can detect intrusions in the first step. We can only work with an asset owner to respond when we know who owns the asset and how we can contain and recover it.

- **Minimized:** This step is the first to directly impact the configuration and posture of assets. Here, we work with stakeholders to reduce the attack surface of their network devices. You can apply this idea to clients, servers, applications, network links, and so on. By reducing attack surface area, you improve your ability to perform all of the other steps, but you can't really implement minimization until you know who owns what.
- **Assessed:** This is a vulnerability assessment process to identify weaknesses in assets. You could easily place this step before minimization. Some might argue that it pays to begin with an assessment, but the first question is going to be: "What do we assess?" I think it might be easier to start disabling unnecessary services first, but you may not know what's running on the machines without assessing them. Also, consider performing an adversary simulation to test your overall security operations. Assessment is the step where you decide if what you've done so far is making any difference.
- **Current:** Current means keeping your assets configured and patched such that they can resist known attacks by addressing known vulnerabilities. It's easy to disable functionality no one needs. However, upgrades can sometimes break applications. That's why this step is last.¹
- **Measured:** In the book, *The Practice of Network Security Monitoring*, Richard also adds the final "measured" step. This step involves having the Blue Team and IT in general measuring their own progress against the previous steps, creating a feedback loop to ensure the task is being completed."

Throughout this module, we will focus on monitoring. Understanding your data collection and monitoring system is the first part of orienting yourself in a SOC and becoming a great analyst! Once you understand what data is collected, how it is collected, and where it comes from, you will be able to intuitively know what sources you must consult to answer any given analysis question.

[1] <https://taosecurity.blogspot.com/2008/01/defensible-network-architecture-20.html>

Two Sides of Monitoring

Monitoring breaks down into **2 main areas**:



Endpoints



Network

Richard Bejtlich on network monitoring:

*"...deploy **Network Security Monitoring sensors** capturing session data (at an absolute minimum), **full content data** (if you can get it), and **statistical data**.¹"*

Two Sides of Monitoring

One of the first concepts in Richard Bejtlich's defensible network definition is "monitored." Monitoring breaks down into two main areas that we will discuss separately. One area is the network—the ability to identify files, services, and any other information being transferred across your infrastructure. The other area is endpoint monitoring. This involves being able to see what is running and happening on the hosts on the network.

[1] <https://taosecurity.blogspot.com/2008/01/defensible-network-architecture-20.html>

What Can You See on Your Network?

Consider your network data collection...

- Can you see high-level bandwidth statistics and traffic flow?
- Do you know which **ports** are actually in use?
- Which **services** are *actually* being used on those ports?
- Do you know which **domains** are being visited and by whom?
- Can you retrieve the **full packet data** from the transaction?
- Can you detect malicious encrypted traffic?

How do we get these answers?

What Can You See On Your Network?

Step back for a moment and consider what information you collect about your network traffic. If you asked about traffic flow volumes between any two IP internal or external addresses, would you be able to produce that information? What about looking at all ports that are in use, or the actual services and programs using those ports? Can you pull a copy of the actual traffic on the network? Do you know which domains are being visited? Are you decrypting or at least collecting transaction data for encrypted traffic? These are some of the things we will need to monitor for effective network-based defense.

Network Security Monitoring

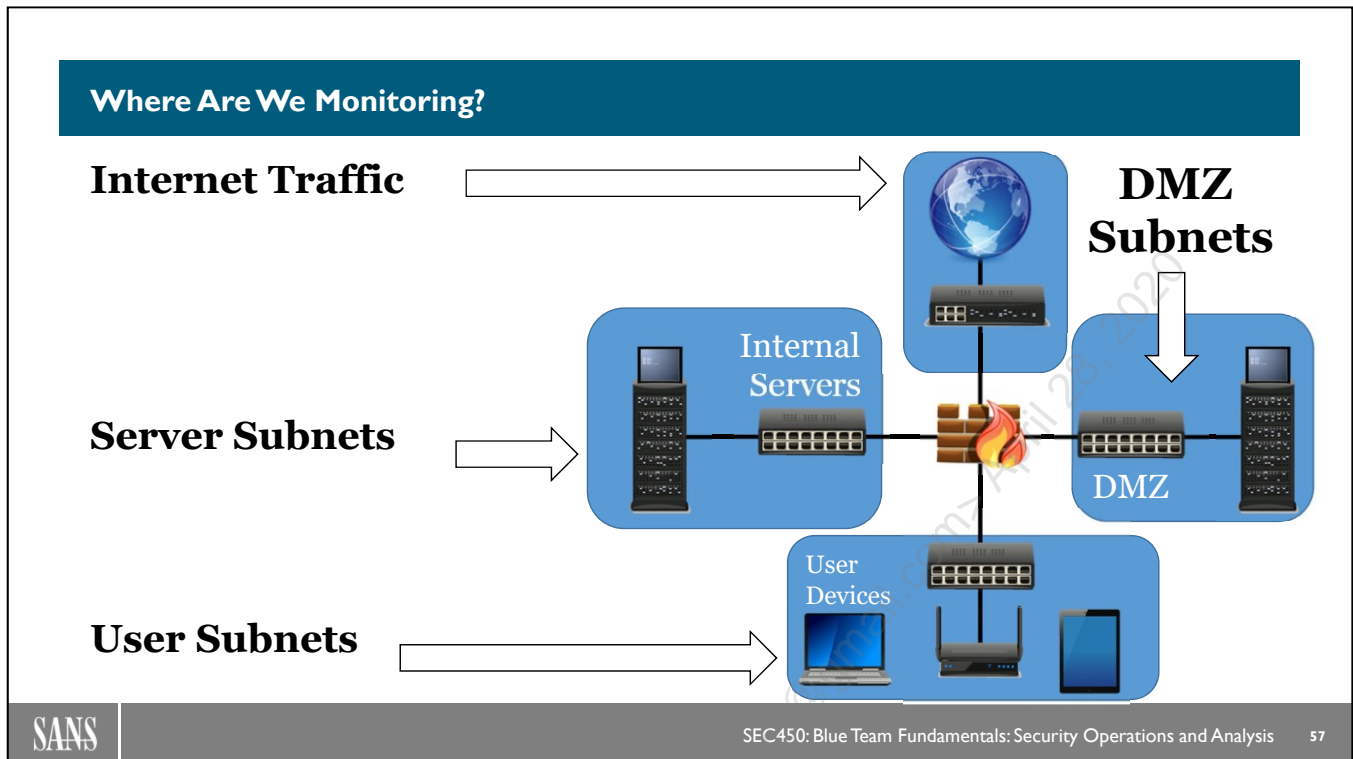
What is **Network Security Monitoring**?

- Analyzing network traffic – "data in motion"
- **Services:** DNS, HTTP(S), SMB, RDP, FTP, SSH, etc.
- Identifying risky / compromise-like behavior
 - Exploit delivery
 - Executable content transfer
 - Internal recon and pivoting
 - Command and control traffic
 - Data exfiltration



Network Security Monitoring

"Monitored" is the first item on Richard Bejtlich's list for good reason. If you can't see what's going on inside the network, it will be extremely hard or impossible to defend it. When discussing monitoring, there's one term you should be familiar with—Network Security Monitoring. Network Security Monitoring (NSM), as a term, is somewhat self-explanatory in that it relates to the analysis of everything as it is crossing the network. Monitoring traffic, recording interactions from machine to machine, and analyzing service logs for HTTP, DNS, and the like all fall under the definition of NSM. NSM is implemented through multiple different appliance types watching the traffic on the network and either recording what it sees, or specifically trying to point out evil. This means that in order to perform NSM tasks, we will need to capture traffic through taps, switches or other devices, and be able to analyze its content.



Where Are We Monitoring?

The real question is where *aren't* we monitoring? Ideally, we have instrumented sensors for NSM present in all our subnets, and endpoint-based info collection occurring for all the endpoints inside them. The slide depicts a very basic network showing the internet and a border router at the top, which leads into the border firewall, which in this case is doing all the traffic for the internal network as well. On the network, there are DMZ servers that are available to the internet as well as internal-only servers that are not. Finally, at the bottom, we have a subnet full of the users and their devices.

In a network such as this, we should ideally have every piece of equipment helping us with our monitoring tasks. The firewalls, routers and switches can send traffic information, and we could even use a tap to collect full PCAP information. The servers in the DMZ and internal segments, as well as the user desktops, should have a strong auditing policy with centralized logging to a SIEM platform. This would be a very well monitored network if every device was sending what it saw or what it logged to a place the Blue Team could search and alert based on that data. The next slide has a depiction of this collection.

NSM by Layer

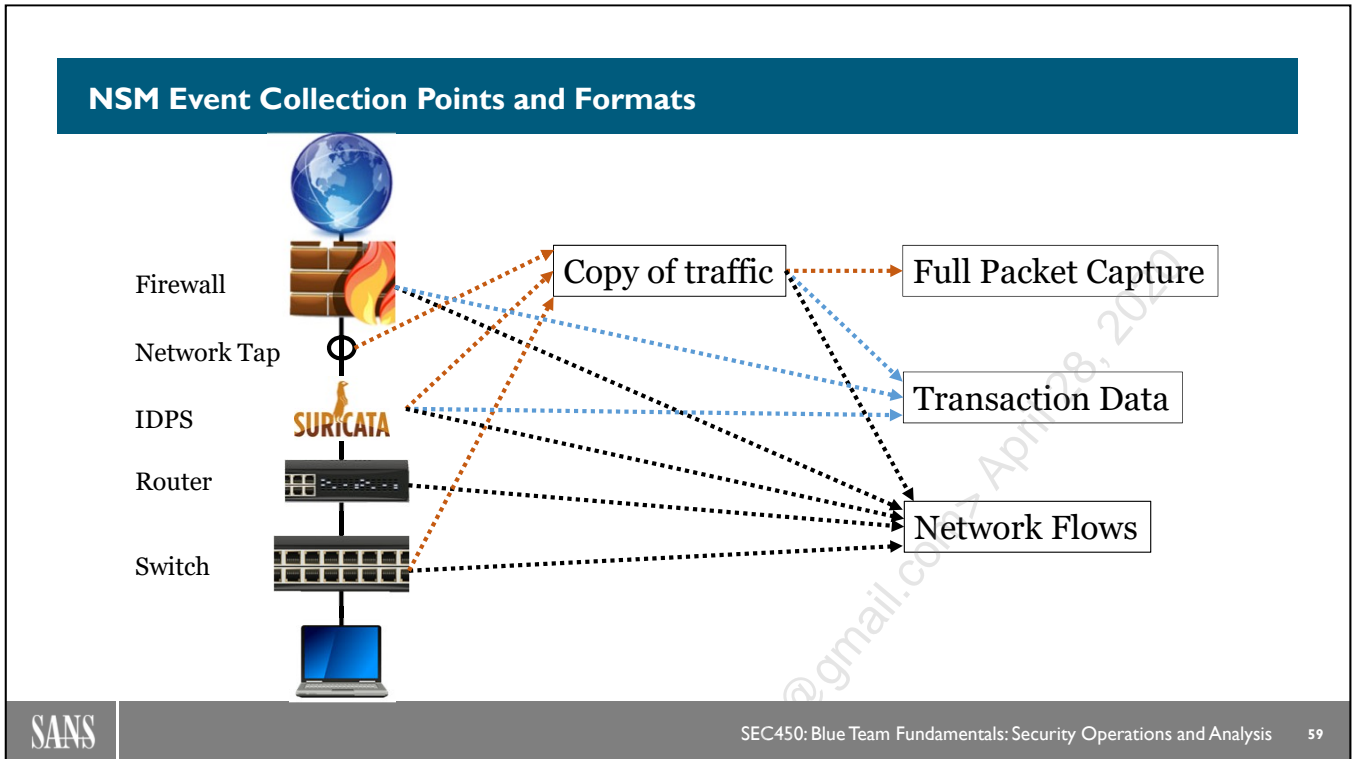
Low layer info can be useful, but we will need more...

- **Layer 3/4 (IP/Port)**
 - NetFlow, Statistical Data, Firewall Logs, almost anything
- **Layer 7 Transaction Data**
 - Service Logs, NSM Sensor Data
- **Layer 7 Full Payload**
 - Packet Capture, IDS alerts

```
▶ Frame 535: 531 bytes on wire (4248 bits), 531 bytes captured (4248 bits) on interface 0
▶ Ethernet II, Src: CadmusCo_69:e9:de (08:00:27:69:e9:de), Dst: RealtekU_12:35:02 (52:54:00:12:35:02)
▶ Internet Protocol Version 4, Src: 10.0.2.15 (10.0.2.15), Dst: 91.230.147.222 (91.230.147.222)
▶ Transmission Control Protocol, Src Port: activesync (1034), Dst Port: http (80), Seq: 1, Ack: 1, Len: 477
▼ Hypertext Transfer Protocol
  ▶ POST /aaa.php HTTP/1.1\r\n
    Host: google-analytics-sv1.com\r\n
    Content-Type: application/x-www-form-urlencoded\r\n
    User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1)\r\n
```

NSM by Layer

Each of our collection sources will provide a different depth of visibility into the traffic it sees. Most network appliances can pick up the Layer 3 and Layer 4 information (ports and IP addresses) associated with a given session, and sometimes that is enough for determining something is out of order (too many sessions or very large sessions, etc.). But for a complete modern defense, we absolutely must go deeper and be able to extract application layer information. Protocols like HTTP are routinely used for malicious purposes. Looking at the traffic purely from an IP and port perspective will not give an analyst enough information to identify it as evil. When HTTP is used maliciously, it will often be in a 100% protocol compliant way with the only clues potentially being URL patterns, hostnames that can be identified as odd, or event packet payloads that contain executables or other unexpected items. Given this, the necessity of having at *least* detailed transaction data captured at Layer 7 should be clear.



NSM Event Collection Points and Formats

Where NSM information can be sourced and what types of formats each produces is shown on this slide. As previously mentioned, nearly every tool can produce a record of network flows that have passed through the device – you switch, router, Intrusion Detection and Prevention System (IDPS) and firewall will likely be able to send you that data directly. Transaction data for layer 7, however, is a different story. To gather it, one option is to use a tool that can produce it directly like the IPDS or a next-gen firewall. The other option is to combine it with full packet capture and extract the data from that (as shown in the intermediate "copy of traffic" box). Once you have a copy of traffic, which is typically gathered through a network tap or switch mirror port, you can save it in full, convert it to transaction data, or convert it down to network flows.

What Can You See On your Endpoints?

Consider your endpoints...

- What ports are listening and why?
- What exploits are you vulnerable to?
- Has anyone installed unauthorized programs?
- Have any critical system files been changed?
- Have any malicious scripts been run?
- Do any systems have unique startup items?

How do we get these answers?

What Can You See On Your Endpoints?

Now that we've discussed some of the concepts of network security monitoring, let's turn our attention to endpoints. Organizations commonly have a decent view of network traffic, but endpoints may be a very different situation. Numerous solutions in the Endpoint Detection and Response (EDR) market have come out to aid in this capability, but they are not widely deployed compared to network traffic monitoring. If you can answer the questions above from the point of view of your hosts, consider yourself lucky and ahead of the curve in the monitoring realm.

Although network monitoring is a crucial part of our strategy, it is the endpoints themselves that truly hold the data that most attackers are after, and therefore, endpoint visibility is as crucial as network visibility, if not more so. Knowing which programs are installed, which services are running, if any important configuration items have changed, scripts have been run, or auto-start items have been added are all important pieces for keeping tabs on the hosts in your network.

Continuous Security Monitoring

What is **Continuous Security Monitoring (CSM)**?

- Looking at endpoint data – "data at rest"
- Configuration and baseline monitoring
- Vulnerability scanning
- File/registry integrity monitoring
- Running processes
- Services
- Autorun items




Continuous Security Monitoring

Continuous Security Monitoring (CSM) is the host-based side of the security equation, a vaguer sounding term than NSM. CSM is about monitoring the rest of the information – data that can be recorded from the endpoints on the network. CSM is concerned with configuration and baseline monitoring as well as analyzing and recording information about file and registry changes, processes, autoruns items and the like. It also includes locating and tracking vulnerabilities present in the environment. The "continuous" term in CSM is a callout to the fact that this is not "periodic, occasional check that things are still the same" but that ideally it is happening on an ongoing basis. This means that the moment that something unexpected happens on the endpoint, such as virus executable runs, you would be notified in near real-time. This type of monitoring is often accomplished via a strong audit and logging policy with the centralization of what is recorded to a SIEM system.

CSM Event Collection Illustrated

CSM collection sources:

- OS/application auth logs
- Sysmon
- Antivirus
- Whitelisting
- EDR
- HIDS/HIPS
- Vuln. Scanner



CSM Event Collection Illustrated

CSM is usually log file oriented and, therefore, is concerned with the reliable and complete collection of endpoint information. Endpoints typically either run a log agent that collects and forwards information to a SIEM or has dedicated software for something like EDR that collects the information in a separate system.

Example: Why CSM Data is Crucial

Investigations will require **network AND endpoint** data

Example: Malware infection occurs...

- Traffic is generated to an IP on port 443
 - Concerned? Unlikely, unless known bad IP
 - With only network data, that's all you see
- What if we know these processes created it!...

```

calc.exe (PID: 2092) 27/68
├── cmd.exe /C mkdir %WINDIR%\system32\zxddnmat\ (PID: 3540)
├── cmd.exe /C move /Y "%TEMP%\orzxxvzi.exe" %WINDIR%\system32\zxddnmat\ (PID: 3432)
├── sc.exe create zxddnmat binPath= "%WINDIR%\system32\zxddnmat\orzxxvzi.exe /d\C:\calc.exe" type= own start= auto DisplayName= "wifi support"
├── sc.exe description zxddnmat "wifi internet conection" (PID: 3016)
├── sc.exe start zxddnmat (PID: 3172)
└── netsh.exe advfirewall firewall add rule name="Host-process for services of Windows" dir=in action=allow program="%WINDIR%\system32\svchost.exe
  
```

Example: Why CSM Data is Crucial

Here's an example of why collection endpoint data is necessary. A random malware analysis was pulled from the site hybrid-analysis.com, which generated what looked like standard traffic to an IP address on port 443.¹ This type of traffic happens all the time in almost everyone's environment, and that traffic alone is very unlikely to trigger any type of alert looked at from a network-only perspective.

Consider, however, if we can see the names of the processes and the command-line parameters that were used to start the program that generated the traffic. With the CSM data, we can see that calc.exe made network traffic (odd by itself), then it started cmd.exe and wrote a randomly named file into the windows\system32 folder, created a randomly named service on the Windows machine disguised as a Wi-Fi service, and added a firewall rule. This type of activity should easily stick out and reveal the malicious nature of this process, but without the endpoint data, it's likely this would go completely unnoticed unless it was caught by antivirus.

[1] <https://www.hybrid-analysis.com/sample/8a1c844413ffa69968758facbc48126a0a9f54234fe251bd1f8cb4e39ac112c0?environmentId=100>

Monitoring Data Sources Overview

NSM Data

- Network extraction
- Routers / Switches
- Network firewalls
- IDS/IPS
- Proxy
- Web Application Firewall
- Service Logs
 - DHCP, DNS, HTTP(S), SMTP, SMB, FTP, SSH, Kerberos, etc.

CSM Data

- Authentication
- Antivirus
- HIDS/HIPS
- Process command line
- Application access logs
- Executables
- Vulnerability scanners
- DLP

Monitoring Data Sources Overview

When it comes to data being collected, it is useful to differentiate it in your head, not based on whether it is literally a text log or not, but whether that data is about network traffic, or activity occurring on an endpoint.

Network data can be sourced from many different devices, some that will record the actual traffic, and others that will turn it into logs describing the traffic that will be collected with a SIEM. Regardless of the way the information is recorded, it is still ultimately telling the analyst its observations about what was seen "on the wire."

Common sources for network data include:

- Network extraction tools like Zeek
- NetFlow from routers and switches
- Firewall logs that may contain block/allow actions or, in the case of "next-gen firewalls", actual Layer 7 application info
- IDS/IPS logs with a rule name that matched and a potential recording of the traffic itself
- Proxy logs that can tell which user was going to which site, including Layer 7 attributes such as URLs and HTTP methods
- Service logs from servers running things like Apache, Windows/BIND DNS, or DHCP

Endpoint data information provides the other piece of the puzzle when combined with network logs. These logs include items such as:

- Authentication data showing successful and failed login attempts
- Antivirus logs showing files identified as malicious
- Host IDS/IPS logs that may show which files were modified, or if any suspicious processes were created


- Process creation logs to describe how each new process was created
- Application and System informational and error logs
- Digital Loss Prevention logs, which may show how users moved files or interacted with sensitive data

Consider that some of these sources may contain the same data, but one may give more detail than another. For example, network firewall logs may show that a host was trying to communicate on port 4444, but the host firewall logs may show which actual process created that data, and the command line used to invoke it.

Licensed To: David Owerbach <0mamaloney0@gmail.com> April 28, 2020

Without Centralized Searching


Investigation without centralized logging



Without Centralized Searching

Log and data centralization is another important topic for a functional SOC. If you are unable to centralize transaction data, metadata, and logs, trying to investigate a potential incident may be an exercise in frustration. Consider that almost all security tools come with their own console interface. That is great when you need to investigate an alert using all the custom features of that appliance, but how do you know which tools may hold information on a given alert in the first place? The SIEM is what steps in to take care of this for us and provides a centralized repository for all the data. Without a SIEM centralizing logs, if you had to investigate activity for a given IP address, you would have to log into every application individually and try to put all the pieces of an incident together—an inefficient way to go about things.

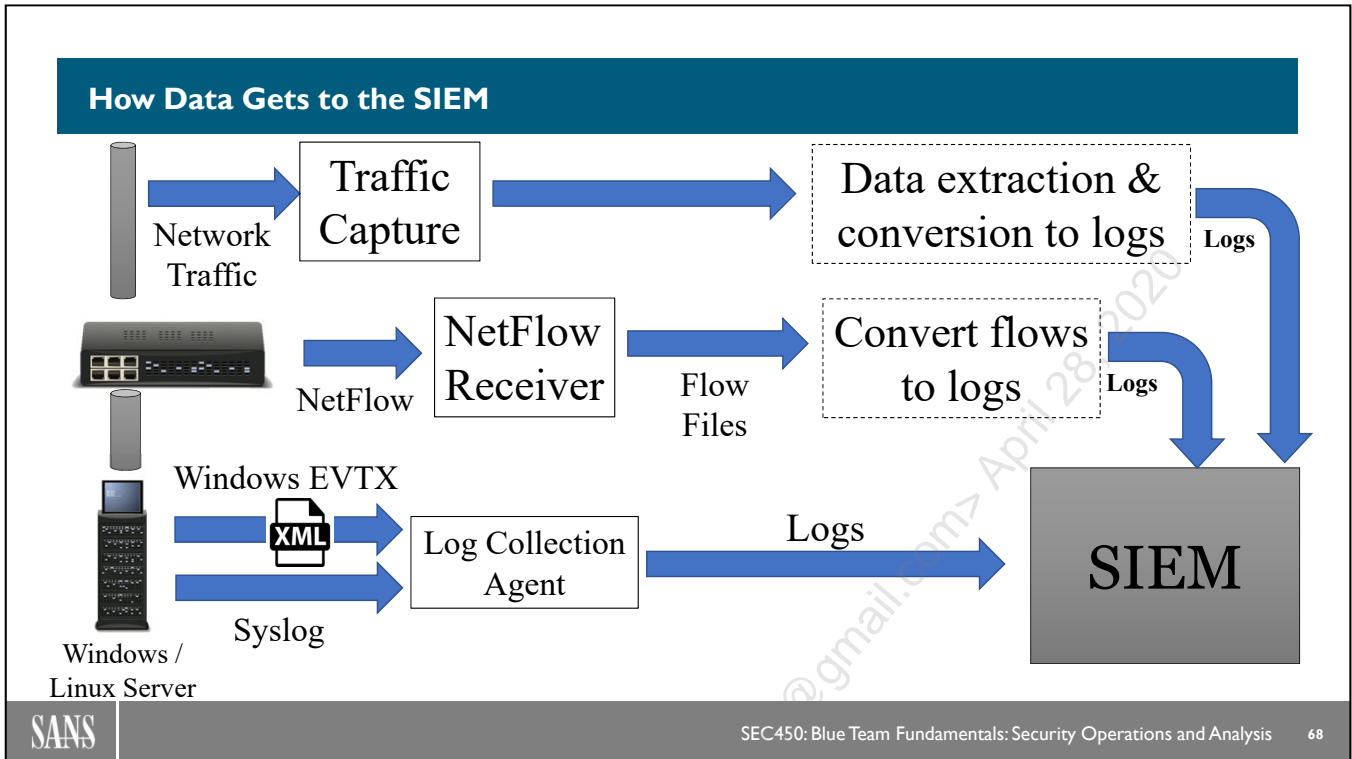
What We Want



One place to search and alert on events

What We Want

Therefore, the ideal situation is log centralization. NSM and CSM information should be centralized in a fashion that allows you to locate it all from the SIEM. In this way, if an alert fires for a malicious IP on an intrusion detection system, for example, you can then run a single search for that IP address across all data sources in the SIEM and see which other tools may have information about it. This may reveal that you have proxy logs, NetFlow, full PCAP, and a malware sandbox hit all at once. Without log centralization, you would have to individually check all those disparate data sources.



How Data Gets to the SIEM

This slide shows the several methods for getting both network and endpoint data into the SIEM. Notice that some data that is recorded is not in the format of text logs by nature and must be converted. NetFlow and PCAP data for example could not go into many SIEMs in the native form, there would need to be an intermediate conversion process that turned them into text that could be parsed and saved by the SIEM. Log files from endpoints on the other hand *are* text-based, but the specific form they are saved in may differ. Windows for example uses the EVTX format – though these logs are ultimately text in XML form, the file itself cannot be directly read, Windows write it in a special binary format that must be read by a program that knows how to interpret it before the text can be extracted. This is what happens when you use Windows Event Viewer and also what the log agent you likely use to get the logs to your SIEM has to do as well.

Again, the point of this to underscore the idea that whatever output your tools are natively generating, ultimately any data that is generated should be converted into a text-based that can be sent to and searched centrally in the SIEM.

Defensible Network Concepts Summary

A defensible network requires:

- **Network Security Monitoring**
 - Traffic to/from the internet AND within the internal network
 - Layer 7 transaction data at *least* for critical traffic
- **Continuous Security Monitoring**
 - Critical log collection from desktops, servers, and appliances
 - Configuration and baseline monitoring
 - Vulnerability information

Defensible Network Concepts Summary

The goal of this section is to introduce concepts that will be referenced throughout the rest of the class. In order to effectively defend a modern network, the events on the network and endpoints must be visible. This can be thought of in two camps: **Network security monitoring** for the "**data in motion**" on the network as well as **continuous security monitoring** for the "**data at rest**" on the endpoints. Architecturally, we are interested in monitoring so that we may ultimately understand what types of traffic are on the network and implement preventive measures in strategic locations that will detect and block as much malicious traffic as possible, making our jobs easier. Of course, not all malicious traffic can be blocked, so the *internal* NSM sensors and CSM log collection will form an important second step for seeking out what has bypassed the perimeter.

Data Centralization Summary

Centralization of data is crucial

- Not every tool natively creates logs or centralizes them
- If we don't centralize all data, must use multiple systems
- Using multiple tools is painful, fault-prone, and inefficient
- Log agents help us pick up files, convert binary formats
- SIEMs take logs as input, get non-text data converted!

Data Centralization Summary

This data, regardless of the format it is initially recorded, data should ultimately be centralized into a SIEM where we can apply searches, visualizations, correlations, and alerting rules with as much fidelity as possible. As an analyst, reading and interpreting this NSM and CSM data this will be your primary task.

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. **Events, Alerts, Anomalies, and Incidents**
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Licensed To: David Owerbach <0mamaloney@gmail.com> April 28, 2020

Understanding Your Tools

In this module:

- Defining events, alerts, and incidents
- Flow of event logs vs. alert logs
- Considerations for event and alert collection
- Alerts vs. anomalies
- Alert outcomes and tuning
- The job of an analyst

Understanding Your Tools

In this module, we will begin to explore the definition of events vs. alerts vs. incidents, and how each is part of our workflow as an analyst. We discuss event collection, the multiple options for alerting, which direction and processes alerts must go through, as well as how incidents are ultimately created. We will also explore the more subtle difference between anomaly and "signature of evil" based alerts.

Events, Alerts, and Incidents

NIST SP800-61 on definitions:



Events: Any observable occurrence in a system or network



Alerts: An event of interest that may be unwanted or unauthorized



Incidents: A violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices

Events, Alerts, and Incidents

First, we must discuss the difference between events, alerts, and incidents as these will be important concepts to understand the flow of data to an incident management system. In information security, the general task is to record events that occur on systems and on the network, flag suspicious ones as alerts, and if confirmed to be malicious, work them as incidents. The definitions we will be using in the class for events and incidents align with NIST SP800-61.¹

- **Events:** Any observable occurrence in a system or network. This could be a user logging in, or someone opening a website.
- **Alerts:** An event that *may* be unwanted or unauthorized. An example would be an IDS alert claiming it saw malware command and control traffic. Note that this is not addressed within NIST 800-61 but will be useful for the discussion of how security tools work together. Alerts are just *events that are found to be interesting* by some definition, such that you would want to notify the SOC that an event occurred.
- **Incidents:** A violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices. These typically are events that will affect the Confidentiality, Integrity, or Availability of business data. In other words, a confirmed incident. *Events* that become *alerts* that are verified as true positives become *incidents*.

Compared to NIST, other sources such as ITIL have different definitions for these terms. We will use the definitions here as guidelines. In short, the volume of items that meet these definitions is like a filter. All observable occurrences are events, and some of those events that are interesting will be alerted upon. Alerts are then collected (often in a SIEM) and triaged by an analyst that can confirm them as *incidents*. Once an alert has been confirmed to be an incident, it will be worked in an *incident management system*.

[1] <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf>

Event Collection

Most logs are records of **events** as transaction data...

```
Network connection detected:
UtcTime: 2018-09-23 16:16:16.054
ProcessId: 8012
Image: C:\Users\user1\AppData\Local\slack\app-3.3.1\slack.exe
User: win10\user1
Protocol: tcp
SourceIsIpv6: false
SourceIp: 10.150.159.161
SourceHostname: win10
SourcePort: 60201
SourcePortName:
DestinationIp: 52.85.81.133
DestinationHostname:
DestinationPort: 443
DestinationPortName: https
```

```
1335543155 315 127.0.0.1 TCP_MISS/301 -1 GET
http://www.imdb.com/title/tt0056869 - NONE/- text/plain
```

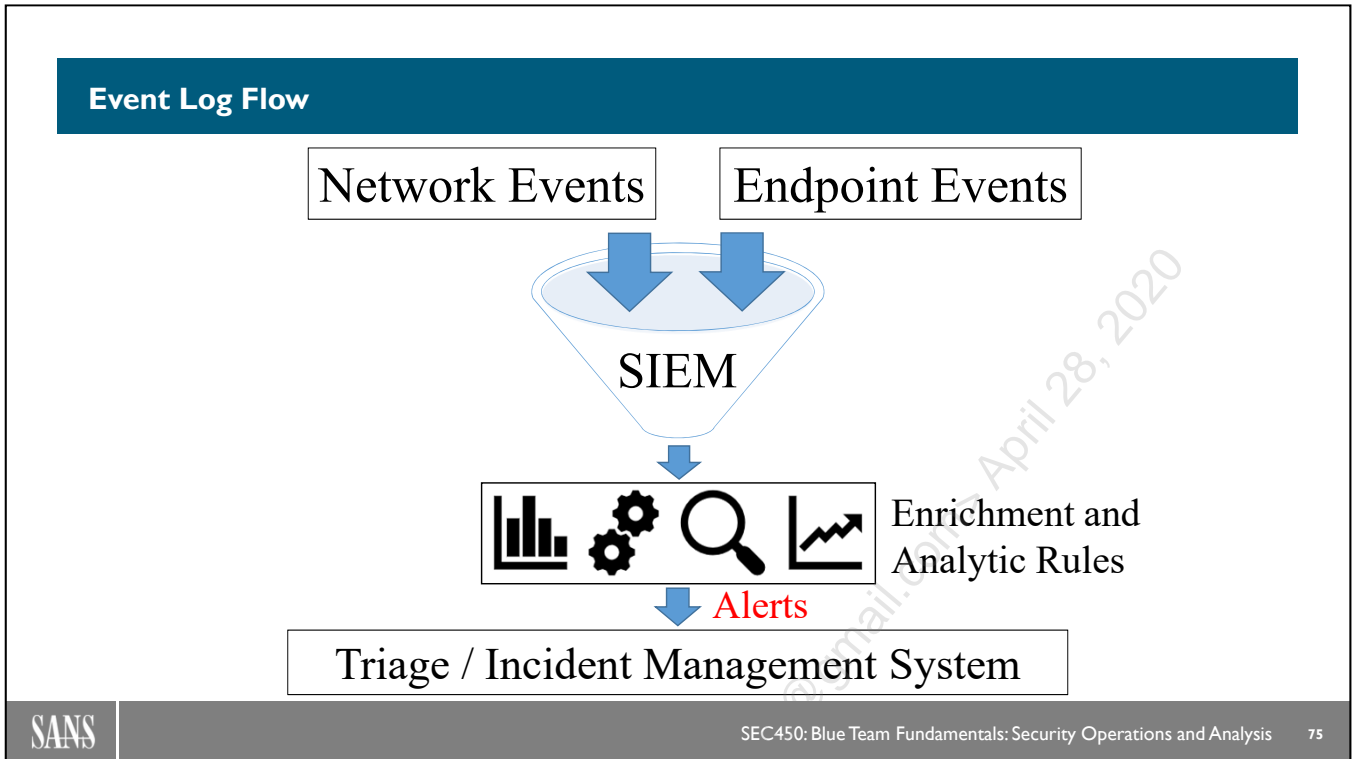
```
117.201.11.139 - - [02/Jan/2017:02:35:54 -0800] "GET /
HTTP/1.1" 200 34374 "-" "Mozilla/5.0 (Windows NT
6.1; WOW64; rv:40.0) Gecko/20100101 Firefox/40.1"
```

```
Feb 27 15:12:28 srv named[8978]: client 1.1.1.2#38595:
query: www.allegro.pl IN A +E
```

```
Sep 23 08:54:56 ubuntu CRON[16355]:
pam_unix(cron:session): session opened for user root by
(uid=0)
```

Event Collection

Most logs that are collected will fall under the definition of events—any observable occurrence. They are simply a log telling us that something happened with no clear evidence that it was bad. It could be that someone visited a website, someone logging in, etc. In this slide, there are log examples from Squid Proxy, an Apache web server access log, a BIND DNS request, a Linux authentication, and a Windows Sysmon network connection. In most environments, most data collected will be these events—a log that someone did something without any inherently interesting content. The log becomes interesting once we centralize it and process it in some way to see if it contains any suspicious information. The tool we often do this with is the SIEM. The SIEM allows us to parse *events* and apply threat intelligence, correlate them with other logs, or match it against a blacklist of malicious IPs, domains or otherwise, and potentially upgrade them to *alerts*.



Event Log Flow

Here is the typical flow of event logs in an event monitoring system. Events occur on endpoints or are recorded by our network appliances and, as discussed in the last section, turned into logs. These logs are centralized and, through a variety of means, are stored in the SIEM. The SIEM specializes in mass parsing, enriching, and inspecting of log data, and it will apply the set of analytics (rules for detection conditions of interest) to surface any events that contain items we may need to react to. Once an event of interest is identified, that individual log can be elevated to an alert, and that log is typically then sent with any additional info that helped identify it as potentially evil to a triage queue where it will then be inspected by an analyst and potentially sent to an incident management system.

A Word on Event Collection

Collecting lots of events is great for detection...or is it?

- Before you "collect it all", consider:
 - SIEMs are licensed on volume and EPS
 - Searching more data takes longer
 - Finding signal is harder in the noise
- Events are not Pokémon
- **Goal: Tactical** detection, not hoarding



VS.



A Word On Event Collection

At this point, it's worth briefly discussing event collection. When it comes to collecting events in your environment for malicious activity detection, should you try to collect everything you can? Not necessarily. Event collection with default settings from many sources will include an incredible amount of noise that will only serve to slow down your SIEM, make it more expensive, and complicate finding the "needle in the haystack." Consider the events you collect and, if you can show that a subset of them are providing no value, don't be shy cutting them out! SANS Instructor and Author Justin Henderson's "SEC555: SIEM with Tactical Analytics" course is a great resource for a deep-dive on this topic for those who are interested in setting up a SIEM in the most efficient way possible. The takeaway here is that hoarding all logs is not necessarily better for malicious activity detection. We should strive to be tactical in the events that we collect; it makes running a SIEM easier for everyone involved.

Alert Collection

Some logs represent **alerts!**

```
12/12-06:22:52.879193 [**] [1:2014126:1] ET CURRENT_EVENTS DRIVEBY Blackhole Likely
Flash Exploit Request /field.swf [**] [Classification: A Network Trojan was Detected] [Priority: 1]
{TCP} 192.168.45.10:1046 -> 78.46.173.138:80
```

```
Oct 30 09:50:48 1,2012/10/30 09:50:48,01606001116,THREAT,url,1,2012/04/10
04:38:23,192.168.0.2,204.232.231.46,0.0.0.0,0.0.0.0,rule1,crusher,,web-
browsing,vsys1,trust,untrust,ethernet1/2,ethernet1/1,forwardAll,2012/04/10
04:38:24,12783,1,59024,80,0,0,0x200000,tcp,block-
url,"onlinebrandsecuritys.com/install/ws.zip", (9999),malware-sites,informational,client-to-
server,0,0x0,192.168.0.0-192.168.255.255,United States,0,
```

Potentially actionable data, must decide where these go:

- Logs only? IDS console? SIEM? Direct to ticketing?

Alert Collection

While many *logs* represent *events*, logs that are collected from *security appliances* like intrusion detection systems will often contain actual *alerts*. (We must make the distinction carefully between logs of events and logs of alerts here). On this slide, the top box shows a Snort IDS *alert* log for a Blackhole exploit kit telling us that whoever had IP 192.168.45.10 at this time has potentially become infected. The bottom alert log is from a Palo Alto firewall's threat log and is signaling that a known malicious URL, onlinebrandsecuritys.com, was contacted by IP address 192.168.0.2, which downloaded a suspicious zip file.

Both event and alert logs are sent to the SIEM, but they are different in that alert logs already have matched against a list of malicious activity. Therefore, these logs need to be handled differently in the SIEM than normal events, and there are multiple ways a SOC may choose to do this. Should we triage these items in an IDS console before sending them to the SIEM? Should we immediately send all alert logs to the SIEM for centralization? Should we send an alert log directly to a ticketing solution to spin up a response? There are multiple possible valid answers here.

Alert Triage

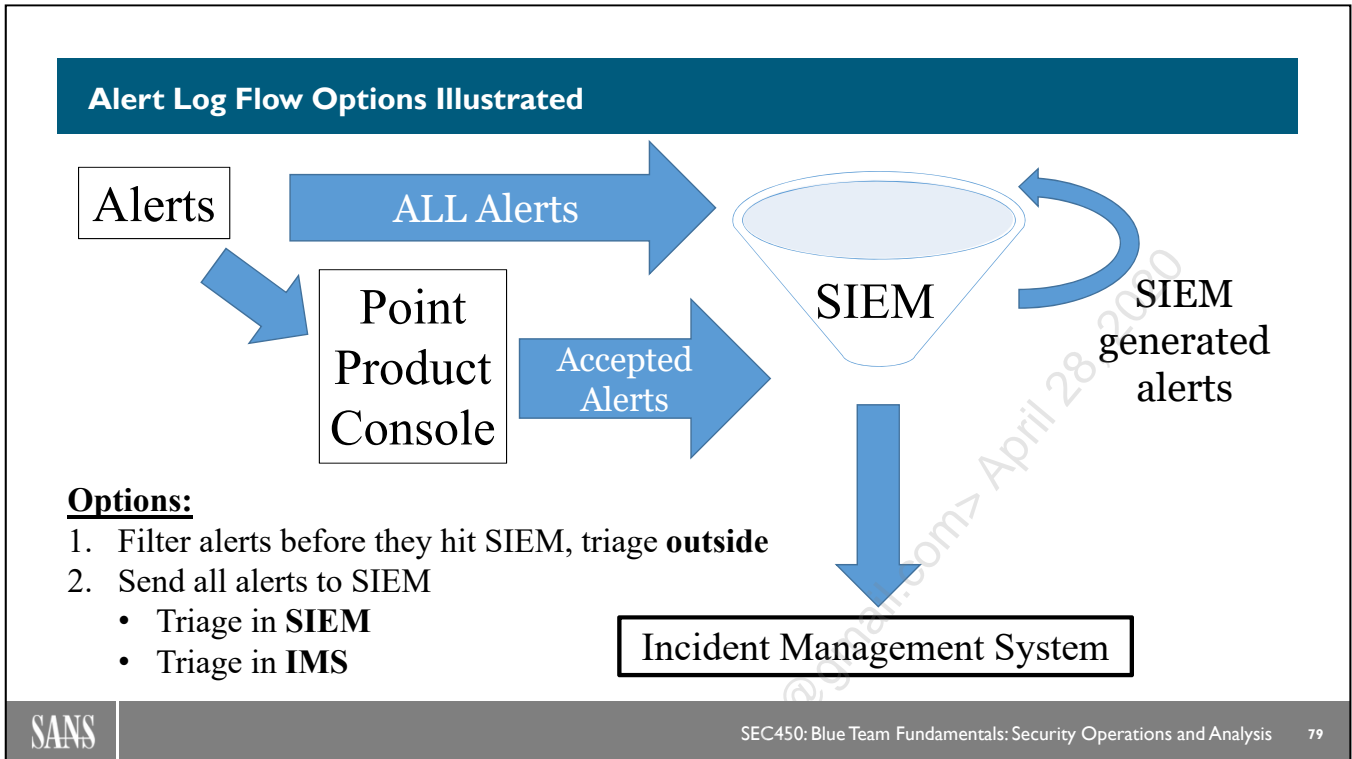
Alerts often end up in a console for triage:

ST	CNT	Sensor	Alert ID	Date/Time	Src IP	SP...	Dst IP	DP...	Pr	Event Message
RT	1	sec-51...	3.114	2017-05-02 2...	10.5.11.57	52...	94.152.8.57	80	6	ET CURRENT_EVENTS Terse alphanumeric executable downloader high
RT	1	sec-51...	3.388	2017-05-08 2...	10.5.11.57	52...	10.5.11.10	53	17	ET DNS Query to a *.pw domain - Likely Hostile
RT	6	sec-51...	3.103	2017-05-02 2...	10.5.11.173	445	10.99.99.8	50...	6	ET EXPLOIT Possible DOUBLEPULSAR Beacon Response
RT	1	sec-51...	3.102	2017-05-02 2...	10.5.11.173	445	10.99.99.8	50...	6	ET EXPLOIT Possible ETERNALBLUE MS17-010 Echo Response
RT	13	sec-51...	3.116	2017-05-02 2...	10.5.11.57	53...	23.88.92.15	80	6	ET INFO GENERIC SUSPICIOUS POST to Dotted Quad with Fake Browse
RT	2	sec-51...	3.390	2017-05-08 2...	10.5.11.57	52...	103.16.76.2...	80	6	ET INFO HTTP Request to a *.pw domain
RT	3	sec-51...	3.140	2017-05-02 2...	10.5.11.57	54...	10.5.11.10	53	17	ET POLICY DNS Update From External net
RT	2	sec-51...	3.200	2017-05-03 1...	10.5.11.52	52...	10.5.11.10	53	17	ET POLICY DNS Update From External net
RT	2	sec-51...	3.223	2017-05-03 1...	10.5.11.44	53...	10.5.11.10	53	17	ET POLICY DNS Update From External net
RT	2	sec-51...	3.229	2017-05-03 1...	10.5.11.85	61...	10.5.11.10	53	17	ET POLICY DNS Update From External net
RT	2	sec-51...	3.389	2017-05-08 2...	10.5.11.57	52...	103.16.76.2...	80	6	ET POLICY Possible HTA Application Download
RT	3	sec-51...	3.398	2018-08-03 1...	0.0.0.0	68	255.255.25...	67	17	ET POLICY Possible Kali Linux hostname in DHCP Request Packet
RT	1	sec-51...	3.406	2018-09-08 1...	10.5.100.93	32...	10.5.11.173	3306	6	ET POLICY Suspicious inbound to MySQL port 3306
RT	3	sec-51...	3.1	2017-04-25 1...	45.76.92.117	123	10.5.11.11	123	17	ET TOR Known Tor Relay/Router (Not Exit) Node UDP Traffic group 400
RT	2	sec-51...	3.112	2017-05-02 2...	10.5.11.57	52...	213.136.26...	80	6	ET TROJAN WS/JS Downloader Mar 07 2017 M1
RT	1	sec-51...	3.397	2018-08-01 1...	10.5.100.94		10.5.11.25		1	GPI ICMP INFO PING *NIX

Alert Triage

One option is to capture alerts like this in the appliance that made them and triage them from a console. This slide shows the IDS console for Squid, which is included in Security Onion and can be used for triaging IDS alerts. We can see that multiple different signatures have fired in the right-side column, and that in some cases, such as the "ET INFO GENERIC SUSPICIOUS POST" alert, that the CNT (count) column reveals that this alert has fired multiple times for the same host.

Triaging alerts in a separate interface before they reach the ticketing system is one way of keeping downstream work related to false positives minimized. If an alert can be dismissed in this console before the SIEM sees it or a ticket is generated, that is potentially less wasted time for the team. The downside of this method is you do not get the enrichment and correlation features of the SIEM to help you decide which alerts are false positives, which can be a major help. Centralizing all alerts allows crucial fidelity improvement and means analysts only must go to one location to see all alerts. This class will visit these concepts in more detail later.



Alert Log Flow Options Illustrated

This slide demonstrates the multiple routes that alert logs can take through your SOC. Keep in mind that more than one of these options may be utilized simultaneously, depending on the fidelity of the rule, the tools involved, and the SOC's designated workflow. Each option is explained on the next page.

Licensed To: David Owerbach <dmaloney0@gmail.com> April 28, 2020

Alert Log Flow Options

Consider your options, combine if needed:


1. Send alerts to the preliminary queue for initial processing
 - For low-fidelity alerts/anomalies, false pos. reduction early in the process
 - **Two queues** to deal with, SIEM enrichment not available at first queue
2. Send all alerts to SIEM (remember, the SIEM makes alerts, too)
 - Gives you the ability to add **enrichment** to **all** alerts
 - You still must decide what tool to use for **triage**:
 1. SIEM built-in alert triage system
 2. Incident management system with alert dashboard

Alert Log Flow Options

One option is to have more than one queue for alerts. Alerts generated by some appliances may go to an IDS console for triage before being passed to a SIEM or to the incident management system. This method may be best used for alerts that are experimental or represent potential anomalies in the environment instead of matching a list of known bad indicators. While it is nice to keep high volume out of your main alert triage system, you will not have the enrichment benefits of the SIEM to help you decide good from bad. Doing this also creates two places analysts must attend to, but also can bring down the burden of dismissing false positives for low-fidelity rules in the "main" point of triage.

The other option is to send all alerts directly to the SIEM. In many cases, this will be the better option because the SIEM can bring in additional context and perform data enrichment on the alerts, making it easier to decide true from false positives. Assuming this is the route you take, you still have options on where to perform the triage. One option is to do alert triage inside the SIEM—some SIEMs provide a built-in tool and area for doing this. The other option is to send all alerts in their enriched and improved form to an incident management system that has its own alert collection and aggregation capability (such as TheHive). This distinction is less important and which option you pick is best informed by the specific tools you have and which you like to use the most.

The Problem With Alerting: Volume



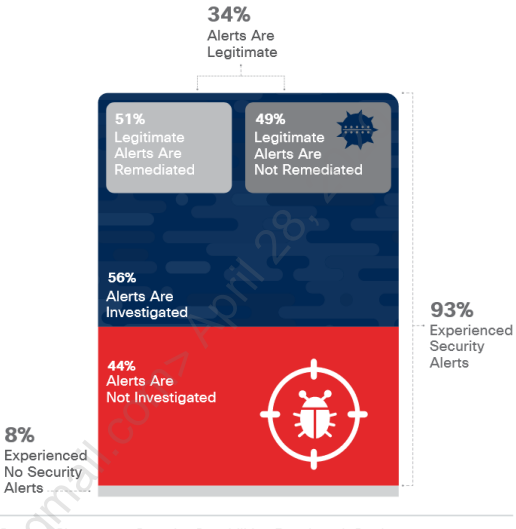
Businesses ignoring half of their security alerts, warns Cisco

Target Ignored Data Breach Alarms

Target's security team reviewed -- and ignored -- urgent warnings from threat-detection tool about unknown malware spotted on the network.

Target confirmed Friday that the hack attack against the retailer's point-of-sale (POS) systems that began in late November triggered alarms, which its information security team evaluated and chose to ignore.

Neiman Marcus Hackers Set Off 60,000 Alerts While Bagging Credit Card Data



34% Alerts Are Legitimate

51% Legitimate Alerts Are Remediated

49% Legitimate Alerts Are Not Remediated


56% Alerts Are Investigated

44% Alerts Are Not Investigated

8% Experienced No Security Alerts

93% Experienced Security Alerts

Source: Cisco 2018 Security Capabilities Benchmark Study



SEC450: Blue Team Fundamentals: Security Operations and Analysis

81

The Problem With Alerting: Volume

If there's one thing most SOCs struggle with, it's the volume of alerts that are generated. It is extremely common to find organizations buried in a massive amount of alert info and unable to figure out where to start on the pile. The 2018 Cisco Cybersecurity Report shows that a shocking 44% of daily alerts are not even investigated in the average organization. Of the ones that are, only 34% are deemed legitimate. Of the legitimate alerts, only 51% are acted upon by some type of remediation.¹ From this, we can gather that:


1. Most companies are overwhelmed with alerts.
2. Many of those alerts are false positives.
3. We don't seem to have enough resources to even deal with the true positives. This is a poor situation for the Blue Team to put themselves in.

We also know that alert fatigue may have played a part in multiple large-scale breaches throughout the years. Both Target and Neiman Marcus made headlines due to the apparent fact that they had missed the alerts that could have keyed them off to their respective breaches. How many of us could honestly say that, if put in the same situation, we wouldn't be in the same position? If Blue Teams can't separate the false from the true positives, and even the true positives have poor naming conventions, it would be next to impossible to triage these situations in the order of true importance. This situation leads to delayed response times and incidents that escalate to the point of breaches before they can be addressed.

[1] 2018 Cisco Cybersecurity Report: <https://www.cisco.com/c/en/us/products/security/security-reports.html>

Two Flavors of Alerts: Signatures and Anomalies

- **Signatures:** Blacklist-based alerts directly indicate badness
 - A malicious protocol, domain, IP, file hash
- **Anomalies: "An odd, peculiar, or strange condition..."** - Dictionary.com
 - User logs in on new computer
 - User logs in from new location
- Do anomalies == malicious?
 - Not necessarily, but the reverse is true!
 - Correlation means it is a valid way to find evil!

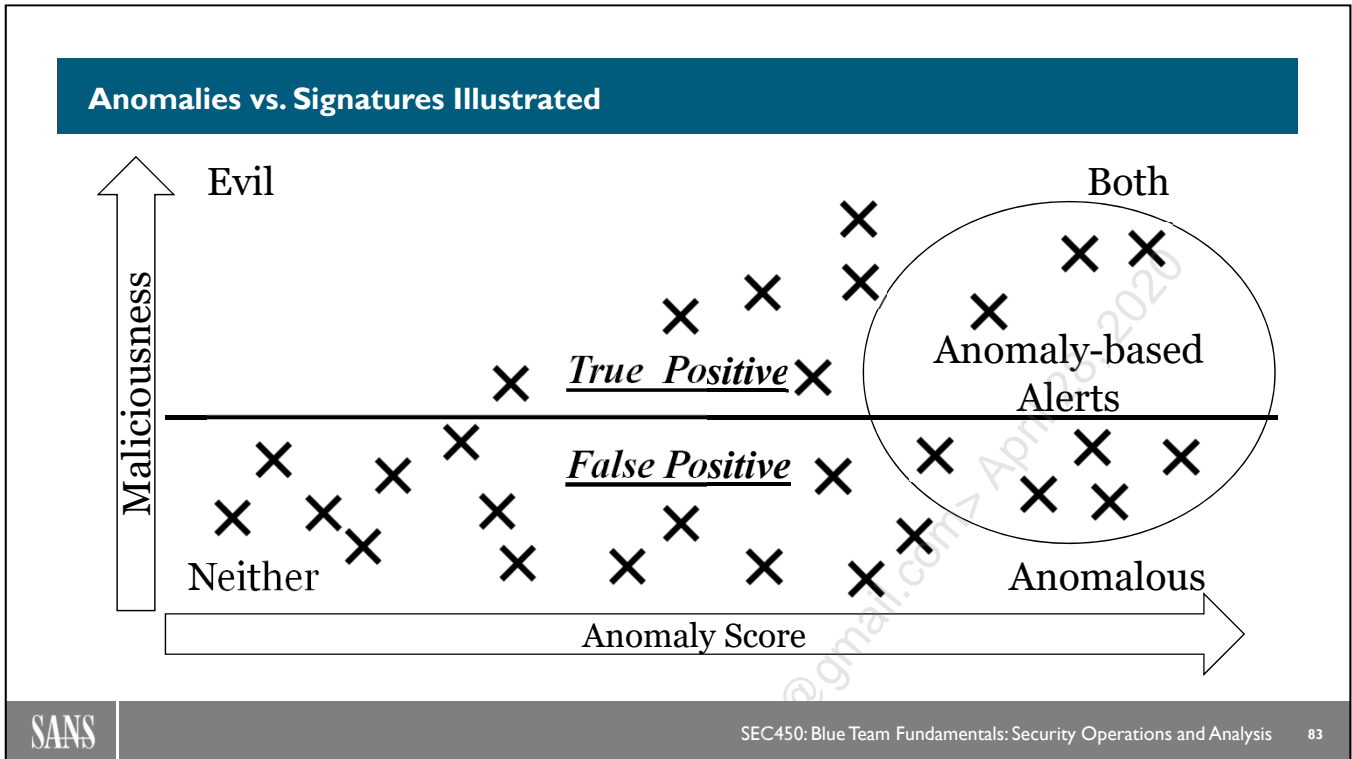


Two Flavors of Alerts: Signatures and Anomalies

There are multiple types of rules that can be written for an IDS. One is signature or blacklist-based, which will alert on any observed indicators such as URLs or IPs that are known bad, or even URL patterns for things like exploit kits. For these alerts, assuming the signature or indicator behind the rule is solid, then these rules tend to reliably point out evil on the network and should be considered a high-fidelity detection.

The other type of alert, however, is the *anomaly-based* alert. Although this is only a proposed distinction, it is nonetheless a helpful model to help separate alerts into these two types in your head. Anomaly-based alerts are not based on any list of known bad indicators of compromise or patterns, but rather on things that would be defined as unusual within the environment.

What is considered odd? That is up to the organization wielding the rule, but many times it is things like moving executables from one unexpected source to another, large file uploads, a user logging in from a new computer that they've never used before, or other similarly interesting, but not necessarily predetermined known-malicious conditions. Depending on the type of anomaly rules you have written, as well as how locked down and defined your environment is, you will likely find that these types of alerts are more likely to produce false positives than "signature of known evil" type alerts.



Anomalies vs. Signatures Illustrated

The reason that anomaly rules are more likely to produce false positives is the assumption behind these types of rules. Specifically, that is that anomalous events are more likely to be malicious than non-anomalous events. Think of these relationships like squares and rectangles. While not all rectangles are squares, all squares are rectangles, and the same is true of anomalies and malicious activity. While it is very possible to have a non-malicious anomaly, all malicious activity should be anomalous in *some* fashion (which is why the diagram shows nothing in the upper left area). The question is whether your team can identify an event as an anomaly or not, doing so might take a data source you don't currently collect. Because of this relationship, in the mathematical sense, we can say anomalous and evil events are correlated. Therefore, it is reasonable to attempt to detect evil in this way.

The best way to improve the fidelity of anomaly-based alerts is to enrich their data with additional context or correlate the activity with other things happening from the same host or user. Doing this may bring to light additional information that can push anomaly-based alerts into the high-fidelity detection territory. Details on doing this will be covered in more detail in Day 5.

Using Signature and Anomaly Alerts

Alerts that identify evil

- Signatures/lists of known bad
 - Known evil domain name
 - C2 traffic pattern
 - Exploit kit URL pattern
- Ruleset requires tuning
 - Poorly sourced list of evil causes false positives
 - Overactive alerts cause fatigue
- Alerts should be investigated
- More likely to find true evil

Alerts that identify anomalies

- Show abnormalities/deviations from normal activity
 - Out-of-hours login
 - Uploading large file
 - Login on new PC
- More difficult in less-defined environment, must know normal
- Triage before ticket generation?
- Enrich to produce higher fidelity

Using Signature and Anomaly Alerts

Keeping in mind the difference between the two types of alerts can help you decide how to design the workflow associated with the generation of each type of alert, and guide how they are triaged. For example, you may choose to separate anomaly-based alerts that are low fidelity to a preliminary queue as discussed earlier. This could prevent jamming up analysts with lower-fidelity alerts.

Highly enriched anomaly-based alerts, or ones that can be correlated with other events, may fall in the "high-fidelity" range. If separating workflows by alert fidelity, these items can be sent down the normal alert path. Further information on this topic is available at the following Microsoft link:

<https://blogs.msdn.microsoft.com/azuresecurity/2016/03/09/evolution-of-useful-results-from-anomaly-detection-systems/>

Discussion: IDS – Signature/List or Anomaly-Based Alert?

Consider the following Snort rule names:

1. ET CURRENT_EVENTS Blackhole 16-hex/a.php Jar Download
2. ET POLICY ZIPPED EXE in transit
3. ET POLICY Wget User Agent
4. ET CURRENT_EVENTS Fake Adobe Flash Player malware binary requested
5. ET POLICY Unusual number of DNS No Such Name Responses
6. ET POLICY TLS/SSL Server Certificate Exchange on Unusual Port
7. ET ATTACK_RESPONSE Possible /etc/passwd via HTTP (linux style)
8. ET CHAT IRC JOIN command
9. ET POLICY PDF File Containing Javascript

Discussion: IDS – Signature/List or Anomaly-Based Alert?

Here are some example rules that may be used by your intrusion detection system. Looking purely at the list of names, which one would you be inclined to respond to first? Which one of these seems to be a signature-based alert vs. an anomaly identification? Generally speaking, it's the rules that have POLICY in their name that would fall under the anomaly category, while most other things in the list indicate that something bad is occurring. The blackhole exploit kit, fake Adobe malware binary, and /etc/passwd warning are all the most concerning, as would be the IRC alert if IRC is not allowed on the network.

Policy rules, in this case, mention files being downloaded, command line tool user-agents, DNS request failures, non-standard SSL ports, and PDF files with JavaScript. None of these items necessarily imply malicious activity, but all of them are hallmarks of certain types of attacks. Whether it is truly evil or not could likely easily be determined by enriching the alert data with information about the domain names involved, how long those domains have existed, or if they themselves have a poor reputation. Like most anomaly alerts, they wouldn't be labeled high fidelity on their own, but could easily be put there with a little additional context.

Signature vs. Anomalies by Tool

Signature-centric: IDS/IPS, Firewall, AV, WAF

- Primarily use **signatures** of known bad traffic and files
- If high-fidelity, might send straight to incident ticket

Anomaly-centric: User Behavior Analysis, ILP, some IDS

- UBA uses data science / heuristics to find outliers
- IDS signatures can be anomaly centric (VNC seen, EXE download)
- Anomaly + domain expertise and context = malicious activity alert

Signature vs. Anomaly Alerts by Tool

Some tools are more likely to produce one type of alert or the other. IDS/IPS, firewalls, antivirus, and web application firewalls are all examples of tools that are generally geared toward the signature and blacklist end of the alerting spectrum, although this is not always 100% as we saw in the last slide. The most common anomaly-centric tools used currently are User Behavior Analysis (UBA), Information Loss Prevention (ILP), and *some* IDS and other security appliance signatures.

UBA is often anomaly-centric in that, in many cases, it uses heuristics and data science to find outliers in data acting differently than the rest of the population. It will also track each entity a user has interacted with before and alert when new items are encountered. It is certainly not the case that the first time a user interacts with another user or device we should assume it's evil, but sometimes it is; therefore, these are anomaly-based alerts and the context of the interaction will have to be used to guide the analyst when triaging such events.

Information loss prevention systems can function much in the same way. A user is moving a concerning number of files at once or performing a large volume transfer—there are many legitimate reasons such an event might occur, but it also may be representative of malicious behavior, so those instances should be reviewed by an analyst with additional context and correlation of surrounding events.

When an Alert Fires

An analyst's job - verifying alerts, taking appropriate action

- Investigate the event / alert
- Consider rule **type**
 - Event identified as bad?
 - Alert log? High-fidelity? Investigative? Anomaly?
- Consider **context** of activity
- **Enrich** with further information
- **Correlate** with other logs and activity
- Make a determination. If evil, *incident* is created...

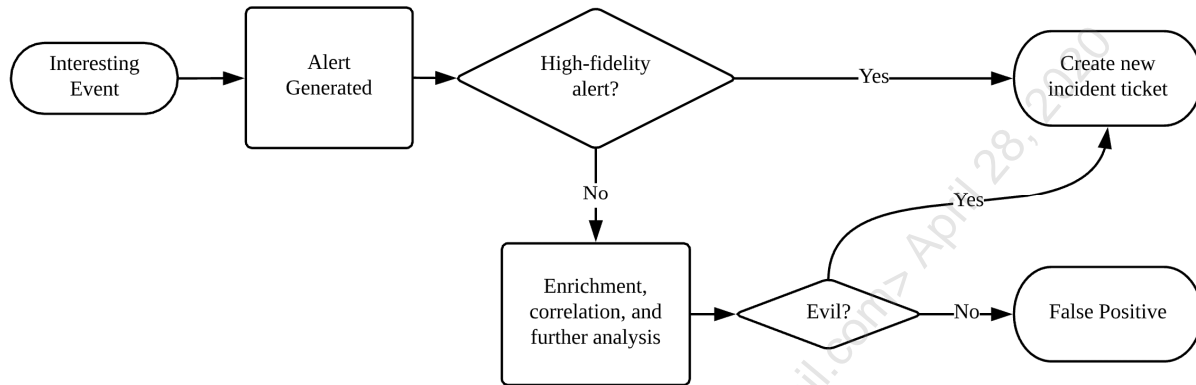
When An Alert Fires

As an analyst, one of your main jobs is to stay on top of any alerts that fire and quickly triage to understand if they are credible or not. Ideally, this would begin with the SIEM or automation tools grabbing all the extra enrichment information for you from the start, but some information you will inevitably need to find yourself. We will discuss the enrichment data types that are available to improve analysis through the course; but, at this point, we will just focus on the information at hand. Day 5 dives heavily into alert design and tuning. For now, we are just throwing out some of the high-level questions.

- Consider the rule type. Is it something that directly indicates evil, such as a known malicious domain being contacted, or is it pointing out an anomaly that needs manual analysis for consideration? In most cases, having analyst eyes set on the alert and its details will be the next step required to determine if it is a false positive or not.
- Consider the context of the alert: If an alert for administrative tool use fires, was it set off by someone in IT who might feasibly use them, or someone in HR who would have no business running the tool?
- Correlate the information with additional logs from that same source as well, pull up all alerts related to that host or user over time and see if they tell a story of compromise, or if there is nothing else to see. Are there any traffic logs available to refute or back up the claims made by the alert?

All these activities fall under the responsibility of the analyst. It is their job to take this information and make a determination if additional action is needed, and if an incident case should be created or not.

High-Level Workflow



High-Level Workflow

Here's the high-level workflow of the items that were discussed in this section. The main idea is that some events will generate alerts. Those alerts may be high-fidelity and, in some cases, may be high-fidelity enough to immediately cut into incident tickets. If they aren't, however, such as in the case of anomaly-based rules, additional context and correlation should be used to verify the legitimacy of the alert. This will prevent false positives from proceeding to the point where it will be more work to clear them out. In addition, it is work tracking ALL alerts that fire in the SIEM, so overactive rules can be identified and tuned, which is another activity that should be done on a regular basis.

For additional information on alerts and anomalies, see the following webcast from Justin Henderson and Tim Garcia:

<http://www.sans.org/webcasts/high-fidelity-alerts-alert-anomaly-sibling-rivalry-107875>

Events Alerts and Incidents Summary

Goal is to understand:

- Events and alerts vs. incidents
- Information and log flow
- Alert types
 - Signature / blacklist-based
 - Anomaly-based
 - How to approach each
- Enrichment and correlation help eliminate needless work
- If evil, make a case and investigate or escalate!



Events Alerts and Incidents Summary

This section covered the differences between events, alerts, and incidents, and discussed the differences between anomaly and signature-based incidents. With these definitions and an understanding of the flow of events and alerts, we are now able to jump into the tools that will manage this information and the incidents created from their output.

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. **Incident Management Systems**
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

SOC Data Organization

SOC solutions required:

- Track all alerts and potential incidents
- Collect and organize threat intelligence
- Collect and search network and endpoint logs
- Log correlation, enrichment and alerting
- Automation of investigation actions
- Making a repository of security team documents and info

SOC Data Organization

Running a SOC takes immense data storage and organization capabilities. Throughout the next several modules, we will investigate the core pieces of SOC technology used in day-to-day operations—the incident management system, threat intelligence platform, security information and event management (SIEM), orchestration and automation tools, and general knowledge databases. Each one plays a crucial part of security operations, so understanding their uses, inputs, and outputs lays an important foundation we will use throughout the rest of the class.

Tools for SOC Data Organization and Search

Incident Management System (IMS)

- Tracking alerts, incident status and associated indicators
- Otherwise known as *Security Incident Response Platforms* (SIRP)

Threat Intelligence Platform (TIP)

- Collection of indicators and higher-level intelligence info

Security Information and Event Management (SIEM)

- Log collection, indexing, search, correlation and alerting

Security Orchestration, Automation and Response (SOAR)

- Automation of common tasks, orchestration of workflow

Knowledge Database / Source Code Repositories

- For all SOC documents/code, playbooks, and use cases

Tools for SOC Data Organization and Search

With respect to SOC tools and data, the tools that are doing the detection are not the end of the story. There is a host of further applications that are typically run to organize and act upon the information that is collected from security appliances. One is the **Incident Management System**, the system where all recorded incidents will be investigated and worked. These often resemble ticketing systems such as those used by the help desk but should contain tweaks and extra features to help customize the experience toward a security team use case.

Another tool is the **threat intelligence platform**. As incidents are worked and information about adversaries is independently collected, there must be a structured place to put it where the data can be saved and operationalized. These systems should hold both the low-level tactical information such as Indicators of Compromise, as well as higher level strategic intelligence on what adversaries are doing, how they are doing it, and why.

The **SIEM** is a tool that doesn't do any direct data production itself, but instead aggregates all the data from elsewhere. It takes the network sensor's information and endpoint logs and makes them all available for search and visualizing. It is the single best source of information in many environments due to its wide-reaching view. Because of this fact, it also frequently implements additional alerting rules (beyond the ones in your IDS for example) that may fire based on the content of the logs it collects.

SOAR is a relative newcomer on the scene and is the name for tools to help Blue Teams automate and orchestrate response or investigation actions based on the inputs from other tools. For example, a SOAR platform could be used to automate the investigation of any machine that had a positive antivirus detection fire. The tool could be used to watch for AV logs, identify the host in question, then run a series of data

collection and test scripts to make sure the system is in the expected state and is not compromised. You can probably see why these tools have caught on fast, they free up analysts to do the hard analysis work as opposed to the purely mechanical and automatable data collection.

Knowledge databases and source code repositories help the SOC keep track of all data that doesn't fit into one of the other categories. This can be general documentation for new team members, case studies, scripts, custom programs, or other information that is of general interest to the SOC.

Licensed To: David Owerbach <0mamaloney0@gmail.com> April 28, 2020

Incident Management Systems (IMS)

Options:

1. Traditional ticketing solutions
2. SIEM built-in solutions
3. Security-tailored ticketing

Traditional Ticketing:

- BMC Remedy
- RT
- Redmine

Commercial, security-oriented:

- Resilient
- Archer
- ServiceNow
- CyberCPR

Open source, security-oriented:

- TheHive
- RTIR – Req. Tracker for IR
- FIR – Fast Incident Response

Incident Management Systems (IMS)

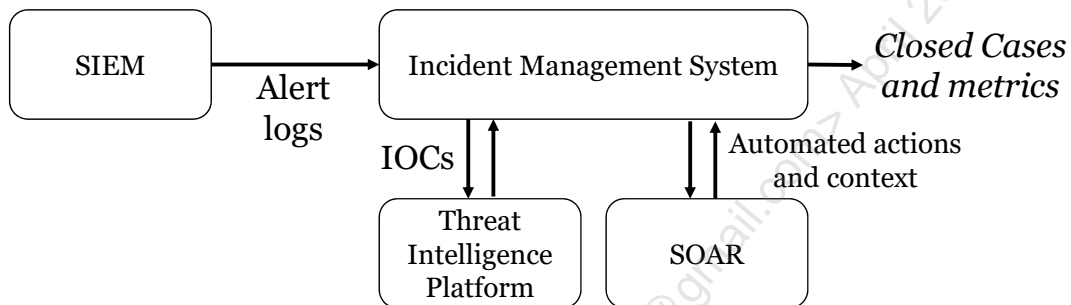
Incident management systems are where Blue Teams record the investigation and actions taken when an incident has occurred. Generally, teams will have acquired a solution from one of three places—a vendor that makes general ticketing software such as what a help desk would use, one that is built into one of their tools like a SIEM, or a dedicated security-focused ticketing system. Dedicated security-focused systems are usually the best as they are designed to work to the goals of a security team. They often will have all the typical features of a ticketing system but will include extra features important to security practitioners such as the ability to record IOCs associated with an investigation, and integrations with threat intelligence platforms.

If you ask a room full of information security professionals what they think of their incident management system, responses are often not extremely enthusiastic. Be aware that many ticketing systems exist, and none are perfect, but some are certainly better at specific tasks and have a much nicer interface to use than others. In the author's opinion, the interface and general experience for the ticketing system should be one of the most highly-considered choices you make. Analysts will be using it all day every day and a poor choice in this area can lead to years of pain and frustration. Do not think that commercial options will necessarily be better than free open source systems either. There are options from both the free and commercial columns that have very painful interfaces or other significant limitations that will negatively impact the analysts. Therefore it is an absolute necessity to give the ticketing system a very thorough test run before selecting it, many issues will not become apparent until you give it a true test drive.

In the author's opinion, one of the strongest free open source incident management systems is TheHive. It seems to do a good job of striking the right balance with detail and ease of use and has great features for data enrichment. This class will use TheHive in the lab VM as a representative incident management system.

Incident Management Systems – Systems View

- IMS receives alerts, creates case in ticket queue
- Cases assigned to analysts by cases or specific task
- Analyst works case to completion or escalates
- Associate observables with cases across time



Incident Management Systems – Systems View

Looking at the IMS as an abstracted system, the input will be alert logs either directly from a security appliance or from the SIEM. These logs will contain multiple fields of useful information pertaining to the potential issue and these fields will make up the initial detail put into a case, ticket, or whatever you have named an individual incident. When processing these cases, incident management systems may be integrated with multiple other tools like the threat intelligence platform or a security orchestration, automation, and response platform. The connection allows the data related to the ticket such as malicious domain names to be stored into the TIP, as well and are used by the SOAR platform for enrichment and additional context gathering.

Once the log has entered the IMS and context has been added either manually or through a SOAR platform, analysts will use the system to work through the case, performing analysis and deciding on if incident response is necessary. Once the incident has been dealt with and detailed notes on the actions taken are complete, the output of the system is the library of closed cases and their related categorizations and metrics.

Incident Management System Features

- Incident management systems vary greatly
- **Test them carefully** before choosing one
- Commercial is not always better than open-source
- This will be one of your main tools; **you MUST enjoy using it**
- Some non-obvious, but important and useful features include:
 - Rich text
 - Inline pictures / tables
 - Indicator database integration
 - Mass close/open/edit actions
 - Hierarchical tickets
 - API
 - Tagging / attack-cycle alignment
 - Keyboard navigation
 - Built in knowledge database/wiki
 - Workflow customization
 - Automation / API access
 - Activity "wall"

Incident Management System Features

Some of the features that may not be initially considered but add a great deal of usefulness to incident management systems are those shown above.

- Rich text, pictures, and tables are desirable as they make note taking and reading a much better and easier experience compared to plaintext entry. The ability to include inline screenshots and text highlighted in red can be invaluable for those reading the analysis after the fact.
- Indicator database integration is a must-have. Without it, tying a specific domain name or hash value to a ticket or correlating it with previous sightings can be overly complicated or impossible.
- Mass close/open/edit actions help in situations where you must open multiple tickets at once or close a bunch that were made in error. There's nothing more painful than having to manually click close on 100 alerts because someone wrote an overly sensitive rule.
- Hierarchical tickets can be great for threat intelligence and organization. Many incident systems allow making one ticket a "parent" or "child" of another ticket. This type of setup can be used to group multiple infections under one parent ticket about the specific virus for example or coordinate the activity of a threat group over time by putting all tickets attributed to that group under a parent item for the threat group. Some incident management systems even implement this style of tracking by design.
- Attack-cycle alignment means some ability to mark an incident with the progress that the adversary made toward the goal. Over time, this allows metrics tracking that is useful for all sorts of conclusions, such as how far most attacks are getting and how. If this is not specifically built in, attack cycle tracking is one of many things you can force to work via a tagging system. Tagging allows the analyst to add an arbitrary data tag to any ticket. In the past, I have seen SOCs use these for delivery vector tracking (USB, web, email, etc., tag), the source of an alert, or any other useful bit of information the incident management system doesn't explicitly support itself.

- Everyone knows keeping your hands on the keyboard is much faster than using a mouse. Analysts should be able to fill in fields and tab to the next one, and, in general, use the mouse as little as possible when filling out information in the system. Trust me when I say this. I know this sounds minor, but it can be a HUGE time saver.
- Knowledge databases are built into some ticketing solutions. If you do not already have a solution to this, having it built into the ticketing system can make it easier to set up and access.
- Workflow customization helps teams mold the solution to the statuses and processes the SOC team has. Ideally, the tool should mold to work with your process, not dictate it.
- Ideally, the ticketing system is compatible with any SOAR platforms so that actions can trigger new ticket creation and processing. API access is another way to develop custom scripts to create metrics or take actions on data that can be pulled from each ticket.

Licensed To: David Owerbach <0mamaloney0@gmail.com> April 28, 2020

Playbooks

- Playbooks mean different things to different SOCs
 - In this class: "a set of expected actions for alert response"
- Often implemented through your IMS or SOAR platform
- Contain required and optional steps for analysis and closure
- Guide analysts toward standardized analysis and completeness
- Unique for each type of case (phishing vs. malware playbooks)



Playbooks

As a means of efficiency, many SOCs strive to have cases of similar types investigated in similar ways. To ensure a complete analysis with a consistent response, lists or workflows of steps that should be taken are made that can be applied to each case of the same type. These steps are often referred to as "playbooks." The idea is that even if an analyst is new to the group, or even information security in general, playbooks will guide them down the necessary path, regardless of whether they could produce the response steps themselves. Therefore, this elevates the capabilities of SOC across the entire group and ensures the best possible response, regardless of the person who assumes responsibility for the ticket.

Playbooks will come in many forms and should align with the items required to thoroughly investigate an incident. For example, the steps of a simple playbook for a malicious email wave are shown on the slide. In some cases, not all steps of a playbook will perfectly apply or be required for a given situation; therefore, in many cases, playbook steps can be labeled optional or mandatory. Once each mandatory step is complete, the case can be closed inside the system. Therefore, once an analyst assigned themselves a case in the IMS, this usually translates to them assigning themselves the set of atomic steps that make up the playbook. These steps must be completed before the case can be closed (unless the playbook supports assigning individual steps to different analysts – TheHive supports this).

TheHive: Incident Management System¹

- Outstanding free IMS: Used for class virtual machine
- Incidents are organized and assigned by **case**
 - Cases follow steps in a pre-made **case template** (playbooks)
 - Cases have **tasks** that to be completed
 - Tasks have associated **worklogs**
 - Cases can have **observables** assigned to them
 - Observables can be enriched by **analyzers** enabled in **Cortex** engine

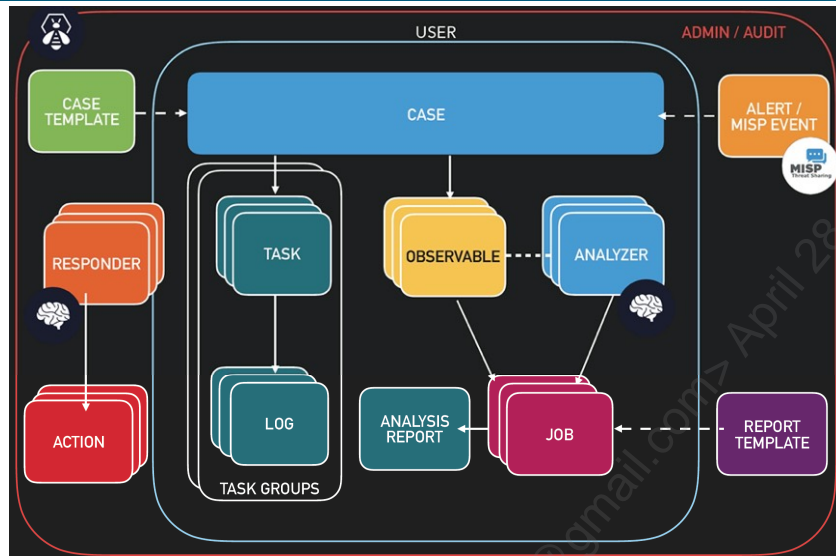


TheHive: Incident Management System

For training in the class, we will use a free, open-source incident response system called TheHive. TheHive is highly representative of most security-oriented incident management systems in that it supports alert input, case creation, playbooks in a form it calls "case templates", and tasks inside those playbooks that each individually have their own "worklog" where notes can be taken. Indicators associated with a case are all stored in what is called "observables" and these can be automatically sent to threat intelligence platforms, as well as enriched through a complementary piece of software called the Cortex engine. We will discuss Cortex further in a bit but, for now, know that it will act as a SOAR system in our in-class virtual machine setup, automated data gathering, enrichment, and context entry.

[1] <https://thehive-project.org/>

TheHive: Workflow Illustrated



TheHive: Workflow Illustrated

This slide shows the picture created by TheHive team to illustrate the entities involved in the system, and how they relate to each other. At the top, we can see the alerts or MISP events (a threat intelligence platform we will also use) are entered into the system and become a case, which has a case template (playbook) applied to it depending on the nature of the alert. Derived from the case template, multiple tasks are created inside the case that must be completed before the case can be closed. In addition, observables like IP addresses, usernames, hostnames, URLs, or hashes can be added to the case. Cortex (represented by the white brain) then applies a modular set of analyzers to these observables, which will further enrich the data and give the analyst additional information to help understand them. Additionally, in the newest versions of TheHive, responders can be created to take automated actions based on the data in the ticket. These can be things like sending an email or kicking off data collection and response tasks. In this way, TheHive and Cortex also act as an automation platform.

https://github.com/TheHive-Project/TheHiveDocs/blob/master/additional-resources/TLP-WHITE-TheHive-MISP_Summit_04v2.pdf

TheHive: Automatic Case Creation with Context

1. **Events** collected in SIEM, items of interest become **alerts**
2. **Alerts** sent to TheHive for triage
 - New **case** created for all accepted alerts
3. **Case** is populated with field from alert:
 - Parses fields from alert
 - IP addresses, domains, usernames, hostnames, etc.
 - Pulls in additional info if available
 - **Tasks** created from designated **case template** (playbook)



TheHive: Automatic Case Creation with Context


For the incident management system, the basic workflow is one most people are familiar with—a ticket comes in and gets assigned to an analyst. That analyst then works it to completion in the order of priority. We want to understand these systems at a deeper level than that, however. For example, what data is being transferred to create a new alert? Consider the actual plumbing for the data moving in and out of the system. Ticket systems inherently involve many fields that need to be filled out, and those fields may come from a variety of sources beyond just the SIEM. Even when they do filter through the SIEM, the fields included in logs from different log sources will contain wildly different information. In the ideal world, regardless of the fields included in a log, whether it's virus name, source IP, or hash – all these items would be automatically pre-populated into a case as observables or otherwise, leaving the analyst to do the non-tedious work. To facilitate this part of the alert to case process, TheHive team has recently produced a solution called Synapse, which is a Python 3 app that uses custom "connectors" to ensure that all relevant fields from any source can make a case with the maximum amount of information and context automatically added. They also have produced Python modules called TheHive4Py that make API interfacing easy.

Without a system equivalent to Synapse for your own IMS, new items would need to be manually filled out. This type of tedious work is what analysts typically hate, and worse, spending their time cutting and pasting information adds no value to the investigation. Modern incident management systems should strive to automatically fill out all fields possible as to make the analyst's life easier, no matter what the log source.

Case and Alert Naming Convention

Case names should be **meaningful** and convey **priority**:

- Cisco SOC / book, *Crafting the InfoSec Playbook*¹
 - { \$UNIQUE_ID } - { HF, INV } - { \$EVENTSOURCE } - { \$REPORT_CATEGORY } : \$DESCRIPTION
 - ID: Most significant digits indicate source
 - "**High-Fidelity**" or "**Investigative**" detection
 - **Event Source**: Correlates with unique ID
 - **Category**: Malware, Policy, APT, etc.
 - **Description** of what the rule attempts to detect



Case and Alert Naming Convention

There are many strategies that exist for an incident ticket naming convention for incident tickets. One of the most practical ones I've seen is from *Crafting the InfoSec Playbook* by Jeff Bollinger, Brandon Enright, and Matthew Valites. Aside from an outstanding look at the best practices discovered and utilized within the Cisco CSIRT, this book contains the suggestion that the below fields should be included when creating incidents (or alerts, depending on workflow).

Format: { \$UNIQUE_ID } - { HF, INV } - { \$EVENTSOURCE } - { \$REPORT_CATEGORY } : \$DESCRIPTION

Where:

- \$UNIQUE_ID
- HF, INV = High fidelity (almost sure it's not false positive) or investigative alert (one that isn't yet matured or isn't very reliable)
- \$EVENTSOURCE = The source of the data that originally found the event. This would be IDS, Antivirus, Firewall, etc.
- \$REPORT_CATEGORY = These can be split any way that works best for your team. They suggest things similar to the Emerging Threats Snort ruleset – Malware, Policy, etc.
- \$DESCRIPTION = A short but meaningful description of what condition the alert identifies

With this convention, it would be easy to manage alert flow since so much information is directly in the title. For example, high-fidelity alerts might get sent straight to IMS for ticket creation, where investigative alerts might not. You can also do interesting statistics like finding which of your data sources are actually giving you the best detection capability, or which is producing the most false positives. With the ability to tie alerts and case names to closure codes, lots of useful metrics may be derived.

[1] <http://shop.oreilly.com/product/0636920032991.do>

TheHive: Working a Case

Assigned tasks show up in "My tasks"

My tasks **2** Waiting tasks **10**

Tasks can be part of **task groups**

- Investigation Questions: Aligned with "kill chain" or other
 - Delivery: "How many malicious emails were delivered? Where did it come from?"
 - Exploit/Install: "What exploit does the file use? What files does it drop?"
 - C2: "What domain is used for command and control?"
- Response Actions
 - "Delete message from all inboxes"
 - "Reset all compromised passwords"
 - "Submit sample to antivirus vendor"

Tasks populated from **case template** (playbook) must be completed

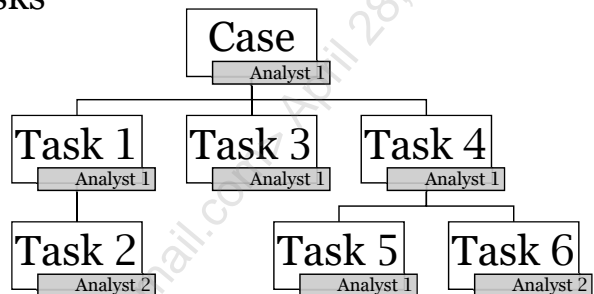
TheHive: Working a Case

In TheHive, at the analyst level, you may be assigned cases and tasks. In order to resolve a case, all tasks within that case must be completed. The specific list of tasks assigned to any given case will be dependent upon the *case template* for that type of case. This is how playbooks are implemented through TheHive. Case types are pre-configured such that any time a ticket for something like phishing comes up, the typical steps required for investigating a phishing event can be put into the case. This ensures a consistent response across all phishing cases and that analysts will not accidentally forget a containment or response actions.

The types of tasks assigned as part of the case can be further logically divided into what TheHive calls "task groups." For instance, tasks could be divided between "investigative questions" and "response actions" as shown on the slide. Or even further into questions aligned with frameworks like the cyber kill-chain—questions about the delivery stage, exploit method, installation, command and control, etc. (the cyber kill-chain will be discussed later on). This division into groups helps logically separate the steps of the investigation in the analyst's head.

TheHive: Case and Task Assignment

- Both **cases** and **tasks** can be assigned to an analyst
 - Not all IMSs do this – more granular assignment
 - Enables easier collaborative / tierless SOC operations
- Newer analyst (1) takes easier tasks
- Analyst 2 takes complex tasks
 - Analyst 1 reads 2's work logs
 - Learns task over time
- Better load balancing



TheHive: Case and Task Assignment

A note on the ability to assign individual tasks: This setup is unique compared to some incident management systems in that it is more granular, allowing multiple people to work on the same incident at the same time. Some solutions only allow assigning of entities at the case level, meaning that once you decide to start a case, you must complete the entire thing by yourself and if you are unsure of how to do a step, you may get stuck. TheHive's capability to assign individual tasks lends itself well to tierless SOCs. With this arrangement, a complicated ticket with advanced steps such as malware or memory analysis may be still be taken by a newer analyst, but the steps that analyst is unfamiliar with can be individually passed on to someone who has the experience to complete the task. Both people can easily work the ticket in parallel and when complete, the original analyst can review the worklog for the tasks completed by others, which over time will act as instructions on how to perform that task.

TheHive: Case Template

M Case # 6 - [Phishing Wave - attachments] Phishing wave reported - "Invoice #4251 - OVERDUE"

Group	Task	Date	Assignee	Actions
Recon	Are the users who received the email related in any way?		Analyst1	▶ Start ⚙
Delivery	Check how many emails are delivered organization wide		Analyst1	▶ Start ⚙
Delivery	Identify source email addresses and IPs		Analyst1	▶ Start ⚙
Exploit	What exploit does the file use?		Analyst2	▶ Start ⚙
Install	What happens if the attachment is opened?		Analyst2	▶ Start ⚙
C2	What domains/IPs are used for command and control?		Analyst2	▶ Start ⚙
Resposne	Block email address of sender		Analyst1	▶ Start ⚙

TheHive: Case Template

Here, we see an example of what the task tab in a newly created case looks like. This list of tasks was prepopulated by the case template called "Phishing Wave – attachments" which is designated with the prefix in the case name. The case template was designed to reflect common investigative questions as well as response actions that would need to occur in any case with a phishing wave that included a malicious file. The tasks are additionally labeled with task groups reflecting the stage in the cyber kill-chain that they address, such as Delivery, Exploit, Install, etc. The task groups are shown on the left side and the individual tasks have their own line with either Analyst1 or Analyst2 assigned to them. Once a task is assigned to someone, they can click the start button to begin work on that individual task.

Observables

Create new observable(s)

Type *

domain

}

Type *

autonomous-system	hash	regexp
domain	ip	registry
file	mail	uri_path
filename	mail_subject	url
fqdn	other	user-agent

Value *

evilsite.com
 reallybadsite.biz
 malware4free.tk

(one observable per line) 3 unique observables

TLP * **Is IOC** **Has been sighted**

WHITE GREEN AMBER RED

★ ☐

Tags **

phishing
Add tags

Description **

Phishing contains MS Word document that attempts to reach out and download 2nd stage installer from these domains.

Observable List (3 of 3)

		Type	Value/Filename
<input type="checkbox"/>	★ <input checked="" type="checkbox"/>	domain	malware4free[.]tk phishing No reports available
<input type="checkbox"/>	★ <input checked="" type="checkbox"/>	domain	evilsite[.]com phishing No reports available
<input type="checkbox"/>	★ <input checked="" type="checkbox"/>	domain	reallybadsite[.]biz phishing No reports available

SANS

SEC450: Blue Team Fundamentals: Security Operations and Analysis

106

Observables

TheHive allows you to associate observables to an incident at the case level. These are meant to be interesting snippets of data found within the course of the investigation, but do not necessarily have to represent malicious infrastructure or files. Most options you will need for observable types are built into TheHive and are shown in the Type dropdown on the slide, adding more is easy as well. Multiple observables of the same type can be entered at once by placing them on separate lines – a very convenient feature for mass observable entry! The "is IOC" star can be used to control whether the individual observable is an "indicator of compromise" or whether it is just a piece of data potentially associated with the case.

In the screenshot, three domains have been added as IOCs and the "has been sighted" checkbox has been switched. This was done because, in this pretend scenario, a malicious phishing email was received with an MS Word document that contained an autorunning macro. That macro would have reached out to the domains in question; therefore, these indicators have been seen in the actual environment as part of an active compromise attempt. If we were investigating something that we hadn't seen in our organization, then this checkbox could be unswitched to signify that it is not something the organization has encountered yet.

Case Closure

Once tasks are completed, cases can be classified and closed

Status * Incident

True Positive False Positive Indeterminate Other

Investigation clearly demonstrates that there is something malicious

Impact *

Yes No

Something altered availability, integrity or confidentiality

Summary *

B I H S Preview

The malicious document was delivered to 1000 people, of those, 2 people opened the file and enabled the macro. The 2nd stage download site was blocked by the proxy but the macro did leave a startup item that had to be manually removed. No damage was done since the callback was contained. Emails were removed from inboxes and recovery is complete.

Additional information

VERIS Delivery Category Email attachment : Email via user-executed attachment

Incident Detection User

System Impact Minimal impact to non-critical services

Type Of System Affected Laptop(s)

Case Closure

Once a case has been worked to completion, which is defined as having completed at least all required tasks in the playbook, the case can be closed. In TheHive as well as other IMS software, this typically entails labeling the case as a true/false positive, marking any categories that apply to the incident, and writing a brief summary of what the ticket was about and how it was resolved. On the slide is an example of a potential closing note for a delivered phishing case. Note the additional information stating the delivery vector, method of incident detection, impact, and type of systems affected. This data can be aggregated over time across all alerts to create reports and show trends in the type of attack the SOC is dealing with. Do not overlook the importance of ticket classification. Over time, this activity will build into your best source of threat intelligence as your organization will know what types of attacks are most common and which cause the most impact. These metrics can then be used to justify spending on bolstering defenses in these areas. They also form a crucial feedback mechanism to the threat intelligence team.

Incident Categorization Frameworks

There are many options for categorization. Here are a few:

1. **VERIS:** Vocabulary for Event Recording and Incident Sharing

- Captures 4 A's: **Actor, Action, Asset, and Attributes** (how affected)
- Yearly DBIR report data collected in this format



2. **US-CERT** Incident Reporting System Categories

- Medium-level detail
- Designates "what" and "how" and impact



3. **Tags**

- Track any arbitrary data

Incident Categorization Frameworks

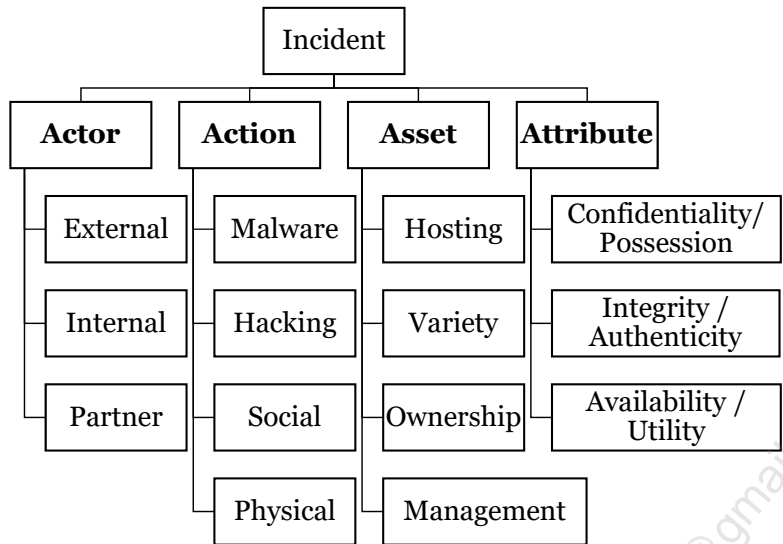
Metrics are one of the most important things your ticketing system can make for you, so incident classification features rank highly in the list of needs for an incident management system. When it comes to classification, there are several different frameworks available, and each has its own strengths and weaknesses. Ultimately, the choice of classification framework will come down to the needs of the organization, and the types of questions you are trying to answer about each incident. Systems like VERIS (Vocabulary for Event Recording and Incident Sharing) allow for highly detailed classification that explains the delivery vector, motivations, and impact of an incident.¹ Simpler frameworks like the US-CERT Incident Notification guidelines make classification slightly simpler but leave some detail out.²

Simple frameworks can be fast and easy to use but will paint a weaker picture. Thorough frameworks are outstanding for getting actionable metrics but can suffer in usability from having too many options. Using a more detailed framework often becomes a challenge to ensure everyone on the team is filling out items in a consistent way, as well as convincing everyone to fill out each item in the framework for all incidents they work. A "middle of the road" approach that strikes a happy medium between detail and usability is best.

[1] VERIS: <http://veriscommunity.net/>

[2] US CERT: <https://www.us-cert.gov/incident-notification-guidelines>

Classification Options: VERIS



Actor: Whose actions affected the asset?

Action: What actions affected the asset?

Assets: Which assets were affected?

Attributes: How was the asset affected?

Classification Options: VERIS

VERIS is one of the most comprehensive incident categorization frameworks out there. This slide shows at a high level some of the categories that could be used for each of the 4 A's. Within each one of these are suboptions that can be used if desired. For example, under action/malware, you can place a "vector" and "variety" statement that tells what type of malware it was and how the malware was delivered. Explanations of the schema can be found on the GitHub page for VERIS.¹

[1] <http://veriscommunity.net/>


US-CERT Incident Categorization

NCISS-based scores from 0-100 based on weighted categories:

- **Functional:** "No impact" to "DoS/Loss of Control"
- **Information:** "No impact" to "Destruction of critical system"
- **Recoverability:** "Regular" to "Not recoverable"
- **Attack Vectors**
 - Web, Phishing, Ext. Media, Impersonation, Improper Usage, Theft, etc.
- **Incident Attributes**
 - Location of Observed Activity – "DMZ" through "Safety Systems"
 - Actor characterization and potential impact

Calculated Score

58



US-CERT Incident Categorization

Another method of categorization is the US Government's system for classification of incidents under the Federal Information Security Modernization Act of 2014 (FISMA). The system is based on the National Cyber Incident Scoring System (NCISS) described in NIST SP800-61 Rev. 2 and uses a weighted arithmetic mean to create an incident score between zero and 100.¹ The incident factors are broken down by categories that are independently weighted and the response to each category has an associated score. The categories are:

- Functional Impact
- Observed Activity
- Location of Observed Activity
- Actor Characterization
- Information Impact
- Recoverability
- Cross-Sector Dependency
- Potential Impact

The answer to each category is multiplied by the category's weight and the resulting calculation is described on a priority level that is either emergency, severe, high, medium, low, or baseline. An interactive demonstration of the scoring system can be found at the link below.² Although the categories may not translate 100% to commercial organizations, the categories and weighting approach can be adapted to create a similar quantitative score.

[1] <https://www.us-cert.gov/NCCIC-Cyber-Incident-Scoring-System>

[2] NCISS Scoring Demo: <https://www.us-cert.gov/nciss/demo>

Classification Options: Tagging

Don't want to get too formal? Use **tags!**
Simple tagging can add useful info for metrics:

- Delivery: USB, Phishing, Web
- Attack Type: Opportunistic, Targeted
- Discovery Method: Threat Hunting, AV, user report, etc.

Allows answers like:

- "What percentage of incidents came from email?"
- "What is threat hunting finding?"

Summary	
Title	Targeted phishing email received
Severity	H
TLP	TLP:RED
PAP	PAP:AMBER
Assignee	student
Date	Thu, Oct 4th, 2018 8:21 -07:00
Tags	phishing targeted

Classification Options: Tagging

Don't get stuck thinking you have to use a formal framework! The tagging option available in most incident management solutions allows adding arbitrary data to incidents that can later be pulled and visualized to answer additional questions. Have you ever been asked what finds threat hunting has produced? Or how bad the phishing or removable device problem is? What about how many targeted attacks you have per month? Simple tags consistently applied to incidents can provide this information. Tagging allows you to wrench in metrics for tracking things your IMS solution may not support out of the box.

Incident Management Systems Summary

Your IMS is one of the most important pieces of software!

- **Test the interface**, or you will be miserable!
- **Cases** should be created from alerting source
- **Context** from alert should be parsed into fields
- **Task level assignment** makes work balancing easy
- **Observables** should be entered – ideally automatically
- **Playbooks** guide you through tasks
- Close with **categorizations** for metrics

Incident Management Systems Summary

One of the most important pieces of advice I can give is that you choose your IMS extremely carefully. Although there is no single perfect solution, the IMS will be a primary tool that will be worked in constantly by all analysts. Having a suboptimal solution, you're forcing to make fit your workflow or one that requires lots of extra clicking, workarounds, or otherwise non-value-added time will wear on team morale. If there is a tool worth spending extra money and time to decide on, this is it.

A good IMS will disappear and become a natural extension of your process, allowing fast and efficient investigation. At a minimum, they should allow the development of customized workflows, track custom metrics, and enable a breakdown of incident types into playbooks with individual steps that can be followed and annotated by analysts as they are completed.

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. **Exercise 1.1: TheHive Incident Management System**
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Licensed To: David Owerbach <0mamaloney@gmail.com> April 28, 2020



Exercise 1.1: TheHive Incident Management System

Exercise 1.1: TheHive Incident Management System

Please go to Exercise 1.1 in the SEC450 Workbook or virtual wiki.

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. **Threat Intelligence Platforms**
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack


This page intentionally left blank.

Licensed To: David Owerbach <0mamaloney@gmail.com> April 28, 2020

What Is Cyber Threat Intelligence?

First of all, what is cyber threat intelligence?

- NOT just a list of bad domains and IP addresses...
- Summarized: **Analyzed cyber threat data** giving a **strategic** and **tactical advantage** over the adversary
 - Helps prioritize defensive resources
 - Drives "Offense informs defense"
- Made up of several pieces
 - Let's break them down...



1

What Is Cyber Threat Intelligence?

We've all heard the term, but there is still a lot of confusion and ambiguity around the term "threat intelligence." Although the term can seem sort of nebulous, many have put forth workable definitions. Regardless of the specific wording, these all seem to center around one main goal—gaining a strategic and tactical advantage over the adversary by understanding their tactics, techniques, and procedures that they will use to accomplish that goal. You've probably heard the term "offense informs defense" when it comes to the cyber threat intelligence space. This is a nod to the fact that knowledge of what our adversary does should be leveraged as much as possible to help drive defensive preparation. That is what cyber threat intelligence is all about.

Before we launch into our discussion on threat intelligence platforms, let's step through the components of cyber threat intelligence to get a good understanding of the concept. We will do so using the Venn diagram from an outstanding book on the topic—*Intelligence-Driven Incident Response* by SANS Instructors Scott J. Roberts and Rebekah Brown¹. Defining each layer will help us understand what threat intel is, and how TIPs help us accomplish our goals in the CTI space.


[1] <http://shop.oreilly.com/product/0636920043614.do>

Intelligence Definition

"Taking in external information from a variety of sources and analyzing it against existing requirements in order to provide an assessment that will affect decision making."

- Scott J. Roberts & Rebekah Brown, *Intelligence-Drive Incident Response: Outwitting the Adversary*¹

- **Examples:**
 - Weather Report: Do I need to bring a coat today?
 - Traffic Report: How much time do I need to get to work?



Intelligence Definition

First, what is intelligence? One concise definition is the quote on the slide from Scott J. Roberts and Rebekah Brown's book.¹ This explanation contains some of the crucial items such as the requirement for analysis, comparing it against an existing requirement, and using it for actionable decision making. Across many definitions of the term intelligence, these are themes that are often encountered.

Sergio Caltagirone, Director, Threat Intelligence and Analytics at Dragos, Inc., sums up intelligence in a slightly more academic way in his blog post referenced below: "Intelligence is the collecting and processing of that information about threats and their agents which is needed by an organization for its policy and for security, the conduct of non-attributable activities outside the organization's boundaries to facilitate the implementation of policy, and the protection of both process and product, as well as persons and organizations concerned with these, against unauthorized disclosure." And: "I propose that cyber threat intelligence is nothing more than the application of intelligence principles and tradecraft to information security."²


[1] <http://shop.oreilly.com/product/0636920043614.do>

[2] <http://www.activeresponse.org/threat-intelligence-definition-old-new/>

Threat Definition

How to define threat?

- "... combination of intent, capability and opportunity." – Rob M. Lee
- **Intent** is a malicious actor's desire to target your organization
- **Capability** is their means to do so (such as specific types of malware)
- **Opportunity** is the opening the actor needs (such as vulnerabilities, whether it be in software, hardware, or personnel)



Threat Definition

When we add the word "threat" in front of intelligence, what type of threat are we talking about? Rob M. Lee, SANS Certified Instructor and CEO and Founder of Dragos, Inc., uses the 3 characteristics on the slide to specify what is and is not a threat to an organization. The requirements state that the entity must have the **intent**, **capability**, and **opportunity** to cause harm. If any items are missing, the adversary will not be able to, or have reason to attack an organization.

Using this definition of threat to answer the question "what is threat intelligence?" then leads Rob to describe it as "The process and product resulting from the interpretation of raw data into information that meets a requirement as it relates to the adversaries that have the intent, opportunity and capability to do harm."¹ This definition aligns with the previous slide in that it highlights the need to turn raw data into information (through analysis) and must meet some requirements laid out as a goal of the analysis.


[1] <http://www.robertmlee.org/intelligence-defined-and-its-impact-on-cyber-threat-intelligence/>

Cyber Threat Intelligence Definition

"Threat intelligence is the analysis of adversaries – their capabilities, motivations, and goals; and cyber threat intelligence (CTI) is the analysis of how adversaries use the cyber domain to accomplish their goals."

- Scott J. Roberts & Rebekah Brown, *Intelligence-Driven Incident Response: Outwitting the Adversary*¹

- Definition themes:
 - Human analysis required
 - Identify TTPs and goals of threat actor
 - Output drives decision making



Cyber Threat Intelligence Definition

So, what then is Cyber Threat Intelligence? It is the combination of terms previously discussed, as applicable to the cyber domain. According to dictionary.com, cyber means "relating to or characteristic of the culture of computers, information technology, and virtual reality." Therefore, cyber threat intelligence relates to threat intelligence in the realm of computers and information technology. The quote from *Intelligence-Driven Incident Response* on the slide above then puts it all together. "Threat intelligence is the analysis of adversaries – their capabilities, motivations, and goals; and cyber threat intelligence (CTI) is the analysis of how adversaries use the cyber domain to accomplish their goals."¹

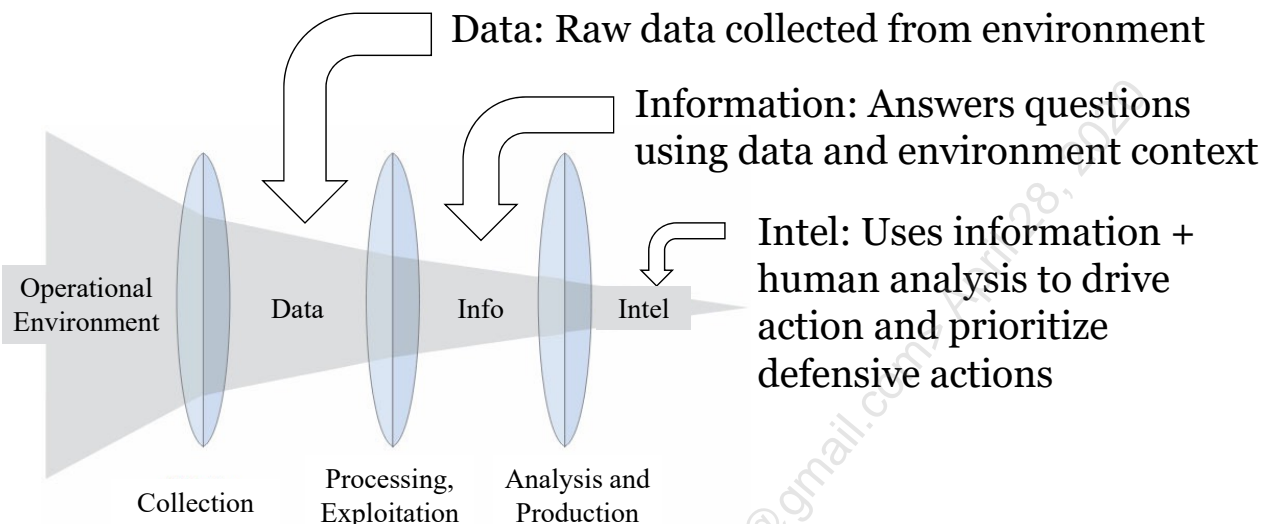
Michael Cloppert, co-author of Lockheed Martin Cyber Kill Chain paper, further breaks this down into Cyber Threat Intelligence Operations and Analysis: "I define Cyber Threat Intelligence Operations as *actions taken in cyberspace to compromise and defend protected information and capabilities available in that domain*; I define Cyber Threat Intelligence Analysis as *the analysis of those actions and the actors, tools, and techniques behind them so as to support Operations*; and **I define the Cyber Threat Intelligence domain as the union of Cyber Threat Intelligence Operations and Analysis.**"²

Regardless of the specific wording, these all seem to center around one main goal—*analyzing* cyber threat *data* to create information about a threat actor that helps us understand their tactics, techniques, and procedures and what their goals might be. This analysis must be taken on with specified requirements to meet an analysis goal, and produce an output formatted for the intended audience as well as assist in decision making or produce otherwise actionable output.

[1] <http://shop.oreilly.com/product/0636920043614.do>

[2] <http://www.activeresponse.org/threat-intelligence-definition-old-new/>

Threat Data vs. Information vs. Intelligence



SANS

SEC450: Blue Team Fundamentals: Security Operations and Analysis

120

Threat Data vs. Information vs. Intelligence

This leaves one item left to be defined: What is threat data vs. threat information vs. threat intelligence? This graphic from the US Department of Defense Joint Publication 2.0 describes how one becomes the next throughout the intelligence generation process.¹

Threat **data** is the raw information and unarguable facts collected from the environment. In the case of network security monitoring, for example, this would be the logs and packets recorded by our sensors on the network. On their own, this does not represent intelligence or even information—merely the raw data that must be analyzed to create threat information or threat intelligence.

Threat **information** is an intermediate step that is all about using analysis to answer a question using the threat data as input. Threat information doesn't necessarily have the requirement to inform action; it is, however, the next step to producing threat intelligence. For example, is a system on the network compromised? We could answer that question using threat data collected from the environment, possibly by looking for signs of malware communication or execution.

Threat **intelligence** requires gathering lots of threat information and aggregating it to make an assessment of some sort. This analysis of multiple bits of threat information drives an organization's security policy, spending, and defensive posture. It attempts to align defensive actions against what appears to be the TTPs of actors that are a threat to the organization or, in other words, "offense informs defense."²

For more information, threat intelligence vendor Recorded Future has a useful blog post that explains these concepts in more detail.

[1]Joint Publication 2.0 – Joint Intelligence: http://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/jp2_0.pdf

[2]Threat Intelligence, Information, and Data: What Is the Difference? <https://www.recordedfuture.com/threat-intelligence-data/>

Threat Intelligence Platforms and You

- There are threat intelligence **producers** and **consumers**
 - Many SOCs contain a threat intelligence group
 - Threat intel group produces intelligence, while analysts consume it
- Analyst jobs require using threat **data, information, and intelligence** to identify and protect against compromise
- **Threat intelligence platforms** help you accomplish this task
 - Knowledgebase of your threat data, information, and intelligence
 - Automates exchange and querying of data from other security tools
 - Do NOT produce intelligence for you!

Threat Intelligence Platforms and You

When it comes to threat intelligence, there are often two sides to the coin—producing it and consuming it. As SOC analysts, most of the work will require the consumption of threat intel while the majority, but not all, of the production of intel may be done by a dedicated threat intelligence group.

In your daily job, it is highly likely that you will be using threat data, information, and analysis. Activities such as validating alerts, looking for threats, and investigating incidents will almost surely involve referencing threat intelligence at some point.

Threat Intelligence Platform Features

Much of the Blue Team alerting will be based on indicator lists

- Known bad IPs, domains, hashes, etc.

Need a solution to:

- **Store analysis** and **threat information** for known indicators
- Perform **automated** / **fast lookups** via API
- Record **context** about stored items (NOT just list)
 - Ex: IP 1.2.3.4 resolved to evilsite.com serving exploit kit on 2018-02-19
- **Find associations** across multiple events
- **Sharing** of indicators with other organizations

Threat Intelligence Platform Features

A threat intelligence platform's main purpose is to store the body of threat information, analysis, and indicators you have collected and then make it available in an easy-to-search way. The database will be manually searched as part of incidents and thus needs to have a user-friendly interface that makes the correlation of data across events easy. It also needs to have good integration capability so that tools such as your IMS and SIEM can send and pull information from it through an API of some sort. This is a crucial feature. Threat intelligence locked up in a proprietary system that can't be leveraged by your other security devices will be of minimal use, so communication across security tools is a highly important capability.

Threat Intelligence Platform Requirements

Indicators or low-level configuration details?

- Most TIPs handle indicators of compromise with ease
 - IP Addresses, filenames, domains, hash values, URLs, etc.
 - Important feature: **easy bulk entry and integration**
- *Some* TIPs do a better job with additional features
 - Are you storing **malware configurations? Non-standard fields?**
 - How do you want to **correlate** across items stored?
 - Is **sharing** a required function?
 - What **volume** of indicators will you be storing?

Threat Intelligence Platform Requirements

One of the items that may drive your decision is the depth of detail of intelligence that you need to store, as well as your interest in sharing your findings. John Bambanek of Bambanek Consulting, a prominent malware researcher, points out in his talk referenced below that many TIPs were not capable of storing the level of detail he needs for his work.¹

Most threat intelligence platforms will be built for, and have no problem storing common indicators such as IP addresses, hashes, URLs, filenames and the like. If this is your main need, it is likely almost any solution will provide what you need, and the choice can be driven based on integration with other tools. If you are doing an analysis that needs to include in-depth, arbitrary field names, however, some solutions are likely better suited than others. According to Bambanek, who has tried many of the solutions listed on the previous slide, MISP was the solution that provided him the required amount of field storage flexibility, as well as the ability to handle high volume indicators, sharing features, and correlation capability he was looking for.

[1] <https://www.youtube.com/watch?v=6k-1QEQFgmI>

De-fanged Indicators

When dealing with indicators of compromise, beware!

- IP addresses, links often become "live" when entered
- Many articles / tools "de-fang" them¹

IOCs

Domains and IPs:

- asushotfix[.]com
- 141.105.71[.]116

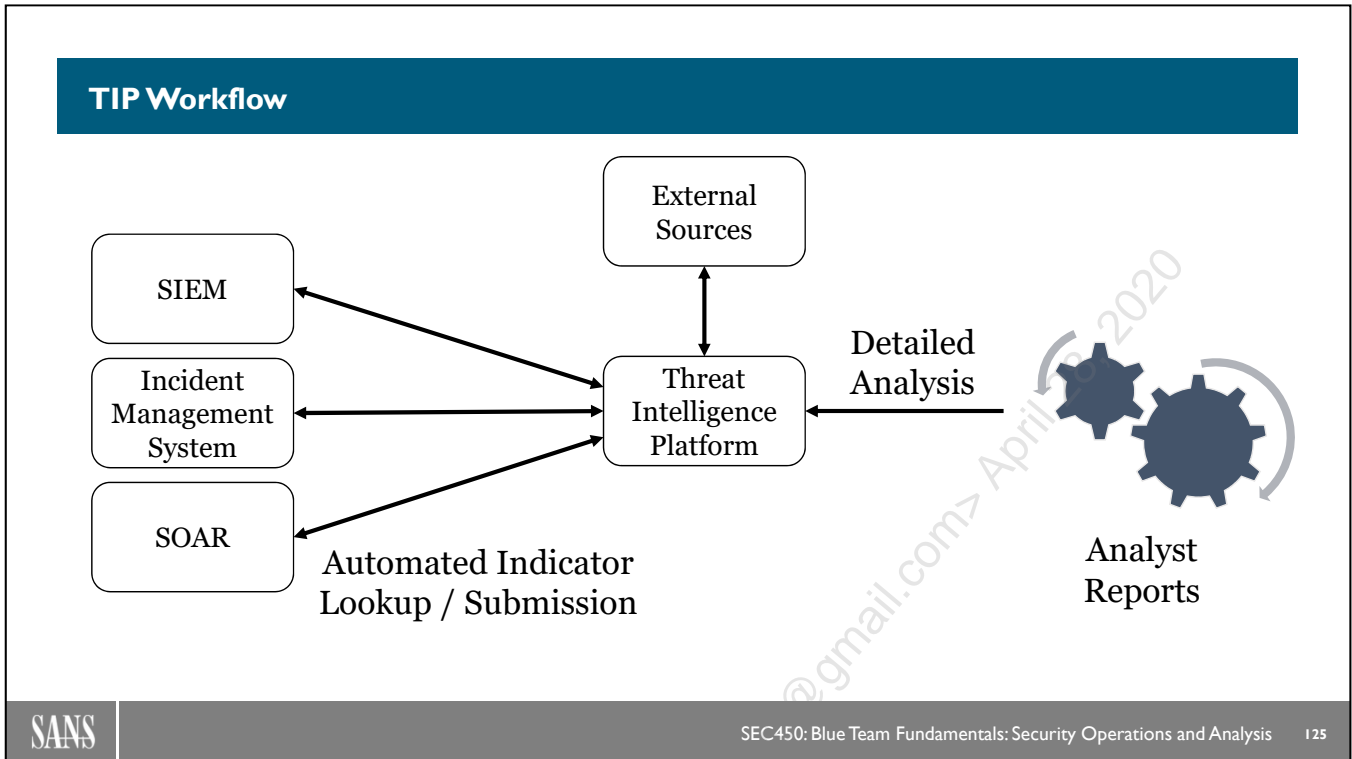
Some of the URLs used to distribute the compromised packages:

- hxxp://liveupdate01.asus[.]com/pub/ASUS/nb/Apps_for_Win8/LiveUpdate/Liveupdate_Test_VER365.zip
- hxxps://liveupdate01s.asus[.]com/pub/ASUS/nb/Apps_for_Win8/LiveUpdate/Liveupdate_Test_VER362.zip
- hxxps://liveupdate01s.asus[.]com/pub/ASUS/nb/Apps_for_Win8/LiveUpdate/Liveupdate_Test_VER360.zip
- hxxps://liveupdate01s.asus[.]com/pub/ASUS/nb/Apps_for_Win8/LiveUpdate/Liveupdate_Test_VER359.zip

De-fanged Indicators

One important detail that often relates to threat intelligence platforms is using "de-fanged" indicators. Since many solutions will take valid URLs and IP addresses and make them "live", analysts often need to take extra precaution that this does not happen inside their own notes and tools. If it did, we may open ourselves up to accidental clicking and unintentional infection. To combat this issue, square brackets are often used before the TLD in URLs and in IP addresses so that links are invalid and will not become clickable. You may also see http or other protocols written with "x" instead to accomplish the same. The picture on this slide is from the Kaspersky writeup blog post on the ShadowHammer attacks.¹

[1] <https://securelist.com/operation-shadowhammer/89992/>



TIP Workflow

Looking at the TIP with a systems mindset, we can see that most of the input and output from the system are indicators and threat information either being pushed to the TIP from security tools, or queries being sent to it. One note is that TIPs may take the source of information from outside of the organization. There are many vendors and industry groups that offer indicator feeds that can be used as a verified external source of threat information. As they say, "sharing is caring", and that applies to threat information as well, the more sources you have of malicious indicators and threat information (ideally tailored for your industry and environment), the more likely you are to catch attackers in your own network.

Threat Intelligence Platform Products

Self-Hosted, Free:

- **MISP** (Malware Information Sharing Platform)
- **CIF** (Collective Intelligence Framework)
- **YETI** (Your Everyday Threat Intelligence)
- **CRITS** (Collaborative Research Into Threats)

Cloud, Commercial:

- ThreatConnect
- AlienVault OTX
- Threat Quotient
- Anomali

Threat Intelligence Platform Products

This slide lists some of the options for threat intelligence products. In this space, there are many great free open source solutions, such as MISP, CIF, or Yeti. There are also many cloud-based systems from vendors such as ThreatConnect and AlienVault OTX. Which one is best for you will depend on the depth of analysis you're performing, and which tools integrate best with your existing systems.

One point to remember—a threat intelligence platform does not *produce* intelligence for you. It enables you to store and query the threat information and intelligence you have created on your own or obtained from others. If you're looking for an actual threat intelligence vendor, that is an entirely different product altogether that consists of the reports written by those company's analysts. Threat intelligence platforms are for the storage of your own intelligence and information.

MISP

In this class, we will use MISP for a TIP

- A free, open-source analyst favorite
- Capability of high-volume indicator storage
- Great web UI and REST API interface
- Classification and sharing functionality
- Flexible indicator storage
- Easy import/export
- Integrates with TheHive for automated storage/analysis



MISP

This class will use MISP as the example threat intelligence platform. MISP was chosen because it is emerging as the de-facto standard of free and open-source threat intelligence platforms and has gained a lot of traction with large organizations over the last few years. In addition, it has built-in integrations with TheHive, which is being used for an IMS and gives us a great idea of the type of integration we're looking to see from our SOC tools.

MISP has lots of attractive features and is a full-featured threat intelligence management tool that competes with most of the commercial options out there. The team behind it has done an outstanding job of keeping it constantly updated with new features and taking input from the community on what should be added. This slide calls out some of the most useful capabilities, and here are a few more from the MISP site.¹

- "An **efficient IoC and indicators database** allowing to store technical and non-technical information about malware samples, incidents, attackers and intelligence.
- **Automatic correlation** finding relationships between attributes and indicators from malware, attacks campaigns or analysis. Correlation engine includes correlation between attributes and more advanced correlations like Fuzzy hashing correlation (e.g. ssdeep) or CIDR block matching. Correlation can be also enabled or event disabled per attribute.
- A **flexible data model** where complex objects can be expressed and linked together to express threat intelligence, incidents or connected elements.
- Built-in **sharing functionality** to ease data sharing using different models of distributions. MISP can synchronize automatically events and attributes among different MISP. Advanced filtering functionalities can be used to meet each organization sharing policy, including a **flexible sharing group** capacity and an attribute level distribution mechanisms.
- An intuitive user-interface for end-users to create, update and collaborate on events and attributes/indicators. A graphical interface to navigate seamlessly between events and their correlations. Advanced filtering functionalities and warning list to help the analysts to contribute events and attributes.

- **Storing data** in a structured format (allowing automated use of the database for various purposes) with the extensive support of cybersecurity indicators along with fraud indicators as in the financial sector.
- **Export:** generating IDS (Suricata, Snort and Bro are supported by default), OpenIOC, plain text, CSV, MISP XML or JSON output to integrate with other systems (network IDS, host IDS, custom tools)
- **Import:** bulk-import, batch-import, free-text import, import from OpenIOC, GFI sandbox, ThreatConnect CSV or MISP format.
- Flexible **free text import** tool to ease the integration of unstructured reports into MISP.
- A gentle system to **collaborate** on events and attributes allowing MISP users to propose changes or updates to attributes/indicators.
- **Data-sharing:** automatically exchange and synchronization with other parties and trust-groups using MISP.
- **Feed import:** a flexible tool to import and integrate MISP feed and any threat intel or OSINT feed from third parties. Many default feeds are included in standard MISP installation.
- **Delegating of sharing:** allows a simple pseudo-anonymous mechanism to delegate publication of events/indicators to another organization.
- Flexible **API** to integrate MISP with your own solutions. MISP is bundled with PyMISP which is a flexible Python Library to fetch, add or update events attributes, handle malware samples or search for attributes.
- **Adjustable taxonomy** to classify and tag events following your own classification schemes or existing taxonomies. The taxonomy can be local to your MISP but also shareable among MISP instances. MISP comes with a default set of well-known taxonomies and classification schemes to support standard classification as used by ENISA, Europol, DHS, CSIRTs or many other organizations.
- **Intelligence vocabularies** called MISP galaxy and bundled with existing threat actors, malware, RAT, ransomware or MITRE ATT&CK which can be easily linked with events in MISP.
- **Expansion modules in Python** to expand MISP with your own services or activate already available MISP modules.
- **Sighting support** to get observations from organizations concerning shared indicators and attributes.
- Sighting can be contributed via MISP user-interface, API as MISP document or STIX sighting documents. Starting with MISP 2.4.66, Sighting has been extended to support false-negative sighting or expiration sighting.
- **STIX support:** export data in the STIX format (XML and JSON) including export in STIX 2.0 format.
- **Integrated encryption and signing of the notifications** via PGP and/or S/MIME depending on the user preferences."

[1] <https://www.misp-project.org/>

MISP Terminology¹

Events: Encapsulations for contextually linked information

- The main entity type you will be creating and adding attributes to

Attributes: Holds indicators (url, hash, IP), links, text

- Child item of events, have a **category** and **type** (md5, link, text), and comment

Classification Features

- **Sightings:** A way to count true/false positives for an attribute
- **Tags:** Additional way to add context to events
 - **Taxonomies:** Add families of pre-made tags
- **Galaxies:** Adds **clusters** of threat actors, tools, or "intelligence"



(0/0/0)

Type:OSINT x tlp:white x

MISP Terminology

To use MISP as an analyst, these are the terms you should be familiar with:

- The *event* is the object type that everything in MISP is centered around. For every report you read, malware you analyze, or incident you want to track, a new event will be created. Events have unique ID numbers associated with them and act as parent containers to hold a group of attributes.
- **Attributes** are the individual bits of data that are being tracked about an event. Options for attributes are the normal indicator-like pieces of data (hash value, filename, URL, domain, IP, etc.) but also can be links to outside articles, explanatory text, or attached files themselves. Attributes that are identical across multiple events will be highlighted for correlation.
- **Categories** for attributes describe how that attribute was used. For example, for a md5 hash, categories available are "payload delivery", "artifacts dropped", "payload installation", and "external analysis."
- **Type** is the type of attribute you are entering—md5, sha1, user-agent, email-subject, mime-type. MISP supports many attribute types for granular classification.
- **Instances** of MISP are a single running copy of the MISP process and are important to keep straight because, by nature, MISP is made to facilitate sharing. One instance of MISP can be linked to other copies for selective sharing of information within an organization, or with external organizations.
- MISP also has detailed classification features for noting the types of indicators seen, as well as the groups and tools associated with them.
- **Sightings** are a simple "thumbs up" and "thumbs down" mechanism developed to lend credibility to each *attribute*. The intention is for analysts to indicate when they've run into the indicator, whether it was a true positive (thumbs up) or a false positive (thumbs down). This way, each attribute can develop a reputation over time and bad or useless attributes can be expired.
- **Tags** are arbitrary data that can be attached at the *event* level to information being stored in MISP. They can be either part of a taxonomy, or any arbitrary text. Tags do not have to be pre-defined.

- **Taxonomies** act as pre-made lists of tags that can be defined and enabled within MISP. For example, there is a kill-chain taxonomy that creates tags called "kill-chain:Delivery" and "kill-chain:Installation". This makes it easy to organize and track the tags that will be suggested to analysts.
- **Galaxies** are similar to tags in that they label an event as part of a larger group but come broken down into families of options called "clusters" which are a pre-set group of values, sort of like a taxonomy of tags. For example, there may be a "threat actor" galaxy intended to track the names of adversary groups. In this galaxy, there would be a cluster for all known threat groups – APT1, APT5, etc. Each individual group can then be assigned synonyms—the "Anunak" threat group would have "Carbanak", "Carbon Spider" and "FIN7" as options, for example, since these names have been used by various vendors for this group. Any time an event is created that can be attributed to this group, the analyst can search for "anunak" or any of its synonyms to attach that cluster to the event. The difference between cluster tags is that each cluster can hold its own attributes and notes, making it a better mechanism for correlation and tracking true "intelligence." The intention of the galaxy/cluster mechanism is, therefore, more geared toward the goal of identifying adversaries, tools, and TTPs across multiple events as compared to tags, which are merely meant for simple labeling. Galaxies can be used in the correlation view the same as the attributes of events to tie multiple events together.

[1] Terminology from: <https://www.circl.lu/assets/files/misp-training/luxembourg2018/misp-training.pdf>

MISP Workflow Overview

Analyst Usage:

- Analyst creates new **event**
- All indicators, links, files, and notes are added as **attributes**
- **Tags** and other classifications (galaxies) applied
- Event reviewed, published to other organizations (if desired)

Automated usage through SOC tools:

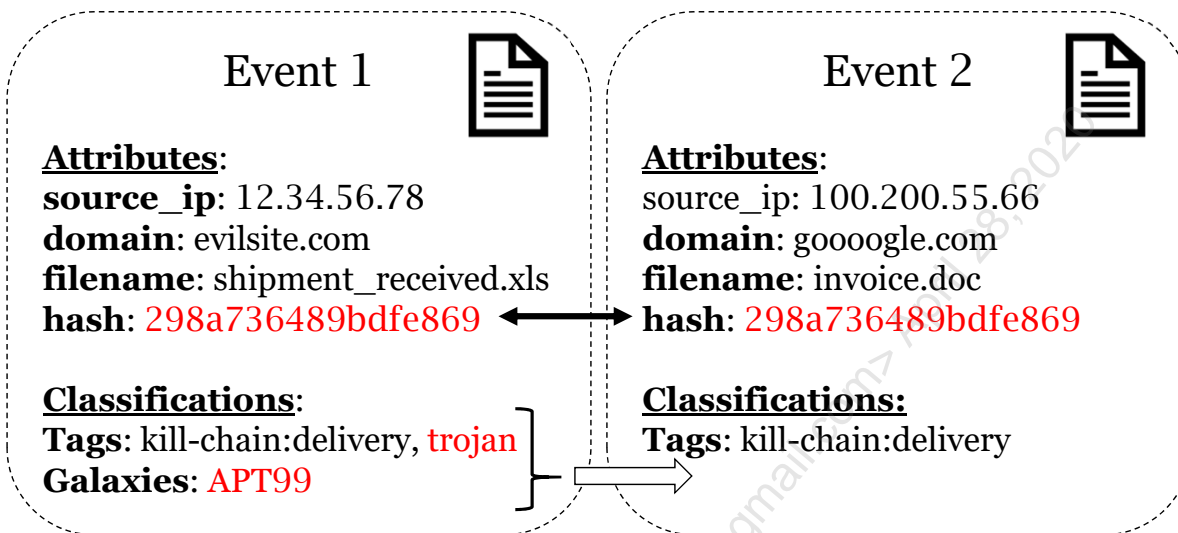
- SIEM, SOAR, IMS use API to look up or push attributes to event
- Subscribed feeds automatically download external event data
- Anytime any of them are seen in live traffic = Alert

MISP Workflow Overview

There are 2 main ways that you will utilize MISP as an analyst. One is manually interacting with it, searching for indicators, and creating new events with attributes and notes. The second is indirectly through lookups in other SOC tools. For example, when you create an incident in TheHive, it is possible to have TheHive automatically create a new MISP event for you and populate the attributes itself.

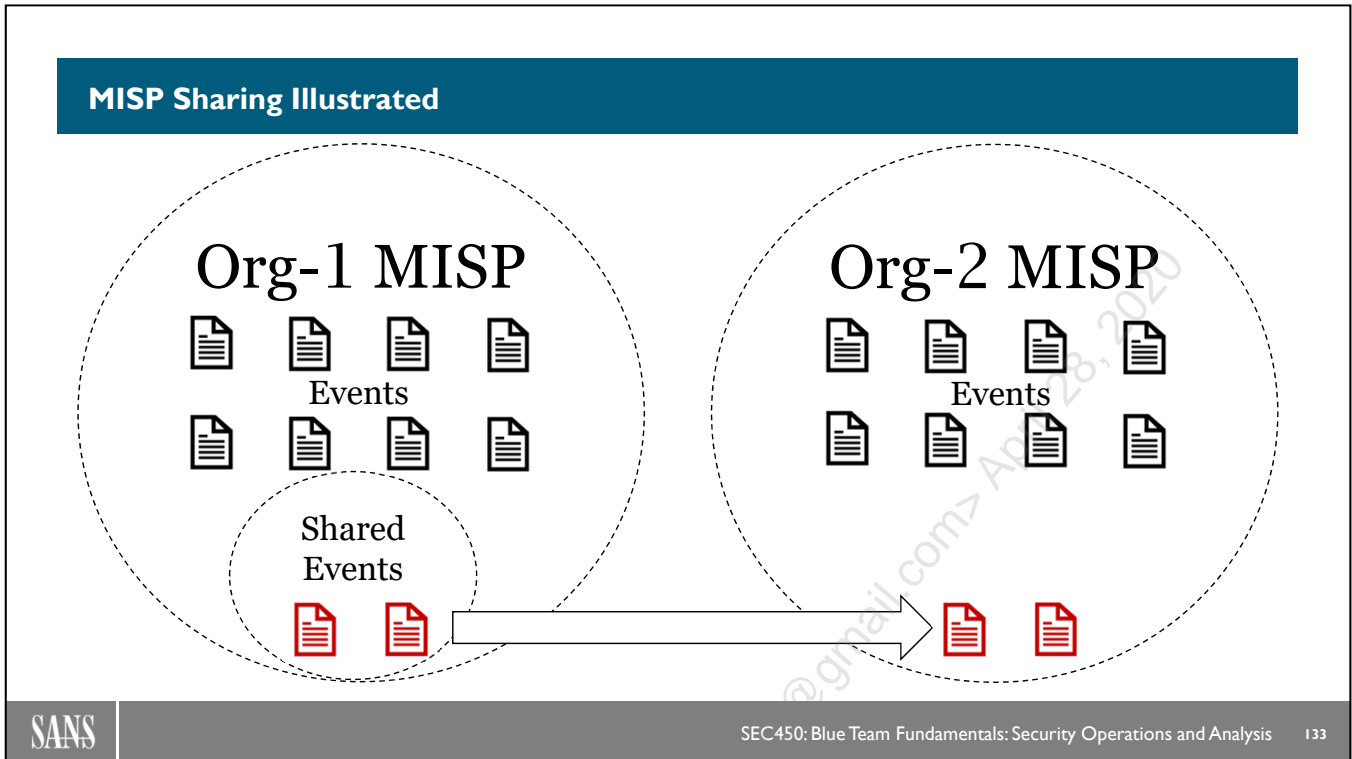
Even if your SOC has a separate threat intelligence team, it is highly likely that you will be putting indicators into a threat intelligence platform at some point. In MISP, this entails creating a new event, adding the appropriate tags, importing attributes from the article or intelligence source as well as the actual source of information, and classifying the intel against a set of tools or threat actors, if possible. Let's walk through a simple example of manually adding an event in MISP.

MISP Events Illustrated



MISP Events Illustrated

This slide shows 2 events that could represent links received via phishing. Each have their own set of attributes and classifications associated with them. But one attribute, the file hash, happens to be the same. Using a platform like MISP lets us easily recognize this fact. As the attributes of event 2 are entered into the system, a threat intelligence platform should highlight any other events that have the same values for indicators, allowing you to tie events 1 and 2 together and conclude that they are likely part of the same campaign from the same attacker, even though their filenames and source is different. Notice that event 1 has already been classified as a trojan from APT99. An analyst entering the items for event 2 would ideally recognize that this hash has been seen before and not have to repeat the work done to identify event 2 as another attempt to deliver a trojan, since it has already been classified during event 1. You would also immediately be able to attribute the newly observed infrastructure (gooooogle.com) to APT99 due to the previous attribution.



MISP Sharing Illustrated

MISP is designed for intelligence sharing, and threat intel is one of the many activities that's better with friends! Each organization that runs their own instance of MISP can designate which events they are willing to share, and link their copy of MISP up with others, allowing bidirectional sharing of events across multiple organizations.

Creating An Event in MISP

Home | Event Actions | Galaxies

List Events | Add Event | Import from...

Add Event

Date: 2017-06-27 | Distribution: All communities | Threat Level: High | Analysis: Completed

Event Info: Petya? I hardly know ya! - an ISC update on the 2017-06-27 ransomware

Owner org: SEC450

Contributors: Email: admin@admin.test | Date: 2017-06-27

All Tags: Taxonomy Library: PAP, kill-chain, malware_classification, tlp

Tags: ttp:white

Creating An Event in MISP

Let's walk through an example of making a simple event out of the SANS ISC post about NotPetya.¹

First, we click the Event Actions > Add Event button and fill in the details on the event we've created. Remember to use the date of the event, not the current date. For a title, we're using the title of the blog post itself, which is the usual convention.

The new event is now created. From the metadata section, we can now add a tag of TLP:WHITE since this is open-source information. Click the + sign next to tags, select the TLP Taxonomy Library, then select tlp:white. This adds a white colored tag to the tags section of the event. The color of each tag can be controlled in the JSON configuration file used to set up tag taxonomies.

Now we need to add attributes to the event...

[1] <https://isc.sans.edu/forums/diary/Petya+I+hardly+know+ya+an+ISC+update+on+the+20170627+ransomware+o+utbreak/22566/>

Adding Event Attributes

Populate using the freetext import tool

Category
Type
Value

Freetext Import Tool

Paste a list of IOCs into the field below for automatic detection.

027cc450ef5f8c5f653329641ec1fed91f694e0d229928963b30f6b0d7d3a745
 64b0b58a2c030c77fdb2b537b2fcc4af432bc55ffb36599a31d418c7c69e94b1
 https://isc.sans.edu/forums/diary
 /Petya+I+hardly+know+ya+an+ISC+update+on+the+20170627+ransomware+outbre

Submit

Value	Similar Attributes	Category	Type
027cc450ef5f8c5f653329641ec1fed91f694e0d229928963b30f6b0d7d3a745	149 184 703	Payload installation	sha256
64b0b58a2c030c77fdb2b537b2fcc4af432bc55ffb36599a31d418c7c69e94b1	149	Payload installation	sha256
https://isc.sans.edu/forums/diary/Petya+I+hardly+know+ya+an+ISC+update+on+the+20170627+ransomware+outbreak/22566/		External analysis	url

Adding Event Attributes

One of the best features of MISP is the ease in which you can add multiple pieces of information with different types. Since the article contains multiple pieces of useful info with various data types, we will use a free text import tool to add attributes to the event.

Parsing through the SANS ISC article (referenced again below), there are several hashes associated with NotPetya we would want to enter into this event. Each item of interest can be entered into its own line for automatic type detection and import. Since we also want to save a link back to the blog article, this is entered in the freetext import tool as well (length caused the URL to word wrap). Once the items are entered, we hit the "Submit" button.

The bottom screen shows the results of the type detection as well as the attribute category and type. At this point, we need to adjust categories to make sure they align with what each indicator represents—in this case, the hashes are payload installation hashes and the URL is to the external analysis of the event (not command and control or some other malicious URL). Once everything is correct, we can hit the "Submit attributes" button (not pictured) to make the attributes submission final.

<https://isc.sans.edu/forums/diary/Petya+I+hardly+know+ya+an+ISC+update+on+the+20170627+ransomware+outbreak/22566/>

Attribute Correlation Example

Petya? I hardly know ya! - a

Event ID	1143
Uuid	5bd471e9-de0c-4f2e-bc58-0abdac
Org	SEC450
Owner org	SEC450
Contributors	
Email	admin@admin.test
Tags	tip:white x +
Date	2017-06-27
Threat Level	High
Analysis	Completed
Distribution	All communities
Info	Petya? I hardly know ya! - an ISC
Published	No
#Attributes	3
Last change	2018-10-27 15:25:22

Date	Org	Category	Type	Value	Tags	Galaxies	Comment	Correlate	Related Events
2018-10-27		Payload installation	sha256	027cc450ef5f8c5f653329641ec1fed91f694e0d229928963b30f6b0d7d3a745	+	Add		<input checked="" type="checkbox"/>	149 184 244 703
2018-10-27		Payload installation	sha256	64b0b58a2c030c77fdb2b537b2fcc4af432bc55ffb36599a31d418c7c69e94b1	+	Add		<input checked="" type="checkbox"/>	149
2018-10-27		External analysis	url	https://isc.sans.edu/forums/diary/Petya+I+hardly+know+ya+an+ISC+update+on+the+20170627+ransomw	+	Add		<input checked="" type="checkbox"/>	

SANS
SEC450: Blue Team Fundamentals: Security Operations and Analysis
136

Attribute Correlation Example

This slide shows a basic manually created example event with 2 attributes—the SHA256 hashes of NotPetya samples. The upper left box shows the event metadata, the bottom box shows the attribute list, and the upper right box shows the correlation graph for the event. The attribute list allows us to easily note and view any attributes associated with an event of interest, which, in this case, were the 2 hash values.

This is where the power of a threat intelligence platform comes in—in addition to all locally added event attributes, this MISP instance was set up to pull in the external CIRCL OSINT threat indicator feed. Doing so causes hundreds of MISP events with their related attributes to be automatically downloaded as well. In the correlation graph view, as well as on the attribute view, you can see that manually entered hashes from event #370 were also present in multiple other events from the external feed (visible by the different icon in the correlation graph view). These attributes that show up across multiple events were automatically identified by MISP, and given this information, we can now correlate this event with activity previously entered by a third party. To get additional context, we can now pivot to the other MISP events and see what the context was around the entered hashes from events 184, 149, 703, and 244.

The screenshot displays the MISP Events List interface. At the top, there is a search bar with the email 'admin@admin.test' and tabs for 'My Events' and 'Org Events'. Below this is a table with the following columns: Published, Org, Owner, Org, Id, Clusters, Tags, #Attr, Email, Date, and Info. The table contains five rows of event data. The first row is for event ID 1143, dated 2017-06-27, with tags 'ttp:white'. The second row is for event ID 1142, dated 2017-04-23, with tags 'ttp:white' and 'circl:incident-classification="malware"'. The third row is for event ID 1141, dated 2016-08-10, with tags 'Type:OSINT' and 'ms-caro-malware:malware-type="RemoteAccess"'. The fourth row is for event ID 1140, dated 2016-12-22, with tags 'osint:source-type="blog-post"' and 'ttp:white'. The fifth row is for event ID 1139, dated 2017-10-25, with tags 'Type:OSINT', 'ttp:white', 'malware_classification:malware-category="Ransomware"', and 'osint:source-type="blog-post"'. The interface also includes a SANS logo in the bottom left and the text 'SEC450: Blue Team Fundamentals: Security Operations and Analysis 137' in the bottom right.

Events List

Here is a screenshot of the main events page in MISP. You can see the NotPetya example event that was previously entered, as well as multiple other events that were pulled in pre-classified from external feeds. The "Org" column shows the source of the information, which, in this case, was the CIRLC OSINT feed. The clusters column shows any galaxy classifications that were applied, and the tags show any labels from taxonomies or otherwise that have been attached.

Threat Intelligence Platforms Summary

Indicators and intelligence managed in TIP

- Indicators should be automatically pulled/added through API integration with other tools
- Manual events created based on individual event analysis

Incidents managed in IMS

- All high-fidelity or triage alerts become cases
- Cases added to queue, assigned to analysts
- A tool you will spend LOTS of time with, choose wisely
- Ideally integrates with automation frameworks and TIP

Threat Intelligence Platforms Summary

Through the last 2 sections, we have covered two main pieces used for SOC data organization: The incident management system and threat intelligence platform. These two tools allow SOC's to keep tabs on all the incidents that have occurred and correlate their details together to get a higher-level picture of items that are connected. This capability allows Blue Teams to grow their threat intelligence and, over time, a picture should emerge of the type of threats you face, the tactics and techniques they'll use, and where you should prioritize defensive spending in order to counter the assault.

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. **Exercise 1.2: MISP Threat Intelligence Platform**
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Licensed To: David Owerbach <0mamaloney@gmail.com> April 28, 2020

Exercise 1.2: MISP Threat Intelligence Platform



Exercise 1.2: MISP Threat Intelligence Platform

Exercise 1.2: MISP Threat Intelligence Platform

Please go to Exercise 1.2 in the SEC450 Workbook or virtual wiki.

Course Roadmap

- **Day 1: Blue Team Tools & Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
- 10. SIEM and Automation**
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Licensed To: David Owerbach <0mamaloney@gmail.com> April 28, 2020

SIEM and Automation

In this module:

- The role of SIEM in detection
- Data flow into the SIEM
- What is a use case?
- The role of SOAR in response
- How SOAR improves the Blue Team

SIEM and Automation

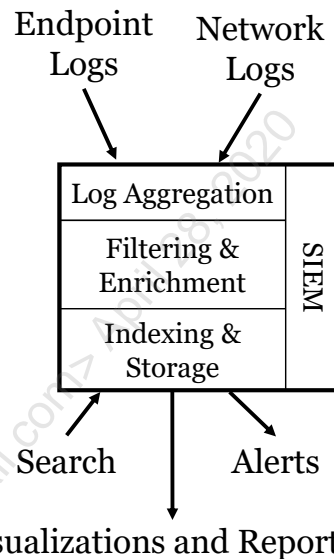
In this module, we review two more crucial tools for security operations. One is the SIEM, which is your central repository for all logs and primary method for searching, alerting, and reporting. The second is automation or, if you have a specific tool for it, SOAR—Security Orchestration, Automation, and Response frameworks. The SIEM can be a complicated seeming appliance, so the goal will be to simplify its operation down to the core functions and look at it with a systems level view. We'll also discuss SIEM use cases and how automation can assist us with their execution, as well as many other common operational tasks.

A SIEMs Job

SIEM duties:

- **Receive** all log data
- **Parse** it correctly
- **Filter** unwanted events
- **Enrich** useful events with additional data
- **Index** log into database
- Fast **searching**
- **Visualization** and **dashboard** creation
- Analytics and correlation for **alerting**

One of the *most* important SOC tools!



A SIEMs Job

At the highest level, the whole point of a SIEM is to be a centralized collection of logs that we can use in various ways to highlight suspicious conditions and detect attacks and/or compliance issues. To accomplish this, a SIEM must first have all logs in the environment forwarded to it, then parse and store them. Assuming all of this has been orchestrated correctly, we then also want the ability to filter out items of interest with a flexible search language, visualize log data, and enrich fields within our logs to give them more context. Additionally, the SIEM should include an alerting engine that will fire when logs of interest come in, and also give the SOC the ability to create periodic reports that can be distributed to stakeholders.

This slide shows the systems level view of a SIEM. Ultimately, the input to the system is logs (and potentially other information like NetFlow if the SIEM supports it), and the output of the SIEM is visualizations, alerts, reports and search results. For a SIEM, the biggest pieces of the puzzle are the inner processes that are required to parse, enrich, filter, and manage the immense number of logs that are generally collected by these devices.

SIEM Features

Features to look for:

- High-performance logging agent for data collection
- **Multi-format log compatibility**
- Message buffering for resiliency
- Methods to **easily parse data**
- High performance ingestion and indexing
- **Multiple types of log enrichment** and correlation capability
- **Fast search** with easy query language
- Multiple **visualization types**
- **Well thought out UI**
- **Flexible and expressive alerting** options
- API for third-party tool integration
- Good documentation and vendor support
- Frequent updates and modular "app" system

SIEM Features

Although there are many SIEM options available on the market, some are much more of a pleasure to use than others. Here are some of the key items you should look for in a good SIEM. The highlighted entries are what many people find to be the most crucial for long-term happiness with a SIEM product. Ultimately, what it comes down to is that you should be able to receive logs into your SIEM using a wide range of formats, easily define a parsing algorithm, and be able to search, visualize and alert on the contents of those logs without wanting to pull your hair out. The best systems make it painless to flexibly search through log content at blazing fast speeds and support many visualizations types and alerting conditions that make identifying evil quick and painless.

SIEM Products

splunk>

QRadar®

elasticsearch

exabeam

ArcSight™

LogRhythm™

graylog

SIEMPLIFY

McAfee
Enterprise
Security Manager

RSA NETWITNESS

EventTracker
Actionable Security Intelligence

Trustwave®

ALIEN VAULT OSSIM

RAPID™
insightIDR

LOGPOINT

solarwinds

SANS

SEC450: Blue Team Fundamentals: Security Operations and Analysis

145

SIEM Products


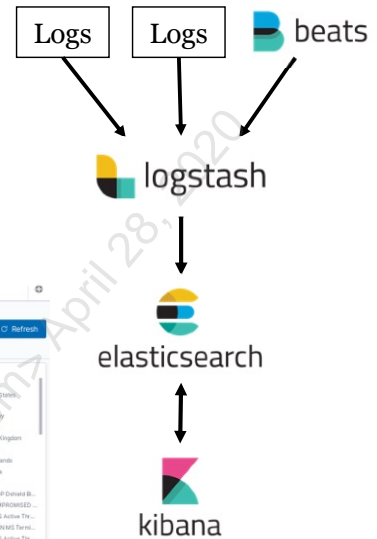
Here are just some of the options you might see in the SIEM market. The most commonly seen large vendors are on the left side with solutions like Splunk, QRadar, LogRhythm, ArcSight, and more, which have been around for many years. When shopping, be aware that there are a lot of lesser-known alternative solutions and niche players as well that are aimed at Managed Security Provider small to mid-size business markets.

When it comes to choosing a solution, be sure to get actual hands-on use with a product to assess it for fit to your situation. Perform log ingestion, agent setup, search and visualization operation, example alert creation and other capability checks. While features can look the same on paper, the actual implementation can be drastically different and, as the saying goes, "the devil is in the details." As always, take vendors' promises with a grain of salt. Most vendors will not volunteer their products' weak points, so make sure you ask them to compare themselves to others and ask what their current customers complain about or wish they could do better.

Lab SIEM

In your class VM: The **Elastic Stack**

- Logs ingested through **Logstash**
- Stored in **Elasticsearch** database
- Searched and visualized in **Kibana**



Lab SIEM

For this class, the software we will use to act as a SIEM is the Elastic Stack. We do not need to go into the details of how it works (SANS offers SEC455 for those who are using the Elastic Stack), but the basic idea is that logs are first collected from various log sources with syslog or software agents (Beats) and are sent to a log aggregator – Logstash. Logstash takes the logs and parses, filters, and enriches them with additional information before sending them on to Elasticsearch to be stored. Elasticsearch acts like a large database, but instead of the typical columns and tables format, Elasticsearch stores every log as a JSON formatted "document" in what is called an "index" (similar to a table in traditional databases).

Kibana is the frontend web-based search interface we will use to query the logs stored in Elasticsearch, which have been pre-ingested for the class virtual machine. You will not need to configure or interact with Beats, Logstash, or Elasticsearch directly for this class, only Kibana. This piece is only mentioned to give context to the architecture behind Kibana, which will be used for labs, and explain the Elastic Stack and how its pieces are similar to that of any SIEM.

The screenshot displays the Kibana Discover interface. At the top, the title 'Searching with Kibana: The Discover Tab' is visible. Below it, the search results show '774,140 hits'. The interface includes a 'Time picker' at the top right, a 'Histogram' showing log activity over time, and a 'Field list' on the left side. The 'Document data' section at the bottom shows individual log entries with various fields like host, geoip, and event_type. Annotations with arrows point to the 'Index' dropdown, the 'Histogram', the 'Field list', and the 'Document data' section.

Searching with Kibana: The Discover Tab

To perform a search in the Discover tab, select a time range with the time picker in the upper right corner to limit the scope of the logs. Next, an index must be chosen to determine which set of logs will be searched. Indexes in Kibana typically exist for each type of data source (Windows logs, Firewall logs, AV logs) and each will have a name that can be selected in the dropdown. If you want to search *all* logs, select the * pattern.

Once the index pattern is selected, a list of all available fields will appear in the field list on the left side of the page. If you click any of these fields, the top 5 values will be shown, as well as an “add” button to specifically show that field as a column. If no specific fields are added as a column, the entire document will be shown (as seen on the slide) with the field names highlighted in bold. Documents can be expanded using the small triangle on the left, which will display all fields in the individual log. If specific fields are added as columns, only the info from those selected fields will be displayed in the data section.

The search bar is at the top of the page and can be used to refine the displayed results. You can run searches for terms in any field or within a specified field, including using wildcards or regular expressions. Once a search is run, the documents that match the time frame and selected index pattern will immediately begin to show up in the histogram. If you want to refine the time range of the search visually, the histogram can be clicked and highlighted to zoom to a specific time of interest.

Kibana Query Language Examples

Open search for "string"

- `string`

response field containing "string"

- `response:string`

destination_port field above 1024

- `destination_port > 1024`

Searching for multiple matches

- `response:string and destination_port:80`

Searching for one of two things

- `response:string or destination_port:80`
- `response:(200 or 404)`

Kibana Query Language Examples

This page demonstrates some of the more common methods to search for data within Kibana. It is not an exhaustive list but covers most of the situations needed for this class.

SIEM Use Cases

What is a SIEM **use case**?

- Answers "**What is the SIEM doing for us?**"
- Some type of output – report, alert, dashboard, etc.
- Documented actionable item produced by the SIEM
 - Brute force login attempted
 - Suspicious volume of upload traffic from user
 - Potentially malicious download
 - User added to administrator group
 - Credentials submitted to phishing site



SIEM Use Cases

The way the Blue Team defines what they want the SIEM (and other tools) to do is the "use case." All Blue Teams have analytics like "if a failed login occurs 10+ times, send us an alert", but this rule itself is not the full use case, only a part of it. A use case is a more formal way of documenting we want our SIEM to identify, and it includes not just the condition itself, but the reasoning behind it, the expected response, traceability to a business or compliance requirement, and more.

You may have heard of common use cases referred to by what they are identifying—brute force logins, malicious downloads, etc.—but the actual documentation for how this is done, and what process should follow, is also included. While there is no hard and fast rule for what should be documented, some best practices have emerged that we will cover.

To Catch a Penetration Tester: Top SIEM Use Cases: <https://www.youtube.com/watch?v=9Ndv0W2Uq0U>

Use Case Development

Use case documentation fields:

- Name
- Description
- Problem Statement
- Goals
- Requirements
- Primary Data Source
 - Secondary Data Sources
- Analytic Logic
- References
- Suggested Analysis Steps
- False Positive Reduction Steps
- Categories and framework
 - MITRE ATT&CK / Kill Chain
 - VERIS
- Compliance / audit support
- Threat group / attribution

Use Case Development

Here are some of the more commonly tracked items in a set of well written use cases. Having the information listed above for each rule helps in day-to-day operations in several ways. For one, and probably most importantly, it helps analysts who are not familiar with a particular alert know what should be done when the rule fires, why the rule was put there in the first place, and how to do analysis and false positive reduction. Without this guidance, analysts both new and experienced would be left to guess the intention of a given rule if they weren't already familiar with it, leading to inconsistent responses.

Another extremely important use for this information is tracking coverage across your set of tools and data sources. With use cases well documented, you can answer questions like "which tool implements most of our detection capability?" and "do we have the proper coverage across the attack kill-chain or MITRE ATT&CK techniques?" (Frameworks that we will discuss in detail later.)

Creating a New Use Case

1. What is the situation we're trying to identify?
2. Define the conditions to detect it
3. Define the data sources that can identify that condition
4. Write logic, analytic, and expected outputs
5. TEST IT to ensure it functions as expected
 - Ideally, automate periodic testing
6. Document details in use case database

Remember: "**Offense informs defense**"

Creating a New Use Case

When crafting a new use case, first step back and consider what the situation is you're trying to detect and specify in as much detail as possible the sources, data fields, and values that will uniquely identify it. These will make up the heart of the use case and drive how the analytic might create false positives. Afterwards, be sure to fully document the reason the use case exists, what a use is supposed to do when it triggers, and any common conditions that can cause a false positive.

Once a use case is fully documented, implement the rule in your SIEM or other security tool and be sure to test it. The only thing worse than not having an analytic, is having a one that you think will work correctly, but in fact doesn't. That doesn't mean one test is enough for all of eternity, either. If possible, the team should use a script or other solution to assist with periodically testing rules to ensure they are still working. Purple teams also work very well for this type of verification. There's no telling when a vendor may change a log and break a detection that once worked perfectly, so repeated testing should be considered a necessity. Once it is been verified and all the details are written down, it's ready to put into the use case database.

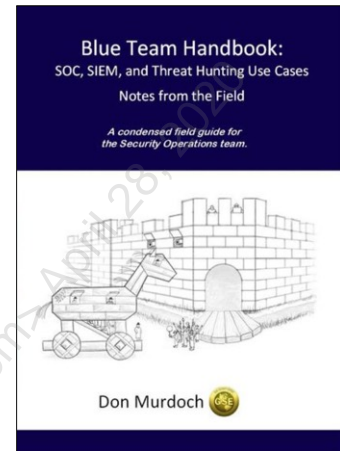
Use Case Databases

All use case info should be tracked closely

- Options: Ticketing systems, Excel, wikis, text files
- **Track alignment**
 - To business requirements and compliance
 - To MITRE ATT&CK or other frameworks
- **Changes** to analytics over time

Helps analysts understand:

- How each detection works and theory behind it
- **How to interpret** data and respond
- **Data sources** involved



Use Case Databases

How do we keep track of all the use cases we have generated over time and which items do we need to include inside them? The answer is a use case database. There are many ways to implement a use case database: Through a ticketing system, Excel spreadsheets, Wiki software, text files, databases, etc. For storage, the choice is yours to pick whatever best lines up with the tools your team already uses. What should be standardized across most Blue Teams, however, is the actual data that's tracked for each individual use case. A well thought out, customizable software solution will allow you to organize your use cases in the way that works best for your team's workflow, and ideally will also allow tracking of changes, and metrics on use cases over time.

One of the best and most thorough sources of information on use cases comes from Don Murdoch, Assistant Director of the Institute for Cyber Security at Regent University, graduate of the SANS Technology Institute program, and GSE #99. His newest book—*Blue Team Handbook: SOC, SIEM, and Threat Hunting Use Cases*—is absolutely filled with low-level detail on use case creation suggestions and best practices.¹ Don has some of the most thoroughly documented use cases out there, and speaks from years of experience running Blue Teams, so if you're new to use case organization and creation, taking his advice can be a great start to getting on the right path.² An alternative, simple Excel-based approach to use case tracking is found in Ryan Voloch's talk, "Simplified SIEM Use Case Development" from DerbyCon 2015.³

[1] Don Murdoch – Blue Team Handbook Volume 2: <https://www.amazon.com/Blue-Team-Handbookcondensed-Operations/dp/1726273989/>

[2] Don Murdoch / SANS Security Operations Summit, 2018 – SecOps, SIEM, and Security Architecture Use Case Development: <https://www.sans.org/summit-archives/file/summit-archive-1533050405.pdf>

[3] Ryan Voloch – Simplified SIEM Use Case Development: <https://www.youtube.com/watch?v=JvYIOPvMyeA>

Automation and Orchestration Definition

- **Automation** accomplishes a specific task
- **Orchestration** chains together automated tasks into workflow
 - Think "running your playbook for you!"
- **Benefits:**
 - **Standardization** of response tasks (implements playbooks)
 - **Immediate response** time
 - **Higher capacity** to address alerts – reduces fatigue
 - **Faster onboarding** for new employees
 - **Focused effort** on things that matter
 - **Happier analysts** that don't have to do repetitive work



Automation and Orchestration Definition

Automation as an enabling technology has caught on in a big way over the last few years and, at this point, having some level of automation is practically a necessity. The amount of data we collect for putting up a modern defense almost requires it.

One could argue whether automation and orchestration tools are truly a new category or just a new spin on an old one, but the point is we now have a wide and mature selection of tools meant specifically to automate tasks in the security space. What, though, is the difference between automation and orchestration? When we talk about **automation**, we are referencing specific tasks that were performed manually now being passed to a script or software doing the work for us.

Orchestration, however, is one level of abstraction higher. Orchestration is taking the tasks that have been automated and deploying them in a series of events defined by a workflow complimented with conditional logic. In a way, these are very similar to playbooks in the IMS. In fact, SOAR is the perfect companion to a playbook because it can take a lot of the steps that would've previously taken analysts time and get them done preemptively or automatically as human analysis is performed, saving lots of time!

SOAR Platforms

Products:

- Phantom (Splunk)
- Demisto
- DFLabs
- NetWitness Orchestrator
- Komand
- Siemplify
- Swimlane
- NSA WALKOFF

Generic flow-based programming:

- Node-RED (Free)
- Total.js Flow (Free)

Paired with TheHive:  Cortex



SOAR Platforms

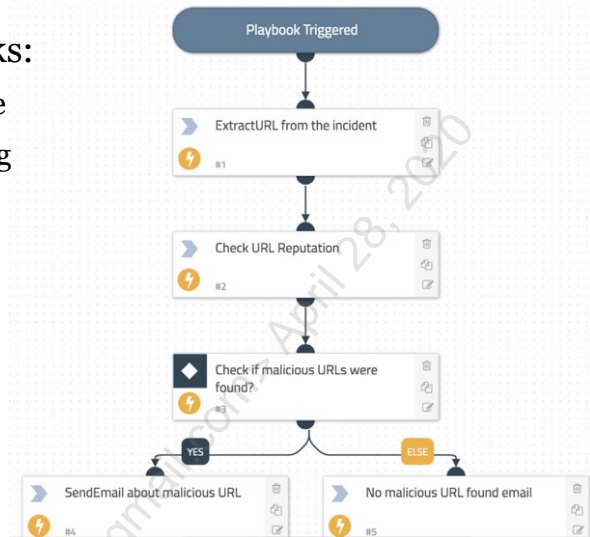
With the demand for automation on the rise, the SOAR product category has appeared to satisfy this need. SOAR stands for Security Orchestration, Automation, and Response. While many new companies spring up around this category, existing product teams have also hustled to add SOAR capability to their offering. Some well-established products have even white-labeled tools from third-party SOAR companies, or outright purchased SOAR vendors for a fast win. Without a doubt, the scramble to offer automation as a feature is on and SOAR is here to stay. It's not hard to see why. SOAR takes all the painful, non-value-added parts of the job people don't like doing and makes it fast and easy! Quite simply, SOAR can directly make your analysts happier!

The unfortunate fact about SOAR platforms is that since they are so new, there aren't many free options yet. The best we have are generic flow-based programming languages like IBM's Node-RED, or Total.js Flow. There is also a new free tool called WALKOFF released by the NSA. We also have Cortex, which is free and built by TheHive's developers as a way of starting to add free automation capabilities to it. Although Cortex is not quite the same as the offerings on this slide, we can use it in many of the same ways to demonstrate the benefits of automation, and we will do so in labs in this class.

SOAR Value-Adds

Automation of initial investigation tasks:

- Spam: Check logs for extent of email wave
- Web-Exploit: Automated domain blocking
- Command and Control: Enrichment of domain with data from Virus Total
- Virus detection: Isolate from network
- Phishing: Force user password expiration
- Enrichment: Look up passive DNS, IP address, Whois, or GeoIP info
- Remediation: Craft and send rebuild request to help desk



SOAR Value-Adds

There are several scenarios where a SOAR platform can help. Consider the average playbook and the steps that are required for analysis. How many of them truly require a human and how many can be fully performed by automation utilizing an API? At the most basic, all the items on this list can be accomplished by taking some value out of an alert like an email sender, domain name, or computer hostname, and contacting one of your various tools with that information to tell it to take action. In most SOAR platforms, these actions are *orchestrated* through a defined workflow (playbook) drawn in a GUI as seen above. Each step represents a task that must be completed and can contain logical checks on the output to drive what happens next.

Cortex

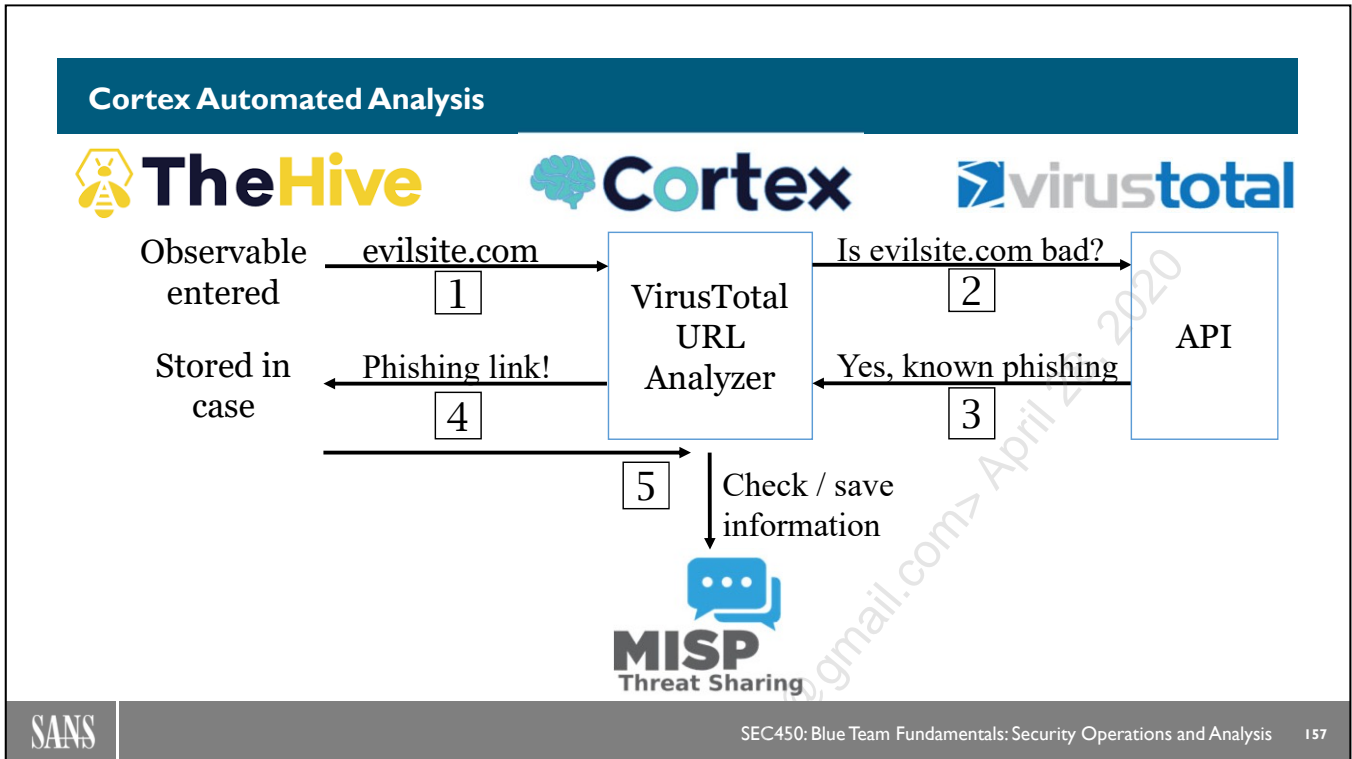
Automated analysis and enrichment engine

- Comes with TheHive and is tightly integrated
- Has its own setup and users
- Uses API key to link to TheHive
- **Analyzers** reduce repetitive enrichment actions for indicators
 - VirusTotal, Passive DNS lookups, etc.
- **Responders** provide automated response actions
 - New functionality more directly aimed at SOAR



Cortex

Cortex is a separate program bundled with TheHive that is meant to work closely with it to perform automated lookups and other enrichment actions. It does not have the GUI workflow design style of other platforms and, at this point, is more focused on the "automation" piece. However, it is being actively developed and I wouldn't be surprised to see it move in the Orchestration direction next. Currently, Cortex contains numerous automatable single actions for data enrichment called "analyzers" and just added a feature scripting automated action in your environment through what they call "responders."

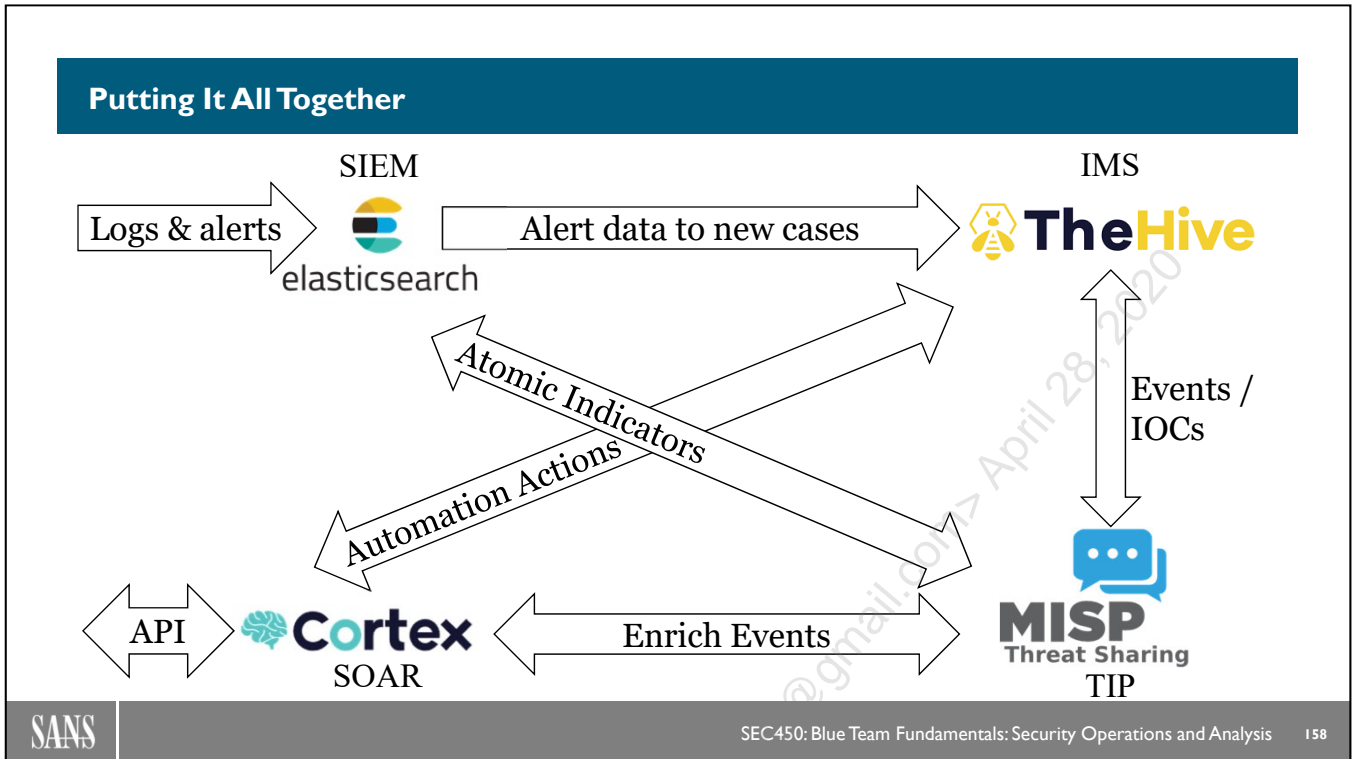


Cortex Automated Analysis

Here's an example of how data can flow when a new "observable" is entered into Cortex.

- First, the domain 'evilsite.com' is entered into TheHive as a new observable, and the analyst is asked if they want to run any analyzers on it.
- The analyst can select a VirusTotal lookup, which will pass the domain name to Cortex and will, in turn, use the VirusTotal API to programmatically get the answer.
- The results are then pulled back into the case in TheHive for context.

What this all means is that as soon as the analyst enters the domain name and tells the analyzers to run, they immediately and automatically have received information about whether a file is malicious or not and were also able to store that information in the threat intelligence platform. This is the magic of automation! Without Cortex, the analyst would have had to manually go to VirusTotal, copy and paste the information to do the look, type the returned results, and manually generate an event in MISP to track the indicators. Automation via Cortex has eliminated all that work and made the analyst much faster at responding!



Putting It All Together

Through the last few sections, we have talked about numerous tools—the IMS, TIP, SIEM, and SOAR. Now that we understand what each tool does, this slide shows the 50,000 ft. view of how each item interacts to form the complete picture. Note that there is more than one valid way to connect all these systems. The diagram above shows one option but how you do this will depend on which tools you use in your SOC, which of them have an API that makes integration easy, and the order of your data flow. At this level, it merely is important to understand what each tool is doing, and the types of information that should be passed between the tools in some method to enable them to do their job in an automated way.

Knowledge and Code Database

For everything else – need general purpose doc storage

- **Wiki's, Notes, Document programs, code**
 - OneNote, SharePoint+Word, Git, DocuWiki, HackMD...
- **Want:**
 - Real-time collaborative editing
 - Easy search
 - Version control and comments
 - Automatic syncing
 - Easy learning curve
 - Rich text and pic support
 - Alerting for doc changes

Knowledge and Code Database

The last SOC information management tool that hasn't yet been discussed is the SOC knowledge database. SOCs need a repository for general text-based information for things like onboarding materials, and training info, or presentations that are developed by the group. They also need a place for more structured items like code that should be version controlled across time. There are many tools to meet both needs. Most SOCs settled on some combination of online collaborative note taking tools such as OneNote or Wikis and use a tool like Git or SharePoint for custom code and other files that need close version control.

The most important feature in knowledge and code databases is that they are simple enough to use that the team will *want* to interact with them. If the solution is even slightly burdensome to use, employees may be deterred from entering information or code as frequently as they should, defeating the entire purpose. Markdown driven wikis like HackMD can strike a nice balance between formatting and usability, although OneNote is the simplest solution commonly available. Many teams like OneNote because it can be used to take collaborative notes during meetings that will instantly sync to all in attendance. It is also easy to manage over time with new pages and notebooks being created for each meeting or period and no infrastructure to worry about.

SIEM and Automation Summary

Logging Info **Searched, Visualized** and **Alerted** on in the **SIEM**

- Must understand what data is available, how to read it

Threat Intelligence Platform is used to hold threat info

- SIEM and IMS query/add information to lists

Alerts become cases in **IMS**

- Cases assigned to analysts, follow playbooks for analysis

SOAR facilitates automation, decreasing response time

- Immediate context gathering, data enrichment

SIEM and Automation Summary

This section has covered numerous tools and how they all work together to make the SOC process. Throughout the class, we will utilize these tools to perform investigations and hopefully gain some inspiration on how these capabilities can make our own process better. If you have these tools but they aren't connected and interacting in an optimal way, consider how the data that flows between them can be improved upon to make your response time faster, and day-to-day life better!

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. **Know Your Enemy**
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Licensed To: David Owerbach <0mamaloney@gmail.com> April 28, 2020

In This Module

To put up a strong defense, we must understand our enemy!

“If you know the enemy and know yourself, you need not fear the result of a hundred battles.

If you know yourself but not the enemy, for every victory gained you will also suffer a defeat.

If you know neither the enemy nor yourself, you will succumb in every battle.”

— Sun Tzu, The Art of War



In This Module

Yes, you may have heard this quote numerous times, but there are good reasons for that—it's absolutely true. Sun Tzu's eternal advice has direct application to cyber defense in that, as we discussed with threat intelligence, we must understand our attacker if we wish to stop them.

This module will dive further into the who, why, and the styles of attack we see from the various opponents the SOC will face over time.

Who's Attacking Us?

Government-backed groups¹

- Attribution lists contain: Russia, China, USA, UK, Iran, North Korea, Vietnam, France, India, Israel, Syria and more



Organized Crime

- Anunak²
- Business Club³
- Carbanak⁴
- The Dark Overlord⁶



Who's Attacking Us?

There are several different groups behind the tidal wave of cyber attacks. Two of the main perpetrators we will talk about first are government-backed groups and organized crime.

Although not the biggest attacker by volume, nations represent the most capable, stealthy, and potentially the most damaging attacker. The list of countries accused of attacks is constantly growing, and their new tools and capabilities constantly impress. These groups are known for using zero-days, rootkits, and other high-complexity attack styles that often break new ground in what defenders thought was possible.¹ A watershed moment in the attribution of attacks to countries was the APT1 report in 2012; and since then, vendors have continued to regularly publish details about the incidents they see, and how they were able to tie the actions back to the attack groups from these countries.

One of the most prolific groups of attackers are those backed by organized crime. Groups such as The Dark Overlord have made hundreds of thousands to millions of dollars over the years through various fraud and extortion techniques. In The Dark Overlord's case, for example, the group was gaining access to sensitive data at healthcare organizations using crimeware-as-a-service offerings.⁵ Following this, they would send a ransom note demanding payment of thousands of dollars in bitcoin for return of the information instead of releasing it on the dark web. On top of it, they made sure to worsen the situation by reaching out to journalists offering to give out the scoop on which companies had been breached, forcing companies to either pay or be left with a PR and IT nightmare.

[1] <https://attack.mitre.org/wiki/Groups>

[2] <https://krebsonsecurity.com/2014/12/gang-hacked-atms-from-inside-banks/>

[3] <https://krebsonsecurity.com/2015/08/inside-the-100m-business-club-crime-gang/>

[4] <https://krebsonsecurity.com/2016/07/carbanak-gang-tied-to-russian-security-firm/>

[5] <https://www.cybereason.com/blog/carbanak-cybercrime-gang-arrest>

[6] https://motherboard.vice.com/en_us/article/mbkex8/dark-overlord-arrest-serbia-netflix-hackers

Hacktivists

Hacktivist Groups:

- Anonymous
- Lulzsec
- Lizard Squad
- Syrian Electronic Army
- Chaos Computer Club
- Level Seven Crew
- globalHell



Hacktivists

On the other side of the coin, we have much more poorly organized groups or even individuals that fall under the umbrella of "**hacktivists**"—people out to make an ideological or political point through hacking. There's also **hackers for hire**—vigilantes willing to do pretty much anything requested of them for the right amount of money, and finally, **malicious insiders**. These three groups represent much different threats than the nation-states and organized crime groups previously discussed.

One of the trends we've seen throughout the last few years is "hacktivism." Anonymous and Lulzsec are some of the most noticeable groups in this category. If you've watched the USA series "Mr. Robot," this is where the "FSociety" group falls. These groups are known for broad-ranging hacks from the Anonymous-led DoS attack on the Church of Scientology in "Project Chanology"¹ to the HBGary hack by Lulzsec in 2011², to more recent attacks like Turkish activists trying to take over social media accounts of U.S. journalists.³ Tracking these groups is possible, but difficult. Those who attempt to find threat intelligence on hacktivists often must resort to infiltrating their online dark-web meeting places or forums and pretending to be one of them. If you're lucky, their plans may even be publicized on social media. For this reason, it is a good idea to have someone watching these arenas if you expect that your organization will have ideological attackers.

[1] https://en.wikipedia.org/wiki/Project_Chanology

[2] <https://arstechnica.com/tech-policy/2012/03/the-hbgary-saga-nears-its-end/>

[3] <https://www.cnn.com/2018/08/24/turkish-hacktivists-make-cyberattacks-on-us-journalists.html>

Who's Attacking Us? Hackers-For-Hire and Insiders


Hackers-for-Hire

- Outsourcing hacking to others
- Ransomware-as-a-service
- Exploit-as-a-service



Malicious Insiders

- Making money
- Stealing intellectual property
- Hurting organization reputation



Who's Attacking Us? Hackers-For-Hire and Insiders

Hackers-for-hire are not necessarily a well-defined group, but more of a catchall term for the individual actors, bot herders, and other attackers out there working in more of a "business-to-business for hackers" type fashion. They may run a business by sending out major phishing campaigns and collecting access to victim computers at as many organizations as possible to sell to the highest bidder. There are also groups that create "ransomware-as-a-service" offerings.¹ These groups create full turn-key solutions for running your own ransomware campaign and will sell you the software to do so.

Finally, one of the biggest wild cards: Malicious insiders. Employees who are disgruntled or otherwise find themselves at odds with the company may take steps to release confidential information, steal money, or cause disruption upon their departure. One example of this is the insider who stole \$1.8B (yes billion) from the Punjab Nation Bank in India by using access to the SWIFT (interbank transfer system) password, an exceptionally damaging attack!² These are often referred to as "black swan" events, as they are so infrequent and hard to predict. In the case of users, we do have software that can help profile employee activity and alert the SOC to anything out of the ordinary. These tools—mainly data loss prevention software (DLP) and user and entity behavioral analysis (UBA/UEBA)—can be fraught with false positives until they are tuned correctly for the environment, as many of their alerts are anomaly-based. Depending on where you live, they may also have legal implications that need to be reviewed with lawyers. Caution is advised when dipping into this space.

[1] <https://nakedsecurity.sophos.com/2017/12/13/5-ransomware-as-a-service-raas-kits-sophoslabs-investigates/>

[2] <https://www.bankinfosecurity.com/mitigating-insider-threat-lessons-from-indian-fraud-case-a-10674>

Cooperation of Groups and Unexpected Motivations

Nation-states, organized crime, etc. using hackers for hire

- Makes it harder to tell who is who
- Great OPSEC, plausible deniability...attribution more difficult
- Makes threat intel complicated
- Examples: Ransomware-as-a-service, hackers selling established access

Many attacks don't end up what they seem to be on the surface...

- North Korean indicted for Sony hack ¹
- HBO hack tied to alleged Iranian government-related hacker / group ²
- Anthem healthcare breach blamed on China ³

Cooperation of Groups and Unexpected Motivations

The tricky part about understanding who's attacking you is that it might also change over the course of a single attack! Hackers for hire may break into your organization, then take the access to a market on the dark web and sell it to the highest bidder. In this case, anyone could end up with it—organized crime, governments, or hacktivists. The bad part: Trying to pin down who was behind it will be difficult due to the blended nature of the attack TTPs. Once the access is handed off, you may see the new owners use a whole host of different tools that are incongruent with what was seen in the early parts of the kill chain, making the story very complex to put together. In a way, it's a brilliant way for many groups to operate. It affords them additional operational security and takes the risk off them getting caught during the initial delivery and exploitation phases. Plus, it confuses the Blue Team!

Groups leaning on each other for access is bad enough, but it gets worse. There have been several attacks that although you would think you could pin them on one type of group at first glance, end up being something entirely different altogether. When you first heard about the Sony hack, did you first think "North Korea", or did you think it was possible that it was a rogue insider or other organized crime group that may not have received the ransom they were looking for? In that case, many explanations were plausible, but it turns out, at least according to the U.S. Government indictment that went out on September 6, 2018, it was indeed North Korean operatives behind the attack. The point here is that obviously not all attacks turn out to be motivated by what you might expect. It's hard to know the specific motivations of attackers but knowing the type of groups you expect to attack you and aligning against their likely TTPs is a good start.


[1] <https://www.nytimes.com/2018/09/06/us/politics/north-korea-sony-hack-wannacry-indictment.html>

[2] <https://www.wired.com/story/feds-indict-iranian-for-hbo-hack-good-luck-arresting-him/>

[3] <https://krebsonsecurity.com/2015/02/china-to-blame-in-anthem-hack/>

Know Your Enemy: Opportunistic Attackers

- **Adversary Goal:** High-volume, untargeted compromise
- **Strategy:** Collect infected computers, control as botnet, use them for financial gain
- **Tactics:**
 - Malicious and scam email
 - Web drive-by downloads
 - Web scanning and exploitation
 - Browser-based social engineering
 - FakeAV / popups, etc.



Know Your Enemy: Opportunistic Attackers

Unlike the specific groups and motivations we discussed earlier, we can break attackers down into two more generalized camps—**opportunistic** and **targeted attackers**. Opportunistic attackers are groups like organized crime (in some cases) and hackers-for-hire. They have no specific interest in who you are as long as your resources can ultimately help them with their mission. They indiscriminately send out large phishing waves, throw exploit kits in advertisements, and will use browser-based social engineering attacks to gather as many compromised individuals as possible.

The uses for a large group of infected computers are many. As discussed before, certain unique infected users may at first be able to be sold for high amounts of money if they reside in a hard-to-infiltrate company. Beyond that, bots can be used to host or redirect other victims through a "fast-flux" DNS attack infrastructure, further assisting the attackers in gaining more bots in their army. In addition, these large groups of infected machines can be used to send more spam, run DDoS attacks, and more.

Opportunistic Attacker Motivation

What opportunistic attackers *ultimately* want? **MONEY!**

How?

- Ransomware
- Webcam Extortion
- Banking Trojans
- Cryptocurrency
- Sending spam
- Virtual Goods
- Click Fraud
- Phishing Sites
- Social Media Spamming
- DDoS Zombie



Opportunistic Attacker Motivation


What is it that most opportunistic attackers are after? Money! The paths to get there are numerous, and some are more direct than others. But rest assured, in many cases, cash is the ultimate goal of these types of attackers. Some of the methods of turning one of the many hacked into a revenue stream are listed on the slide and even more are named in Brian Krebs' article on the scrap value of a hacked PC.¹ While moves like ransomware, cryptocurrency theft, or direct extortion can be used to gather money directly from the victim, other methods such as DDoS, spamming or phishing sites hosting can bolster a bot herder's capabilities to generate money from others. DDoS capacity, for example, is often used for "protection money" schemes. Owners of millions of PCs in a botnet can email sites such as casinos or ticket brokers and demand to be paid lest they unleash a torrent of traffic that will ruin their ability to sell time-sensitive bets or tickets to their customers. Rest assured there is almost no limit to the creative ways one can turn a fleet of hacked computers into a way to generate money.

[1] The scrap value of a hacked PC:

<https://krebsonsecurity.com/2012/10/the-scrap-value-of-a-hacked-pc-revisited/>

Opportunistic Attack Case Study: Mirai

- Started as a Minecraft booter...
- Written by 3 college-aged people in the US
 - One owned DDoS mitigation company
- Source code publicly released to avoid attribution
- Evolved to millions of compromised IoT devices
- Used to launch (then) record-breaking 660Gbps DDoS!
- Also used to attack DNS services at Dyn
 - Downed major internet services for the East coast and more (pictured)



Opportunistic Attack Case Study: Mirai

The Mirai botnet was one case of an opportunistic attack that we saw play out on the world's stage not long ago. This malware was used by its botmasters to gain control of millions of smart home devices with weak security such as DVRs, routers, IP cameras, and other smaller IoT devices (and still is). Although each individual device does not have much power, the enormous number of infected devices were able to source an unbelievable amount of traffic in aggregate, making the botnet extremely powerful. Once the botnet had gained strength, it was used to unleash a record-breaking DDoS attack against Brian Krebs¹, generating a 660Gb/sec flood of packets that took his site offline for an extended period.

What was the inspiration for creating such a botnet? The money, of course. In December 2017, an owner of a DDoS mitigation service (clearly making money on both ends) from New Jersey and two others were arrested by the FBI and pled guilty for their role in authoring the malware. From the investigation, the details emerged unsurprisingly, that the criminals were using it to sell DDoS capabilities out to other cybercriminals and perform click fraud.² In a surprising twist ending, the defendants were not given jail time, but a large fine, five years of probation, and a 2,500-hour community service mandate. One of the ways they will be helping the community is by working with the FBI on fighting cybercrime and other cybersecurity matters!³


[1] <https://krebsonsecurity.com/2018/09/mirai-botnet-authors-avoid-jail-time/>

[2] <https://www.justice.gov/usao-nj/pr/justice-department-announces-charges-and-guilty-pleas-three-computer-crime-cases>

[3] <https://www.wired.com/story/mirai-botnet-creators-fbi-sentencing/>

Know Your Enemy: Targeted Attackers (APTs)

- **Goal:** Target people/organizations, get specific info, access specific systems, or disrupt YOUR business
- **Strategy:** Attacks tailored to you and your organization
- **Tactics:**
 - Spear-phishing
 - Watering holes (strategic web compromise)
 - People/physical attack (social-engineering, USB)
 - Targeted service-side exploitation and zero-days



Know Your Enemy: Targeted Attackers (APTs)

Targeted attacks represent the more dangerous type of motivation for an attacker to have. These individuals and groups have decided that YOU specifically have something they are interested in. That fact alone means they are much more likely to be determined to compromise you than an opportunistic attacker that might move on after a rejection. Because they are very determined, they are also much more likely to tailor their attacks to your employees and your network infrastructure and software. This means they will likely pull out tactics like well-crafted spear-phishing, social engineering, targeted exploits customized to your exact software, and potentially even watering holes or supply chain compromise.

If you have never seen a targeted attack before, go back to the incidents you have dealt with and consider, were any of them suspiciously well put together? Did anything infer that more than a trivial amount of time was spent crafting the attack and personalizing it to your company? These are some of the clues that can potentially point to an incident heading in this direction. As we will discuss later, it's very important to get a sense of what could be a targeted attack and tailor your response accordingly.

Attack Group Naming and Convention

Why don't they line up?

Factors according to Florian Roth¹:

1 Human

- Malware or operation name put forth as group name
- Vendors miss links across campaigns

2 Technical

- Each vendor sees different details
- Threat actors join/split up
- Toolsets shared

3 Operational

- Vendors hesitant to names from other org / admit research is better

You've heard the names...

- Emissary Panda
- Lazarus Group
- Comment Crew
- Dark Hotel
- Fancy Bear
- Equation Group
- Sandworm
- Crouching Yeti
- APT [#]
- FIN [#]

"What's in a name?"

That which we call APT

*By any other word
would be as persistent"*

- Cyber William Shakespeare



Attack Group Naming and Convention

How are these attack groups named? Unfortunately, the answer is "inconsistently." Each vendor might use different names for the same threat actor, mislabel them based on an operation name, use the name of malware as the name, or fail to see a link to previous incidents, all of which leads to this problem. There are a whole host of issues that could lead to confusion about who is who. Researcher Florian Roth has an even more in-depth explanation in his blog on why this is.¹ Florian breaks the issues down that cause this problem into "human, technical, and operational" factors, and explains why it does actually make sense for vendors to do this, even though it leads to misunderstanding among analysts.

The good news is, in the same article, Florian has made an awesome attempt to clear up the problem.² In a public Google Spreadsheet, he tracks all the known threat actors by country, lists all the names they've been called, references to where that name was used, and which operation or tools they correspond to. This is an outstanding effort that can be used as a reference for creating MISP clusters such that when a name is entered from one vendor, it can be correlated to other names you may have already seen despite the different label.

[1] <https://medium.com/@cyb3rops/the-newcomers-guide-to-cyber-threat-actor-naming-7428e18ee263>

[2] <https://apt.threattracking.com>

Targeted Attack Case Study 1: Equifax 2017



Timeline:

- March 6 – Patch released for Apache Struts 2 (CVE-2017-5638¹)
- March 10 – Attackers find unpatched server, test exploit
- May 13 – Attackers come back, steal info, achieve root access
 - Attacker remain persistent, pivot internally, slowly extract data from 51 more databases over next 76 days
- July 29 – Breach discovered, vulnerable site removed 1 day later
- Data on 150M people stolen

Targeted Attack Case Study 1: Equifax 2017

The Equifax breach in 2017 is an interesting case study because we know the vulnerability that was exploited was only days old upon its initial usage. Truth be told, calling the Equifax breach a targeted attack is somewhat of an assumption. Most people know that Equifax holds sensitive information on almost every single person in the U.S. However, there are several other credit agencies that, in theory, hold the same data. We do know that Equifax was targeted likely with getting this information as the goal. What we don't know is if the attacker would have, or did, try these exploits on Experian or TransUnion, the 2 other major credit bureaus in the U.S.

The initial intrusion into Equifax's systems came on March 10, just 4 days after the patch was released by Apache for CVE-2017-5638. Interestingly, the attacker did not seem to do anything but test that the exploit worked on this day. It wasn't until 2 months later that they came back and started acting toward their goal.

This is an important example of one of the modern cyber defense mindset tenets—it's not over until the attacker achieves their objectives. Had Equifax detected this initial step in the kill chain at the intrusion in March, it's likely they would've patched immediately and none of this would have happened. As we stated earlier, however, the cost and complication of cleanup expands over time. After May 13, attackers were able to achieve root on the server they launched the exploit against, grab credentials for other systems, and begin to exfiltrate data. This exfiltration went unnoticed for 76 days while the intruders moved around to 51 different databases within the environment, pillaging info and shipping it out. It wasn't until July 29 when the intrusion was discovered, well into the adversary having achieved their goal, that the breach was noticed and the whole operation ground to a halt when the compromised server was taken offline. In the aftermath, the US Government Accountability Office published a report² with the details described above and analyzing the situation that led to the breach in the first place. The fallout was a record-breaking estimated \$439 million in damages³, only some of which are being covered by insurance. If this isn't a perfect lesson on the importance of fast patching, I don't know what is!

[1] <https://cwiki.apache.org/confluence/display/WW/S2-045>

[2] <https://www.gao.gov/products/GAO-18-559>

[3] <https://www.reuters.com/article/us-equifax-cyber/equifax-breach-could-be-most-costly-in-corporate-history-idUSKCN1GE257ff>

Targeted Attack Case Study 2: Sony Pictures 2014

- **"100TB" exfil'd** over 2 months
- November 24, 2014: Employees find this picture on all computers!
- PCs' **hard drives wiped**
- "Guardians of Peace" claim responsibility for hack
 - FBI analysis shows link to North Korean malware
 - Later attributed to "**Lazarus Group**"
- Embarrassing **confidential employee data, email and unreleased movies dispersed online**



Targeted Attack Case Study 2: Sony Pictures 2014

One of the most catastrophic and also interesting breaches occurred in November 2014 when the "Guardians of Peace" attacked Sony Pictures. Although the compromise wasn't known until the final goal was achieved, the attackers had apparently been in the network for months collecting damaging information from internal company servers, including VIP email, sensitive personal information, and digital copies of movies that had not yet been released. After exporting what the adversary claimed (but was never verified) to be 100TB of data, they let loose the final chapter of the breach—destructive malware that wiped every machine it could reach. On November 24, Sony Pictures employees found that when they came to work and turned on their PCs, they were presented with the skull picture shown above, as well the ransom note.

Business disruption attacks like this are rare, but not as rare as many believe. The ones that do occur tend to be kept as silent as possible, but this attack, due to its nature, leaked and made a media splash in a big way. It didn't help either that the attackers made good on their promise to release the data they had stolen, which led to private email conversations being made public that painted several people in a very unfavorable light.

Ultimately, this attack was blamed on North Korea, and the U.S. Government released an indictment of this attack, among others, that detailed how the attack was tied to North Korean spy Park Jin-Hyok.¹ The most thorough and definitive investigation available to the public of what happened is a joint report prepared by Novetta in cooperation with several other security vendors.² It also contains malware analysis and independent investigation that aligns with the possibility of the North Korean explanation.

[1] <https://int.nyt.com/data/documenthelper/274-park-jin-hyo-complaint/7b40e5ed5b185f141e1a/optimized/full.pdf>

[2] <https://www.operationblockbuster.com/wp-content/uploads/2016/02/Operation-Blockbuster-Report.pdf>

Targeted Attack Case Study 3: OPM 2015

Timeline

- Nov. 2013: First adversarial activity begins
- Mar. 2014: Investigation begins after **US-CERT warns OPM of activity**
- May 2014: Attacker uses **VPN access** from contractor to install **PlugX** malware
- May 2014: "Big Bang" initiated to kick attackers out, fails to stop activity from May 7
- Jun. 2014: Attackers get **mainframe access** and control of background investigation servers
- July 2014: Data exfiltration begins
- Dec. 2014: **Exfil of 4.2M personnel records**
- Apr. – Jun. 2015: Incident response begins
 - Tipoff: SSL cert seen for "opmsecurity.org"
 - Registered to Steve Rogers (Captain America)
 - Then opm-learning.org: Tony Stark (Iron Man)



Outcome:

- 21.5M people compromised
- Created "generation-long" national security issue
- Costs approaching \$1B

Targeted Attack Case Study: OPM 2015

One of the most disastrous compromises of personal information was the U.S. Office of Personnel Management breach. If you thought Equifax losing sensitive personal information was bad, this breach went even further. Included in the OPM data was not just Social Security numbers and addresses, but far more sensitive information that must be disclosed by anyone applying for a US Government security clearance. Information about family members, acquaintances, and roommates, psychological information, drug history, even fingerprints for some individuals. This information obviously cannot be changed and, as a result, those who are working undercover could potentially be exposed via biometrics despite fake names and other measures that might have been taken.

The OPM breach was the biggest and potentially most significant breach ever of government data and has been said to have "jeopardized our national security for more than a generation" by the official in-depth report on the incident.¹ As a result of this breach, both the Director of OPM and the CIO resigned. Costs are still piling up and are approaching \$1B for the remediation and ongoing monitoring of information required.

[1]

https://archive.org/stream/ReportFromTheCommitteeOnOversightAndGovernmentReformOnTheOPMBreach/Report%20from%20the%20Committee%20on%20Oversight%20and%20Government%20Reform%20on%20the%20OPM%20Breach_djvu.txt

Know Your Enemy Summary

- Actors range from governments to script kiddies
 - Consider **threat actor**: Who is attacking me?
 - Consider **motivations**: What do they want? Are they likely to attack us? Do they have the **capability** to do harm?
 - Consider **attack paths**: How will they get it?
 - Consider **defensive measures**: How to best monitor/protect against those threats, depending on group
- Keep this in mind during your analysis and response
- We will cover specifics in the coming days!

Know Your Enemy Summary

Throughout your career in information security, you will face the whole gamut of threat actors in your day-to-day job. Some will be easy to repel; others will show you things you've never seen before and stretch you to produce detailed plans and creative solutions for remediation.

We will dive much deeper into this topic throughout the course; but for now, be aware that you should constantly be trying to identify (at least between opportunistic and targeted) the type of group perpetrating any attack you encounter and consider what their motivations might be. It's important to read reports on these groups in your free time so that you are cognizant of the threat landscape and who might be after organizations like yours. It will give you an important perspective on who's out there, what type of techniques they're using, and, as we'll discuss later, the way you should act during remediation.

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. Know Your Enemy
- 12. Day 1 Summary**
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Day 1 Summary

- **Your goal:** Identify and halt cyber attacks before they become a true disaster
- This requires a thorough understanding of:
 - Tools at your disposal
 - Understanding your enemy and how to appropriately react
 - Network protocols and how they are misused
 - Signs of an attacked host, commonly weaponized file formats
 - Performing high quality analysis, and automating response
- The rest of this course will focus on these items!

Day 1 Summary

That brings us to the end of Day 1. As you've seen, Blue Teamers have a complicated challenge laid out ahead of them. But given the wealth of tools and some persistence, we can prevail over the attackers. To do so will require a thorough understanding of our tools, networks and hosts, and our enemy's motivations. Throughout the rest of this course, we'll explore these topics and more!

Coming Up...

SEC450.2 – Understanding Your Network

SEC450.3 – Understanding Endpoints, Logs, and Files

SEC450.4 – Triage and Analysis

SEC450.5 – Continuous Improvement, Analytics,
and Automation

SEC450.6 – Capture the Flag Contest!

This page intentionally left blank.

Course Roadmap

- **Day 1: Blue Team Tools and Operations**
- Day 2: Understanding Your Network
- Day 3: Understanding Hosts, Logs, and Files
- Day 4: Triage and Analysis
- Day 5: Continuous Improvement, Analytics, and Automation

Blue Team Tools and Operations

1. Welcome to the Blue Team!
2. Exercise 1.0: Virtual Machine Setup
3. SOC Overview
4. Defensible Network Concepts
5. Events, Alerts, Anomalies, and Incidents
6. Incident Management Systems
7. Exercise 1.1: TheHive Incident Management System
8. Threat Intelligence Platforms
9. Exercise 1.2: MISP Threat Intelligence Platform
10. SIEM and Automation
11. Know Your Enemy
12. Day 1 Summary
13. Exercise 1.3: SIEM with the Elastic Stack

This page intentionally left blank.

Licensed To: David Owerbach <0mamaloney@gmail.com> April 28, 2020

Exercise 1.3: SIEM with The Elastic Stack



Exercise 1.3: SIEM with The Elastic Stack

Exercise 1.3: SIEM with The Elastic Stack

Please go to Exercise 1.3 in the SEC450 Workbook or virtual wiki.