SIEM/SOAR TESTBED IMPLEMENTATION FOR USE IN UNDERGRAD AND POSTGRAD MODULES

Research Manual

Name: Darragh Murphy

Student ID: C00135047

Project Supervisor: James Egan

24/04/2022

Abstract

SIEM and SOAR technology solutions are being adopted by a much larger portion of business across the world. Developing the skills required to use these tools will become essential in the years to come. The goal of this project is to create an environment where students can begin to develop those skills. This document contains the research that has been completed on the types of technologies and their inner workings, some software solutions that can be used to throughout the education lifecycle of future students.

Contents

Introduction

The world we live in today is plagued by cyber-attacks. It feels that nearly every day another company announces that they have been hit by a ransomware attack or suffered a data breach. So far in 2021, the number of cyber-attacks performed against companies has increased by 29% and ransomware attacks specifically have increased by 93%[1].

In 2019, the average cost of a data breach was $3.92 million, this decreased to $3.86 million in 2020[2]. This decline would lead you to believe that cybercrime and the costs associated with it are reducing however, this is misleading.

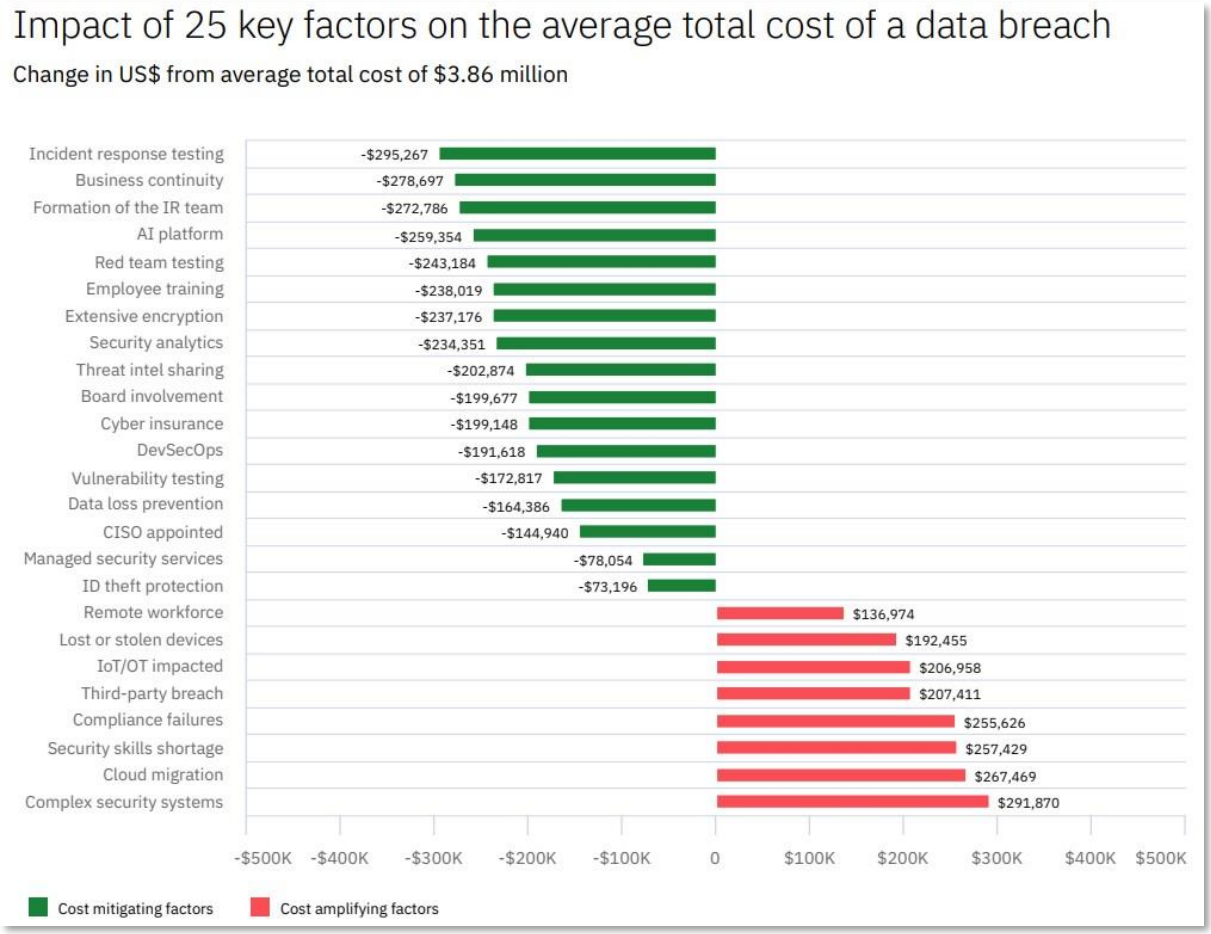

Figure 1: Key Factors for calculating the cost of a data breach.

[1] (Check Point Software Technologies Ltd., 2021)

[2] (Ponemon Institute LLC, 2020)

Many organizations are choosing to invest their time and money into their security teams. Organizations are developing Security Operations Centers (SOC) and they are implementing many robust security tools, as well as hiring staff to monitor these tools. Companies are choosing to adopt an approach of integrating Security Information and Event Management (SIEM) and Security, Orchestration, Automation and Response (SOAR) tools with other parts of their security infrastructure. Adopting these software solutions is a major step forward in protecting your company from cyber-attacks.

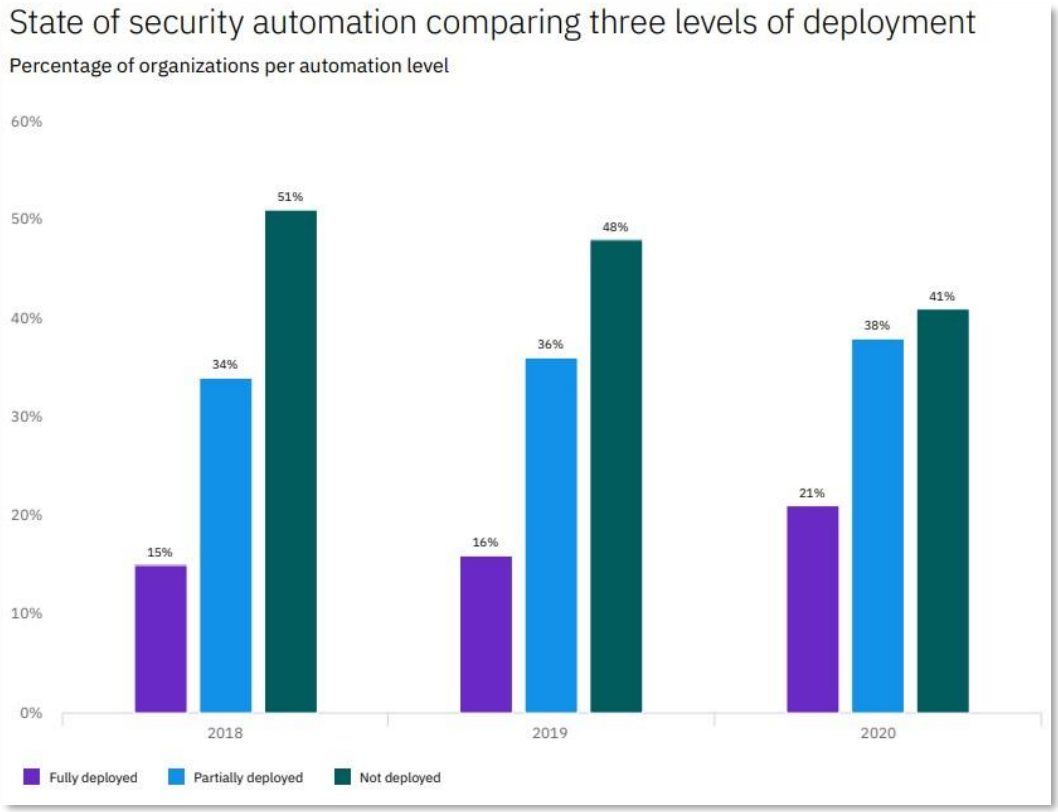

Figure 2: IBM Cost of a Data Breach Report 2020

It has been found that only 21% of companies worldwide in 2020 have fully implemented security automation solutions, compared to a figure of 16% in 2019. The average total cost of a data breach for a company with a fully implemented security automation security solution in 2020 was $2.45 million dollars, compared to $6.03 million dollars for a company without security automation. This is a difference of a staggering $3.58 million dollars. It is easy to see why companies are migrating to these new methods[3].

[3] (Ponemon Institute LLC, 2020)

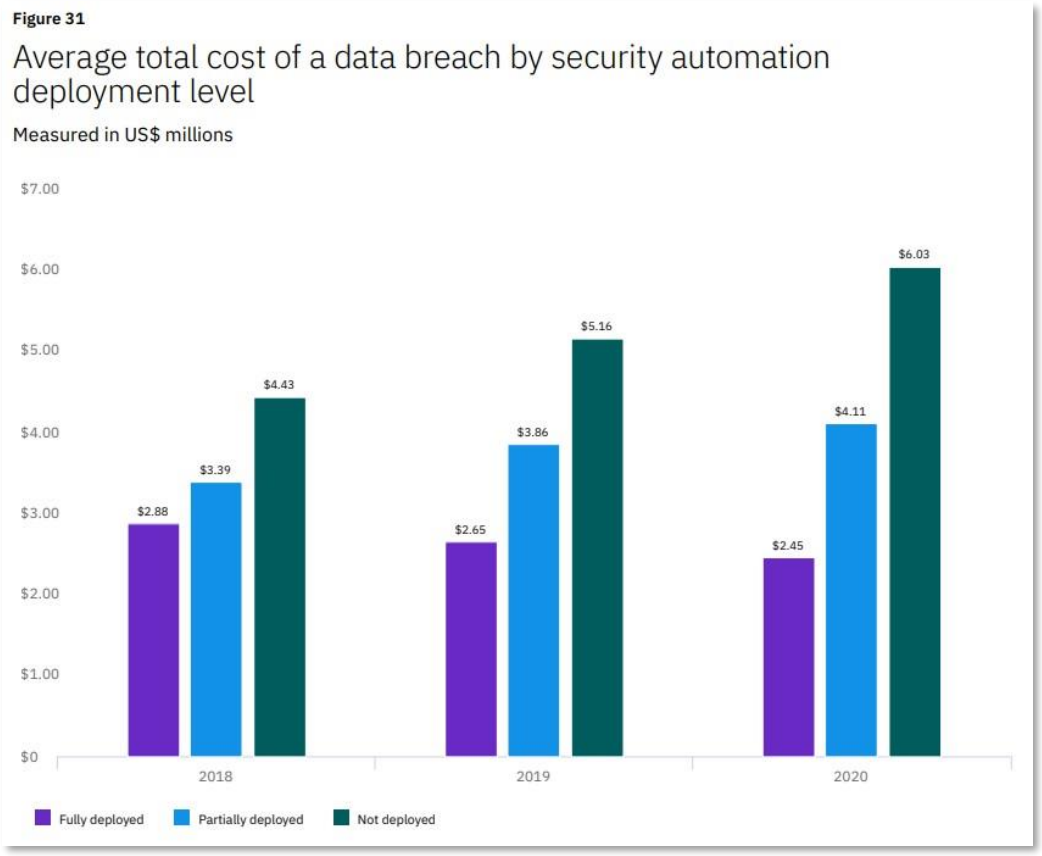

Figure 3: IBM Cost of a Data Breach Report 2020

This document will focus on my research into various different SIEM/SOAR technologies and help determine the correct combination of tools that will allow the creation of test environment in a computer lab. The aim of this project is to investigate how a SIEM/SOAR technology solution can be established within IT Carlow for future study.

Students will have the opportunity to learn about SIEM/SOAR technologies in a practical setting. This experience will be very beneficial to students once they begin their careers.

SIEM

Security Information and Event Management (SIEM) is a technology solution that offers organizations tools that perform detection, analysis, and response to threats from malicious actors.

- A SIEM tool provides real-time analysis of threats by utilizing data analysis capabilities and generating alerts based upon preconfigured sets of rules.
- Log data is collected from a variety of sources across an organization's infrastructure, including network devices, endpoints, and servers.
- This data is collected into one centralized platform, and it can be displayed to analysts via Dashboards and Visualizations.

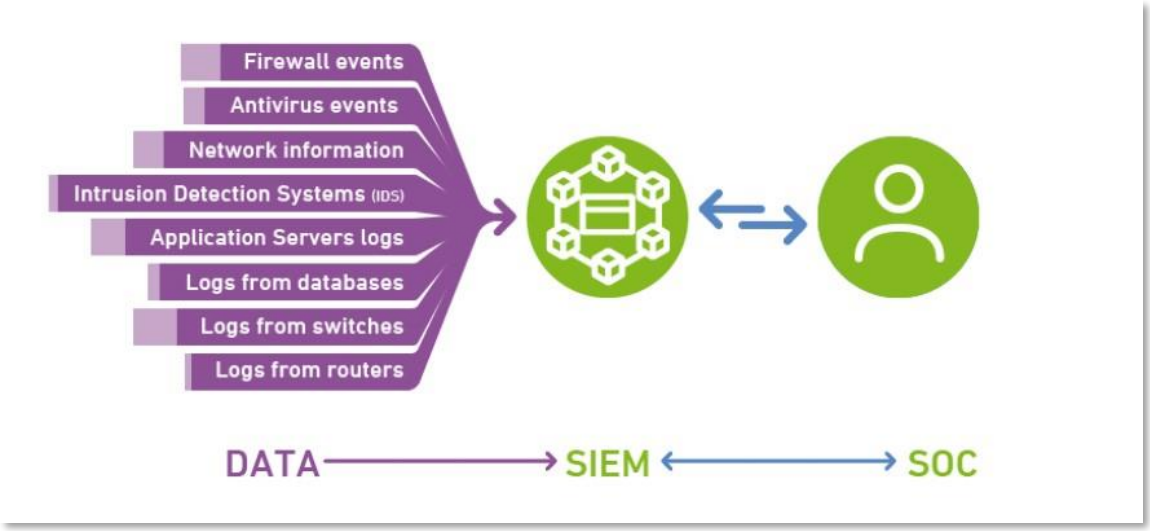

Figure 4: How a SIEM works.[4]

History of SIEM

The first use of the term SIEM was in a Gartner Report published in May 2005, titled "*Improve IT Security With Vulnerability Management*". This report stated that Security Information and Event Management was one of the four main categories of technology that could be used to provide automation in the vulnerability management process, along with

[4] (Leaseweb, 2021)

vulnerability assessment, security configuration management and policy compliance and IT security risk management.[5]

[5] (Williams & Nicolett, 2005)

SIEM technology is built upon two pre-existing management practices. These were Security Information Management (SIM) which provides log storage, log analysis and reporting capabilities and Security Event Management (SEM) which provides data aggregation, establishes correlation between security events and provides real-time monitoring of an organization's security infrastructure.

How does a SIEM work?

A fully functional SIEM tool has many capabilities, however These five operations are:

1. Data Collection
2. Data Storage
3. Policies and Rules
4. Data Consolidation and Correlation
5. Real-time Monitoring and Alerts

Data Collection

Data Collection occurs using by deploying 'agents' deployed to the target system to enable log collection. Throughout this stage logs are collected from all available sources within an organization's infrastructure, such as endpoints, servers, operating systems, firewalls, anti-virus software, and many more. Logs are also collected via protocols such as syslog, SNMP and WMI. Some SIEM tools offer a feature where logs can be processed and filtered at the collection agent. The collected logs are then forwarded to a centralized location for storage at regular intervals.

Data Storage

Data Storage is an important part of SIEM implementation. Today companies are ingesting an extraordinary amount of data ranging from customer data to network logs. For a SIEM to operate effectively, it needs to be able to access and analyze 100% of your data.

In the past, it was difficult and costly to store such volumes of data. Data centers were a necessary investment for organizations. However, the use of Data Lake technology (e.g., Amazon S3) allows an organization to store as much data as required in one centralized repository and this data storage is scalable.

Policies and Rules

The creation of policies and rule sets are an essential part of SIEM configuration. Individual profiles can be created for systems within an organization's infrastructure. These profiles help define how a system behaves under normal circumstances.

A SIEM will provide a default set of rules, alerts but these can be configured and tuned to better integrate with your organization. These set of rules can be applied to the profiles that were created for each system.

If the SIEM detects behavior that does not conform with the policy and rules for a given system, it will generate an alert that a security analyst may need to investigate. The SIEM will assign the alert a specific threat level based upon pre-determined sets of rules.

Rules can also be created that relate to specific Tactics, Techniques and Procedures (TTPs) used by malicious actors. Indicators of Compromise (IOC) are shared between many organizations, and they may also be posted publicly throughout the investigation of an attack.

Data Consolidation and Correlation

After the logs are collected, a SIEM tool can begin the final part of the Log Management Process. This involves the consolidation and correlation of data.

Data Consolidation involves using data aggregation techniques which parse the logs from a number of sources, extracts structured data, and normalizes it, so it can be displayed in a format that is easily readable. This also ensures the data is searchable by analysts as they are investigating an alert generated by the SIEM.

Data Correlation involves combining security events into a meaningful security alert. It does this based upon certain security events that may form a pattern that relates to a TTP or IOC from a known attack.

Real-time Monitoring and Alerts

As the logs are analyzed and alerts will trigger based upon pre-configured rules. Analysts will receive a notification and Data correlation provides analysts with greater context for the alert, and it allows analysts to drill down further into the logs. This enables an analyst to get

a clearer picture of what has occurred to cause the alert. Analysts will be more confident making decisions during an investigation when they are able to gain a greater insight into the logs.

SIEM Use Cases

SIEM technology has been adopted by many organizations over wide range of different industries and as such there are many use case examples relating to security. However, if you wish to build the use cases yourself, ensure to focus on three areas:

1. Insight
2. Data
3. Analytics

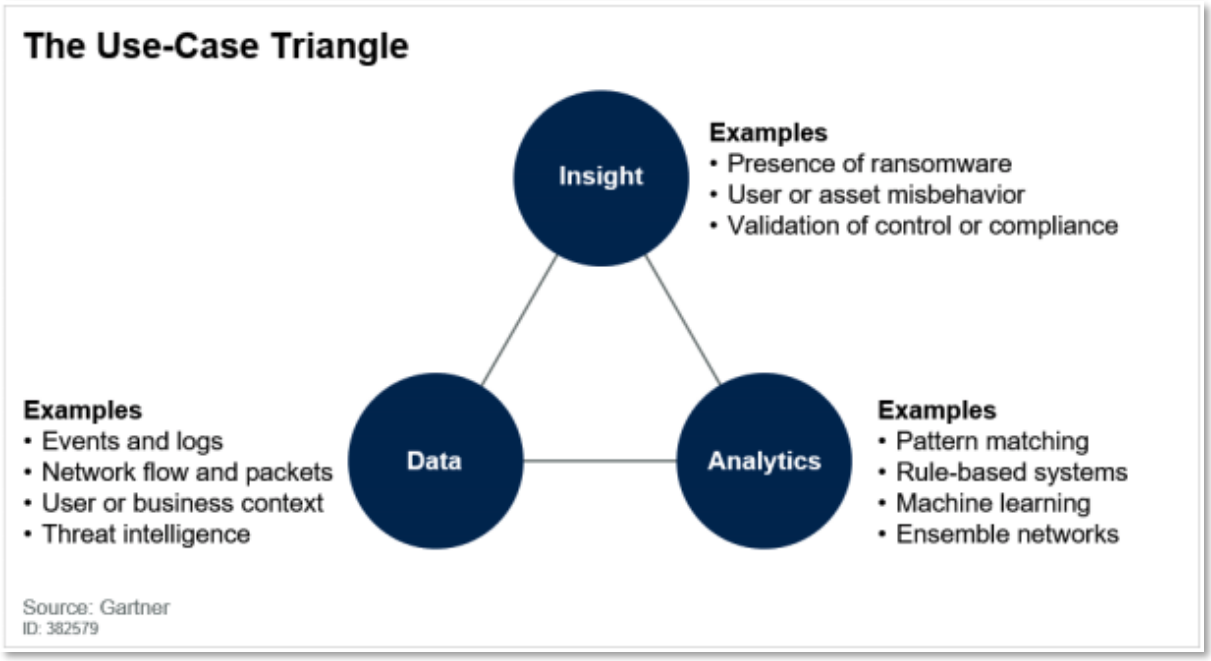

Figure 5: The Use-Case Triangle[6]

The first step that you should take when creating a Use-Case is to define what data is required. As discussed before, storing data for a SIEM is expensive and not all companies can afford to store all their data indefinitely. Defining what data is required for each use-case is an important step in this process.

[6] (Babacamp, 2020)

The second step is to apply analytics to the data. Depending on your SIEM, this could involve machine learning, or insight can be gained by using correlation rules. Analyzing the data can help find patterns in security events, that when combined form a sequence of events that will fire an alert.

The final step in use-case creation is organizing and prioritizing use-cases. The amount of use-cases that a security teams needs to manage depends on the organization. However, no matter how many use-cases are deployed, they should be organized.

Gartner suggest organizing use-cases by category. This allows for easy navigation and will help avoid duplicating use-cases if the naming system is unclear.

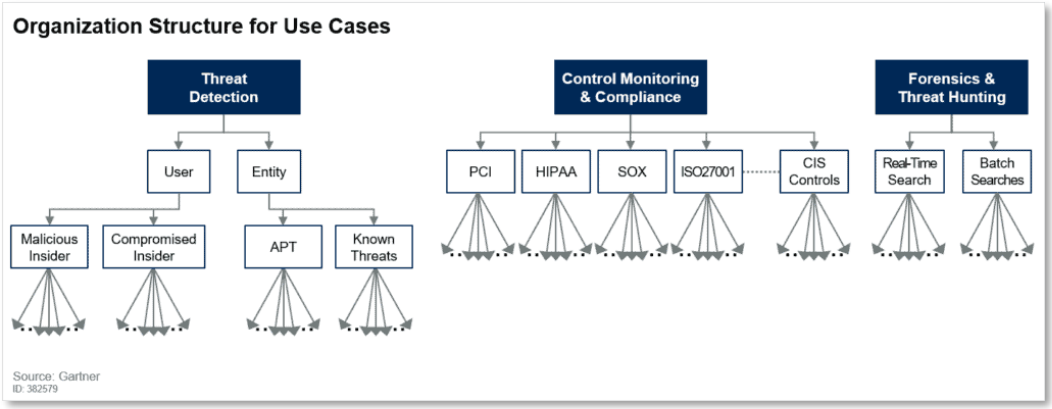

Figure 6: Gartner's Organization for Use Cases[7]

[7] (Babacamp, 2020)

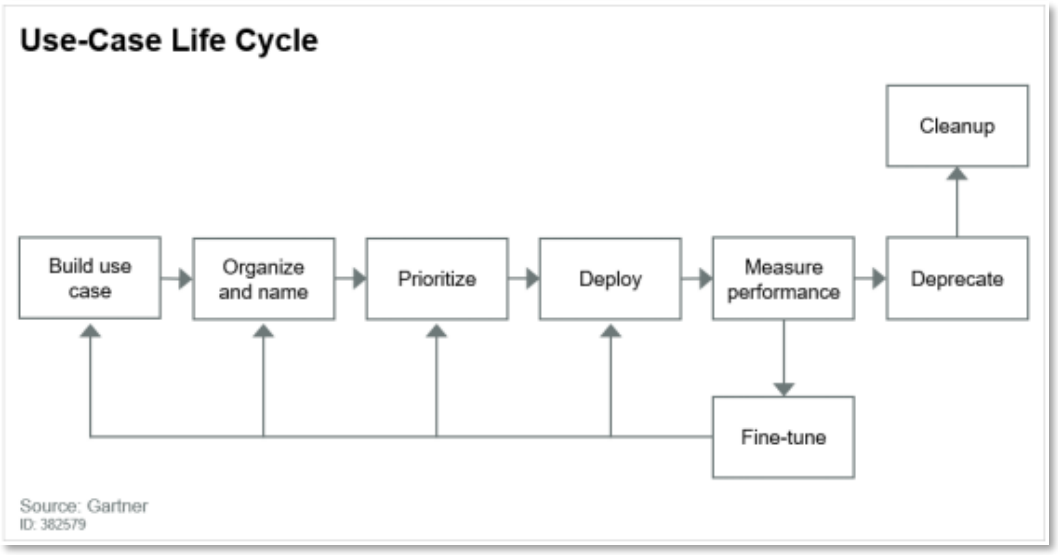

Figure 7: Gartner's Use-Case Life Cycle [8]

There are many examples of use-cases, however most of common use-case examples fall into these three categories:

- Compliance
- Insider Threats
- Threat Hunting

SIEM Use-Case: Compliance

The increasing level of government regulation and the need for compliance has led to a large increase in IT infrastructure for major organizations across the world. Organizations often have to deal with regulations for many different governments, depending on where they want to operate. The GDPR is a good example of this type of regulation. Any business that wants to operate within the EU must abide by the GDPR.

Organizations investing in their IT infrastructure has led to an increase in the deployment of SIEM tools to help with such regulations. A SIEM tool can help a company comply with the following regulations:

- PCI DSS (Payment Card Industry Data Security Standard)
- GDPR (General Data Protection Regulation)

[8] (Babacamp, 2020)

- HIPAA (Health Insurance Portability and Accountability Act)
- SOX (Sarbanes-Oxley Act)

SIEM Use-Case: Insider Threats

As an organization grows, so does its surface for attack. One of the most vulnerable areas for a company are its employees. Employee access needs to be managed efficiently, with many different privilege levels and group policies applied to profiles.

There are three categories of insider threats:

- **Unintentional Threat:** This can occur if an employee is negligent or through accidental disclosure of credentials or sensitive information.
- **Intentional Threats:** This can occur through disgruntled employees that are looking for personal gain or revenge on the company.
- **Other Threats:** This relates to employees and/or that may be compromised and have been recruited by cybercrime gangs in order to enable a data breach.

According to a Verizon Data Breach Investigation Report, the cause of three out of five security breaches were related to an insider threat and that these threats could go undetected for months or even years.

A SIEM tool can help detect compromised accounts, if an employee is exfiltrating data, lateral movement and even detect network communication to a C2 (Command and Control) server.

Deploying User Entity Behavioral Analytics (UEBA) with your SIEM tool can reduce this risk. UEBA will build profiles for your employees based upon their normal activity, and then it can detect unusual behavior.

SIEM Use-Case: Threat Hunting

According to CrowdStrike, threat hunting is the practice of proactively searching for cyber threats that are lurking undetected in an organization or network[9]. Performing a threat hunt requires access to log data, and this is where a SIEM comes into play.

[9] (Taschler, 2018)

Log data is collected and stored in a centralized location and is easily searchable by the SIEM. The SIEM allows a user to search for a known IOC or TTPs during a threat hunt, which significantly speeds up an investigation. Investigators can gain a good overview of a threat this way.

The ability to create alert rules within a SIEM will allow analysts to create a rule that fires an alert if the SIEM detects known IOCs from threat intelligence feeds is another important factor in this use-case example.

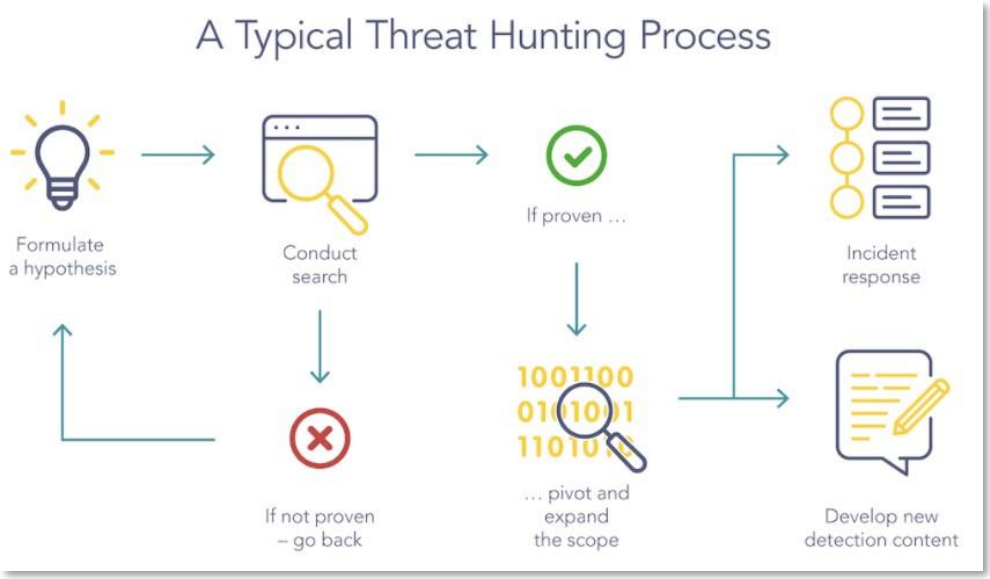

Figure 8: Log Point Threat Hunt Process

The Future of SIEM

As technology and threat actors advanced, so did the companies who were deploying SIEM tools as part of their defense. Companies were soon collecting more data than they knew what to do with and they may have come under constant attack from malicious actors.

The limitations of SIEM technology soon became evident. SIEMs had limitations in the following areas:

- Due to the increasing amount of data, SIEM tools were unable to process and analyze all of the relevant data. This severely limits the effectiveness of the SIEM.
- SIEM technology can create alerts, many of which are false positives. This in turns wastes an analyst's time, when they could have spent that time investigating an actual threat.

Figure 9: Exabeam SIEM Evolution Timeline

Advances in technology have helped mitigate these limitations and these advances have greatly improved the efficiency of a SIEM in preventing threats. Some of the advances are as follows:

- The use of "Big Data Architecture" allows for a smoother integration with other parts of your organization’s infrastructure, including Cloud Services, on site machines, and BYOD.
- It is possible to integrate threat intel feeds from many sources including commercial and open-source entities.
- SIEMs provide many visualization tools that allow users to prioritize alerts in real time.

- The use of behaviour analytics in combination with SIEM tools can help give better context to an event, and highlight a change in user behaviour.
- Security teams can build custom workflows based upon specific events and unique situations.

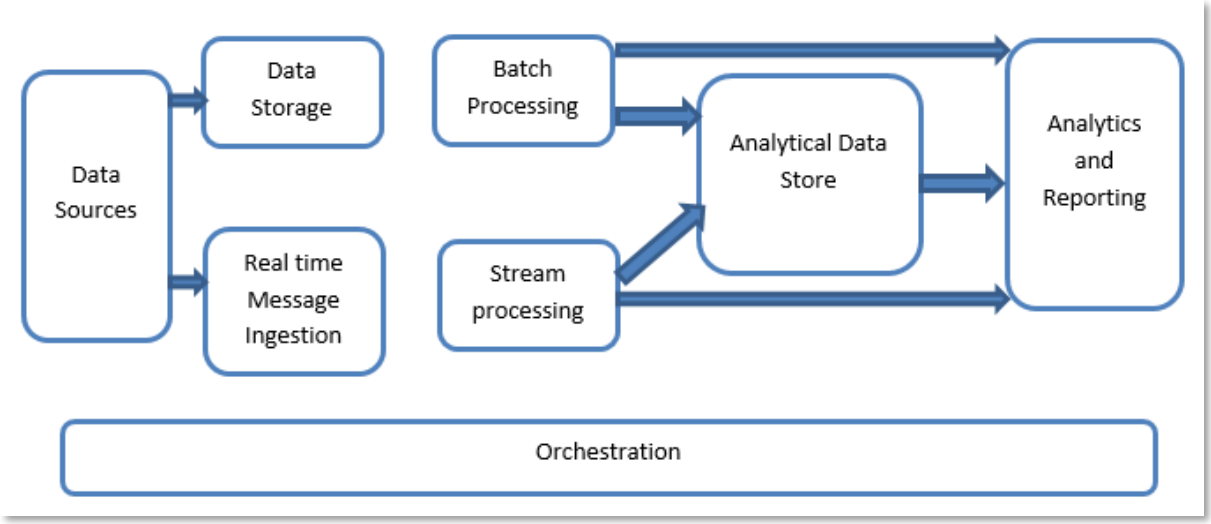

Figure 10: Big Data Architecture Diagram [10]

SOAR

Security, Orchestration, Automation and Response (SOAR) is a collection of software solutions that allow organizations to use machine learning to analyze data and gain better insights into the data that is being collected. This allows organizations to become more efficient when analyzing alerts.

The three main areas where SOAR will have the most effect are:

1. Threat and Vulnerability Management
2. Incident Response
3. Operations Automation

[10] (OmniSci, 2021)

An article in Infosecurity Magazine from 2018 claimed that companies were experiencing approximately 174,000 alerts per week[11]. Hidden in that noise are some very real threats that could be missed by security analysts if they were required to manually work each alert.

Automating tasks that are repetitive and that are required to be completed regularly can help free up time that a member of the security team would normally spend on those tasks. In a busy environment like a SOC, every minute counts.

Orchestration allows organizations to chain task together, across multiple systems, in order to create larger processes and workflows.

Deploying a SOAR solution within your organization can bring many benefits to your company. It allows for easier integration of across your infrastructure, from Operations to Security and any relevant tools. SOAR technologies also reduce the mean time to detect (MTTD) and mean time to response (MTTR) by using various automation techniques.[12]

Security Automation

Security Automation is defined by Splunk as the machine-based execution of security actions with the power to programmatically detect, investigate and remediate cyberthreats with or without human intervention by identifying incoming threats, triaging, and prioritizing alerts as they emerge, then responding to them in a timely fashion.[13]

As an organization grows, so does its attack surface. New developments in technology, and the threat from Covid-19 have exacerbated this issue. In the past most companies required their employees to work on site, however more employees are required to work from home due to the lockdowns that governments across the world are implementing.

This change in working habits has required companies to provision laptops, desktops, tablets, and many other types of devices to their employees to enable them to be able to perform their job in a remote setting.

[11] (Zurkus, 2018)

[12] (Splunk, 2021)

[13] (Splunk, 2021)

Without automation, a security team would be required to monitor these additional devices manually. The additional time that analysts spending monitoring new devices, the less time an analyst spends investigating a genuine alert, compared to a false positive.

The key areas where Security Automation can benefit a company are as follows[14]:

- **Monitoring and Detection:** Automation can provide visibility across an organization's infrastructure and can notify analysts when it detects a threat. This eliminates the need for an analyst to monitor the entire security process.
- **Data Enrichment:** Automation can help provide extra content for a security alert by gathering any logs that may be associated with an event. This will help the analyst perform a more detailed investigation.
- **Incident Response:** Automation can provide a response to a threat immediately after a threat is detected. It does this by implementing response playbooks which are created based upon a set of rules.
- **User Permissions:** Automation ensures that the permissions are issued to accounts smoothly. This feature can also be used to automatically remove permissions from a user account. This can be useful if an insider threat is discovered or in the event where a user account is compromised, and the malicious actor attempts to escalate their privileges.
- **Business Continuity:** Automation ensures that in the event of an attack, the relevant servers containing backups and important data and services are taken offline immediately.

Security Orchestration

Security Orchestration is defined by Splunk as the machine-based coordination of a series of interdependent security actions across a complex infrastructure[15].

What this means is that the tools that are used across an organization are integrated so that they function together as one entity. With the increasing number of tools that companies use to monitor their network, orchestration seems like the next logical step to take. The

[14] (Rapid7, 2021)

[15] (Splunk, 2021)

ability to link tasks from multiple services and tools removes the limitation to automating tasks from a single service. Orchestration can streamline the whole workflow for more complex security processes[16].

The ability to automate more complex processes with SOAR frees up time that security analysts would normally spend working on those tasks. Analysts can better focus on other areas of security that would be beneficial to the company, such as performing threat hunts.

SOAR Use Cases

Before implementing a SOAR solution, an organization should identify the areas where SOAR could have the most impact. This can be done by creating relevant Use Cases to your organization.

Depending on what industry your organization is a part of some companies may have use cases that are more specific for them. However, some of the most common use cases across all industries for SOAR technologies are as follows:

1. Incident Response
2. Threat Hunting
3. Vulnerability Management

SOAR Use-Case: Incident Response

One of the most prevalent use-cases for SOAR technologies is the area of incident response. Attacks can vary in techniques and procedures. As such, there are many different use-cases that fall under the incident response category.

One of the most common issues that companies face is phishing. Malicious actors attempting to steal login credentials or to deploy malicious software bombard organizations with emails that may or may not appear to be legitimate. Phishing emails can be investigated using a playbook, which parses the email and pulls relevant signatures from the email. The signatures can then be checked against online databases such as VirusTotal to confirm if they are legitimate or malicious.[17]

[16] (Rapid7, 2021)

[17] (Anon., 2022)

SOAR Use-Case: Threat Hunting

Another top SOAR User-Case is using the features to enable deeper threat hunting capabilities. Automating the log collection and gathering of important data, including any IOCs, enables the analyst to determine the threat level and if it required immediate attention.[18]

SOAR Use-Case: Vulnerability Management

SOAR is also very useful when it comes to applying it to Vulnerability Management. The automation features of SOAR allow in depth log collation surrounding the detected vulnerability and allows for a quick resolution to the problem at hand.[19]

Disadvantages of SOAR

The benefits of using a SOAR technology solution to aid in bolster the security of your organisation is very real, and sometimes costly to implement.

However, human input should not be removed. These technology solutions are code based, and in the end, they are tools to be used. Human insight is still an important part of managing an organization's security.

A machine will not be able to provide the same rational as a human might in some situations. Any decisions the automation makes, is based on a given set of rules.

For example, your organisation has received an alert relating to a user that failed to authenticate with your servers. This user has inputted an incorrect password many times. Was this just a genuine mistake or is there a malicious actor trying to brute force this account?

This is where human insight comes into play. An analyst will investigate this alert and make a final decision on whether this was a malicious alert or not.

[18] (Cyware, 2020)

[19] (Cyware, 2020)

SOC

A Security Operations Centre (SOC), also known as a Cyber Security Operations Centre (CSOC), is the central point of an organization's security team. Here analysts monitor the available security tools and respond to alerts. Analysts monitor the queues using dashboards created by a SIEM tool and action alerts as the fire in a queue.

A SOC operates 24 hours per day, 7 days a week and 365 days a year.

The security team that works in a SOC consists of analysts and engineers who monitor activity with an organization. Engineers are tasked with configuring and tuning the security tools, while analysts monitor for alerts.

Employees within a SOC are categorized into 4 separate areas based upon skillsets and expertise. These are as follows:

- **Level 1:** Incident Responders. Monitor for alerts using the available security tools, investigate alerts and determine if they need to be escalated.
- **Level 2:** Incident Responders. These analysts would have more experience than level 1 incident responders.
- **Level 3:** Expert security analysts. Performing vulnerability scanning, and threat hunts to actively look for threats.
- **Level 4:** Managers, CISO. This level oversees the operations within the SOC. Staff at this level work in the field for years.

UEBA

Rapid7 defines User and Entity Behavior Analytics (UEBA) or User Behavior Analytics (UBA) as the process of gathering insights into the network activity and the events generated by users every day.[20]

UEBA can be deployed in an enterprise to monitor device and user behavior. Through analytics techniques, this before is logged over time, and normal behavior can then be defined. UEBA can then identify when there is abnormal activity on the network, be it a

[20] (Rapid7, 2022)

malware executing remotely to an insider who may be exfiltrating data. Some activity may generate some false positives, and as with any security tool, there may be some tuning involved.[21]

Researched Technologies

Atomic Red Team

Atomic Red Team is an open-source project, developed by the cyber security company Red Canary, that provides a library of tests that a security team can execute to test an organizations defence. [22]

The tests are mapped to the MITRE ATT&CK framework, and they are created in such a way that they can easily be deployed and reproduced. Atomic Red Team does not require many dependencies and that paired with the fact that it is easily deployable, and the tests are reproducible makes it an ideal candidate to use to trigger events that will produce log telemetry for analysis. [23]

Red Canary are a cyber security company based in Denver, Colorado. On top of developing open-source tools (Atomic Red Team), Red Canary provides managed detection and response that help augment an organizations security posture.

Invoke-AtomicRedTeam

The Atomic Red Team developed a PowerShell Execution Framework that enables automation of the testing. When installed, the atomics folder contains configuration files that define the attack procedure. The test configuration files also contain clean up commands to reverse the system changes and remove created files after a test has been executed. [24]

Using this module will allow me to execute the relevant tests quickly and due to the fact that the tests can be executed with simple commands

[21] (Fireeye, 2022)

[22] (Red Canary, 2022)

[23] (Atomic Red Team, 2022)

[24] (Roberts, 2022)

System Requirements

Due to the changes that the Atomic Tests can make to your environment, it is best to perform these tests in a non-production environment.

The tests can also be performed across multiple OS types, once PowerShell has been installed.

Before running any tests, ensure that you have a tool configured to collect the log data generated during testing.

Cuckoo Sandbox

Cuckoo Sandbox is an open-source automated malware analysis system. A sandbox is an isolated virtual machine, where malware can be detonated safely, without infecting your IT infrastructure. Analysts may use this tool as part of the process of an alert investigation.

The Cuckoo Sandbox is able to analyze all malware files that have been engineered to affect a specific operating system. It is able to analyze Windows, Linux, Android and MacOS files. It is also able to analyze many different file types.

When a file is detonated, Cuckoo Sandbox will pull all relevant IOCs and signatures that it detects. This includes file hashes, IP addresses, process activity, network activity and any other relevant signatures.

Once the file is analyzed Cuckoo Sandbox will generate a detailed report that an analyst can use to help gain a better idea of what they are dealing with. The report will provide:

- **Summary:** Information about the file including file hashes, file path, and file size.
- **Information on Execution:** Information about the file execution including when the file was executed, how long it took, and routing information.
- **Signatures:** Information about any signatures detected. Signatures are colour coded to indicate risk level. Blue indicates low risk and Red indicates high risk such as beacon activity, keylogging, and log erasure to hide activity.
- **Screenshots:** Any screenshots taken during file analysis will be here. This can include command line activity, opening programs, and ransom messages.

- **Domain & IP Information:** Any domains and IP addresses that were detected during file analysis will be summarised here.
- **Static Analysis:** Static malware analysis is performed in the sandbox on the file and a summary is provided in the report. Details include listing DLLs and APIs that the malware imported, sections of the executable and any other resources found.
- **Behavioural Analysis:** Information about processes that were created by the file that is being analyzed. It is possible to view the process tree.
- **Network Analysis:** Information about network activity is provided here. This includes host information, DNS information and any other network activity detected.

Cuckoo Sandbox was developed by Claudio Guarnieri, a security researcher who was featured in Forbes magazine's "30 under 30" back in 2016. Claudio was part of a team that uncovered government surveillance against political activists, in countries that had poor human rights. Claudio has worked with Rapid7 and Citizen Lab in the past. He is the current "Head of Security Lab" for Amnesty International.

System Requirements

It is recommended that Cuckoo is installed on a Linux Operating System.

Software Requirements:

- Python 2.7
- MondoDB
- PostgreSQL
- tcpdump
- Volatility
- M2Crypto
- guacd

Minimum System Requirements[25]:

- 2 GB RAM
- 40 GB HDD

Recommended System Requirements[26]:

[25] (Oktavianto, 2013)

[26] (Oktavianto, 2013)

- Quad Core CPU
- 4 GB RAM
- 320 GB HDD

Detection Lab

The Detection Lab is a virtual environment developed by Chris Long, who is a security engineer based in California and has worked with some of the world's largest companies. [27]

The Detection Lab is a repository hosted on GitHub that contains a series of Scripts that when executed, spin up a number of virtual machines to simulates a production environment.

The Virtual Machines are configured with many security features and tools and according to Chris Long, he built this to be used by Defenders. The environment is designed to be purposely vulnerable in order to gain better insights through the log data. [28]

The Detection Lab comes preconfigured with many different security tools and has a series of Group Policies applied to the machines, including some that enable log collection for command line arguments. This is a significant feature, that will allow for mor accurate log collection. [29]

The scripts that are contained within the Detection Lab Repository will deploy four virtual machines:

- An Ubuntu Server configured with Splunk, Suricata, Zeek and a number of other tools that enable log collection.
- A Windows Server machine configured as a Domain Controller.
- A Windows Server configured as a Windows Event Forwarder. This machine collects logs from the DC and the Windows 10 machine, and forwards them to the Ubuntu Server for analysis.
- A Windows 10 machine that will simulate a user.

[27] (Long, 2022)

[28] (Long, 2017)

[29] (clong, 2022)

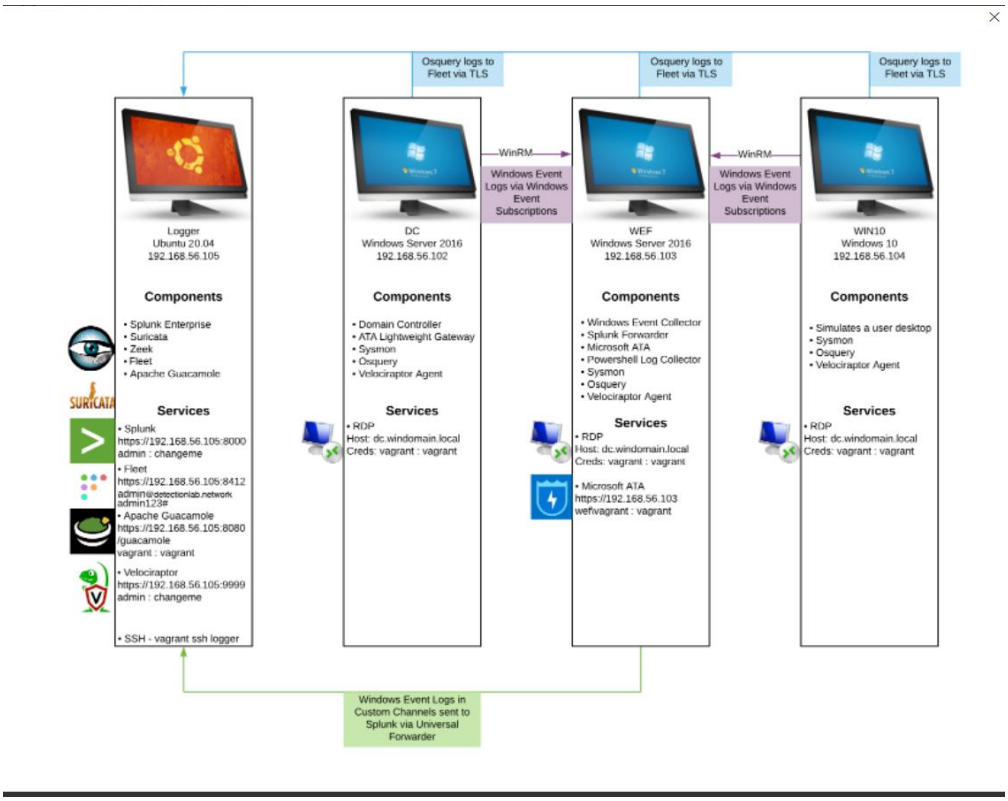

Figure 11: Detection Lab

System Requirements

Due to the number of Virtual Machines required, the system requirements to implement the Detection Lab is quite high. [30]

- 55 GB+ of Hard Disk Space
- 16 GB+ of RAM
- 8+ CPU Cores
- VirtualBox 6.0+
- Vagrant 2.2.9+

ELK Stack

The ELK Stack is a combination of three opensource projects, Elasticsearch, Logstash and Kibana, that can be leveraged to ingest logs, analyse the data, and display this information back to a user.

[30] (Detection Lab, 2022)

When used in conjunction with each other, the ELK stack may be used in place of a SIEM tool to analyse log telemetry data.

Elasticsearch

Elasticsearch provides the data analysis service that a SIEM tool can provide. Once the data has been ingested and indexed, a user can use Elasticsearch as a search engine for your logs. Analysis may be performed in order to discover trends within your data, such as user behaviour. [31]

Elasticsearch stores the data in JSON format and using a data structure called an inverted index, it allows the data to be searched in an extremely fast manner and according to the documentation, this occurs within 1 second. An inverted index is one which lists every unique word that appears in a document, and then identifies all the documents that that word occurs in. [32]

Elasticsearch also provides a feature that allows data aggregation that will help a company gain valuable insights into their data, such as the number of failed sign-in attempts on a device within a specific time period. [33]

On the backend, Elasticsearch consists of a series of nodes that are connected to each other in what is called a cluster. A node is a single server and can be assigned a specific task, such as data storage or the creation of further nodes. [34]

Logstash & Beats

Logstash and Beats are two significant components of the ELK stack that provide log ingestion and forwarding capabilities.

[31] (Elastic, 2022)

[32] (Elastic, 2022)

[33] (Elastic, 2022)

[34] (Abueg, 2020)

Logstash is a data collection engine that allows you to ingest data from a number of sources across an organisation. It can provide structure to unstructured data, identify certain fields (e.g. IP addresses) and even anonymise sensitive data.[35]

Beats is tool that allows you to forward log data to Elasticsearch directly or forward it through Logstash first to enable further data enhancement before it reaches Elasticsearch. Beats is installed as an agent on the target device and can provide forwarding for many log sources. [36] For example:

- Network Traffic is called a Packetbeat.
- Windows Event Logs are called Winlogbeat.

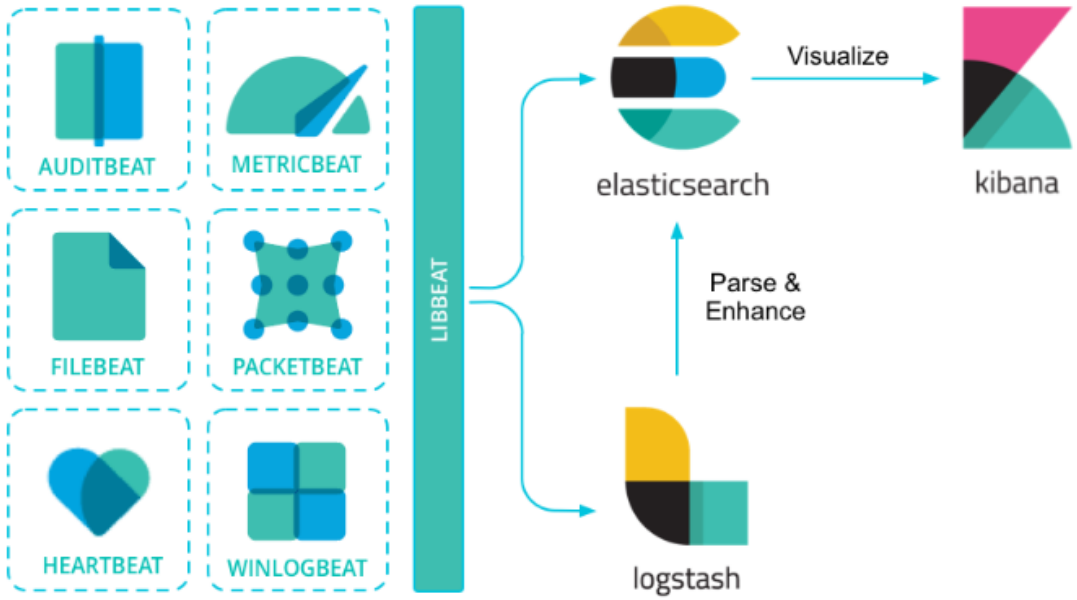

Figure 12: Beats Agents forwarding log data.

Kibana

Kibana is the front end to the ELK stack, that provides a user with data visualisation for the ingested log data. Kibana also allows you to perform data management techniques on your ingested data, such as setting an expiry date for data retention purposes. Kibana also has a

[35] (Elastic, 2022)

[36] (Elastic, 2022)

feature that enables firing an alert for certain detections. This is a useful aspect for any tool that is used to analyse log data. [37]

Pricing

The ELK stack is open source, and therefore can be deployed within your environment for free. However, they do have subscription services that offer varying degrees of services. The pricing for their cloud services seems reasonable, beginning for as little as $95 per month for the "Standard" offering, to as little as $175 per month for the "Platinum" service. As with most SaaS (Software-as-a-Service), the more premium the service, the better the features that are available. [38]

Elastic's pricing model is based on what resources you use to store the data, not the total amount of logs ingested per month, or other arbitrary measures to generate revenue. [39]

Guacamole

The Guacamole Project had its beginnings as a JavaScript Telnet client called RealMint, however over the years, it has evolved into a HTML5 web application built upon the Guacamole API that can provide remote access to Desktops using a number of different protocols, such as VNC and RDP. By configuring a Guacamole Server on your network, Guacamole can be configured to provide access to any device on your network remotely by using a web browser. [40]

As I mentioned above, Guacamole is a web application built upon on API of the same name. [41]

- A user opens the Guacamole Web Application in their Web Browser.
- The Guacamole Web Application is hosted on the Guacamole Server and is presented to the user.

[37] (Elastic, 2022)

[38] (Elastic, 2022)

[39] (Elastic, 2022)

[40] (Apache Guacamole, 2021)

[41] (Apache Guacamole, 2021)

- The Web Browser communicates with the Guacamole server using the *Guacamole Protocol* over HTTP.
- The user selects the device the want to connect to, and this request is forwarded the request to *guacd*.
- *guacd* facilitates connecting to the desktop selected by the user.

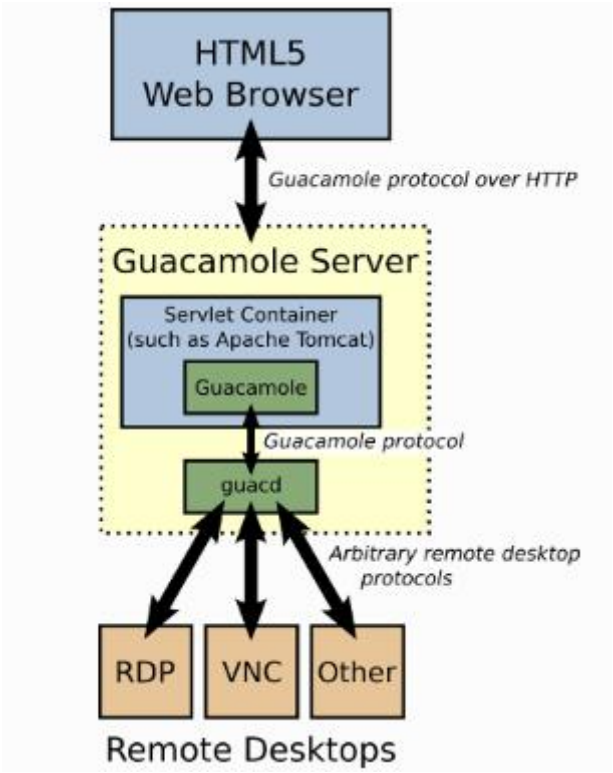

Figure 13: Guacamole Architecture

The Guacamole Protocol

The protocol consists of a series of instructions that form a message that is sent from the client to the server. The instructions relate to the device we are attempting to connect to and requests for screen size, audio and mouse and keyboard control. During the handshake phase, the client machine responds with the requested information. [42]

Guacd

Guacd is the server proxy that is used by the Guacamole web Application.[43]

[42] (Guacamole, 2022)

[43] (Guacamole, 2022)

MISP

MISP or the Malware Information Sharing Platform is an open-source project that falls into Threat Intelligence Platform category of security tools.

MISP gathers, stores and correlates IOCs from well-known attacks, threat intelligence feeds, indicators of financial fraud, and many other sources. The tool is designed to be used by security analysts across the world and to enable the sharing of information more efficiently.

The project has been funded by the European Union through their “Connecting Europe Facility” and also by the Computer Incident Response Centre Luxembourg (CIRCL). The project has many developers from across the world, including NATO.

MISP has many features. Some of the most important are listed below:

- An IOC database, which allows an organization to store data relating to malware sample, incidents, attackers, and intelligence.
- Automatic correlation between attributes and indicators from malware, attack campaigns and log analysis.
- A built-in sharing feature, which allows syncing of events across multiple MISP instances. This can also be used to share data with other organizations using MISP.
- A GUI is provided so that the user can navigate the results easily. The user is able to view the relationship between events, attributes, and objects. These relationships can be displayed visually, so the user may get a clearer picture.
- It is possible to export data in a variety of different formats which allow MISP to be integrated with many different tools.

Threat intelligence has become such an integral part of how an organization detects and prevents attacks.

MITRE ATT&CK

The MITRE ATTACK framework is a database that tracks Tactics, Techniques and Procedures that are used during an attack. This behaviour is tracked to provide analysis on adversary behaviour.[44] The tactics are categorised into the different phases of an attack, which are as follows:

- Reconnaissance
- Resource Development
- Initial Access
- Execution
- Persistence
- Privilege Escalation
- Defense Evasion
- Credential Access
- Discovery
- Lateral Movement
- Collection
- Command and Control
- Exfiltration
- Impact

The MITRE ATTACK framework has many use cases in emulating adversary behaviour, red teaming exercises, threat hunting, alert creation and much more. Each technique is broken down into its associated tactics, mitigations and detections and sections include where the technique has been observed out in the wild.[45]

[44] (Strom, et al., 2020)

[45] (Strom, et al., 2020)

Lab Exercises

Research was performed on the top 10 techniques that were exploited in 2021 and published by Red Canary in their 2022 Threat Detection Report.[46] This was going to form the basis of the labs that were to be created.

Top 10 Techniques:

- T1059: Command and Scripting Interpreter (53.4%)
- T1218: Signed Binary Proxy Execution (34.8%)
- T1047: Windows Management Instrumentation (15.4%)
- T1003: Credential Dumping (18.3%)
- T1105: Ingress Tool Transfer (20.4%)
- T1055: Process Injection (21.7%)
- T1053: Scheduled Task/Job (14.7)
- T1027: Obfuscated Files or Information (19.4%)
- T1036: Masquerading (22.1%)
- T1574: Hijack Execution Flow (8.4%)

The lab exercises I hope to create will contain theory section from five of the above techniques and include a testing part of the lab that can be Implemented and testing using the Atomic Red Team tests which are mapped to the MITRE ATTACK framework.

[46] (Red Canary, 2022)

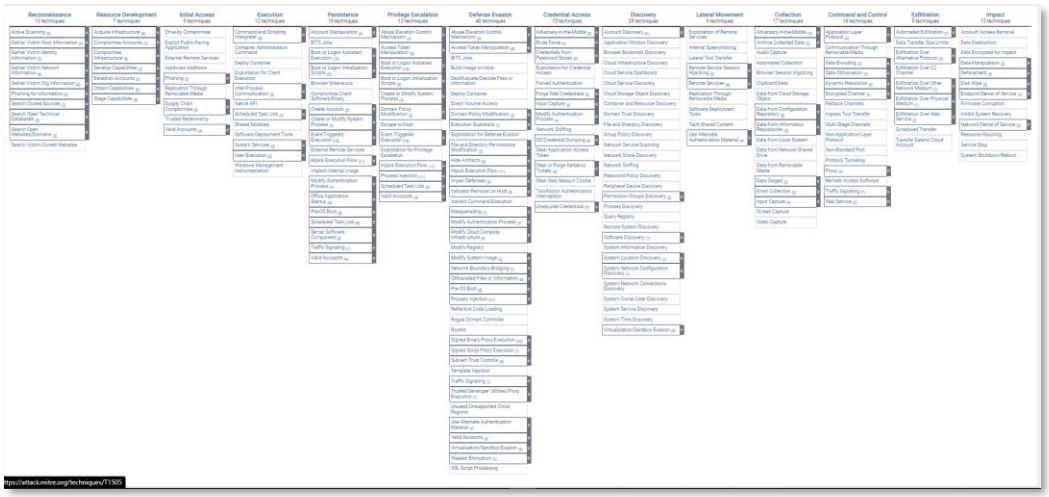

Figure 14; Mitre Attack Framework[47]

Osquery

Osquery was created by Facebook in 2014 as an open-source project to monitor security across Mac and Linux devices. The Linux Foundation has managed osquery since the stewardship was handed over to them in 2019. [48]

OSquery is a tool that allows a user to query an operating system by writing SQL queries. The tool is cross platform, and functions on Windows, Linux, MacOS and FreeBSD and it comes with native packages for all operating systems which makes deploying this across your infrastructure a little easier. [49]

An operating system on a device is broken down into individual components and stored in separate tables in order to facilitate using SQL to perform queries. For basic usage, the *'SELECT'* feature of SQL is the only command that will work to query the operating system. [50]

[47] (MITRE ATTACK, 2022)

[48] (Knowles, 2020)

[49] (osquery, 2022)

[50] (osquery, 2022)

Fleet

Fleet is the most commonly used osquery manager. It provides a GUI that allows you to manage your devices through a web browser. The GUI allows you to group endpoints based on attributes from queries.

It is possible to scale this up, depending on the devices in your organisation. According to the documentation, stress testing has been performed with up to 150,000 hosts online at one time. [51]

Siemplify

Siemplify is an Israeli security firm formed in 2015. The organisation was formed with the goal of developing an efficient method of managing security operations.[52] Their primary product is the Siemplify Security Operations Platform, one of the world's leading SOAR tools. The tool is vendor-agnostic and can be integrated with the most common security tools that are used in industry today.

The Siemplify SOAR platform automates many tasks, reducing the need for analysts to dig through 1000s of logs to find the relevant logs for the alert. The tools *threat-centric* approach to alert investigation groups related logs into the one case. This reduces the time an analyst spends searching and allows them an opportunity to perform a more thorough analysis. The tool also come pre-packaged with playbooks to allow your business to quickly deploy them to assist with automation and response. Further playbooks can be built by engineers as your business requires.[53]

There are two versions of the product available:

- Community Edition (Free)
- Enterprise Edition (Paid)

[51] (fleet, 2022)

[52] (Siemplify, 2022)

[53] (Siemplify, 2022)

	Community Edition	Enterprise Edition
IDE	✔	✔
Investigator	✔	✔
Integrations	✔	✔
Crisis Management	✘	✔
Remote Agent	✘	✔
BI	✘	✔
HA/DR	✘	✔
Multi-Tenancy	✘	✔
Users	1	Multiple
Alerts per day	25	Unlimited
Case history	90 days	Unlimited
Playbooks	5	Unlimited
Support	Online Community	24/7 Dedicated Support

Figure 15: Siemplify Product Offerings

Siemplify hosted a webinar that discusses the creation of an open-source SOC which provided the inspiration for this project. These are alternatives to the paid solutions offered by many companies and with the right team of engineers, this would make a great addition to a training lab in the college where students could gain extended knowledge of how the tools operate.[54]

[54] (Loos, 2022)

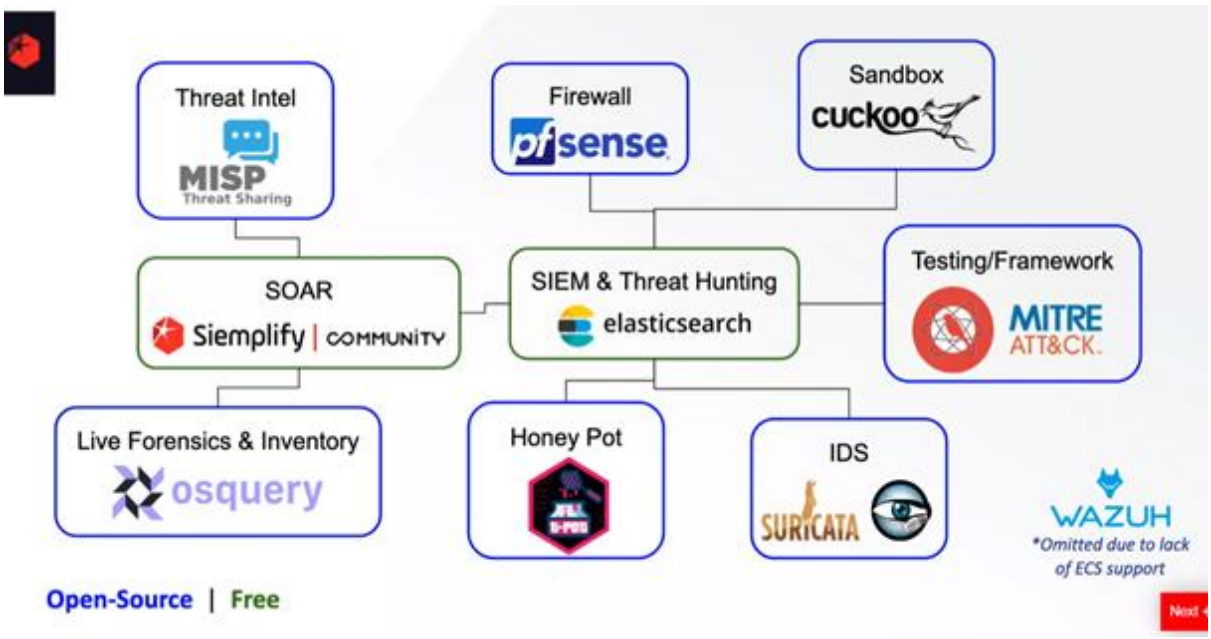

Figure 16: Slide from Webinar

Suricata

Suricata is a high-performance network intrusion detection **(IDS)**, intrusion prevention **(IPS)**, and network security monitoring engine. It is open source and is controlled by the Open Information Security **Foundation** (OISF), a non-profit community-run foundation. The OISF is responsible for the development of Suricata.[55]

Suricata is a rule based IPS/IDS that monitors network traffic and warns the system administrator when suspicious occurrences occur using externally generated rule sets and is designed to function alongside current network security components.[56]

This tool is recommended for use in both the Siemplify Open-Source SOC and is included with the Detection Lab. This is a valuable tool for Network Analysis and I would certainly recommend it's use.

Splunk

Splunk are a cybersecurity company that was founded in 2003 and provides a number of different services across that field. Their primary tool they offer is a SIEM tool which is

[55] (Suricata, 2022)

[56] (Redmine, 2022)

available in two offerings: Splunk Cloud and Splunk Enterprise. Their software also boasts the fact that it has over 2400+ Splunk apps that can readily and easily be integrated with the other security tools that your business operates to ensure that you gain full visibility into your infrastructure. Splunk is one of the market leaders when it comes to SIEM tool offerings.[57]

The Splunk software is a tool that can be used to search and analyze the data that your business produces on a daily basis. Splunk provide ingesting and storage services as well as log analysis, visualisations and much more. The pricing for Splunk Enterprise is generally based off of the number of logs that are ingested. However, the service provides the ability for your business to scale up its data needs as required. Splunk Enterprise provide a free service that allows 500MB of ingestion for free each day.[58]

Tines

Tines is an Irish start-up founded in 2018.[59] The founders have spent many years working within the cybersecurity field before creating this company. Tines is an automation platform that allows anyone, regardless of skill level, to automate any manual task. There is no need for apps, plugins, or custom code. The tool presents the automation tasks in storyboards, and these can be reused across your business as needed. This seems like a very interesting tool and I would certainly be interested in using.

Vagrant

Vagrant is a tool used to easily configure virtual machines and enables a user to consistently deploy the machine. Vagrant boxes can be stored in the cloud, which allows you to deploy your virtual environment from anywhere in the world as long as you have access to the internet and the virtual machine repository that contains your configuration files.[60]

The Vagrantfile contains all the relevant configuration information required to deploy the machine. This vagrant file is written in the Ruby programming language. Virtual machines

[57] (Splunk, 2022)

[58] (Splunk, 2022)

[59] (Tines, 2022)

[60] (Vagrant, 2022)

are deployed via the command line and it executes the vagrant file, including any scripts that are necessary for the device's configuration.[61]

Detection Lab uses Vagrant to deploy the virtual machines used for that lab.

Velociraptor

Velociraptor is an open-source endpoint monitoring, digital forensic, and cyber response platform. It was created by professionals in Digital Forensic and Incident Response (DFIR) who needed an efficient way to hunt for any tracks an adversary may have left on an endpoint.[62]

Velociraptor operates using Velociraptor Query Language (VQL), which is a framework developed to allow the efficient collection and monitoring of devices across a network. VQL was influenced by SQL and it uses the same basic syntax.[63]

This would be an excellent addition to my project, as it would provide much needed telemetry into the endpoint device.

Zeek

Zeek is an open-source network traffic analyzer that operates passively. Zeek is widely used by operators as a network security monitor (NSM) to aid in the investigation of suspicious or malicious behaviour. Zeek also offers a wide range of traffic analysis functions, such as performance monitoring and troubleshooting.

Zeek collects a vast amount of log data including all HTTP sessions with their requested URIs, key headers, MIME types, and server responses; DNS requests with replies; SSL certificates; key content of SMTP sessions. This list is not exhaustive. The log data is stored as a JSON file which enables processing by a variety of software including SIEM tools.[64]

[61] (Vagrant, 2022)

[62] (Velociraptor , 2022)

[63] (Velociraptor, 2022)

[64] (Zeek, 2022)

Zeek was also included in the Open-Source SOC and is packaged with the Detection Lab. I feel this would be a valuable tool to include in any demonstration due to the significance of network logs throughout an investigation into an alert.

Conclusion

Glossary

Agent: A program deployed to a system that performs various actions continuously and autonomously on behalf of an individual or an organization.

Amazon S3: Amazon S3 is object storage built to store and retrieve any amount of data from anywhere. It's a simple storage service that offers industry leading durability, availability, performance, security, and virtually unlimited scalability at very low costs.

API: Application Programming Interface

Big Data Architecture: It is a system used to manage large amounts of data so that it can be analyzed for business purposes, steer data analytics, and provide an environment in which big data analytics tools can extract vital business information from otherwise ambiguous data.

CSOC: Cyber Security Operations Center

Data Aggregation: Data Aggregation refers to the process by which large amounts of data is compiled and presented in a human readable format.

Data enrichment: Data enrichment refers to the process of appending or otherwise enhancing collected data with relevant context obtained from additional sources.

Data Lake: A data lake is a centralized repository that allows you to store all your structured and unstructured data at any scale. You can store your data as-is, without having to first structure the data, and run different types of analytics—from dashboards and visualizations to big data processing, real-time analytics, and machine learning to guide better decisions.

DFIR: Digital Forensic and Incident Response

DMZ: Demilitarized Zone

IOC: Indicator of Compromise

IDS: Intrusion Detection System

IPS: Intrusion Prevention System

Normalise: Normalization consists of breaking each field of a raw event into variables and combining them into views that are relevant to security administrators.

OISF: Open Information Security Foundation

Programmatically: Programmatically is used to refer to tasks that can be done in an automated way using a computer program.

RDP: Remote Desktop Protocol

SaaS: Software as a Service

SecOps: The practice of aligning Security, IT and Development teams through a shared set of data and tools.

SEM: Security Event Management

SIM: Security Information Management

SIEM: Security Information and Event Management

SNMP: Simple Network Management Protocol

SOAR: Security Orchestration, Automation, and Response

SOC: Security Operations Center

SQL: Structured Query Language

Syslog: System Logging Protocol – a standard used to send system log messages to a server.

TTP: Tactics, Technique and Procedure

UEBA: User and Entity Behaviour Analytics

VNC: Virtual Network Computing

VQL: Velociraptor Query Language

WMI: Windows Management Instrumentation

References

Abueg, R., 2020. *Elasticsearch: What It Is, How It Works, And What It's Used For.* [Online]
Available at: https://www.knowi.com/blog/what-is-elastic-search/
[Accessed 20th April 2022].

Andrew Froehlich, W. G. N., 2021. *SOAR vs. SIEM: What's the difference?.* [Online]
Available at: https://www.techtarget.com/searchsecurity/answer/SOAR-vs-SIEM-Whats-the-difference
[Accessed 11th November 2021].

Anon., 2016. *Phantom and ESG Research Finds Companies Ignore Majority of Security Alerts.* [Online]
Available at: https://www.businesswire.com/news/home/20160315005555/en/Phantom-ESG-Research-Finds-Companies-Ignore-Majority
[Accessed 11th November 2021].

Anon., 2021. *Security Orchestration Automation and Response (SOAR) Tools and Solutions.* [Online]
Available at: https://www.rapid7.com/solutions/security-orchestration-and-automation/
[Accessed 11th November 2021].

Anon., 2022. *4 Stages of Vulnerability Management: A Process for Risk Mitigation.* [Online]
Available at: https://www.exabeam.com/information-security/vulnerability-management/
[Accessed 2nd February 2022].

Anon., 2022. *Incident Response Automation and Security Orchestration with SOAR.* [Online]
Available at: https://www.exabeam.com/siem-guide/incident-response-and-automation/
[Accessed 24th April 2022].

Apache Guacamole, 2021. *Implementation and architecture.* [Online]
Available at: https://guacamole.apache.org/doc/gug/guacamole-architecture.html
[Accessed 21st April 2022].

Apache Guacamole, 2021. *Introduction.* [Online]
Available at: https://guacamole.apache.org/doc/gug/introduction.html
[Accessed 21st April 2022].

Atomic Red Team, 2022. *Atomic Red Team.* [Online]
Available at: https://github.com/redcanaryco/atomic-red-team
[Accessed 22nd April 2022].

Atomic Red Team, 2022. *Atomic Red Team.* [Online]
Available at: https://atomicredteam.io/atomicredteam
[Accessed 20th April 2022].

Babacamp, O., 2020. *Building a sustainable SIEM use cases..* [Online]
Available at: https://medium.com/@olucampbell_64749/building-a-sustainable-siem-use-cases-8f58408db375
[Accessed 26th November 2021].

Boden, P., 2016. *The Emerging Era of Cyber Defense and Cybercrime.* [Online]
Available at: https://www.microsoft.com/security/blog/2016/01/27/the-emerging-era-of-cyber-defense-and-cybercrime/
[Accessed 11th November 2021].

Brook, C., 2020. *What is Security Orchestration?.* [Online]
Available at: https://digitalguardian.com/blog/what-security-orchestration
[Accessed 11th November 2021].

BROOKS, K. J., 2021. *U.S. has almost 500,000 job openings in Cybersecurity.* [Online]
Available at: https://www.cbsnews.com/news/cybersecurity-job-openings-united-states/
[Accessed 11th November 2021].

Check Point Software Technologies Ltd., 2021. *Cyber Attack Trends: Mid Year Report 2021,* Tel Aviv: Check Point Software Technologies Ltd..

Check Point, 2021. *What is a Security Operations Center (SOC)?.* [Online]
Available at: https://www.checkpoint.com/cyber-hub/threat-prevention/what-is-soc/
[Accessed 11th November 2021].

Chef, 2022. *bento.* [Online]
Available at: https://github.com/chef/bento
[Accessed 21st April 2022].

Chrisander, M., 2020. *What is Cyber Threat Hunting? A simple guide to Threat Hunting.* [Online]
Available at: https://www.logpoint.com/en/blog/threat-hunting/
[Accessed 26th November 2021].

clong, 2022. *Detection Lab.* [Online]
Available at: https://github.com/clong/detectionlab
[Accessed 20th April 2022].

CompTIA, 2021. *What Is a Security Operations Center?.* [Online]
Available at: https://www.comptia.org/content/articles/what-is-a-security-operations-center
[Accessed 11th November 2021].

Crowdstrike, 2022. *VULNERABILITY MANAGEMENT?.* [Online]
Available at: https://www.crowdstrike.com/cybersecurity-101/vulnerability-management/

Cybersecurity & Infrastructure Security Agency (CISA), 2021. *DEFINING INSIDER THREATS.* [Online]
Available at: https://www.cisa.gov/defining-insider-threats
[Accessed 26th November 2021].

Cyware, 2020. *SOAR Use Cases.* [Online]
Available at: https://cyware.com/educational-guides/security-orchestration-automation-and-response/soar-use-cases-c670
[Accessed 24th April 2022].

Detection Lab, 2022. *INTRODUCTION.* [Online]
Available at: https://detectionlab.network/introduction/
[Accessed 21st April 2022].

Detection Lab, 2022. *PREREQUISITES.* [Online]
Available at: https://detectionlab.network/introduction/prerequisites/
[Accessed 20th April 2022].

Elastic, 2022. *Data In: Documents and Indices.* [Online]
Available at: https://www.elastic.co/guide/en/elasticsearch/reference/current/documents-indices.html
[Accessed 20th April 2022].

Elastic, 2022. *Elastic pricing.* [Online]
Available at: https://www.elastic.co/pricing/
[Accessed 20th April 2022].

Elastic, 2022. *Information Out: Search and Analyze.* [Online]
Available at: https://www.elastic.co/guide/en/elasticsearch/reference/current/search-analyze.html
[Accessed 20th April 2022].

Elastic, 2022. *Kibana - Your Window into Elastic.* [Online]
Available at: https://www.elastic.co/guide/en/kibana/current/introduction.html
[Accessed 20th April 2022].

Elastic, 2022. *Logstash Introduction.* [Online]
Available at: https://www.elastic.co/guide/en/logstash/current/introduction.html
[Accessed 20th April 2022].

Elastic, 2022. *Pricing the Elastic way.* [Online]
Available at: https://www.elastic.co/pricing/philosophy
[Accessed 20th April 2022].

Elastic, 2022. *What are Beats?.* [Online]
Available at: https://www.elastic.co/guide/en/beats/libbeat/current/beats-reference.html
[Accessed 20th April 2022].

Elastic, 2022. *What is Elasticsearch?.* [Online]
Available at:

https://www.elastic.co/guide/en/elasticsearch/reference/current/elasticsearch-intro.html
[Accessed 20th April 2022].

Exabeam, 2021. *SIEM Use Cases in a Modern Threat Landscape.* [Online]
Available at: https://www.exabeam.com/siem-guide/siem-use-cases/
[Accessed 26th November 2021].

Exabeam, 2021. *What is SIEM?.* [Online]
Available at: https://www.exabeam.com/siem-guide/what-is-siem/
[Accessed 26th November 2021].

FireEye, 2021. *What is SIEM and how does it work?.* [Online]
Available at: https://www.fireeye.com/products/helix/what-is-siem-and-how-does-it-work.html
[Accessed 26th November 2021].

FireEye, 2022. *What is SIEM and how does it work?.* [Online]
Available at: https://www.fireeye.com/products/helix/what-is-siem-and-how-does-it-work.html

FireEye, 2022. *What is SOAR?.* [Online]
Available at: https://www.fireeye.com/products/helix/what-is-soar.html

Fireeye, 2022. *What is UEBA?.* [Online]
Available at: https://www.fireeye.com/products/helix/what-is-ueba.html
[Accessed 24th April 2022].

fleet, 2022. *Using Fleet FAQ.* [Online]
Available at: https://fleetdm.com/docs/using-fleet/faq
[Accessed 20th April 2022].

Gast, K., 2021. *What is SIEM? And How Does it Work?.* [Online]
Available at: https://logrhythm.com/what-is-siem/
[Accessed 26th November 2021].

Gast, K., 2021. *What is SIEM? And How Does it Work?.* [Online]
Available at: https://logrhythm.com/blog/what-is-siem/
[Accessed 11th November 2021].

Glennon, N., 2022. *Clicks over bricks: The surge in businesses moving online.* [Online]
Available at: https://www.irishexaminer.com/business/economy/arid-40071515.html

Gonzalez, C., 2022. *User and Entity Behavior Analytics.* [Online]
Available at: https://www.exabeam.com/ueba/user-and-entity-behavior-analytics/
[Accessed 24th April 2022].

Guacamole, 2022. *guacamole/guacd.* [Online]
Available at: https://hub.docker.com/r/guacamole/guacd
[Accessed 24th April 2022].

Guacamole, 2022. *The Guacamole protocol.* [Online]
Available at: https://guacamole.apache.org/doc/gug/guacamole-protocol.html
[Accessed 24th April 2022].

Hanna, K. T. & Posey, B., 2021. *Definition: Insider Threat.* [Online]
Available at: https://www.techtarget.com/searchsecurity/definition/insider-threat
[Accessed 26th November 2021].

Hartong, O., 2018. *Endpoint detection Superpowers on the cheap — part 1.* [Online]
Available at: https://medium.com/@olafhartong/endpoint-detection-superpowers-on-the-cheap-part-1-e9c28201ac47

Hartong, O., 2022. *sysmon-modular.* [Online]
Available at: https://github.com/olafhartong/sysmon-modular

IBM Security, 2020. *Cost of a Data Breach,* New York: Ponemon Institute.

IBM, 2021. *What is SIEM?.* [Online]
Available at: https://www.ibm.com/topics/siem
[Accessed 26th November 2021].

IBM, 2021. *What is threat hunting?.* [Online]
Available at: https://www.ibm.com/topics/threat-hunting
[Accessed 26th November 2021].

Imperva, 2021. *PCI DSS Certification.* [Online]
Available at: https://www.imperva.com/learn/data-security/pci-dss-certification/
[Accessed 26th November 2021].

Imperva, 2021. *Security information and event management (SIEM).* [Online]
Available at: https://www.imperva.com/learn/application-security/siem/
[Accessed 26th November 2021].

Intersoft Consulting, 2022. *Art. 5 GDPR: Principles relating to processing of personal data.* [Online]
Available at: https://gdpr-info.eu/art-5-gdpr/

Jessica Groopman, K. I., 2021. *Top 6 SOAR use cases to implement in enterprise SOCs.* [Online]
Available at: https://www.techtarget.com/searchsecurity/tip/Top-6-SOAR-uses-cases-to-implement-in-enterprise-SOCs
[Accessed 11th November 2021].

Knowles, M., 2020. *Osquery: What It Is, How It Works, and How To Use It.* [Online]
Available at: https://www.uptycs.com/blog/osquery-what-it-is-how-it-works-and-how-to-use-it
[Accessed 20th April 2022].

Kumar, A., 2021. *Key Architectural Components of a Data Lake.* [Online]
Available at: https://vitalflux.com/key-architectural-components-of-a-data-lake/
[Accessed 11th November 2021].

Leaseweb, 2021. *Managed Cyber Security - SIEM SOC.* [Online]
Available at: https://www.leaseweb.com/cyber-security/managed-cyber-security
[Accessed 26th November 2021].

Logpoint, 2021. *Top 10 SIEM use cases to implement.* [Online]
Available at: https://www.logpoint.com/en/understand/top-10-use-cases-implement/
[Accessed 11th November 2021].

Long, C., 2017. *Introducing: Detection Lab.* [Online]
Available at: https://medium.com/@clong/introducing-detection-lab-61db34bed6ae
[Accessed 20th April 2022].

Long, C., 2022. *Chris Long.* [Online]
Available at: https://www.linkedin.com/in/chris-long-4057b410/
[Accessed 20th April 2022].

Long, C., 2022. *DetectionLab.* [Online]
Available at: https://github.com/clong/detectionlab
[Accessed 22nd April 2022].

Loos, A., 2022. *Building an Open-Source Security Operations Center.* [Online]
Available at: https://www.siemplify.co/resources/webinar-building-an-open-source-security-operations-center-watch/
[Accessed 24th April 2022].

MacDonald, M. R., Pike, N. D. & Pike, R. E., 2020. *Exploring Depth in Cybersecurity Education,* California: Cal Poly Pomona.

Miller, J., 2022. *WHAT IS SIEM AND MSSP AND HOW ARE THEY RELATED.* [Online]
Available at: https://www.bitlyft.com/resources/what-is-siem-and-mssp-and-how-are-they-related

MITRE ATTACK, 2020. *Ingress Tool Transfer.* [Online]
Available at: https://attack.mitre.org/techniques/T1105/
[Accessed 25th April 2022].

MITRE ATTACK, 2021. *Command and Scripting Interpreter.* [Online]
Available at: https://attack.mitre.org/techniques/T1059/
[Accessed 25th April 2022].

MITRE ATTACK, 2021. *Obfuscated Files or Information.* [Online]
Available at: https://attack.mitre.org/techniques/T1027/
[Accessed 25th April 2022].

MITRE ATTACK, 2021. *Windows Management Instrumentation.* [Online]
Available at: https://attack.mitre.org/techniques/T1047/
[Accessed 25th April 2022].

MITRE ATTACK, 2022. *ATTACK.* [Online]
Available at: https://attack.mitre.org/
[Accessed 24th April 2022].

MITRE ATTACK, 2022. *Signed Binary Proxy Execution.* [Online]
Available at: https://attack.mitre.org/techniques/T1218/
[Accessed 25th April 2022].

Morgan, S., 2020. *Cybercrime To Cost The World $10.5 Trillion Annually By 2025.* [Online]
Available at: https://cybersecurityventures.com/cybercrime-damages-6-trillion-by-2021/
[Accessed 26th November 2021].

Multibit, 2022. *Multibit-Legacy / multibit.* [Online]
Available at: https://github.com/Multibit-Legacy/multibit

Oktavianto, D., 2013. *Cuckoo Malware Analysis.* [Online]
Available at: https://www.oreilly.com/library/view/cuckoo-malware-analysis/9781782169239/ch01s05.html
[Accessed 26th November 2021].

OmniSci, 2021. *Big Data Architecture.* [Online]
Available at: https://www.omnisci.com/technical-glossary/big-data-architecture
[Accessed 26th November 2021].

osquery, 2022. *SQL Introduction.* [Online]
Available at: https://osquery.readthedocs.io/en/stable/introduction/sql/
[Accessed 20th April 2022].

osquery, 2022. *Welcome to osquery.* [Online]
Available at: https://osquery.readthedocs.io/en/stable/
[Accessed 20th April 2022].

Palantir, 2022. *Windows Event Forwarding Guidance.* [Online]
Available at: https://github.com/palantir/windows-event-forwarding
[Accessed 22nd April 2022].

Petters, J., 2020. *What is SIEM? A Beginner's Guide.* [Online]
Available at: https://www.varonis.com/blog/what-is-siem/
[Accessed 26th November 2021].

Ponemon Institute LLC, 2020. *Cost of a Data Breach Report,* Traverse City, Michigan : IBM Security.

Rapid7, 2016. *What is Security Orchestration?.* [Online]
Available at: https://www.rapid7.com/blog/post/2016/04/05/what-is-security-orchestration/
[Accessed 11th November 2021].

Rapid7, 2021. *Security Automation.* [Online]
Available at: https://www.rapid7.com/fundamentals/security-automation/
[Accessed 20th November 2021].

Rapid7, 2021. *Security Orchestration.* [Online]
Available at: https://www.rapid7.com/fundamentals/security-orchestration/
[Accessed 20th November 2021].

Rapid7, 2022. *User and Entity Behavior Analytics (UEBA).* [Online]
Available at: https://www.rapid7.com/fundamentals/user-behavior-analytics/
[Accessed 24th April 2022].

Recorded Future, 2019. *9 SOAR Use Cases for Effectively Mitigating Cyber Threats (Part 2).* [Online]
Available at: https://www.recordedfuture.com/soar-use-cases/
[Accessed 11th November 2021].

Red Canary, 2022. *2022 Threat Detection Report,* s.l.: Red Canary.

Red Canary, 2022. *2022 Threat Detection Report.* [Online]
Available at: https://resource.redcanary.com/rs/003-YRU-

314/images/2022_ThreatDetectionReport_RedCanary.pdf

[Accessed 22nd April 2022].

Red Canary, 2022. *Atomic Red Team.* [Online]

Available at: https://redcanary.com/atomic-red-team/

[Accessed 20th April 2022].

Redmine, 2022. *What is Suricata.* [Online]

Available at:

https://redmine.openinfosecfoundation.org/projects/suricata/wiki/What_is_Suricata

[Accessed 24th April 2022].

Robbins, J., 2009. *Announcing Chef.* [Online]

Available at: https://www.chef.io/blog/announcing-chef

[Accessed 21st April 2022].

Roberts, C., 2022. *Home.* [Online]

Available at: https://github.com/redcanaryco/invoke-atomicredteam/wiki

[Accessed 20th April 2022].

Romero, A., 2020. *How to Build Security Use Cases for Your SIEM.* [Online]

Available at: https://logrhythm.com/blog/how-to-build-security-use-cases-for-your-siem/

[Accessed 11th November 2021].

Shoffner, M., 2022. *SIEM Cost Breakdown and Tips.* [Online]

Available at: https://www.peerspot.com/articles/siem-cost-breakdown-and-tips

Siemplify, 2021. *Security Orchestration & Automation.* [Online]

Available at: https://www.siemplify.co/security-orchestration-automation/

[Accessed 11th November 2021].

Siemplify, 2022. *About Us.* [Online]

Available at: https://www.siemplify.co/about-us/#

[Accessed 24th April 2022].

Siemplify, 2022. *Siemplify Security Operations Platform for Enterprises.* [Online]

Available at: https://www.siemplify.co/enterprise/

[Accessed 24th April 2022].

Siemplify, 2022. *The Siemplify Security, Orchestration, Automation and Response (SOAR) Platform.* [Online]
Available at: https://www.siemplify.co/soar-platform-overview/
[Accessed 24th April 2022].

Slintel, 2022. *Market Share of Splunk.* [Online]
Available at: https://www.slintel.com/tech/security-information-and-event-management-siem/splunk-market-share
[Accessed 22nd April 2022].

Sobers, R., 2021. *134 Cybersecurity Statistics and Trends for 2021.* [Online]
Available at: https://www.varonis.com/blog/cybersecurity-statistics
[Accessed 11th November 2021].

Splunk, 2021. *What Is a Security Operations Center (SOC)?.* [Online]
Available at: https://www.splunk.com/en_us/data-insider/what-is-a-security-operations-center.html
[Accessed 11th November 2021].

Splunk, 2021. *What Is Security Automation?.* [Online]
Available at: https://www.splunk.com/en_us/data-insider/what-is-security-automation.html
[Accessed 20th November 2021].

Splunk, 2021. *What Is SOAR?.* [Online]
Available at: https://www.splunk.com/en_us/data-insider/what-is-soar.html
[Accessed 20th November 2021].

Splunk, 2022. *About Splunk.* [Online]
Available at: https://www.splunk.com/en_us/about-splunk.html
[Accessed 24th# April 2022].

Splunk, 2022. *Choose Volume or Compute-Based Pricing Plans for Ultimate Flexibility.* [Online]
Available at: https://www.splunk.com/en_us/software/pricing/enterprise-platform.html
[Accessed 24th Aptil 2022].

Splunk, 2022. *Splunk Enterprise Security.* [Online]
Available at: https://splunkbase.splunk.com/app/263/#/overview
[Accessed 24th April 2022].

Splunk, 2022. *What Is SIEM?.* [Online]
Available at: https://www.splunk.com/en_us/data-insider/what-is-siem.html

Stefanova, D., 2022. *Using SIEM for Regulatory Compliance: Importance, Best Practices, Use Cases.* [Online]
Available at: https://logsentinel.com/blog/using-siem-for-regulatory-compliance-importance-best-practices-use-cases/

Strom, B. E. et al., 2020. *MITRE ATT&CK: Design and Philosophy,* s.l.: The MITRE Corporation.
.

Suricata, 2022. *What is Suricata.* [Online]
Available at: https://suricata.readthedocs.io/en/suricata-6.0.0/what-is-suricata.html
[Accessed 24th April 2022].

Taschler, S., 2018. *What is Cyber Threat Hunting?.* [Online]
Available at: https://www.crowdstrike.com/cybersecurity-101/threat-hunting/
[Accessed 13th November 2021].

Tines, 2022. *Tines.* [Online]
Available at: https://www.linkedin.com/company/tines-io/about/
[Accessed 24th April 2022].

Unum, 2021. *Splunk Senior Infrastructure Engineer.* [Online]
Available at: https://www.irishjobs.ie/Jobs/Splunk-Senior-Infrastructure-Engineer-8648280.aspx
[Accessed 22nd April 2022].

Vagrant, 2022. *Introduction to Vagrant.* [Online]
Available at: https://www.vagrantup.com/intro
[Accessed 24th April 2022].

Vagrant, 2022. *Vagrantfile.* [Online]
Available at: https://www.vagrantup.com/docs/vagrantfile
[Accessed 24th April 2022].

Velociraptor , 2022. *Velociraptor Overview.* [Online]
Available at: https://docs.velociraptor.app/docs/overview/
[Accessed 24th April 2022].

Velociraptor, 2022. *VQL Fundamentals.* [Online]
Available at: https://docs.velociraptor.app/docs/vql/
[Accessed 24th April 2022].

Williams, A. T. & Nicolett, M., 2005. *Improve IT Security With Vulnerability,* Stamford, CT: Gartner.

Zeek, 2022. *About Zeek.* [Online]
Available at: https://docs.zeek.org/en/master/about.html#what-is-zeek
[Accessed 24th April 2022].

Zurkus, K., 2018. *174,000 Alerts per Week Besiege Security Teams.* [Online]
Available at: https://www.infosecurity-magazine.com/news/174000-alerts-per-week-besiege/
[Accessed 20th November 2021].