Aki Haakila

Implementing Security Monitoring at Small and Medium sized Businesses

Metropolia University of Applied Sciences

Master of Engineering

Information Technology

Master's Thesis

31 May 2022

Abstract

Author:	Aki Haakila
Title:	Implementing Security Monitoring at Small and Medium sized Businesses
Number of Pages:	49 pages
Date:	31 May 2022
Degree:	Master of Engineering
Degree Programme:	Information Technology
Professional Major:	Networking and Services
Supervisors:	Kimmo Saurén, Principal Lecturer

The internet is not a safe place, threat actors are exploiting vulnerabilities to compromise companies and either stealing their data or demanding ransom to release encryption keys to crypto locked files. This is an all too true situation to many a company, regardless of size or reputation. The threat actors range from opportunistic script kiddies in poor nations looking to bring food to the family table to nation state groups with seemingly endless budgets and technical resources.

Luckily security monitoring is becoming more and more ubiquitous. The aim of security monitoring is to detect the first steps of an intrusion, to be able to act before the attempt becomes a breach. And if that fails, to detect the successful breach before irreversible damage can happen.

Planning and implementing a security monitoring project with no previous experience can be a hard and daunting task. Consultants and Managed Service Providers (MSP) will likely be involved, and this thesis aims to balance the knowledge gap between the SMB and MSP, so that the reader will be able to discuss the topic on an even footing and make better decisions. Resulting in a security monitoring system that will best suit the company's environment and risk appetite.

Keywords: Cybersecurity, Information Security, SIEM

Tiivistelmä

Tekijä:	Aki Haakila
Otsikko:	Implementing Security Monitoring at Small and Medium sized Businesses
Sivumäärä:	49 sivua
Aika:	31 Toukokuu 2022
Tutkinto:	Master of Engineering
Koulutusohjelma:	Tietotekniikka
Suuntautumisvaihtoehto:	Tietoverkot ja Palvelut
Ohjaaja:	Lehtori Kimmo Saurén

Internet ei ole turvallinen paikka. Uhkatekijät murtautuvat yrityksiin käyttäen hyväkseen haavoittuvuuksia ja ryöstävät näiden datan tai vaativat lunnaita kiristyshaittaohjelmien avulla. Tämä on turhankin todellinen tilanne monelle yritykselle riippumatta koosta tai maineesta. Uhkatekijöitä ovat nuoret harrastelijat jotka opportunistisesti yrittävät tienata rahaa ruokkiakseen perhettään aina valtioiden tukemiin tutkimusryhmiin joilla on lähes loputtomat resurssit käytössään.

Onneksi tietoturvamonitorointi on yleistymässä nopeasti. Parhaimmillaan Tietoturvamonitorointi auttaa huomaamaan jo ensi askeeleet tietomurrosta, antaen yritykselle mahdollisuuden reagoida ennen kuin hyökkäys onnistuu, tai ennen kuin onnistunut tietoturvamurto aiheuttaa peruuttamatonta vahinkoa.

Tietoturvamonitoroinnin rakentaminen ei ole helppoa, varsinkaan ilman aikaisempaa kokemusta jota ei välttämättä löydy Pienistä ja Keskisuurista (PK) yrityksistä. Konsulttitalot astuvat todennäköisesti mukaan kuvaan ja tämä lopputyö auttaa tasoittamaan tietotaito epä-tasapainoa PK-yrityksen ja konsulttitalon välillä, jotta lukija pystyy keskustelemaan asiasta tasavertaisesti ja tekemään parempia päätöksiä. Tavoitteena on rakentaa tietoturvamonitorointi joka parhaiten sopii yrityksen ympäristöön ja riskienhallinta tavoitteisiin.

Keywords: Kyberturvallisuus, Tietoturva, SIEM

Contents

List of Abbreviations

APT	Advanced Persistent Threat
BEC	Business Email Compromise
C&C	Command & Control
DDoS	Distributed Denial of Service
DHCP	Dynamic Host Control Policy
DNS	Domain Name System
DiD	Defence in Depth
EDR	Endpoint Detection and Response
GDPR	General Data Protection Regulation
IDS	Intrusion Detection System
IOC	Indication of Compromise
IP	Internet Protocol
IPS	Intrusion Prevention System
IT	Information Technology
IoT	Internet of Things
MAC	Media Access Control

MFA	Multi Factor Authentication
ATT&CK	Adversarial Tactics, Techniques and Common Knowledge
MSP	Managed Service Provider
MitM	Man in the Middle
NIST	National Institute of Standards and Technology
NMS	Network Management System
NSA	National Security Agency
RNA	Ribonucleic Acid
SHA	Secure Hash Algorithm
SIEM	Security Information and Event Management
SMB	Small and Midsize Business
SOC	Security Operations Centre
TI	Threat Intelligence
TLS	Transport Layer Security
UBA	User Behaviour Analytics
URL	Uniform Resource Locator
VPN	Virtual Private Network

1 Introduction

This thesis gives an overview of the main considerations when planning and implementing security monitoring. The intended audience are IT workers who are interested in security monitoring or have been tasked with implementing security monitoring at their environment, but do not have previous experience with security monitoring.

This thesis is not an attempt at a complete guide for each step, as most sub chapters are worth a book on their own. Instead, it will help the reader by giving an overview of some of the common practicalities that have to be considered, thus reducing unknown unknowns to make it easier to further familiarize oneself with security monitoring - you can't research a topic unless you have relevant keywords to use.

The thesis consists of four main sections:

- Figuring out what you are attempting to protect by doing an inventory of all assets in the environment.
- Figuring out who you are protecting yourself from by doing a threat assessment.
- Defining the sources of information that you have available and identifying the ones that are useful for security monitoring.
- Collecting the useful logs to a SIEM and properly taking the system to use

Additional sections further describe activities to consider after the security monitoring has been taken into use to effectively respond to security alerts, as well as alternatives to in house security monitoring.

2 Preparations

Before you can decide what to monitor, you need to know what you have.

It is not uncommon for even tightly run companies to gather undocumented or forgotten hardware and software assets as time goes on. Some of this is due to loss of knowledge as people change companies, and some can be an accumulation of Shadow IT. Shadow IT is the use of information technology systems, devices, software, applications, and services without explicit IT department approval [1.].

2.1 Inventory of Assets

The hardware asset list should in addition of physical devices, also include virtualized systems. Do not forget network devices like routers, firewalls etc. or IoT devices like TV’s, security cameras and fish tanks.

Listing 1 describes a security breach at a casino that happened because the fish tanks IoT water pump was not accounted for and appropriately protected.

```
A North American casino recently installed a high-tech fish
tank as a new attraction, with advanced sensors that
automatically regulate temperature, salinity, and feeding
schedules.

Anomalous activity detected:
 • Transfer of 10GB outside the network

The data was being transferred to a device in Finland
where an attacker had managed to gain control over the
tank. This was a clear case of data exfiltration, but far
more subtle than typical attempts at data theft.

By targeting an unconventional device that had recently
been introduced into the network, the attack managed to
evade the casino’s traditional security tools
```

Listing 1. [2. pp. 8]

Software asset compilation can feel like an endless quagmire, but the effort spent here will pay dividends when defining patch management and vulnerability management policies. In addition to the name of the software, the

asset list should also include the vendor, the version, and the person responsible for it. Look for ways to automate gathering and updating this information in a systematic way.

Human assets should not be forgotten. According to Deloitte: “*91% of all cyber-attacks begin with a phishing email to an unexpected victim*” [3.].

2.2 Inventory of Critical Assets

Once all the assets have been named and compiled, it is time to decide which ones are the most important. While not all non-critical assets require security monitoring, all assets defined as critical should be included. Additional monitoring can be implemented on case-by-case basis. For example, the email boxes of employees authorized to transfer money can be monitored more closely, than the email boxes of normal employees.

2.3 Network Topology

There are many different ways and tools to create network topology maps. You should try different tools and methods to find out which best suits your needs. The map should be human readable and show the way that traffic moves into and out from the organization.

After the creation of both logical and physical network topology maps, you will be able to decide which points on the network need to be added to security monitoring. The maps can also be used to create efficient containment policies to prevent lateral movement and malware outbreaks.

Find all systems that are reachable from the public internet, any vulnerabilities in hardware or software on these systems will be quickly exploited. Consider adding these systems to the list of critical assets.

3 Threat Assessment

Threat assessment attempts to figure out who you need to protect your environment from, and how to best accomplish that. Not a single company in the world has an unlimited budget, so cost/efficiency has to be considered. A local pizza place probably does not need to consider Chinese espionage as a credible threat, but a pharmaceutical company working on RNA vaccinations would very much have to consider nation states and organized cybercrime as credible threats and plan their defences accordingly.

Everyone who has a presence on the Internet has to take into account opportunistic attackers. These can range from script kiddies testing out Kali Linux to Advanced Persistent Threat (APT) groups looking to add more nodes to their botnets via automated scanning and exploitation of recently released vulnerabilities. Phishing waves sent out based on publicly available email lists and Business Email Compromise (BEC) via leaked credentials also fall into this category. Even a ransomware attack caused by an unpatched server is nothing personal, just the current state of the Internet.

Companies that have name recognition in the public will also have to account for more directed attacks. Spear phishing where the attacker devotes time to stalking the company's personnel on social media, and then creates a custom phishing email that aligns with the recipient's interests can be a highly successful way of attack. Well-known companies can also expect to have their public internet presence probed more often and more deeply, so staying on top of the vulnerability => patch whack-a-mole game is important.

Among the most targeted sectors are Defence, Finance, R&D, and Infrastructure. These sectors have to consider high interest from APT groups when planning their defences.

Once you have defined your adversaries you can start selecting the defences that best suit your needs. The industry best practise is known as Defence in

Depth (DiD) and National Institute of Standards and Technology (NIST) defines DiD as follows:

> The application of multiple countermeasures in a layered or stepwise manner to achieve security objectives. The methodology involves layering heterogeneous security technologies in the common attack vectors to ensure that attacks missed by one technology are caught by another.

[4.]

In practice, this could be described as a Firewall acting as the outmost layer of defence, an Intrusion Detection System (IDS) or an Intrusion Prevention System (IPS) as the second layer, an Endpoint Detection and Response (EDR) client as the third layer and a User Behaviour Analytics (UBA) tool as the innermost layer of defence. The number of layers and their sophistication is only limited by your budget.

But even with an almost limitless budget, technological solutions are not enough when you are dealing with APT actors. Consider what NSA did to companies ordering networking gear from Cisco:

> A document included in the trove of National Security Agency files released with Glenn Greenwald's book *No Place to Hide* details how the agency's Tailored Access Operations (TAO) unit and other NSA employees intercept servers, routers, and other network gear being shipped to organizations targeted for surveillance and install covert implant firmware onto them before they're delivered.
>
> These Trojan horse systems were described by an NSA manager as being "some of the most productive operations in TAO because they pre-position access points into hard target networks around the world."

[5.]

Even the most loyal employees can become insider threats when their or their family's safety is threatened. Figure 1 is a famous webcomic strip from XKCD, that illustrates how the bad guys are not going to give up on the first sign of proper use of technical security measures.

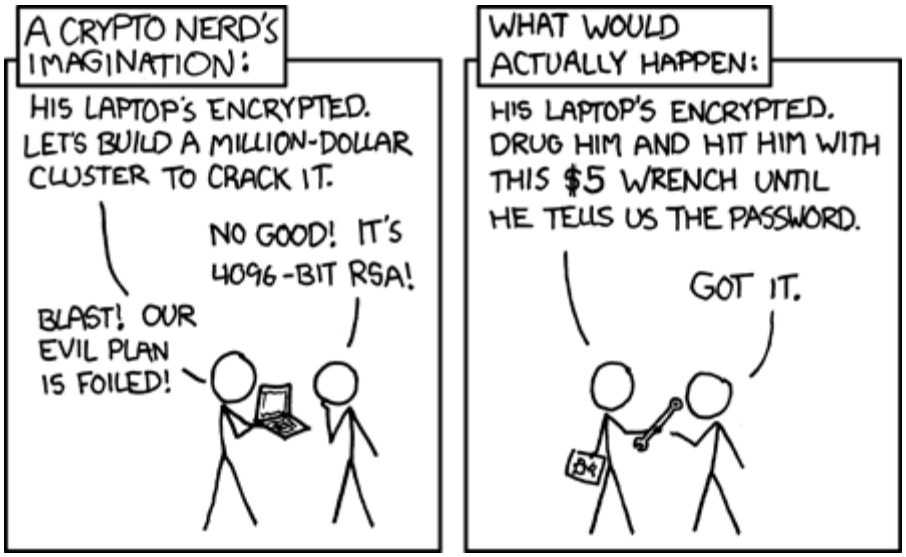

Figure 1. XKCD webcomic strip regarding security [6.].

And even if your security is flawless, there is always the Inception option: the APT actor can attempt to buy the whole company to get your data directly or buy a company that your company depends on for services to get a way inside your company to steal the data. A direct buyout like this would not have the company's future as a high priority and could be considered a national security risk, depending on the data your company has or has access to.

4 Systems to Monitor

Once the inventory of systems and threat assessment have been completed, it is time to select the log sources that will give the optimal visibility to the happenings in your environment. Every facet of the attack surface identified in the threat assessment phase should have actionable log events coming in and useless log events should be filtered out to reduce the strain on both your security monitoring tools database as well as on the security analysts investigating the alerts.

Below are descriptions of log sources that are commonly considered useful for security monitoring, as well as discussion on how to utilize them and how their events commonly look like.

4.1 Anti-Virus

Historically anti-virus programs are a common building block on computers, both at home and enterprise environments. The most basic functionality of anti-virus is to keep an updated list of hashes for known bad files. This list is then compared to the files on the computers disk via a scan, and any detections are alerted on.

Listing 2 shows the content of a sample Trend Micro alert even.

```
CEF:0|Trend Micro|Deep Security Agent|<DSA
version>|4000000|Eicar_test_file|6|cn1=1 cn1Label=Host ID dvchost=hostname
cn2=205 cn2Label=Quarantine File Size
filePath=C:\\Users\\trend\\Desktop\\eicar.exe act=Delete msg=Realtime
TrendMicroDsMalwareTarget=N/A TrendMicroDsMalwareTargetType=N/A
TrendMicroDsFileMD5=44D88612FEA8A8F36DE82E1278ABB02F>
TrendMicroDsFileSHA1=3395856CE81F2B7382DEE72602F798B642F14140
TrendMicroDsFileSHA256=275A021BBFB6489E54D471899F7DB9D1663FC695EC2FE2A2C4538AA
BF651FD0F
```

Listing 2 [7.].

The interesting fields from listing 2 are:
dvchost=hostname
cn2Label=Quarantine
filePath=C:\\Users\\trend\\Desktop\\eicar.exe
act=Delete
msg=Realtime
TrendMicroDsFileSHA256=275A021BBFB6489E54D471899F7DB9D1663FC695EC2FE2A2C4538AABF651FD0F

While the rest of the log can also be helpful, these fields are the most interesting. They have all the information required to investigate the alert.

Regardless of the Anti-virus client that you are using, you will need to dig into the documentation to review the log events that the software generates and decide which ones will be forward to the central log collector. Additionally, you should review which fields in each log event type need to be parsed for investigation purposes.

4.2 Endpoint Detection and Response

Endpoint Detection and Response (EDR) can be thought of as the next generation of anti-virus. EDR uses behavioural analysis as its base and can thus detect threats that are never written to disk or are not yet even publicly known. It is enough that the activity is identified as known bad or suspicious. A classic example of this would be a Word document that spawns a PowerShell process that then executes commands / script. There are very few instances where something like this would be considered normal.

Figure 2 show the Gartner Endpoint Protection Platform Magic Quadrant 2021, which gives a quick look on who the main players on the EDR field are.

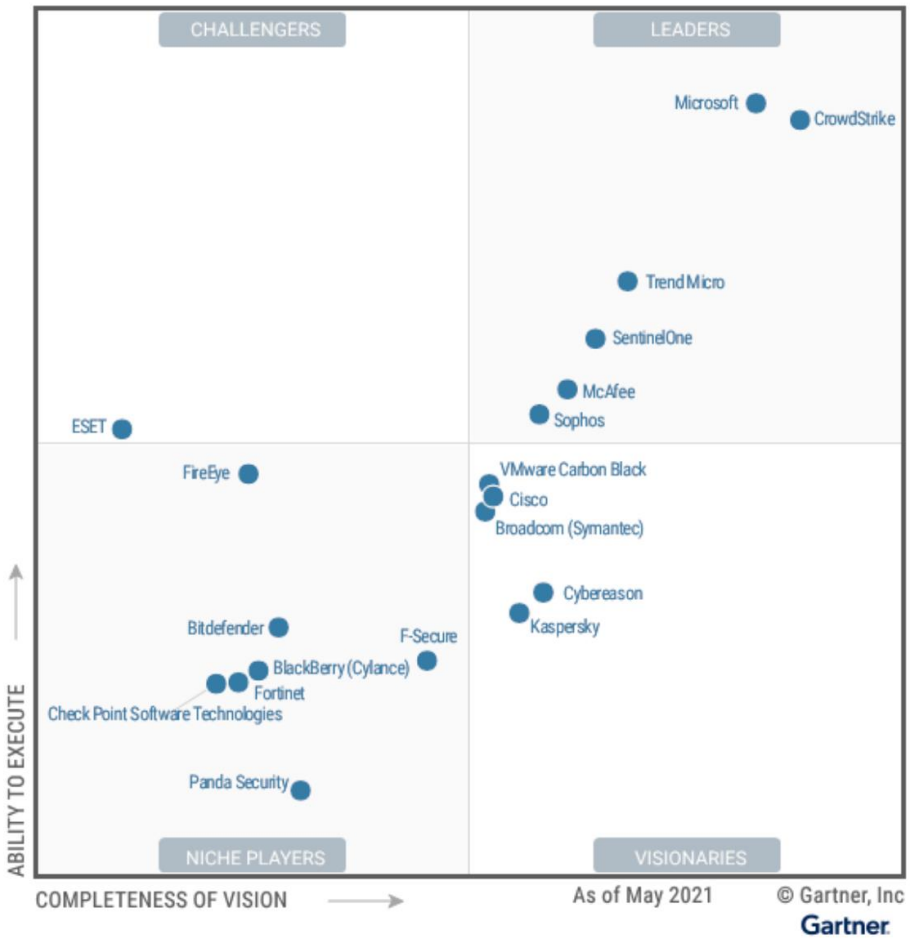

Source: Gartner (May 2021)

Figure 2. Gartner Magic Quadrant for EDR's 2021 [8.].

VMWare Carbon Black used to be the among the top choices some years ago but has since fallen to the middle of the pack due to stagnation on new features. Meanwhile both Microsoft and CrowdStrike have risen to the top.

Microsoft has the very real advantage of creating the systems around the EDR, like the Operating System and O365 cloud services. This in theory gives them a better understanding of what is happening under the hood and possibly access that is not available to other vendors. The downside of Microsoft is the usual Microsoft thing of constantly renaming and changing functionality - even whole Portals, so that you can never be quite sure where some specific view or option was.

The general weakness of EDR solutions is their ability to view, log and search network connections. Let’s say that you have malware on your host, and it is attempting to connect to “bad.domain.com”, but this is blocked by your firewall, as it already knows that this domain is bad. Due to the connection not going through, it is quite possible that this connection attempt will be invisible to the EDR. Due to this limitation, it is important to complement EDR with robust network monitoring.

Listing 3 shows us a sample CrowdStrike alert event:

```
{
   {
    "metadata": {
        "customerIDString": "xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx",
        "offset": 14947764,
        "eventType": "DetectionSummaryEvent",
        "eventCreationTime": 1536846439000,
        "version": "1.0"
    },
    "event": {
        "ProcessStartTime": 1536846339,
        "ProcessEndTime": 0,
        "ProcessId": 38684386611,
        "ParentProcessId": 38682494050,
        "ComputerName": "CS-SE-EZ64",
        "UserName": "demo",
        "DetectName": "Process Terminated",
        "DetectDescription": "Terminated a process related to the deletion of
backups, which is often indicative of ransomware activity.",
        "Severity": 4,
        "SeverityName": "High",
        "FileName": "explorer.exe",
        "FilePath": "\\Device\\HarddiskVolume1\\Windows",
        "CommandLine": "C:\\Windows\\Explorer.EXE",
        "SHA256String":
"6a671b92a69755de6fd063fcbe4ba926d83b49f78c42dbaeed8cdb6bbc57576a",
        "MD5String": "ac4c51eb24aa95b77f705ab159189e24",
        "MachineDomain": "CS-SE-EZ64",
        "FalconHostLink":
"https://falcon.crowdstrike.com/activ...xxxxxxxxxxxxx",
        "SensorId": "ec86abd353824e96765ecbe18eb4f0b4",
        "DetectId": "ldt:ec86abd353824e96765ecbe18eb4f0b4:38655257584",
        "LocalIP": "xx.xx.xx.xx",
        "MACAddress": "xx-xx-xx-xx-xx",
        "Tactic": "Malware",
        "Technique": "Ransomware",
        "Objective": "Falcon Detection Method",
        "PatternDispositionDescription": "Prevention, process killed.",
        "PatternDispositionValue": 16,
        "PatternDispositionFlags": {
            "Indicator": false,
            "Detect": false,
            "InddetMask": false,
            "SensorOnly": false,
            "Rooting": false,
            "KillProcess": true,
            "KillSubProcess": false,
            "QuarantineMachine": false,
            "QuarantineFile": false,
            "PolicyDisabled": false,
            "KillParent": false,
            "OperationBlocked": false,
            "ProcessBlocked": false
        }
    }
}
```


Listing 3 [9.].

CrowdStrike log event seen in listing 3 is a great example of a good log event. Almost all the fields are relevant to investigation, and the information is well presented. At a glance, we can see where this happened (ComputerName, UserName, LocalIP), what was attempted (DetectName, DetectDescription, Tactic, Technique), how it was attempted (CommandLine) and what the end result was (PatternDispositionDescription). Additionally, FalconHostLink is provided for quick access to the alert process tree in Falcon Portal.

In figure 3 we can see a sample investigation process tree from CrowdStrike Falcon.

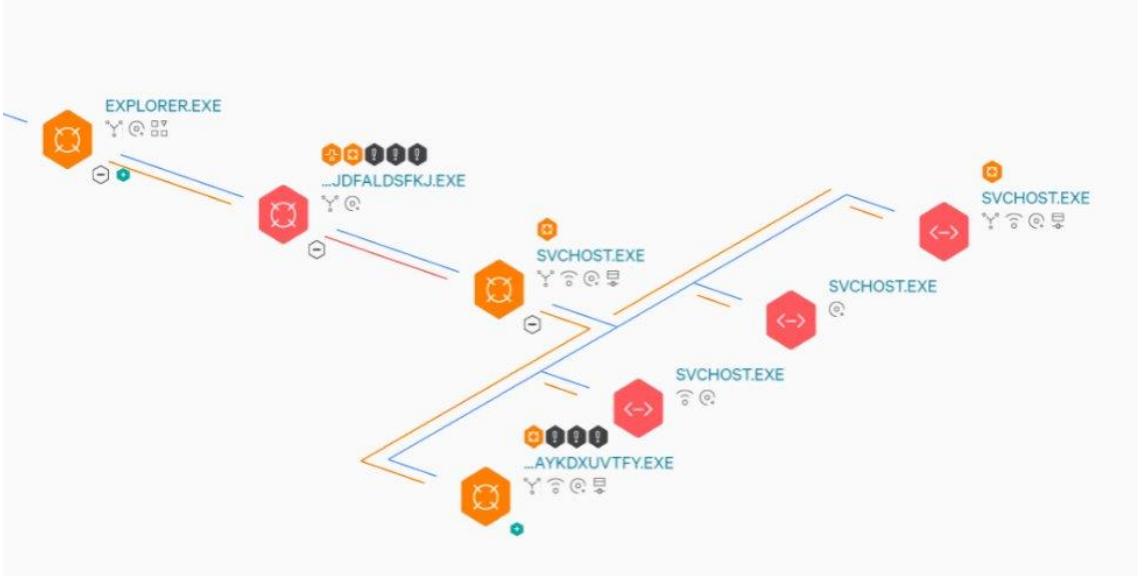

Figure 3. CrowdStrike Falcon EDR investigation process tree [10.].

One of the key things in the above log file is the amount of context information provided. Not only are the raw event details provided but the event is also mapped to the MITRE ATT&CK framework, to better understand it. MITRE ATT&CK is framework that maps attacker techniques and tactics to various phases of the cyberattack. The framework is valuable in making defensive operations more structured and science based.

4.3 Firewall

A stateful firewall is a basic building block of any network. Usually placed at the edge and protects the company’s internal network from malicious traffic originating from the internet or preventing data exfiltration and compromised internal hosts from connecting to Command & Control (C&C) servers on the Internet.

Modern firewalls have additional capabilities in detecting threats via Intrusion Detection (IDS) and Intrusion Prevention (IPS) capabilities, as well as the ability to Man-in-the-Middle (MitM) TLS connections. This is usually called TLS inspection. A firewall that does not support TLS inspection should be considered obsolete.

In figure 4 we can see the Gartner Magic Quadrant for Network Firewalls 2021.

Figure 4. Gartner Magic Quadrant for Network Firewalls 2021 [11.].

Gartner provides a good yearly overview of the firewall market, and as can be seen from figure 4, the firewall market is currently dominated by three key players: Palo Alto, Fortinet and Check Point. Cisco, Juniper, Sophos and Forcepoint (previously Stonesoft) also have products that are worth a look.

Listing 4 shows a Checkpoint Firewall log event sample:

```
Sep 3 15:10:54 192.168.99.1 Checkpoint: 3Sep2007 15:10:28 accept 192.168.99.1
>eth2 rule: 9; rule_uid: {11111111-2222-3333-8A67-F54CED606693}; service_id:
```

```
domain-udp; src: 200.14.120.9; dst: 192.168.99.184; proto: udp; product: VPN-1
& FireWall-1; service: 53; s_port: 32769;
```

Listing 4 [12.].

The firewall log shown in listing 4 shows the source IP, Port and Protocol as well as the destination IP, Port and Protocol. This can be enhanced with information regarding the matched rule and related interface.

Listing 5 shows a Cisco IDS/IPS module log event sample:

```
Sep 1 10:38:36 10.10.10.1 615: *Sep 1 17:36:34.307: %IPS-4-SIGNATURE: Sig:5123
Subsig:0 Sev:5 WWW IIS Internet Printing Overflow [192.168.100.11:59633 ->
10.10.10.10:80]
```

Listing 5 [13.].

As seen from listing 5, firewall IPS/IDS logs will usually include the matched malicious traffic signature, that can be used to find additional context information from other sources to investigate the detection.

4.4 DNS

Domain Name System (DNS) logs will give a wealth of information and should not be ignored. Compromised hosts will quite often attempt to retrieve additional payloads or get further instructions from their C&C sources. DNS log files are an excellent source of information to find these compromised hosts, assuming that your Threat Intel (TI) data is up to date.

DNS data can be logged at the client end, from network traffic via IDS or at the DNS server. If you are forwarding logs from the DNS server, make sure that the log events will allow you to identify the original source of the DNS query.

Listing 6 shows a sample Zeek IDS DNS log event:

```
{"ts":1591367999.305988,"uid":"CMdzit1AMNsmfAIiQc","id.orig_h":"192.168.4.76",
"id.orig_p":36844,"id.resp_h":"192.168.4.1","id.resp_p":53,"proto":"udp","tran
s_id":19671,"rtt":0.06685185432434082,"query":"testmyids.com","qclass":1,"qcla
ss_name":"C_INTERNET","qtype":1,"qtype_name":"A","rcode":0,"rcode_name":"NOERR
```

```
OR","AA":false,"TC":false,"RD":true,"RA":true,"Z":0,"answers":["31.3.245.133"]
,"TTLs":[3600],"rejected":false}
```


Listing 6 [14.].

From the log event shown in listing 6 the following fields are important for investigations:
id.orig_h Source IP
id.orig.p Source Port
id.resp_h DNS responder IP
id_resp_p DNS responder Port
query: DNS query
answers: DNS server response to the requester

What the above tells us is, who requested what, from where. This will allow us to quickly move to the correct source host on the EDR side, to see which process was responsible for the query, and what the full process tree for that was. Any anomalies in the DNS query and responder would also raise concerns for data exfiltration via DNS or C&C traffic.

4.5 DHCP

Dynamic Host Configuration Protocol (DHCP) logs will match an IP address to MAC address and will normally also match the IP and MAC address to a hostname.

Listing 7 shows a sample DHCP event capture from Zeek NMS:

```
04:14:39.158211 IP (tos 0x0, ttl 128, id 44415, offset 0, flags [none], proto
UDP (17), length 365)
    0.0.0.0.68 > 255.255.255.255.67: [udp sum ok] BOOTP/DHCP, Request from
3c:58:c2:2f:91:21, length 337, xid 0xfd9859a7, Flags [none] (0x0000)
          Client-Ethernet-Address 3c:58:c2:2f:91:21
          Vendor-rfc1048 Extensions
            Magic Cookie 0x63825363
            DHCP-Message Option 53, length 1: Request
            Client-ID Option 61, length 7: ether 3c:58:c2:2f:91:21
            Requested-IP Option 50, length 4: 192.168.4.152
            Server-ID Option 54, length 4: 192.168.4.1
            Hostname Option 12, length 15: "3071N0098017422"
            FQDN Option 81, length 27: "3071N0098017422.fcps.edu"
            Vendor-Class Option 60, length 8: "MSFT 5.0"
            Parameter-Request Option 55, length 14:
              Subnet-Mask, Default-Gateway, Domain-Name-Server, Domain-Name
              Router-Discovery, Static-Route, Vendor-Option, Netbios-Name-
Server
              Netbios-Node, Netbios-Scope, Option 119, Classless-Static-Route
              Classless-Static-Route-Microsoft, Option 252
            END Option 255, length 0
```

Listing 7 [15.].

When conducting an investigation using a source IP, the “timestamp”, ”Requested-IP” and ”Hostname” fields in the above log event allow us to identify the hostname during the time of interest. From here we can then pivot to EDR and investigate the host for interesting events.

4.6 VPN

VPN services have quickly become ubiquitous due to working from home becoming an accepted way of working. Working from home can also increase risk of security incidents due to home networks being less protected than the office network.

Listing 8 shows a sample PaloAlto Global Protect VPN log event:

```
Feb 17 13:58:01 server.pa01 1,2021/02/17
13:58:00,011901013191,GLOBALPROTECT,0,2305,2021/02/17 13:58:00,vsys1,gateway-
connected,connected,,IPSec,domain\user.a1,BR,NOTE01,192.168.93.210,0.0.0.0,10.
10.1.10,0.0.0.0,es11-3120-f2g9-g4e7,NOTE01,5.1.5,Windows,"Microsoft Windows 10
Pro , 64-bit",1,,,"",success,,0,,0,SSLVPN,3533509,0x0
```

Listing 8 [16.].

The Palo Alto Global Protect documentation shown in listing 9 shows that the log format is the following:

```
Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type,
FUTURE_USE, Generated Time, Virtual System, Event ID, Stage, Authentication
Method, Tunnel Type, Source User, Source Region, Machine Name, Public IP,
Public IPv6, Private IP, Private IPv6, Host ID, Serial Number, Client Version,
Client OS, Client OS Version, Repeat Count, Reason, Error, Description,
Status, Location, Login Duration, Connect Method, Error Code, Portal, Sequence
Number, Action Flags
```

Listing 9 [17.].

There are multiple interesting fields in listing 9, but the most relevant for investigation purposes are the following:

```
Source User (srcuser)      The username of the user who initiated the session.

Machine Name (machinename) The name of the user’s machine.

Public IP (public_ip)      The public IP address for the user who initiated the
session.

Private IP (private_ip)    The private IP address for the user who initiated
the session.
```

Using the above information VPN logs can be used to match an IP address to a user, making it easier to identify the source of a security alert that only has an IP address as the starting point.

4.7 IDS / IPS

Intrusion detection systems (IDS) and intrusion prevention systems (IPS) can be components of a firewall, or their own separate appliances that are in-line with the network. The difference between them is that an IDS will only detect and alert, while an IPS can also automatically prevent or stop the activity by dropping connections and traffic.

SNORT is an example of software that can be used as either IDS or IPS and can be found in both firewalls as a module or as a separate device attached to the network.

The difference between a firewall, IDS and an IPS can be confusing, but figure 5 can be used to clarify the situation.

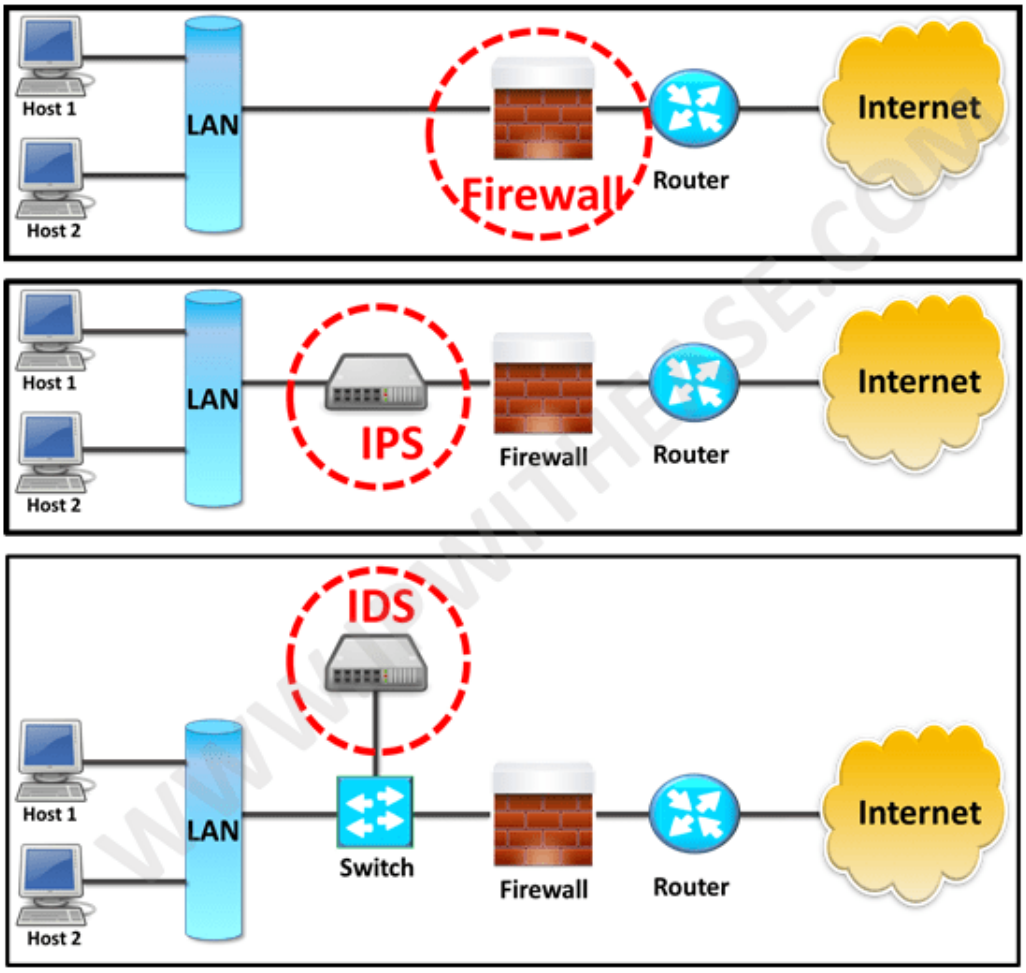

Figure 5. The differences between Firewall, IDS & IPS [18.].

4.8 Cloud Services

The utilization of cloud services is normal operations today. Software as a Service, Infrastructure as a Service etc. models all rely on “someone else” to operate the hardware, and for you the customer to just use the provided service. The most common cloud service that companies use is Microsoft Office 365 and the most common way that data breaches happen today is via the compromise of a business email account credentials.

According to the 2020 Data Breach Investigations Report, over 80% of security compromises is driven and enabled by credential theft to date. When a business email account is compromised, the entire organization can suffer far-reaching harm.

[19.]

It is important to be able to detect when business email credentials have been compromised, as the situation can quite quickly spiral out of control.

A common attack pattern is as follows:

- A wave of generic phishing emails hits employees.
- A few users enter their credentials on the phishing site.
- Malicious actors use the credentials to login to the users O365 account.
- Emails are read, interesting data is harvested.
- Existing email discussions are identified and hijacked to spread malware and compromise further accounts.
- If the captured user account belongs to an authority figure, money related scams can be attempted.

Some threat actors specialize in gaining the initial access but will not themselves do anything with it. Instead, they will sell the credentials on black market web sites for other interested parties.

The absolute best way to combat business email compromise (BEC) is to enable multi factor authentication (MFA). Without it, you will be constantly recovering from successful BEC’s.

MFA failure alerts, as well as Geolocation and Authorized Device alerts fall under conditional access blocks in O365. When dealing with conditional access

block alerts the most important thing to keep in mind is, that it is done **AFTER** the credentials have been validated. This means that every single conditional alert block had the correct credentials entered and the credentials should now be treated as compromised.

4.9 System Logs

System logs will give a good overview on what has been happening on the specific host, and will help in both pre-emptive detection and containment, as well as in after incident investigation. For example, unexpected shadow volume deletion is a good indication of a ransomware malware getting ready to encrypt the files on the host.

4.10 Internal Applications

If your company is creating and utilizing home grown applications, developers should take security and logging into account. At minimum proper audit logging should be incorporated.

5 Monitoring

Security alert monitoring is done with a tool called Security Information and Event Management (SIEM), this tool takes in all the logs and parses them for actionable key-value pairs. For this to work properly the key-value pairs inside the events will have to be normalized i.e., standardized across the different log sources.

Baselining the environment is also an important step in preventing false positive alerts from overwhelming the security analysts and causing loss of morale and effectiveness via alert fatigue. Properly created alert rules and a well-defined policy for tuning badly working alert rules will help to raise the efficiency of the monitoring team.

Alert rules are compromised from logic and indicators of compromise (IOC), both of which are enriched and refreshed with threat intelligence (TI). TI will give context to alert logic and detections while also removing obsolete IOC's and bringing in new ones.

5.1 SIEM

Security information and event management (SIEM) is the heart of the security monitoring. SIEM is where all to logs come together, and the magic happens.

> Security information and event management (SIEM) technology supports threat detection, compliance and security incident management through the collection and analysis (both near real time and historical) of security events, as well as a wide variety of other event and contextual data sources. The core capabilities are a broad scope of log event collection and management, the ability to analyze log events and other data across disparate sources, and operational capabilities (such as incident management, dashboards and reporting).

[20.]

Finding the SIEM that best suits your environment and personnel is one of the most important choices to make when starting security monitoring. All the security alerts that you receive, are supposed to be generated by your SIEM. All

the security investigations that you do, will be mostly done in SIEM. If you find yourself constantly looking up log events from other systems, it is a sign that you are not sending all the required logs to the SIEM.

Figure 6 shows the current state of the SIEM market according to Gartner. The recommendation would be to start the SIEM selection process form the products that are shown in the “LEADERS” quadrant, as these products should be the most complete in features and usability.

Source: Gartner (June 2021)

Figure 6. Gartner Magic Quadrant for SIEM’s 2021 [21.].

The cost and effort to run a SIEM can vary wildly. For example, Splunk licenses are based on data volume, and the price tag for a Splunk SIEM that has a long enough log retention time to meet requirements set by compliance can be eye wateringly high.

Another thing to consider is the upkeep. IBM QRadar is quite easy to use and has a low learning curve but keeping the system up and running can be comparatively difficult. Updating the system is quite often a full reinstall with just the configuration files re-used.

Both of these issues can be alleviated somewhat. For Splunk it becomes important to pay attention to the log events that are sent and filter out as much of useless data as possible. The danger here then becomes that something important is not gathered. For IBM QRadar the update issues can be completely mitigated by choosing the cloud version of the product if your compliance requirements allow for that, making the upkeep truly someone else's problem.

5.2 Event Normalization

After all the relevant logs from the disparate systems across the whole environment are gathered in SIEM, they need to be normalized to better support investigations and rule creation. In normalization the SIEM is taught which keys across the different logs refer to the same thing. For example, one log might refer to usernames by the key: "usr", while another uses the key: "user". After the SIEM is informed of this usernames from both log sources can be searched simultaneously with just one keyword.

5.3 Baseline

Baselining is an activity that is done before the security monitoring system is taken into production use. It aims to identify what is normal in your specific environment and tune the alerts to ignore these.

An example of baselining would be to identify that the high data volume connection going to an external IP address every day is a remote backup job, and not an indication of data exfiltration by a malicious third party.

5.4 Whitelisting

Whitelisting works in conjunction with baselining and alert rules. Any organization specific false alert triggering findings from baselining should be handled by either whitelisting the activity or modifying the alert rules to ignore it.

An example of whitelisting would be to add the remote backup IP address from our previous example to a list of safe IP addresses inside the SIEM, that the SIEM alert rules then use to make the decision to not trigger an alert.

5.5 Indicator of Compromise

Indicators of compromise (IoC) are small bits of data, that security monitoring systems depend on to do their base detection work. Fancy rules and machine learning can be used to detect abnormal or hostile actions, but the bread and butter of detection is still the IoC. IoC's can be hash values of known bad files, or IP addresses and domain names that investigation has detected a malicious file connecting to.

By investigating an IoC you are following a trail of breadcrumbs, where more and more IoC's are uncovered and adding context around the overall picture of the malicious activity that you are investigating.

5.6 Alert Rules

SIEM's use alert rules to trigger alerts. Some of these will be built into the system and can be tuned to better suit the environment. The built-in rules will not be enough to protect the environment, and more rules must be created by hand.

The rules can be simple like:
Login to: O365
From: NOT Finland
Action: Trigger Alert

All the way to very complex multi detection rules that combine into a single alert rule.

Threat intelligence is an important part of alert and detection rules, as without constant IoC refreshing the rules will go stale. This will have an adverse impact in the detection capabilities of the rules.

Inventory of services and systems also comes to play when creating rules, as your SIEM will populate its internal rule component lists with the services and systems it recognizes to make them easier to use. In our earlier “O365 login from NOT Finland” example, the SIEM is aware of O365 being in use, and offers it as an option for a “login” detection rule component.

5.7 Alert Fatigue

Alert fatigue is something that anyone that has worked in a Security Operations Centre (SOC) can relate to. It is caused by an excess of low-quality security alerts, and over time causes a significant drop in morale and efficiency.

The best way to combat alert fatigue is to incorporate alert tuning deep into the security alert handling processes. When a rule is triggering false positive alerts, the cause should be investigated and corrected. Common causes for false positive alerts are inadequate baselining, a change in the environment, addition of new detection rules, or the addition of low-quality threat intel data.

5.8 Threat Intelligence

Security Monitoring lives and dies based on the quality of its threat intelligence. If you do not receive updated threat intelligence, your IoC's will over time decline in quality, false positives will start going up, and real malicious activity will go undetected.

There are both free and paid threat intelligence services, and it is important to incorporate at least some kind of an threat intelligence update process to your security monitoring. Automating the importing of data might be convenient, but incorrectly categorized data is quite common, and without any data validation you will soon find that some very common services that users depend on suddenly start alerting and possibly even become blocked.

When dealing with threat intelligence updates, it is important to consider how long the data is going to be considered valid, as this will vary by the type of data that you are receiving. For example, IP address based threat intelligence can be quite short lived, while a domain that has been created for a phishing campaign can be expected to remain valid forever.

Make it easy to remove invalid threat intelligence from your system. Prevent automation from re-applying the removed invalid data on sync.

Figure 7 shows the four types of threat intelligence and their descriptions.

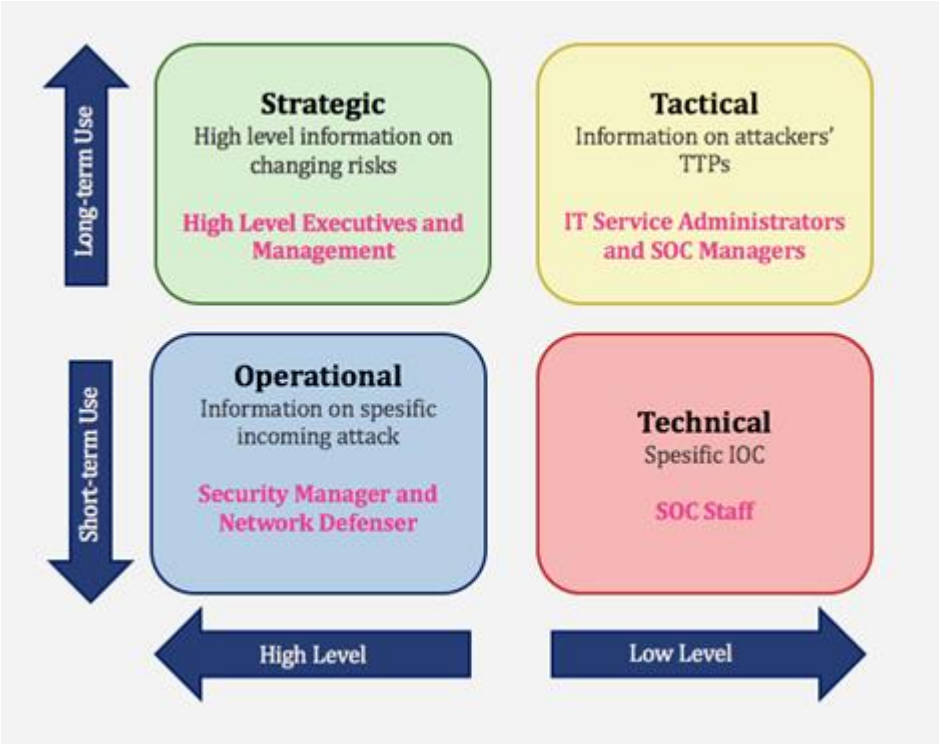

Figure 7. The four types of Threat Intelligence [22.].

Technical threat intelligence will include IoC data that can be imported to SIEM blacklists and alerting rules, like domain names that are known to be used by malicious actors, or IP addresses that have been involved in recent malicious activity.

Tactical threat intelligence will give context around IoC's, attempting to answer questions like: "Why is this considered bad", "Who are the threat actors related to this" and "What other information links to this".

Operational threat intelligence looks at the current situation and identifies the vulnerabilities and weaknesses that are currently exploited in the wild and used to attack targets.

Strategic threat intelligence attempts to look at the global picture and identify changes in threat actor behaviour. For example, when Russia invaded Ukraine one of their first acts was to activate a malware that rendered VIASAT satellite terminals inoperable [23.]. An attack on satellite infrastructure by a nation state

was unexpected but makes sense in the context of a war. The spill over from this action also affected other parties using VIASAT satellite terminals, for example causing wind energy company Enercon to lose contact with their wind turbines.

5.9 Log Source Anomaly Monitoring

It can be quite embarrassing to find out during a BEC investigation that your O365 logs stopped coming in a week ago.

To prevent this from happening to you, it is important to create low event volume alerts for your log sources. The low event volume value is specific to your environment and each log source, and identifying the value is part of the baselining process.

A high event volume alert is also useful to identify misconfigurations and other special circumstances, like Distributed Denial of Service (DDoS) attacks.

6 Use Cases

Use cases are instructions that you will write to cover specific scenarios. For example, what to do when a user has fallen victim to a phishing message and entered their credentials on the phishing page? It makes sense to create a general working instruction that is tailored for your environment, that all security analysts can follow instead of re-inventing the wheel every time the alert triggers. The use cases should not be overtly strict and should allow the security analyst to use good judgement in doing things differently if needed.

The following are two barebones examples of use cases. To make them usable, environment specific instructions should be added to enrich each step.

6.1 Business Email Compromise

Scenario: A user fell victim to a phishing message and entered his credentials to a phishing page. The users O365 email account was compromised and is being used to send phishing emails to other people inside the company.

Recovery steps to take:

- Disable user account
- Reset all current sessions
- Change user password
- Enable MFA for user
- Check for malicious email and mailbox rules
- Check for malicious O365 apps
- Enable user account
- Check sent emails
- Check trash folder
- Notify any recipients of sent phishing emails that the situation has been corrected and emails can be considered safe again
- Check if confidential or GDPR data was accessed and possibly exfiltrated

Investigate the original phishing email and landing page, find the URL that the credentials were POST'd to. Check if anyone else accessed the POST URL. For any other users that sent their credentials away, perform the same recovery steps. Perform the same investigation for the phishing emails that were sent from the compromised email account.

6.2 Suspected Malware

Scenario: A user received an email with an encrypted zip file. Sender claims that he did not want to send confidential material unencrypted in an email. User has asked for help verifying if the file contained inside the password protected zip file is safe.

Investigation steps to take:

- Transfer the zip to a Linux host
- Unzip the zip
- Run "file" command on the extracted file
- Take the SHA 256 hash of the extracted file
- Run the file in internal sandbox
- Check the SHA 256 hash for information on virustotal.com

If the file is still suspicious but the company considers it important, consider sending it to a third-party security vendor for further analysis, else the recommendation would be to throw the file in the trash.

If on the other hand you would like to conduct further analysis yourself, Didier Stevens has created a lot of valuable tools [24.] that can be used. I would recommend taking some time to read and watch information regarding these tools before attempting to use them.

7 Computer Forensics

Computer forensics is a task that you should not ever attempt to do yourself, if there is even a tiny chance that the situation will involve the courts. The chance of making the evidence inadmissible is too great.

If the incident has no chance of going to court and you are looking to gain some experience on the subject, here is a quick overview of the process.

Additional information on forensic steps can be found in NIST publication SP 800-60 “Guide to Integrating Forensic Techniques into Incident Response” [25.]. The document should be fully read before attempting to conduct a forensic investigation.

7.1 Acquire

Take a bit-by-bit copy of the full disk image of the system you are investigating. Compare the hash value of the disk image to the original, if they do not match, retry the image acquisition. Once the hashes match, take a second copy of the image, verify that its hash also matches, and store the second copy safely as your back-up image.

If the computer is still powered on and has not been restarted after the incident happened, a memory dump can also be useful.

7.2 Examination

Attempt to catalogue files and data of interest on the disk image, to take to the next step of investigation. This can be a tedious task but doing it properly will make the Analysis phase more streamlined. Any information that you have of the incident can be used to whittle down the data mountain to a more manageable size. Text patterns, file types, email addresses etc. both inclusively

to take to the analysis phase and exclusively to filter out from the analysis phase.

7.3 Analysis

After the interesting data has been extracted, it can be analysed with software purpose made for it. One such software tool is Autopsy by Brian Carrier and Basis Technology.

> Autopsy development is led by Brian Carrier, whose team builds easy-to-use tools for cyber first responders to intrusions, crime scenes, and war zones. Our team also develops Cyber Triage, fast and affordable incident response software any organization can use to rapidly investigate compromised endpoints.

[24.]

Attempt to find data relevant to the investigation and correlate it with other data sources that are available to you, for example IPS/IDS logs or Firewall logs. Comparing system settings and files to a known safe computer can also bring to light malicious changes done to the system.

7.4 Report

As the audience of the final report in this forensics exercise is the analyst himself, the report can be quite light, but should include the investigation steps taken, and their results, even if negative. Learning to write proper reports is a valuable skill that should not be overlooked.

8 Additional Activities

In addition to security monitoring there are various other activities that can be done to improve the overall security posture of the environment. These can be bought as a service or done internally if the company has the competencies to conduct the activities. The attackers in these scenarios are called the red team, while the defenders are called the blue team.

8.1 Threat Hunting

In threat hunting the security analyst uses his knowledge of the environment in combination with threat intelligence to look for malicious activity in historic log files. The average time that malicious actors spend undetected in the environment can be quite high, for example Mandiant reports:

> Back in 2011, we reported a 416-day global median dwell time, indicating that attackers were operating undetected on a system or network for over a year on average. This time, from Oct. 1, 2019 through Sept. 30, 2020, the median dwell time has decreased to only 24 days. This means—for the first time in M-Trends history—the median dwell time has dropped to under one month.

[25.]

Better tools and practices at the defending organizations have reduced the detection times drastically, but the current average of nearly a month still gives the threat actors enough time to complete their goals before detection.

8.2 Penetration Testing

A penetration test is a pre-agreed and scoped attempt to breach the environment. This could be limited to a certain application or system and start from an authorized device inside the company’s network or from the public internet and have the whole environment as the scope. The aim is to cover as much of the in-scope attack surface as possible so that the identified security issues can be corrected.

8.3 Red Teaming

In contrast to penetration testing, red teaming does not attempt to cover the whole attack surface, and instead attempts to reach the pre-decided targets as efficiently as possible. The Blue team are also kept in the dark regarding the activity, to better see what their capabilities and response processes are in reality. Red teaming activity can cause critical security alerts to be triggered, and these should be handled as if they were real incidents until the person in charge of the red teaming exercise blows the whistle.

8.4 Purple Teaming

Purple teaming consists of the red and blue teams working together to strengthen the overall security posture of the environment. The blue team is kept in the loop regarding the actions taken by the red team and can verify that their rules and alerts cover the activities being conducted by the red team. Purple teaming is also an efficient way to identify visibility gaps in the logs being brought in to the SIEM.

9 Managed Service Providers

Managed Service Providers (MSP's) can be used to fill in gaps in resourcing or knowledge. If the company does not have the resources to do security monitoring them self or finds that using an MSP is more cost-effective, the service can be bought from a third-party company.

There are currently many companies offering Security Operations Centre (SOC) services and picking the correct one can be hard. Pay attention to the "security analysts on duty" to "customers they are monitoring" ratio, as this will give a rough estimate on how much time they can dedicate per customer. Some SOC's will also offer 24/7 coverage but verify that there will really be a security analyst at the keyboard 24/7, instead of the night hours being covered in some other manner.

When using an outside service to do security monitoring, you will have to send your logs outside of your organization, verify that this is not prohibited by laws or contracts covering you. For example, it is quite normal that not all logs can be exported outside of EU, due to the data contained within them. Verify from your prospective MSP where they would store the logs, and from where their employees would access the logs.

10 Conclusions

Implementing security monitoring can be a daunting task, but I am certain that after reading this thesis, the reader will have the knowledge required to discuss the task on an even footing. Projects like this usually involve bringing in consultants or MSPs, and you need to be aware of what you can and should demand from them.

Cost-effectiveness analysis can turn the scales to favour buying security monitoring as a service from an MSP instead of building your own, but some companies might be loath to allow an external entity access to their most valuable data. Creating your own security monitoring system allows for tailoring exactly to the company's needs and has all the pertinent silent knowledge available to the security analysts in the form of their colleagues. These considerations are something that must be discussed and taken into account when making the decision to go with an MSP or building your own.

I am happy that I have been able to use my experience and knowledge to help others to increase the security of their environments. Hopefully at least one security breach will be prevented by the application of knowledge from this thesis.

References

1. Forcepoint. What is Shadow IT?, Article.
<https://www.forcepoint.com/cyber-edu/shadow-it>.
Accessed 26 May 2022.

2. Darktrace. Global Threat Report 2017, Presentation.
<https://cdn2.hubspot.net/hubfs/2784256/1_nat_2017_recap/Presentations/Darktrace%20-%20Global%20Threat%20Report%202017.pdf?t=1528334118161>.
Accessed 26 May 2022.

3. Deloitte. 91% of all cyber attacks begin with a phishing email to an unexpected victim, Press release.
<https://www2.deloitte.com/my/en/pages/risk/articles/91-percent-of-all-cyber-attacks-begin-with-a-phishing-email-to-an-unexpected-victim.html>.
Accessed 26 May 2022.

4. NIST. Glossary – defense in depth. Glossary.
<https://csrc.nist.gov/glossary/term/defense_in_depth>.
Accessed 26 May 2022.

5. Arstechnica. Photos of an NSA upgrade factory show cisco router getting implant, Article.
<https://arstechnica.com/tech-policy/2014/05/photos-of-an-nsa-upgrade-factory-show-cisco-router-getting-implant/>.
Accessed 26 May 2022.

6. XKCD. Security, Webcomic.
<https://xkcd.com/538>.
Accessed 31 May 2022.

7. Trend Micro. Deep Security Help Center – Syslog message formats, Article.
<https://help.deepsecurity.trendmicro.com/10/0/Events-Alerts/syslog-

parsing.html>.
Accessed 26 May 2022.

8. SentinelOne. SentinelOne is a Leader in the 2021 Gartner Magic Quadrant for Endpoint Protection Platforms. Here's Why., Article.
<https://www.sentinelone.com/blog/sentinelone-is-a-leader-in-the-2021-gartner-magic-quadrant-for-endpoint-protection-platforms-heres-why>.
Accessed 26 May 2022.

9. Sumologic. Collect logs for the CrowdStrike Falcon Endpoint Protection App, Article.
<https://help.sumologic.com/07Sumo-Logic-Apps/22Security_and_Threat_Detection/CrowdStrike_Falcon_Endpoint_Protection/Collect_logs_for_the_CrowdStrike_Falcon_Endpoint_Protection_App>.
Accessed 26 May 2022.

10. CrowdStrike. Digging into BokBot's Core Module, Article.
<https://www.crowdstrike.com/blog/digging-into-bokbots-core-module>.
Accessed 26 May 2022.

11. Palo Alto Networks. The Next-10 Firewall, Article.
<https://start.paloaltonetworks.com/2021-gartner-mq-for-firewalls.html>.
Accessed 26 May 2022.

12. Ossec. Log samples for Checkpoint, Documentation.
<https://www.ossec.net/docs/log_samples/firewalls/checkpoint.html>.
Accessed 26 May 2022.

13. Ossec. Log samples for the Cisco IDS/IPS module for IOS., Documentation.
<https://ossec.net/docs/log_samples/cisco/cisco_ids.html>.
Accessed 26 May 2022.

14. Zeek. dns.log, Documentation.
<https://docs.zeek.org/en/master/logs/dns.html>.
Accessed 26 May 2022.

15. Zeek. Dhcp log, Documentation.
<https://docs.zeek.org/en/master/logs/dhcp.html>.
Accessed 26 May 2022.

16. Splunk Community. How to correlate events from PaloAlto VPN logs and Windows authentication per user, comparing src_ip and machine_name?, Community Forum.
<https://community.splunk.com/t5/Splunk-Search/How-to-correlate-events-from-PaloAlto-VPN-logs-and-Windows/m-p/540287>.
Accessed 26 May 2022.

17. Palo Alto Networks. GlobalProtect Log Fields for PAN-OS 9.1.3 and Later Releases, Documentation.
<https://docs.paloaltonetworks.com/pan-os/9-1/pan-os-admin/monitoring/use-syslog-for-monitoring/syslog-field-descriptions/globalprotect-log-fields/globalprotect-log-fields-for-pan-os-913-and-later-releases#id5795bd71-1dc5-4f82-872d-a9ba6cb7cedf>.
Accessed 26 May 2022.

18. IPWITHEASE. IDS vs IPS vs Firewall – Know the Difference, Article.
<https://ipwithease.com/firewall-vs-ips-vs-ids>.
Accessed 26 May 2022.

19. F-Secure. Business Email Compromises – What Are the Most Common Threats?, Article.
<https://blog.f-secure.com/business-email-compromises-what-are-the-most-common-threats>.
Accessed 26 May 2022.

20. Gartner. Security Information And Event Management (SIEM), Documentation.
<https://www.gartner.com/en/information-technology/glossary/security-information-and-event-management-siem>.
Accessed 27 May 2022.

21. Exabeam. 2021 Gartner Magic Quadrant for SIEM, Article.
<https://www.exabeam.com/library/2021-gartner-magic-quadrant-for-siem>.
Accessed 27 May 2022.

22. SOCRadar. What is Strategic Cyber Intelligence and How to Use it, Article.
<https://socradar.io/what-is-strategic-cyber-intelligence-and-how-to-use-it>.
Accessed 27 May 2022.

23. Recorded Future. Viasat confirms report of wiper malware used in Ukraine cyberattack, Article.
<https://therecord.media/viasat-confirms-report-of-wiper-malware-used-in-ukraine-cyberattack>.
Accessed 27 May 2022.

24. Didier Stevens. My Software, Web Page.
<https://blog.didierstevens.com/my-software>.
Accessed 31 May 2022.

25. NIST. Publication SP 800-60 “Guide to Integrating Forensic Techniques into Incident Response”, Article.
<https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-86.pdf>.
Accessed 31 May 2022.

26. Autopsy. Autopsy Digital Forensics, Web Page.
<https://www.autopsy.com>.
Accessed 27 May 2022.

27. Mandiant. M-Trends 2021: A View From the Front Lines, Article. <https://www.mandiant.com/resources/m-trends-2021-a-view-from-the-front-lines>.
Accessed 27 May 2022.