Detection of Advanced Persistent Threats Using System and Attack Intelligence

Alberto Redondo-Hernández, Aitor Couce-Vieira, Siv Hilde Houmb

Secure-NOK AS
Hamar, Norway
Email: {albertoredondo,aitorcouce,sivhoumb}@securenok.com

Abstract—**Cyber attacks have evolved from being mostly harmless to sophisticated and devastating Advanced Persistent Threats (APT), such as the Stuxnet or Aurora attacks. APTs have the capabilities to stop business operations and cause physical damage to plants and equipment. This is a serious threat to Industrial Control Systems common in critical infrastructures such as pipelines, refineries, electrical grids or nuclear plants. This paper discusses why existing cyber attack detection technologies and solutions are not able to detect APTs, and makes use of a flawed detection paradigm based on prior knowledge of attacks. This paper also introduces a novel approach to detect APTs that is based on deep monitoring over large time intervals combined with correlation and analysis of monitored events over these time periods to detect indications of a cyber attack. The paper also provides an example of using the proposed approach to detect Stuxnet.**

Keywords–Malware; APT; Monitoring System; Intrusion Detection Systems; Intrusion Prevention Systems; Cybersecurity.

I. Introduction

The computerization of industrial environments has introduced new cybersecurity problems [1]. Cybersecurity breaches, espionage, insiders, and threats to privacy continue to increase in frequency, impact and sophistication [2]. Indeed, their impact on the global economy has been estimated at more than $400 billion in annual cost, or around 0.8% of the global Gross Domestic Product (GDP) (in comparison, drug trade represents 0.9%, and international crime, 1.2%) [3]. World leaders are raising their concerns on cyber attacks and the serious menace they pose to critical infrastructure and intellectual property [4]. Governments, in coordination with the industry, are developing strategies and guidelines to improve critical infrastructure cybersecurity and prevent the increasing social and economic impact of attacks.

The challenge is that today's Industrial Control Systems (ICS) and critical infrastructure rely on outdated security models and invalid assumptions. At the same time, the frequency and sophistication of cyber attacks against ICS are increasing and these critical assets are becoming prime targets both by criminal and terrorist organizations. These sophisticated attacks are difficult to detect and they operate covertly; they typically start with seemingly benign activities that do not trigger any warning, as was the case with the Stuxnet [5] and Aurora [6].

These attacks are called Advanced Persistent Threats (APT) [7], characterized as attacks that remain unnoticed until the consequences become visible in the system or its environment. APTs cannot be detected using conventional security tools and represent a significant challenge and risk to industrial environments and critical infrastructure.

The remainder of the paper provides an evaluation of current detection paradigms and proposes a new paradigm based on event analysis. Section II provides an overview of APTs and their phases. Section III discusses why existing detection solutions are not designed to detect APTs and use a flawed paradigm based on prior knowledge of attacks. Section IV introduces a novel detection approach tailored to the nature of APTs and based on (1) deep monitoring over large time intervals combined with (2) the analysis of monitored events over such periods to detect indications of a cyber attack. Section IV provides an example of the proposed approach applied to detect the early phases of Stuxnet.

II. Advanced Persistent Threats

APTs [7] works in the background conducting espionage or sabotage actions that could result in considerable monetary, environmental or safety loses. The steps taken by these threats usually go unnoticed until they have reached their goal or have penetrated large parts of the infected systems, making their removal costly and difficult. Table I outlines a selection of APTs attacks happened in the last decade [8][9]:

TABLE I. Important attacks within the last six years.

Attack	*Entry Method*	*Date*	*Classification*
Aurora operation	Malware	2007	Espionage
Stuxnet	Malware	2009	Sabotage
Energetic Bear	Malware	2011	Espionage
Flame	Malware	2012	Espionage
Shamoon	Malware	2012	Sabotage
Heartbleed	Malware	2014	Espionage

A. APT Characteristics

There has been significant research and analysis of APTs and experts have defined their main features [10][7][11] as follows:

Targeted: APTs target organizations with the purpose of stealing specific data or causing damage.

Persistent: APTs play out varied phases over a long period of time. To steal data, the attacker must identify vulnerabilities, evaluate existing security controls, gain access to privileged hosts within the target network, find the target data and, finally, ex-filtrate or manipulate them.

Evasive: APTs are designed to evade traditional security products gaining, for instance, privilege access in hosts within the target network while avoiding firewalls, antivirus and other security protective mechanisms.

Complex: APTs apply a complex mix of attack methods adapted to the multiple vulnerabilities that the attacker identifies in the targeted system.

B. The APT Process

The APT process includes three dominant phases which may take place over a period of several months [12].

1) Phase 1: The attacker performs reconnaissance, identifies vulnerabilities, launches the attack and infects target hosts.

- **Recognize:** Attackers look for entry points, vulnerabilities, key individuals, and key assets.
- **Launch:** Common methods to gain access to a privileged host may include email traps with hidden malware, malicious websites aimed at extracting passwords, or social engineering to get access to specific accounts.
- **Infect:** Code is installed into a targeted host and the malware reports back to a Command and Control (C&C) fueling the attack.

2) Phase 2: The attacker controls infected hosts, updates the malware, spread it to other machines, and collects data.

- **Control:** The attacker remotely controls infected hosts with a C&C service on the Internet, often on a dynamic Domain Name System host. C&C provides the attacker with remote control.
- **Discover:** The infected host downloads additional components to identify target data on the infected hosts, on mapped network drivers, and at other network locations.
- **Persist:** An important difference between traditional malware and an APT is the ability to persist. Traditional malware will often remove itself or be removed by an antivirus program. However APT operations are designed to go on in silence and persist by downloading new code to avoid being detected.

3) Phase 3: Once the attackers have taken control of one or more hosts within the target network, they may establish access credentials to expand their reach. In case of exfiltration, the attacker send the data out of the network through the C&C server or a previously unused server. At this point, the consequences can result in the public disclosure or selling of sensitive information, blackmail, or the share of the attack methods to other attackers that may repeat the attack. In addition, if the purpose is sabotage, the attackers may manipulate the data of specific targets in order to alter the normal operations and processes that the attacked system supports.

C. The Challenge of Detecting APTs

APTs are methodical, adaptive and efficiently covering tracks while carefully penetrating the network, ceasing the attack or staying ”under the radar” for days to avoid raising suspicion, gaining knowledge of the system, or taking advantage of zero-day vulnerabilities [13] in the underlying operating system that cannot be defined through patterns or signatures.

This adaptive behaviour is hard to detect. Standard antivirus systems are not able to detect these attacks and perimeter defences; even the most sophisticated ones, are frequently breached [14]. Therefore, monitoring and detection mechanism need to implement a meticulous surveillance strategy focused on tracking the footprints of cyber attacks to be able to detect data-thefts and other loses using both system and attack intelligence.

III. Related work and Existing Cybersecurity Solutions for Industrial Control Systems

There exist multiple solutions and techniques aimed at detecting sophisticated cyber attacks, such as Tofino and Industrial Defender.

A. Tofino

Tofino is a device which provides attack detection using Deep Package Inspection (DPI) [15], with a simple installation and rugged hardware design. Whilst this method is efficient for many purposes, it is not efficient in preventing against APT because these attacks are able to change their behaviour according to the purpose of the attack. For example, an APT may, on the fly, adapt itself to a form that is not detectable as an attack using DPI.

1) Strengths: The Tofino configuration is a relatively **simple and straightforward** network monitoring solution that can be configured in various manners using the Tofino Configurator software. Tofino also includes an unique **Test Mode** that allows firewall testing with no risk to current operations, as well as being **compatible** with all Distributed Control Systems (DCS), Programmable Logic Controllers (PLC), Supervisory Control And Data Acquisition (SCADA systems), networking and software products.

2) Weaknesses: Tofino is a perimeter defence and, as we discussed in the previous section, their main problem is that they are breachable. Another issue of Tofino is that it is hardware-based and needs to be installed in a separate module. As some ICS are setup in a cabinet or a limited space type of location it might be infeasible to add additional hardware components to the network.

B. Industrial Defender

Industrial Defender [16] is a platform-agnostic monitoring system for ICS. It is employed to monitor security events, check configurations, collect data, and identify and protect the system. However, Industrial Defender is not very effective in detecting APTs, as this solution monitors the Industrial Components behaviour and not the attacker behaviour. APT attacks do not need to change the component behaviour to perpetuate the attack and therefore are not detected.

1) Strengths: Industrial Defender is comprised of a group of applications depending on the specific objective. One of its biggest benefits that it has a **single** and **unified** view of all assets within the automation systems environment, **actionable security intelligence** and that the system is able to **automate** tedious manual change management processes.

2) Weaknesses: An aspect to take into account is that Industrial Defender does not use DPI, as it obtains information from the control system itself. In addition, Industrial Defender does not provide a method to detect smart attacks, as it focuses solely on the behaviour of the controllers and does not take attack intelligence into consideration.

C. *Other Cybersecurity Solutions*

- **Darktrace Cyber Intelligence Platform (DCIP):** Evans [17] designed a cyberdefence system based on Bayesian methods with self-learning capabilities tracking evolving patterns of operations and behaviour.
- **Wurldtech Technology & Professional Services:** Ferris and Gilthorpe [18] designed a system to discover operational vulnerabilities in distinct products assessing the root cause.
- **Websense Security Labs:** McCormack [19] has designed the **TRITON** architecture, which is a set of shared security analytics, deployment platforms and management services aimed at identifying infected hosts and data extrusion attempts or prevent infection in APT phase 1.

IV. How to Detect APTs Using System and Attack Intelligence

A. *Proposed method*

Traditional antivirus products have been proven to be ineffective mitigating APT attacks due to the evasive nature of these threats, as discussed in [20]. The same is the case with other types of attack detection technologies and solutions, as discussed in Section III. Therefore, it is necessary to develop a more sophisticated approach tailored to monitor the behaviour of the system and correlate it with system and attack intelligence. For this purpose, we have designed an approach comprised of the following steps:

1) Monitor relevant events: We monitor events in the hosts that may be related with APTs. The monitoring system must be able to record information on various events in the hosts and network. These events represent the possible movements of the attacker and are used to build the attack patterns, including various steps of the attack. These include: insertion or removal of Universal Serial Bus (USB) devices, activation or deactivation of processes and critical processes, activation or deactivation of firewalls and antivirus, and increase or decrease of the usage of the Central Processing Unit (CPU) or the Random Access Memory (RAM).

2) Check behaviour patterns of different attacks: The second step is to check for behaviour patterns of different kind of attacks such as Denial of Service (DoS) or spyware. An example attack pattern for malware infection could be comprised of the steps illustrated in Figure 1: an attacker inserts an USB, then several processes are activated on the host and finally, the attacker removes the USB device. To separate between false positives and actual APTs, the monitoring system needs to continuously analyse against common patterns of different attacks and adopt a pessimistic algorithm for issuing alarms.

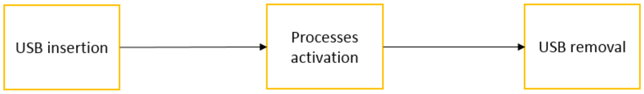

Figure 1. Malware infection pattern.

3) Raise an alarm in case of patterns detected: The monitoring system raises an alarm whenever it detects an attack pattern. This could be done by using alarm notification and by changing for example the alarm colour in a visualization module from green to orange for the affected monitored host as shown in Figure 2. The purpose of this is to display the variation of the risk level on the monitored host and separate between what might be an attack and a false positive.

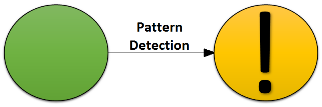

Figure 2. Raising an alarm in monitored host.

B. *Example: Stuxnet attack*

Stuxnet is a malware that was discovered in June 2010. It was tailored to attack Programmable Logic Controllers (PLCs) in a nuclear facility in Iran. The attack comprised the following six main phases:

1) **Infection** through an USB with a valid certificate.
2) **Search** for targeted machines.
3) **Update** itself with the latest version.
4) **Compromise** with “zero day” vulnerabilities.
5) **Control** of the systems.
6) **Deceive and Destroy.**

We assume that we have a monitor system that detect events such as USB insertion (x_1), USB removal (x_2), process activation (x_3), process deactivation (x_4), firewall deactivation (x_5), firewall activation (x_6) and increased CPU usage (x_7) and that these events are monitored continuously. The following demonstrates how to detect Stuxnet in its two first phases:

C. *Phase 1: Infection through USB*

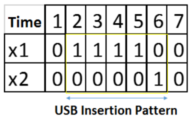

Time	1	2	3	4	5	6	7
x1	0	1	1	1	1	0	0
x2	0	0	0	0	0	1	0

USB Insertion Pattern

Figure 3. How a monitoring system might detect Stuxnet in APT phase 1.

As shown in Figure 3, at instant 2, a user inserts an USB device and, after 3 periods of time, and in instant 5, the user removed an device. At this point, the system recognizes the pattern and checks whether it coincides with a common attack pattern behaviour. Finally, the monitoring system may raise an alarm if a pattern is detected or record this behaviour for later analysis.

D. *Phase 2: Search targeted machines*

As is shown in Figure 4, at instant 9, Stuxnet is trying to search targeted machines generating process activation, firewall deactivation and increased CPU usage. This behaviour will last five periods and will finish in instance 14. At this point, the system will check again whether this behaviour represents an attack action. Finally, the system recognizes the series of events as part of the behaviour pattern of the Stuxnet attack and issues an alarm to inform about the risk.

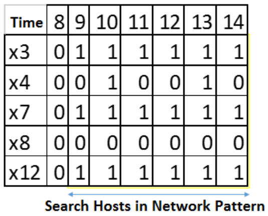

Time	8	9	10	11	12	13	14
x3	0	1	1	1	1	1	1
x4	0	0	1	0	0	1	0
x7	0	1	1	1	1	1	1
x8	0	0	0	0	0	0	0
x12	0	1	1	1	1	1	1

Search Hosts in Network Pattern

Figure 4. How a monitoring system might detect Stuxnet in APT phase 2.

V. Conclusion and Future Work

APT is an emerging threat that has already caused devastating consequences, such as, with the Stuxnet and Aurora attacks. Antivirus and other type of perimeter defences does not provide sufficient protection against these sophisticated threats for various reasons. The same is the case with existing monitoring and detection technologies and solutions, such as Tofino and Industrial Defender. This paper discussed a sophisticated dynamic attack pattern and behaviour approach to detect APTs. It performs monitoring, detection and analysis of cybersecurity events taking both attack and system intelligence into consideration and is able to detect the behavior of evasive threats such as APTs as they are emerging. In order to mitigate false positives created by the proposed approach the system continuously analyses against common attack patterns. This may be improved by applying stochastic processes to model attack pattern and behaviour.

In the future, we plan to test our patterns with other APT attacks and improving the patterns algorithms adding new events to monitor.

Acknowledgment

This work is supported by the Regionalt Forskingsfond Vestlandet project 245291, “Cybersecurity Incident Response Framework for the O&G Industry to Optimize Oil Production and Prevent Safety and Environmental Disasters”, Norway.

References

[1] A. Couce Vieira, S. H. Houmb, and D. Rios Insua, “A graphical adversarial risk analysis model for oil and gas drilling cybersecurity,” in Proceedings First International Workshop on Graphical Models for Security, ser. EPTCS, vol. 148, 2014, pp. 78–93.

[2] Symantec, “Internet security threat report 2014,” 2014, retrieved: May, 2015. [Online]. Available: http://www.symantec.com/security_response/publications/threatreport.jsp.

[3] Mcafee, “Net losses: Estimating the global cost of cybercrime,” 2014, retrieved: May, 2015. [Online]. Available: http://csis.org/files/attachments/140609_rp_economic_impact_cybercrime_report.pdf.

[4] “Crashing the system,” The Economist, July 10 2014, retrieved: May, 2015. [Online]. Available: http://www.economist.com/news/special-report/21606419-how-protect-critical-infrastructure-cyber-attacks-crashing-system.

[5] M. Kenney, “Cyber-terrorism in a post-stuxnet world,” Orbis, vol. 59, no. 1, 2015, pp. 111–128.

[6] M. Zeller, “Myth or reality—does the aurora vulnerability pose a risk to my generator?” in Protective Relay Engineers, 2011 64th Annual Conference for. IEEE, 2011, pp. 130–136.

[7] C. Tankard, “Advanced persistent threats and how to monitor and deter them,” Network Security, vol. 2011, no. 8, 2011, pp. 16 – 19. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1353485811700861

[8] S. McClure, “Operation cleaver,” Cylance, 2012, retrieved: May, 2015. [Online]. Available: http://www.cylance.com/assets/Cleaver/Cylance_Operation_Cleaver_Report.pdf.

[9] C. Five, “Advanced Persistent Threats: A Decade in Review,” Command Five PTY LTD, 2011, pp. 1–13, .

[10] P. Bondarenko, “Security Indicents Management in Offshore Drilling Rigs,” Master i informasjonssikkerhet IMT4882 Specialization Course, 2013, pp. 0–28, .

[11] Websense, “Advanced persistent threats and other advanced attacks,” A Websense White Paper, 2011, retrieved: May, 2015. [Online]. Available: https://www.websense.com/assets/white-papers/whitepaper-websense-advanced-persistent-threats-and-other-advanced-attacks-en.pdf.

[12] M. Cobb, “Advanced persistent threats - attack and defense,” Hacking, June 13 2013, retrieved: May, 2015. [Online]. Available: http://resources.infosecinstitute.com/advanced-persistent-threats-attack-and-defense/.

[13] L. Bilge and T. Dumitras, “Before we knew it, an empirical study of zero-day attacks in the real world,” Symantec Research Labs, 2012, retrieved: May, 2015. [Online]. Available: http://users.ece.cmu.edu/~tdumitra/public_documents/bilge12_zero_day.pdf.

[14] Norton, “Cibercrime report 2012,” Cibercrime Report 2012, 2012, retrieved: May, 2015. [Online]. Available: http://uk.norton.com/cybercrimereport/promo.

[15] E. B. E. Schweigert and M. Thomas, “Securing ethernet/ip control systems using deep packet inspection firewall technology,” Tofino Security, 2014. [Online]. Available: https://odva.org/Portals/0/Library/Annual20Meeting202014/2014_ODVA_Conference_Byres_Schweigert_Thomas_Securing_EtherNetIP_with_DPI_FINAL.pdf.

[16] M. Brian, “Industrial defender solutions,” Lockheed Martin’s, 1996, retrieved: May, 2015. [Online]. Available: http://www.wurldtech.com/.

[17] J. Evans, “Darktrace cyber intelligence platform (dcip),” GCHQ, retrieved: May, 2015. [Online]. Available: http://www.darktrace.com/.

[18] L. Ferris and H. Gilthorpe, “Wurltech technology & professional services,” Wurltech, 2006, retrieved: May, 2015. [Online]. Available: http://www.wurldtech.com/.

[19] J. McCormack, “Websense security labs,” Websense, 1994, retrieved: May, 2015. [Online]. Available: http://www.websense.com.

[20] D. Goldman, “Hacker hits on u.s. power and nuclear targets spiked in 2012,” The Cybercrime Economy, January 9 2013, retrieved: May, 2015. [Online]. Available: http://money.cnn.com/2013/01/09/technology/security/infrastructure-cyberattacks/.