

**599.1**

# Purple Team Tactics & Kill Chain Defenses

**SANS**



# Purple Team Tactics & Kill Chain Defenses

© 2019 Erik Van Buggenhout & Stephen Sims | All Rights Reserved

Welcome to SANS Security SEC599: Defeating Advanced Adversaries.

In this course, you will build essential skills required to fend off today's advanced cyber attacks. The course will be highly hands-on, as we help you develop skills by exercising them in hands-on, realistic lab settings. Although this is not a penetration testing course, we will have sufficient attention for the offensive side of the spectrum. We will provide you with a deep technical understanding of how advanced adversaries work, as this will help us be more efficient defenders. Likewise, we will inform you on how to respond to cyber security attacks but will primarily focus on how to prevent and detect them.

Our goal is to keep the course as interactive as possible. If you have a question, please let the instructor know. Discussions about relevant topics are incredibly important in a class like this, as we have numerous attendees with various levels of skill coming into the class. Share your insights and ask questions. The instructor does reserve the right, however, to take a conversation offline during a break or outside of class in the interest of time and applicability of the topic.

As course authors, we welcome any comments, questions, or suggestions pertaining to the course material. We would also like to extend our thanks to Didier Stevens (a SANS ISC handler), whose contributions greatly helped improve the course.

Erik Van Buggenhout  
erik.van.buggenhout@gmail.com  
<https://www.nviso.eu/>

Stephen Sims  
ssims@sans.org  
<https://www.sans.org/>

## Course Outline

Day 1: Introduction & Reconnaissance

Day 2: Payload Delivery & Execution

Day 3: Exploitation, Persistence and Command & Control

Day 4: Lateral Movement

Day 5: Action on Objectives, Threat Hunting & Incident Response

Day 6: APT Defender Capstone

Throughout the week, a large focus on hands-on exercises that illustrate the Purple Team concept!

### Course Outline

SEC599 has six days of content:

#### Day 1: Introduction & Reconnaissance

In Day 1, we will explain what purple teaming is and what the current threat and attack landscape looks like. We will explain what techniques are being used by our adversaries, so we can prepare ourselves to prevent, detect and respond to them. We will also zoom in on the importance of knowing one's own environment. Finally, we will cover how the attacker takes his first steps: How does he perform reconnaissance and what can we do to hinder it? The courseware will cover technical controls but will also touch upon "soft topics" such as security awareness.

#### Day 2: Payload Delivery & Execution

After reconnaissance is performed and vulnerabilities are spotted, the adversary will weaponize the payload and deliver it to the target. We will analyze how delivery of the payload can be detected and blocked. We will cover a variety of techniques, including mail-based controls (e.g. SMTP file and URL carving, sandboxing...) and web-based controls (access controls using web proxies). We will also introduce YARA as a payload signature language!

#### Day 3: Exploitation, Persistence and Command & Control

Day 3 will explain how exploitation can be prevented or detected. Attendees will obtain an in-depth understanding of current exploitation tactics. We will introduce effective security controls to stop exploitation attempts. We will also zoom in on persistence techniques typically employed by adversaries and how command and control is established.

#### **Day 4: Lateral Movement**

Once an initial foothold is obtained, the adversary can start performing lateral movement, where they pivot throughout the environment looking to accomplish their objectives (e.g. steal sensitive data). We will focus on how lateral movement is done in a typical AD environment. We will also introduce cyber deception techniques to slow down adversaries!

#### **Day 5: Action on Objectives, Threat Hunting & Incident Response**

Day 5 focuses on stopping the adversary during the final stages of the attack:

- How does an adversary dominate the AD environment?
- How can data exfiltration be detected and stopped?
- How can we do threat hunting?
- How can threat intelligence aid defenders in the APT Attack Cycle?
- How can defenders perform effective incident response?

#### **Day 6: APT Defender Capstone**

Day 6 concludes with a hands-on Capstone challenge, applying all the skills you've learned in a friendly, competitive, environment!

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.1

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim

### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

## Introduction

The key goal of the course is to help you improve how you prevent, detect (and respond) to cyber security attacks by advanced adversaries. In order to implement effective security controls, we are convinced you first need to **learn how the adversary operates**, so we can "stop them in their tracks"

The course authors (with a combined 20+ years' experience in red teaming, penetration testing & exploit development) created the course together with SANS ISC handlers, providing a **unique mix of offensive and defensive skills** bundled in one course!

The course will structure effective security controls using industry standards such as **MITRE ATT&CK** and the **Cyber Kill Chain**, which describe how adversaries operate.

### Introduction

The key goal of SEC599 "Defeating Advanced Adversaries" is to help you improve how you prevent and detect cyber security attacks by advanced adversaries. We will also cover techniques for effective incident response, although not in-depth, as SANS has dedicated courses that cover this topic (such as FOR508 | Advanced Incident Response, Threat Hunting, & Digital Forensics).

The course authors (with a combined 20+ years' experience in red teaming, penetration testing & exploit development) created the course together with SANS ISC handlers, providing a unique mix of offensive and defensive skills bundled in one course!

In order to implement effective security controls, we are convinced it is vital to first understand how adversaries operate. We will thus first explain offensive security techniques, explaining how organizations are currently being compromised. Based on this understanding, we will structure attacks according to industry standards such as MITRE ATT&CK and the Cyber Kill Chain, which describe how adversaries operate and how effective controls can be implemented.

## Introducing the SANS Integrated Lab Platform

In order to simulate realistic scenarios, we needed to create an enterprise-like environment. We cannot just set this all up during a 6-day course, so we prepared this environment for you.

Systems are **preconfigured** with the tools and settings needed to complete lab exercises. Key focus is on learning experience, not troubleshooting prerequisites / compatibility issues

Individual client and server targets: No one else can interfere with your lab experience

- Integrated system access and step-by-step directions for completing the exercises
- Key knowledge areas called out as you complete the lab

Virtualized labs accessed through your **web browser**.

### Introducing the SANS Integrated Lab Platform

SANS is committed to providing a superior course with skills that you can use immediately when you get back to the office. A significant part of this course experience is the use of hands-on lab exercises designed to reinforce the topics we cover during lecture.

The SANS Integrated Lab Platform is an integrated lab and workbook environment, providing consistent and easy access to the client systems and server targets through your web browser. Through this platform, you simply browse to a URL and login, then you will be able to access all the client and server systems and see the lab step-by-step directions needed to complete the lab exercises, in a single browser window.

The systems you will access through this platform have been preconfigured with the tools, software, and files needed to complete all the exercises. This allows you to focus on applying the learning objectives for the lab instead of spending valuable time configuring your laptop, troubleshooting network or conflicting software settings, and clicking Next, Next, Next, Next, Next, Finish over and over again.

Another benefit of the SANS Integrated Lab Platform is that you have individualized access to client and server systems. When you start a lab, the servers supporting the platform spin up a duplicate copy of the server and client systems needed to complete the lab, uniquely accessible to you. This stops other people in the classroom from interfering with your lab experience (intentionally or unintentionally), making the lab exercises more consistent and accessible. Instead of flipping back and forth between a printed lab workbook and your laptop, the SANS Integrated Lab Platform integrates both the client and attacker system view with the step-by-step exercises.

The step-by-step directions in the lab call out key knowledge areas that are important to recognize, as well as alerts to make you aware of the need for caution when using a tool or completing a specific lab step, and screenshots to help you stay on track with the exercises.

Fundamentally, the SANS Integrated Lab Platform is a way for SANS to deliver a consistent lab experience that focuses on helping you build your skills while minimizing system setup needs.

## Getting Started

<https://labplatform.sans.org>

- Bookmark this URL for easy access
- Use the instructor-supplied card to log in with your username and password



Username:   
Password:

Please keep your credential card handy!  
You will use it each day!

### Getting Started

At the beginning of class, your instructor will hand out a login card with your username and password information needed to access the lab server. Please keep this card handy, as you'll use it each day for labs.

Simply browse to the URL on this page (and printed on the login card). When prompted, enter your username and password information, then click Sign In.



## Lab Assignments



Admin My Labs

Welcome Erik Van Buggenhout [Sign Out](#)



Lab and Assignments  
Erik Van Buggenhout

[Details](#) [Search](#)

All times shown in Central Standard Time

[Lab Series Assignments](#) [Lab Instances](#)

▶ Running and Saved Labs (0)

▶ Checked Out (0)

▶ Assignments (1)

Series	Start	Expires
SEC599-C01	5/16/2017	6/30/2018

Click here

▶ Recent Lab Instances (0)



SEC599 | Defeating Advanced Adversaries – Purple Team Tactics & Kill Chain Defenses

8

### Lab Assignments

When you log in to the system, you will see the "My Labs" page. In the Assignments group, you will see your course assignment. Click the course assignment link to see the exercises.

## Launching Lab Exercises (1)



Admin My Labs

Welcome Erik Van Buggenhout [Sign Out](#)



Lab Series Assignment

Erik Van Buggenhout

SEC599-C01

[Edit](#) [Delete](#) [Search](#)

Student: Erik Van Buggenhout  
Series: SEC599-C01  
Starts: Tuesday, May 16, 2017  
Expires: Saturday, June 30, 2018

- Labs



1 Exercise - One click is all it takes [Details](#)  
SEC599-1.1

Status: Cancelled  
Started: Tuesday, May 16, 2017 9:43 AM (Central Standard Time)  
Ended: Tuesday, May 16, 2017 9:45 AM (Central Standard Time)

Click here

Launch



SEC599 | Defeating Advanced Adversaries – Purple Team Tactics & Kill Chain Defenses

9

### Launching Lab Exercises (1)


After clicking on your lab assignment, you will see a list of all the exercises in the lab assignment. Click the Launch button to start the desired exercises. The exercise will open a new window and kick off the virtual machines needed for the exercise automatically.

## Launching Lab Exercises (2)

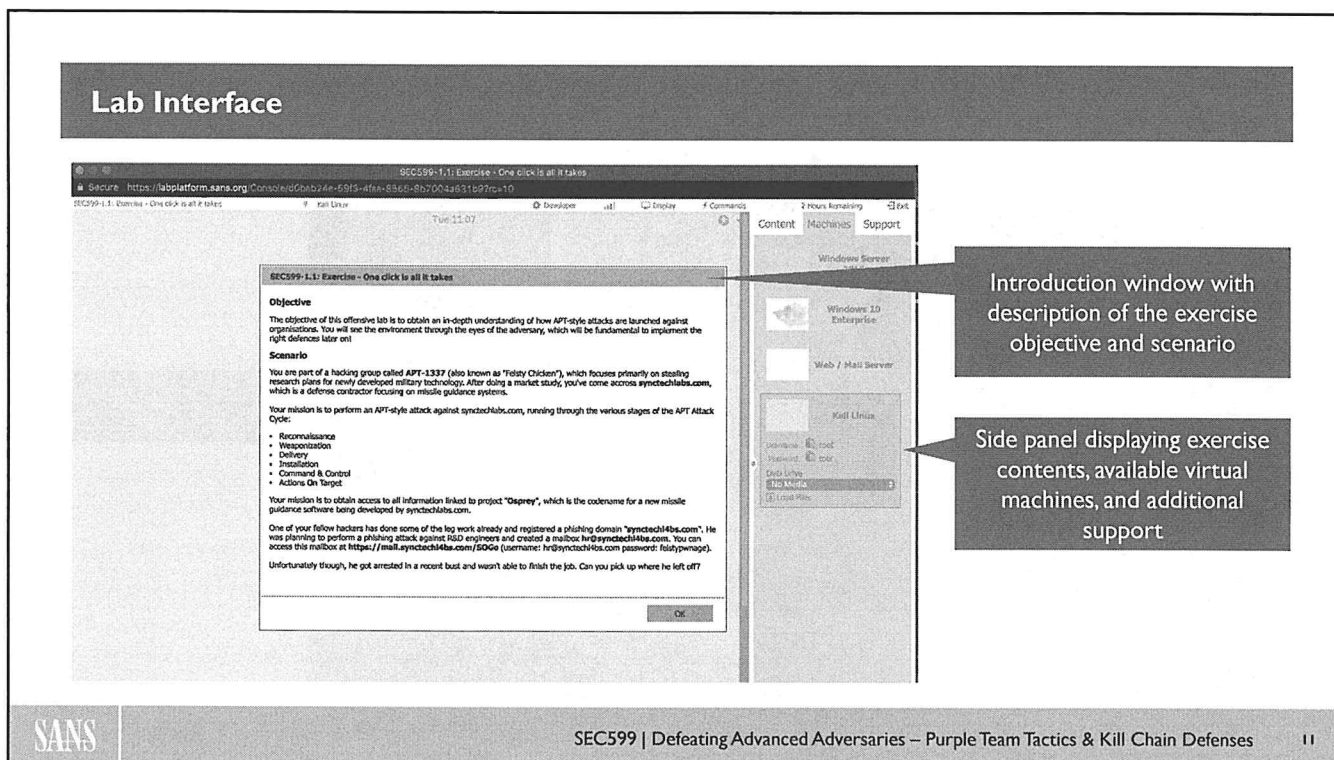


Your lab environment is being built

Your lab will be available in about 2 minutes and 30 seconds.



This page intentionally left blank.



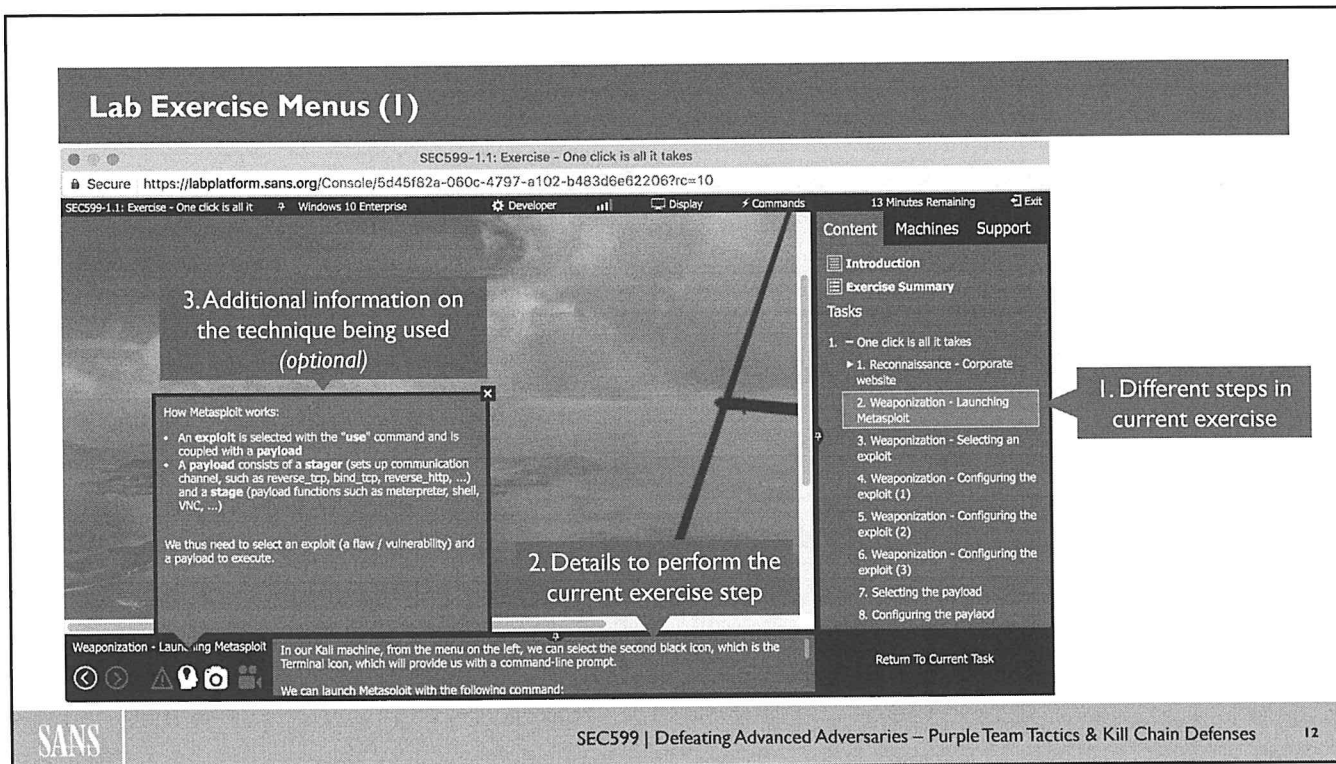
## Lab Interface

The window that opens when you click the Launch button will provide the interface for access to one or more virtual machines, and the step-by-step directions for the exercise.

First, you will see the objective and scenario information for the lab. Read this material, then click the OK button. Next, you will see an introduction to this specific exercise (a lab can have more than one exercise). Read through this material, then click Next\*\*\*.

On the right-hand side of the overall display, you have three different titles in the menu: "Content" – "Machines" – "Support"

- Content: A step-by-step overview of the activities to be concluded in this exercise
- Machines: A listing of the different machines available to you in this exercise
- Support: Support information on the SANS Integrated Lab environment



### Lab Exercise Menu (1)

Now you are ready to begin the exercise. Let's look at a few of the elements shown on this page. On the right side is a list of the step-by-step directions for the exercise. You can click to jump ahead and explore any of the steps as desired.

In the bottom of the window are the detailed directions for the selected step. As you change to the next step, the detailed instructions will update as well. These detailed instructions tell you what to do to complete the selected step.

When you complete the instructions in the selected step, you can click the Done button to mark the step as completed. The progress bar in the lower-right corner of the window will show you how many steps are completed, and how many remain.

The main portion of the browser window is your access to the virtual machine that you'll use for this exercise. You can click on this portion of the window and interact with the system like you would for your local system.

## Lab Exercise Menus (2)

The screenshot displays the LODS platform interface. The main window shows a terminal with the following output:

```
root@kali:~# curl https://metasploit.com
https://metasploit.com

Tired of typing 'set RHOSTS'? Click & pin with Metasploit Pro
Learn more on http://rnc147.com/metasploit

- Metasploit v4.14.2 show
+ ---+ 2621 exploits - 332 auxiliary - 282 post
+ ---+ 472 payloads - 48 encoders - 9 nops
+ ---+ Free Metasploit Pro Trial: http://r-7.co/rxmp

msf > search openoffice
[*] Module database cache not built yet, using slow search
Matching Modules
=====
Name                               Disclosure Date  Rank  Description
-----
msf5/windows/office/openoffice_document_importer 2008-04-17      normal  OpenOffice OLE Importer DocumentSummaryInfo
msf5/windows/office/openoffice_macro_stream_handling_overflow
```

At the bottom of the terminal, there is a "Screenshot" button. A callout box points to this button with the text: "4. You can ask for a screenshot of the command to execute / technique to use".

On the right side of the interface, there is a "Content" menu with a list of tasks:

1. One click to all it takes
2. Reconnaissance - Corporate website
3. Weaponization - Launching Metasploit
3. Weaponization - Selecting an exploit
4. Weaponization - Configuring the exploit (1)
5. Weaponization - Configuring the exploit (2)
6. Weaponization - Configuring the exploit (3)
7. Selecting the payload
8. Configuring the payload
9. Running the exploit
10. Move the payload
11. Access the phishing mailbox
12. Login into the phishing mailbox
13. Creating a new mail message
14. Crafting the phishing mail
15. The waiting game
16. Interacting with our

4. You can ask for a screenshot of the command to execute / technique to use

SANS

SEC599 | Defeating Advanced Adversaries – Purple Team Tactics & Kill Chain Defenses

13

### Lab Exercise Menus (2)

Finally, the bottom menu also has a "screenshot" button, which you can click to obtain a detailed screenshot of what task is expected of you.

All in all, the LODS platform was set up to be a highly intuitive platform that can help you to complete labs without any prerequisite issues. Should you have any further questions or remarks as we go through the different exercises, please don't hesitate to get in touch with your Instructor / TA.

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.1

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim

### Reconnaissance


Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

## What's Happening Out There?


More and more bad guys are using "cyber space" for their malicious means:



**Cyber Crime**



**Sabotage**



**Espionage**

Not all attacks are technically sophisticated, but that doesn't make them any less effective! "Advanced" doesn't have to mean **technologically advanced!**


### What's Happening Out There?

Over the past few years, more and more organizations have fallen victim to a variety of cyber attacks. Current trends include:

- *Cybercrime*: Attacks focused on earning money / generating revenue for a malicious (group of) perpetrators. Some of the most common attack methods we see these days are ransomware (both against organizations and individuals) and denial of service attacks. With ransomware, business-critical data is encrypted, after which a ransom is asked to allow the data to be recovered. With denial of service, a typical attack technique would be to disrupt the online presence of an organization, after which a ransom is asked to stop the denial of service attack.
- *Sabotage*: Attempts to disrupt your (online) operations. Sabotage is typically executed by hacktivists (e.g. politically motivated attacks against organizations that have different ideological views) or nation-states (e.g. sabotaging the critical infrastructure of other countries in times of war).
- *Espionage*: Could typically include both industrial and political espionage. State-sponsored attacks are not uncommon, as evidenced by a wide variety of discovered APT campaigns. The goal is often to steal data that could result in a commercial advantage (e.g. stealing R&D plans or strategy documents). Next to industrial espionage, political espionage can be focused on obtaining access to sensitive diplomatic intelligence or military technology. It is a "public secret" that most nations are developing offensive cyber capabilities and are using them to their benefit.




## Zooming in on Cyber Crime



**Cyber Crime**



**Sabotage**



**Espionage**

### **Zooming in on Cyber Crime**

First, let's have a look at Cyber Crime!



## Key Driver for Cyber Crime: \$\$\$



Online **banking Trojans** such as Zeus, SpyEye, Citadel...  
Tailored malware against POS & **ATM systems**

Attacks such as "Carbanak (2015)" and the "Bangladesh Hack (2016)" targeted against **backend banking services**

Ransomware targeted at **anyone** (Individuals, commercial companies, government organizations...)

### Key Driver for Cyber Crime: \$\$\$

Monetary gain is THE key driver for cyber crime. This makes the attacks somewhat predictable: They are coming for the money... Additionally, it makes the adversaries less persistent: Cyber criminal adversaries are looking for the path of least resistance; they will go where the money is easiest to obtain / steal. In order to fend off these adversaries, an age-old safari axiom can be of interest:

*"You don't have to be the fastest, just don't be the slowest."*

Some interesting attack techniques we've seen over the past couple of years:

- Online banking Trojans such as Zeus, Citadel and Dridex that attempt to infect online / mobile banking users. The idea here is to infect as many users as possible and transfer relatively small amounts per infection. Furthermore, tailored malware is being written that attacks POS (Point of Sales) and ATM systems.
- Somewhat more advanced attacks against banks themselves, where they attempt to infect business users involved in the creation, signing, and approval of larger fund transfers. Key examples of this include the "Carbanak" attack revealed in 2015 and the "Bangladesh Bank Heist" that occurred in 2016. In both cases, adversaries had obtained a foothold in the internal bank networks and were monitoring the environment to understand how fund transfer approval flows worked and how large fraudulent transactions could be executed.
- Finally, since 2015, we see a very strong rise in the use of ransomware that is targeted against anyone with data. Ranging from individuals, commercial companies to top-secret government organizations: If you are willing to pay to retrieve your data, you are a target.

#### Reference:

<https://securelist.com/the-great-bank-robbery-the-carbanak-apt/68732/>



## Online Banking Trojans – Zeus



The King of banking Trojans, Zeus first made a name for itself in 2007 during a **credential-theft attack** against the US Department of Transportation



Zeus **source code** was made public in 2011, after which it inspired a whole new generation of banking Trojans (e.g. the equally infamous Citadel)



An innovative example for many, it was one of the first Banking Trojans to introduce a “mobile brother” in **ZitMo** (Zeus in the Mobile)

### Online Banking Trojans – Zeus

In order to describe how Banking Trojans work, we will zoom in on the King of banking Trojans, namely Zeus (aka Zbot).

Zeus is a versatile Trojan horse malware, capable to perform many malicious acts. It is often used as man-in-the-browser. Man-in-the-browser is a word play on man-in-the-middle: In a man-in-the-browser attack, malicious code is injected into the internet browser process to steal information and to tamper with the rendering of web pages, for example by adding forms for phishing purposes. Man-in-the-browser malware can be written as a browser plugin, or as stand-alone code directly injected into the browser process.

When used in man-in-the-browser mode, Zeus can perform keylogging and form grabbing (stealing data entered into forms). Zeus is also used to spread ransomware (CryptoLocker). It first became known in 2007 when it was used (and detected) in a credential-theft attack against the US Department of Transportation, although its widespread use began in 2009. Another activity of the Zeus authors was the facilitation of tech support scams. When running on a Windows machine, this variant of Zeus would display a pop-up message alerting the user to the simulated presence of a computer virus. The message would instruct users to call a phone number (often claiming to be Microsoft support), where scam artists would “help” users to check for errors with the Windows event viewer (there are always error events in the viewer), claim that this was caused by a virus and get the victim to pay for a fake anti-virus solution.

In 2011, the Zeus source code (a Microsoft Visual Studio project) was made public, spawning many new banking Trojans. It is always easier to copy something than to start from fresh. Initially, the Zeus source code was reused by criminals with minor modifications, requiring little skill: Just be able to compile in Visual Studio and replace some strings like IP addresses and domain names. Later on, substantial changes were made resulting in completely new banking Trojans like Citadel.


Zeus-in-the-mobile (ZitMo) appeared in late September 2011. Working together with Zeus on Windows, ZitMo intercepts and steals mobile transaction authentication numbers (mTAN codes) sent to the victims phones. They can subsequently be used to perform fraudulent transactions.




## Carbanak or "The First APT against Banks"

Carbanak (also "Anunak") was one of the first targeted / advanced attacks against financial institutions (discovered in 2015)

- Previous efforts were typically aimed at banking customers (e.g. online or mobile banking malware); now the targets are the bank's own systems



**Step 1** – Phishing emails toward bank employees (not customers) infect workstation with Trojan



**Step 2** – Initially compromised machine is used for further exploration (looking for transactional systems)



**Step 3** – Behavior of users on transactional systems is monitored (learn how funds can be transferred)



**Step 4** – Steal funds through a variety of techniques (e.g. SWIFT transactions, ATM cash-out...)

Total losses reported to be about \$2 to \$10 million per victim bank with a total of up to \$1 billion

### Carbanak or "The first APT against Banks"

Carbanak is the name of an APT attack and associated malware, performed against financial institutions and discovered in 2015 by antivirus company Kaspersky. The Carbanak gang managed to steal at least \$500 million from financial institutions and their clients, through various means. The malware was often delivered via phishing emails.

Via phishing emails with executable attachments like CPLs (Malicious Control Panel items) or Word documents with exploits, a backdoor (Carbanak) is installed on the victim's machine. This malware is based on the Carberp malware. The malware is designed to support the following functions: Espionage, data exfiltration, and remote control. Once they gained access to the victim's machine, the criminals used this beachhead in search of computers that could help them perform fraudulent financial transactions, like computers operated by administrators. This lateral movement led them to computers that could perform financial transactions.

Often, the criminals behind the Carbanak gang would not have financial knowledge and procedures of the bank they were targeting but would quickly learn by recording the screens and keyboard strokes of the compromised machines and learn via videos how to operate the financial systems. Armed with this knowledge and credentials, they would perform operations to obtain money.

Money would be obtained through different scenarios, depending on the environment they discovered at the bank they targeted. They are known to have:

- Programmed ATMs to cash out money without any interaction.
- Transferred money to mule accounts.
- Used the SWIFT network to inject financial transactions.
- Create fake bank accounts with a high balance.

The losses were between \$2 million and \$10 million per bank and could be as high as \$1 billion in total.

**References:**

<https://securelist.com/the-great-bank-robbery-the-carbanak-apt/68732/><https://en.wikipedia.org/wiki/Carbanak>




## The Bangladesh Bank Heist

In 2016, a cyber attack occurred against "Bangladesh Bank"



In 2016, adversaries obtained access to the SWIFT payment system and instructed the New York Federal Reserve bank to transfer money from BB's accounts to accounts in the Philippines



Highly targeted, possibly state-sponsored, attack that manipulates SWIFT transaction messages and attempts to hide itself. It is suspected that the Dridex malware was used as part of the attack.

Transactions for up to \$951 million were attempted, but "only" \$63 million was eventually stolen

### The Bangladesh Bank Heist

The Bangladesh bank heist of 2016 is a notorious digital attack via the SWIFT network on the Bangladesh Bank account at the Federal Reserve Bank of New York.

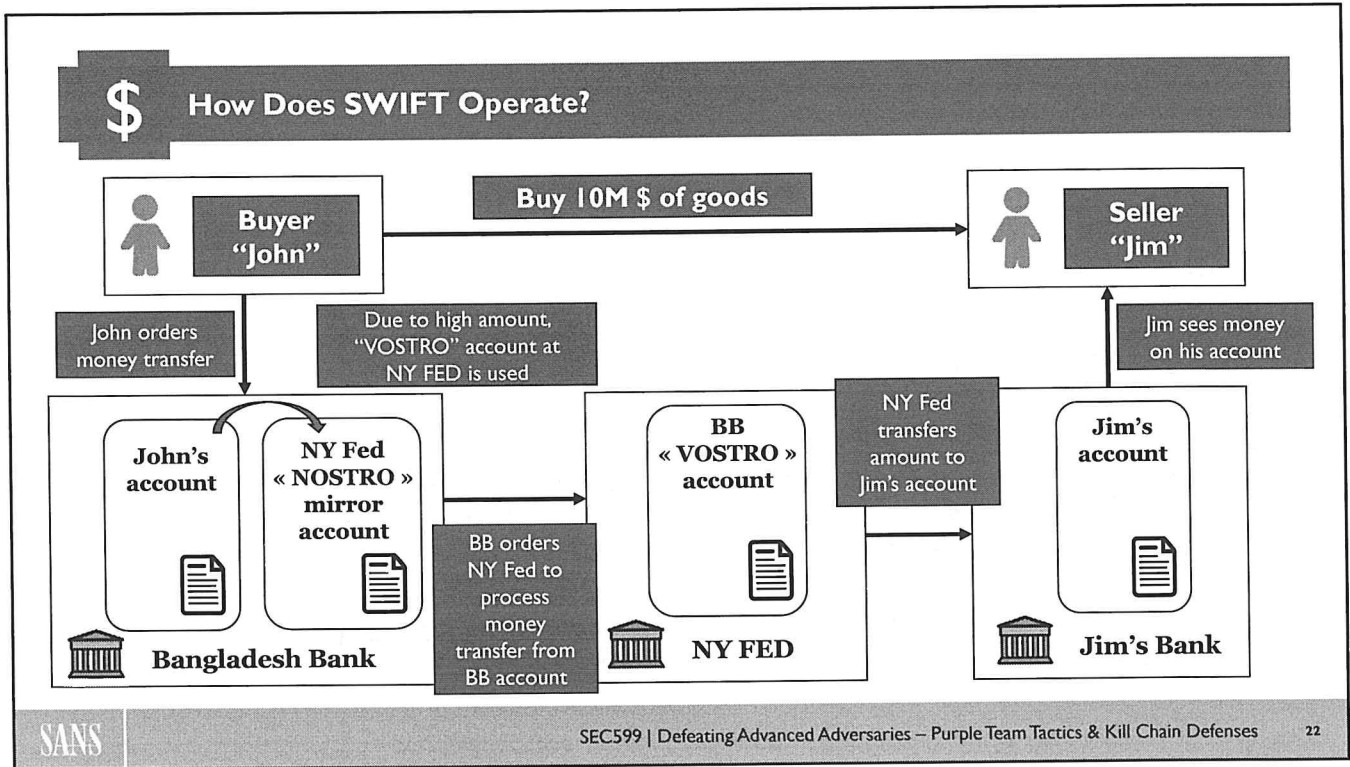
Adversaries outside of Bangladesh used the Dridex malware to compromise computer systems of the Bangladesh Bank, possibly with the help of insiders. Dridex gave them the capabilities to observe the operations of the bank regarding international payments and money transfers. Adversaries install SysMon on SWIFT systems as a reconnaissance tool, helping them to understand how the SWIFT network operates and how the bank employees operated the SWIFT network to execute financial transactions.

The installed malware would manipulate SWIFT messages through PRT files and Printer Command Language, allowing the adversaries to generate 35 fraudulent SWIFT messages for a total of 951 million USD. The Federal Reserve Bank of New York blocked 30 suspicious SWIFT transactions but let five of them through. These remaining five transactions resulted in the loss of 101 million USD of the Bangladesh Bank account at the Federal Reserve Bank of New York: 20 million USD were transferred to Sri Lanka and 81 million USD to the Philippines. The money transferred to Sri Lanka was blocked through a typo, after which 18 million USD from the Philippines were also recovered.

In total, the criminals managed to successfully steal 63 million USD (81 million USD – 18 million USD recovered). Most of this money was quickly laundered through casinos in the Philippines.

#### References:

- <https://www.nettitude.com/wp-content/uploads/2016/12/Nettitude-SWIFT-Threat-Advisory-Report-client.pdf>
- [https://en.wikipedia.org/wiki/2016\\_Bangladesh\\_Bank\\_heist](https://en.wikipedia.org/wiki/2016_Bangladesh_Bank_heist)
- <https://baesystemsai.blogspot.com/2016/04/two-bytes-to-951m.html>
- <https://www.niceideas.ch/roller2/badtrash/entry/deciphering-the-bengladesh-bank-heist>



### How Does SWIFT Operate?

In order to understand the attack that took place against Bangladesh Bank, we first need to understand how SWIFT plays a role in the international payment system. Let's assume we have a buyer, John, who is buying something from a seller, Jim. John wants to pay Jim 10M \$ for this. Unless John is a mobster, he would most likely not pay Jim 10M \$ in cash. So what happens? We will illustrate this using the examples of Bangladesh Bank and NY Fed to make it relevant to the case.

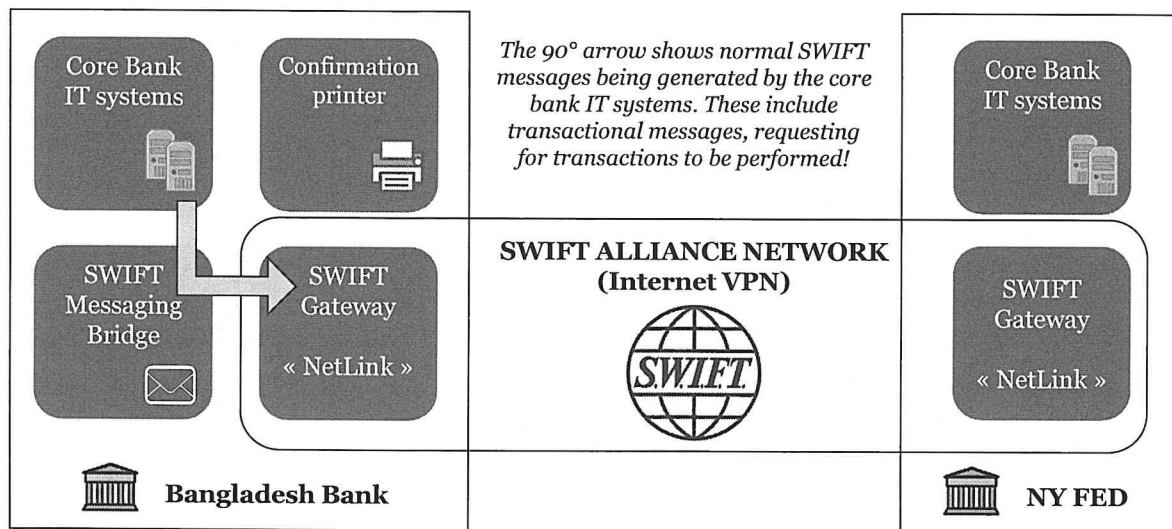
The following steps would take place:

1. As a first step, John will ask his bank (Bangladesh Bank) to perform the transfer of 10M \$
2. Due to the high amount involved, this transaction would be performed through a "VOSTRO" / "NOSTRO" setup, where an intermediary bank is used to handle the transaction. There's a few reasons why such an intermediary bank setup is used:
  - Because the domestic bank has limited access to foreign financial markets;
  - The intermediary bank acts as an intermediary between banks in different countries or as an agent to process local transactions for customers abroad;
  - The intermediary bank accepts deposits, processes documentation and serves as transfer agents for transactions.
3. Through the "VOSTRO" / "NOSTRO" setup, Bangladesh Bank (John's bank) instructs NY FED (intermediary) to process a money transfer from BB's account to Jim's bank;
4. NY FED debits the BB VOSTRO account and transfers the transaction amount to Jim's Bank;
5. Jim sees that money was added to his account.

The intermediary setup is where a provider such as SWIFT comes into play!



## SWIFT Technical Architecture – Bangladesh Bank Case (I)



### SWIFT Technical Architecture – Bangladesh Bank Case (I)

In this first diagram, we can see how Bangladesh Bank is connected to the NY Fed. Please note that this is a slightly simplified diagram, as we don't want to lose ourselves too much in banking terminology. In Bangladesh's environment, let's assume there are 4 main components:

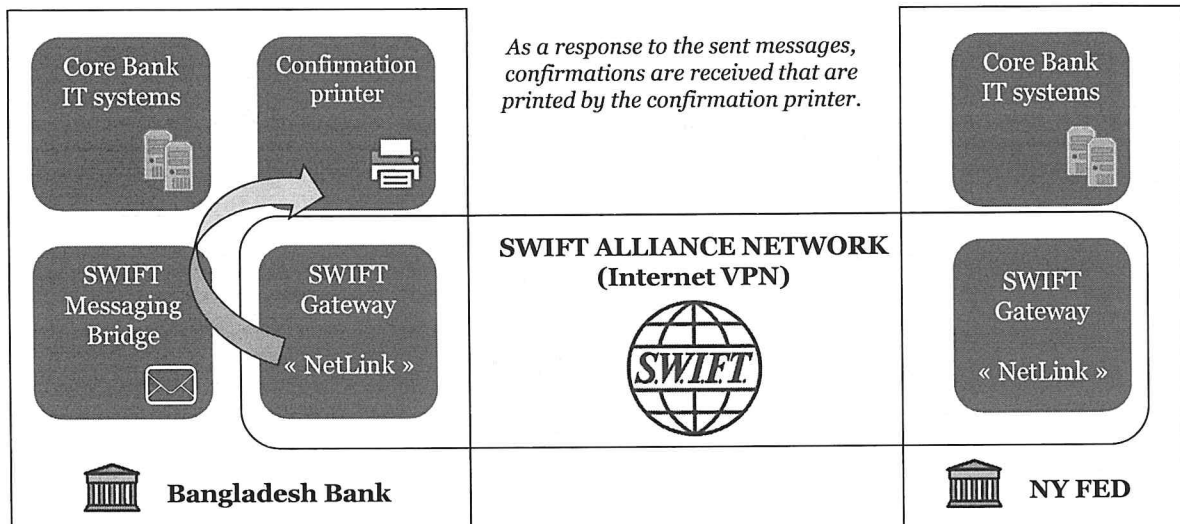
- The core bank IT systems, which are traditional transactional systems of the bank.
- The SWIFT Messaging bridge, which generates SWIFT messages that are sent across the SWIFT network (e.g. requests for transactions).
- The SWIFT Gateway, which is used to set up connectivity toward other banks using SWIFT by using the NetLink software.
- A confirmation printer, which is used to print out SWIFT messages (e.g. transaction confirmations), which can be used for validation.

The SWIFT network between the different banks is set up using VPN connectivity over the internet to connect the different SWIFT Gateways. The 90° arrow on the diagram shows normal SWIFT messages being generated by the core bank IT systems. These include transactional messages, requesting for transactions to be performed!





## SWIFT Technical Architecture – Bangladesh Bank Case (2)

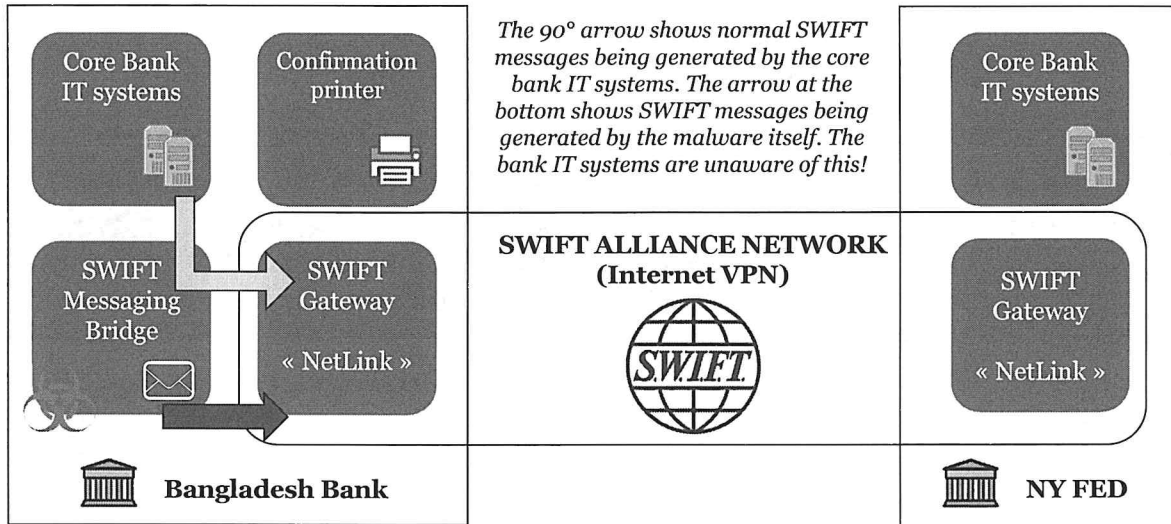


### SWIFT Technical Architecture – Bangladesh Bank Case (2)

In a continuation of the previous flow, a response is now received from the SWIFT network, which is passed through the SWIFT Messaging Bridge and finally printed by the confirmation printer.



## The Bangladesh Bank Heist – The Intrusion (1)

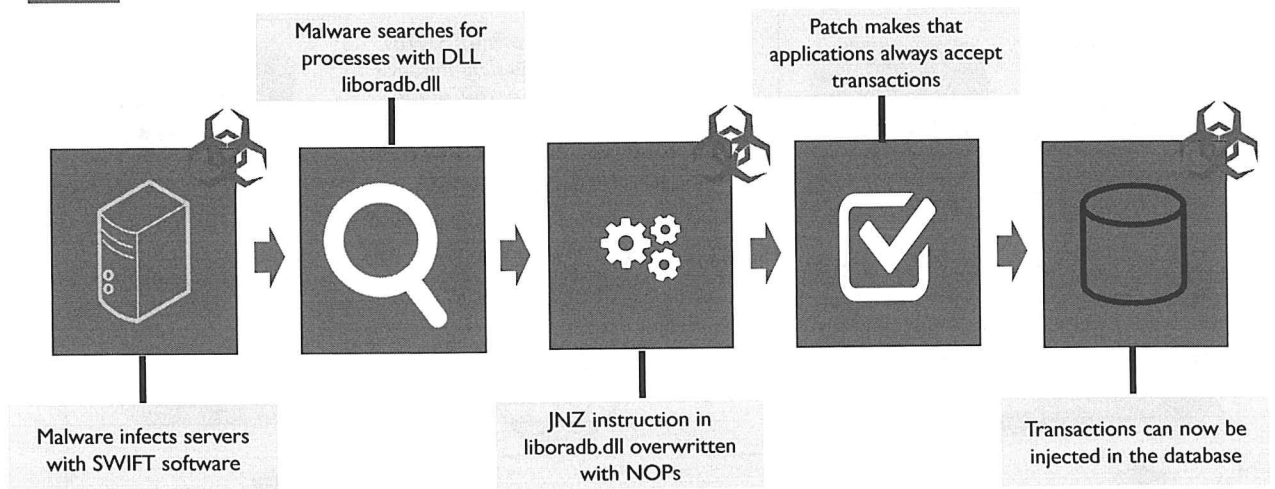


### The Bangladesh Bank Heist – The Intrusion (1)

During the attack, adversaries managed to compromise systems on which the SWIFT Messaging Bridge software was running, where they injected fraudulent SWIFT messages. It's important to note that the bank's own IT systems are "blissfully" unaware of this. This is a direct injection of transactions into the SWIFT network. We will zoom in on how the malware pulled this off in the next slide.



## The Bangladesh Bank Heist – Zooming in on the Malware (1)



### The Bangladesh Bank Heist – Zooming in on the Malware (1)

The malware is used to infect Bangladesh Bank's servers running SWIFT Alliance software. This software is responsible for the processing and managing of SWIFT messages. It is complex software that performs many checks to validate transactions. The malware will change the behavior of the validations of the SWIFT software.

When executing on the server, the malware will check all processes running on the Windows OS and enumerate all modules loaded by processes. Modules are .exe files, .dll files and data files. The malware looks for a particular dll loaded inside a process: liboradb.dll. This DLL is part of SWIFT's Alliance software and performs the following tasks:

- Reading the Alliance database path from the registry
- Starting the database
- Performing database backup and restore functions

In each process that loads DLL liboradb.dll, the malware will patch the DLL in memory by replacing a particular JNZ instruction with 2 NOP instructions. Due to this change, the checks performed by the SWIFT software will always succeed: Counterfeit transactions will now be accepted. Patching the DLL in memory has the advantage for the adversaries that it will not be detected by doing an integrity check of the software's files and that it does not invalidate SWIFT's digital signature of the DLL.

Once the SWIFT software has been patched in memory, the criminals can create counterfeit SWIFT messages and inject them into the database without having to get all the details and checks right.



## The Bangladesh Bank Heist – Zooming in on the Malware (2)

Original code in DLL liboradb.dll: The validation function returns 0 upon success and 1 upon failure.

```
85 C0      test eax, eax ; important validation
75 04      jnz failed  ; if failed, jump to label failed
33 C0      xor  eax, eax ; otherwise, set result to 0 (success)
EB 17      jmp  exit   ; and then exit

        failed:
B8 01 00 00 00 mov  eax, 1  ; set result to 1 (failure)
        exit:
C3                ret                ; return to caller
```

Patched code in DLL liboradb.dll: The validation function always returns 0 (success)

```
85 C0      test eax, eax ; important validation
90 90      nop, nop    ; 'no operation' replacing 0x75
33 C0      xor  eax, eax ; always set result to 0 (success)
EB 17      jmp  exit   ; and then exit

        failed:
B8 01 00 00 00 mov  eax, 1  ; never reached: set result to 1 (fail)
        exit:
C3                ret                ; return to caller
```

### The Bangladesh Bank Heist – Zooming in on the Malware (2)

In this slide, we see assembly code similar to code we would find in DLL liboradb.dll. On the top, we see the original code and at the bottom, the patched code.

The bytes we see at the left of the listings are the bytecodes of the assembly instructions. After that comes the assembly instructions themselves, followed by comment (everything starting with the semicolon).

For example, on the first line, we have instruction "test eax, eax". This is an x86 instruction to perform a test on the value in register eax. This test instruction is encoded with 2 bytes: 0x85 and 0xC0.

On the second line, we have a conditional jump instruction: Jump if Not Zero. The Zero Flag is set by the previous test instruction (can be set to 0 or 1), and the jnz instruction will jump to label failed (4 bytes further) if the zero flag is not set and will not jump (e.g. move on to the next instruction, xor) if the zero flag is set.

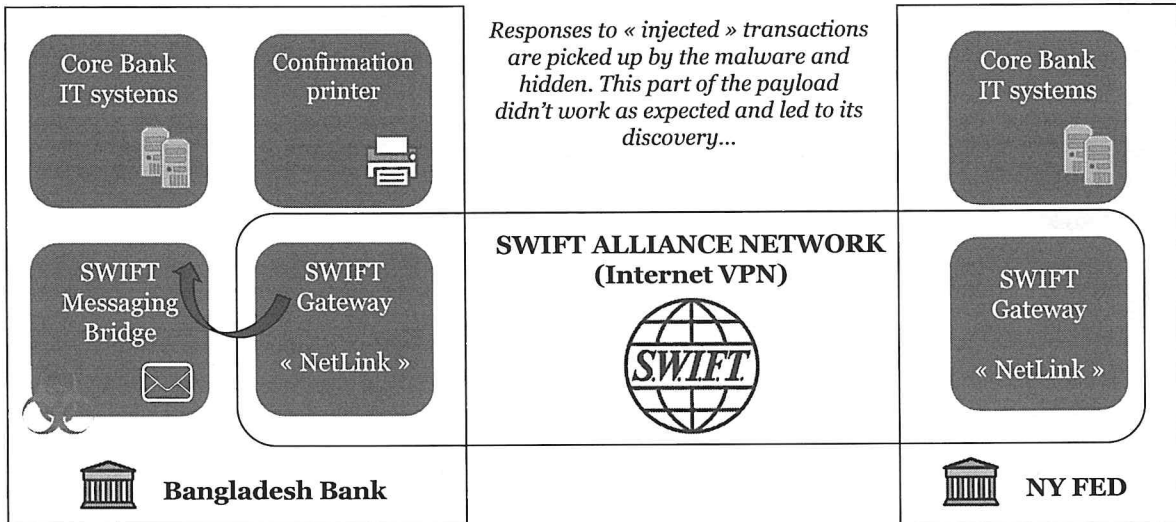
Instruction "xor eax, eax" is a trick to set register eax to 0 with a shorter instruction (xor) than a move instruction (mov). mov eax, 0 is valid too but takes 5 bytes instead of 2 for xor eax, eax. Compilers typically use xor when they have to set a register to 0, and not mov.

So, to make that this function always returns success (0), the malware authors have to remove the jnz instruction. But just deleting those 2 bytes is a problem, as this would imply that all subsequent bytes have to be shifted 2 positions, and this would also break jump locations. What is typically done to remove instructions in machine code without changing the position to the remaining instructions, is to replace the instructions with instructions that do nothing. The x86 instruction set has an operation just for that No Operation, NOP. This instruction is 1 byte long (0x90). Hence to replace "jnz failed", with instructions that do nothing, we have to replace its 2 bytes (0x75 and 0x04) with 2 nop instructions (0x90 0x90).

Patching machine code by replacing instructions with nop instructions is a popular technique.

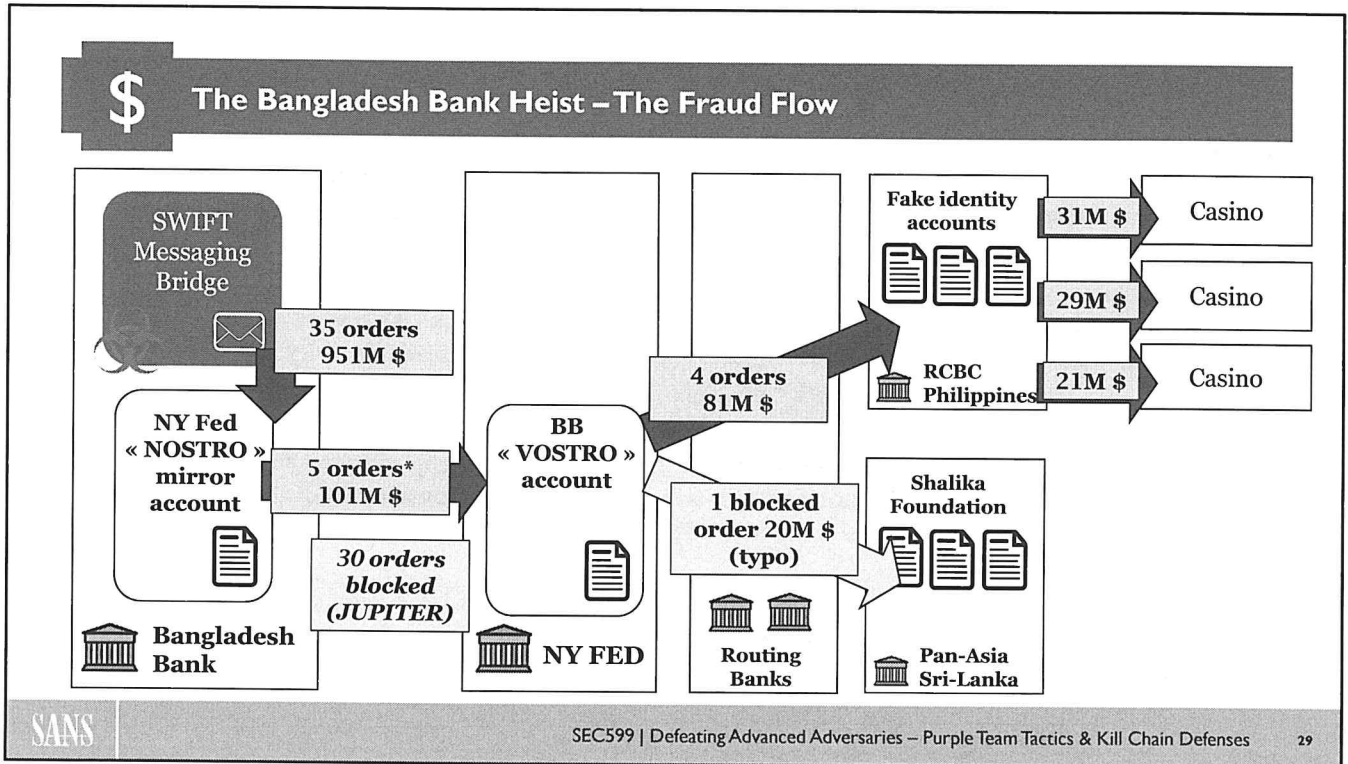


## The Bangladesh Bank Heist – The Intrusion (2)



### The Bangladesh Bank Heist – The Intrusion (2)

As a next step in the intrusion chain, any confirmations that are received from the SWIFT gateway are intercepted by the malware infection and are thus not printed. It's interesting to note that this part of the attack did not go as planned, as the confirmation printer malfunctioned (and did not print any transactions at all), which was a reason for suspicion / investigation. Once the printer was "fixed", it started printing out the backlog of transactions, including the injected ones. Due to the careful planning and timing of the attacks, though, the adversaries still managed to get away with a number of transactions.



### The Bangladesh Bank Heist – The Fraud Flow

When looking at the overall picture of the fraud attempts, here's a few interesting items to note:

- The adversaries initially attempted to inject 35 transactions, for a total value of 951M USD.
  - Of these 35 transactions, 30 were immediately blocked because they included the keyword "Jupiter". This keyword was part of the address of one of the Philippines banks, but was blocked by the NY FED. It was never the intention of NY FED to block this specific bank; this was more of a "lucky shot." By chance, "Jupiter" was also the name of an oil tanker that was banned by US sanctions against Iran. Due to this hit, the transactions were heavily scrutinized and finally rejected.
- 5 transactions for a total value of 101M USD were executed by the NY FED.
  - Of these 5 transactions, 4 transactions succeeded, and were used to transfer money to 3 accounts created at Rizal Commercial Banking Corporation (RCBC) in the Philippines. These 3 accounts had been created in 2015 already but were never really used before.
  - Of these 5 transactions, 1 transaction was blocked, as it included a typo ("Shalika foundation" was written as "Shalika fandation"). This made Deutsche Bank block the transaction and request further details from Bangladesh Bank.
- Finally, the 81M USD that was transferred to RCBC was further forwarded to different accounts that belonged to Casinos. From there, withdrawals were made to extract the money.



## The Bangladesh Bank Heist – Key Takeaways

The overall cyber security posture of the Bangladesh Bank network environment was shockingly low for a financial services institution (no segmentation, use of low-cost (secondhand) network infrastructure...).

A secure network such as "SWIFT" could still be exposed to attacks, as it's only as secure as its weakest link(s). This was the main reason for SWIFT to start its Customer Security Program (CSP). Note that the Bangladesh Bank Heist did not involve a "direct attack" against SWIFT.

The Bangladesh Bank heist was well-prepared, as the malware was highly targeted and the fraudulent transactions were performed at specific times where response times of the involved banks would be lower (close to the weekend, Bangladesh bank holiday, Philippines bank holiday...)

The adversaries had obtained access to the Bangladesh Bank network for quite some time before executing their attack. This is one of the main reasons why the initial infection vector is hard to identify. Improved monitoring would have allowed detection of the intrusion (much) sooner!

### The Bangladesh Bank Heist – Key Takeaways

Here's a few interesting takeaways with regards to the Bangladesh Bank Heist:

- The overall cyber security posture of the Bangladesh Bank network environment was shockingly low for a financial services institution (no segmentation, use of low-cost (secondhand) network infrastructure...).
- A secure network such as "SWIFT" could still be exposed to attacks, as it's only as secure as its weakest link(s). This was the main reason for SWIFT to start its Customer Security Program (CSP). Note that the Bangladesh Bank Heist did not involve a "direct attack" against SWIFT.
- The Bangladesh Bank heist was well-prepared, as the malware was highly targeted and the fraudulent transactions were performed at specific times where response times of the involved banks would be lower (close to the weekend, Bangladesh bank holiday, Philippines bank holiday...).
- The adversaries had obtained access to the Bangladesh Bank network for quite some time before executing their attack. This is one of the main reasons why the initial infection vector is hard to identify. Improved monitoring would have allowed detection of the intrusion (much) sooner!



## Cyber Crime – Some Ransomware Families

New ransomware families are popping up on a frequent basis; we list some interesting examples with particular behavior below:

Name	First appeared?	Specifics?
Locky	2016	Can leverage different exploit kits, highly flexible
Cerber	2016	Includes non-typical ransomware features like DDoS attacks
Jigsaw	2016	Both steals and encrypts your data, focuses on "victim service"
Crysis / LeChiffre	2015	Uses RDP brute forcing to obtain access to target systems
Goldeneye / Petya / HDDCryptor	2016	If run with administrative privileges, will encrypt entire drive and overwrite Master Boot Record
Popcorn Time	2016	"Infect-a-friend" in exchange for decryption key
Wcry	2017	Uses ShadowBrokers SMB exploit to spread in worm-like fashion
(Not)Petya	2017	Uses ShadowBrokers SMB exploit, Mimikatz-like features + PSEXEC & WMIC

### Cyber Crime – Some Ransomware Families

Due to its high effectiveness, new ransomware families are popping up on a very frequent basis these days. We list some interesting variants in this slide:

- Locky is one of the most popular ransomware samples out there. It's highly flexible and can be delivered using multiple exploit kits (drive-by downloads), or just using the traditional phishing scheme.
- Cerber is not your typical ransomware: It also includes other features / attack methods such as DDoS support.
- Jigsaw (themed like the movie "Saw") doesn't limit itself to only encrypting your data: It also steals it!
- The Crysis & LeChiffre ransomware variants have something interesting in common: They use brute force attacks against Windows RDP (Remote Desktop Protocol) to obtain access to victim systems (instead of the usual phishing techniques).
- Goldeneye, Petya and HDDCryptor attempt to not only encrypt individual files: When run with administrative privileges, they will attempt to encrypt the entire hard drive and overwrite the Master Boot Record.
- Popcorn Time implements the interesting "infect-a-friend" function, where victims receive the decryption key for free provided, they infect a number of other users / friends.
- Wcry caused a major impact in May 2017, holding several large organizations hostage. The "innovative" part of the attack was the use of an SMB exploit (published by the ShadowBrokers) to spread the ransomware throughout victim networks.
- (Not)Petya rose to stardom in June 2017, as it impacted several large organizations. While also relying on an SMB exploit (published by the ShadowBrokers) it coupled this with a highly effective combination of Mimikatz-like techniques (to steal credentials) and PsExec / WMIC to perform lateral movement. Several experts claim that the ransomware part of the malware was only a distraction of its actual intent, which was to cause as much downtime / damage as possible.






## Cyber Crime – Crypto-Mining



After the booming rise of cryptocurrency in late 2017, we started seeing an interesting new trend: Malware that abuses system resources to mine cryptocurrency.



As opposed to ransomware, crypto-mining malware is stealth and non-disruptive, thereby hoping to remain on victim systems for as long as possible.



Although seemingly "innocent", crypto-mining malware families exist that exhibit worm-like behavior, password stealing and general information stealing.

### Cyber Crime – Crypto-Mining


While ransomware is a highly disruptive (and obvious) type of malware, another type of malware has started to rise. After the booming rise of cryptocurrency in late 2017, we started seeing an interesting new trend: Malware that abuses system resources to mine cryptocurrency.

As opposed to ransomware, crypto-mining malware is stealth and non-disruptive, thereby hoping to remain on victim systems for as long as possible. The malware authors obviously try to have their malware running for as long as possible. Two main categories of crypto-mining malware exist: Browser-based and host-based.

- Browser-based crypto-mining is typically implemented using JavaScript and is hosted on a website that is visited by the victim. Once the victim leaves the website, the crypto-mining activity stops. Websites that often exhibit this behavior include streaming websites, as visitors generally keep these websites open for a long time.
- Host-based crypto-mining usually involves a typical malware delivery scheme (e.g. a phishing mail), after which a crypto-mining malware is persisted on the system (e.g. using PowerShell).

Although seemingly "innocent", crypto-mining malware families exist that exhibit worm-like behavior, password stealing and general information stealing.


## Zooming in on Sabotage



**Cyber Crime**



**Sabotage**




**Espionage**

### **Zooming in on Sabotage**

As a second attack motivator, let's see how sabotage is typically performed in cyber space!



## Stuxnet – The World's First Digital Weapon



In 2010, a Windows-based virus was identified that targets Siemens industrial control systems. The virus was labeled "Stuxnet" and mainly appeared in organizations related to Iran's uranium enrichment infrastructure.

Stuxnet reportedly targeted five different Iranian organizations and destroyed about 10% of the country's enrichment centrifuges

Due to its highly complex nature and its specific target, its development is believed to have been supported or coordinated by **nation states**

### Stuxnet – The World's First Digital Weapon

Since the 1950s, Iran has pursued a nuclear program with the support of Western countries such as the United States of America. The goal of this program is the production of electricity via nuclear energy. After the Iranian Revolution in 1979, Western countries started to express doubts that the Iranian nuclear program had solely peaceful goals and thus revoked its support.

The nuclear material used in power plants and in nuclear weapons is different but can be produced in similar nuclear facilities. In 1968, Iran signed the Nuclear Non-Proliferation Treaty, thereby accepting inspections from the International Atomic Energy Agency. As the same nuclear facilities can be used to produce weapons-grade nuclear material, the IAEA inspects these facilities to ascertain that they are not misused to produce illegal nuclear material suitable for weapons. But in 2003, the IAEA launched an investigation after it received information of illegal activities in Iran's nuclear facilities. Iran opposed these inspections and has since then been in conflict with the IAEA and western countries.

Stuxnet is malware developed to disrupt Iran's nuclear program (59% of all Stuxnet infections were in Iran). It is generally agreed the virus' development was supported or even coordinated by nation states.

#### References:

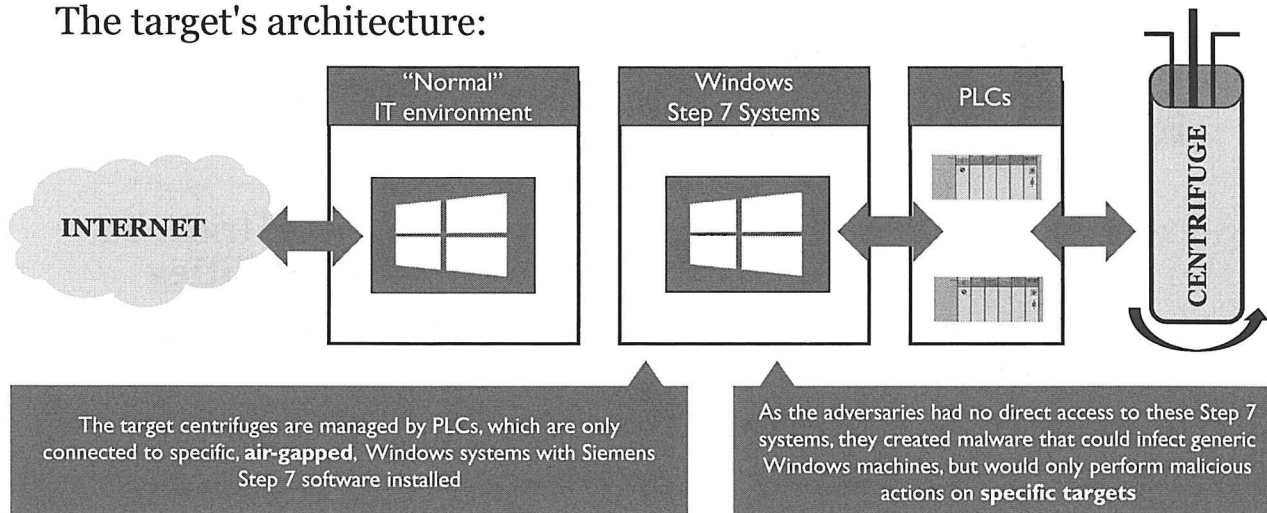
<https://en.wikipedia.org/wiki/Stuxnet>

<https://www.wired.com/2014/11/countdown-to-zero-day-stuxnet/>



## Stuxnet – Zooming in on the Attack (I)

### The target's architecture:



### Stuxnet – Zooming in on the Attack (I)

In order to enrich / weaponized uranium, the nuclear centrifuges have to spin at the exact right frequency. In order to control or disrupt the uranium enrichment process, the adversaries would thus have to obtain control over the PLCs (Programmable Logic Controllers). These PLCs are not connected to the internet and cannot be attacked directly. The PLCs are programmed via Siemens' Step 7 software on Windows computers. For obvious reasons, these Windows computers also are not connected to the internet.


Because the adversaries had no physical access or internet connection to the Windows computers programming the PLCs, they decided to follow an elaborate plan to infect thousands of computers in the hope to reach the target computers.

It is not known how Stuxnet was initially propagated to Windows computers to start the chain of infections, but the Windows component has several attack vectors (including infection capabilities using USB drives, network connectivity ...). Stuxnet will achieve persistence on Windows computers but will remain dormant unless the infected Windows machine runs Siemens' Step 7 PLCs programming software.

If the Windows computer runs Step 7, Stuxnet will proceed to infect the attached PLCs. But again, it does not do this indiscriminately; Stuxnet will only infect Siemens PLCs of a particular model and with particular modules attached to it.



## Stuxnet – Zooming in on the Attack (2)



As explained in the previous slides, the key target for Stuxnet will be Windows-based systems. So, what techniques did Stuxnet use for its infection?

Stuxnet does not only rely on simple "social engineering" tricks but also relies on **four (4!) zero-day Windows vulnerabilities**

This is uncommon, as the use of zero-days is considered to be an expensive investment

This further supports the theory that Stuxnet was developed by a highly determined, well-funded, adversary

### Stuxnet – Zooming in on the Attack (2)

To reach the target computers in Iran's nuclear enrichment facilities, Stuxnet infected computers worldwide. More than half of the infected computers were located in Iran, but many computers were infected in countries like Indonesia and India, too.

Stuxnet used several exploits to infect computers; several of them were zero-days. A zero-day is an exploit that is not publicly known and for which there is no patch. Usually, Windows zero-day exploits that achieve code execution are valuable, and it is rare to see them used in common malware. As reliable zero-day exploits can command prices from tens of thousands to hundreds of thousands of dollars, malware authors tend to use them sparingly. Using a zero-day in malware exposes it to discovery and ultimately patching, thereby significantly reducing its utility to the adversaries and thus its worth.


That is why the use of zero-days in malware is remarkable. Using 4 Windows zero-days, like Stuxnet did, is unprecedented. Many researchers believe that this indicates that the adversaries had vast resources at their disposal and were very determined to achieve their goal. This is another argument for attributing Stuxnet to a nation-state actor like the United States of America and Israel.

Two zero days are used as a propagation vector: The .lnk file vulnerability and the printer spooler vulnerability. Both achieve code execution without user interaction. The other zero-day exploits achieve privilege escalation, allowing the malware to run with the highest privilege and infect the Windows kernel.



## Stuxnet – Zooming in on the Attack (3)


### Presenting Stuxnet's zero-day arsenal:




MS10-046: Allow automatic execution of DLLs on USB sticks through malformed .lnk files



MS10-061: Vulnerability in the Print Spooler Service allows Remote Code Execution over the network



MS10-073: Vulnerabilities in Windows Keyboard Layout allow local privilege escalation



MS10-092: Vulnerability in the Task Scheduler allows privilege escalation to SYSTEM

### Stuxnet – Zooming in on the Attack (3)

#### Presenting Stuxnet's four zero-days:

- **MS10-046: Allow automatic execution of DLLs on USB sticks through malformed .lnk files**  
Infecting computers via portable media like USB sticks is a common practice but has become less practical since Microsoft started to change the autorun behavior of Windows. On old versions of Windows, removable storage could be configured to execute programs stored on the medium automatically upon connection of the removable media with the computer. This behavior has changed in modern versions of Windows, and the user is always warned before programs autorun, with the option to prevent execution.

Stuxnet exploits a vulnerability in the parsing of Windows Shortcut files (.lnk) to achieve code execution without user interaction. Due to a bug in Windows Explorer, DLLs present on the USB stick can be loaded and executed inside the Windows Explorer process when they are referenced in a particular way in the .lnk file. A DLL is a Windows library with executable code. By putting a malicious DLL file on a USB stick together with a malformed .lnk file exploiting this vulnerability, adversaries could achieve code execution on a Windows computer merely by inserting a USB stick and viewing the drive in Windows Explorer.

- **MS10-061: Vulnerability in the Print Spooler Service allows Remote Code Execution**  
Although the target Windows computers were air-gapped from the internet, they were nevertheless connected via an IP/Ethernet network. When a Windows computer is connected to a trusted IP network, it will expose many services to be consumed by peer computers on the network. One of these services is the print spooler service, designed to share an attached printer with other Windows computers.

The zero-day vulnerability inside the print spooler service allowed adversaries to have an infected computer connect to the print spooler of another Windows computer on the network and achieve remote code execution. Malware that self-propagates via networked computers without any user interaction is called a worm and is considered very potent malware that can spread blindingly fast.

The zero-days described above would only provide limited, unprivileged, access to Windows systems. In order to reach its full potential, Stuxnet also included two zero-days that could help it escalate privileges:

- MS10-073: Vulnerabilities in Windows Keyboard Layout allow local privilege escalation. A privilege escalation vulnerability existed due to the way that the Windows kernel-mode drivers maintain the reference count for an object. An attacker who successfully exploited this vulnerability could run arbitrary code in kernel mode. An attacker could then install programs; view, change, or delete data or create new accounts with full user rights.

Stuxnet abused this flaw to escalate its privileges on infected systems.

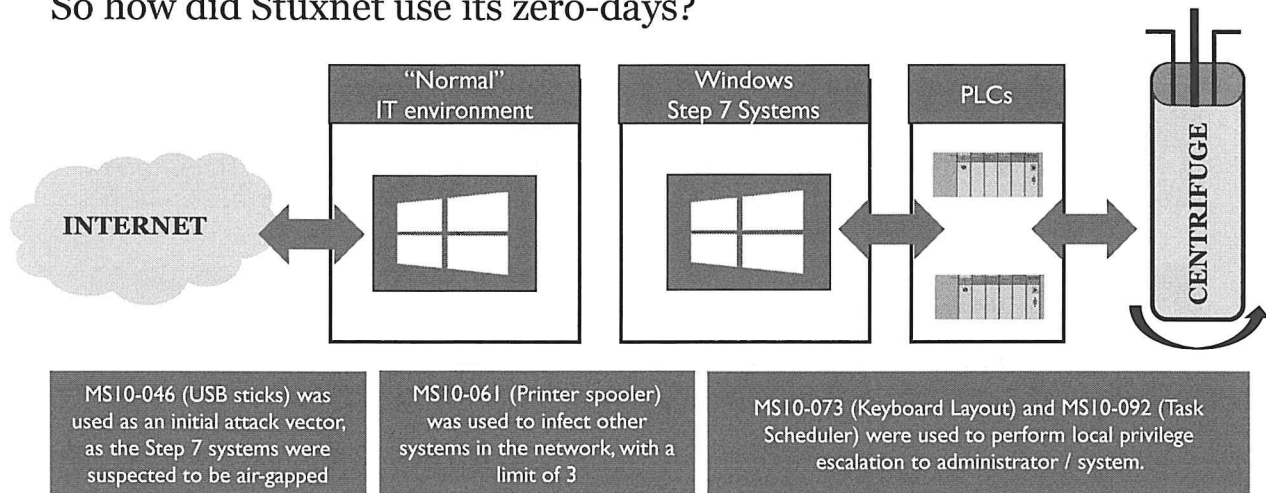
- MS10-092: Vulnerability in the task scheduler allows privilege escalation to SYSTEM. When processing task files, the Windows Task Scheduler only used a CRC32 checksum to validate that the file has not been tampered with. Also, in a default configuration, normal users can read and write the task files that they have created. By modifying the task file and creating a CRC32 collision, an attacker can execute arbitrary commands with SYSTEM privileges.

Stuxnet abused this flaw to escalate its privileges on infected systems.



## Stuxnet – Zooming in on the Attack (4)

So how did Stuxnet use its zero-days?



SANS

SEC599 | Defeating Advanced Adversaries – Purple Team Tactics & Kill Chain Defenses

39

### Stuxnet – Zooming in on the Attack (4)

Stuxnet's zero-day arsenal was used in the following way:


- In order to overcome the "airgapped network" problem, Stuxnet used MS10-046 in an attempt to infect target systems through infected USB sticks.
- In order to further infect other systems (e.g. in the airgapped network), MS10-061 was used to create a worm that would infect Windows systems through the typically exposed Printer Spooler Service. This worm-like behavior could spiral out of control, which is something the Stuxnet developers took into account: Stuxnet would not spread to more than 3 machines and erased itself after 24 June 2012.
- After achieving code execution on a Windows machine, Stuxnet needs to obtain full permissions on the Windows machine and achieve persistence. When code is executed via the .lnk vulnerability or the printer spooler vulnerability, it is not running with full permissions. The code is running in the context of a restricted user and needs to run in system context to fully compromise the host Windows machine. To obtain system-level access on infected hosts, Stuxnet used the MS10-073 and MS10-092 zero-days.





## Stuxnet – Zooming in on the Attack (5)

### Making sure it sticks... Achieving persistence!



In order to achieve persistence, malware typically relies on rootkits that hide its presence. Stuxnet used malicious device drivers for this purpose!



The device drivers used by Stuxnet were digitally signed with a digital code-signing certificate that was first stolen from a number of Taiwanese companies (amongst others, Realtek). This is another clear artifact showing the persistence and expertise of the Stuxnet developers.

#### Stuxnet – Zooming in on the Attack (5)

Upon obtaining SYSTEM-level access to target machines, Stuxnet proceeds to install rootkits (user and kernel rootkits) so that it could hide its presence on the infected machine. A rootkit is malware designed to conceal the presence of malware. For example, a rootkit might hide certain files when a user executes a directory listing command.

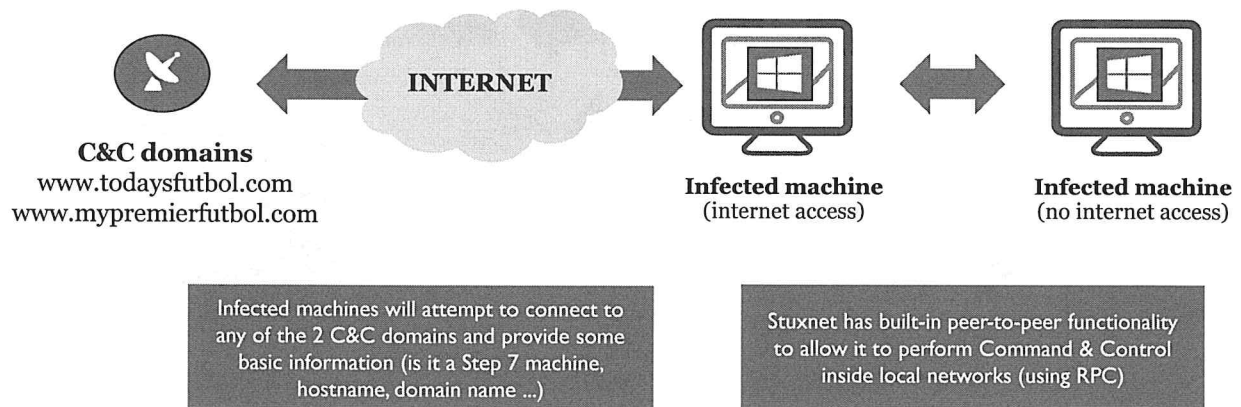
The kernel rootkit was installed via device drivers. On Windows, device drivers need to be digitally signed before they can be installed. The authors of Stuxnet obtained 2 stolen digital code-signing certificates from Taiwanese companies (amongst others, the rather well-known Realtek). These certificates were used by the Stuxnet developers to sign their own malicious, device drivers.

At this stage, Stuxnet looks for Step 7 software on the infected machine and the actual “attack” can start!



## Stuxnet – Zooming in on the Attack (6)

So how did Stuxnet perform Command & Control?



### Stuxnet – Zooming in on the Attack (6)

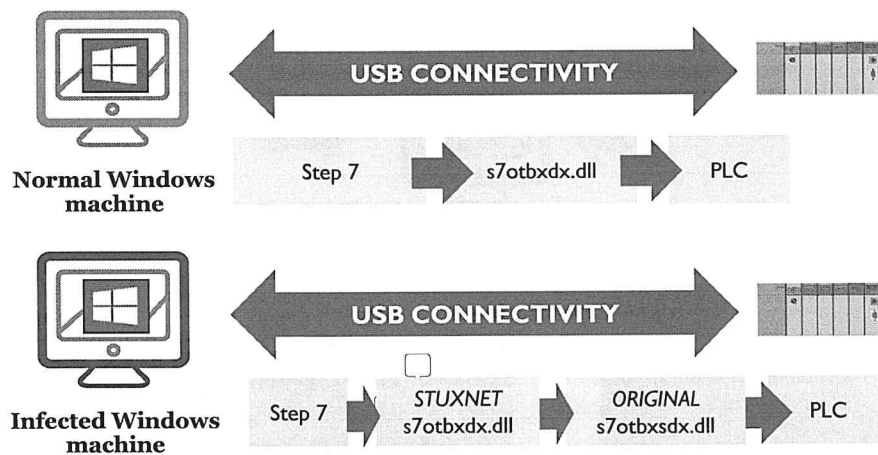
Stuxnet was designed to operate without a fine-grained C&C infrastructure (as it needs to operate in air-gapped networks). That doesn't mean it has no C&C infrastructure:

- Two main C&C domains were used (www.todaysfutbol.com and www.mypremierfutbol.com), to which infected systems would send some initial information. This included, for example, the hostname & domain name of the system, but also a Boolean to indicate the system had Step 7 installed or not;
- Stuxnet-infected systems will also run an RPC server, which is used for peer-to-peer communications between infected hosts. This allows machines that are not connected to the internet to receive updates and exfiltrate information (if it has peer-to-peer connectivity with other systems that are connected to the internet).



## Stuxnet – Zooming in on the Attack (7)

### Moving from Windows to the ICS world:



#### DID YOU KNOW THAT?

Stuxnet hijacked the used communication library to install hidden malware on the PLC:

- Add own malicious code (STL) by changing project files
- Rename the original communication library and insert its own

When engineers read the STL code from the PLC, the inserted communication library will hide the malicious parts

### Stuxnet – Zooming in on the Attack (7)

Siemens' PLCs need to be connected via a data cable to a Windows machine running Step 7 software to be programmed. When connected, Step 7 will communicate with the PLC via a communication library (DLL s7otbxdx.dll).

On Windows computers with Step 7, Stuxnet will modify Step 7 project files to inject code and hijack the communication library to install hidden malware on the PLC:

- By modifying the project files, Stuxnet can inject its own STL code (Siemens' PLC programming language, Statement List) into the PLC.
- By renaming the original communication library s7otbxdx.dll to s7otbxsdx.dll, and inserting its own malicious communication library as s7otbxdx.dll, Stuxnet can interfere with the communication between Step 7 and the PLC.

Under normal circumstances, Step 7 can read an STL code block from the PLC by calling a function in the communication library s7otbxdx.dll to read a particular code block. This allows Step 7 to retrieve the program code of a PLC and have a programmer inspect and/or modify the code.

As this would potentially reveal malicious code installed on the PLC, the Stuxnet developers wanted to prevent this. Therefore, they inserted their own communication library between Step 7 and the original communication library. When Step 7 would want to retrieve a particular STL code block, it would call a function in communication library s7otbxdx.dll (the adversaries' library), which would pass it on to the original communication library which in turn would retrieve it from the PLC over the data cable. If the retrieved STL code block would contain malicious code (implanted by Stuxnet), the adversary's communication library would modify the STL code block to hide the malicious code before returning it to Step 7.



## Stuxnet – Not Everyone's Friend

In order to remain stealth, Stuxnet was highly targeted:

- Stuxnet only attacks the right Siemens PLC's (S7-300)
- On these specific PLC's, a number of particular modules had to be present (variable frequency drives)
- Spinning frequency of the attached motors had to be exactly between 807 Hz and 1210 Hz
- Previously infected systems were identified and not "double infected"

### DID YOU KNOW THAT?

The main target of Stuxnet was the nuclear facility / plant of Natanz.

Stuxnet was discovered "by accident", due to a mistake, an engineer who was connected to the centrifuges got infected.

When he returned home and connected his machine to the network, it started further spreading and "escaped"!

### Stuxnet – Not Everyone's Friend

To achieve its goal while remaining stealth, Stuxnet would only infect specific PLCs that could possibly be used in Iran's uranium enrichment facilities to drive centrifuges. This is another strong indication that the adversaries were well prepared and disposed of specific information through reconnaissance.

Stuxnet would only infect Siemens' S7-300 PLCs. All other PLCs were left untouched. Targeted S7-300 PLCs would have to be configured with modules connected to variable-frequency drives of particular make and model (Iran and Finland). Variable-frequency drives control the speed of motors. These infected PLCs would be programmed with malicious code to monitor the speed used to drive the centrifuge motors, and only interfere with the operation if specific criteria are met. For example, the spinning frequency of the attached motors had to be between 807 Hz and 1210 Hz, all to avoid interfering with PLCs that are not used for the Iranian nuclear program.


Furthermore, Stuxnet did not interfere with systems that were already infected.

When all conditions were met, Stuxnet would periodically modify the frequency of the drivers to alter the speed of the centrifuge motors, while reporting the original frequency back to the monitoring systems. This is the first documented case of a rootkit on a PLC.



## BlackEnergy – Lights Out in Ukraine (1)

Context: Ukraine has been in a large conflict with Russia since 2014



On 23 December 2015, the power of 200,000 Ukrainian citizens was cut for periods ranging from 1 to 6 hours. The power outages were the result of a successful cyber attack on at least 3 Ukrainian power distribution companies.

The malware delivered during the attack was named BlackEnergy and is believed to originate from an APT group called "Sandworm"

As opposed to the sophistication of the tools used by The Equation Group, the malware used during the Ukraine attack was highly unsophisticated, though highly effective!

### BlackEnergy – Lights Out in Ukraine (1)

On December 23, 2015, an estimated 200,000 inhabitants of the Ukraine were left without electricity for periods varying between 1 and 6 hours. These power outages were the result of a successful digital attack on at least 3 Ukrainian power distribution companies. BlackEnergy is the name of the malware used in this attack.

Involved in a conflict with Russia since 2014, the digital attack on Ukraine is believed to have originated in Russia and security researchers have attributed such attacks to a Russian APT group with the name Sandstorm.

Although the Idaho National Laboratory demonstrated in 2007 that it was possible to physically destroy an electricity generator just using a program (the Aurora Generator Test), this attack on the Ukrainian power grid is believed to be the first successful digital attack on a power grid.

#### References:

[https://en.wikipedia.org/wiki/December\\_2015\\_Ukraine\\_power\\_grid\\_cyber\\_attack](https://en.wikipedia.org/wiki/December_2015_Ukraine_power_grid_cyber_attack)


<https://www.wired.com/2016/01/everything-we-know-about-ukraines-power-plant-hack/>



## BlackEnergy – Lights Out in Ukraine (2)



**Step 1** – Phishing emails toward employees to deliver the virus through malicious Office documents



**Step 2** – BlackEnergy is written to disk and is used to harvest VPN credentials and identify SCADA systems



**Step 3** – Access to SCADA systems is used to open circuit breakers at 230 substations of the power grid



**Step 4** – Infected workstations are wiped and DDoS attack against call centers is launched

### DID YOU KNOW THAT?

The BlackEnergy malware was not specifically targeted against SCADA systems, as it featured mainly standard Trojan-like behavior

For defenders, it would have been fairly easy to close the circuit breakers again upon cutting of the power!

This was, however, hindered by the adversaries by wiping infected workstations (used for management of the SCADA systems) and a DDoS attack against the victims' call centers

### BlackEnergy – Lights Out in Ukraine (2)

Using spear-phishing, malicious documents were delivered to key staff in the power distribution companies. Spear-phishing is a form of phishing (using fake emails) where the recipients are a small, carefully selected group to maximize the success rate of the phishing campaign. Through reconnaissance, the adversaries identified key people in the target companies and obtained their email addresses. The malicious documents were Microsoft office documents using Visual Basic for Applications to deliver the BlackEnergy payload.

BlackEnergy is a modular, 32-bit Windows malware family that is not particularly designed to attack SCADA systems. It has been used for various purposes, like stealing information, remote access and compromising home banking transactions. Through the remote access feature, the adversaries seized control over Windows workstations connected to SCADA systems.

Via those SCADA systems, the adversaries managed to open circuit breakers at 230 substations of the power grid, thereby cutting electricity to 200,000 people. Once power was lost and the power companies were alerted, it would have been simple to restore power by closing the circuit breakers again. However, the adversaries took additional steps to prevent this simple operation and thereby prolonging the duration of the power cuts. By wiping the workstations infected with BlackEnergy with the KillDisk program, the adversaries prevented power grid company staff to remotely close the circuit breakers. Staff had to be dispatched to the different substations to manually close the circuit breakers.

At the same time, a denial of service attack on the power grid company's call centers was executed. By overloading the exchanges with phone calls, adversaries prevented customers from calling in and reporting power cuts. Denying staff access to its control systems and information resulted in power cuts taking up to six hours.



## NotPetya



The NotPetya ransomware struck a variety of organizations in June 2017. Although it had all the symptoms of a traditional ransomware attack aimed at stealing money, it's generally agreed the main goal was sabotage (as the malware also overwrote the Master Boot Record).



Prior to the attack, the Ukraine software developer "Linkos Group" was compromised in a supply chain attack. The compromise was used to backdoor the MEDoc software, in which the NotPetya software was deployed. This was used as the initial infection vector of NotPetya.



NotPetya became famous very quickly because of its large impact (encryption of data + overwriting of the Master Boot Record) and fast spread method, where its combined exploitation of a Windows vulnerability (MS017-10) and typical lateral movement strategies used by APT's (Mimikatz-variant to dump credentials and PSEXEC + WMIC to connect to remote machines).

### NotPetya

Let's look at another example of an attack mainly aimed toward sabotage. The NotPetya ransomware struck a variety of organizations in June 2017. Although it had all symptoms of a traditional ransomware attack aimed at stealing money, it's generally agreed the main goal was sabotage (as the malware also overwrote the Master Boot Record).

There's a few interesting things to note with regards to the NotPetya ransomware:

- Prior to the attack, the Ukraine software developer "Linkos Group" was compromised in a supply chain attack. The compromise was used to backdoor the MEDoc software, in which the NotPetya software was deployed. This was used as the initial infection vector of NotPetya. To date, NotPetya remains one of the most devastating Supply Chain Attack examples (next to, for example, the CCleaner example);
- NotPetya became famous very quickly because of its large impact (encryption of data + overwriting of the Master Boot Record) and fast spread method, where it combined exploitation of a Windows vulnerability (MS017-10) and typical lateral movement strategies used by APT's (Mimikatz to dump credentials and PSEXEC + WMIC to connect to remote machines). The combination of both MS017-10 and Mimikatz allowed NotPetya to first infect an unpatched machine, but afterwards spread laterally to infect also patched machines (using stolen credentials).

#### Reference:

<https://www.wired.com/2016/01/everything-we-know-about-ukraines-power-plant-hack/>

<https://www.wired.com/story/inside-the-unnering-supply-chain-attack-that-corrupted-ccleaner/>



## NotPetya – Attack Diagram




**Step 1** – MeDoc software repository is compromised to deploy NotPetya in source code\*



**Step 2** – MeDoc software update infects “patient zero” in the internal network.

**MS17-010**

MS17-010: “EternalBlue” used to infect unpatched systems on the network.



Reuses stolen administrative credentials to infect patched systems on the network.

### On infected machines, NotPetya performs the following actions:

- Dump credentials from LSASS (tokens + passwords & hashes)
- Encrypt local files
- Overwrite Master Boot Record
- Reboot

## NotPetya – Attack Diagram

In the NotPetya ransomware attack, the following key attack steps were completed throughout the campaign:

1. MeDoc software repository is compromised to deploy NotPetya in source code. The initial intrusion at Linkos Group occurred early in 2017, the actual delivery of the NotPetya payloads was in June 2017.
2. Through the software update, a “patient zero” is infected in a corporate network. Once a system is infected, NotPetya performs the following operations:
  - a. It dumps credentials in LSASS memory from the local machine (it uses a Mimikatz variant for this)
  - b. It encrypts local files
  - c. Overwrites the Master Boot Record
  - d. Reboots the computer
3. The further lateral movement through the corporate environment is achieved using two main methods:
  - a. For unpatched systems, NotPetya uses the EternalBlue and EternalRomance vulnerabilities to spread (TCP ports 139 and 445)
  - b. For patched systems, NotPetya attempts to reuse credentials it stole in the previous attack (Using PSEXEC & WMIC to run its payload)


### References:

<https://www.cybereason.com/blog/blog-the-two-actor-theory-behind-the-notpetya-attack>

<https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/>




## Zooming in on Espionage



**Cyber Crime**



**Sabotage**



**Espionage**


### **Zooming in on Espionage**

Finally, let's turn our eyes to the interesting world of cyber espionage!



## Key Drivers for Espionage

TOP SECRET



Espionage is aimed at obtaining unauthorized access to sensitive data / information that could benefit the adversary

Adversaries include commercial competitors, but also hostile (and even friendly) nations

Although spying is an old tradecraft, spies often are the earliest adopters of new technology

### Key Drivers for Espionage

Espionage is the act of spying. Spies obtain secret information without permission and knowledge from their targets. Spying is done by many actors and at many levels.


Nation states have always had espionage agencies, like the CIA in the USA and KGB in the Soviet Union. These agencies collect information about their targets to further the cause of the nation. This is military information to be better prepared in case of an armed conflict, and political information to have an advantage in political negotiations with insider information. For example, if you know to what level your opponent is willing to make concessions during negotiations, you can use this to your advantage to get a better deal.

Companies and organizations also spy on each other, within and outside national borders. This industrial espionage is performed to have an advantage over the competition. Research and development is a very costly activity companies engage in, and it doesn't always yield expected results. But R&D is vital for a company's growth. A cheaper way is to obtain research from a competitor and develop new products before the competitors do. This is illegal in most countries, but it will not stop companies from doing this. Sometimes, the cost and risk of espionage outweigh the cost of research, and companies engage in industrial espionage and factor in the fines they will have to pay if they get caught.

Spying is a millennium-old practice; however, it is not behind in its use of technology. Spies are often early adopters of new technology to improve their practices. With the digitalization of technology came the need to have the capability to steal information from digital devices. This resulted in espionage actors adopting sophisticated technology to infiltrate these digital devices.



## Espionage – Looking at an Active Campaign – Turla



Turla, also known as Snake or Uroburos, is one of the most sophisticated ongoing cyber-espionage campaigns (on a similar technical level as Regin)

Targets of the campaign include government entities (e.g. Ministry of Foreign Affairs, Intelligence Agencies), embassies, military, research and education organizations and pharmaceutical companies

The name of the campaign refers to Ouroboros, which is the ancient symbol of a snake or dragon biting its own tail; it is widely accepted to be of Russian origin

### Espionage – Looking at an Active Campaign – Turla

Turla, also known as Snake or Uroburos is one of the most sophisticated ongoing cyber-espionage campaigns (on a similar technical level as, for example, Regin).

Turla is Windows malware (32-bit and 64-bit) that was first discovered in 2013 but has been targeting Western government and military organizations since at least 2008. Due to language and strings in the executables, encryption keys used and behavior, G Data attributes Turla to Russia. The malware seems to be related to malware Agent.BTZ that was used in 2008 during an attack on the United States of America. Turla checks for the presence of Agent.BTZ on a machine it tries to infect and remains inactive if found.

The Turla malware contains many references to snakes. Filenames containing the word snake and strings in the code like Ur0BuR()s. Ouroboros is an ancient symbol of a snake or dragon biting its own tail.

#### References:


<https://en.wikipedia.org/wiki/Ouroboros>

[https://public.gdatasoftware.com/Web/Content/INT/Blog/2014/02\\_2014/documents/GData\\_Uroburos\\_RedPaper\\_EN\\_v1.pdf](https://public.gdatasoftware.com/Web/Content/INT/Blog/2014/02_2014/documents/GData_Uroburos_RedPaper_EN_v1.pdf)

[https://en.wikipedia.org/wiki/Turla\\_\(malware\)](https://en.wikipedia.org/wiki/Turla_(malware))



## Espionage – Looking at an Active Campaign – Turla













The group has launched many different campaigns over the years and continuously updates & innovates its toolkit. Some of the more interesting feats they have achieved:

- Persistence using COM object hijacking (one of their trademarks)
- Using satellite connectivity for C&C (to hinder C&C identification / take-down)
- C&C using steganography in Instagram pictures
- C&C using a custom Outlook backdoor (identified in 2018)

### Turla victims by geography *(not exhaustive)*

Turla often relies on typical social engineering tricks for initial intrusion. This includes, for example, lure documents, watering hole attacks,... Their use of 0-days has been limited and mostly focused on client-side software (e.g. Flash)

 Belgium	 Poland
 China	 Romania
 France	 Russia
 Germany	 Switzerland
 Netherlands	 USA

### Espionage – Some Active Groups and Campaigns

The group has launched many different campaigns over the years and continuously updates and innovates its toolkit. Some of the more interesting feats they have achieved:

- Persistence using COM object hijacking (one of their trademarks).
- Using satellite connectivity for C&C (to hinder C&C identification / take-down).
- C&C using steganography in Instagram pictures.
- C&C using a custom Outlook backdoor (identified in 2018).

We will zoom in on some of these techniques in the next few slides.

While Turla can definitely be classified as an "APT", they often rely on typical social engineering tricks for initial intrusion. This includes, for example, lure documents, watering hole attacks,... Their use of zero-days has been limited and mostly focused on client-side software (e.g. Flash).

Turla victims have been identified in the following countries (not exhaustive):

- Belgium
- China
- France
- Germany
- The Netherlands
- Poland
- Romania
- Russia
- Switzerland
- United States of America



## Espionage – Looking at an Active Campaign – Turla: COM Object Hijacking

### COM

COM (Component Object Model) is described by Microsoft as “platform-independent, distributed, object-oriented system for creating binary software components that can interact.” The purpose of COM is to provide an interface to allow developers to control and manipulate objects of other applications. All COM objects are defined by a unique ID called CLSID.

COM hijacking is a “stealthy” persistence technique that allows adversaries to run payloads in the context of trusted processes.

Two commonly used techniques include “Phantom COM objects” abuse and “COM Search Order Hijacking”, where the adversary hijacks a COM object to run a payload of his choosing.

Hundreds of COM objects exist that can be used to run payloads in the context of core Windows processes such as explorer.exe, svchost.exe, chrome.exe,...

COM hijacking can be performed with regular user privileges and is much stealthier as opposed to, for example, code injection techniques.

### Espionage – Looking at an Active Campaign – Turla: COM Object Hijacking

COM (Component Object Model) is described by Microsoft as “platform-independent, distributed, object-oriented system for creating binary software components that can interact.” The purpose of COM is to provide an interface to allow developers to control and manipulate objects of other applications. All COM objects are defined by a unique ID called CLSID.

COM hijacking is a “stealthy” persistence technique that allows adversaries to run payloads in the context of trusted processes. This in itself is not a new technique / objective used by adversaries, but it was formerly implemented using, for example, code injection techniques. Many of today’s current security products, however, look for code injection techniques, rendering this strategy noisy and detectable.

Hundreds of COM objects exist that can be used to run payloads in the context of core Windows processes such as explorer.exe, svchost.exe, chrome.exe,... Research by Cyberbit in July 2018 exposes many of these as vulnerable to COM Object hijacking.

Two commonly used techniques include “Phantom COM objects” abuse and “COM Search Order Hijacking”, where the adversary hijacks a COM object to run a payload of his choosing:

- Phantom COM objects are instances of COM objects that exist in registry, but don’t have an “implementation file” on disk. This implementation file could thus be created by an adversary.
- Machine-wide COM objects are stored in the HKLM registry hive, while user-wide COM objects are stored in the HKCU registry hive. It’s important to note that user-wide COM objects take precedence over machine-wide objects. Adversaries could thus hijack existing machine-wide COM objects using user-wide COM objects.


#### Reference:

<https://www.cyberbit.com/blog/endpoint-security/com-hijacking-windows-overlooked-security-vulnerability/>





## Espionage – Looking at an Active Campaign – Turla: Instagram C&C



A variant of Turla communicated by reading and writing comments on Britney Spears' Instagram account.

### Example comment

"asmith2155:#2hot make loved to her, uupss #Hot #X"

- The malware would read comments posted for this picture and calculate a custom hash. If the custom hash is 183, then the malware knows the comment is from the C&C server and that it must decode it.
- The comment contains a hidden special UNICODE character: \200d. This is a non-printable character called "Zero Width Joiner". Like its name implies, it does not take up space, but it is present.
- This character, together with a couple of other characters like # and @, are used as a prefix to the encoded URL:

```
<200d>#2hot ma<200d>ke lovei<200d>d to <200d>her, <200d>uupss <200d>#Hot <200d>#X
```

This gives us the following string: 2kdhuHX, or URL [http://bit\(.\)ly/2kdhuHX](http://bit(.)ly/2kdhuHX)!

## Espionage – Looking at an Active Campaign – Turla: Instagram C&C

A variant of the Turla malware uses Instagram as a command & control channel. Not via Instagram accounts that would communicate directly with each other, but via Instagram accounts that would post comments on pictures posted by a very popular Instagram user: Britney Spears.

To send a message to the malware, the command & control channel would post a link (for example to pastebin) shortened with a URL shortening service like Bitly. This URL would not be posted directly to Britney Spears' picture as a comment, but it would be encoded with steganography in a seemingly innocuous comment.

An example of such a comment would be: #2hot make loved to her, uupss #Hot #X

The malware would read comments posted for this picture and calculate a custom hash. If the custom hash is 183, then the malware knows the comment is a command from the command & control server and that it must decode it.

It can't be seen in the text of the comment, but the comment contains a special UNICODE character: \200d. This is a non-printable character called "Zero Width Joiner". Like its name implies, it does not take up space, but it is present.

This character, together with a couple of other characters like # and @, are used as prefix to the encoded URL.

The malware extracts the URL from the message by taking the character that follows the prefix character.

Here is the message with the special UNICODE character made visible (<200d>):

```
<200d>#2hot ma<200d>ke lovei<200d>d to <200d>her, <200d>uupss <200d>#Hot <200d>#X
```


This gives us the following string: 2kdhuHX, or URL [http://bit\(.\)ly/2kdhuHX](http://bit(.)ly/2kdhuHX)

### Reference:

<https://www.welivesecurity.com/2017/06/06/turlas-watering-hole-campaign-updated-firefox-extension-abusing-instagram/>




## Espionage – Looking at an Active Campaign – Turla: Outlook Backdoor



In 2018, a new Turla backdoor was identified that purely relies on an Outlook backdoor to persist and set up command & control connectivity. Although it primarily targets Outlook, the backdoor can also infect "The Bat!", a mail client that is very popular in Eastern Europe.



The backdoor allows for both stealth command execution, download / upload of files. It thus serves for both Command & Control and as an exfiltration channel.



The backdoor only uses standard, valid, email communication, including said PDF documents. This is not uncommon in corporate environments. 😊



The backdoor uses steganography to hide commands and exfiltrated inside images in specifically crafted PDFs.



Persistence of the backdoor is established using their well-known COM object hijacking technique! Never change a winning team!

### Espionage – Looking at an Active Campaign – Turla: Outlook Backdoor

In 2018, a new Turla backdoor was identified that purely relies on an Outlook backdoor to persist and set up command & control connectivity. Although it primarily targets Outlook, the backdoor can also infect "The Bat!", a mail client that is very popular in Eastern Europe.

There's a few interesting items to note with regards to this backdoor:

- The backdoor allows for both stealth command execution, download / upload of files. It thus serves for both Command & Control and as an exfiltration channel.
- The backdoor uses steganography to hide commands and exfiltrated inside images in specifically crafted PDFs.
- The backdoor only uses standard, valid, email communication, including said PDF documents. This is not uncommon in corporate environments. 😊
- Persistence of the backdoor is established using their well-known COM object hijacking technique! Never change a winning team!

The Outlook backdoor observed in 2018 shows that, while Turla further continues to innovate, it reuses techniques that were previously effective (not reinvent the wheel, maximize ROI).

#### Reference:

<https://www.welivesecurity.com/wp-content/uploads/2018/08/Eset-Turla-Outlook-Backdoor.pdf>





## Espionage – Some Active Groups and Campaigns

The current cyber espionage landscape includes actors from all parts of the world (“everybody is doing it”); we listed some of the most known groups below:

Name	Other Names	Known Campaigns	Main Targets
APT-28	Sofacy, Fancy Bear	Grizzly Steppe, DNC Hack, Bundestag	US and European Governments
APT-29	Dukes, Cozy Bear	Grizzly Steppe	US and European Businesses
Turla	Snake	Satellite Turla, Epic Turla	US and European Governments and Businesses
Sandworm	TEMPNoble	Black Energy	Eastern European Utilities
APT-1	Comment Panda	ShadyRAT	English-speaking high-tech firms
APT-3	Gothic Panda	Clandestine Fox, Double Tap	Worldwide defense contractors
APT-27	Emissary Panda	Operation Iron Tiger, A Tale of 2 Targets	US Government and defense contractors
Charming Kitten	Parastoo	Stonedrill, Shamoon	Saudi & US Interests (focus on utilities & defense contractors)
Copy Kitten	Slayer Kitten	Matryoshka	Israeli Interests in the Middle East
Rocket Kitten	TEMPBeanie	Operation Woolen Goldfish	Saudi Arabian, Israeli and US Interests in the Middle East
Equation Group	Tilded Team	Stuxnet, Regin	Worldwide

### Espionage – Some Active Groups and Campaigns

The current cyber espionage landscape includes actors from all parts of the world. It's safe to assume that the vast majority of states are developing and using cyber espionage capabilities. The table in the slide lists a number of different groups, coupled with some of their best-known campaigns and the types of organizations they usually target.

It's important to note that attribution based on pure "technical facts" is often difficult. For example, "the source IP address is from Russia" or "the command & control server is hosted in China" are highly unreliable elements in attribution (as they can be easily forged). Some more interesting elements that are currently being used for attribution:

- Similarities in coding style (or even copy / paste work).
- Similarities in tools that are used (e.g. the use of Mimikatz during post-exploitation).
- Exploitation of the same vulnerabilities (e.g. zero-days that are used in different campaigns).
- Sophistication of the malware.
- Artifacts identified during analysis (e.g. PDB path, compilation times...).
- The target of the attack (correlated with the current geopolitical situation).

Still, attribution remains a difficult topic in today's world.

#### Reference:

[https://docs.google.com/spreadsheets/d/1H9\\_xaxQHpwaa4O\\_Son4Gx0YOIzlcBWMsdvePFX68EKU/edit#gid=1864660085](https://docs.google.com/spreadsheets/d/1H9_xaxQHpwaa4O_Son4Gx0YOIzlcBWMsdvePFX68EKU/edit#gid=1864660085)

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.I

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim

### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

## Introducing SYNCTECHLABS

SYNCTECHLABS is an organization that is focused on a wide variety of different services and products in critical industries, making them a possible target for persistent adversaries.

Your name is Alan Marshall and you were recently hired to:

- Do adversary emulation;
- In true "purple" fashion, help implement security controls!



During the course, you will work on the above two tasks as part of Alan's daily activities!

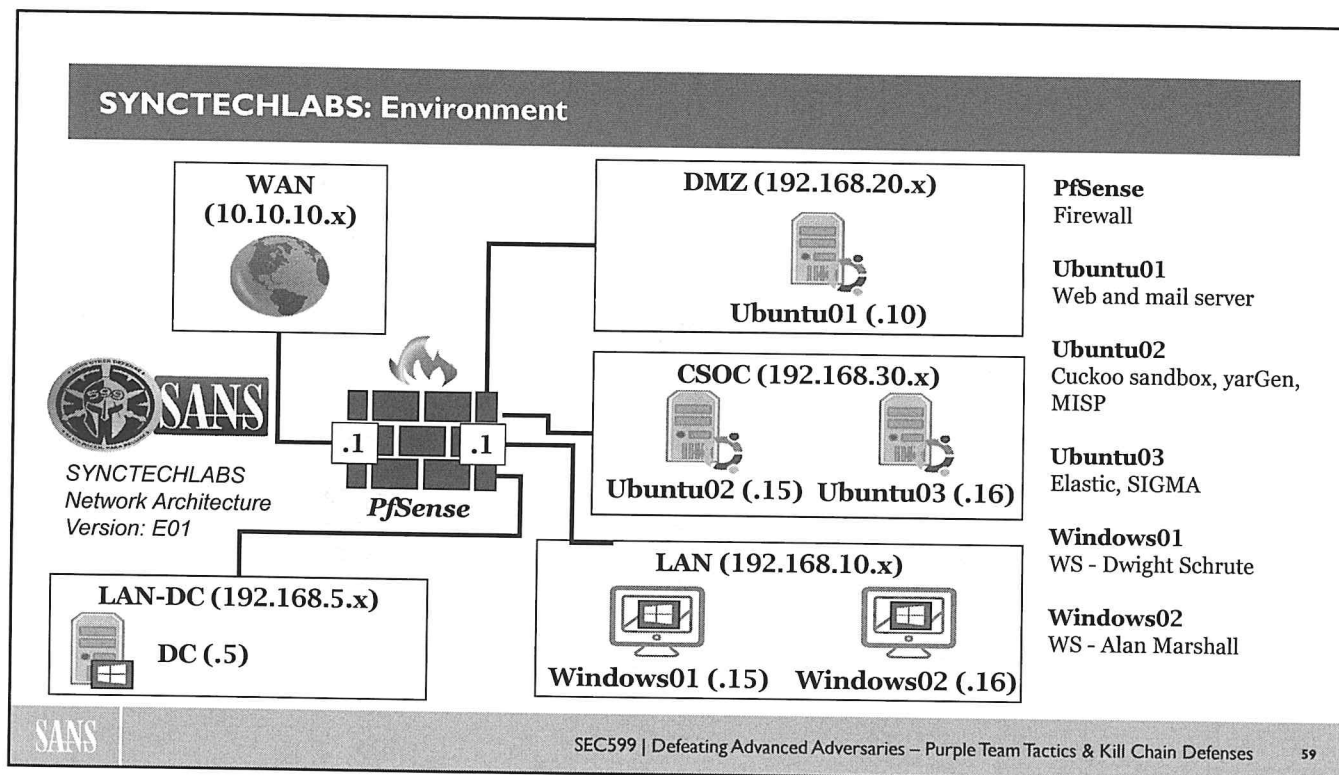
## Introducing SYNCTECHLABS

SYNCTECHLABS is an organization that is focused on a wide variety of different services and products in critical industries, making them a possible target for persistent adversaries.

Your name is Alan Marshall and you were recently hired to:

- Do adversary emulation according to a defined standard such as MITRE ATT&CK.
- In true "purple" fashion, they want you to also help improve things by suggesting and implementing security controls!

During SEC599, you will take up Alan's role and work on both emulation and implementation of security controls.



## SYNCTECHLABS: Environment

The above is what the full SYNCTECHLABS environment looks like. Here's a quick rundown:

### Network zones

- WAN: This is the "outbound" network (internet).
- DMZ: This is where internet-accessible systems are hosted.
- CSOC: This is where Alan Marshall hosts security monitoring and analysis tools.
- LAN: This is where the workstations reside.
- LAN-DC: This is a segmented network zone for domain controllers.

Note that the WAN network is simulated and does not include actual internet outbound connectivity. There is a WebNet interface available (which should not be restricted) that provides actual outbound internet connectivity.

### Systems

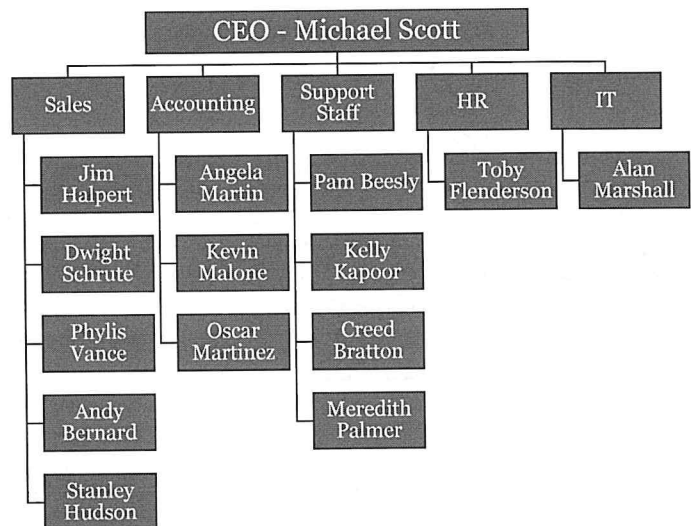
Throughout the environment, we have the following systems:

- The PfSense firewall is a PfSense firewall
- The Ubuntu01 server is a Ubuntu server providing mail & web services (IP 192.168.20.10)
- The Ubuntu02 server is a Ubuntu server hosting a Cuckoo sandbox, MISP and the yarGen tool (IP 192.168.30.15)
- The Ubuntu03 server is a Ubuntu server hosting an Elastic stack and SIGMA (IP 192.168.30.16)
- The Windows01 is a Windows 10 workstation primarily used by Dwight Schrute (IP 192.168.10.15)
- The Windows02 is a Windows 10 workstation primarily used by Alan Marshall (IP 192.168.10.16)
- The DC is a Windows Server 2016 which serves as the SYNCTECHLABS domain controller (IP 192.168.5.5)

## SYNCTECHLABS: Organigram

In order to get familiar with the different roles in SYNCTECHLABS, we've provide you with an organigram on the right.

Please keep this in mind when performing your work. It appears Dwight Schrute is rather IT savvy and he's tried to get around some of IT's security controls already...



### SYNCTECHLABS: Organigram



In order to get familiar with the different roles in SYNCTECHLABS, we've provided you with an organigram on the right. Please keep this in mind when performing your work. It appears Dwight Schrute is rather IT savvy, and he's tried to get around some of IT's security controls already...

Please have a look and familiarize yourself with key personnel!

## SYNCTECHLABS: Red Team Test

As a first step, you are asked to execute a **red team test** against SYNCTECHLABS. You will be provided with a Kali Linux machine that you need to use to try to gain access to the internal SYNCTECHLABS environment...

The "flag" defined in the red team engagement is to try to obtain "**Domain Administrator**" access to the environment!



### SYNCTECHLABS: Red Team Test

As a first task in your young career, you are asked to execute a red team test against SYNCTECHLABS. You will be provided with a Kali Linux machine that you need to use to try to gain access to the internal SYNCTECHLABS environment... The "flag" defined in the red team engagement is to try to obtain "Domain Administrator" access to the environment!

How will you get started? You cannot use your own "Alan Marshall" account, so you have to start from scratch as an external user.

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.I

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim


### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

**Exercise: One Click is All It Takes...**



Please refer to the workbook for further instructions on the exercise!

This page intentionally left blank.



# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599. I

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim


### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

## Adversary Emulation



Adversary emulation is an activity where security experts emulate how an adversary operates. The ultimate goal, of course, is to improve how resilient the organization is versus these adversary techniques. Both red and purple teaming can be considered as adversary emulation.

### TTP

Adversary activities are described using TTPs (Tactics, Techniques & Procedures). These are not as concrete as, for example, IOCs, but they describe how the adversary operates at a higher level. Adversary emulation should be based on TTPs. As such, a traditional vulnerability scan or internal penetration test that is not based on TTPs should not be considered adversary emulation.

### ATT&CK

Adversary emulation should be performed using a structured approach, which can be based on a kill chain or attack flow. MITRE ATT&CK is a good example of such a standard approach.

### Adversary Emulation

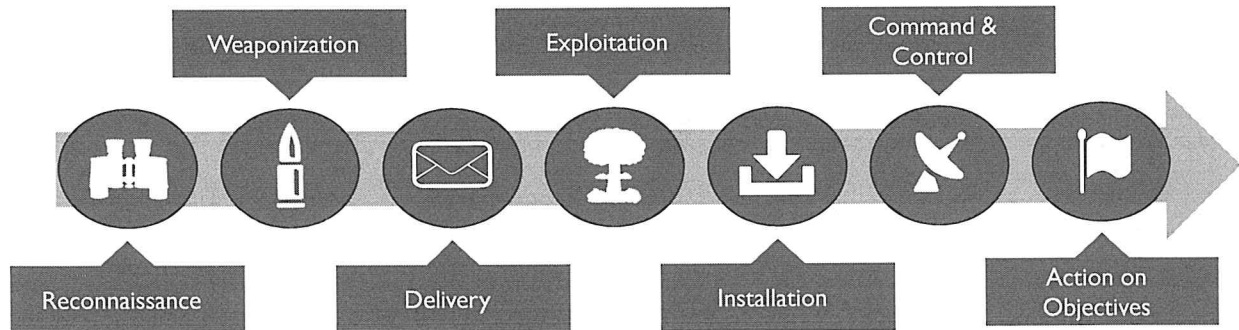
Adversary emulation is an activity where security experts emulate how an adversary operates. The ultimate goal, of course, is to improve how resilient the organization is versus these adversary techniques. Both red and purple teaming can be considered as adversary emulation.

One of the primary properties of adversary emulation is the use of TTPs. Adversary activities are typically described using TTPs (Tactics, Techniques & Procedures). TTPs are used by both red / purple teams (when emulating attacks) and by blue teams (when analyzing actual attacks that are taking place). These are not as concrete as, for example, IOCs, but they describe how the adversary operates at a higher level. Adversary emulation should be based on TTPs. As such, a traditional vulnerability scan or internal penetration test that is not based on TTPs should not be considered adversary emulation.

Adversary emulation should be performed using a structured approach, which can be based on a kill chain or attack flow. MITRE ATT&CK is a good example of such a standard approach.

## The Cyber Kill Chain®

One of the first examples of a structured description of attacks was the Cyber Kill Chain® by Lockheed Martin:



### The Cyber Kill Chain®

Different groups and organizations have worked on documenting adversaries' methods in a digital kill chain. Lockheed Martin developed the trademarked "Cyber Kill Chain®", which has risen in popularity to become one of the most used frameworks to describe cyber attacks. An alternative, slightly adopted variant is Dell SecureWorks' "Cyber Kill Chain". Both chains have more steps than the military kill chain.

Lockheed Martin: Reconnaissance, Weaponization, Delivery, Exploitation, Installation, Command & Control, Actions On Objectives.


Dell SecureWorks: Target Defined, Recon, Development, Weaponization, Delivery, Exploitation, Installation, Command & Control, Actions on Objectives, Objective Met

### References:

<https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html>

<https://www.secureworks.com/resources/wp-breaking-the-kill-chain>

## Limitations of the Cyber Kill Chain



It assumes the adversary leverages a certain "payload" that is delivered from the outside.

Key components of a typical attack (e.g. lateral movement) do not have the required focus.

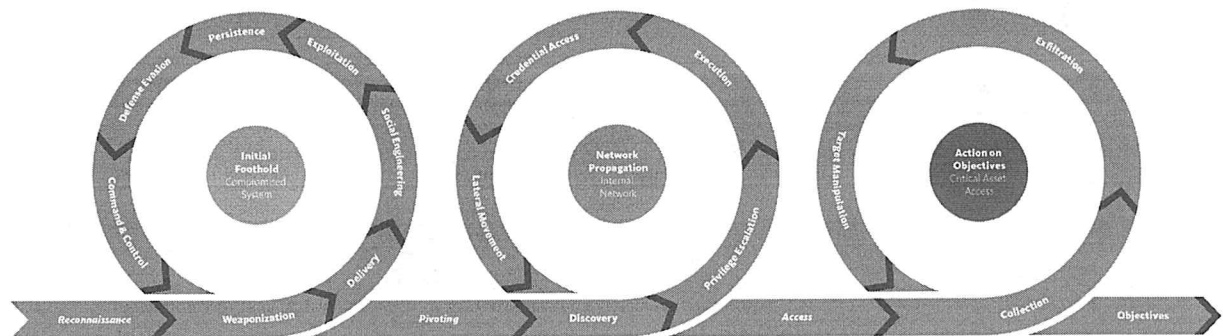
### Limitations of the Cyber Kill Chain

While the Cyber Kill Chain is not a bad model, it was initially developed / proposed in 2011 and could be further adapted to accommodate modern adversary tactics. Some of the most prominent limitations of the traditional Cyber Kill Chain include:

- It assumes the adversary leverages a certain "payload" that is delivered from the outside.
- Key components of a typical modern attack (e.g. lateral movement) do not have the required focus. They could be considered to be part of "Action on Objectives", but that would strongly "inflate" this phase.

Throughout the years, several adaptations have been proposed, in order to modernize / update the Cyber Kill Chain. The concept of a kill chain remains a fundamentally good way of approaching cyber security!

## Alternatives to the Cyber Kill Chain: Unified Kill Chain



**The Unified Kill Chain is a good answer to some of the Cyber Kill Chain limitations!**

### Alternatives to the Cyber Kill Chain – Unified Kill Chain

The Unified Kill Chain responds to the most important limitations / criticisms of the traditional Cyber Kill Chain. It uses three main phases in an attack:

- Initial Foothold
- Network Propagation
- Action on Objectives

Using this structure, different threat perspectives (e.g. an insider threat who is already on the network), can be modelled as well. You might recognize the majority of traditional Cyber Kill Chain steps in the "Initial Foothold" phase!

#### Reference:

<https://www.csacademy.nl/images/scripties/2018/Paul-Pols---The-Unified-Kill-Chain.pdf>

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.I

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim

### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

## What Is "Red Team" & "Blue Team"?



### What Is "Red Team" & "Blue Team"?

Traditionally, information security efforts have been split in a red team and blue team. The red team is mostly focused on offensive activities, where they attempt to test / bypass security controls implemented by an organization. On the other hand, the blue team focuses on implementing security controls, monitoring their effectiveness and responding in case of incidents. However, both share a common goal: Improving the defenses of the organization!

The list below is by no means exhaustive, but includes a number of typical red team activities:

- Vulnerability assessments, where a report is delivered that provides insights.
- Penetration tests, where red teams attempt to abuse identified vulnerabilities.
- Social engineering, where the overall security awareness of the organization's staff is tested.

Likewise, the list below is by no means exhaustive, but includes a number of typical blue team activities:

- Implementing security controls in order to prevent adversaries (and the red team) from obtaining unauthorized access to the organization's environment;
- Perform security monitoring and threat hunting to detect unauthorized access and to improve the security controls that have been implemented;
- Incident response, where the blue team responds to detected incidents.

In most organizations, at least a part of these blue and red team capabilities are delivered by third parties (e.g. contractors, specialized firms ...)

## So, Why a Purple Team?

Although they share a common goal, blue and red teams are often not well-aligned, which leads to organizations not leveraging the full value of their team's expertise. They typically report to a different hierarchy, which leads to different objectives:

RED	BLUE
Report with many vulnerabilities = Well done!	No alerts mean that preventive controls are working!
Success is measured by # of failed controls	Large volume of alerts means detection controls are working (though may need fine-tuning)
No big incentive to help blue team, as blue team failure = red team success!	No big incentive to help red team, as red team failure = blue team success!

While we don't necessarily think a purple team is a "third" team, think of it as a concept aimed at bringing the red and blue teams together. Red and blue teams should be encouraged to work as a joint team, share insights and create a strong feedback loop.

## So, Why a Purple Team?

Although they share a common goal, blue and red teams are often not well-aligned, which leads to organizations not leveraging the full value of their team's expertise. They typically report to a different hierarchy, which leads to different objectives... Even if they "get along", they might not have an intuitive nature to work together... Robert Wood and William Bengston delivered an interesting presentation at RSA 2016 ("The Rise of the Purple Team"), where they highlight this problem.

Consider the following:

- For the red team, an assessment report with many vulnerabilities means they have done a good job. Indeed, the success of the red team activities is measured by the number of failed (or bypassed) security controls.
- For the blue team, a lack of alerts means that preventive controls are working, while a large volume of alerts indicates that detection controls are working (although they may need fine-tuning).

Given the above, there is no big incentive for either the red team or the blue team to help "the other side"... While we don't necessarily think a purple team is a "third" team, think of it as a concept aimed at bringing the red and blue teams together. Red and blue teams should be encouraged to work as a joint team, share insights and create a strong feedback loop.



## The Purple Team Feedback Loop

What is happening now?:



But how about some of the following actions?:

- Red Team sharing TTPs of new actors with Blue Team + thinking together how they can be defended against (prevention + detection)
- Red Team helping with vulnerability management process (not just assessment), thereby prioritizing what vulnerabilities are most critical
- Blue Team testing out Red Team techniques themselves, so they can continuously improve on prevent and detection capability
- Blue Team sharing monitoring tactics & playbooks in place with Red Team, so Red Team can improve tactics for next exercise

### The Purple Team Feedback Loop

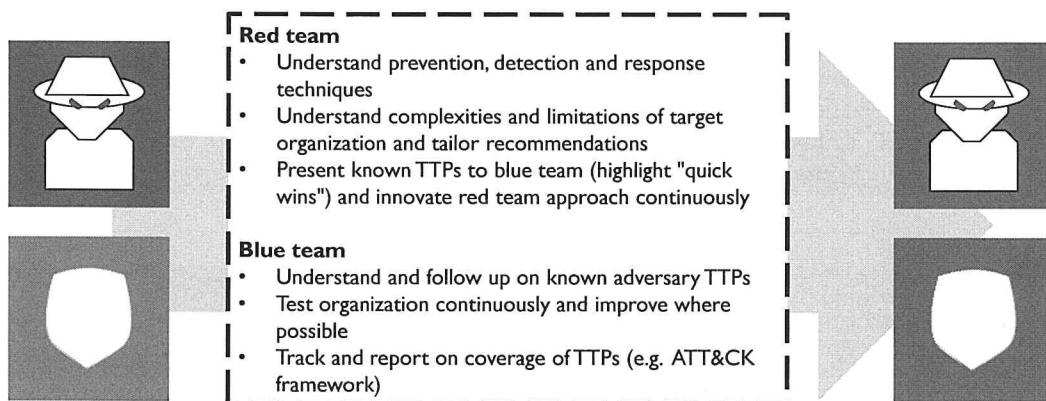
We would like to further expand on the Purple Team feedback loop. It's vital to stress that a feedback loop is NOT just the red team delivering a penetration testing report full of vulnerabilities to the blue team. This is what we observe in many organizations, but it misses out on a variety of other improvement possibilities:

- How about the red team sharing new TTPs (Tactics, Techniques & Procedures) with the blue team, so they can prepare for new red team exercises?
- How about the red team helping with the overall vulnerability management process and thus adapting its reporting style to the blue team vulnerability management process? (creating tickets or bug reports as opposed to writing a 120-page PDF).
- How about the blue team sharing existing monitoring in place with the Red Team, to help them better prepare their next red team engagement?
- How about the blue team sharing the existing incident response playbooks with the red team?

This list is again not meant to be exhaustive, but it provides an interesting insight in additional knowledge-sharing opportunities that could be leveraged as part of a strong feedback loop. Overall, we should "reward" the red team and blue team on the knowledge increase and progress that is made by both teams! Offense must inform defense and defense must inform offense!

## How to Approach This?

Let's make blue more "red" and make "red" more blue:



### How to Approach This?

So how can we approach this? Without making any drastic changes, let's make the red team a bit bluer and make the blue team a bit more red. There's quite some activities that could be picked up by the red and blue teams.

For the red team, consider the following:

- Understand prevention, detection and response techniques.
- Understand complexities and limitations of target organization and tailor recommendations.
- Present known TTPs to blue team (highlight "quick wins") and innovate red team approach continuously.


For the blue team, consider the following:

- Understand and follow up on known adversary TTPs.
- Test organization continuously and improve where possible.
- Track and report on coverage of TTPs (e.g. ATT&CK framework).


We will follow this approach throughout the course!

## So, How Do We Practically Do This? What about Our Yearly Red Team?

Does this mean "Purple" is better than "Red"? The answer is not that simple. ☺ Depending on your objectives, either could offer value. Here's an idea for a setup:



Organize a yearly **red team** to **assess** the actual state of security in the organization. Offer feedback only after the exercise ends, as the exercise is typically meant to be stealth (realistic adversary emulation)...  
**VALUE: Periodic assessment of organization resilience**



Perform continuous **purple teaming** to **improve** the state of security in the organization. Blue team members simulate focused attack techniques as part of their operations to immediately test effectiveness of detection and prevention controls.  
**VALUE: Continuous improvement of organization resilience**

## So, How Do We Practically Do This? What about Our Yearly Red Team?

Does this mean "Purple" is better than "Red"? The answer is not that simple! Depending on your objectives, either could offer value. Here's a proposed approach:


- Organize a yearly red team to assess the actual state of security in the organization. Offer feedback only after the exercise ends, as the exercise is typically meant to be stealth (realistic adversary emulation)... This could be compared to "traditional" penetration testing / red teaming. The value of such an exercise is that you can periodically assess the organization's resilience against cyber threats.
- In parallel to the above red team testing, perform continuous purple teaming to improve the state of security in the organization. Blue (or red) team members could simulate focused attack techniques as part of their operations to immediately test effectiveness of detection and prevention controls. During such exercises, the ultimate goal is not to assess the resilience of the organization, it is to improve!

We will illustrate these components throughout the course!

## SEC599: A True Purple Team Course

SEC599 was built from the ground up by seasoned cyber security experts with vast experience in both blue team and red team operations:

We illustrate what the adversaries are **ACTUALLY** doing with real-world examples and lab exercises



We provide **EFFECTIVE** strategies to combat adversary tactics throughout all parts of the Kill Chain!

## SEC599: Defeating Advanced Adversaries – Purple Team Tactics & Kill Chain Defenses

SANS

SEC599 | Defeating Advanced Adversaries – Purple Team Tactics & Kill Chain Defenses

75

### SEC599: A True Purple Team Course

SEC599 was authored with these problems in mind... How do we bridge the "gap" between the red and blue teams? In order to accomplish our goal, we assembled a team of experienced instructors with a key focus on red team operations (Stephen Sims and Erik Van Buggenhout) and added a number of our most well-known SANS ISC (Internet Storm Center) handlers (including malware analyst Didier Stevens). Together, they provide a unique mix of offensive and defensive skills bundled in one course!

So how do we approach cyber security concepts throughout the course?

- In order to implement effective security controls, we are convinced it is vital to first understand how adversaries operate. We will thus first explain offensive security techniques, explaining how organizations are currently being compromised. These will be hands-on, and students will use the latest techniques to compromise a target network.
- Once completed, we will deliver an in-depth explanation of what defensive controls can help mitigate the attack technique used above. We will also deploy the mitigating control in our infrastructure, after which we test its successful implementation!

## Purple Team: An Example Using Mimikatz and CredentialGuard

An example: As part of our "Lateral Movement" section, we will demonstrate how adversaries are stealing credentials from Windows systems using the infamous Mimikatz hacking tool.

1

During the first phase of the lab, you will dump credentials from the LSASS process memory, thereby experiencing the attack (and its impact) firsthand!

2

During the second phase of the lab, we will provide an in-depth explanation of how Windows 10 CredentialGuard works, after which it will be implemented.

3

During the third and final part of the lab, we will re-attempt dumping credentials from the LSASS process memory, thereby confirming effectiveness of our control!

### Purple Team: An Example Using Mimikatz and CredentialGuard

You may find that this does sound slightly theoretical. So, how does this work in practice? An ideal example to explain this approach is how we handle the lateral movement section of the course. One of the most popular techniques used by adversaries is to leverage the infamous Mimikatz hacking tool to dump hashes from the LSASS process.

How do we approach this in SEC599?

1. During the first phase of the lab, you will dump credentials from the LSASS process memory, thereby experiencing the attack (and its impact) firsthand!
2. During the second phase of the lab, we will provide an in-depth explanation of how Windows 10 CredentialGuard works, after which it will be implemented.
3. During the third and final part of the lab, we will reattempt dumping credentials from the LSASS process memory, thereby confirming effectiveness of our control!

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.I

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim

### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

## Introducing MITRE ATT&CK

ATT&CK Matrix for Enterprise

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Exfiltration	Command and Control
Drive-by Compromise	AppleScript	birth_profile and bashrc	Access Token Manipulation	Access Token Manipulation	Account Manipulation	Account Discovery	AppleScript	Audio Capture	Automated Exfiltration	Commonly Used Port
Exploit Public-Facing Application	CMSTP	Accessibility Features	Accessibility Features	BITS Jobs	Bash History	Application Window Discovery	Application Deployment Software	Automated Collection	Data Compressed	Communication Through Removable Media
Hardware Address	Command Line Interface	Account Manipulation	AppCert DLLs	Binary Patching	Brute Force	Browsers Bookmark Discovery	Distributed Component Object Model	Clipboard Data	Data Encrypted	Connectors Proxy
Replication Through Removable Media	Compiled HTML File	AppCert DLLs	AppCert DLLs	Bypass User Account Control	Credential Dumping	File and Directory Discovery	Exploitation of Remote Services	Data Staged	Data Transfer Size Limits	Custom Command and Control Protocol
Spearghishing Attachment	Control Panel Items	AppCert DLLs	Application Shimming	CMSTP	Credentials in Files	Network Service Scanning	Login Scripts	Data from Information Repositories	Exfiltration Over Alternative Protocol	Custom Cryptographic Protocol
Spearghishing Link	Dynamic Data Exchange	Application Shimming	Bypass User Account Control	Clear Command History	Credentials in Registry	Network Share Discovery	Pass the Hash	Data from Local Systems	Exfiltration Over Command and Control Channel	Data Encoding
Spearghishing via Service	Execution through API	Authentication Package	DLL Search Order Hijacking	Code Signing	Exploitation for Credential Access	Network Shifting	Pass the Ticket	Data from Network Shared Drives	Exfiltration Over Other Network Mediums	Data Obfuscation
Supply Chain Compromise	Extension through Module Load	BITS Jobs	Cydia Hijacking	Compiled HTML File	Forced Authentication	Password Policy Discovery	Remote Desktop Protocol	Data from Removable Media	Exfiltration Over Physical Medium	Domain Fronting

MITRE has developed the ATT&CK Matrix as a central repository for adversary TTP's. It is used by red teams and blue teams alike. It is rapidly gaining traction as a de facto standard!

### Introducing MITRE ATT&CK

MITRE has developed the ATT&CK Matrix as a central repository for adversary TTP's. It is used by red teams and blue teams alike. It is rapidly gaining traction as a de facto standard! According to its official website:

*"MITRE ATT&CK™ is a globally-accessible knowledge base of adversary tactics and techniques based on real-world observations. The ATT&CK knowledge base is used as a foundation for the development of specific threat models and methodologies in the private sector, in government, and in the cybersecurity product and service community.*

*With the creation of ATT&CK, MITRE is fulfilling its mission to solve problems for a safer world — by bringing communities together to develop more effective cybersecurity. ATT&CK is open and available to any person or organization for use at no charge."*

In the screenshot above, the "tactics" are described as column headers with a gray background. The "techniques" are all the entries below it in the matrix. Every technique has a specific ID, which can be easily referenced. In order to facilitate integration, MITRE has provided for example an API to get MITRE's ATT&CK information in STIX format.

#### Reference:

<https://attack.mitre.org/>

## Analyzing a MITRE ATT&CK Entry

The screenshot shows the MITRE ATT&CK website page for 'Component Object Model Hijacking'. The page title is 'Component Object Model Hijacking' and the ID is 'T1122'. The tactic is 'Defense Evasion, Persistence'. The platform is 'Windows'. The permissions required are 'User'. The data sources are 'Windows Registry, DLL monitoring, Loaded DLLs'. The defense bypassed is 'Autonomous Analysis'. The contributors are 'ENDGAME' and the version is '1.0'. The page also includes a sidebar with categories like 'Initial Access', 'Execution', and 'Persistence'.

As an example, let's have a look at one of Turla's favorite techniques: COM object hijacking. In MITRE's ATT&CK framework, this technique is known as T1122, and it's part of the "Defense Evasion" and "Persistence" tactics for Windows.

For every one of these techniques, MITRE includes a dedicated entry with, amongst others:

- Technique information
- Known adversaries that use it
- Detection opportunities
- Prevention opportunities

### Analyzing a MITRE ATT&CK entry

As an example technique, let's have a look at one of Turla's favorite techniques: COM object hijacking. In MITRE's ATT&CK framework, this technique is known as T1122, and is part of the "Defense Evasion" and "Persistence" tactics for Windows. As you can see in the screenshot on the slide, MITRE ATT&CK has a dedicated entry, which includes a lot of information on the technique:

- An overall description of the technique.
- The platforms for which it's relevant.
- The type of access required by an adversary to use the technique.
- Detection opportunities
- Prevention opportunities



## Leveraging ATT&CK: The Navigator

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Exfiltration	Command And Control
10 items	32 items	58 items	28 items	63 items	19 items	20 items	17 items	13 items	9 items	23 items
Drive-by Compromise	AppleScript	Jobs_profile and bashrc	Access Token Manipulation	Access Token Manipulation	Account Manipulation	Account Discovery	AppleScript	Audio Capture	Automated Exfiltration	Commonly Used Port
Exploit Public-facing Application	CMSTP	Accessibility Features	Account Manipulation	Binary Padding	Bash History	Application Window Discovery	Application Deployment Software	Automated Collection	Data Compressed	Communication Through Removable Media
Hardware Additions	Command-Line Interface	Account Manipulation	Accessibility Features	BITS Jobs	Brute Force	Application	Clipboard Data	Data Encrypted	Connection Privacy	Custom Command and Control Protocol
Replication Through Removable Media	Compiled HTML File	AppCert DLLs	AppCert DLLs	Bypass User Account Control	Credential Dumping	Browser Bookmark Discovery	Distributed Component Object Model	Data from Information Repositories	Data Transfer Size Limits	Custom Cryptographic Protocol
Spearpishing Attachment	Control Panel Items	Applet DLLs	Applet DLLs	Clear Command History	Credentials in Files	File and Directory Discovery	Exploitation of Remote Services	Data from Local System	Exfiltration Over Alternative Protocol	Data Encoding
Spearpishing Link	Dynamic Data Exchange	Application Shimmming	Application Shimmming	CMSTP	Credentials in Registry	Network Service Scanning	Logon Scripts	Data from Network Shared Drive	Exfiltration Over Command and Control Channel	Data Obfuscation
Spearpishing via Service	Execution through Module Load	Authentication Package	Bypass User Account Control	Code Signing	Exploitation for Credential Access	Network Share Discovery	Pass the Hash	Data from Removable		
		BITS Jobs	DLL Search Order Hijacking	Compiled HTML File	Forced Authentication					
		Bootkit		Component Firmware						

As part of many tools, MITRE has developed the ATT&CK Navigator. It's a web application that represents the ATT&CK techniques in a dynamic fashion. It can be used to select specific techniques based on a threat group (e.g. select all APT-28 techniques), after which modifications and annotations can be made. It can be an excellent central resource for an organization to obtain a "quick look" at how well they are doing versus the different techniques. It's open-source and can thus be self-hosted!

### Leveraging ATT&CK: The Navigator

As part of many tools, MITRE has developed the ATT&CK Navigator. It's a web application that represents the ATT&CK techniques in a dynamic fashion. It can be used to select specific techniques based on a threat group (e.g. select all APT-28 techniques), after which modifications and annotations can be made. It can be an excellent central resource for an organization to obtain a "quick look" at how well they are doing versus the different techniques. It's open-source and can thus be self-hosted!

A demo instance of the ATT&CK navigator can be found here:

<https://mitre.github.io/attack-navigator/enterprise/>

The source code can be found here:

<https://github.com/mitre/attack-navigator>

## Adversary Emulation Tools – The Usual Suspects: Metasploit, Empire ...



Metasploit is an exploitation framework used by virtually all penetration testers. It has both a free community edition and a commercial edition available. Its main focus is on "standardization" of exploit development and usage.



Empire is primarily a post-exploitation tool. It has both Windows support (using a pure PowerShell2.0 agent) and Linux / OS X support (using a pure Python 2.6/2.7 agent). It is the result of the merger of PowerShell Empire and Python EmPyre!

### Adversary Emulation Tools – The Usual Suspects: Metasploit, Empire ...

So how can we emulate these adversary techniques? Traditionally, some of the most commonly used frameworks during penetration testing and red teaming have been Metasploit and Empire:

- Metasploit is an exploitation framework used by virtually all penetration testers. It has both a free community edition and a commercial edition available. Its main focus is on "standardization" of exploit development and usage.
- Empire is primarily a post-exploitation tool. It has both Windows support (using a pure PowerShell2.0 agent) and Linux / OS X support (using a pure Python 2.6/2.7 agent). It is the result of the merger of PowerShell Empire and Python EmPyre!

These are highly effective tools, yet they are mostly focused on the "red team" side of adversary emulation and less on a purple approach. Let's have a look at some other tools, which are more geared toward purple team approaches.

## Keeping It Simple: APT Simulator

```
Select Administrator: Command Prompt - APTSimulator.bat

APT Simulator
Florian Roth, Nextron Systems, v0.6.0
Select the test-set that you want to run:
[0] RUN EVERY TEST
[1] Collection
[2] Command and Control
[3] Credential Access
[4] Defense Evasion
[5] Discovery
[6] Execution
[7] Lateral Movement
[8] Persistence
[9] Privilege Escalation
[A] Apply AV Exclusions in Registry
[S] Settings
[E] Exit
Your selection (then press ENTER): A_
```

APTSimulator is a Windows-based tool that makes a system look like it was victim of a targeted attack. (Key focus is thus on the endpoint)

It supports a wide variety of the ATT&CK tactics, as described in the screenshot to the left.

As it's primarily built on BAT files, it is highly customizable / extendable.

### Keeping It Simple: APT Simulator

APT Simulator was written by Florian Roth (Nextron Systems) and was designed to keep things simple. It is not a full-blown adversary emulation tool. Some of its use cases include:

- Test your endpoint detection and response tools.
- Test your security monitoring capabilities.
- Test your security response effectiveness.
- Prepare an environment for DFIR labs or trainings.

From its official documentation:

*"APT Simulator is a Windows Batch script that uses a set of tools and output files to make a system look as if it was compromised. In contrast to other adversary simulation tools, APT Simulator is designed to make the application as simple as possible. You don't need to run a web server, database or any agents on set of virtual machines. Just download the prepared archive, extract and run the contained Batch file as Administrator. Running APT Simulator takes less than a minute of your time."*

#### Reference:

<https://github.com/NextronSystems/APTSimulator>

## Keeping It Simple: Network Flight Simulator (flightsim)

```
bash-3.2# ./flightsim-darwin-amd64 run dga
AlphaSOC Network Flight Simulator™ v1.0.4 (https://github.com/alphasoc/flightsim)
The IP address of the network interface is 172.20.0.27
The current time is 15-Nov-18 07:26:39

Time      Module  Description
-----
07:26:39 dga      Starting
07:26:39 dga      Generating list of DGA domains
07:26:39 dga      Resolving teovhmk.com
07:26:40 dga      Resolving teovhmk.biz
07:26:41 dga      Resolving teovhmk.info
07:26:42 dga      Resolving yjdsnbi.com
07:26:43 dga      Resolving yjdsnbi.biz
07:26:44 dga      Resolving yjdsnbi.info
07:26:45 dga      Resolving ijatwtr.com
07:26:46 dga      Resolving ijatwtr.biz
07:26:47 dga      Resolving ijatwtr.info
07:26:48 dga      Resolving dpmqgk.com
07:26:49 dga      Resolving dpmqgk.biz
07:26:50 dga      Resolving dpmqgk.info
07:26:51 dga      Resolving fgexvbf.com
07:26:52 dga      Resolving fgexvbf.biz
07:26:53 dga      Resolving fgexvbf.info
07:26:54 dga      Resolving puaklce.com
07:26:55 dga      Resolving puaklce.biz
07:26:56 dga      Resolving puaklce.info
07:26:57 dga      Resolving tkaiymp.com
07:26:58 dga      Resolving tkaiymp.biz
07:26:59 dga      Resolving tkaiymp.info
07:27:00 dga      Resolving wkppnes.com
07:27:01 dga      Resolving wkppnes.biz
07:27:02 dga      Resolving wkppnes.info
07:27:03 dga      Resolving lhqallt.com
07:27:04 dga      Resolving lhqallt.biz
07:27:05 dga      Resolving lhqallt.info
07:27:06 dga      Resolving sywfedm.com
07:27:07 dga      Resolving sywfedm.biz
07:27:08 dga      Resolving sywfedm.info
07:27:09 dga      Finished

All done! Check your SIEM for alerts using the timestamps and details above.
bash-3.2#
```

Where APTSimulator is mostly focused on endpoint controls, Network Flight Simulator was developed by alphasoc to focus on network-level detection. It supports a variety of suspicious network connectivity.

```
bash-3.2# ./flightsim-darwin-amd64
AlphaSOC Network Flight Simulator™ v1.0.4 (https://github.com/alphasoc/flightsim)

flightsim is an application which generates malicious network traffic for security teams to evaluate security controls (e.g. firewalls) and ensure that monitoring tools are able to detect malicious traffic.

Usage:
  flightsim [command]

Available Commands:
  help      Help about any command
  run       Run all simulators (default) or a particular test
  version   Print version and exit

Flags:
  -h, --help  help for flightsim

Use "flightsim [command] --help" for more information about a command.
bash-3.2#
```

### Keeping It Simple: Network Flight Simulator (flightsim)

Where APTSimulator is mostly focused on endpoint controls, Network Flight Simulator was developed by alphasoc to focus on network-level detection. It supports the simulation of a variety of suspicious network connectivity:

- DNS tunneling
- Domain Generation Algorithms (DGA)
- Known bad domains
- Tor
- ...

According to its official documentation:

*"flightsim is a lightweight utility used to generate malicious network traffic and help security teams to evaluate security controls and network visibility. The tool performs tests to simulate DNS tunneling, DGA traffic, requests to known active C2 destinations, and other suspicious traffic patterns."*

#### Reference:

<https://github.com/alphasoc/flightsim>

## Keeping It Simple: Atomic Red Team

### Atomic Test #1 - System Service Discovery

Identify system services


Supported Platforms: Windows

#### Inputs

Name	Description	Type	Default Value
service_name	Name of service to start stop, query	string	svchost.exe

Run it with `command_prompt` !

```
tasklist.exe
sc query
sc query state= all
sc start ${servicename}
sc stop ${servicename}
wmic service where (displayname like "${servicename}") get name
```



Red Canary has developed "Atomic Red Team", which is a series of "simple" tests that can be used to emulate the behavior of adversaries in the environment.

All tests are fully linked to MITRE ATT&CK!

### Keeping it simple – Atomic Red Team

Atomic Red Team (by Red Canary) is the first of a few tools we'll introduce that are linked to the MITRE ATT&CK framework. The goal of Atomic Red Team is to "allow every security team to test their controls by executing simple "atomic tests" that exercise the same techniques used by adversaries (all mapped to Mitre's ATT&CK)." This is very much in line with a purple team approach: Empower the blue team to test prevention and detection of various adversary techniques!

From its official GitHub page:

*"Three key beliefs made up the Atomic Red Team charter:*

- Teams need to be able to test everything from specific technical controls to outcomes. Our security teams do not want to operate with a "hopes and prayers" attitude toward detection. We need to know what our controls and program can detect, and what it cannot. We don't have to detect every adversary, but we do believe in knowing our blind spots.*
- We should be able to run a test in less than five minutes. Most security tests and automation tools take a tremendous amount of time to install, configure, and execute. We coined the term "atomic tests" because we felt there was a simple way to decompose tests so most could be run in a few minutes. **The best test is the one you actually run.***
- We need to keep learning how adversaries are operating. Most security teams don't have the benefit of seeing a wide variety of adversary types and techniques crossing their desk every day. Even we at Red Canary only come across a fraction of the possible techniques being used, which makes the community working together essential to making us all better."*

#### References:

<https://github.com/redcanaryco/atomic-red-team>  
<https://atomicredteam.io/>

## Going All the Way: MITRE CALDERA

hostname	Status
win7x01	active
win7x04	active
win7x02	active
win7x03	active
win2012xdc	active

CALDERA is a tool built by MITRE, with the express purpose of doing adversary emulation. It requires a bit of setup (as a server needs to be installed) and it will actively "attack" target systems by deploying custom backdoors. CALDERA's attack steps are fully linked to the ATT&CK framework techniques!

### Going All the Way: MITRE CALDERA

While the previously mentioned tools were rather "simple" to set up and configure, CALDERA is a bit different! It requires a bit of setup (as a server needs to be installed) and it will actively "attack" target systems by deploying custom backdoors. CALDERA's attack steps are fully linked to the ATT&CK framework techniques!

From its official documentation:

*"CALDERA is an automated adversary emulation system that performs post-compromise adversarial behavior within Windows Enterprise networks. It generates plans during operation using a planning system and a pre-configured adversary model based on the Adversarial Tactics, Techniques & Common Knowledge (ATT&CK™) project. These features allow CALDERA to dynamically operate over a set of systems using variable behavior, which better represents how human adversaries perform operations than systems that follow prescribed sequences of actions.*

*CALDERA is useful for defenders who want to generate real data that represents how an adversary would typically behave within their networks. Since CALDERA's knowledge about a network is gathered during its operation and is used to drive its use of techniques to reach a goal, defenders can get a glimpse into how the intrinsic security dependencies of their network allow an adversary to be successful. CALDERA is useful for identifying new data sources, creating and refining behavioral-based intrusion detection analytics, testing defenses and security configurations, and generating experience for training."*

#### Reference:

<https://github.com/mitre/caldera>

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.1

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim



### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

## Prevent or Detect?



“The prevent-only model has failed” is a solid statement. However, it does not say “The prevent model has failed.” We still need to deploy “preventive” controls in order to stop low hanging fruit and reduce the overall noise. In addition to these preventive controls, we should accept some attacks will succeed and thus complement preventive controls with detective controls.

### Prevent or Detect?

Work under the assumption that persistent adversaries with enough resources will succeed in the initial intrusion. When adversaries have months to prepare and execute the initial intrusion phase of the attack, it is safe to assume they will succeed, regardless of how good your defenses are. Complex systems like your network and computers always have vulnerabilities (through bugs, configuration errors or even a lack of security awareness of your staff) and a persistent attacker will have the time to discover and exploit them.

Prevention is important. However, working under the assumption that you will not be able to prevent all attacks, detection is even more important. This can be the detection of the attack itself or the actions of the adversaries after the initial intrusion (like lateral movement). Even though the attacker could have successfully completed the first steps of the kill chain, we might be able to prevent a more damaging phase, such as sensitive data exfiltration, from happening.

Additionally, controls aimed at detection are more forgiving than prevention mechanisms. A strict prevention control causing a lot of false positives will have a negative impact on business operations, as legitimate actions will be blocked. In case of a strict detection control, there might still be a large number of false positives reported, but the operational impact will be limited. As a result, it is a good idea to test a new control in detection mode first and replicate it in prevention mode once it has proven its worth and the false positive rate is reduced.



## Where Do We Start...?

### CIS

The CIS controls are a set of 20 top controls that can prevent adversaries against (advanced) adversaries. They include both technical and organizational controls and are a great "start" for organizations to start with.

As 20 controls can sometimes still be considered "daunting", the Australian Signals Directorate (ASD) further brought this down to 4 key controls. These controls would, according to their study, mitigate 85% of the intrusion techniques the ACSC (Australian Cyber Security Centre) respond to. They include:

1. Application Whitelisting

3. Patch Applications

2. Patch Operating Systems

4. Restrict Administrative Privileges

### Where Do We Start...?

The CIS controls are a set of 20 top controls that can prevent adversaries against (advanced) adversaries. They include both technical and organizational controls and are a great "start" for organizations to start with. We will walk through many of these controls throughout the course.

As 20 controls can sometimes still be considered "daunting", the Australian Signals Directorate (ASD) further brought this down to 4 key controls. These controls would, according to their study, mitigate 85% of the intrusion techniques the ACSC (Australian Cyber Security Centre) respond to. They include:

1. Application whitelisting
2. Operating system patching
3. Application patching
4. Restriction of administrative privileges


#### References:

<https://www.cisecurity.org/controls/>

<https://www.cyber.gov.au/publications/strategies-to-mitigate-cyber-security-incidents-mitigation-details>

## Security Architecture for Detection and Prevention

In order to facilitate detection and prevention, following are some key architectural principles to take into account:



Segment internal network to complicate lateral movement. Private VLANs are a great concept!



Ensure logs are generated, centralized and analyzed (both for internal and perimeter devices)



Apply similar privileges that allow you to leverage cloud architectures! Centralize "outbound" traffic!



Increased hardening and monitoring on employee workstations (usually initial point of intrusion)

**Accept "loss" of the perimeter and architect accordingly!**

### Security Architecture for Detection and Prevention

In order to facilitate detection and prevention, following are some key architectural principles to consider:


- Not a single large enterprise has a flat network architecture, where all systems are connected to each other without any limitations. Complex networks need to be structured in a way to make them efficient, manageable, and secure. There are several network architectures, and most of them include a perimeter segregating the enterprise systems from the systems on the internet. The "gates" separating these different network segments should be configured as control points similar to your environment's entry and exit points. They can be seen as the entry/exit points for your different network segments, where certain segments might have stricter security requirements than others.
- To allow detection and monitoring, all devices located in the network, such as routers, firewalls, servers, workstations,... should generate log information. Some of these devices will be serving as control points in your network, which makes logging especially important on these devices. As we have seen in the previous slides, having adequate logs available will facilitate investigations in case of a breach.
- The cloud doesn't have to lead to decreased security. When implemented well, cloud-based services can be leveraged to obtain better security. Apply similar privileges that allow you to leverage cloud architectures (e.g. a cloud-based proxy like Zscaler)! Centralize "outbound" traffic, even if it means you centralize it in the cloud!
- Workstations should be a key focal point for hardening, monitoring and response. Note that workstations are quite often an initial point of intrusion, thus they should be given the required care. Consider an Endpoint Detection & Response (EDR) tool to facilitate fast detection and response of incidents.

In the next slides, we will illustrate a couple of key architecture principles.

## Architecture Principle 1 – Web Proxy Setup

In a mature environment, the web proxy is one of few (or the only) means of allowed outbound connectivity for end-systems:

- Security decisions (e.g. blacklist-specific URLs or categories) can be made in one central location
- Facilitated security controls and monitoring



### Architecture Principle 1 – Web Proxy Setup

Implementing web filters is "easy" in an enterprise if web proxies are used. The proxy server can operate as a central location where all security decisions are made and where all the filtering can be done. If no proxy is present, then a less efficient form of filtering can be implemented in DNS and firewalls.

In case the internal network is split up into segments, for example, grouping workstations separately from servers, the proxy can be used as an in-line system providing functionality for both segments. As mentioned before, it can serve as a web proxy for your users, implementing URL filtering, blacklisting, or even SSL interception.


For servers, it can serve as a reverse proxy, potentially as a reverse proxy cache or to perform pre-authorization or load balancing and providing a centralized point for monitoring.

## Architecture Principle 1 – Web Proxy Setup Leveraging Cloud

An increasing number of employees work from outside the boundaries of our corporate perimeter (e.g. when they are travelling or working from home)

As we want to **protect employees from wherever they work**, the traditional approach required employees to set up a VPN and access the internet through the corporate infrastructure.

**Cloud-based proxies** provide an interesting solution: The web proxies are no longer hosted in our own perimeter but are hosted on the internet (in the cloud) themselves. Commercial solutions include Symantec, Zscaler, Cisco Umbrella...



## Architecture Principle 1 – Web Proxy Setup Leveraging Cloud

Implementing centralized egress points in the internet perimeter sounds like an excellent idea (and it is!). There is, however, one slight problem: An increasing number of employees work from outside the boundaries of our corporate perimeter. This can include people who work from home or people who are travelling for work.

As an organization, we want to protect employees from wherever they work! The traditional approach required employees to set up a VPN and access the internet through the corporate infrastructure. But what if they haven't enabled their VPN (because they are not connecting to corporate services)? We still want to make sure no malware reaches their systems!

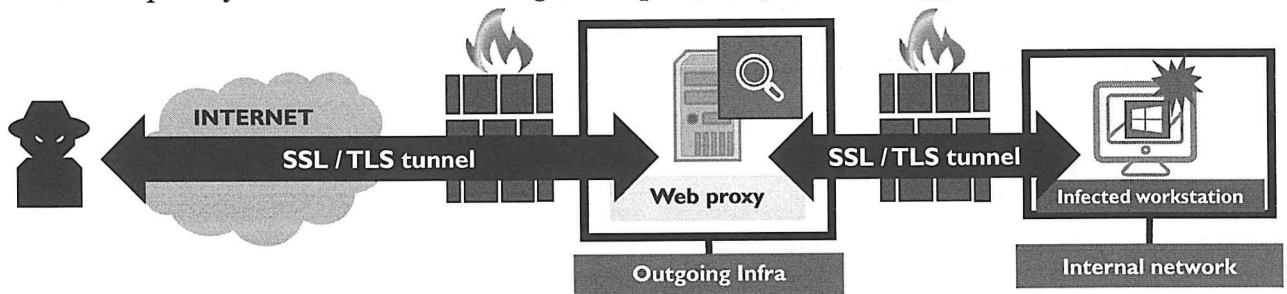
Cloud-based proxies provide an interesting solution: The web proxies are no longer hosted in our own perimeter but are hosted on the internet (in the cloud) themselves. This allows organizations to enforce the use of such proxies wherever employees are located (and regardless whether or not they have their VPN enabled!). There are many different solutions out there, including Symantec, Zscaler, Cisco Umbrella...

## Architecture Principle I – Web Proxy with SSL / TLS Interception



During SSL interception, one of the systems between the internet and the internal systems will "break" the connection. Instead of allowing the internal system to create a tunnel to the internet system, the interceptor will set up two tunnels: A tunnel between the internet and itself and a tunnel between the internal workstation and itself.

SSL interception is typically performed by web proxies (which will, of course, require the root CA to be accepted by the clients that are being intercepted) or by dedicated appliances.



SANS

SEC599 | Defeating Advanced Adversaries – Purple Team Tactics & Kill Chain Defenses

92

### Architecture Principle 1 – Web Proxy with SSL / TLS Interception

Because encrypted HTTP traffic becomes the norm on the internet, network devices that perform deep inspection of network traffic are becoming "blind." A solution to this problem is TLS interception.

A proxy that supports TLS interception works as follows:

- When the client starts a TLS connection, the TLS intercepting proxy does not forward the TLS connection to the web server, but it establishes a TLS connection with the client (TLS 1).
- The client accepts this TLS connection from the proxy because the proxy uses a certificate (with matching private key) that is trusted by the client. This can be done by using a certificate generated by a corporate PKI that has its root certificate installed on all corporate machines.
- Next, the proxy will establish its own TLS connection with the web server (TLS 2).
- One set of keys is used to encrypt TLS connection 1, and another set is used to encrypt TLS connection 2.
- Since the proxy is now an endpoint for both TLS connections, it can decrypt the traffic from both sides.
- The proxy decrypts the traffic from one channel and sends it encrypted in the other channel.

This allows the proxy to inspect the "encrypted" traffic, and it can also create a network tap for use by other network devices.

## Architecture Principle 2 – Securing DNS Traffic

### DNS

Stopping or analyzing traffic toward known malicious hosts at the web proxy level is one thing. We could, however, implement controls at an earlier stage: Before any domain name is contacted, it will have to be resolved. We can thus implement security controls at the DNS level that can help detect or prevent the resolution of known malicious domains.

- DNS logs are essential to have when doing investigations. These can be logged from your corporate DNS server, or could also be parsed from network streams (PassiveDNS, Zeek, Suricata,...). Don't forget to also capture NXDOMAIN (which is often not the default setting)!
- Some online services like Quad9 (by the Global Cyber Alliance) offer a free DNS service that automatically checks domain names against threat intelligence feeds, thereby attempting to prevent known malicious domains from resolving. Since Q1 2018, Cloudflare offers a similar service as "1.1.1.1", with a key focus on privacy and security!
- Another service is offered by OpenDNS (now part of Cisco), which has free plans for personal users and offers premium services for enterprise customers.


### Architecture Principle 2 – Securing DNS Traffic

Stopping or analyzing traffic toward known malicious hosts at the web proxy level is one thing. We could, however, implement controls at an earlier stage: Before any domain name is contacted, it will have to be resolved. We can thus implement security controls at DNS level that can help detect or prevent the resolution of known malicious domains.

- DNS logs are essential to have when doing investigations. These can be logged from your corporate DNS server, or could also be parsed from network streams (PassiveDNS, Zeek, Suricata,...). Don't forget to also capture NXDOMAIN (which is often not the default setting)!
- Some online services like Quad9 (by the Global Cyber Alliance) offer a free DNS service that automatically checks domain names against threat intelligence feeds, thereby attempting to prevent known malicious domains from resolving. Since Q1 2018, Cloudflare offers a similar service as "1.1.1.1", with a key focus on privacy and security!
- Another service is offered by OpenDNS (now part of Cisco), which has free plans for personal users and offers premium services for enterprise customers.

These services can be easily consumed by adapting the upstream DNS server in your organizations! This obviously has a privacy impact, as these services will now see all of your outgoing DNS resolutions... Of course, this can also be something you configure yourself based on intelligence gathered from external sources and investigations performed internally.

## Architecture Principle 3 – Network Monitoring



Network monitoring is a term that has been around for ages. There are, however, a number of fundamental questions that are to be answered before a network monitoring capability can be set up:

- Should we monitor only NetFlow or full packet capture?
- At what location(s) of the network should we perform monitoring? How about cloud applications? Remote workers? ...
- How do we handle SSL/TLS?
- How long should we retain NetFlow / full packet captures?
- How relevant is network monitoring in the age of SSL/TLS?

### Architecture Principle 3 – Network Monitoring

Network monitoring is a term that has been around for ages. There are, however, a number of fundamental questions that are to be answered before a network monitoring capability can be set up... When advanced adversaries attack our infrastructure, they will be forced to use our network. Even if they have physical access to one machine, they will need to use the network to move laterally. That is why in our network architecture design, we apply segmentation so that we have chokepoints where we can monitor network traffic.

So, let's agree we need to monitor network traffic. But, how? Should we do full packet capture, or should we just focus on NetFlow? Furthermore, how do we handle SSL / TLS encrypted protocols? Because we are bound to miss some attacks, it is a good idea to keep some kind of "history" of network traffic, so that this history can be consulted post-facto.

But even if we would monitor all network traffic, we would still miss detecting some attacks. For example, these attacks "hide in plain sight" by using existing network traffic and injecting their own messages. A very difficult technique to detect, for example, is steganography: Steganography conceals messages inside other messages. The typical example of steganography is an existing picture where some bits are altered to encode the hidden message. These altered bits do not change the overall aspect of the picture. The same principle can be applied to network traffic by adding whitespace characters to the headers of an HTTP request.

Even if we discuss this topic while addressing Command & Control, it's important to note that network monitoring can be highly useful to detect several other stages of the APT Attack Cycle as well!

## Architecture Principle 3 – Network Monitoring – FPC



Full Packet Capture is the concept of sniffing and storing all network traffic using, for example, a network tap. As opposed to NetFlow, Full Packet Capture will capture the full packet payload / content.

- Large enterprise networks generate A LOT of traffic, so storage could be an issue!
- The advantages of FPC over plain NetFlow could be limited if there's a lot of encrypted protocols...
- If feasible in your organization, it can be a good "catch-all" or "flight recorder"
- Some critical government institutions will collect, and store full packet captures for a large period of time (e.g. 1 to 2 years). This will allow them to retroactively hunt their logs / packet captures for attack campaigns that are identified in the future!

### Architecture Principle 3 – Network Monitoring – FPC

Where NetFlow (and similar technologies) will capture metadata, full packet capture will capture and store the complete content of network packets. This means that for a TCP connection, we not only have all the metadata like NetFlow would provide, but also the content of the TCP transmission itself: All the data that was transmitted and received by the computer.

Full packet capture will generate a lot of data, gigabytes to terabytes of data per day, depending on the network. This is why it is not possible to perform full packet capture on all computers. When full packet capture is implemented, it is typically done by tapping the network connections at key points in the corporate network infrastructure. These key points are, for example, the chokepoints where traffic flows between network segments. A good starting point for full packet capture is at the perimeter: Capture all traffic between the corporate network and the internet.

Not only does full packet capture help analysts to investigate attacks post-facto, but it also allows for the detection of attacks post-facto. This is done by processing all captured data of a given period (for example the last month) with an IDS like Snort or Suricata. The IDS needs to be provided with the detection rules and IOCs of the latest attack methods so that it can detect attacks that happened in the past and were not detected.

A good example is the Heartbleed vulnerability: This is a vulnerability in OpenSSL that allowed an adversary to obtain private data from SSL/TLS webservers. This private data could contain the encryption keys, for example. After Heartbleed was discovered and disclosed, several IDS rules were created to detect Heartbleed attacks. By searching all captured network traffic for evidence of a Heartbleed attack, a corporation could determine (post-facto) if it had been attacked and if data was leaked.

Although full packet capture produces a huge amount of data, it must be taken into account that large capacity magnetic hard disks have become cheap.



### Architecture Principle 3 – Network Monitoring – IDS & NSM

Many commercial and open-source tools exist that can meet a variety of our needs:

- Full packet capture
- IDS alerting
- Protocol anomaly detection
- Network Security Monitoring



Some commercial examples include the typical firewall vendors: Palo Alto, Juniper, Cisco, Fortinet, Sophos, ... Open-source alternatives for network monitoring include Snort, Suricata, Zeek (formerly Bro), ... They can typically be configured in IDS (passive) or IPS mode (inline).

For both commercial and open-source tools, the architecture and configuration of the device will be essential. If they are not correctly positioned in the network, they will not fulfill their full potential.

### Architecture Principle 3 – Network Monitoring – IDS & NSM

Many commercial and open-source tools exist that that can meet a variety of our needs... We are typically looking for the following types of functions:

- Full packet capture
- IDS alerting
- Protocol anomaly detection
- Network Security Monitoring...

Some commercial examples include the typical firewall vendors: Palo Alto, Juniper, Cisco, Fortinet, Sophos, ... Open-source alternatives for network monitoring include Snort, Suricata, Zeek (formerly Bro), ... They can typically be configured in IDS (passive) or IPS mode (inline).

In inline mode, the IDS/IPS has at least two network interfaces: One network interface is connected to the corporate network and the second network interface is connected to the internet. This way, the network traffic flows through the IDS/IPS device and is available for inspection, while the IPS can block traffic for which alerts are generated. The disadvantage of inline mode is availability; when the IDS/IPS device goes down, network traffic no longer flows through the device and the corporate network is severed from the internet.

Passive mode does not have this inconvenience. In passive mode, the IDS/IPS device will receive network traffic on one interface via a network tap (just like full packet capture). If the device goes down, the network flow is not impacted. The IPS can no longer block the traffic in passive mode, but there is a solution for this problem: The IPS device can send an instruction to the network device with the network tap to block the traffic detected by the IPS.

For both commercial and open-source tools, the architecture and configuration of the device will be essential. If they are not correctly positioned in the network, they will not fulfill their full potential.

### Architecture Principle 3 – Network Monitoring – Suricata

Suricata is a free, open-source, network threat detection / prevention engine (can be in IDS or IPS mode). It was first introduced in 2010.

- Maintained by the Open Information Security Foundation (OISF)
- One of the few free IDS engines that has multi-threading
- Has a large community and user base
- Main use cases include IDS, IPS, NSM and offline PCAP processing
- Standard input / output formats (e.g. YAML, JSON...) allow easy integration
- Biggest use is the “IDS” engine, but it can also generate application logs by parsing network traffic. Although not as powerful as Zeek (formerly Bro) in this regard, it still supports HTTP, SSL, TLS, SMB, DCERPC, SMTP, FTP, SSH, DNS, Modbus, ENIP/CIP, DNP3, NFS, NTP, DHCP, TFTP, KRB5, IKEv2 (as of version 4.1)



### Architecture Principle 3 – Network Monitoring – Suricata

Suricata is an open-source network detection and prevention engine. Suricata is the Latin name for a meerkat, a small mammal standing on its hind legs always looking out for signs of danger. Suricata is free and runs on different operating systems like Linux and Windows.

It is developed and maintained by the Open Information Security Foundation, which first introduced it in July 2010. It is an open-source project with many developers and contributors.

Suricata can be used for many purposes:

- As Intrusion Detection System
- As Intrusion Prevention System
- As Network Security Monitor
- An engine to process network capture files (PCAP files) offline.

It can be easily integrated with other products because of its openness. It uses well-known, standard formats for input like YAML and JSON. YAML (YAML Ain't Markup Language) is a structured data format used for configuration files, which can easily be read by humans, too.

JSON (JavaScript Object Notation) is another structured data format used to serialize objects. Objects are data structures with properties. JSON contains the names of objects, properties, and their values.

Suricata has built-in functionality to support detection / blocking of several steps in the APT Attack Cycle. Some examples include:

- Detecting the activity of exploit kits in HTTP traffic;
- Detecting Command & Control communications (beaconing...);
- Detecting malicious payloads being delivered via email (SMTP);
- ...

### Architecture Principle 3 – Network Monitoring – Suricata IDS Rulesets

- Emerging Threats is an organization focused on threat intelligence (acquired by Proofpoint in 2015)
- They distribute a community IDS ruleset ("ET") and a paid IDS ruleset ("ET Pro")
- ET rulesets are easily deployed in both Snort and Suricata
- An alternative to ET are the Talos / Snort (part of Cisco) rulesets:  
=> Both free community and paid professional versions are available



TALOS

```
alert udp $HOME_NET any -> any 53 (msg:"ET TROJAN Sofacy DNS Lookup hotfix-update.com"; content:"|01
00 00 01 00 00 00 00 00|"; depth:10; offset:2; content:"|0d|hotfix-update|03|com|00|";
fast_pattern; nocase; distance:0; reference:url,fireeye.com/resources/pdfs/apt28.pdf;
classtype:trojan-activity; sid:2019570; rev:2;)
```

### Architecture Principle 3 – Network Monitoring – Suricata IDS Rulesets

One source of free and paid-for rules is Emerging Threats. Emerging Threats is an organization that focuses on threat intelligence (they were acquired by Proofpoint in 2015). They observe threat actors and analyze their methods of operation in detail. With this threat intelligence, they are able to create rules that can be used in IDS/IPS devices to detect attacks and other actions by threat actors.

Emerging Threats distributes a free set of rules (the community IDS ruleset called "Emerging Threats") and a paid ruleset (called "Emerging Threats Pro"). These rulesets are updated daily: Existing rules are modified or retired, and new rules are added. This is why the commercial offering by Emerging Threats works with a subscription model: Paying for a subscription gives the right to the subscriber to download new rules (Emerging Threat Pro) daily. These rulesets are available for Snort and Suricata (the syntax of these rules can vary slightly), and also for different versions of Snort and Suricata, as more recent versions of these engines contain features that rules cannot use in older versions.

Another source of rules are the Snort rulesets. Snort was acquired by Cisco and is not part of Talos. Talos both offers free community and paid professional rulesets. Professional rulesets are often based on threat intelligence that is not public knowledge. For more information on the free Emerging Threats ruleset, consult <https://doc.emergingthreats.net/>. This Wiki gives an overview of all rules in the free ruleset.

The above example is a rule that inspects DNS queries. This can be deduced from the type of traffic it analyses (UDP) and the direction the traffic takes (from the internal network, on any port, to any IP address on port 53—that is the DNS port). It will look for byte patterns that indicate a DNS query (01 00 00 01 00 00 00 00 00) and then it will check if the domain is hotfix-update dot com. If this is the case, the rule will produce an alert.

This alert will include the name of the rule ("ET TROJAN Sofacy DNS Lookup hotfix-update com") and the ID of the rule (2019570), together with information regarding the traffic like source and destination addresses and ports. The classtype of the rule is trojan-activity. Emerging Threats has many categories, like Trojan, DoS, SCADA, Worm, Exploit, SMTP ...

## Architecture Principle 3 – Network Monitoring – Suricata Output

```
root@suricata:~# nano /etc/suricata/suricata.yaml
GNU nano 2.5.3 File: /etc/suricata/suricata.yaml
- eve-log:
  enabled: yes
  filetype: regular
  filename: eve.json
  #prefix: "@cee: "
  # the following are
  #ident: "suricata"
  #facility: locals
  #level: Info ##
  #redts:
  # server: 127.0.0.1
  # port: 6379
  # nodes: list ##
  # key: suricata ##
  # Redis pipelining
  # 'batch-size' eve
  # connection at the cost of some memory. There is no flushing implemented
  # so this setting as to be reserved to high traffic suricata.
  # pipelining:
  #   enabled: yes ## set enable to yes to enable query pipelining
  #   batch-size: 10 ## number of entry to keep in buffer
  types:
  - alert:
    # payload: yes # enable dumping payload in Base64
    # payload-buffer-size: 4kb # max size of payload buffer to output in eve-log
    # payload-printable: yes # enable dumping payload in printable (lossy) format
    # packet: yes # enable dumping of packet (without stream segments)
    http: yes # enable dumping of http fields
    tls: yes # enable dumping of tls fields
    ssh: yes # enable dumping of ssh fields
    smtp: yes # enable dumping of smtp fields

1 {"timestamp": "2017-02-16T17:18:36.504419+0100", "flow_id": 1564303607707941, "pcap_cn
t": 63, "event_type": "alert", "src_ip": "96.126.103.196", "src_port": 80, "dest_ip": "10.1
0.10.16", "dest_port": 63923, "proto": "006", "alert": {"action": "allowed", "gid": 1, "sign
ature_id": 2803027, "rev": 4, "signature": "ETPRO WEB_CLIENT Microsoft Excel
malformed Selection (type 0x1D) BIFF record", "category": "Attempted User
Privilege
Gain", "severity": 1}, "http": {"hostname": "example.com", "url": "\/files\/temp\/merged.
xls", "http_user_agent": "Mozilla\/5.0 (Windows NT 6.3; WOW64) AppleWebKit\/537.36
(RHTML, like Gecko) Chrome\/56.0.2924.87
Safari\/537.36", "http_content_type": "application\/vnd.ms-excel", "http_method": "GET
", "protocol": "HTTP\/1.1", "status": 200, "length": 54918}}
```

EVE JSON

Example of eve.json with an alert

### Suricata output types

Suricata supports a variety of output types & logs. One of the most useful ones is the EVE JSON log format, which provides a logging format for all of its different event types (alert, http, tls, smtp ...)

## Architecture Principle 3 – Network Monitoring – Suricata Output

When analyzing traffic, Suricata can be configured to produce different types of output. This is done through the configuration file `suricata.yaml`. A very useful output log is the EVE JSON log format (file `eve.json`). This will contain alerts and other information in JSON and can then be integrated into other applications.

In the example above, the `eve-log` is configured to produce an `eve.json` file with alerts.

One alert is displayed in the `eve.json` file. We know it is an alert because of key/value entry `"event_type": "alert"`.

### Architecture Principle 3 – Network Monitoring – Zeek NSM



The Zeek NSM (Network Security Monitor) is a network analysis framework that goes beyond being an IDS; it allows for more general network traffic analysis as well. Zeek has various protocol analyzers available, which allows you to perform analysis at the application layer.

Zeek is used in a variety of open-source network monitoring solutions, including, for example, "Security Onion." It supports (among others) parsing of the following network protocols:

HTTP	DHCP	DNS	FTP	ICMP	IMAP
ModBus	MySQL	NetBIOS	NTLM	NTP	POP3
RADIUS	RDP	SIP	SMB	SMTP	SNMP
SSH	SSL	Teredo	...		

**We will use Zeek to obtain network visibility in different future labs!**


#### Architecture Principle 3 – Network Monitoring – Zeek NSM

The Zeek NSM (Network Security Monitor) is a network analysis framework that goes beyond being an IDS; it allows for more general network traffic analysis as well. Zeek has various protocol analyzers available, which allow you to perform analysis at the application layer.

Zeek is used in a variety of open-source network monitoring solutions, including, for example, "Security Onion." It was built with a huge selection of protocol parsers, including highly interesting ones such as RDP, SMB... , which we don't find in an IDS like Suricata (although Suricata is working on adding SMB parsing).

In order to build additional functionality, anyone can write "recipes" for Zeek, which allows for complex logic and functionality to be built. RITA is an example of a project that consumes Zeek logs. The "Data Exfiltration Framework" is an addition on top of Zeek that looks for signs of data exfiltration (we will discuss data exfiltration on Day 5).

## Architecture Principle 4 – Network Segmentation



Overall, network segmentation is a key “defense-in-depth” principle! Many organizations struggle to implement and maintain this correctly. Still, this is an important part of your security controls. Not having a segmented network complicates further preventive and detective security controls!

Follow a structured process:  
Design -> Implement -> Maintain -> Monitor

Segment with key attack strategies in mind:  
How can you break lateral movement strategies?

Start simple: Guest, Test, Development,... and further segment from there

Implementation can be done through a variety of means, don't forget some quick wins!

### Architecture Principle 4 – Network Segmentation

Overall, network segmentation is a key "defense-in-depth" principle! Many organizations struggle to implement and maintain this correctly. Still, this is an important part of your security controls. Not having a segmented network complicates further preventive and detective security controls! Different components of a typical attack flow will be complicated / prevented by proper network segmentation. If it doesn't stop the adversary, it will at least slow them down.

Here's a few good pointers to get started with segmentation:

- Follow a structured process to start with:
  - Design what the network should look like (this is the phase where most of the efforts will probably be).
  - Implement the created design.
  - Maintain the implement segmentation when new systems are added to the network.
  - Monitor the environment for wrongly configured systems and adapt where required.
- Start simple by segmenting "easy" parts initially: Segment the Guest, Test and Development systems.
- Segment with key attack strategies in mind: How can you break lateral movement strategies? Isolate workstations, domain controllers,...
- Implementation of network segmentation can be done through a variety of means, don't forget some quick wins! This could include, for example, the use of host-based firewalls for workstations or private VLANs.

## Architecture Principle 4 – Network Segmentation – Private VLANs



The problem: All hosts within the same physical network are able to communicate with each other, without network-based restrictions. VLANs provide a logical segregation of networks within the same physical network (i.e. multiple networks with different security requirements connected to the same switch). Private VLANs are a specific type of VLANs that can be used to achieve client isolation.

Lateral movement often relies on stealing local credentials, which are then used to move to different other workstations (e.g. looking for administrative domain credentials). Properly configured private VLANs could render this attack strategy ineffective!

Private VLANs can be used to configure every client (workstation) in its own private VLAN, so it cannot connect to workstations in other VLANs.

Client isolation is a common practice in wireless networks; why not extend it to wired networks?

### Architecture Principle 4 – Network Segmentation – Private VLANs

The problem: All hosts within the same physical network are able to communicate with each other, without network-based restrictions. VLANs provide a logical segregation of networks within the same physical network (i.e. multiple networks with different security requirements connected to the same switch). Private VLANs are a specific type of VLAN that can be used to achieve client isolation.


Even when the overall network hasn't been fully segmented just yet, private VLANs could be a quick win. Lateral movement often relies on stealing local credentials, which are then used to move to different other workstations (e.g. looking for administrative domain credentials). Properly configured private VLANs could render this attack strategy ineffective.

Private VLANs can be used to configure every client (workstation) in its own private VLAN, so it cannot connect to workstations in other VLANs. While client isolation is a very common feature in wireless networks, it's much less frequent in wired networks. Let's reverse this trend!

#### Reference:

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/CLIConfigurationGuide/PrivateVLANs.pdf>

## Architecture Principle 4 – Network Segmentation – Host-Based Firewalls



In traditional network environments, firewall decisions are made on a network zone level, while host-based firewalls are often ignored (or even disabled). The rise of the “mobile worker” and the disappearing perimeter force us to rethink this approach!

You still want to protect your endpoints, even if the employee is not physically in the network.

Host-based firewalls are another “quick win” that can help you achieve endpoint isolation!

The Windows Firewall can be easily centrally managed using GPOs; this should be an essential part of your group policies!

Quick wins include disable in-bound SMB, blocking NetBIOS, blocking commonly abused processes from establishing outbound connectivity,...

### A pretty interesting article:

<https://medium.com/@cryps1s/endpoint-isolation-with-the-windows-firewall-462a795f4cfb>

## Architecture Principle 4 – Network Segmentation – Host-Based Firewalls

In traditional network environments, firewall decisions are made on a network zone level, while host-based firewalls are often ignored (or even disabled). The rise of the “mobile worker” and the disappearing perimeter force us to rethink this approach! Even while employees are not physically present, you still want to secure their endpoints. This is another consequence of the fading perimeter, where we move security controls from the perimeter to the endpoint themselves.

Host-based firewalls are usually already included on your machines (thus no extra investment) and just need to be configured correctly. They can be a good quick win to achieve endpoint isolation.

The Windows Firewall is available on the vast majority of Windows systems and can be easily centrally managed using GPOs; this should be an essential part of your group policies!

Although you will need to assess what kind of rules can be enforced in your environment, there's a few quick wins that can be achieved:

- Disable in-bound SMB connectivity on workstations.
- Disable out-bound SMB connectivity to non-domain systems (careful (!)).
- Blocking NetBIOS
- Blocking commonly abused processes from establishing outbound connectivity.
  - mshta.exe
  - wscript.exe
  - cscript.exe
  - ..

### Reference:


<https://medium.com/@cryps1s/endpoint-isolation-with-the-windows-firewall-462a795f4cfb>



## Hardening Considerations



Hardening should be performed on different types of devices: Workstations and servers, but also smartphones and tablets, network devices, VOIP devices ...



Different operating systems need to be considered for hardening: Windows, Linux, OSX, iOS, Android ... Next to the OS, key applications are to be hardened as well (e.g. browsers).



Various checklists and tools are available detailing what configuration changes to make to harden an operating system or application.

### Hardening Considerations

OS and application hardening are not limited to computers.

It should be clear that we have to harden desktops, laptops, and servers. But these are not the only devices in our enterprise network that constitute our exposed attack surface.

We also have network devices, like switches, routers, wireless access points. Enterprise network devices have various options that can be configured and must be considered for hardening.

Other network-connected devices are VOIP phones and network printers/scanners. These are commodity appliances that are often forgotten when it comes to security. VOIP phones and networked printers are often configured open so that they can serve as many clients as possible. But this enlarges the attack surface and should be reduced by hardening. As an example of a large attack surface, we want to mention that network printers often have an ftp server to facilitate printing of uploaded documents, and that such ftp servers have been used as pivot points in attacks.

Smart devices like smartphones and tablets should also be included in an OS hardening program.

OS hardening is not limited to Windows. Other operating systems also have features and options that dictate the attack surface, and that can be reduced by disabling features and options. Linux and OSX should be hardened for workstations, but operating systems for smart devices like iOS and Android should not be left out of the picture.

## NIST Checklist

The National Institute of Standards and Technology (NIST) publishes checklists that can be used to harden operating systems and applications:



- These security checklists are published in the National Checklist Program Repository
- The checklists are available in various formats, varying from text for humans to a formalized format for programs
- Security Content Automation Protocol (SCAP) is a formalized format

### NIST Checklist

One of the most known providers of security checklists to harden operating systems and applications is the US National Institute of Standards and Technology: NIST.

This organization has a long tradition of publishing recommendations to configure operating systems and applications, to make them harder to attack. This started with detailed instructions written in a guide that system administrators would read and apply to their system. To get an idea of the level of detail, such guides would mention registry keys and values to be configured in Windows.

The security checklists are stored in the repository of the National Checklist Program, which can be accessed by going to <https://nvd.nist.gov/ncp/repository>.

Providing guides written for humans was a good starting point, but due to the amount of technical detail in these guides, applying the recommendations to an operating system involved a lot of work that was error-prone. Scaling was another problem: Applying the recommendations of a NIST checklist to a corporate network involved thousands of machines.

That is why NIST evolved its checklist repository to include formats that support automation: Checklists can be downloaded in a machine-readable form, that can be applied by configuration applications to harden operating systems and applications.

The Security Content Automation Protocol (SCAP) is a format designed to automate vulnerability assessment and management. The NIST security checklists are available in SCAP format so that they can be used for automatic processing.

Checklists can be used to check the configuration of a system, and to change the configuration according to the recommendations of the checklist.

## NIST's Repository

NIST's Repository

NIST National Institute of Standards and Technology U.S. Department of Commerce

Computer Security Resource Center National Vulnerability Database

NIST Security Configuration CHECKLISTS <http://checklists.nist.gov> CSO

NCP > Checklists

**National Checklist Program Repository**

The availability of publicly available information on the NIST Security Configuration Checklist Program (SCAP). SCAP enables validated information page or the glossary of terms.

Search for Checklists using the fields below. The keyword search will search across the name, and summary.

Tier: Any..... Category: Any..... Search

Type: Any..... Authority: Any..... Reset

Target Product: Any..... Keyword:

There are 436 matching records. Displaying matches 1 through 20.

### NIST's Repository

The security checklist repository of NIST can be accessed by visiting <https://nvd.nist.gov/ncp/repository>.

As can be seen above, this will direct us to a website that enables us to specify the checklists we look for.

We select Target Product "Windows 10" as an example.

By clicking on the search button, we get four results.

The first result is "Windows 10 STIG" produced by the Defense Information Systems Agency.


STIG stands for Security Technical Implementation Guide. It covers Windows 10 systems member of a domain and, therefore, covers Windows 10 Enterprise.

It supports the SCAP format, and also OVAL. This is an XML format meant for machine consumption, but it can also be rendered in human-readable form, although with a bit of practice, information can be obtained from its raw form. For example, this:

```
<criteria operator="AND">
  <criterion test_ref="oval:mil.disa.fso.windows:tst:388900" comment="Verifies Telnet Client feature is not installed" />
</criteria>
```

With this XML chunk, we know that one of the checks is for the presence of Telnet. Telnet is a cleartext, unauthenticated remote access service, and should never be used.



## Security Technical Implementation Guide (STIG)



Security Technical Implementation Guides (STIG) are made available by a variety of sources. One of the most well-known are the DoD General Purpose STIG, which provides tools, checklists,... for facilitated implementation of the STIG. They are available for a wide variety of Operating Systems.

STIG standards provide an "easy-to-refer-to" standard and can be used as a baseline for your hardening

The STIG can be further fine-tuned and adapted according to what is required in your organization



**Key message: Do not reinvent the wheel; reuse and adapt!**

### Security Technical Implementation Guide (STIG)

Security Technical Implementation Guides (STIG) are made available by a variety of sources. One of the most well-known are the DoD General Purpose STIG, which provides tools, checklists,... for facilitated implementation of the STIG. They are available for a wide variety of Operating Systems.

STIG standards provide an "easy-to-refer-to" standard and can be used as a baseline for your hardening.

The STIG can be further fine-tuned and adapted according to what is required in your organization

#### Reference:

<https://iase.disa.mil/stigs/Pages/index.aspx>

## Microsoft Recommended Security Baseline

Microsoft also publishes recommended security baselines; these can be checked and applied with the free **Security Compliance Toolkit**. However, the SCT is only available as of Windows 10 and Windows Server 2012 & 2016.

From Microsoft's official website:

*“This set of tools allows enterprise security administrators to download, analyze, test, edit and store Microsoft-recommended security configuration baselines for Windows and other Microsoft products, while comparing them against other security configurations.”*

The SCT enables administrators to effectively manage their enterprise's Group Policy Objects (GPOs). Using the toolkit, administrators can compare their current GPOs with Microsoft-recommended GPO baselines or other baselines, edit them, store them in GPO backup file format, and apply them broadly through Active Directory or individually through local policy.

### Microsoft Recommended Security Baseline

Another source for guidelines to harden Microsoft Windows is Microsoft itself.

Microsoft publishes recommended security baselines and security guides.

These baselines can be automatically checked and applied with tools offered by Microsoft.

One of these tools is the Security Compliance Manager.

The SCM is a free tool from Microsoft. It is not part of a default Windows install but needs to be downloaded and installed.

SCM 4.0 supports the latest Windows versions (Server 2016 and Windows 10).

The tool not only allows managing Microsoft produced security baselines, but these can be modified to better suit your enterprise's environment. New templates can be created, too, for custom configurations.

SCM can be used to deploy configurations to a stand-alone machine, too: Machines that are not members of a domain.

Baselines can be exported to Microsoft-supported formats like GPO policies, but also open formats like SCAP used by NIST, so that Microsoft Security Baselines can also be converted SCAP and audited with SCAP tools.

## Hardening Browsers – Not Only IE / Edge – ADMX Templates!



- Traditionally, Microsoft browsers such as Internet Explorer and Edge have been the browsers of choice in enterprise environments, as they can easily be managed (and hardened) using GPOs
- Many end-users, however, use a different browser personally (Firefox, Chrome...) and thus often look for a way to use these browsers (e.g. portable apps)
- As an interesting solution, ADMX templates can be used to perform enterprise hardening of browsers such as Chrome or Firefox!

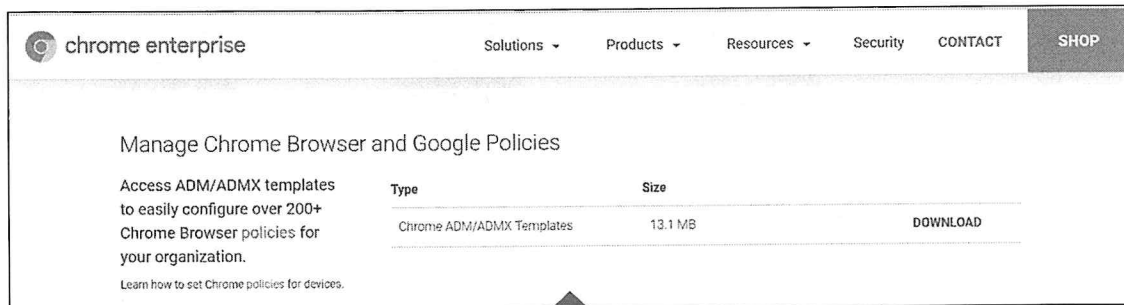
### Hardening Browsers – Not Only IE / Edge – ADMX Templates!

As nearly all applications are "web-based" (or will be in the future), browser hardening is a big topic! Traditionally, Microsoft browsers such as Internet Explorer and Edge have been the browsers of choice in enterprise environments, as they can easily be managed (and hardened) using GPOs.

These days, however, many alternatives have surfaced, which are providing an interesting alternative. Many end-users use a different browser personally (Firefox, Chrome...) and thus often look for a way to use these browsers (e.g. portable apps). Some of these browsers (e.g. Chrome) have an excellent reputation for security controls, and they are definitely worthwhile considering as a corporate browser.

However, the issue is often linked to manageability... To what extent can these "third-party" browsers be centrally managed in an AD environment? It turns out several options exist! ADMX templates can, for example, be used to perform enterprise hardening of browsers such as Chrome or Firefox!

## Hardening Browsers – Chrome Enterprise



The screenshot shows the Chrome Enterprise website interface. At the top, there is a navigation bar with the Chrome Enterprise logo and links for Solutions, Products, Resources, Security, CONTACT, and SHOP. The main content area is titled "Manage Chrome Browser and Google Policies" and includes a description: "Access ADM/ADMX templates to easily configure over 200+ Chrome Browser policies for your organization." Below this is a table with columns for Type and Size, and a DOWNLOAD button.

Type	Size	
Chrome ADM/ADMX Templates	13.1 MB	DOWNLOAD

### Chrome Enterprise


Google has worked hard to promote Chrome as an enterprise-grade browser! It has a dedicated website where it lists, among others, the security controls implemented in Chrome, but also support options for Chrome. They also offer a free download of a set of ADMX templates that can be used to harden Google Chrome using GPOs!

### Hardening Browsers – Chrome Enterprise

Let's take the example of Google Chrome. Over the past few years, Google has worked hard to promote Chrome as an enterprise-grade browser! It has a dedicated website where it lists, among others, the security controls implemented in Chrome, but also support options for Chrome. It very much highlights the different security features available! Furthermore, they also offer a free download of a set of ADMX templates that can be used to harden Google Chrome using GPOs!

We will use these ADMX templates in our lab environment!

## Linux Hardening – Ansible Automation



Ansible is open-source software that automates software provisioning, configuration management, and application deployment. Using its configuration-management feature, we can use it to audit and enforce security settings on a variety of Operating Systems.

Ansible uses playbooks to change / audit configuration of the OS and any application installed on the OS.

For security purposes, you could create a standard configuration role for your "Linux Web Server" (example), after which it can be used to deploy / enforce / audit systems.



<https://github.com/openstack/ansible-hardening>

Ansible hardening roles have been created by OpenStack and are available on GitHub!

### Linux Hardening – Ansible Automation

Ansible is open-source software that automates software provisioning, configuration management, and application deployment. Using its configuration management feature, we can use it to audit and enforce security settings on a variety of Operating Systems. Ansible was acquired by RedHat in 2015.

Ansible uses playbooks to change / audit configuration of the OS and any application installed on the OS. For security purposes, you could create a standard configuration role for your "Linux Web Server" (example), after which it can be used to deploy / enforce / audit systems.

The guys from OpenStack have done some amazing work and have drafted roles to secure a wide variety of Linux flavors, including:

- CentOS
- Debian
- Fedora
- OpenSUSE
- Red Hat
- SUSE
- Ubuntu

#### Reference:

<https://github.com/openstack/ansible-hardening>



# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.1

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim


### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

**Exercise: Hardening Our Domain Using SCT and STIG**



Please refer to the workbook for further instructions on the exercise!

This page intentionally left blank.

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.I

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim

### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

## A Fundamental Detection Capability



Throughout the course, we will walk through the kill chain and focus on a variety of different security controls that can help stop (advanced) adversaries in their attempts to penetrate your environment. We should, however, understand that a "prevent-only" approach is not sufficient, especially when dealing with targeted attacks.

### So, what do we require for a proper detection capability?



A central logging platform that can parse, index and visualize collected information



Network device logs; key focus areas include DNS, web proxy, firewall, IDS ...



Endpoint (workstation and server) visibility using real-time log and periodic data collection



Optionally, a full packet capture solution that acts as a "flight recorder" for egress & ingress traffic

**Depending on your environment, some log sources might be more important than others!**

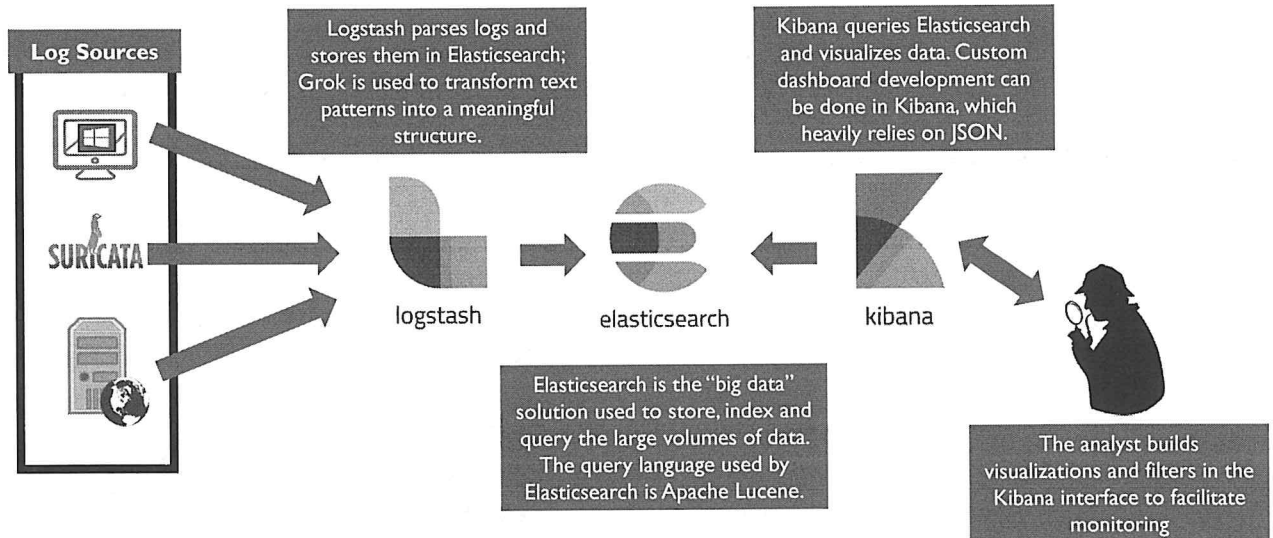
## A Fundamental Detection Capability

Although this is not a purely "detection-driven" course, we will spend time explaining what a fundamental detection capability looks like. Throughout the course, we will walk through the kill chain and focus on a variety of different security controls that can help stop (advanced) adversaries in their attempts to penetrate your environment. We should, however, understand that a "prevent-only" approach is not sufficient, especially when dealing with targeted attacks. By their nature, these adversaries are equipped to bypass existing secure controls and we should thus ensure we can at least observe them bypassing our controls.

### So, what do we need for a proper detection capability?

- A central logging platform that can parse, index and visualize collected information. Purely log centralization solutions include Splunk, Graylog, ELK ... Furthermore, many different vendors offer solutions that can implement correlation on top of this (typical SIEM solutions).
- Network device logs, of which the key focus areas should be DNS logs, web proxy logs, firewalls and IDS infrastructure.
- An often-overlooked part is endpoint visibility. Endpoint (workstation and server) visibility can be attained using either real-time logs or periodical data collection. Built-in tools can be used for this (e.g. Windows Events, Sysmon ...) or organizations can choose to deploy additional commercial or open-source tooling (EDR tooling, OSQuery ...).
- Optionally, an interesting addition could be the implementation of a full packet capture solution at your perimeter that acts as a "flight recorder" for egress and ingress traffic.

## A Central Logging Platform – Introducing the Elastic Stack



### A Central Logging Platform – Introducing the Elastic Stack

The Elastic stack (formerly "ELK") consists of three products working together, namely Elasticsearch, Logstash, and Kibana.


ElasticSearch is the big data solution and is used to store, index, and query the large volumes of data. Its functionality is similar to Splunk. However, some of the underlying technologies used are different. Elasticsearch makes use of Apache Lucene for information retrieval, originally completely written in Java, but meanwhile ported to C++ and Python, among others.

Logstash is used for parsing logs submitted to the stack and stores the results in Elasticsearch. Logstash uses Grok to transform text patterns into a meaningful structure. Grok is perfect for syslog logs, Apache and other web server logs, mysql logs, and in general, any log format that is written for humans and not computer consumption.


Kibana takes care of the graphical component of the stack and visualizes data that it queries from Elasticsearch. Kibana can be used to implement custom dashboards, which heavily relies on JSON. Kibana has all the classics such as histograms, line graphs, and pie charts. Next to all that, Kibana is also able to create geo maps, time series and to analyze relationships or anomalies using machine learning.

Pre-installed and pre-configured Elastic appliances are available online from a variety of sources (e.g. Bitnami...). Furthermore, several online services offer "Elastic stacks as a service" (e.g. <https://logz.io/>)

## A Central Logging Platform – A Note about Security



elasticsearch



logstash



kibana

One important point to raise is that a default Elastic stack install requires a number of security controls to be implemented:

- Kibana dashboards are by default accessible over HTTP without authentication  
Remediation: Implement HTTPS + authentication on web server
- Elasticsearch engine cluster can be interacted with on TCP port 9200 without authentication  
Remediation: Network segmentation or only listen on local interface

### A Central Logging Platform – A Note about Security

Since we want to avoid wrapping all our data and information in a nice gift for adversaries to pick up, we need a number of security controls. By default, Kibana dashboards are reachable through HTTP without any form of authentication. This would allow anyone that is able to reach the server to view the dashboards. The solution is to implement HTTPS and authentication on the web server.

Additionally, the Elasticsearch engine cluster has a service running that can be interacted with on TCP port 9200 without authentication, which could allow an outsider to read data or shut down the cluster. This problem can be solved by network segmentation or only listening on a local interface, which would prevent access from another device on the same network.


## Logstash Configuration and Parsers



Depending on the type of logs we want to analyze, Logstash configuration can be trivial or complex: Structured data formats such as XML or key-value pairs like JSON are easy to parse, while raw logs will require some analysis in order to understand how they are to be parsed. Grok provides endless flexibility to build structure in a big dump of raw data...

Logstash parsing configurations for the majority of known log formats have been developed by the community and are freely available online. Some examples include:

- <https://github.com/cvandeplas/ELK-forensics>
- <https://github.com/philhagen/sof-elk>
- <https://github.com/HASecuritySolutions/Logstash>
- <https://github.com/CuBoulder/logstash>



### Logstash Configuration and Parsers

Depending on the type of logs we want to analyze, Logstash configuration can be trivial or complex.

Structured data such as XML, or key-value pairs like JSON are easy to parse since these data types are already well-structured and aimed at automated processing. Raw logs require some analysis in order to find a certain structure and understand how they should be parsed.

The community has already developed parsing configurations for the majority of known log formats. These configurations are available online for free and can thus be reused in your Logstash configuration. We have included a nice overview of available community work on the next slide.

A useful plugin to help with structuring logs is Grok. It is currently the best way to parse unstructured log data into something that is structured and queryable. Grok is perfect for syslog logs, Apache and other web server logs, mysql logs, and in general, any log format that is written for humans and not computer consumption. It works by combining text patterns into something that matches your logs.


The following repositories contain some useful, free Logstash configurations:

- <https://github.com/cvandeplas/ELK-forensics>
- <https://github.com/philhagen/sof-elk>
- <https://github.com/SMAPPER/Logstash-Configs>
- <https://github.com/CuBoulder/logstash>

There are configurations for Bluecoat, Checkpoint, IIS, McAfee, ESXi, and many more. One of the repositories is used in another SANS course, namely FOR572, which is aimed at advanced network forensics and analysis.

It should be noted that these configuration files could have minor differences, depending on the Logstash version you deploy. Some minor adaptations could thus be required.

## Fundamental Endpoint Logs – Windows Event Logs



Windows event logs are log files kept to keep track of all sorts of events. They are typically stored locally in C:\Windows\System32\winevt\Logs\. Windows event logs are a GREAT way of obtaining visibility on what is happening on your endpoints

- Logs can be viewed locally with the event viewer (eventvwr.msc) or with various command line tools (wevtutil.exe, PsLogList ...). Windows event logs can be forwarded and centralized to a central repository such as Splunk, QRadar, ArcSight, Elastic,...
- Many Windows components have dedicated logs in the Windows event log to which they write information (e.g. Defender, Exploit Guard, PowerShell, Applocker,...)
- Windows event logs are an excellent, free, built-in means of obtaining visibility on Windows endpoints (both workstations and servers); they do need to be centralized to be useful!

### Fundamental Endpoint Logs – Windows Event Logs

Since the introduction of Windows NT in 1993, the Windows operating system has kept logs of events that took place while the operating system was running. Events are very important for managing and troubleshooting Windows but are also very important for security. For example, to detect intrusions while they are happening, or for post-facto digital forensic investigations.

Applications and the Windows kernel can generate events. For example, Windows will generate several events when it is started, and many applications will generate events when they encounter an error.

Windows events have a well-defined format and are stored in files called event logs, depending on their "source." For example, events that originate from the Windows kernel when Windows is started are stored in the System event log, while error events from applications (without dedicated event log) are stored in the Application event log.


By default, event logs are stored locally in folder C:\Windows\System32\winevt\Logs\. Up to Windows XP, a proprietary, binary file format was used to store event logs: The EVT format. Starting with Windows Vista, the format used for Windows XML EventLog is EVTX. This is a binary XML format. By switching to XML, Microsoft opens up its event logs for even more interoperability with third parties, but XML is a very verbose text format that would negatively impact performance. Therefore, Microsoft opted for a binary XML format.

There are two native Windows tools to view event logs. There is the event viewer, a graphical user interface tool that is a snap-in for the Microsoft Management Console (MMC). Wevtutil.exe is a command line tool.

Event logs can also be viewed with PowerShell and many third-party tools, like Sysinternals' PsLogList.



## Fundamental Endpoint Logs – Standard Windows Event Logs



"Spotting the Adversary with Windows Event Log Monitoring" is the reference on event log monitoring, published by the NSA/CSS.

Windows Vista and above Events	General Event IDs
Account and Group Activities	4624, 4625, 4648, 4728, 4732, 4634, 4735, 4740, 4756
Application Crashes and Hangs	1000 and 1002
Windows Error Reporting	1001
Blue Screen of Death (BSOD)	1001
Windows Defender Errors	1005, 1006, 1008, 1010, 2001, 2003, 2004, 3002, 5008
Windows Integrity Errors	3001, 3002, 3003, 3004, 3010 and 3023
EMET Crash Logs	1 and 2
Windows Firewall Logs	2004, 2005, 2006, 2009, 2033
MSI Packages Installed	1022 and 1033
Windows Update Installed	2 and 19
Windows Service Manager Errors	7022, 7023, 7024, 7026, 7031, 7032, 7034
Group Policy Errors	1125, 1127, 1129
AppLocker and SRP Logs	865, 866, 867, 868, 882, 8003, 8004, 8006, 8007
Windows Update Errors	20, 24, 25, 31, 34, 35
Hotpatching Error	1009
Kernel Driver and Kernel Driver Signing Errors	5038, 6281, 219
Log Clearing	104 and 1102
Kernel Filter Driver	6
Windows Service Installed	7045

Table 1 of "Spotting the Adversary with Windows Event Log Monitoring" provides a very good overview of essential event IDs that should be monitored.

### Fundamental Endpoint Logs – Standard Windows Event Logs

"Spotting the Adversary with Windows Event Log Monitoring" is a good reference on Windows event log monitoring, published by the NSA/CSS. This lengthy document (around 50 pages) is a must read for blue teams. It was last reviewed 16 July 2015. The following topics are covered in detail:

- Deployment
- Hardening Event Collection
- Recommended Events to Collect
- Event Log Retention
- Final Recommendations

Deployment will not only cover the configuration of Windows event logs but also centralization of these logs using Microsoft's publisher/subscriber model. It must be said that this is not the only way to centralize event logs. There are many other systems, for example, Splunk comes to mind.

<https://apps.nsa.gov/iaarchive/library/ia-guidance/security-configuration/applications/spotting-the-adversary-with-windows-event-log-monitoring.cfm>

Table 1 of "Spotting the Adversary with Windows Event Log Monitoring" provides a very good overview of essential event IDs that should be monitored.

The event IDs themselves are not explained in much detail in this document, however. Let's illustrate this with event ID 4648. This ID is mentioned first in table 1 (under account and group activities), and second in table 9. From table 9, we know that this event is from the Security event log. The description of this event is "Account login with explicit credentials." However, the document does not provide more information about this event, and why it is important.



## Fundamental Endpoint Logs – Malware Archaeology Mapping to ATT&CK

Tactic	Technique Name	Technique ID	Data Source 1	Data Source 2	Data Source 3	Data Source 4	Data Source 5	Data Source 6
Collection	Audio Capture	T1123	4688 Process Execution	4663 File monitoring	API monitoring			
Collection	Automated Collection	T1119	4688 Process CMD Line	4663 File monitoring	Data loss prevention			
Collection	Clipboard Data	T1115	API monitoring					
Collection	Data from Information Repositories	T1213	Application Logs	Authentication logs	Data loss prevention	Third-party application logs		
Collection	Data from Local System	T1005	4688 Process Execution	4688 Process CMD Line	200-500, 4100-4104 PowerShell logs	4663 File monitoring	S861 WMI	
Collection	Data from Network Shared Drive	T1039	4688 Process CMD Line	4688 Process Execution	S140/S145 Share connection	4663 File monitoring		
Collection	Data from Removable Media	T1025	4688 Process Execution	4688 Process CMD Line	4657 Windows Registry	4663 File monitoring	S140/S145 Net Shares	
Collection	Data Staged	T1074	4688 Process CMD Line	4688 Process Execution	4663 File monitoring			
Collection	Email Collection	T1114	4688 Process Execution	5156 Firewall Logs	4624 Authentication logs	4663 File monitoring		
Collection	Man in the Browser	T1185	4624 Authentication logs	4688 Process Execution	API monitoring	Packet capture		
Collection	Screen Capture	T1113	4688 Process Execution	4663 File monitoring	API monitoring			

### ATT&CK Mapping

Another testament to the "rise of ATT&CK" is the mapping created by Malware Archaeology of Windows event IDs to the MITRE ATT&CK framework.

It includes a coding scheme for most relevant event identifiers as well!

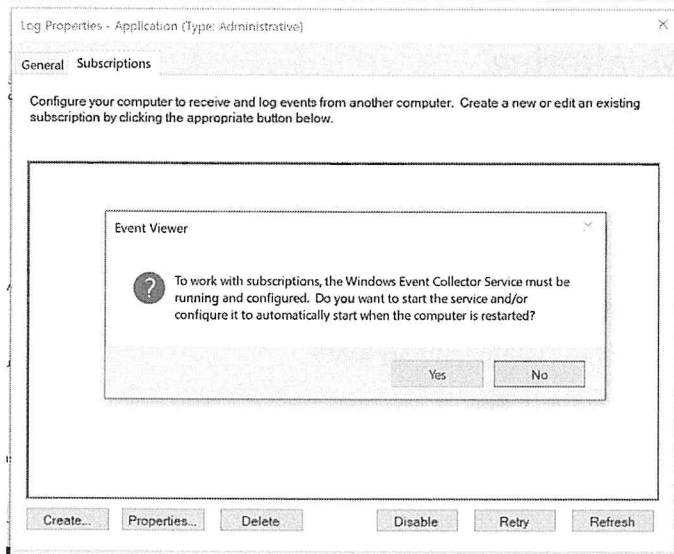
It's updated regularly and can be found at <https://www.malwarearchaeology.com/cheat-sheets>.

### Fundamental Endpoint Logs – Malware Archaeology Mapping to ATT&CK

Another testament to the "rise of ATT&CK" is the mapping created by Malware Archaeology of Windows event IDs to the MITRE ATT&CK framework. It includes a coding scheme for most relevant event identifiers as well! It's updated regularly and can be found at <https://www.malwarearchaeology.com/cheat-sheets>.

We will not add more detail or information about this right now, but it's important to note that these types of guides exist and can be used to configure how standard Windows logging can be fine-tuned without installation of additional (commercial) tools.

## Fundamental Endpoint Logs – Centralizing Windows Event Logs



In an enterprise environment, however, we cannot query systems individually, so we need to centralize the event logs for storage and analysis.

“Subscriptions” are a Microsoft technology to handle this. To start using subscriptions, the Windows Event Collector & Windows Remote Management services must be running.

When the Windows Event Collector Service is started, subscriptions can be created.

Subscriptions can be pull (collector initiated) or push (source computer initiated) subscriptions.

### Fundamental Endpoint Logs – Centralizing Windows Event Logs

Event logs can be centralized. There are different third-party solutions for this, like Splunk, NXLog or Winlogbeat (Elastic), but it can also be done using Microsoft's technology that comes out-of-the-box with Windows. This is called Windows Event Forwarding and is done with subscriptions.

Once the Windows Event Collector Service and Windows Remote Management Service are started, subscriptions can be created. Each subscription requires a name (to be chosen by the administrator), and an optional description. Subscriptions can be pull (collector initiated) or push (source computer initiated) subscriptions.

In a collector-initiated subscription (pull), the source computers have to be selected. These have to be domain members. In a source computer-initiated subscription (push), the source computers groups have to be selected. These groups can be domain members or non-domain members. In the case of non-domain members, certificates have to be provided for authentication and encryption.

The advantage of a source computer-initiated subscription is that event logs from non-domain computers can be collected, too. Forwarding all events from many computers would create too much data; therefore, events can be filtered prior to forwarding.

Windows Event Forwarding requires no additional software and can be configured with a few simple steps. We will illustrate this via the eventvwr dialog, and with a command-line configuration, but of course, this configuration is not something you would do manually on each machine. This is something you would script with Active Directory GPOs.

## Fundamental Endpoint Logs – Extending Windows Event Logs Using Sysmon

Sysmon is short for System Monitoring.

Sysmon installs a Windows service and a device driver.

These components monitor activity on a system:

- Creation and termination of processes
- Loading of executable images
- Network connection establishing
- ...

Activity is recorded with events.

### Sysmon v8.0

05/22/2017 • 12 minutes to read • Contributors

By Mark Russinovich and Thomas Garnier

Published: July 5, 2018



[Download Sysmon \(1.4 MB\)](#)

### Fundamental Endpoint Logs – Extending Windows Event Logs Using Sysmon

Sysmon is a system monitoring tool that is part of the Sysinternals Suite.

It is a tool that was originally developed for Microsoft. It is deployed on many of their servers and workstations to monitor system activity. In case of incidents, Sysmon provides a valuable log of system activities that can help forensic investigators to reconstruct an incident.

Sysinternal tools are stand-alone tools that don't come with an installer (like setup.exe or install.msi). For Sysmon, there is Sysmon.exe and Sysmon64.exe.

Sysmon.exe is a 32-bit version that embeds the 64-bit version, too. Sysmon64.exe is a 64-bit version only; it is provided for Windows systems that only support 64-bit executables, and not 32-bit (Windows Servers without 32-bit subsystem).

If the 32-bit version is executed on a 64-bit OS, it will extract the 64-bit version and run that instead. When Sysmon is installed on a Windows machine, it installs a Windows service and a device driver. These components are necessary to detect and record system activities like the creation and termination of processes, loading of executable images, creation of network connections, loading of drivers, ...

All this activity is logged in a dedicated Windows event log.

## Fundamental Endpoint Logs – Sysmon – Additional Event Types

Since Sysmon version 6.10, it records 22 different types of events:

Event ID	Description	Event ID	Description
1	Process creation	12	RegistryEvent (Object create and delete)
2	A process changed a file creation time	13	RegistryEvent (Value Set)
3	Network connection	14	RegistryEvent (Key and Value Rename)
4	Sysmon service state changed	15	FileCreateStreamHash
5	Process terminated	16	Sysmon Configuration Changed
6	Driver loaded	17	Pipe Created
7	Image loaded	18	Pipe Connected
8	CreateRemoteThread	19	WmiEventFilter activity detected
9	RawAccessRead	20	WmiEventConsumer activity detected
10	ProcessAccess	21	WmiEventConsumerToFilter activity detected
11	FileCreate	255	Sysmon internal error

### Fundamental Endpoint Logs – Sysmon – Additional Event Types

Since Sysmon version 6.10, it records 22 different types of events. These events monitor objects like processes, files, registry objects, ... But also, events to monitor changes to Sysmon itself (IDs 4 and 16) can be logged. This can indicate tampering attempts.

Other event IDs that can be indicative of tampering by malicious actors are changes in file creation times (ID 2), creation of remote threads (ID 8), often used for code injection, opening of processes (ID 10),...

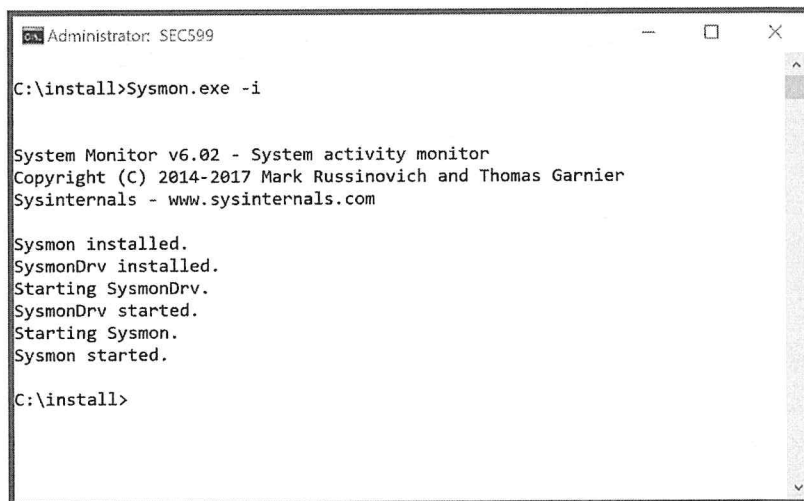
## Fundamental Endpoint Logs – Sysmon – Installation

### Confirmation

After acceptance and installation, Sysmon will report which components were installed.

In this case, the default configuration is used.

Sysmon is immediately and completely active, no reboots are required.



```
Administrator: SEC599
C:\install>Sysmon.exe -i

System Monitor v6.02 - System activity monitor
Copyright (C) 2014-2017 Mark Russinovich and Thomas Garnier
Sysinternals - www.sysinternals.com

Sysmon installed.
SysmonDrv installed.
Starting SysmonDrv.
SysmonDrv started.
Starting Sysmon.
Sysmon started.

C:\install>
```

### Fundamental Endpoint Logs – Sysmon – Installation

To install Sysmon on a machine, start command “sysmon.exe -i” from an elevated, administrative command prompt. Before the installation of the Service and driver takes place, a dialog with the EULA will be presented to the user. This EULA has to be accepted to proceed with the installation. This EULA is presented the first time sysmon is executed. This is the case for all Sysinternals tools: After acceptance of the EULA, the dialog is no longer displayed.

As this dialog hinders unattended deployment of a tool like sysmon.exe, the EULA can also be accepted by providing option --accepteula on the command line. With this option, the dialog box with the EULA will never be displayed. Remark that typing --accepteula is considered equivalent to clicking on the Agree button: You accept the EULA.

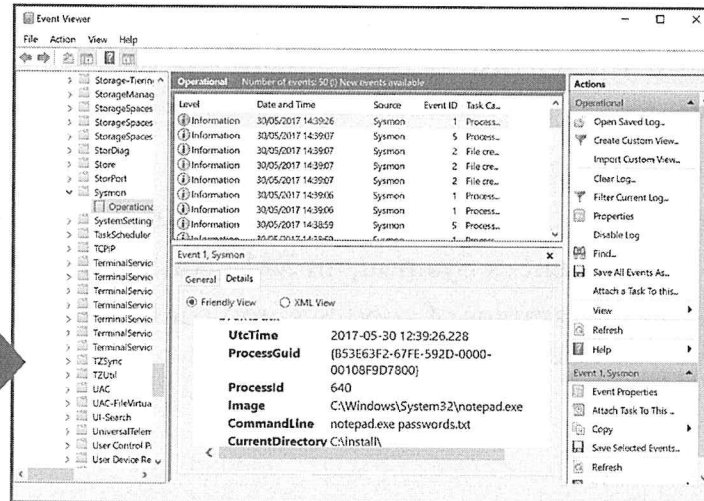
After accepting the EULA, Sysmon will install its service (Sysmon) and its driver (SysmonDrv). In the screenshot above, we can see that installation was successful: The service and the driver were installed and started. No reboot is required. Without configuration file or configuration option, Sysmon will be installed with the default configuration and will run with these settings until they are changed.

Sysmon can be deployed as an application with group policies on all the machines in your environment. If you need custom configurations (which we recommend), it can be easier to package Sysmon and the configuration file in an MSI file and deploy that through group policies.

## Fundamental Endpoint Logs – Sysmon – Default Configuration

With the default configuration, Sysmon will start logging process creation and termination to a dedicated Windows event log.

In this event, we can see notepad.exe was executed to edit a file with the name passwords.txt



SANS

SEC599 | Defeating Advanced Adversaries – Purple Team Tactics & Kill Chain Defenses

127

### Fundamental Endpoint Logs – Sysmon – Default Configuration

Once Sysmon is installed, it will start monitoring system activity and create events in a dedicated Windows event log (Sysmon / Operational). These events contain a lot of information that can help us reconstruct what happened on a machine.

In the example above, we see a Process Creation event (Event ID 1): This event is created each time a new program is launched. The program that was launched is notepad.exe.

From the commandline property, we can see what file was edited with notepad.exe: passwords.txt

Sysmon not only records the Process ID, but also a Process GUID. Process IDs are a 16-bit integer. Because of this limitation (65536 possible combinations), Process IDs can be reused: For example, after a couple of days, depending on the amount of activity of your system, a new process can be created that will also have the PID 640.

Since this can be a problem for log correlation, Sysmon will create a unique Process GUID per process. Another useful piece of information in a Process Creation event is the cryptographic hash of the content of the executable (notepad.exe in our case). By default, this is a SHA1 hash, but Sysmon can be configured to use more hashes, too.

By default, Sysmon will not generate all types of events. Sysmon needs to be configured, but it can take up too many resources if it is configured to record all events.



## Fundamental Endpoint Logs – Sysmon – Tailor Configuration

Sysmon needs to be properly configured to maximize the amount of useful information without overloading the system or producing too many events:

- It can be useful to increase the event log size.
- Centralizing this event log prevents tampering by adversaries.
- For a large deployment of Sysmon, an XML configuration file can be used.
- There are several examples of good Sysmon configuration files on the internet, like this one from @SwiftOnSecurity.

### Fundamental Endpoint Logs – Sysmon – Tailor Configuration

As Sysmon can create a huge amount of events, that can impact performance of the machine, or flood the event log with events thereby flushing out older events, a system needs to be configured properly for Sysmon to operate well.

By configuring Sysmon via a configuration file, filtering can be applied. Filtering denotes which type of events need to be included or excluded. For example, one could configure Sysmon to not log network connections established by Internet Explorer, but to do log network connections established by Word. This will already significantly reduce the amount of events.

Increasing the size of the Windows event logs is a good idea to increase the retention period (once an event log is full, old events are discarded to make a place for new events). Centralizing event logs can also alleviate the lack of space and is also a solution to combat log tampering.

The Sysmon configuration file is an XML file. This is very useful to deploy and configure Sysmon on many systems.

Good examples of Sysmon configuration files can be found on the internet, like internet security legend @SwiftOnSecurity. This configuration file has been deployed on thousands of systems and has been in use for more than a year.

## Fundamental Endpoint Logs – Sysmon – Tailor Configuration – SwiftOnSecurity

```
66 <!--SYSMON EVENT ID 1 : PROCESS CREATION [ProcessCreate]-->
67 <!--COMMENT: All process launched will be included, except for what matches a rule below. It's best to be as specific as
68 avoid user-mode executables imitating other process names to avoid logging, or if malware drops files in an existin
69 Ultimately, you must weigh CPU time checking many detailed rules, against the risk of malware exploiting the blindr
Beware of Masquerading, where attackers imitate the names and paths of legitimate tools. Ideally, you'd use both fi
code signatures to validate, but Sysmon does not support that. Look into Windows Device Guard for whitelisting supp

<!--DATA:UtcTime, ProcessGuid, ProcessID, Image, FileVersion, Description, Product, Company, CommandLine, CurrentDirectory
<ProcessCreate onmatch="exclude">
  <!--SECTION: Microsoft Windows-->
  <CommandLine condition="begin with">C:\Windows\system32\DllHost.exe /Processid</CommandLine> <!--Microsoft:Windows-
  <CommandLine condition="is">C:\Windows\system32\SearchIndexer.exe /Embedding</CommandLine> <!--Microsoft:Windows: S
  <Image condition="is">C:\Windows\system32\CompatTelRunner.exe</Image> <!--Microsoft:Windows: Customer Experience Im
  <Image condition="is">C:\Windows\system32\audiodg.exe</Image> <!--Microsoft:Windows: Launched constantly-->
  <Image condition="is">C:\Windows\system32\conhost.exe</Image> <!--Microsoft:Windows: Command line interface host pr
  <Image condition="is">C:\Windows\system32\musNotification.exe</Image> <!--Microsoft:Windows: Update pop-ups-->
  <Image condition="is">C:\Windows\system32\musNotificationUx.exe</Image> <!--Microsoft:Windows: Update pop-ups-->
  <Image condition="is">C:\Windows\system32\powercfg.exe</Image> <!--Microsoft:Power configuration management-->
  <Image condition="is">C:\Windows\system32\sndVol.exe</Image> <!--Microsoft:Windows: Volume control-->
```

### SwiftOnSecurity

SwiftOnSecurity's Sysmon configuration file is very neatly documented and whitelists a large number of default Windows entries that could otherwise generate noise in the environment!

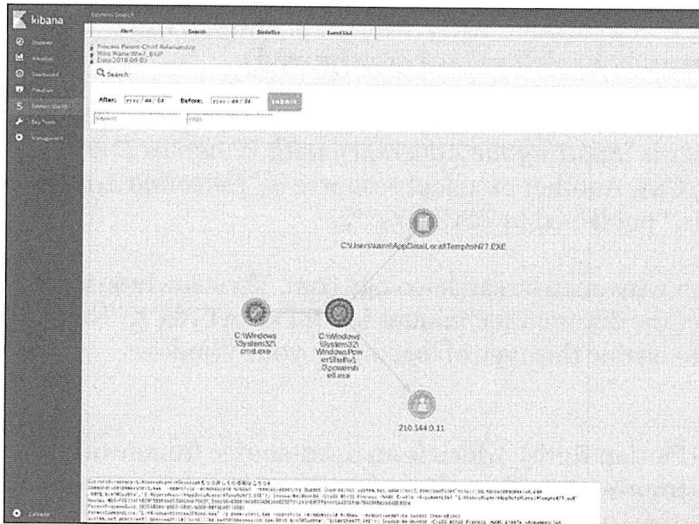
### Fundamental endpoint logs – Sysmon – Tailor Configuration – SwiftOnSecurity

The SwiftOnSecurity Sysmon configuration file attempts to limit the noise that gets generated by Sysmon's detailed logging capabilities. A key way of achieving this is by whitelisting large amounts of default Windows entries.

It's a widely popular Sysmon configuration file that is used by many different organizations! We will use it during the upcoming lab!



## Fundamental Endpoint Logs – Sysmon – SysmonSearch Visualization



### SysmonSearch (JPCERT)

JPCERT/CC has developed and released a system "SysmonSearch", which consolidates Sysmon logs to perform faster and more accurate log analysis. SysmonSearch aims to easily visualize the relationship between different events, thereby facilitating in-depth analysis by monitoring analysts / incident responders.

In the screenshot on the left, you can see a PowerShell command that was launched, which connected to an IP address and dropped an executable.

## Fundamental Endpoint Logs – Sysmon – SysmonSearch Visualization

JPCERT/CC has developed and released a system "SysmonSearch", which consolidates Sysmon logs to perform faster and more accurate log analysis. SysmonSearch aims to easily visualize the relationship between different events, thereby facilitating in-depth analysis by monitoring analysts / incident responders. In the screenshot on the slide, you can see a PowerShell command that was launched, which connected to an IP address and dropped an executable.

SysmonSearch is built on Elastic and can search Sysmon logs with the following conditions:

- Date
- IP address
- Port number
- Host name
- Process name
- File name
- Registry key
- Registry value
- Hash value

Please refer to the full article for additional details:

<https://blog.jpcert.or.jp/2018/09/visualise-sysmon-logs-and-detect-suspicious-device-behaviour--sysmonsearch.html>

## Fundamental Endpoint Logs – What Event IDs to Enable and Collect?

### The number of logs and event IDs available in Windows is staggering (especially when Sysmon is installed and configured)

- Several white-papers and resources exist that attempt to explain how event IDs can be used to detect attacks. An example of such a paper is "*Spotting the Adversary with Windows Event Log Monitoring*," published by the NSA/CSS. Another excellent resource is "*Detecting Lateral Movement through Tracking Event Logs*," published by JPCERT/CC.
- Several good cheat sheets are available on [www.malwarearchaeology.com](http://www.malwarearchaeology.com). These include what Windows event IDs are required to detect the different techniques in MITRE ATT&CK. They also host configuration files and scripts to enable this type of logging on your Windows endpoints.
- Another interesting initiative is SIGMA (Florian Roth), which is an open-source framework that wants to help security professionals improve how they share detection use-cases in a vendor-neutral way. You can find it on <https://github.com/Neo23x0/sigma>

### Fundamental Endpoint Logs – What Event IDs to Enable and Collect?

As stated before, there are a lot of Windows event logs, and therefore even more event IDs. It requires a lot of practice and experience to understand the most important event IDs and to know how to respond adequately.


To help with the understanding of Windows event logs, several good resources are available on the internet. Microsoft's own documentation is a bit lacking when it comes to event logs. Not all event IDs are clearly documented and explain exactly what they mean.

Several white-papers and resources exist that attempt to explain how event IDs can be used to detect attacks. In addition to "*Spotting the Adversary with Windows Event Log Monitoring*," another excellent resource is "*Detecting Lateral Movement through Tracking Event Logs*," published by JPCERT/CC.

Several good cheat sheets are available on <https://www.malwarearchaeology.com/>. These include what Windows event IDs are required to detect the different techniques in MITRE ATT&CK. They also host configuration files and scripts to enable this type of logging on your Windows endpoints.

Another interesting initiative is SIGMA (Florian Roth), which is an open-source framework that wants to help security professionals improve how they share detection use-cases in a vendor-neutral way. You can find it on <https://github.com/Neo23x0/sigma>.

## Endpoint Visibility – Introducing SIGMA



Sigma is a project by Florian Roth that tries to provide a generic, vendor-neutral, rule format that can be used to describe suspicious or malicious behavior. Most SIGMA rules are also mapped to MITRE's ATT&CK framework. As part of the project, several "converters" have been written that allow you to convert the SIGMA rules to some of the major technologies (including Splunk, Elastic, QRadar, ArcSight, RSA NetWitness,...)

### Endpoint Visibility – Introducing SIGMA

Sigma is a project by Florian Roth that tries to provide a generic, vendor-neutral, rule format that can be used to describe suspicious or malicious behavior. Most SIGMA rules are also mapped to MITRE's ATT&CK framework. As part of the project, several "converters" have been written that allow you to convert the SIGMA rules to certain technologies. Supported technologies include (but are not limited to):

- Splunk
- Elastic
- Windows Defender ATP
- ArcSight
- QRadar
- RSA NetWitness
- Logpoint

From Florian's GitHub page:

*"Sigma is a generic and open signature format that allows you to describe relevant log events in a straight forward manner. The rule format is meant to be flexible, easy to write and applicable to any type of log file. The main purpose of the project is to provide a structured form in which researchers or analysts can describe their once developed detection methods and make them shareable with others."*

#### Reference:

<https://github.com/Neo23x0/sigma>

## Endpoint Visibility – An Example SIGMA Rule

25 lines (23 sloc) | 628 Bytes

```
1 title: Office Macro Starts Cmd
2 status: experimental
3 description: Detects a Windows command line executable started from
4 references:
5 - https://www.hybrid-analysis.com/sample/465eabe132ccb949e75b8a
6 author: Florian Roth
7 logsource:
8   product: windows
9   service: sysmon
10 detection:
11   selection:
12     EventID: 1
13     ParentImage:
14       - '*\WINWORD.EXE'
15       - '*\EXCEL.EXE'
16     Image: '*\cmd.exe'
17   condition: selection
18 fields:
19   - CommandLine
20   - ParentCommandLine
21 falsepositives:
22   - unknown
23 level: high
24
```

```
7 logsource:
8   product: windows
9   service: sysmon
10 detection:
11   selection:
12     EventID: 1
13     ParentImage:
14       - '*\WINWORD.EXE'
15       - '*\EXCEL.EXE'
16     Image: '*\cmd.exe'
17   condition: selection
18 fields:
19   - CommandLine
20   - ParentCommandLine
21 falsepositives:
22   - unknown
23 level: high
24
```

In the example SIGMA rule on the left-hand side, we see an attempt to detect Office Macros that start an additional payload by running cmd.exe.

The description of the rule includes metadata (such as the author name and reference) a description of the required log source, and a clearly defined set of conditions.

### Endpoint Visibility – An Example SIGMA Rule

As stated before, there are a lot of Windows event logs, and therefore even more, event IDs. It requires a lot of practice and experience to understand the most important event IDs and to know how to respond adequately.

To help with the understanding of Windows event logs, several good resources are available on the internet. Microsoft's own documentation is a bit lacking when it comes to event logs. Not all event IDs are clearly documented and explain exactly what they mean.

This situation has led to the rise of dedicated resources for Windows event logs. There are several websites that try to define as much Windows event IDs as possible. One of these sites is EventId.net. This is a subscription-based service. Another one, <https://www.ultimatewindowssecurity.com/securitylog/default.aspx>, will provide some information without registration.

Often, the difficulty is not to understand a specific Windows event ID but understanding in what context a combination of Windows event IDs can be relevant. For example: "What combination of event IDs indicate an adversary generated a golden ticket?" Third parties have also collected all their knowledge and recommendations on Windows event logs in a single resource. An excellent example is the document, "Spotting the Adversary with Windows Event Log Monitoring," produced by the NSA/CSS.

## Further Extending Endpoint Visibility – OSQuery



osquery (by Facebook) allows you to easily ask questions about your Linux, Windows, and macOS infrastructure. It is used by organizations for a wide variety of use cases: Intrusion detection, threat hunting, operational monitoring ...

osquery gives you the ability to query and log things like running processes, logged in users, password changes, USB devices, firewall exceptions, listening ports, and more. The query language is “generic” across different OS environments!

It supports ad-hoc queries, but querying can also be scheduled. As an optional feature, it also allows you to perform file integrity monitoring. Several add-on solutions (e.g. for fleet management) have been developed.

### Further Extending Endpoint Visibility – OSQuery


osquery (by Facebook) allows you to easily ask questions about your Linux, Windows, and macOS infrastructure. It is used by organizations for a wide variety of use cases: Intrusion detection, threat hunting, operational monitoring ...

osquery gives you the ability to query and log things like running processes, logged in users, password changes, USB devices, firewall exceptions, listening ports, and more.

It supports ad-hoc queries, but querying can also be scheduled. As optional features, it also allows you, for example, to perform file integrity monitoring.



## Further Extending Endpoint Visibility – OSQuery – Example Queries



```
osquery> SELECT uid, name FROM listening_ports l, processes p WHERE l.pid=p.pid;
```

Some examples queries that could be useful for security purposes are listed below:

**SELECT \* FROM processes WHERE on\_disk=0**

⇒ Detect running processes that do not have an executable stored on disk

**SELECT name FROM processes WHERE pid=(SELECT parent FROM processes WHERE parent!=(SELECT pid from processes where name='services.exe') AND LOWER(name)='svchost.exe');**

⇒ Detect a "svchost.exe" process that doesn't have services.exe as a parent process

Interesting set of queries mapped to ATT&CK: <https://github.com/teoseller/osquery-attck>

### Further Extending Endpoint Visibility – OSQuery – Example Queries

For those familiar with SQL (Structured Query Language), OSQuery will prove to be easy to learn. OSQuery relies on a table structure that allows analysts to easily draft and execute queries. Here are a few examples of queries that could be interesting from a security perspective:

*SELECT \* FROM processes WHERE on\_disk=0*

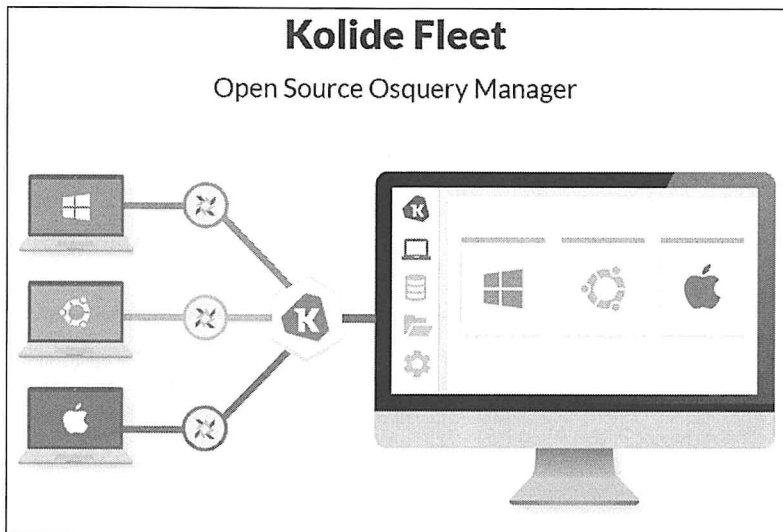
This query will retrieve all running processes that do not have an image ("executable") on disk.

*SELECT name FROM processes WHERE pid=(SELECT parent FROM processes WHERE parent!=(SELECT pid from processes where name='services.exe') AND LOWER(name)='svchost.exe');*

Detect a "svchost.exe" process that doesn't have services.exe as a parent process

Palantir provides an interesting pack of queries available on GitHub: <https://github.com/palantir/osquery-configuration>. For a full overview of all tables available in OSQuery, please refer to <https://osquery.io/schema/3.3.2>. Furthermore, an interesting set of queries mapped to ATT&CK is available on GitHub here: <https://github.com/teoseller/osquery-attck>.

## Further Extending Endpoint Visibility – OSQuery – Kolide Fleet



By itself, OSQuery does not provide a centralized management interface for enterprise use

Kolide Fleet is an interesting "Fleet Manager" for OSQuery, providing a centralized console with extensive management capabilities

Apart from Kolide Fleet (which we will use in an upcoming lab), several other solutions exist!

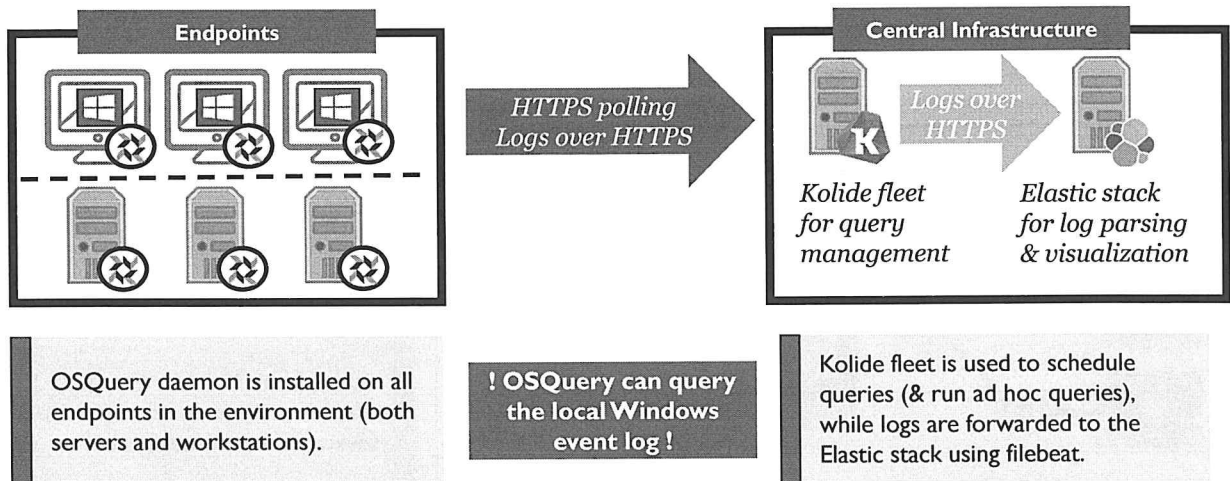
### Further Extending Endpoint Visibility – OSQuery – Kolide Fleet

For those of you already familiar with OSQuery, you may know that the stand-alone / executable is far from an enterprise solution that can easily be managed centrally. However, there is help!

Kolide Fleet is an interesting "Fleet Manager" for OSQuery, providing a centralized console with extensive management capabilities. The software is provided free of charge and is available on GitHub. You will get hands-on with Kolide Fleet during the upcoming lab and even more during the rest of the week!

Apart from Kolide Fleet, several other solutions exist, although Kolide Fleet is one of the most known and popular ones.

## Further Extending Endpoint Visibility – OSQuery – Kolide Fleet Architecture



### Further Extending Endpoint Visibility – OSQuery – Kolide Fleet Architecture

So, what does the overall architecture look like? We can use the following architecture as an example:

- OSQuery is installed on all endpoints. It's configured to use the Kolide Fleet server API's for configuration and query scheduling. Polling happens using predetermined intervals;
- The results of the queries are forwarded over the same HTTPS link to the Kolide Fleet engine, from where it can be forwarded to an Elastic stack.

Although this is a local agent install, we'd like to indicate that OSQuery is a lightweight agent that cannot be compared to more intrusive tooling such as AV engines or EDR tools: OSQuery only runs query on-demand (or when scheduled). Of course, similar architectures exist for commercial environments, but we believe this to be an excellent baseline using only free tools.

## Further Extending Endpoint Visibility – Endpoint Detection & Response (EDR)

### EDR

Endpoint Detection & Response (EDR) tools can further extend your endpoint visibility, while also adding a capability to respond to detected incidents. They are deployed using an agent on the system and provide a central console for monitoring, management and response. Typical tools include Windows Defender ATP, Carbon Black, CrowdStrike Falcon, Endgame, SentinelOne,...

Regardless of the chosen product, endpoint protection suites are on the rise and will likely become an essential element of our toolkit. Some of the EDR capabilities can be picked up by separate (free) tools (e.g. monitoring with Sysmon, hunting with OSQuery,...), but modern EDR solutions will aim to consolidate all functions ("next-gen" AV, increased detection, hunting, response,...) in **one agent**.

So what EDR tool is best? An interesting comparison of different EDR tools and their capabilities is available at [http://www.hexacorn.com/edr/IR\\_EndPointSolutions.xlsx](http://www.hexacorn.com/edr/IR_EndPointSolutions.xlsx) (non-commercial overview)

### Further Extending Endpoint Visibility – Endpoint Detection & Response (EDR)

Endpoint Detection & Response (EDR) tools can further extend your endpoint visibility, while also adding a capability to respond to detected incidents. They are deployed using an agent on the system and provide a central console for monitoring, management and response. Typical tools include Windows Defender ATP, Carbon Black, CrowdStrike Falcon, Endgame, SentinelOne,...

Regardless of the chosen product, endpoint protection suites are on the rise and will likely become an essential element of our toolkit. Some of the EDR capabilities can be picked up by separate (free) tools (e.g. monitoring with Sysmon, hunting with OSQuery,...), but modern EDR solutions will aim to consolidate all functions ("next-gen" AV, increased detection, hunting, response,...) in **one agent**.

So what EDR tool is best? It's an interesting question that will depend on your environment and exact needs. An interesting comparison of different EDR tools and their capabilities is available at [http://www.hexacorn.com/edr/IR\\_EndPointSolutions.xlsx](http://www.hexacorn.com/edr/IR_EndPointSolutions.xlsx) (non-commercial overview)

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.1

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim


### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

**Exercise: Kibana, ATT&CK Navigator, and FlightSim**



**Please refer to the workbook for further instructions on the exercise!**

This page intentionally left blank.

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.1

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim

### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.



## Understanding & Limiting Your Internet Footprint

So, what are you and your employees sharing with the outside world?



### INFORMATION

Employee contact information

Open job positions (and relevant technology)

Commercial documentation

### TECHNOLOGY

Publicly available services you host yourself

Services accessible to partners, customers, vendors...

Publicly available services hosted in the cloud

**During reconnaissance, adversaries will mine this information!**

### Understanding & Limiting Your Internet Footprint

An "internet footprint" can have more than one interpretation.

All the information available on the internet about your company or organization defines the "information" footprint. It should be obvious that more information available means a bigger footprint. Information about your company or organization, available on the internet, is not always under your control. You have your company's resources sharing information (like your websites), or even exposing information that you would rather not. And you have third parties sharing information. For example, social media, where employees, clients, providers... share information.

All your resources facing the internet define the "technical" internet footprint of your company or organization. These resources are web servers serving your websites, file servers, VPN concentrators, email servers, DNS servers... An individual service can have a small or a large footprint, depending on the type and amount of resources it makes available on the internet. For example, the number of ports, the number of web pages, the number of files...

Your internet footprint is discussed in this course because it defines your exposure. Simply put, your footprint is a measure of your exposure. Another term you might be more familiar with is the "attack surface." The bigger your footprint is, the bigger your exposure and the bigger your attack surface.





## Understanding & Limiting Your Internet Footprint – Information

The information footprint defines what information on your company is available on the internet

- This will include information you (and your employees) control, but also information out of your control (e.g. that has been shared or disclosed by third parties)

### Information you control

- Security awareness
- Data classification
- DLP solutions (?)
- Monitor

### Information you do not control

- Monitor
- Respond

### Understanding & Limiting Your Internet Footprint – Information

The information footprint defines what information on your company is available on the internet. To do business, your company has to share information. Most of that information is publicly available on the internet, but not everything!

For our purposes, we can distinguish two main types of online information: Information you (and your employees) control, but also information out of your control (e.g. that has been shared or disclosed by third parties).

For information you control, you can focus on a number of solutions, including:

- Increasing security awareness of your staff: They should be made aware and trained about information classification guidelines and publishing policies. Knowledge about scammers should be shared with your staff.
- Data classification: It's difficult for staff to understand what they can share if they don't know how data is classified. Ensure all data in the organization is classified and clear rules exist on what these classifications mean. Top Secret information is most likely not intended to be shared on social media.
- You could consider implementing DLP solutions that attempt to stop classified information from leaving the internet.
- Monitor the internet to see what type of information is available on your corporate web page, social media accounts, partner websites.

For information you do not control, there's not much we can do except for monitoring the internet and responding to information that is exposed.



## An Example of Information You Don't Control: Breached Credentials!

Pa\*\*

Throughout the past couple of years, a highly worrying trend has been the increase of data breaches that include credentials. Although you may not be the victim of the data breach, some of your employees could be. These breaches could still impact you (e.g. if the employee reuses a compromised password in your environment)!

- Free online services that allow you to check your exposure include HaveIbeenpwned?, Ghostproject... These can be used on an individual basis by employees. Note that some of these offer the passwords in cleartext (!).
- Commercial organizations that offer this type of monitoring for companies are on the rise. Some of these include additional services such as a link with your Active Directory to force password resets when a password is discovered in a breach online!

### An Example of Information You Don't Control: Breached Credentials!

Throughout the past couple of years, a highly worrying trend has been the increase of data breaches that include credentials. Although you may not be the victim of the data breach, some of your employees could be. These breaches could still impact you (e.g. if the employee reuses a compromised password in your environment) and are thus worth monitoring.

In order to monitor your exposure to such leaked data, there's a few options available:

- Free online services that allow you to check your exposure include HaveIbeenpwned? and Ghostproject... These can be used on an individual basis by employees. Note that some of these offer the passwords in cleartext (!). Please see the reference section of this slide for their links.
- Commercial organizations that offer this type of monitoring for companies are on the rise. Some of these include additional services such as a link with your Active Directory to force password resets when a password is discovered in a breach online! Examples include SpyCloud, auth0...

#### References:

<https://haveibeenpwned.com/>

<https://ghostproject.fr>



## Understanding & Limiting Your Internet Footprint – Technical (I)

The technical footprint defines what resources are exposed to the internet

- Most organizations understand this and have started strongly limiting what they expose online.
- Perimeter vulnerability scanning is periodically done by the vast majority for organizations.
- Don't forget those test / development systems.

But what about those AWS systems the marketing team set up last month?  
As security professionals, we have to limit the "Shadow IT" in our organization!

### Understanding & Limiting Your Internet Footprint – Technical (1)

The technical footprint of your company on the internet encompasses all devices connected to the internet. Servers, routers, network devices... all connected to the internet to make your company reachable to customers and potential customers over the internet, are also reachable to adversaries. An open network port on your server is not in itself a risk, but the server application that opened the port to listen to incoming connections can be a problem. Adversaries will connect to the server application via that open port and interact with it to look for vulnerabilities. Vulnerabilities are bugs in the application, or misconfigurations, that allow adversaries to take over control of the application, for example with remote code execution.

An inventory of all internet-facing resources will help you to understand your company's technical footprint. Your company probably already has a list of all its assets used in production and how they are connected to the internet. But this list must also include the services running on those assets: Server applications, open ports, protocols supported... Test, development and staging servers should also be included in this list. It won't be the first time nor the last time that a company has its IT infrastructure compromised because of an internet-facing development server that was "forgotten" and not protected like production servers.

Producing this list is not just an "accounting" exercise: Besides compiling a list by collecting information from the different IT teams, it is important to also approach this problem from a practical point. Your opponents will scan all your IP addresses to map your internet footprint; this is something that is good to do, too. It will allow you to discover services that escaped control from your IT teams. These scanning exercises should be conducted on a regular basis, as your internet footprint is dynamic. It changes over time. Exercises like these can be performed by staff or contracted to third parties specialized in scanning services.

The attack surface is all the resources and services you expose to the internet and that adversaries can use to try to enter your systems and compromise your operations.

An increasing problem here is that cloud-based services allow virtually anyone to easily launch / host online services. We have seen many organizations where different departments are setting up IT systems in the cloud outside of the IT department's control. This is a serious risk we call "Shadow IT"; no one controls these systems:

- Are they properly configured / hardened?
- What data is stored on these platforms?
- Are they patched?
- Is there logging and monitoring in place?

As security professionals, we have to limit this as much as possible!



## Assessing Your Own Footprint

We will shortly introduce a few ways on how your adversaries can easily assess your external footprint:



Using public search engines look for interesting information (e.g. using search directives)



Scan the external IP address range that is assigned to your organization (this would not include third-party hosting)



Monitor social media to understand what type of information is "trending" about your organization

These techniques are not "exclusive" to the attacker: As defenders, we can analyze our own footprint by regularly performing the same assessments

### Assessing Your Own Footprint

As every service exposed to the internet will be attacked, it stands to reason to disable unnecessary services. All internet-exposed devices are permanently scanned by countless scanners under control of criminals and computers infected with malware. Disabling unnecessary services is not just closing ports, but also configuring services to limit the features they offer that are required for the operation of your company's internet presence. For example, if you need a file server just to enable clients to upload documents, don't enable all features of the file server. Only enable uploading of files; don't enable listing of files or downloading of files. If the file server application contains unknown vulnerabilities in its file listing functions, for example, then these vulnerabilities cannot be exploited because the feature cannot be accessed from the internet.

All software contains bugs, and many bugs lead to vulnerabilities that can be exploited. Software vendors and open-source projects that maintain their projects will fix bugs they discover or are reported to them. Keeping your software up-to-date makes sure that known vulnerabilities are removed. Applications are not maintained forever. Besides patching, you need to keep up with major releases because old software that is end-of-life is no longer maintained.

Besides disabling features and services, you can also protect from the internet by filtering the network traffic directed to them with firewalls, web application firewalls, intrusion prevention services, ...



## Assessing Your Own Footprint – Google Search Operators

Search engines such as Google have an amazing index of what is published on the internet

- We can leverage this by using Google's search operators; some examples:

Operator	Comment	Operator	Comment
site:	Find results on a given domain	filetype:	Find specific file extensions
link:	Find links to a certain domain	location:	Find by physical location
inurl:	Find results with this in the URL	+	Include keywords from results
related:	Find related information	-	Exclude keywords from results
daterange:	Find within specific date range	AND / OR	Combine operators

An interesting overview of useful Google search operators is the Google Hacking Database (GHDB) at <https://www.exploit-db.com/google-hacking-database/>

### Assessing Your Own Footprint – Google Search Operators

Another method to scope your internet footprint is to look up what information about your internet-facing devices other actors have collected. There are numerous indexing services on the internet that constantly spider the internet to index information. Well-known ones like Google and Bing index the content of web servers.

This is a tool that is also available to you. Why wouldn't you spend some time to use Google and see what information it has gathered about your company or organization? You can refine your searches with "search operators." We have listed a sample of search operators above, but there's a few more.

- site: Find results on a given domain.
- link: Find links to a certain domain.
- inurl: Find results with this in the URL
- related: Find related information.
- daterange: Find within specific date range.
- filetype: Find by specific file extensions.
- location: Find by physical location.
- + Include keywords from results.
- Exclude keywords from results.
- AND / OR Combine operators

An interesting overview of useful Google search directives is the Google Hacking Database (GHDB) at <https://www.exploit-db.com/google-hacking-database/>



As an enterprise, we can implement a number of useful open-source and free tools to monitor Pastebin for a specific string or regular expression:

- <https://github.com/leapsecurity/Pastepwnd> (Pastepwnd by LeapSecurity)
- <https://github.com/cvandeplas/pystemon> (Pystemon by Christophe Vandepas)
- <https://github.com/xme/pastemon> (Pastemon by Xavier Mertens)
- <https://github.com/CIRCL/AIL-framework> (CIRCL AIL - Analysis of Information Leaks)



PASTEBIN

```
Executable File | 985 lines (883 sloc) | 38.9 KB
1  #!/usr/bin/env python
2  # encoding: utf-8
3
4  ***
5  @author:  Christophe Vandepas <christophe@vandepas.com>
6  @copyright:  AGPLv3
7  http://www.gnu.org/licenses/agpl.html
```

### Assessing Your Own Footprint – Automating Pastebin Monitoring

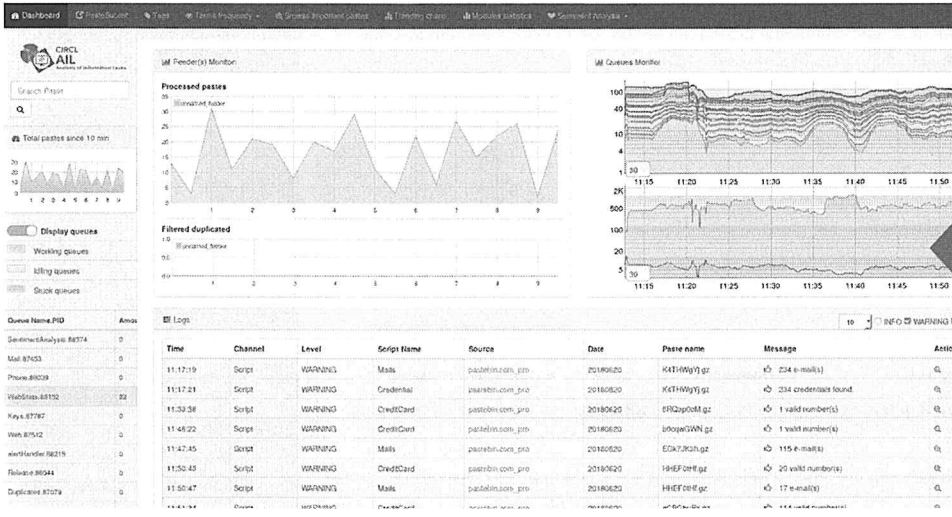
As an enterprise, we are looking for solutions that are automated and do not require too much manual user interaction. It would be a bit overkill to have someone check Pastebin manually on a periodic basis. Several security experts have created some useful scripts where you can define regular expressions or strings that need to be matched on Pastebin:

- Pastepwnd by LeapSecurity (available at <https://github.com/leapsecurity/Pastepwnd>)
- Pystemon by Christophe Vandepas (available at <https://github.com/cvandeplas/pystemon>)
- Pastemon by Xavier Mertens (available at <https://github.com/xme/pastemon>)
- CIRCL AIL-framework (available at <https://github.com/CIRCL/AIL-framework>)

Depending on whether you have a Pastebin Pro account, these scripts can scrape the website or directly query the Pastebin API.



## Assessing Your Own Footprint – CIRCL AIL Framework



The AIL (Analysis for Information Leaks) framework was built by CIRCL (it's one of their many great projects and initiatives) and aims to be a central framework that organizations can leverage to identify possible information leaks. It was built to allow easy, centralized, follow-up and further extension!

### Assessing Your Own Footprint – CIRCL AIL Framework

The AIL (Analysis for Information Leaks) framework was built by CIRCL (it's one of their many great projects and initiatives) and aims to be a central framework that organizations can leverage to identify possible information leaks. It was built to allow easy, centralized, follow-up and further extension through the addition of modules! The central web application can be self-hosted (the code is available on GitHub).

From CIRCL's GitHub page:

*"AIL is a modular framework to analyze potential information leaks from unstructured data sources like pastes from Pastebin or similar services or unstructured data streams. AIL framework is flexible and can be extended to support other functionalities to mine or process sensitive information (e.g. data leak prevention).*


### Reference

<https://github.com/CIRCL/AIL-framework>





## Assessing Your Own Footprint – SpiderFoot



SpiderFoot is an open-source reconnaissance project that aims to visualize reconnaissance activities. Due to its open-source nature, it's highly adaptable and can be customized to include additional modules!

Built packages are available in Linux and Windows!

### Assessing Your Own Footprint – SpiderFoot

SpiderFoot is a tool aimed at "facilitating" a lot of the previously mentioned steps. It only requires a domain name as input and will start "spidering" from there to detect any possibly relevant data. From its official web page:

*"SpiderFoot is a reconnaissance tool that automatically queries over 100 public data sources (OSINT) to gather intelligence on IP addresses, domain names, email addresses, names and more. You simply specify the target you want to investigate, pick which modules to enable and then SpiderFoot will collect data to build up an understanding of all the entities and how they relate to each other."*

*So what is OSINT? OSINT (Open Source Intelligence) is data available in the public domain which might reveal interesting information about your target. This includes DNS, Whois, Web pages, passive DNS, spam blacklists, file meta data, threat intelligence lists as well as services like SHODAN, HaveIBeenPwned? and more. [Click here to see the full list of data sources SpiderFoot utilises.](#)"*

Built packages are available for both Linux and Windows! We will use SpiderFoot in an upcoming lab!



## Assessing Your Own Footprint – Masscan & Scans.io

Initiatives such as Masscan & Scans.io are the reason why you should never leave any system that is unpatched connected to the internet

Masscan is an open-source internet scanning tool; it can scan for a specific service on the entire IPv4 public address range in under 5 minutes (will depend on available bandwidth)

Results of these internet-wide scanning activities are made available by websites such as Scans.io (download full datasets) & Censys (censys.io – query scan results using API)

Using similar tools, adversaries are scanning your perimeter on a continuous basis!

### Assessing Your Own Footprint – Masscan & Scans.io

Initiatives such as Masscan & Scans.io are the reason why you should never leave any system that is unpatched connected to the internet.

Masscan is an open-source internet scanning tool; it can scan for a specific service on the entire IPv4 public address range in under 5 minutes (will depend on available bandwidth). It is commonly used to scan the internet to understand the impact of newly identified vulnerabilities (e.g. what systems on the internet are running this service). Although it appears to offer the same functionality as port scanning tools like NMAP, it uses a highly optimized scanning algorithm and its own custom TCP/IP stack. You could download and install Masscan (or even NMAP) to scan your own external IP ranges.

Instead of scanning your perimeter yourself, you could also leverage public websites that expose this type of information: Results of these internet-wide scanning activities are made available by websites such as Scans.io (download full datasets) & Censys (censys.io – query scan results using API).

Using similar tools, adversaries are scanning your perimeter on a continuous basis!



## Assessing Your Own Footprint – Shodan (1)

Another interesting example of an online scanner is Shodan



### Assessing Your Own Footprint – Shodan (1)

Besides well-known search engines like Google and Bing, there are specialized search engines that index internet-facing resources like webcams, routers, ICS devices... Shodan (<https://www.shodan.io/>) is the most popular scanner for these devices.

While classic indexing services like Google and Bing will index the content of web servers, Shodan will capture information about the services (metadata). Shodan will scan ports of web servers, ssh servers, ftp servers, telnet servers... establish a connection, and index the metadata shared by the service. These are called "service banners", and typically announce the implementation and version of the service. This is, for example, a service banner returned by OpenSSH running on a server:

```
SSH-2.0-OpenSSH_5.3
```

By indexing this information, Shodan offers its users the capability to look for particular services on the internet. For example, it can be used to search for older, vulnerable version of OpenSSH present on the internet.

Shodan is a free service. It returns up to 10 results for searches performed without registration, up to 50 results for searches performed with a free account, and more with paying accounts. It also has a Windows GUI application: Shodan Diggity.



## Assessing Your Own Footprint – Shodan (2)

**SHODAN** country:BE port:5900

Exploits Maps Images Share Search Download Results Create Report

**TOTAL RESULTS**  
3,137

**TOP COUNTRIES**

Belgium	3,137
---------	-------

**TOP CITIES**

Brussels	259
Antwerpen	69
Gent	58
Brussel	45
Liege	34

**TOP ORGANIZATIONS**

Telenet N.V.	1,012
Skywet Belgium	758
Proximus Skywet	572
VOD	152
Universite Libre de Bru...	80

**81.242.120.69**  
IP: 103.174.217.100/19 - ip.belgium.be  
Skywet Belgium  
Added on: 2017-04-04 04:35:55 GMT  
Belgium, La Louviere  
Details

**212.166.5.188**  
IP: 173.101.100.119 - vnc.ssh.wire.be  
WM S.A.  
Added on: 2017-04-04 04:12:01 GMT  
Belgium, Verviers  
Details

**172.111.197.3**  
IP: 173.101.100.119 - vnc.ssh.wire.be  
Secure Internet LLC  
Added on: 2017-04-04 04:11:05 GMT  
Belgium, Brussels  
Details

**178.119.192.186**  
IP: 173.101.100.119 - vnc.ssh.wire.be  
Telenet N.V.  
Added on: 2017-04-04 03:42:31 GMT  
Belgium, Coxmeer  
Details

**212.166.44.117**  
IP: 103.174.217.100/19 - ip.belgium.be  
WM S.A.  
Details

### Shodan search directives

Once you sign up for an account, Shodan also supports search directives

In our example, we are looking for all systems listening on port 5900 (VNC) and hosted in Belgium

### Assessing Your Own Footprint – Shodan (2)

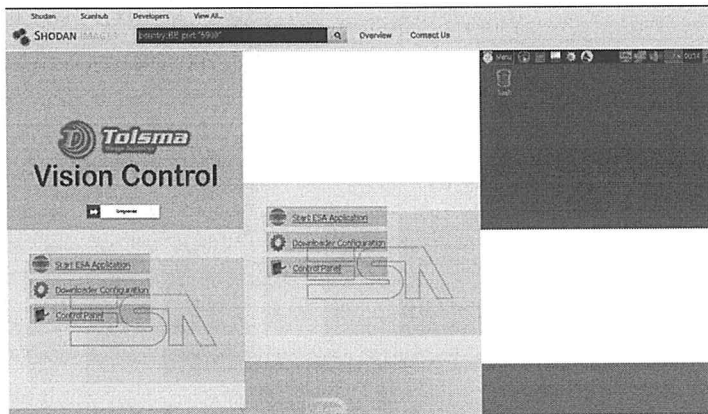
In the above example, we are running a Shodan search directive "country:BE port:5900", resulting in 3,137 results! We are looking for all systems listening on port 5900, which are hosted in Belgium. We can, of course, further tailor this to include specific IP ranges of your organization (or ISP).

VNC is an often insecurely configured remote administration protocol, so it's likely something we'd need to further investigate...



## Assessing Your Own Footprint – Shodan Image Search

It gets scarier. Using Shodan image search, we can see screen captures of what was running on identified network ports:



### Shodan image search

Shodan image search is only available for paid accounts, but provides some very interesting insights: We can actually visualize what is presented to users upon connecting to the exposed services

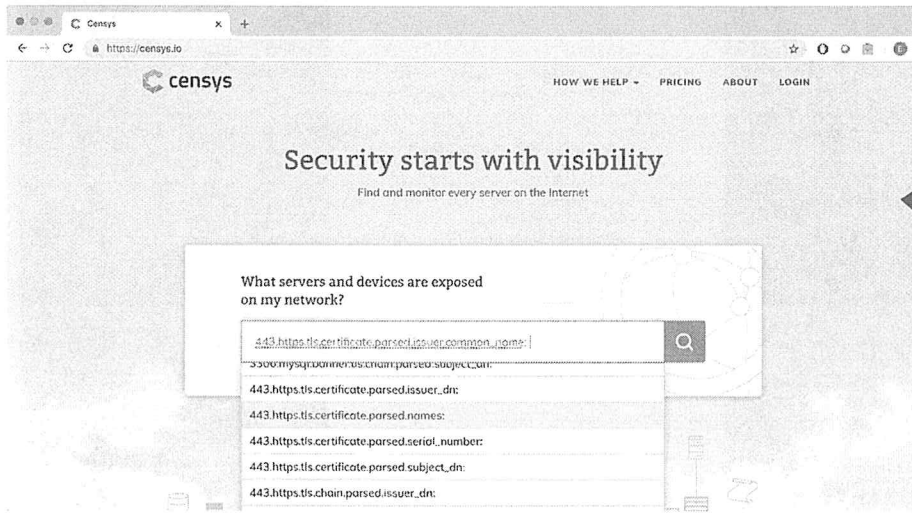
### Assessing Your Own Footprint – Shodan Image Search

It gets scarier. Using Shodan image search, we can see screen captures of what was running on identified network ports! Note that Shodan image search is only available for paid accounts but provides some very interesting insights: We can actually visualize what is presented to users upon connecting to the exposed services!

This is an excellent tool if we want to show people with less technical expertise what Shodan is doing and what you are exposing as an organization!



## Assessing Your Own Footprint – Censys.io



### Censys.io

Censys.io is similar to Shodan, although not the same. Censys has doubled the available filters that can be used to identify servers / hosts / services of interest.

In the attached screenshot, we could be looking to find services running HTTPS with an SSL/TLS certificate that matches a certain keyword in the "common name" field.

### Assessing Your Own Footprint – Censys.io

How about systems that are not running in your own network environment (and are thus not covered by your ASN)? How do you keep the visibility on these? Censys.io aims to provide this type of visibility. Censys.io is similar to Shodan, although not the same. They have doubled the available filters that can be used to identify servers / hosts / services of interest.

In the attached screenshot, we could be looking to find services running HTTPS with a SSL/TLS certificate that matches a certain keyword in the "common name" field. Please note that the web interface is freely available for ad hoc lookups, but they offer a commercial API access to obtain this visibility on a continuous basis.



# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.I

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim

### Reconnaissance


Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.



## Exercise: Automated Reconnaissance using SpiderFoot



Please refer to the workbook for further instructions on the exercise!

This page intentionally left blank.

# Course Roadmap

- **Day 1: Introduction & Reconnaissance**
- Day 2: Payload Delivery & Execution
- Day 3: Exploitation, Persistence and Command & Control
- Day 4: Lateral Movement
- Day 5: Action on Objectives, Threat Hunting & Incident Response
- Day 6: APT Defender Capstone

## SEC599.I

### Course Outline and Lab Setup

Course objectives and lab environment

What's happening out there?

Introducing SYNCTECHLABS

Exercise: One click is all it takes...

### Adversary emulation and purple team

Introducing the extended kill chain

What is the Purple Team?

MITRE ATT&CK framework and "purple tools"

Key controls for prevention and detection

Exercise: Hardening our domain using SCT and STIG

Building a detection stack

Exercise: Kibana, ATT&CK Navigator, and FlightSim

### Reconnaissance

Reconnaissance – Getting to know the target

Exercise: Automated reconnaissance using SpiderFoot

This page intentionally left blank.

## Conclusions for 599.1

That concludes 599.1! Throughout this section, we've touched upon the following topics:

- Explain the concept of Purple Teaming and how it can benefit your organization
- Explain the cyber threat landscape and what adversaries are doing in the wild
- Explain how we can structure these attacks according to the Kill Chain / APT attack cycle
- Deep-dive case studies on recent advanced attacks
- An offensive exercise to get you familiar with how the adversary operates
- Explained the need for a baseline security architecture and monitoring capability
- Started assessing our own environment by doing vulnerability scanning

In the next section of the course (SEC599.2), we will start investigating techniques to prevent initial delivery and execution of payloads!

## Conclusions for 599.1

So, that concludes the first day of SEC599 (599.1)! Throughout this section, we've attempted to illustrate both how you work yourself, but also how your adversaries operate. More specifically, we've touched upon the following topics:

- Explain the concept of Purple Teaming and how it can benefit your organization
- Explain the cyber threat landscape and what adversaries are doing in the wild
- Explain how we can structure these attacks according to the Kill Chain / APT attack cycle
- Deep-dive case studies on recent advanced attacks
- An offensive exercise to get you familiar with how the adversary operates
- Explained the need for a baseline security architecture and monitoring capability
- Started assessing our own environment by doing vulnerability scanning

In the next section of the course (SEC599.2), we will start investigating techniques to prevent initial delivery and execution of payloads!

## Course Resources and Contact Information



### **AUTHOR CONTACT**

Erik Van Buggenhout  
evanbuggenhout@nviso.be  
Stephen Sims  
ssims@sans.org



### **SANS INSTITUTE**

11200 Rockville Pike  
Suite 200  
North Bethesda, MD 20852  
301.654.SANS (7267)



### **CYBER DEFENSE CONTACT**

Stephen Sims  
ssims@sans.org



### **SANS EMAIL**

GENERAL INQUIRIES: [info@sans.org](mailto:info@sans.org)  
REGISTRATION: [registration@sans.org](mailto:registration@sans.org)  
TUITION: [tuition@sans.org](mailto:tuition@sans.org)  
PRESS/PR: [press@sans.org](mailto:press@sans.org)

This page intentionally left blank.

