

INSTITUT FÜR INFORMATIK
DER LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN



Bachelor's Thesis

Monitoring and Mitigating attacks in the environment of Software Defined Networks

Christoph Girstenbrei
Alex Marczinek



Bachelor's Thesis

Monitoring and Mitigating attacks in the environment of Software Defined Networks

Christoph Girstenbrei
Alex Marczinek

Aufgabensteller: Prof. Dr. Dieter Kranzlmüller

Betreuer: Dr. Nils Gentschen Felde
Tobias Guggemos
Daniel Migault (Ericsson)

Abgabetermin: 12. Dezember 2016

Hiermit versichere ich, dass ich die vorliegende Bachelorarbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel verwendet habe.

München, den 12. Dezember 2016

.....
(Unterschrift des Kandidaten)

Preface

This thesis is the result of a collaborative work. As such, the source code authored during this work is the result of joint effort.

In this thesis, chapters written by the other contributor are clearly marked as direct citations. As chapters 1,3 and 7 have been written by both authors, the author of each paragraph is marked on the margin.

Abstract

In recent times, Software Defined Network (SDN) have quickly gained popularity. The opened capabilities and addressing of shortcomings in traditional networks has sparked a lot of interest. But considering the extent of testing regular networks are over, SDNs have to make up a lot of ground. Especially due to new attacks on network clients and networks itself, a reliable and adaptive strategy is necessary. Having insight into the network and access to reliable information is one of the most important properties of a modern network. The intent of this thesis is to address the first step of monitoring data, making it available to other components and alert on possible attacks to enable a reactive and flexible network in context of an SDN environment. Different techniques of gathering, analyzing and output data will be proposed and shown with an example implementation. Comparison between these techniques will present their advantages and disadvantages, providing information on which tool to choose for which type of data or attack.

Abstract

Software Defined Networks are a technology that is quickly becoming popular. It enables new network functionality through its standardized interface and the decoupling of the data and control plane.

Yet, also attacks on networks are as popular as ever, with attacks disrupting the operation of even large and well prepared companies. These attacks often are Denial of Service attacks. The architecture of Software Defined Networks also has exploitable vulnerabilities

To help overcome these issues, a detection and mitigation mechanism based on SDN within a reproducible test environment is proposed. This builds a foundation for research on monitoring and mitigation in software defined environments.

This test setup is then used for the implementation and evaluation of TCP SYN Flood, Distributed TCP SYN Flood and Flow Flooding attacks. The results show that by monitoring the network through an Intrusion Detection System and using statistics of the flows, the attacks can be detected and this information can be used to create flows that are capable of mitigating the attacks.

Contents

1. Introduction	1
2. Background	3
2.1. Software Defined Networking	3
2.1.1. Traditional Infrastructure	3
2.1.2. Why SDN?	4
2.1.3. SDN Infrastructure	5
2.1.4. Flows	6
2.2. Attacks	8
2.2.1. IP Spoofing	8
2.2.2. SDN specific attacks	8
2.2.3. Denial of Service	10
3. Related Work	13
4. Test Setup	15
4.1. Concept	15
4.2. Information flow	16
4.3. Implementation overview	17
4.3.1. Network design	17
4.3.2. Monitoring	18
4.3.3. Mitigation	18
4.3.4. Manager	18
4.3.5. Visualization	19
4.3.6. Component communications	19
4.4. External software components	19
4.4.1. Virtual Box	19
4.4.2. Linux Container	20
4.4.3. Vagrant	20
4.4.4. OpenVSwitch	20
4.4.5. Ryu	20
4.4.6. Bind9	21
4.4.7. Suricata	21
4.4.8. Redis	21
4.4.9. ELKG-Stack	21
4.5. Custom software components	22
4.5.1. Statshandler	22
4.5.2. Alerthandler	22
4.5.3. Controller	22
4.5.4. SynFlood	23

4.6.	FlowFlood	23
4.7.	Usage	23
4.7.1.	Setup VM	23
4.7.2.	Controlling the VM	24
4.7.3.	Visualization	25
4.7.4.	Starting attacks	25
4.7.5.	Development guide	25
4.8.	Advantages and disadvantages	25
4.8.1.	Difference to a production environment	26
5.	Monitoring	29
5.1.	Data Gathering	30
5.1.1.	Packet Capture	30
5.1.2.	Network Metadata	30
5.1.3.	Active Client Monitoring	31
5.1.4.	Comparison	31
5.2.	Data Analysis	32
5.2.1.	Limit	32
5.2.2.	Standard Deviation	32
5.2.3.	Triple Exponential Smoothing	33
5.2.4.	Comparison	34
5.3.	Alerting	34
5.4.	Implementation	35
5.4.1.	Packet Capture	35
5.4.2.	Network Metadata	36
5.5.	Attack Analysis	37
5.5.1.	SYN Flood	38
5.5.2.	Distributed Syn Flood	40
5.5.3.	Flow Table Flooding	41
5.6.	Monitoring conclusion	45
6.	Mitigation	47
6.1.	Mitigation concept	47
6.1.1.	Alerts	48
6.1.2.	Alert Handler	48
6.1.3.	Mitigation	49
6.1.4.	Controller Northbound API	49
6.2.	Implementation	50
6.2.1.	Alert Handler	50
6.2.2.	Mitigation	50
6.2.3.	Controller Northbound API	51
6.3.	Mitigation Strategies	52
6.3.1.	TCP SYN flood	52
6.3.2.	Distributed TCP SYN Flood	54
6.3.3.	Flow flooding	58
6.4.	Conclusion	60

7. Conclusion and future work	61
A. Controller	63
B. Mitigation component	81
B.1. Documentation	81
B.2. Source Code Alert Handler	95
B.3. Source Code Mitigation	97
C. StatsHandler Documentation	103
D. SynFlood Documentation	119
E. FlowFlood Documentation	129
F. VM Vagrant File	139
G. LXC's Vagrant File	143
Acronyms	151
List of Figures	153
List of Tables	155
Bibliography	157

1. Introduction

Traditionally, networking technology is a field dominated by the vendors.[JMD14] describes that each vendor supplies its own proprietary hardware and closed soft- and firmware, offering little in the way of flexibility, innovation and adaptation to the steadily changing system and service requirements of enterprises and customers alike. As such, the evolution of networking technology has been slow, and we have yet to see a major contender arise to topple the status quo.

One contender that might have the potential to change the current state of networking, is Software Defined Networking. The defining point of SDN is its decoupled nature, separating control and packet forwarding functionalities into separate layers and centralizing network intelligence and state, thereby allowing for increased functionality, automation, control and programmability [MBR16; ONF14].

OpenFlow is a SDN protocol that has lead to the technology receiving strong interest from various large companies, such as Facebook, Google and Microsoft[GB14], as well as sparking interest from research communities [JMD14]. OpenFlow provides standardized interfaces to communicate with the forwarding devices [McK+08], thus being the basis for an architecture with a centralized controller.

With this new architecture, the forwarding and routing decisions of the devices can be programmed by Network Applications through the controller. They run fully in software and allow every forwarding device to effectively resume responsibilities traditionally handled by specialized middleboxes, resulting in a very flexible network that can better meet the demands of current networks.[GB14]

Security in our digital world is a topic everybody has already heard of. Reports of people managing to compromise or take out the networks of large companies make it to major newspapers regularly. Examples are the attack on Sony Pictures Entertainment in 2014 [BBC14] or, more recently, the DDoS attack on Dyn, a company providing DNS services, which brought down large parts of America's Internet including sites such as Twitter, Netflix and Reddit, in October this year[Gua16].

Therefore the advent of a rather new networking technology has served as motivation for exploring the field of security in Software Defined Networks. By using Flows to influence the way traffic is forwarded through the network and by taking advantage of the advanced statistics OpenFlow devices can collect, new techniques for detecting threats to the network are investigated. At the same time Flows provide a vast amount of possibilities to block malicious packets at the forwarding devices.

To this end, a virtual, OpenFlow based testing environment is created, presented and subjected to a number of attacks. With a Virtual Machine (VM) as an automated testing environment, it is possible to set up the same environment repeatedly and reliably. Modeling a real network as close as possible, Linux Containers (LXCs) are used to simulate network clients. This allows, in comparison to other approaches like Mininet, for a fully functional system stack including system services and maximum separation between the simulated hosts, including the file system. After setting up an environment, attacks can be conducted

1. Introduction

and observed, data gathered and the effects measured. After that, a defense strategy can be established and tested to prove its effectiveness.

To investigate capabilities in an SDN, this thesis takes a look at two different kinds of attacks. The first kind is attacks commonly found in regular network environments. As the new network technology has to withstand already existing attacks, it can prove its resistance to those. The second kind is new attacks that specifically target the infrastructure of the SDN environment. As the network itself has vulnerabilities, SDNs will have to face and defend against such attacks.

An attack commonly found in all networks is DDoS (Distributed Denial of Service). The aim is to try to prevent legitimate users from accessing a service, by consuming all the available bandwidth or resources of a system and therefore rendering it incapable of serving the requests of legitimate clients. There is a plethora of tools like Low Orbit Ion Cannon or Stacheldraht readily available online that allow anyone with hardly any prior knowledge, to execute a DDoS attack.[Lau+00; ZJT13] Given its ease of execution and effectiveness, it comes as little surprise that DDoS currently is one of the most common attacks on networks. [Sec16].

An example for the second category is the so called flow flooding attack. Aiming to modify and flood the flow tables administered by an SDN controller, this attack carries the possibility of a widespread impact, resulting in an unstable and unusable network. As it is a new attack, the tools are not as available compared to Distributed Denial of Service (DDoS) tools. But as SDN grows in prominence, tools will also be more readily available and mature. SDN will have to be prepared for these challenges.

The goal of this thesis is to develop an SDN aware application stack, capable of detecting and mitigating a variety of attacks. In order to achieve this, intermediate steps include the development and implementation of a test setup to simulate a network. Next, an application stack is proposed to handle monitoring and mitigation of the same network. Then, different attacks are executed against the setup to test and prove its capabilities and show its weak points.

The thesis is organized as follows. In chapter 2 Software Defined Networks and Attacks in the network are explained in more detail. In chapter 3 related background literature in the field is discussed. In chapter 4 the implementation of a virtual test environment is discussed in detail and the tools used in later chapters are introduced. Chapter 5 explores techniques to monitor traffic in SDN and the detection of attacks. In chapter 6 strategies to defend against attacks are introduced and mitigation strategies for a selection of Denial of Service attacks are implemented and evaluated in the environment of the test setup. Finally, chapter 7 draws conclusions.

2. Background

This chapter provides a more detailed overview of the core concepts used in this thesis.

More specifically section 2.1 discusses Software Defined Networking and explains why there is a demand for this technology, how it differs from the networks predominantly in use today and how those differences are implemented and can improve the status quo.

In section 2.2 attacks in the environment of Software Defined Networks are explored. Hereby attacks that can also be found in traditional networks are introduced, as well as attacks that target vulnerabilities of the SDN environment specifically.

2.1. Software Defined Networking

Born from an academic research project, SDN has already had a huge impact on the networking industry. Its real success has started with the release of the OpenFlow protocol. With large companies such as Deutsche Telekom, Facebook, Google, Microsoft, Verizon, and Yahoo! being responsible for the OpenFlow standard, the technology is advancing rapidly and while other protocols exist and are being developed, such as ForCES, OpenFlow still remains the most widely used.[GB14]

The differences, in the infrastructure between an SDN device and a traditional Internet switch, already show an integral part of the SDN design. To this end, the following gives a quick overview of the architecture of a traditional switch.

2.1.1. Traditional Infrastructure




Figure 2.1.: Roles of the control, management, and data planes. From [GB14]

2. Background

The switching functions can be divided into three categories. As these categories are capable of horizontal communication with other entities within the same category as well as vertical communication with other categories, they are usually represented as three planes as shown in Figure 2.1: the data plane, the control plane and the management plane.

Data Plane

As Goransson et al. mentions, the data plane is comprised of its various ports plus a forwarding table and is in control of forwarding the received packets, as well as packet buffering, scheduling and header modifications. The forwarding table contains the rules dictating how a packet has to be forwarded. Packets, whose header information can be found in the forwarding table, can be forwarded, and its header, if necessary, modified without the help from the other planes. However, in the case of yet unknown packets and in the case of protocol packets, the involvement of the control plane becomes necessary.[GB14]

Control Plane

The main function of the control plane is to keep the forwarding table, which resides in the data plane, up to date. In order to implement new network policies, the exchange of management information between devices is necessary. Thus the control plane implements many control and routing protocols, such as Spanning Tree Protocol (STP), Shortest-Path Bridging (SPB) and Routing Information Protocol (RIP), necessary for managing the forwarding table and thus also the topology of the network. One of the most basic but without a doubt most important is MAC-Learning. On the arrival of a packet the switch learns, based on the source Media-Access-Control (MAC) address, on which port the device exists and creates a rule in the forwarding table. Based on this rule, future packets sent to that MAC address, can be forwarded to the appropriate port. [GB14]

Management Plane

To change the behavior of the aforementioned protocols, the management plane can interact with the control plane, as well as trigger changes in the data plane, such as changing or adding entries in the forwarding table. Further it can query statistic and status reports from the control plane. To make use of these functions, a network administrator uses some sort of network management protocol. A well known example for such a protocol is the Simple Network Management Protocol (SNMP).[GB14]

2.1.2. Why SDN?

In comparison to this architecture, where the control plane and data plane is tightly coupled together inside the device, which was a design choice due to the urge of having very resilient networks, SDN takes a different approach. SDN breaks this vertical integration by taking the control plane off the device.

The reasons behind this decision lie in the issues that the old highly decentralized architecture causes. As a study conducted by Benson et al. shows, the networks have become increasingly complex to manage [BAM09]. Given this complexity, it is only a natural consequence that misconfigured devices are a common occurrence.

As a single misconfigured device can already have a severe impact on network performance, network administrators have tried to facilitate the management by using specialized hardware, so called middleboxes. A middlebox performs a single task, e.g. a firewall or deep packet inspection.[Kre+15] This takes away the burden of having to implement the policy by configuring the network device directly. However, as discussed in [She+12], they have to be configured through vendor specific commands which, combined with the large amount of middleboxes used, only shifts the problem instead of solving it.

2.1.3. SDN Infrastructure

As mentioned previously, the control plane has been separated from the data plane.

This results in an architecture with three separate layers, as depicted in Figure 2.2: The forwarding devices, the SDN controller and the network applications.

In the following, the architecture is presented in more detail in a bottom-up approach.

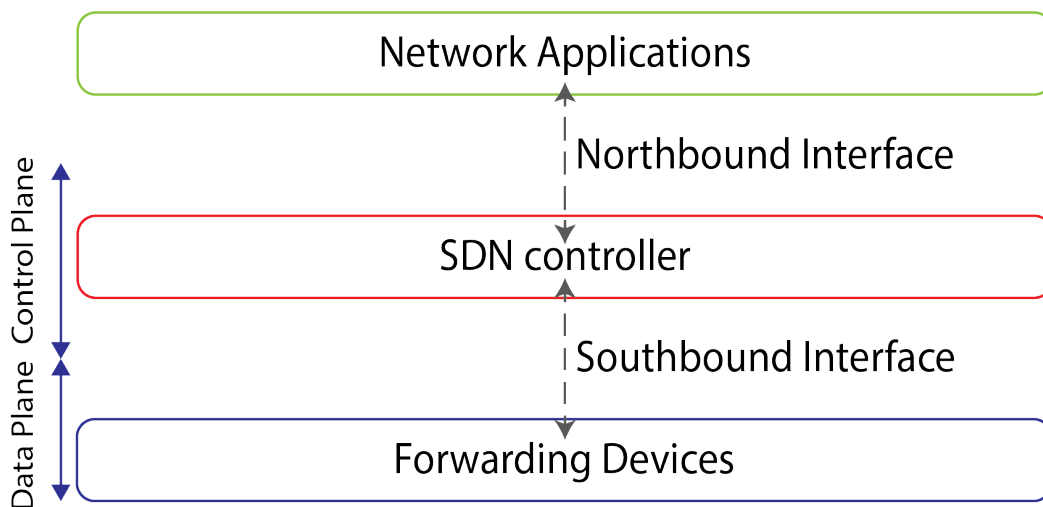


Figure 2.2.: SDN architecture adapted from [Kre+15]

Forwarding Devices

The devices are simple forwarding devices that only know how to handle incoming packets by extracting the characteristics, such as MAC address or IP address, from the header of an incoming packet. This information gets matched against the forwarding table and, if a matching entry is found, the associated action or actions, such as to forward the packet to a specific port, sending it to the controller for further inspection, duplicating it and sending it out on various ports or to simply drop the packet, are executed.[GB14]

Interfaces

The Southbound Interface presents a uniform API for communication between the forwarding devices and the controller. OpenFlow is currently the most well-known Southbound API and standardizes the way the controller should interact with the devices.

The Northbound Interface allows the network applications to talk to the controller. Details of the API are strongly dependent on the controller used. Commonly a REST or an ad-hoc

2. Background

API is used.

SDN controller

The controller resides on a remote machine in the network. While logically always a single entity, he can actually be distributed over multiple machines, allowing for redundancy and adding stability.

To provide another layer of abstraction, the controller uses a Network Operating System (NOS) facilitating access to the underlying low-level devices and resources, comparable to the abstraction an Operating System provides. This is a clear step forward compared to traditional network management that still relies on low-level interaction with the network device for configuration and management.

There are various Network Operating Systems available with different properties, using different programming languages, such as OpendayLight, which is programmed in Java and uses a distributed controller, or POX, which uses Python as the high-level language and uses a centralized architecture. However as POX currently does not support the newer OpenFlow versions, Ryu provides by far more possibilities to configure the forwarding devices and allows access to most features.

The roles of the controller are the discovery of devices joining or leaving the network and providing network topology information, as well as managing the flows in the forwarding tables. These high-level languages, combined with the provided services, make it much easier for the SDN developer to manage the network and implement new network policies.[Kre+15]

Network Applications

To add additional functionality to the network, services such as MAC-Learning, Routing or an Intrusion Detection System (IDS) can be added. They are fully implemented in software and can run on any machine in the network and communicate with the controller via the Northbound API, through which it can receive alerts, such as a device failure or information about the load, from the controller. Yet can also receive alerts from external sources. Usually the network applications reacts based on these alerts and sends commands to the controller to change the flows of the devices. Thus e.g. an application can receive an alert about suspicious activity in the network and decide to redirect the traffic to an application specialized in detailed inspection. [GB14]

2.1.4. Flows

The data structures in an SDN device are the flow tables and they replace the forwarding table in a traditional device. While at least one flow table is mandatory, there are usually multiple flow tables available. Flow tables resemble traditional forwarding tables in functionality, but are more generic than their counterpart.

The individual entries in a flow table are called flows. They each have a priority associated to them and consist of match fields and actions, as seen in Figure 2.3.

These attributes describe the way a packet or set of packets takes from one endpoint to another. Packets arriving at the device are matched against the entries of the first flow table from the highest to the lowest priority. The first entry, whose match field criteria are fully satisfied is chosen and the actions of that entry are executed.

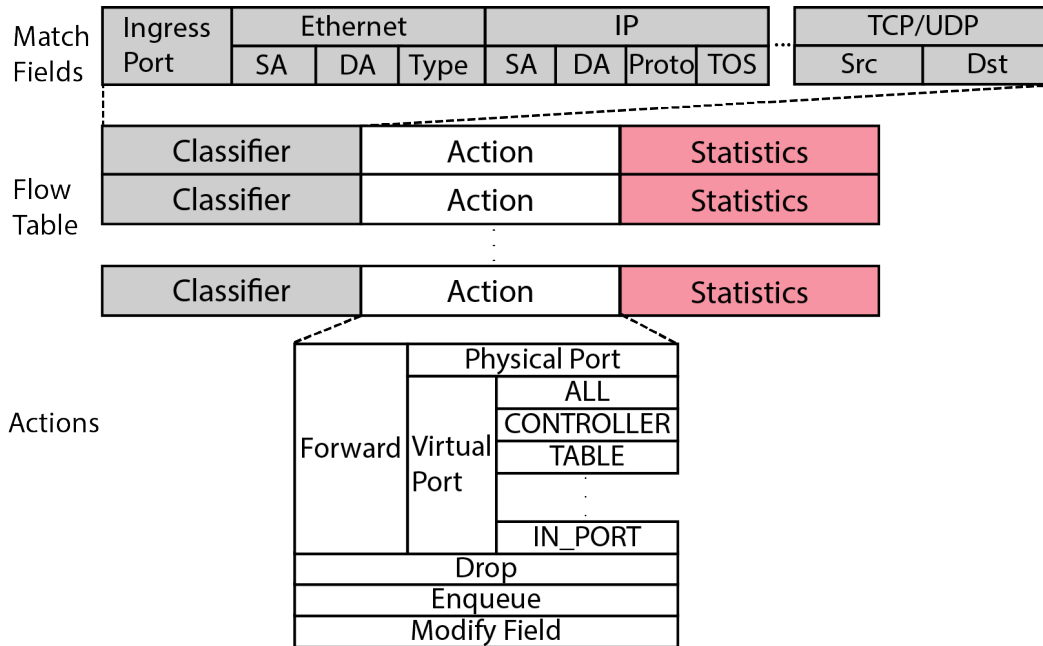


Figure 2.3.: Flow Table and its match fields and actions. Based on [NG13]

By default all possible match criteria are wild-carded, this means any value fulfills it. The fields of interest can then be set to only match a specific value or can be partially wild-carded to match a range of values.

Though the available match fields, e.g. source MAC address and destination IP address, are depending on the implemented protocol, today most devices offer an extensive amount of match fields, giving the SDN developer a lot of flexibility and possibilities in forwarding the traffic. [GB14]

Should no entry match, it is a miss and the packet is discarded. However it is common practice to include a fully wild-carded entry at the lowest priority, sending a miss to the controller instead.

Other possible actions of flow entries, amongst others, are forwarding the packet to an outgoing port, to drop the packet, to send it to another flow table or to a special table like the group table that provides additional functionality, such as duplicating a packet and forwarding it to multiple outgoing ports.

Additionally a flow saves statistics about its usage, such as the amount of packets that has matched to this entry, that can later be queried by the controller. [Kre+15]

2.2. Attacks

Software Defined Networks can, just as traditional networks, be the target of attacks. Denial of Service and Man in the Middle attacks are some of the most prominent examples. But while these attacks apply to SDNs just as much as to traditional networks, the new architecture also opens up new attack vectors.

One thing that stands out in SDNs is the controller. With all the management decisions and the installation of the flow rules relying on the controller, it seems like a prominent target. It helps that the design of SDN allows for a logically centralized but physically distributed controller. However, given its remote nature, it still proves to be a potential bottleneck in the network that can be exploited.

Another security concern is the integrity of the communication on the Southbound and Northbound Interface. Attackers that manage to tamper with the control packets between the controller and the switches or even create fake packets from scratch, can gain full control over the network. Equally important is the communication between the network applications and the controller on the Northbound. While this interface, depending on the implementation of the controller, is potentially not quite as powerful, as the Southbound one, an attack can have the very same consequences.

But not only the management components themselves can be the target of an attacker. The devices, in particular the flow tables, are of interest too. By tampering packets, an attacker can attempt to trick the controller into creating malicious flows that destabilize the whole network. [AX15; LMK16]

In subsection 2.2.2 attacks using these attack vectors are introduced and in subsection 2.2.3 Denial of Service attacks are treated.

2.2.1. IP Spoofing

A security concern, due to the uncontrolled nature of the Internet, is IP Spoofing. As the source of a packet is not verified[YBX11], it is very easy to forge the IP Address and pretend to be someone else, just by changing the source IP Address field in the header.[Tan03]

This enables attacks, like the notorious DNS amplification attack, where the attacker pretends to be the actual target of the attack, causing the DNS server to send large replies to the victim instead of the actual source of the query.[VE06]

On the other hand it conceals the origin of an attack. Being unable to associate malicious packets to an IP Address makes it very hard to trace and stop an attack[Tan03].

As such IP Spoofing is a technique, that many attacks rely on, either as a basic requirement for the attack or as a tool to improve the effectiveness of the attack.

2.2.2. SDN specific attacks

The following shows how the attacks vectors, mentioned previously, can be exploited. To this end an attack, which targets the flow table of a switch, attacks targeting the services provided by the controller, and an attack targeting the communication channel is introduced.

Flow Table Flooding

Technically speaking this attack can be classified as a Denial of Service attack. But what makes this attack unique, is that it targets the flow tables used by the SDN forwarding

devices.

By trying to find a pattern in the way the controller installs new flow entries, the attack tries to launch a stream of packets that generates as many new flow entries as possible. As the amount of flow entries a flow table can store is limited, this will ultimately lead to the flow table overflowing and flow entries that are vital to the forwarding of the traffic, can no longer be installed.




Figure 2.4.: Packet drop due to flow flooding at 100 Mbps. From [KKS13]

Figure 2.4 shows the amount of packet dropped during a flow table flood in regard to the soft timeout for flows, which determines the amount of time a flow may be inactive before it is removed automatically. [KKS13]

Controller

[Hon+15] shows an attack that can seriously compromise one of the core features of the SDN controller. By using a kind of spoofing attack, it misleads the Topology Management Service of the controller. As the controller automatically checks for hosts that have migrated to a new location without further verification, packets can be spoofed to trick the controller into thinking that the target has changed location. Thus future traffic is sent to the attacker's location where, for example in the case of a web server, the website could be impersonated.

Another service that is vulnerable is the Link Discovery Service. As the authentication of the Link Layer Discovery Protocol (LLDP) packets can easily be decoded or reconstructed from the source code of the controller, fake LLDP packets can be crafted. Through these packets actually non-existing links can be created between two switches. As the discovery of a new link also leads to an update of the Shortest Route, this can allow the attacker to introduce a compromised host into the fake link and perform a Man-In-The-Middle attack.[Hon+15]

Communication Channel

As discussed in [Shi+13], the communication channel between the switches and the controller is vulnerable to a saturation attack. As every packet without a matching flow entry is sent to the controller, this channel can quickly become a bottleneck. A DDoS attack designed to cause flow table misses can cause the controller to fall behind in handling these misses and cause the buffer to overflow, resulting in dropped packets. The result is very similar to what happens in a Flow Table Flood, once the flow table is full.

By placing a device between the controller and the switches, for example through a LLDP vulnerability, as discussed earlier, the attacker can gain full control of the network. If Transport Layer Security (TLS) is not implemented the attacker can freely modify or insert flow rules. As the messages from the switches to the control can be intercepted as well the controller will not even notice that something is amiss. Even if TLS is used potentially large parts of the network can be taken down by dropping the packets from the switch and controller. [BCS13]

2.2.3. Denial of Service

A Denial Of Service (DoS) attack is characterized by the explicit attempt to, as the name implies, deny access to a targeted system and prevent legitimate users from using the desired services.

The most common type of DoS attack performed is to flood the target with superfluous requests until its available bandwidth or resources are fully consumed and it can no longer handle regular requests, therefore making the service inaccessible.

Another less common method is a vulnerability attack. This attack abuses a vulnerability in the targeted system and sends malformed packets that lead to the services being unable to perform their intended purposes.[Lau+00]

Distributed Denial of Service

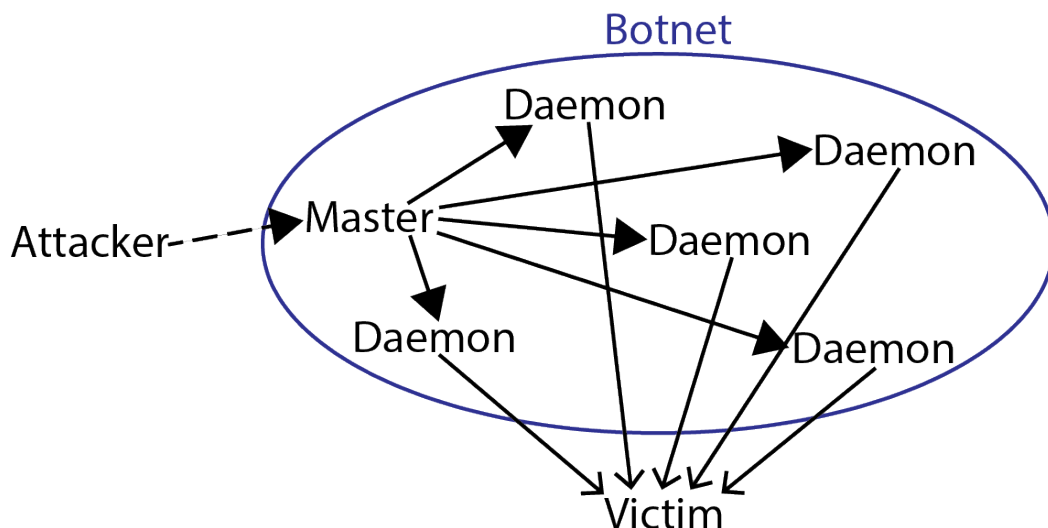


Figure 2.5.: Components of a DDoS attack. Based on [Lau+00]

Most commonly a DoS attack is executed in a distributed fashion and is then called DDoS.

The distributed version is usually executed by a botnet, comprised of computers infected with malware and orchestrated remotely, under the command of the attacker. As seen in Figure 2.5, the attacker coordinates the attack, by sending the attack-command to the master system in the botnet. As a result, the master, who has an overview of the complete botnet, spreads the command to all the other infected systems, usually called Daemon, who then execute the actual attack.

The main benefit of using the master and daemon setup is the concealment of the real culprit behind the attack. While it appears simple to execute the attack, the real work consists of infiltrating enough systems to build a large botnet and searching for vulnerabilities or bottlenecks in the victim's network.

The size of the botnet is often the deciding factor for the success of the attack, as the bandwidth the attack can generate scales linearly with the size of the botnet. This of course directly correlates with the amount of resources the attack can consume.[Lau+00]

The main reason why DDoS is even possible, is that the Internet is not designed to control traffic, but moving packets from source to destination efficiently. This means the intermediate network only does the minimum effort necessary for transporting the packets. However, this design allows for DDoS attacks being possible if only one member in the communication misbehaves, as malicious packets can reach the target without being checked for integrity.[MR04]

TCP Syn Flood

The TCP protocol relies on a three-way handshake to establish a connection, as shown in Figure 2.6. The following shows how this three-way handshake can be exploited to make a system unreachable.




Figure 2.6.: TCP three-way handshake.

On receiving the TCP SYN packet, the receiver replies with a SYN,ACK and waits for the corresponding ACK to complete the establishment of the connection. However, in the meantime he has to allocate resources to save the state of the communication with the initiator.

An attacker can abuse this by never sending the final ACK and thus creating a half-open connection that binds resources of the victim. In a TCP SYN Flood the attacker sends as many SYN requests as possible in order to exhaust all the available resources, rendering the

2. Background

victim unable to answer legitimate connection attempts.[Lem+02]

As shown in [Lem+02], a machine that can have up to 1024 incomplete connections per socket and uses the standard TCP time out of 511 seconds until it drops unsuccessful connections, can be exhausted very easily.

As the limit of 1024 incomplete connections is reached, the machine drops an old incomplete connection to make room for the new connection attempt. If the average time for a packet to make a round trip is 100ms, a Syn Flood attack with a bandwidth of just 4MB/second is enough to completely clear and refill the saved open connections in the mean time. Thus once the ACK arrives there is no record of the corresponding SYN and the connection cannot be established.

Distributed TCP Syn Flood

While the attack remains exactly the same as the non distributed TCP Syn Flood, this attack is far more challenging to defend against.

First of all the use of a botnet makes attacks with an enormous amount of traffic possible. For example Verisign, a company that offers DDoS protections, reports an attack with a peak of approximately 60 Gigabits per second at 150 Million packets per second in Q3 2016 [Ver16]. This makes even systems that have a large pool of resources available susceptible to this attack.

But even more important, the attack packets, especially if combined with IP Spoofing, can no longer be grouped by source IP Address, where a large amount of SYN requests from a single source is a very suspicious behavior. This requires a different defense strategy to mitigate the attack.

3. Related Work

Security in Software Defined Networks is a topic that has garnered a lot of interest recently. The following presents an overview of the literature in the field, related to the work in this thesis.

The setup used in this thesis is heavily based on LXC's to simulate and separate network components within the VM. In [GDK], the method of encapsulated VMs is used. This allows for a more complete component separation, especially by removing the dependency on a shared kernel across LXC's. As the separation capabilities of LXC technology are sufficient to simulate multiple network clients and the kernel dependency does not pose a problem in this setup, the advantages of LXC technology, namely a lightweight isolation approach and therefore fast startup and reaction times, outweigh the advantage of paravirtualization or full virtualization.

Introduced by the monitoring concept of OpenFlow itself, new problems arise. One of them is due to the pull based nature of the protocol, since a gathering component has to actively collect information from the network. This asks for a balance between detailed monitoring and network overhead, a possible solution for this problem is proposed in [Cho+14]. Implementing an abstraction layer between the SDN controller and network monitoring enables for a more intelligent and adaptive monitoring of components than the one currently implemented in the setup of this thesis.

Seen from this angle, scalability and distributability of the monitoring logic in large scale network is necessary. Evaluating the OpenFlow (OF) native approach in comparison with a strategy based on sFlow is done in [Gio+14]. Demonstrating the problem by overloading the SDN control plane with monitoring traffic and comparing the results of both methods with real and high volume traffic, the paper uses an entropy based method to analyze data. As the network used in this setup is considered small, the issue of overloading the system with monitoring traffic is not present, especially due to the low network base load. To implement a similar network on a larger scale successfully, these problems have to be taken into consideration.

A major problem for the security of networks is IP Spoofing. Yao et al. have proposed a mechanism based on the OpenFlow architecture that tries to detect spoofed IP Addresses called VAVE (Virtual source Address Validation Edge). This is an improvement of the system proposed by IETF, SAVI (Source Address Validation Improvements) [BG13]. By forming a perimeter of OpenFlow devices around the network, every packet that comes from a legacy device or from outside of the network is checked. Packets that arrive from outside of the network with an IP Address that is in the IP Address range of the network can then be safely declared as spoofed and dropped. While this technique is useful to prevent attackers from spoofing IP Addresses of systems inside of the network, it cannot prevent any attacks that rely on mostly spoofing random IP Addresses. [YBX11]

AVANT-GUARD[Shi+13] uses SDN switches as SYN proxies. When receiving SYN requests from a new source, the switch starts the TCP handshake using SYN Cookies. Only if the client successfully completes the handshake, the switch installs a flow to forward

3. *Related Work*

subsequents packets from that source and completes the handshake with the original destination[Amb+15]. But [Amb+15] also analyzes that AVANT-GUARD is susceptible to a buffer saturation attack, as the switch has to save state to migrate the TCP handshake to the destination.

Cui et al. [Cui+16] propose a system based on 4 modules: Attack detection trigger, detection, traceback and mitigation. They propose a new method to trigger the detection of an attack. Instead of periodically checking, they monitor the amount of packets sent to the controller to start the attack detection, which only then starts gathering the flow stats and searches for malicious entries which trigger mitigation. For their test environment they use mininet. While mininet is even more lightweight than the virtualization via LXC, mininet uses a shared file system and PID space. To guarantee a more flexible host environment, LXC was chosen for this thesis.

4. Test Setup

This chapter aims to provide on one hand a high level overview enabling an understanding of conceptual parts, component interactions and information flow. It shows the basis and concepts on which the setup is implemented on. After providing a design, it will then on the other hand show implementation details giving insight into important parts of components. The parts and software used will be explained and placed in context of the design concept. At the end, advantages and disadvantages of the setup will be compared.

4.1. Concept

To develop a concept for the test setup, 3 basic goals have to be met. The first one is detecting and mitigating attacks in an SDN environment. The second goal is grouping every functionality into encapsulated and movable components following the Single Responsibility Principle (SRP) [Mar03, pp. 95-98]. This allows a more complex environment and separates concerns. Third, all communication paths shall be defined clearly, use Ethernet/Internet Protocol (IP) and be shaped appropriately for the type of information transmitted. This is especially important in a network environment, as the components are not necessarily inside the same virtual or physical machine. Communication can happen across different base media, but the standardization in place by using the IP protocol ensures available communication everywhere and, if necessary, in between every component. Figure 4.1 shows the underlying idea in the application design.

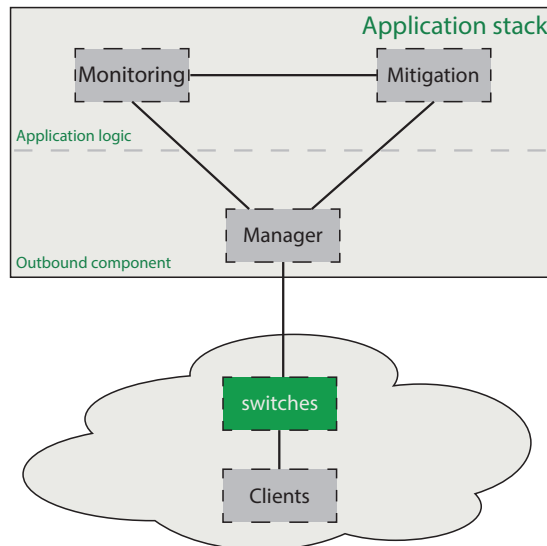


Figure 4.1.: Application stack concept

Responsibilities are split into three base parts. The manager actively communicates on

4. Test Setup

its southbound interface with switches, which provide the SDN network to clients, and are considered the outbound components. Components executing application logic do not communicate with clients directly, only the manager is facing the network side. On its other side, on its northbound interface, the manager provides communication to the application logic, making information actively or passively available to dedicated parts and receiving instructions on how to modify the network. The monitoring components responsibility is to decide whether available data is an anomaly or regular and legitimate traffic. It has to select appropriate tools to analyze data provided by the management component and if an attack is found, an identifiable alert has to be sent to the mitigation component. In summary, its goal is detection. Mitigation reacts upon a received alert with the goal of providing a mitigation strategy. Matching reactions have to be available for different alerts and the actions have to be communicated to the managing component for it to implement changes in the network.

A fourth component can be visualization. This part is not necessary for a working application stack, but it can be used to get insight into the running system thus providing easier use and a smaller feedback loop for human testing. An overview of components and responsibilities can be seen in table 4.1.

Component	Responsibility
Monitoring	Analyze data to generate alerts if necessary
Mitigation	Decide upon a received alert how to mitigate an attack
Manager	Provide a single access point to the network
Visualization	Give insight into events generated by the system

Table 4.1.: Responsibilities of application stack components

Combining all components achieves the primary goal of detecting and mitigating attacks in the SDN environment. Leveraging the three basic concepts keeps components in loose coupling and high cohesion [SMC74], ensuring distributability and portability.

4.2. Information flow

As stated in section 4.1, communication paths and therefore information flow is a vital part of the application. In a large network environment, the sheer amount of information possibly available requires consideration on how to deal with it. In this limited test environment, data flow management is easier compared to a large one. Figure 4.2 depicts which paths information takes to enter the system and how it moves inside the system.

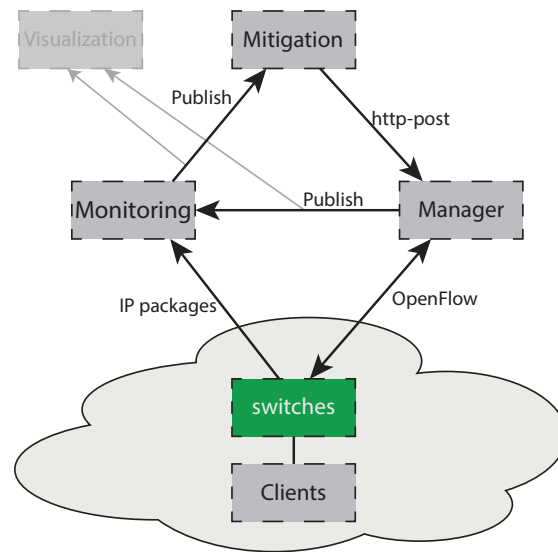


Figure 4.2.: Information flow through the application stack

4.3. Implementation overview

Information flowing into the application stack is always initiated by the manager. This can be statistical information pulled from a switch or IP packet forwarded by one through the network. There are two types of data presented to the application stack. In the first case, IP packets are sent directly to the monitoring component being forwarded by switches. Second, statistical information is gathered by the management component and relayed to the monitoring instance. After data entered the application stack, monitoring decides whether to alert on available information or not to do so. If an alert is generated, it will be received by the mitigation component. This instance will decide on whether the attack is mitigatable and if so, it will inform the managing component about what to do, in order to achieve mitigation. The manager will then install flows according to commands by the mitigation client to appropriate switches. The logical flow of information and the transportation techniques are depicted in figure 4.2.

4.3.1. Network design

The design is chosen with a production network in mind, but aiming for simplicity and ease of use while testing different attacks. Therefore at least two switches are necessary to test how multiple network entry points, with a number of clients attached, behave. Although in a production environment there won't necessarily be a second network available only to connect the monitoring and mitigation stack, this can be achieved by different techniques like VLAN and Quality-of-Service controls similar to a non-SDN -network. In this setup an ideal configuration can be chosen due to no limits concerning available hardware switches and physical paths. Therefore a complete network separation in two different segments is implemented, providing a public network to connect all clients to and a private network providing connectivity to the monitoring and mitigation stack. Though it may be more

4. Test Setup

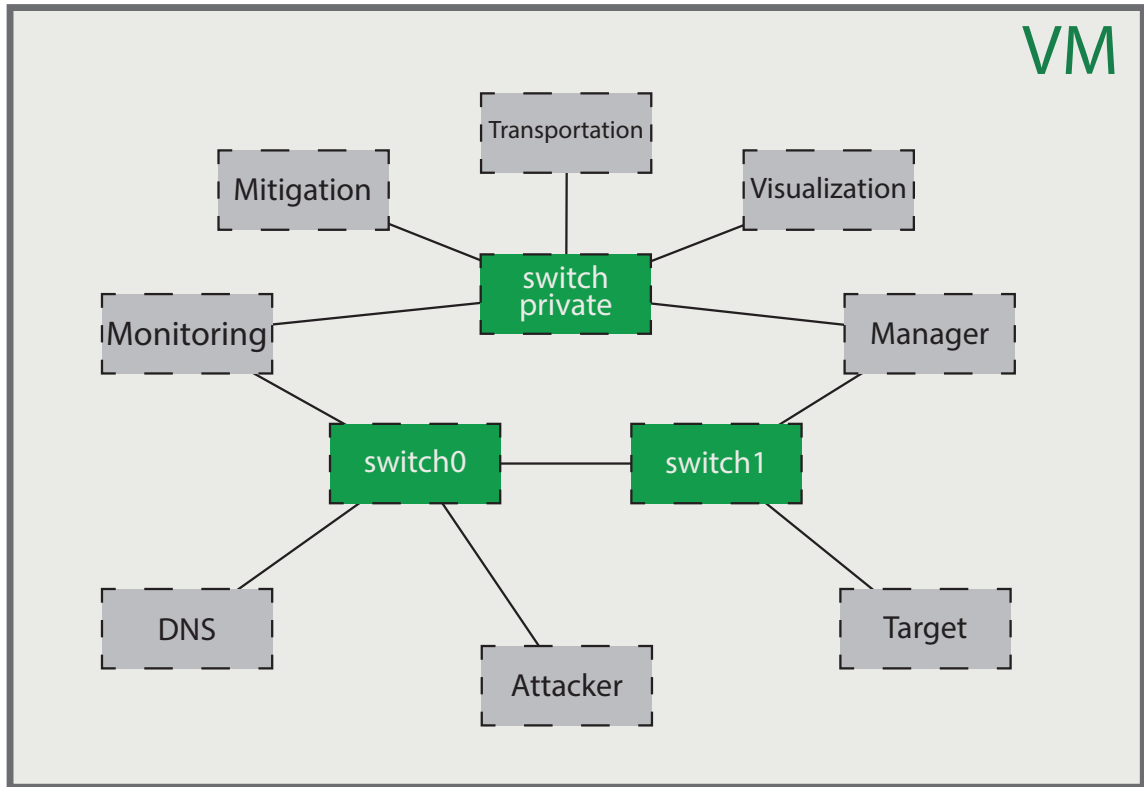


Figure 4.3.: Test-setup of the virtual network

difficult, network separation between the public network and the monitoring and mitigation stack is advisable to minimize the impact in case a system is under attack or compromised.

4.3.2. Monitoring

To analyze any event occurring within the client network, all necessary information gets forwarded to a monitoring client. The same can then decide to alert if specific conditions are met to inform a mitigation client. In this setup, two services are working inside the component. The first one is Suricata (4.4.7) as an IDS. The other one is the Statshandler(4.5.1). More details about monitoring internals are described in chapter 5.

4.3.3. Mitigation

After an attack is discovered by the monitoring component, mitigation can communicate with the manager, trying to mitigate the effects of an attack. Running inside the component are two instances of the Alerthandler(4.5.2), each one a counterpart to the services running in monitoring. This is explained in detail inside chapter 6.

4.3.4. Manager

This component provides an interface to the network. Running Ryu (4.4.5) as an SDN controller, it contains the network application. Responsible for sending instructions to switches,

receiving statistics and requests from them on one hand, it is providing a Hypertext Transfer Protocol (HTTP) interface for mitigation to control the network as needed on the other.

4.3.5. Visualization

On top of that, a visualization component provides almost real time insights to events happening within the network. An Elasticsearch-Logstash-Kibana-Grafana-stack(4.4.9) allows for collection, detailed inspection and visualization of occurring events of almost any type and source. In contrast to monitoring and mitigation, which are absolutely necessary to system functionality, visualization is purely for the ease of use and a faster feedback cycle while performing tests in this system. It can also be used to see and notice events, making it easy for a human viewer to correlate between events on different systems. Especially trying to identify the characteristics to monitor while testing an attack and rechecking if a mitigation strategy works is much faster, if data is presented almost instantaneously in graph form.

4.3.6. Component communications

Communication between these components happens in two different ways. First, Redis(4.4.8) enables via its publish/subscribe model high speed data exchange between a single data source and multiple recipients. This is used to communicate between the monitoring, mitigation and visualization component. As a format to exchange data, Extensible Event Format (EVE) is used. The format is natively supported by Suricata(4.4.7), and its JavaScript Object Notation (JSON) base can be easily parsed. Second, a REST-full API is used to instruct the network operating system what to do in case of a mitigatable alert. This ties in with the already provided API and capabilities of Ryu.




Figure 4.4.: Flow of alert messages

4.4. External software components

The following components are used in conjunction to form the VM setup of this thesis on which all tests will be executed. As many components as possible are open source and all can be used without license fees.

4.4.1. Virtual Box

To ensure encapsulation and portability of the whole test environment, all tests are performed within a Virtual Machine. Enabling simulation of an arbitrary number of devices,

4. Test Setup

multiple topologies can be tested inside the VM. Contained are LXC's, switches (OpenVSwitch (OVS)) and connections between containers via virtual network adapters. To support a broad range of host platforms, VirtualBox v. 5.0.26 is chosen as a platform independent virtualization provider. The base operating system is Ubuntu 16.04 LTS (64bit), both for the virtual machine and the LXC clients. It is installed automatically.

4.4.2. Linux Container

To simulate multiple network clients within a single machine, LXC's are used. They get provisioned via Vagrant and connected to OVS's with virtual network adapters. This enables limiting client resources like available CPUs and RAM via Linux cgroups. Limiting CPU usage is done in the test setup by constraining a container to one CPU. This works also on machines with fewer than 9 cores despite having 9 LXC's. Although it is not possible to truly separate all running processes due to the nature of a single host machine, separation into one container per functionality allows for the best control possible and follows a micro service approach.

4.4.3. Vagrant

To achieve repeatable and reliable tests, an automated setup process for the test environment is necessary. Therefore, Vagrant (1.8.4) is used to automatically provision all virtual machines and containers. The process handles setup of the VM, all LXC's within the machine and connection of all parts to a virtual network. During this setup process, different base images are downloaded from multiple repositories, so Internet access is mandatory. This implies that all containers are connected to an outside network during their setup. This connection can be shut down to ensure network encapsulation and to guarantee all packets sent during the tests will stay within the virtualized network, even if a configuration error is present. Additionally it is possible to commit the infrastructure as code to a version control system and use all advantages like branching and resetting to a previous commit.

4.4.4. OpenVSwitch

OpenVSwitch (v. 2.5.0) is used to simulate OpenFlow capable network switches. Every client is connected to a virtual switch during the setup process. The Networking Operating System is running as a client to control every switch.

4.4.5. Ryu

The Networking Operating System (NOS) controlling all OpenFlow switches is Ryu. It is fully open source and provides a well defined API for communication with the switches in Python. Since Ryu, in the version used for this project (4.5), supports the Open Flow Protocol up to version 1.5, it is the most suitable choice. In the configuration provided by the setup process, Ryu controls the two client-facing switches and provides basic layer 2 network functionality by installing appropriate flows. Additionally, some are installed to duplicate every packet entering the network and send it to a monitoring instance for further analysis.

4.4.6. Bind9

Bind9 is used as a Domain Name System (DNS) server. It runs inside the DNS container and provides its service to all clients connected to the public switches, as known from regular environments.

4.4.7. Suricata

Suricata is used as a packet analyzer and Intrusion Detection System (IDS). It runs as a system service inside the monitoring component and listens on a network interface for incoming IP traffic. It is capable of analyzing different aspects of an IP packet, like its origin, a TCP/UDP port or its contents. Based on the analyzed content, it can generate alerts if predefined criteria are met. It is manually compiled and not installed from standard packages as they don't support the output of alert information into Redis at the time of this thesis. It was chosen due to its similarities to the widely known IDS Snort and its extended capabilities compared to the same. More details about how Suricata is used in this setup are available in subsection 5.4.1. More information about Suricata itself can be found in its user guide [Her+16].

4.4.8. Redis

Redis is an open source (BSD licensed), in-memory data structure store used as database, cache and message broker. - [JM16]

As this setup does not only provide monitoring and mitigation functionality, additional message distributing problems are introduced. Namely, both visualization and mitigation need to receive all messages. This raises the problem of a single source and multiple destinations for a given message. Solving this is the publish/subscribe model provided by Redis. It enables a Redis client to publish messages to a channel without having to deal with the responsibility of distributing it to all subscribed clients.

The system is configured to forward all messages in memory. This implicitly means, none are stored on disk by Redis and there is no replay capability. This is not needed in this test setup, but may be necessary in a more complex environment.

4.4.9. ELKG-Stack

This component stack consists of Elasticsearch, Logstash, Kibana and Grafana. Logstash is the first component in line, receiving input messages, formatting and storing them in Elasticsearch. Though there is no capability of doing this in Redis, all transmitted ones are saved in Elasticsearch. All messages are indexed and made available to query for in Kibana and Grafana. Kibana can be used with its 'discover' tab to view the different fields available within the data. Grafana is configured to set up a dashboard automatically at start up in order to visualize the change and relationships of different parameters in the system. More information about Elasticsearch, Logstash and Kibana can be found at the website of Elasticsearch. Details about Grafana are available at <http://grafana.org/>. Instructions how to use the visualization component can be found in subsection 4.7.3.

4.5. Custom software components

Many tasks in the application stack can be handled by pre-written, third-party software. But some components listed in the following chapter are specifically designed, built and tested for this thesis. All services use Python as their programming language, as it was already used by Ryu and therefore necessary for the controller. To ensure consistency and ease of extensibility, all other components were written in Python 3, too. The communication between components is shown in Figure 4.4.

4.5.1. Statshandler

The Statshandler is an analysis service inside the monitoring component responsible for detecting anomalies in statistical data collected from switches. It is running as a system service and can be controlled as such. Its source code is located at `./VM/lxcs/main/monitoring/statistical`. More detailed information about its functionality is found in the monitoring chapter, subsection 5.4.2.

4.5.2. Alerthandler

The Alerthandler is responsible for receiving alerts sent by the monitoring components and deciding how to mitigate an attack, based on available information. Like Statshandler, it is running as a system service. It runs inside the mitigation container and its source code can be found at `./VM/lxcs/main/mitigation`. Its functions and internals are described in more detail in the mitigation chapter, subsection 6.1.2

4.5.3. Controller

The controller is communicating in two different directions. On its northbound API it enables communication to other systems within the network. In this setup, there are two main communication channels northbound. Redis is used by the controller itself to insert statistical data gathered from switches into the application stack. HTTP is used as a second way of communication by the manager to send instructions to the controller in case an alert is triggered.

Southbound, the controller connects to all switches using the OpenFlow protocol. Every time a switch starts, it will announce itself to the controller and waits for instructions from it in the form of flows on how to handle network traffic. The setup process can be seen in listing 4.1. More information about the concept of a north-/southbound API and how it is used in OpenFlow environments is described in section 2.1.3.

```

1  return _rest_command
2
3
4 class SwitchController(ControllerBase):
5     _SWITCH_LIST = {}
6     _LOGGER = None
7
8     def __init__(self, req, link, data, **config):
9         super(SwitchController, self).__init__(req, link, data, **config)
10        self.waiters = data['waiters']
11
12    @classmethod

```

Listing 4.1: Register switch

4.5.4. SynFlood

There are multiple tools available capable of executing a Transport Control Protocol (TCP) SYN flood, but there is a common problem to them. They are built to execute a blunt flood and don't provide much control over how the exact procedure of flooding is done or when to stop. A similar tool, Hping3 can flood a target with SYN packets, but is not able to stop after a specific amount of packets in flooding mode. These shortcomings are addressed with this tool, which is capable of executing a precise number of packets sent to a target with or without spoofed source IP. Though being implemented in Python, the tool was not slower in this test setup than the in C implemented Hping3.

4.6. FlowFlood

As flow flooding is, compared to SYN flooding, a fairly new attack, no tool was available to produce packets in exactly the manner necessary to trick the controller into installing flows. This problem is solved by the FlowFlood implementation. Working in a master/slave fashion, the program is executed on two different network clients, establishing a TCP connection and sending a small amount of data via the same. Then, the MAC address of the master is changed and a new TCP exchange performed. Every time this is done, a new flow is installed by the system.

4.7. Usage

The virtual machine can be used to simulate the attacks and their effects live. Therefore, a user and development guide is provided here.

4.7.1. Setup VM

The first step in setting up the testing VM is to ensure VirtualBox is installed. Also, Vagrant and the Vagrant-vbuest plugin have to be installed. Almost every part of the installation process needs to download files, so an Internet connection is mandatory.

After changing the working directory in the used shell to `./VM`, the installation process can be started by issuing the Vagrant native command `vagrant up`. If the shell used during the installation process is capable of displaying color, the output will be colored. As some

4. Test Setup

of the installation process happens inside the VM by another instance of Vagrant, the color coding can be wrong, but in general it is normal to see a lot of green lines, some will be red and any lines marking different sections within the installation will be purple. Red lines not necessarily indicate critical errors. For example, output of the curl command will be marked red, but is not an error. If the `vagrant up` command exits with an error code, a critical exception occurred. Additionally, at the end of the installation process, every service will be checked and displayed as 'running' or 'broken'. The installation process should, after installing Vagrant, VirtualBox and switching to `./VM`, only involve the following two commands:

```
1 vagrant plugin install Vagrant-vbguest
2 vagrant up
```

Listing 4.2: Setting up VM

All components should be already set up and run after the installation process completes. This is checked once with a script at the end of the installation process. All necessary connections and services must be marked with 'Internet' or 'Running' and the IP for this installation is displayed. To check live status of components, either logging into the VM and rerunning the test script located `/home/vagrant/lxcs/test-main.sh` or looking at visualization of Grafana is possible.

4.7.2. Controlling the VM

To get shell access to all components, the Vagrant command `vagrant ssh` can be used. This connects the used shell in an ssh equivalent fashion to the Virtual Machine (VM). Inside it, the `sudo` command is available without password authentication and necessary for all further LXC commands.

All standard Linux/Ubuntu 16.04-compatible commands are available. Additionally, TMUX and VIM are installed. To ensure that containers are in fact started, `lxc-ls --fancy` displays a list of running LXCs and their IP addresses. To connect to one of them, the command `lxc-attach -n container-name` is available. When logged into the VM, all Linux/Ubuntu 16.04 commands are available inside all LXCs. Some example use cases from within the VM are shown in listings 4.3 to 4.5.

```
1 lxc-stop -n monitoring
```

Listing 4.3: Stopping a container

```
1 lxc-attach -n mitigation
2 systemctl restart alerthandler-stats
```

Listing 4.4: Restart a service

```
1 lxc-attach -n monitoring
2 journalctl -u statshandler -f
```

Listing 4.5: Following the log output of a service

4.7.3. Visualization

To visualize all information produced by the system, two services are in operation. Kibana is the standard visualization software used in many Elasticsearch-Logstash-Kibana installations. Using Kibana, all unique fields are visible from the 'Discover' tab. Kibana is available via port 81 of the VM, login data is 'kibanaadmin' as user and '123' being the password. Grafana is used to visualize multiple graphs. It can be used to see time correlated events and build almost real time insight graphs to display multiple source fields from Elasticsearch. Grafana is available via port 80 with the login 'admin' as user and password.

More information about how to configure and use Grafana can be found in its documentation: <http://docs.grafana.org/>. Equally, usage information about Kibana can be found in Kibanas documentation: <https://www.elastic.co/guide/en/kibana/current/index.html>.

4.7.4. Starting attacks

An attack is always started from within one of the LXC attackers. All scripts starting an attack are located in `/vagrant/attacker/` and can be executed without any parameters to use the predefined targets. As these are simple bash scripts, the target can be changed by modifying the script. For example, to execute a simple DOS attack one would issue the following commands:

```
1 lxc-attach -n attackone
2 cd /vagrant/attacker
3 ./synflood.sh
```

Listing 4.6: Starting DOS attack

As long as the command is running, a packet flood is visible in Grafana.

4.7.5. Development guide

Extensions and changes to the project should follow the already available directory structure. All information necessary for a component to offer its service should be located in the components folder. In case of an LXC this would be `./VM/lxcs/container-name/`. To modify the network behavior, changes to the controller can be made in `./VM/lxcs/main/manager`, either directly to `controller.py` or by creating new modules.

To create or modify attacks, files in `./VM/lxcs/main/attacker` can be added or changed. Changes to monitoring and mitigation would then happen in `./VM/lxcs/main/mitigation` and/or `./VM/lxcs/main/monitoring`.

The development of custom Python modules follows the style guide PEP8 with a line width of 120 characters. Python 3 is preferred over Python 2 as all current modules are written in Python 3.

Bash is the used shell scripting language.

4.8. Advantages and disadvantages

This testing environment was chosen due to its advantages. The most prominent is repeatability. Due to the capabilities of Vagrant, the VM can be set up from scratch again and

4. Test Setup

again with one simple command, providing a fresh install of the test setup every single time. Ensuring the same test environment not only in between test runs, but also across a development team, makes testing repeatable and development easy. Even if something inside the VM stops working, e.g. the virtual hard disk fills up due to too many packet capture files, the whole machine is simply disposable and set up from scratch in a few minutes. Additionally, by specifying architecture in text like files, version control is simply done in git, allowing version controlled architecture and the merging of system features. Using LXCs inside the VM allows for a full network setup while keeping the most separation possible inside a Linux machine. With this container based solution, every single LXC runs like a separate machine, including the possibility used in this setup of running system services via SystemD. Again using Vagrant, to set everything up, makes it easy to spawn new containers. In summary, the VM is as invisible as possible to the test setup and every function is run in a container.

Disadvantages of this setup include an incomplete separation of network components. Though separated, all containers still run on the same hardware sharing the same VM. This leads to the possibility of unwanted influence from one container to another, e.g. in disk IO operations. These will still read and write from and to the same disk. Another disadvantage is a possible scaling issue. This setup is not intended to be used with a lot of clients, and this problem could only be solved to some degree by providing a larger VM host. Nevertheless, solutions to the problem would be splitting the VM onto multiple hosts and connecting them via hardware. Additionally, there are some security concerns discussed in the following section 4.8.1.

4.8.1. Difference to a production environment

As this setup is a testing environment, it is not production ready. The intent in building this setup was to provide ease of use and proof of concepts. Therefore, the network and application design is not transferable to a production environment as there are no physical distance limitations within a VM. Additionally, some security measures found in a production environment are not implemented in this setup. To run a similar stack in production, the setup process and the implemented security measures would have to be adapted to the real network and the security guidelines in place. A minimal set of changes necessary are the following.

The setup files now used to set up all LXCs have to be adjusted for at least another provider, if Vagrant is still used as a setup manager and only the provider changes. After adjusting the IP addresses and domain names of all containers, the shell part of the setup process can be used to repeat the installation on a variety of hosts capable of running Linux/Ubuntu 16.04 and executing a shell script.

All set passwords should be set to secure ones. Neither the VM base machine used as template, nor the LXC base follows any security guidelines other than the defaults provided by the Ubuntu operating system. The `sudo` command can be used without password. Additionally, Vagrant mounts shared volumes to all components for the ease of use of sharing files. These may be replaced by NFS shares if necessary.

There are currently no authentication and authorization processes in place. The communication via HTTP from mitigation to manager is not secured, a switch to Hypertext Transfer Protocol Secure (HTTPS) would be necessary. Interaction with Redis also happens unencrypted and without authorization by all clients. Likewise, the connection between switches and manager is also insecure. Secure and proven communication possibilities are available

for all these channels so a production ready and security guideline compliant setup can be built.

An overview of minimum security changes is shown in Table 4.2.

Component	Current Setup	Secure Setup
Redis	Every client with a connection to the secure switch has access to Redis	Use stunnel to encrypt communication[Hab14]
OpenFlow	Plain TCP messages to the controller	Switch to secure mode for all switches and TLS for all communication[Nyg+14]
HTTP communication	Simple HTTP messages between mitigation and manager	Use TLS for all HTTP traffic
Visualization	Plain HTTP and simple passwords	Use TLS and change passwords to secure ones
Identity provider	Vagrant SSH with generated keys and sudo without password	Implement deployment specific security guidelines

Table 4.2.: Example changes to secure setup

5. Monitoring

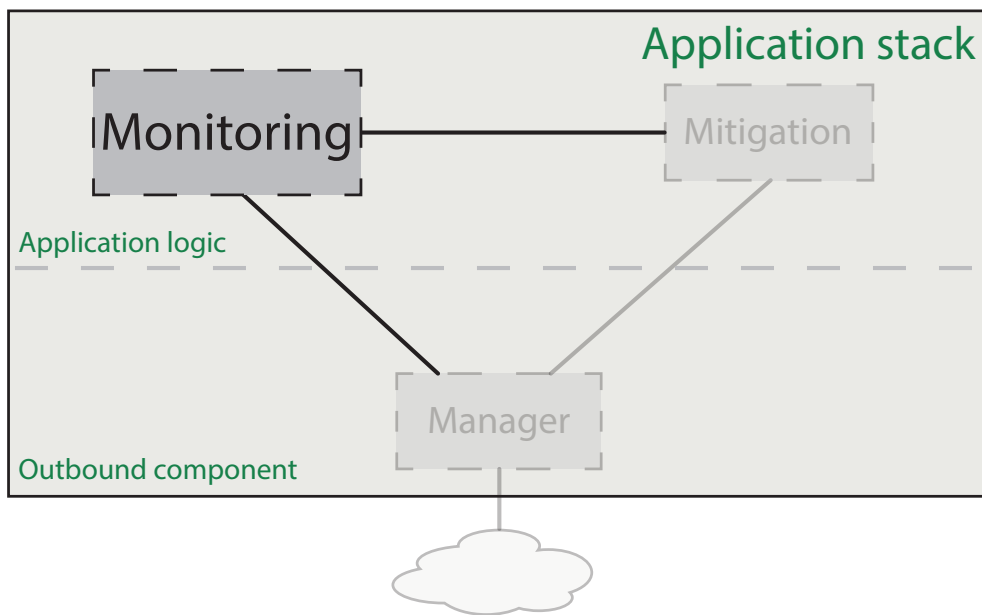


Figure 5.1.: Monitoring in the concept

As a first step in the pipeline to a successful mitigated attack, a working monitoring system is key. All further actions require reliable information and alerts to act upon. At the same time, to keep latency between the start of an attack and mitigation of the same low, the monitoring system has to operate in real time. In contrast, to take the best informed decision possible, as much information as possible needs to get analyzed. This means getting as much relevant information as possible to the component responsible for analyzing it and having all of it analyzed fast enough to keep up with the network.

In itself, the monitoring component has to fulfill three steps. First in line is to gather data from external sources to a decoding component and making it available locally. Then the application logic within monitoring can analyze presented information to decide whether the last step of alerting mitigation shall be performed. An overview to this behavior is given in figure Figure 5.2.

The following sections will give insight into different methods and techniques to perform the steps necessary to generate valid alerts from raw data. First, the concept of the three phases used will be described, to be followed by a description of what is implemented in the VM. Then, the three attacks described in section 2.2 are executed and an analysis of the events is provided.




Figure 5.2.: Monitoring concept

5.1. Data Gathering

Regardless of the type of data gathering used or the type of data collected, all data has to arrive at the monitoring component. The following compares three different methods of collecting data.

5.1.1. Packet Capture

Packet capture can catch every packet injected into the network and forwarded to monitoring. After decoding the input, statistics can be computed using signatures in the analysis phase.

To generate alerts in a reliable way, the component has to analyze as many potentially important packets as possible and therefore receive a large number of them. If there is no filtering previously to decoding the packets, this results in the disadvantage of having to deal with all packages available. Though it is possible with SDN to do so, it can be a problem not only for the component itself, but also for other parts of the network involved, to reroute many packages to a single destination. To minimize this problem, the component can be placed at already existing network constraints like routers, firewalls or other forwarding devices.

Physical limitations like bandwidth can cap the amount of packets a component can receive, especially in situations where traffic floods traverse the network (e.g. DoS attacks). A client can have direct influence on how many packets have to be processed by monitoring, so it may be suitable to distribute package capture.

Another issue emerging from the number of packets captured in combination with how complex and resource intensive the analyzing logic is, is that performance can be a problem. Due to the constraint of having to keep up with the network traffic, thorough analysis of every packet may result in the component not being able to capture all following packets, resulting in missed information. An advantage of having this much source material in its raw form, namely whole packets, is the flexibility to decide in the analysis phase on any signatures detectable inside the packet, including payload of any kind. This enables picking a select number of packet features independent of other ones like e.g. alerting on payload content independent from source IP address. A sample of possibilities can be seen in the Suricata documentation [Her+16, 4. Suricata Rules] In summary, packet capture provides a lot of raw information to the next phase. This is an advantage and a disadvantage at the same time.

5.1.2. Network Metadata

Network metadata is all data that can be retrieved by gathering information from the network infrastructure itself. In an SDN environment this is possible using the OpenFlow protocol.

Every switch can gather data about its ports and installed flows. Available counters can be seen in the OpenFlow specifications [Nyg+14, p. 32]. This information can be relayed via the management component to monitoring and processed to detect anomalies.

A fact to consider when setting up network monitoring via metadata is, that information is implicitly grouped and filtered by the flows installed. As in OpenFlow a packet is assigned to a flow via its match fields, these determine also the granularity of monitoring information available. In OpenFlow v. 1.5 these can match many header fields including an Ethernet frame, all available are listed in the OpenFlow specifications [Nyg+14, pp. 77-78]. As a result of the assignment to a flow, high coupling is present between routing information on one side and monitoring data on the other. Additionally, the matching capabilities are limited, the highest OSI layer [ISO, p. 28] protocol fields matchable according to the specification are TCP/UDP fields on layer 4, meaning it is not possible to monitor anything contained within TCP/UDP packets or non standard header fields used by other protocols.

An advantage of network metadata monitoring is the offloading of a part of the decoding process. In such an environment, the local load on the monitoring component is reduced, as the network already provides information preprocessed and the component itself can focus mainly on analysis of the presented information. Tying into that advantage is a lower network load, as fewer packets containing just the information provided by switches, have to be forwarded to monitoring. Additionally, control over how many packets are sent is with the network provider, as polling intervals and the number of packets gathered and relayed to monitoring can be configured by the provider and cannot be influenced directly by a client. As a summary, monitoring network metadata is receiving statistical data from the network itself and deciding based on that data if an alert is necessary.

5.1.3. Active Client Monitoring

Monitoring can also happen on the client itself. This requires control over the same and additional software in place to forward data to the monitoring component. All information about a single client is thereby implicitly grouped together.

With this method, some of the monitoring workload can be spread across the clients, so that only important information has to be forwarded to the monitoring component. Detecting anomalies involving only a single client could happen within the same, not crossing the network boundary and without any dependencies to a network. To monitor distributed information, a solution is to send its data to a centralized monitoring component.

This can happen independently from the network solution chosen, whether it is a regular network or an SDN approach, as it uses the network only as the transport media and does not rely on any other features. A drawback of active monitoring is, that in some cases, not every network component is accessible and it is not possible to install additional software on it. Should a system be compromised, the information coming from this client may not be reliable and report inaccurate information to the monitoring component.

In summary, monitoring via an active client setup is a distributed approach with the possibility to leverage processing power of clients to analyze only preprocessed data centrally.

5.1.4. Comparison

Each method of data collection has its advantages and drawbacks, which shall be compared in a short summary in Table 5.1 .

5. Monitoring

	Pro	Contra
Packet capture	Fine grained overview of all data in the network	Costly to analyze lots of data
Network Metadata	Overview about data flowing through the network	Only minimal data content analysis possible
Active Client Monitoring	All information from a specific source combined	Every client has to be setup separately

Table 5.1.: Comparison of monitoring methods

5.2. Data Analysis

The decision whether to alert on a set of given data or not do so is the main responsibility and challenge of the analysis phase. Independent of any network technology, this module has to compare the presented data to values the data should have. The data is already decoded and there are multiple ways to compute a range of data in which case to generate an alert. The goal of this phase has to be to generate alerts for as many happening attacks as possible while not issuing them for valid traffic.

5.2.1. Limit

A rather simple method is a fixed limit of x , converting to a range from $\{0 \dots x\}$ or an explicit set range $\{x \dots y\}$. For a number of values, a fixed range may be enough. A number of monitoring tools widely used like Nagios or Icinga provide this functionality [Tea00]. For example, if in a known network the number of IP addresses originating from a switch port is monitored, a range of zero to one IP address per port can be a reasonable limit, if this is a port known to have only one client attached to it.

5.2.2. Standard Deviation

A slightly more complex method is computing an allowed range for the next value. This can be done by computing the standard deviation over a number of last received data points n , and generate an alert, as soon as a new value is outside of the range computed by the equation from Figure 5.3. However, the technique can not work without any data or with rapidly and suddenly changing values. The data rate on a main connection link, forwarding a lot of packets, could be monitored this way. If an almost constant rate with only marginal deviations is expected, then this method can give a lower and upper bound on which to alert on.

$$top = x + 2 * \sigma_n \quad (5.1)$$

$$bottom = x - 2 * \sigma_n \quad (5.2)$$

x : Last measured value, σ : Standard deviation, n : Number of values to calculate σ over

Figure 5.3.: Range by standard deviation

5.2.3. Triple Exponential Smoothing

A third possible way is to use Triple Exponential Smoothing (TES) [Fil+13, p. 6.4.3.5]. This method combines three smoothing equations as seen in Figure 5.4 to compute a forecast of values.

$$S_t = \alpha \frac{y_t}{I_{t-L}} + (1 - \alpha)(S_{t-1} + b_{t-1}) \quad \text{OVERALL SMOOTHING} \quad (5.3)$$

$$b_t = \gamma(S_t - S_{t-1}) + (1 - \gamma)b_{t-1} \quad \text{TREND SMOOTHING} \quad (5.4)$$

$$I_t = \beta \frac{y_t}{S_t} + (1 - \beta)I_{t-L} \quad \text{SEASONAL SMOOTHING} \quad (5.5)$$

$$F_{t+m} = (S_t + mb_t)I_{t-L+m} \quad \text{FORECAST} \quad (5.6)$$

y : Observation	α : Estimated overall smoothing constant
S : Smoothed observation	β : Est. trend smoothing constant
b : Trend factor	γ : Est. seasonal smoothing constant
I : Seasonal index	t : Index denoting time period
F : Forecast at m periods ahead	

Figure 5.4.: TES. Based on [Fil+13][6.4.3.5]

A range of accepted values can then be computed in multiple ways, a simpler one would be e.g. by allowing a fixed percentage deviation as shown in Figure 5.5.

$$top = F_{t+1} * 1,1 \quad (5.7)$$

$$bottom = F_{t+1} * 0,9 \quad (5.8)$$

F : forecast at time period t

Figure 5.5.: Range by fixed percentage

The advantage of this method is that it encompasses multiple ways a value could change over time and adapts to those changes. Then again, this is the most resource intensive method presented here. A second disadvantage is, that in order to compute an initial trend required for the forecast, at least one complete season of data has to be available. Depending upon a season length, this may be quite an amount of data necessary.

5.2.4. Comparison

As each data gathering method, every analyzing technique has its advantages and disadvantages too. Table 5.2 shows a summarized comparison between the presented possibilities.

	Pro	Contra
Limit	Simple to set, fast to implement, few resources required	Rigid, unresponsive to change
Standard Deviation	Adaptive, few prerequisites	Only suitable for slowly changing values
TES	Most adaptive and comprehensive	One season of data necessary, most computing resources used

Table 5.2.: Comparison of analyzing methods

As some methods could counteract disadvantages of others, a combination of approaches is possible. For example, when using TES to monitor the bandwidth used on a connection, it may be feasible to set a maximum as a first check to alert on. The limit could alert if used capacity is higher than 95% of the link capacity and save computing resources in this case, while, if this limit is not hit, TES could still check for an acceptable bandwidth use.

5.3. Alerting

As the last step in this component, the alerting phase is responsible for communicating results. As there is only communication happening in the case of an alert, it has to fulfill the following criteria.

- Reliability
- Speed
- Lucidity

As the communication channel for alerting is empty unless there is an alert, it is crucial to receive every sent packet and not miss a single one. A reliable and proven protocol has to be chosen to ensure a solid communication path.

To keep latency low, speed is an important part too. Especially in a situation where an attack is already detected, and presumably still active, the information about this attack has to be received by the mitigation component as fast as possible. This is true for all kind of attacks, but notably in any flooding attack a malicious network client uses every chance it has got to insert packets.

A last but nonetheless important criteria is lucidity. This encompasses, that a message received by another component has to be identifiable unambiguously, must be parsable fast and contain every bit of information necessary without being overloaded. The recipient should be able to decode a message without the necessity to compute additional information in an intermediate stage. With the only input being the alert message, the recipients should be able to use the presented input in combination with its own information.

Existing software, like e.g Prelude OSS, uses Intrusion Detection Message Exchange Format (IDMEF) [Deb+07] to communicate alerting information [Sys15]. This format leverages

Extensible Markup Language (XML) via TCP as its transportation provider to communicate a range of information.

5.4. Implementation

The VM described in chapter 4 implements a selection of the depicted functionality. This chapter explains in detail the implementation concerning the monitoring component, its Linux Container (LXC) and how it interacts with its surroundings.

To keep a separation of concerns, both monitoring channels are used as independent system services and can operate without any dependency to each other. It is possible, without changes to the application logic itself, to extrude one of the services into another location and reroute its traffic without affecting the other one. A depiction of the implementation internals of the monitoring component is shown in Figure 5.6.




Figure 5.6.: Monitoring implementation

5.4.1. Packet Capture

Suricata is used to implement packet capture and analysis. It is used as an IDS, reading every packet arriving at the Network Interface Card (NIC) of the LXC. All three steps depicted in Figure 5.2 are handled by Suricata.

To get packets to the LXC, every switch in the network is, at startup, instructed to duplicate every packet and send one to its destination and the other to the monitoring instance. The subsequent decoding process is done from raw packets, making data available for the analysis phase. Suricata as a signature based IDS uses rules to decide when to alert. An example rule is shown in Listing 5.1.

```

1 alert udp 192.168.127.15 any -> $HOME_NET 1 (msg:"Malicious Host"; GID:1;
  sid:60000003; rev:002; flow: stateless; threshold: type both, track by_src,
  count 500, seconds 10;)

```

Listing 5.1: Example Rule

The rule generates an alert, if the host 192.168.127.15 sends more than 500 packets within 10 seconds from any port to any host in the 'HOME_NET' at port 1. After the generation of an alert, it is published in JSON format to the Redis instance making it available for consumption by other components.

This implementation was chosen due to its simplicity. Suricata is able to handle all phases of the monitoring component. It is easy to configure and proves itself fast enough in the attack analysis section 5.5. The output format JSON was chosen due to its wide spread use.

5. Monitoring

The output path of Redis enables this test setup to send the alert messages to more than one recipient and interfaces with the statistics Suricata is configured to produce. A drawback to this way of implementing the forwarding is, that internally the packet number is doubled. An attacker therefore can generate double the pressure to the network, although packets have different target clients. Another problem is that IP spoofing is not detected. If the attacker uses a spoofed IP, the monitoring instance generates an alert containing this false IP. In this implementation, Suricata has no information about which OVS port a packet is sourced from.

5.4.2. Network Metadata

OVS supports the collection of metadata as specified in the OpenFlow specifications. All client network switches are probed by the manager for statistics, and they deliver back values about every port and the flows installed. A sample excerpt from a statistics message as sent by a switch can be seen in Listing 5.2.

```
1      "1": {
2        "rx_packets": 1909,
3        "tx_bytes": 100862,
4        "tx_errors": 0,
5        "properties": [
6          {
7            "OFPPortStatsPropEthernet": {
8              "collisions": 0,
9              "rx_crc_err": 0,
10             "type": 0,
11             "rx_frame_err": 0,
12             "rx_over_err": 0,
13             "length": 40
14           }
15         }
16       ],
17       "tx_packets_delta": 0,
18       "rx_dropped": 0,
19       "rx_errors_delta": 0,
20       "length": 120,
21       "tx_bytes_delta": 0,
22       "tx_dropped_delta": 0,
23       "duration_sec": 500,
24       "tx_errors_delta": 0,
25       "rx_bytes_delta": 0,
26       "duration_nsec": 212000000,
27       "rx_dropped_delta": 0,
28       "rx_errors": 0,
29       "rx_packets_delta": 0,
30       "rx_bytes": 5999847,
31       "tx_packets": 1644,
32       "tx_dropped": 0
33     },
```

Listing 5.2: Sample port statistics

In case of this implementation the controller additionally calculates delta-values to previous fetched statistics as seen in lines 17 to 24 at Listing 5.3. As these values are already available inside the NOS, it is easier to calculate them within the controller and then pass them on to the monitoring instance to easily analyze differences in historical values.

```

1      # Default value for mitigation flows
2      table_id = 0
3      if REST_COOKIE in data:
4          cookie = data[REST_COOKIE]
5      else:
6          cookie = 0
7      if REST_COOKIEMASK in data:
8          cookie_mask = data[REST_COOKIEMASK]
9      else:
10         cookie_mask = 0
11     if REST_IDLETIMEOUT in data:
12         idle_timeout = data[REST_IDLETIMEOUT]
13     else:
14         idle_timeout = 0
15     if REST_HARDTIMEOUT in data:
16         hard_timeout = data[REST_HARDTIMEOUT]
17     else:
18         hard_timeout = 0
19     if cmd == 'OFPPC_DELETE':
20         self.ofctl.delete_flow(priority=priority, table_id=table_id,
match=match, command=command_obj,
21                                 cookie=cookie, cookie_mask=cookie_mask)
22     else:
23         self.ofctl.add_flow(priority=priority, table_id=table_id, match=match,
action_list=action_list,
24                                 hard_timeout=hard_timeout, idle_timeout=idle_timeout,
command=command_obj,
25                                 cookie=cookie, cookie_mask=cookie_mask)
26         self.logger.info(
27             "Added mitigation flow: {}{}{}{}{}{}".format(priority, table_id, match,
action_list,
28                                                         hard_timeout, idle_timeout), extra=self.sw_id)
29         details = 'Performed {}s' % cmd

```

Listing 5.3: Statistics gathering

The analysis phase is then executed by the Statshandler running within the monitoring LXC. There, thresholds are implemented that get tested against received values. If a value surpasses a limit, an alert is published. This is done via Redis in JSON format again for the same reasons mentioned before. As there was no software available solving this exact problem, the implementation was done manually. This allowed for the most flexibility and a seamless integration into the application stack. A second drawback is a time issue. The Statshandler component does only get data every 5 seconds. This means, in a worst case, the delay between the start of an attack and its detection is at least 10 seconds. Compared to an IDS, which gets packets as fast as the network can forward them, this is slower.

5.5. Attack Analysis

The following sections will analyze three different attacks and for each one multiple ways to detect the it. Each of the methods will have its advantages and drawbacks highlighted to be able to compare them against each other.

All attacks implemented are performed on a freshly set up VM. Inside of it, they will be launched and monitoring data will be gathered. This data is used as a source for shown figures and to compare methods.

5. Monitoring

5.5.1. SYN Flood

Starting simple, the first example looked at is a DoS attack originating from a single host. After connecting to the LXC 'attackone', it can be started with the command shown in Listing 5.4.

```
1 /vagrant/attacker/ping-flood/SynFlood.py -s 10.10.10.3 -t 10.10.10.4 -n 5000000
  -p r -P 80
```

Listing 5.4: Starting a DoS attack

Suricata

Figure 5.7 shows an unmitigated dos attack originating from a single host machine, attackone in the test setup. The attack was performed with 5 000 000 packets sent in 72.30 seconds ($\approx 69\,000pps$) by the attacker. In this test, the mitigation component was shut down, to allow the full attack to spread through the network.

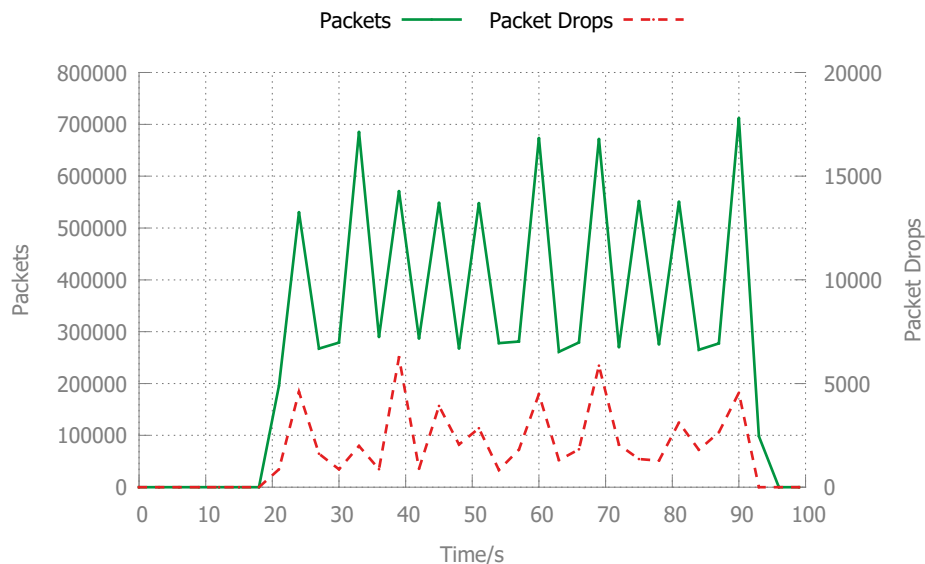


Figure 5.7.: DoS attack without any mitigation

If the attack stays unmitigated, Suricata starts losing packets in its analysis phase. The monitoring instance is flooded with packets and starts dropping them in order to analyze new ones as can be seen in Figure 5.7. Detecting a SYN flood is possible using Suricata. The goal in catching the attack is, to identify the traffic and host from which it originates. The rule used to achieve this is shown in Listing 5.5. In summary, this rule generates an alert, if the following conditions are met: There where more than 500 packets from the same source IP address within the last 10s having any source and destination IP, any TCP ports and their TCP SYN flag set. In case all conditions are met, an alert with the id 10 000 003 is generated.

```
1 alert tcp $HOME_NET any -> $HOME_NET any (flags: S; msg:"Syn_flood_detected";
  SID:1; sid:10000003; rev:001; flow: stateless;threshold: type both, track
  by_src, count 500, seconds 10;)
```

Listing 5.5: DoS Rule attack

With a rule like in Listing 5.5, it is possible to pin a single attacker precisely. The alert generated contains every identifying value of the attack. These are source ip, destination ip, TCP source port and TCP destination port, in Listing 6.2, a full alert message can be seen. The method using a simple limit is sufficient for this set up. Regular clients pulling a test page from the web server running on the 'target' LXC do not surpass this limit. Though it is not adaptive, it is enough as the regular traffic is known.

Statshandler

A second method in detecting a DoS attack is using statistical analysis. The Statshandler service within the instance is monitoring all switch ports. This allows the detection of an unusual behavior of client ports, but gives no insight in what traffic is sent. As seen in Listing 5.2, a number of different metrics are available, but none of them contain any information about what was transmitted other than the number of valid Ethernet frames. To get this information, a flow based analysis has to be used. The flow has to match the attack characteristics as close as possible. It would resemble the Suricata rule loosely without containing information about the rate of packets. A possible implementation would be, to leverage multiple tables in the OVSs. Then, the first table can contain matching rules, some catching every SYN packet per port and one, matching every other packet. All rules forward the packets to table two, which outputs all packets to their destinations. This setup would allow counting all TCP SYN packets within one flow per port. An overview to this setup is shown in Figure 5.8.




Figure 5.8.: Flow of alert messages

With this information present, the statistic analyzer can compare the number of SYN packets on a single port to its logic and, if necessary, generate an alert containing the source OVS port and the information that the attack is a SYN flood. The alert therefore does not contain every information about the attack. TCP source, TCP destination port, and both source IP and destination IP are missing. This is an advantage compared to a solution monitoring ports only, as the link between a heightened packet count and TCP SYN packets is present. Compared to using Suricata, this can not detect the target IP or the TCP destination port used. These informations are used by mitigation in subsection 6.2.2 to generate an alert.

5.5.2. Distributed Syn Flood

A distribution across multiple hosts brings more challenges to the network and to the monitoring component. The main difference for catching such a distributed attack is the missing characteristics of a fixed source IP address and a single OVS source port.

Suricata

In the setup present, Suricata is fast enough to issue multiple alerts in the same way as done for a single source SYN flood, as there are only a few hosts. But as soon as these attackers are using spoofed IPs, this is equal to a lot more hosts, at least in the informations available to Suricata. Additionally, for the IDS to detect the attack, a threshold of 500 packets per 10 seconds has to be surpassed. Using the rule shown in Listing 5.5, this may never be the case if attackers never surpasses the limit with one of the spoofed IPs. Additionally, this generates a lot of load on Suricata, as it tries to count SYN packets issued for every IP individually. To circumvent this problem, the rule shown in Listing 5.6 is used. The two changes compared to Listing 5.5 are 'sid' and 'track by_dst'. This instructs Suricata to count sent SYN packets for every IP destination instead of every IP source address.

```
1 alert tcp $HOME_NET any -> $HOME_NET any (flags: S; msg:"Syn_flood_detected";
  GID:1; sid:10000004; rev:001; flow: stateless;threshold: type both, track
  by_dst, count 500, seconds 10;)
```

Listing 5.6: DDoS Suricata Rule

The effects on Central Processing Unit (CPU) usage and an unmitigated attack using spoofed IPs is shown in Figure 5.9. In both cases, using `track by_src` and `track by_dst`, the CPU usage increases to 100%, shifted in time as Suricata processes packets. As the IDS reaches full load, it starts to drop packets and is not able to keep up with the flood of information. During the process, packet count drops down to 0, as the monitoring instance cannot handle the load and produces inaccurate packet counts. After the attack ends, CPU usage levels drop back to normal again.


The important difference is that when using `track by_dst`, Suricata is fast enough to detect the attack and issue multiple alerts before it starts to drop packets and run out of computing resources. In case of this implementation, Suricata therefore is fast enough to detect the attack multiple times. This can be seen in Figure 5.10.

In summary, Suricata is capable of detecting the attack, but cannot keep up with it for a large amount of time.

Statshandler

One way to solve the problems of Suricata is using the Statshandler component. As it receives its information from the OVS infrastructure, it doesn't have to decode all packets first. Presented with a statistics summary, the decoding process is limited to parsing the JSON message and it can start its analysis phase right away. In Figure 5.11 an attack executed with the same parameters as in Figure 5.9 is shown, but without Suricata and the mitigation stage running.

This shows, that Statshandler can generate alerts more reliable than Suricata in case of a distributed SYN flood attack with spoofed IPs. The instance is not involved with the packets transmitted directly and can act through the attack constantly. It is slower though,

Figure 5.9.: SYN Flood with `track by_dst`

as it has to wait for switch statistics to arrive while Suricata is able to react as soon as viable packets arrive.

The implementation uses a similar fixed limit as Suricata does. An alert is generated, if more than 500 packets are received and transmitted on a port within the 5 seconds time window after which new statistics are received from the controller. To detect ports that insert many packets into the network, this limit is sufficient. Even if not only IPs are spoofed and there truly are multiple clients inserting packets, the Statshandler relays more than one alert with a limit-surpassing port.

In this use case enough, it bears some problems. As the Statshandler can not differentiate between the type of packets, it is possible that if multiple ports send many packets, the Statshandler will not only report the ports sending SYN packets, it will also state all other ones.

Combination

To achieve the goal of detecting attacking hosts precisely, the two methods stated earlier can be combined. If Suricata and the Statshandler run in parallel, information from both services is available. Combining the destination IP, destination TCP port and set TCP SYN flag from Suricata and source OVS port from Statshandler narrows the packets identified as part of the attack further down. It still is possible that legitimate traffic is blocked if it is inserted into the network at the same port as used by an attacker and its target is also the attacked host. But communication flowing from a non-involved source, spiking network traffic at a OVS port to to be forwarded to a different destination will not be blocked. Combining both approaches is therefore the best solution, adding up the advantages from both methods.

5.5.3. Flow Table Flooding

As a, in contrast to DoS and DDoS, unique attack to SDN, flow flooding is more challenging than other attacks. The following sections will show that Suricata has its difficulties in

5. Monitoring




Figure 5.10.: SYN Flood with alerts

dealing with the attack, due to its SDN nature. The Statshandler can work more efficiently due to its capability of interfacing with the information available in the SDN infrastructure.

Suricata

As Suricata can not access the information available in the SDN directly, it can perform two classes of actions. One approach, is to inspect every packet in the network for content that possibly results in more flows. On one hand this depends heavily on the implementation of how flows are generated and on the other hand on how this implementation is exploited. If a common denominator can be found to identify packets generating flows, this characteristic can be searched for and an alert can be generated if suspicious behavior is detected. As in this implementation, the attacker and its partner are sending packets back and forth having the content `ping` and `pong`, Suricata can match for this package content and issue an alert as soon as these packets are detected. This is achieved with the rules shown in Listing 5.7.

```
1 alert tcp any any -> any any (msg:"Flow_flood_detected"; GID:1; sid:10000006;  
  rev:001; content:"|70_69_6E_67|");  
2 alert tcp any any -> any any (msg:"Flow_flood_detected"; GID:1; sid:10000007;  
  rev:001; content:"|70_6F_6E_67|");
```

Listing 5.7: Flow flood rule

This method relies solely on packet content and a known signature of the attack. If this is the case, it can generate alerts as soon as a matching packet is detected. In this implementation, this can be as soon as the first packet containing `ping` arrives at the monitoring instance.

A second way of detecting a flow flood with an IDS would be by watching for missing traffic. If a known host has a constant amount of traffic and manipulation is happening inside the flow tables, communication to this host may break down. This difference in packets transmitted can be seen by an IDS with the rule proposed in Listing 5.8. As Suricata is currently missing a feature for matching below instead of above a threshold in its rule set,



Figure 5.11.: SYN Flood with Statshandler

this is not currently implemented.

```
1 alert tcp $HOME_NET any -> $HOME_NET any (msg:"Flow_flood_detected"; GID:1;
  sid:10000008; rev:001; threshold: type both, track by_dst, count !500,
  seconds 10;)
```

Listing 5.8: Flow flood rule

In summary, Suricata is barely able to detect a flow flooding attack. It is missing the capabilities to measure values directly linked to the attack or has to rely on detecting a very specific and already known signature.

Statshandler

Having access to the internal statistics of the SDN will be a key advantage of Statshandler to detect flow flooding as shown in this section. As with every statistics message sent from the manager to monitoring, a list of installed flows and their statistics is provided, the main value used in this implementation is the number of flows installed to a switch. In case of this rather static test setup, the number of active flows is not varying more than by a few. Therefore, a simple limit on how many more flows are active during the last time period is sufficient. In ??, a flow flooding attack can be seen with alerts generated by the statistic component.

This implementation uses the `OFFAggregateStats` provided by Ryu to receive a compact representation of active flows. Sorted by switch and port, this allows easy detection of maliciously used OVS ports. An example part of such a statistic message received is shown in Listing 5.9

5. Monitoring

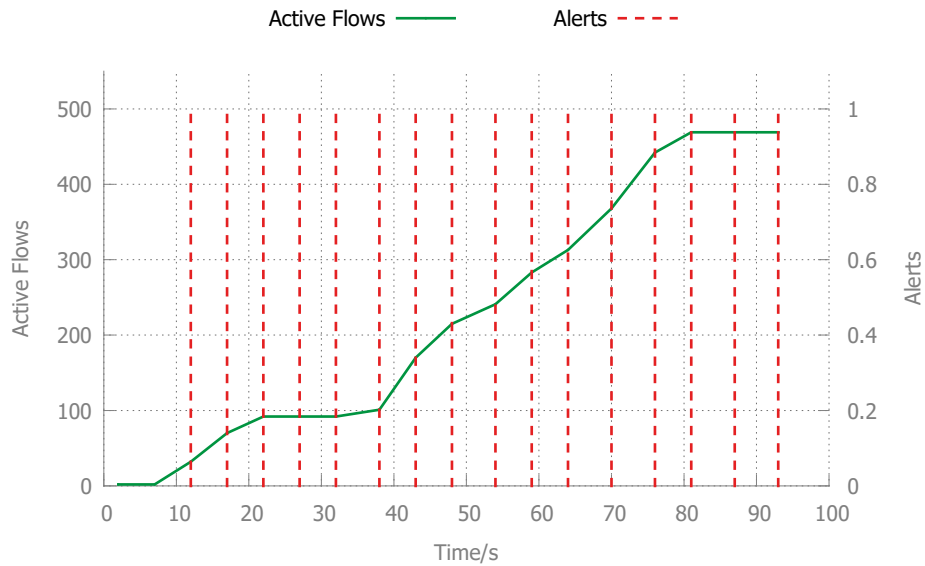


Figure 5.12.: Flow flood with Statshandler

```
1 "port_aggregate": {
2   "0000362b2e9f7f44": {
3     "1": [
4       {
5         "OFPAggregateStatsReply": {
6           "flags": 0,
7           "type": 2,
8           "body": {
9             "OFPAggregateStats": {
10              "byte_count": 0,
11              "flow_count": 0,
12              "packet_count": 0
13            }
14          }
15        }
16      ]
17    }
18  }
19 }
```

Listing 5.9: Port aggregate statistics message

In comparison to Suricata, this is a superior approach. The Statshandler uses fewer resources because there is no package analysis done. Therefore no packages have to be rerouted, the SDN controller only requests statistical information every time interval. There is no dependency to an attack implementation, as soon as a port misbehaves, it is detected. Additionally, the OVS port is specified precisely, therefore even techniques like IP spoofing don't impact the detection. In summary, the uncovering of attacks targeted at SDN specific components requires methods specifically built for this purpose, which Statshandler is.

5.6. Monitoring conclusion

Though the concept of SDN heavily influences every phase of the monitoring process, the single largest contact point of this component with the network is its gathering phase. The way data flows to the component and the type of data that is available to it is specific to a Software Defined Network (SDN) and sufficient to monitor the network. Data present in regular networks like packets per port are also available.

Additionally, SDN provides additional data to give more insight on what paths are used in the network. Both of these are available from a single access point or the single point can be configured to sent information to the monitoring instance. This ease of use is an advantage compared to regular setups.

In its analysis phase, proven techniques can be used to analyze the available standard information and new data can be used to detect anomalies. IDSs already in use still have their place in the network, providing additional information and further analysis of packets. As can be seen with the implemented attacks, the component can analyze and detect more traditional attacks like DDoS as well as new and SDN specific ones like flow flooding.

To output results, the SDN provides a solid foundation and a reliable and configurable way to get the information the next component, the mitigation instance.

6. Mitigation

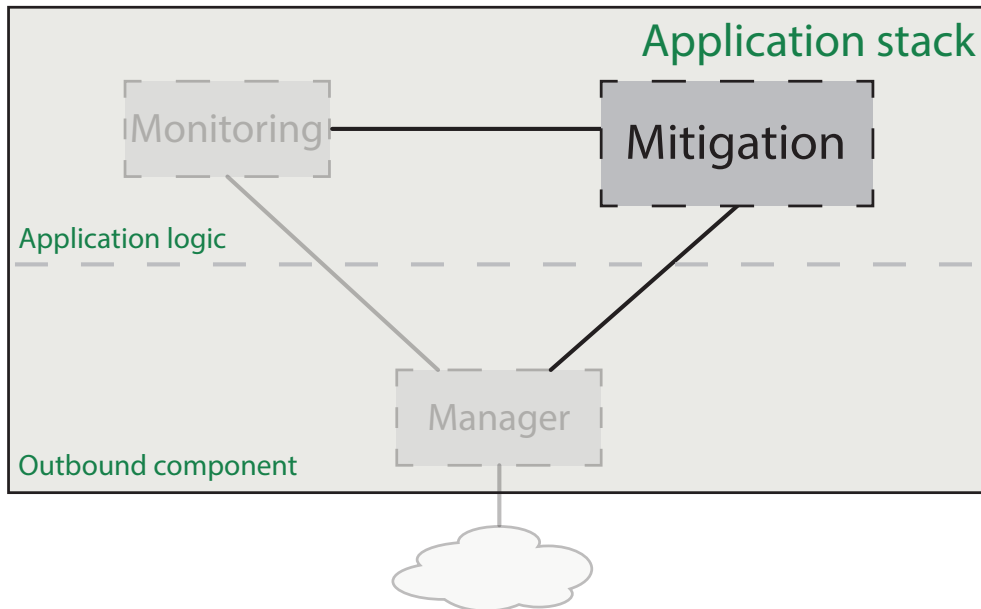


Figure 6.1.: Mitigation in the concept

For mitigation there exist two very different approaches: A proactive mitigation strategy tries to reduce the probability of a successful DDoS attack, by employing techniques like Ingress filtering to preemptively block an attack. On the other hand there is a reactive mitigation strategy that tries to detect an ongoing attack and actively mitigate that particular attack.[DBP05]

Flows are a great way to mitigate attacks. However proactively blocking traffic in a network is for most attacks just not feasible. Especially when considering that a network will most of the time be operating normally. Therefore, as the goal was to explore the capabilities of SDN, the proposed mitigation component uses a reactive approach that can take full advantage of SDN.

As shown in figure 6.1 it fits conceptually between the Monitoring and Management component.

6.1. Mitigation concept

Figure 6.2 shows the process of the mitigation.

The first phase of a successful reactive mitigation, is a trigger that starts the mitigation process and provides the information necessary to handle the ongoing attack. A central database manages all the alerts generated in the network and an Alert Handler, as introduced

6. Mitigation




Figure 6.2.: Concept of the mitigation component

in subsection 6.1.2, subscribes to this database to receive the alerts of interest. These alerts serve as the trigger, as well as the source for all the information needed.

Once such an alert is received, the second phase starts. Based on the information provided the appropriate countermeasures have to be chosen from the set of implemented mitigation strategies. In section 6.3 a selection of such mitigation strategies is introduced and implemented.

Finally in the last stage the countermeasures that are implemented via instructions for the creation of flow entries are sent to the controller, who based on these instructions will implement the flow entries on the required devices. In subsection 6.1.4, this is introduced in more detail.

6.1.1. Alerts

First and foremost an alert must be assignable to a problem or an attack in the network. To this end a unique ID, called the Signature ID, is embedded into every alert. Additionally the system that has issued the alert is added. To reliably reconstruct when the alert was raised a timestamp is included as well.

With just this information it is already clearly described what the issue is. However, it is yet unclear where the problem is and who is involved. Based on the Signature ID and source of the alert, there is a set of information the Mitigation module expects to be informed about. This can be information such as the ID of the switch, where an issue has been detected, the IP Address of the victim or the attacker, the port of the switch under attack, or the TCP port that is targeted.

6.1.2. Alert Handler

The Alert Handler is a module that takes care of receiving the alerts for the Mitigation module.

The alerts are divided into multiple channels. As such for example the Intrusion Detection Systems publishes its alerts to a different channel than the system doing statistical monitoring. For each channel the Mitigation module wants to receive the alerts from an instance of this module is created. This instance is responsible for creating a connection to the database and subscribe to the alert channel. Now every alert that is published to the subscribed alert channel is also sent to that instance of the Alert Handler.

After receiving a message it is checked, if it is an alert, and the alert is forwarded to the mitigation module.

6.1.3. Mitigation

In the mitigation module all the alerts coming from the various Alert Handlers are handled.

The Signature ID is extracted from every alert and looked up in the set of available mitigation strategies. Naturally, as this module is specialized in handling attacks, based on a predefined mitigation strategy, only alerts with a known Signature ID can be handled.

If a mitigation strategy is available further actions depend on this strategy. It can be that the information from a single alert is enough to immediately deploy countermeasures, but the information provided from a single alert might also be insufficient for that type of attack. In this case the information contained in the alert is stored together with the timestamp. When the alerts containing the missing information, potentially from different systems and alert channels, are received, the time difference between the arrival of the alerts is calculated. To assure that these alerts are related to the same event only if this falls within the specified range of tolerance the countermeasures are constructed. The countermeasures consist of instructions for the creation or deletion of flows.

The first step to add flows is to create the match fields. This specifies the criteria for which packets to act upon. The values for these fields are mainly supplied by the information gathered from the alerts. Common criteria are related to the target of the attack or the source of the attack, such as the MAC Address or IP Address, or the origin of the attack in the network, such as the port of the switch.

The next step is to define the action of the flow. For example the packets can be dropped immediately or sent to another system for further inspection. Next additional parameters can be specified, such as the soft or hard timeout. A soft timeout deletes the flow automatically after a set time of inactivity, while a hard timeout automatically deletes the flow after the set time, independent of activity. The priority of the flow is another parameter of interest. Usually the highest priority is desired, as this assures that the mitigation flow is the first to match the packet. If it was the case that a flow with a higher priority matches the packet first, the packet would be handled by that flow and would never match the mitigation flow, rendering it completely ineffective.

Finally the switches, the flow should be installed on, are specified. This can either be the ID of a switch or a keyword to specify a set of switches, such as every switch or the switches on the perimeter of the network.

These instructions are then handed over to the Northbound API of the controller.

6.1.4. Controller Northbound API

The Northbound API of the controller provides an API for Network Applications to interact with the controller. As such it also enables the Mitigation module to implement flows on the switches.

To install flow entries on the flow tables the instructions provided have to be translated into OpenFlow commands that can be sent out on the Southbound API of the controller. As for the mitigation of the attacks a powerful API is necessary the translation should be as generic as possible. To simplify the creation of the instructions default values for parameters that are not explicitly set are defined here.

Also the keywords for the set of switches, the flow entries are to be installed on, have to be translated into the corresponding set of actual switches.

6.2. Implementation

This chapter shows how the concept introduced above is realized.

The mitigation component is run as a service on a LXC that is dedicated to mitigation and the modules are implemented fully in python and run as a system service.

6.2.1. Alert Handler

The central database that manages the alerts is Redis that was introduced in subsection 4.4.8.

```
1 [Unit]
2 Description=Alerthandling service
3 After=syslog.target network.target
4
5 [Service]
6 Type=simple
7 ExecStart=/usr/bin/python3 /vagrant/mitigation/AlertHandler.py -r 10.20.0.6 -k
   switch-stats
8
9 [Install]
10 WantedBy=multi-user.target
```

Listing 6.1: Alert Handler Service Configuration File

For every instance of the alert handler a configuration file, as shown in listing 6.1, is provided. Line 7 shows the command line arguments that are used for the configuration. "-r 10.20.0.6" specifies the IP Address of the Redis server to connect to and "-k switch-stats" the name of the alert channel to subscribe to. Currently there are two channels used and as such two instances of this module are started. One for the "switch-stats" channel that provides the alerts from the statistical monitoring and the other for the "suricata" channel that provides alerts from the Intrusion Detection System Suricata.

Through the subscribe/publish method used by Redis alerts, sent to one of the channels, are immediately passed to the subscribed Alert Handler. Afterwards they are sent via a RESTful API to the Mitigation module.

The full source code for this module can be found in section B.2.

6.2.2. Mitigation

The mitigation module sets up a HTTP Server that receives the alerts from the Alert Handler.

```
1 {'host': 'monitoring', 'event_type': 'alert',
2  'timestamp': '2016-11-26T13:05:25.720206+0100',
3  'alert': {'signature': 'Syn_flood_detected', 'severity': 3, 'gid': 1, 'rev': 1,
4           'signature_id': 10000003},
5  'proto': 'TCP',
6  'src_ip': '10.10.10.3', 'src_port': 1234,
7  'dest_ip': '10.10.10.4', 'dest_port': 21}
```

Listing 6.2: Extract from an alert generated by Suricata

Listing 6.2 shows an alert generated by Suricata. As the alerts are in JSON format they are first translated into a python dictionary.

From that dictionary the Signature ID of the alert is extracted and the appropriate mitigation strategy is found and finally the flow instructions are created.


```

1 # Create action: Output packet on a port
2 action['cmd'] = 'OFPActionOutput'
3 # Paramaters of the action
4 params['port'] = 2
5 action['params'] = params
6 # Add action to the list of actions
7 actions['action'] = action
8 # parameters of the OFPMatch object
9 match['in_port'] = 1
10 match['eth_dst'] = '00:00:00:00:00:02'
11 # Put dictionary together
12 data['cmd'] = 'OFPPFC_ADD'
13 data['actions'] = actions
14 data['match'] = match
15 # Specify additional settings
16 data['priority'] = 1
17 data['table_id'] = 0
18 # Send to controller for all switches
19 send_post(data, 'http://10.20.0.8:8080/switch/all')

```

Listing 6.3: Flow installation on the Northbound API

In listing 6.3 the creation of a flow, such as used in MAC-Learning, via the Northbound API is depicted. This flow matches all packets arriving at port 1 with the destination MAC-Address "00:00:00:00:00:02" and forwards them to port 2. Also the command to add the flow "OFPPFC_ADD", the priority "1" and the number of the flow table "0" is specified. Finally it is sent to the controller via a RESTful API. To do this the dictionary from above is translated into JSON and sent via a POST to the URL "http://10.20.0.8:8080/switch/all", with 10.20.0.8 being the IP Address of the controller. If it shall only be installed on a specific switch the "all" keyword has to be replaced with the ID of the switch.

The full source code for this module can be found in section B.3.

6.2.3. Controller Northbound API

On the controller the keyword or ID is resolved into the set of switches the instructions are to be installed on. As the controller keeps a dictionary with all the switches in the network it is just a matter of looking up the ID or in the case of the "all" keyword to just use every switch in the dictionary. As more advanced keywords were not needed for the mitigation strategies presented later on the "all" keyword is the only one implemented.

```

1 actions = [parser.OFPActionOutput(2)]
2 instruction = [parser.OFPInstructionActions(ofproto.OFPIT_APPLY_ACTIONS, actions)]
3 match = parser.OFPMatch(in_port=1, eth_dst=00:00:00:00:00:02)
4 mod = parser.OFPFlowMod(datapath=self.dp, priority=1, table_id=0, match=match,
5   instructions=instruction)
6 self.dp.send_msg(mod)

```

Listing 6.4: Flow installation on the Southbound API

Then these instruction are parsed into OpenFlow commands. For example listing 6.3 is translated to the commands shown in listing 6.4.

Again a packet that arrives on port 1 with the MAC Address "00:00:00:00:00:02" as destination shall be forwarded to port 2. To achieve this a list of actions is created. In this case the only action is "OFPActionOutput(2)": Send the packet out on port 2. Then this

6. Mitigation

list of actions is passed to an instruction function that causes these actions to be applied immediately once the flow is matched.

To configure what the flow matches to a `OFPMatch(in_port=1, eth_dst=00:00:00:00:00:02)` object is created: Match every packet that arrives on port 1 with the destination MAC Address 00:00:00:00:00:02.

Finally a `FlowModification` object is created and information, such as the priority of the flow and which table to insert the flow into, is added. Here the default values for MAC Learning flows are chosen.

This object is then sent to the switch, where the flow will be inserted into the flow table.

For the interested reader: Appendix A line 379 shows the method used to parse the dictionary into the OpenFlow commands.

6.3. Mitigation Strategies

In the following the mitigation of three Denial of Service attacks is explored and evaluated. As Denial of Service attacks are one of the main security concerns in any network and are sufficient to demonstrate the mitigation concept introduced, the mitigation of other attacks is left open for future work.

First the concept for the mitigation of the attack is introduced, then the concept is implemented and finally the effectiveness is evaluated.

6.3.1. TCP SYN flood

A technique unique to SDN is the use of flows to apply reactive packet filtering. While a traditional firewall can filter packets as well, it has to be placed on the ingress point of the network. As a result it cannot filter any traffic originating from within that network, as those packets do not pass the firewall. As such it is completely useless against attackers with physical access to the network.[Hu+14] With SDN every switch or router can be used to enforce the packet filtering rules, eliminating any entry points that remain unguarded.

A simple TCP SYN flood originates from only one attacker and can be stopped quite easily. Once the attack has been detected and the IP Address of the attacker has been identified a rule to drop these packets is enough to stop it.[Edd06]

In a SDN environment this is achieved, by setting up a flow with a higher priority than the routing flows, to make sure this flow is the first to get matched and the packet does not get forwarded beforehand. The matching criteria are set to any TCP packet with the source IP Address of the attacker and the destination IP Address of the victim. As action dropping the packet is set.

Implementation of the mitigation strategy

Figure 6.3 shows the flow of events from when a TCP SYN Flood attack is launched against the network until it is mitigated.

Once Monitoring has detected the attack an alert from Suricata, such as the one shown in listing 6.2, is sent to the Mitigation module. This alert contains the information that there is a TCP SYN Flood going on and the IP Address of the attacker and the victim.

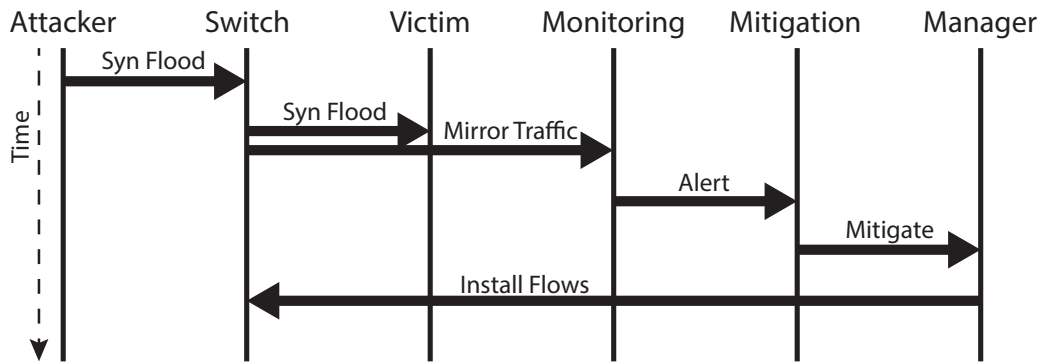


Figure 6.3.: Process of a TCP SYN Flood Mitigation




Figure 6.4.: Mitigation flow for a TCP SYN Flood

As this is already enough information to mitigate this attack a flow entry, such as shown in figure 6.4 that stops the attack is assembled immediately and sent to the controller, who propagates it to all the switches in the test setup.

In order to remove the flow again and not keep the flow in the flow table forever, potentially even accumulating too many flow entries at some point, a hard timeout is added. This value specifies the amount of time after which the flow is removed from the switch again automatically. Another reason for not keeping the flow forever is that the IP Address might be reassigned and thus it might end up blocking legitimate traffic. For test purposes a hard timeout of 20 seconds was chosen. But in a real environment a longer timeout is probably more beneficial.

Once the flow has been implemented the first switch of the test setup, the packets arrive at, drops all the packets of the attacker and the SYN Flood packets no longer arrive at the victim.

Evaluation

To evaluate the mitigation strategy a TCP SYN Flood script is started on the attacker LXC with the target LXC as destination. During the attack there is no additional traffic generated in the network.

In order to get data, to compare the performance of the mitigation strategy to, the mitigation module is disabled first. Afterwards the same attack is started once again. But this time the mitigation module is enabled.

Figure 6.5 shows the amount of packets arriving at the Monitoring component during the attack. As the resolution of the packet logging is 3 seconds the graph shows the amount of packets that arrived in the last 3 seconds.

The curve with active mitigation shows clearly that after a short burst of packets arriving the attack is detected and successfully mitigated, as all the packets are dropped.

Every approximately 20 seconds there is small spike, where a few packets can be detected

6. Mitigation




Figure 6.5.: Network load during a TCP SYN flood with and without mitigation

again. These spikes align perfectly with the hard timeout of the mitigation flows. Thus representing the packets arriving during the time the mitigation flow is automatically removed and the attack being detected and mitigated once more, as can be seen clearly in figure 6.6.

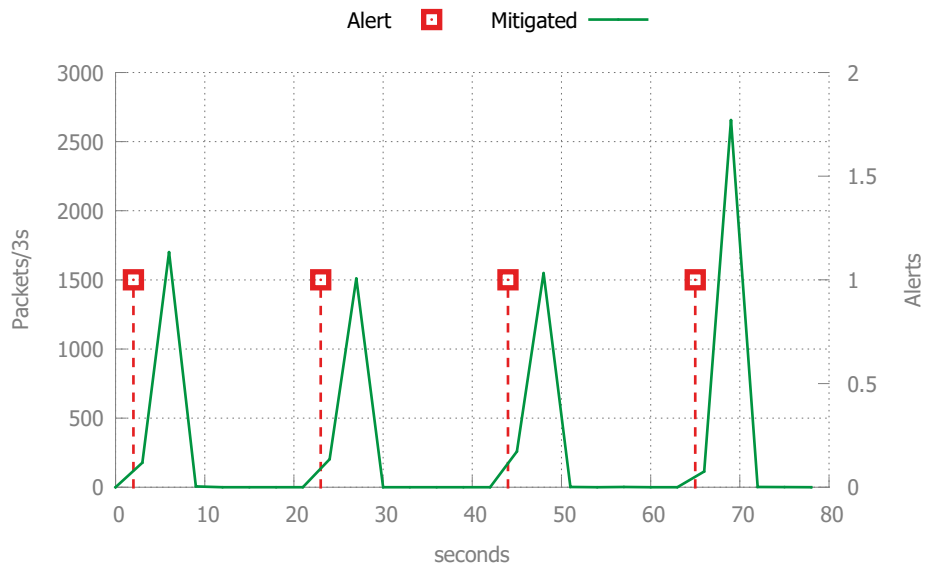


Figure 6.6.: Alerts during mitigated TCP SYN Flood

6.3.2. Distributed TCP SYN Flood

Once the attacker uses a large enough botnet, or especially if he is using IP Spoofing, to forge the source IP Address, the mitigation strategy presented above is no longer capable of defending against the attack.

With spoofed IP Addresses the alerts used previously do no longer trigger, as it relied

on many SYN packets coming from a single IP Address. But while the IDS cannot provide information about the source of the attack it can still detect the destination IP Address and the TCP ports under attack.

However, the monitoring component is able to detect the switch and port, the attack is originating from, by monitoring the statistics of the flow tables similar to the technique used in [Cui+16]. This information can be used to drop the traffic at the ingress switch. It is important to install the flow only on the ingress switch, as the flow rules are not as finely granulated, as is the case with a TCP SYN Flood originating from only a single source.

As a result installing such a flow on all devices might block a lot of legitimate traffic to the target.

Implementation of the mitigation strategy

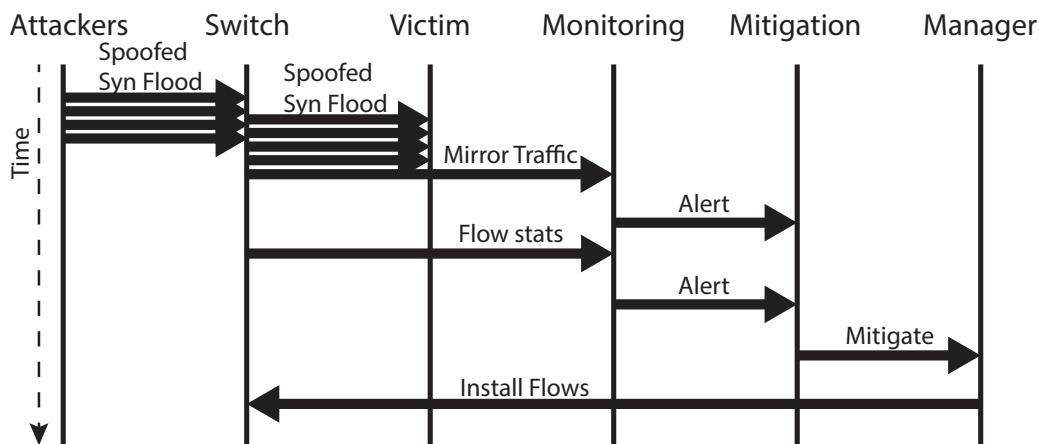


Figure 6.7.: Process of a Distributed TCP SYN Flood Mitigation

In comparison to the simple TCP SYN Flood, shown in figure 6.3, the main difference in the process of mitigating the distributed attack, depicted in figure 6.7, lies in the alerts triggering the mitigation.

The alert, received from Suricata, is pretty similar to the one displayed in listing 6.2. The only real difference is that the Signature ID now signals a Distributed TCP SYN Flood and, while there is still a source IP Address listed in the alert, it is only one of the many IP Addresses used for the attack.

```
1 {'pps': 0.0, 'event_type': 'alert', 'critical': 500, 'sensor': 'statshandler',
2  'switch': '0000ee132ef17348', 'alert': {'signature_id': '10000004'},
3  'port': '4', 'timestamp': '2016-11-30T17:27:44.418396+0100'}
```

Listing 6.5: Alert generated by statistical monitoring

The other part of the information is provided through the alert generated by the statistical monitoring. Such an alert, generated during an attack, is listed in listing 6.5. From this alert the information on which switch, identified by the unique ID "0000ee132ef17348", and on which port the attack (4) is arriving at can be extracted.

Figure 6.8 depicts the workflow for creating the Mitigation flow. Once a corresponding alert arrives the alert is saved. If an alert from Suricata and from the statistical monitoring has arrived within a predefined time frame, for this implementation a time frame of 5 seconds

6. Mitigation

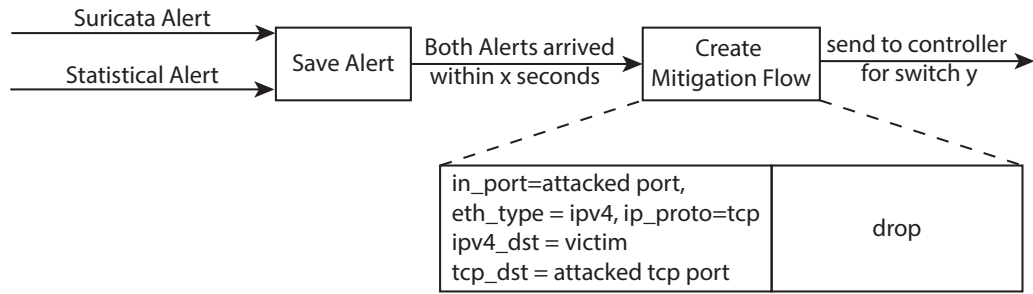


Figure 6.8.: Workflow for the mitigation of a Distributed TCP SYN Flood

was chosen, the mitigation flow is assembled from the information contained in both alerts and sent to the identified switch.

Evaluation

To simulate a Distributed TCP SYN Flood the attack is launched from 2 machines simultaneously. Additionally all the source IP Addresses of the SYN packets are spoofed. While the amount of machines used is not representative for a botnet that might be used in a real attack, given the limited resources available in a virtual machine and by using spoofed IPs, it should nevertheless be representative.

To test if the attack is able to prevent clients from establishing a TCP connections with the victim, a webserver is installed on the LXC of the victim.

```
vagrant@vagrant: ~  
AlexM@AlexM MINGW32 ~/SDN4DOS/VM (alex)  
$ vagrant ssh  
Welcome to Ubuntu 16.04.1 LTS (GNU/Linux 4.4.0-31-generic x86_64)  
* Documentation:  https://help.ubuntu.com  
* Management:    https://landscape.canonical.com  
* Support:       https://ubuntu.com/advantage  
Last login: Thu Dec  1 13:39:42 2016 from 10.0.2.2  
vagrant@vagrant:~$ curl 10.10.10.4  
curl: (7) Failed to connect to 10.10.10.4 port 80: Connection timed out  
vagrant@vagrant:~$ curl 10.10.10.4  
curl: (7) Failed to connect to 10.10.10.4 port 80: Connection timed out  
vagrant@vagrant:~$ curl 10.10.10.4  
curl: (7) Failed to connect to 10.10.10.4 port 80: Connection timed out  
vagrant@vagrant:~$
```

Figure 6.9.: Querying the webserver during an attack

Figure 6.9 shows attempts to query the webserver during the attack. None of the requests manage to successfully establish a connection, proving that the attack was successful.

Figure 6.10 shows the amount of packets originating from the attack. Without any mitigation it stays at about 10000 packets per 3 seconds. The dashed line shows the attack with the mitigation strategy discussed above implemented. As the mitigation module has to wait for the alert from the statistical monitoring, which only receives statistical information from the switches every 5 seconds, the mitigation takes a bit longer.

This can be seen better in figure 6.11. Once the attack starts, Suricata is able to detect the target of the attack very quickly and raises an alert. But the alert from the statistical alert still takes some time to arrive. The hard timeout for the mitigation flow is 20 seconds again, as can be seen by looking at the duration between the 2 peaks, where packets arrive again. As the attack has been executed from two machines that are connected to different

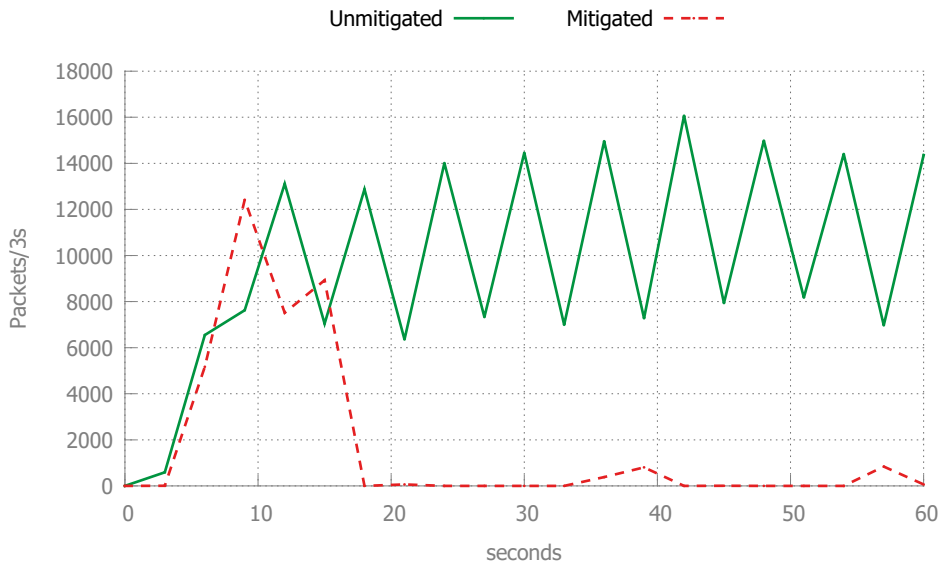


Figure 6.10.: Network load during a distributed TCP SYN flood with and without mitigation

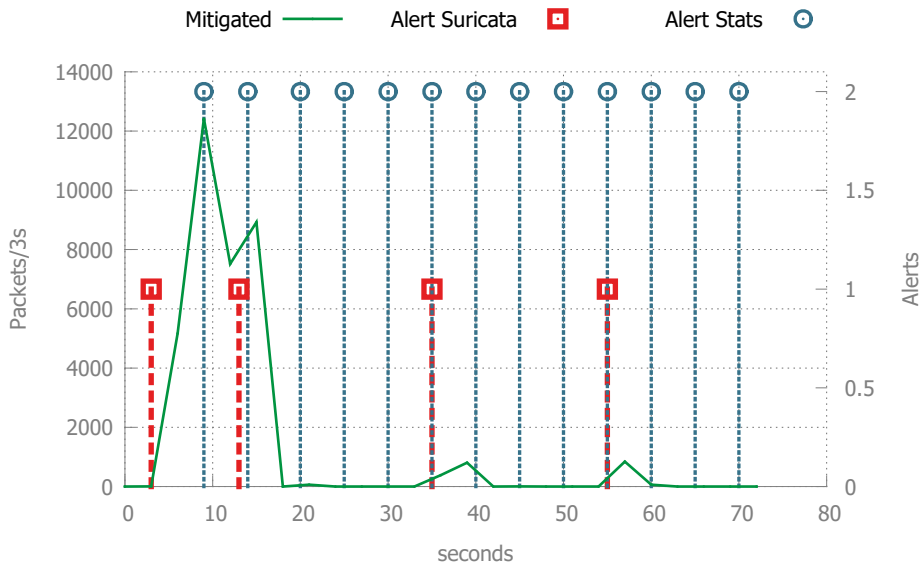


Figure 6.11.: Alerts during mitigated Distributed TCP SYN Flood

ports on the switch.

As the switch still receives packets on the attacked port, monitoring continues to generate alerts even though the packets are dropped. This has the advantage that if the attack is still ongoing, while the flow is deleted through the hard timeout, the alert coming from Suricata can immediately reinstall the flow.

It is to be noted that after the attack has been mitigated the webserver is reachable once again. Thus showing that the mitigation was effective.

6.3.3. Flow flooding

As a Denial of Service attack, this attack shares some inherent similarities with a Distributed TCP SYN Flood. To overflow the flow tables the attacker manipulates packets, usually through or combined with spoofing.

Therefore to mitigate the attack these malicious packets have to be dropped before they can cause a flow rule to be created. As most likely malicious flows have already been created before the mitigation was active these flows have to be removed from the flow table again.

The mitigation component can detect the switch under attack and determine the port the attack is originating from. However, as the attacker does not necessarily target a specific device it is difficult to narrow down the attack. As such before the attack destabilizes the whole network, a first step is to drop all packets arriving on that port. As the rules for the flow creation that are abused may vary from system to system a further analysis of abusable flow creation rules is advisable and the mitigation and detection can be custom tailored to drop only packets that could abuse these vulnerabilities. Thus reducing the amount of legitimate traffic being blocked

After closing the port all flows originating from that port are removed.

Implementation of the mitigation strategy

The detection for this attack is handled by the statistical monitoring. The alert that is received for this kind of attack contains only the information that a Flow flood was detected, the flow table of which switch is filled with malicious flows, and from which port the attack is coming from. Similar to the alert in listing 6.5.

Based on this information the flow to drop the packets is created and sent to the controller. The flow matches every packet that arrives on that port and drops it. As a hard timeout for the flow 20 seconds is chosen again. Thus effectively closing the port for 20s.

Now that the creation of additional superfluous flow entries has been halted the flow entries with `in_port="closed port"` are to be removed. To prevent the mitigation flow from being deleted a cookie has to be used.

```
1 data['cookie'] = 0x0000000000000001
2 data['cookie_mask'] = 0x000000000000000f
```

Listing 6.6: Setting cookies on flows

Listing 6.6 shows how to set the cookie. To remove the other flows instead of the `"OF-PFC_ADD"` command `"OF-PFC_DELETE"` is used and the same matching criteria are used again (`in_port="closed port"`). But the cookie `"0x0000000000000000"` is set, while the mask stays the same. This results in all flows without a cookie set and a match containing the port to be removed from the flow table.

Evaluation

First a flow flooding attack was started to take a look at the effect of the attack on the network. Figure 6.12 shows that as the maximum number of 184 flows was reached the response time went up steadily. A look at the CPU load of the controller also showed that as soon as the flow flood was started the load went up to 85%, reaching a steady 100% as the flow table was full. This effect was expected as the controller itself has to handle every

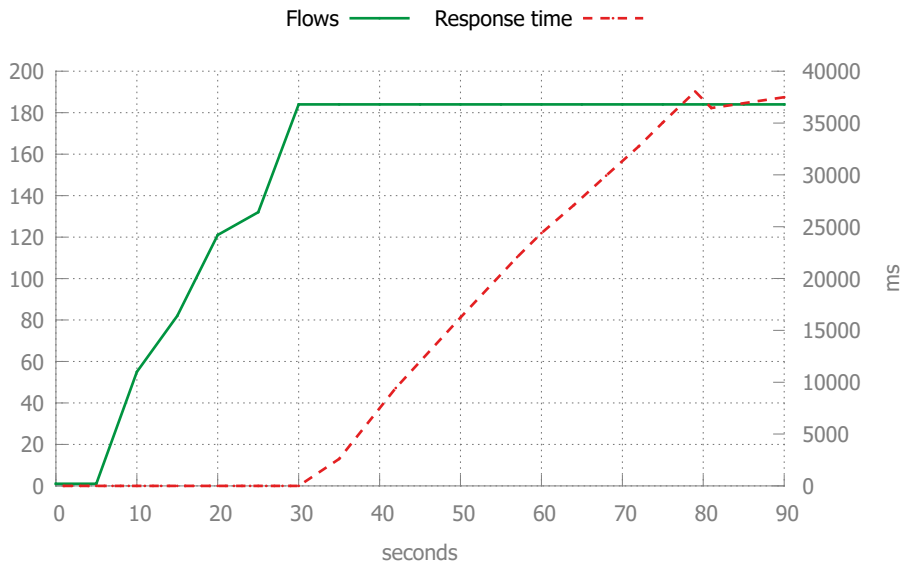


Figure 6.12.: Response time during flow flooding attack

single packet, when no flow rule can be installed. Thus eventually being overwhelmed by the amount of packets that arrive. Even without any load on the network the response time in the network went up from averaging 0.13 ms to 18 ms, showing the overhead introduced through having the controller handle the packets.

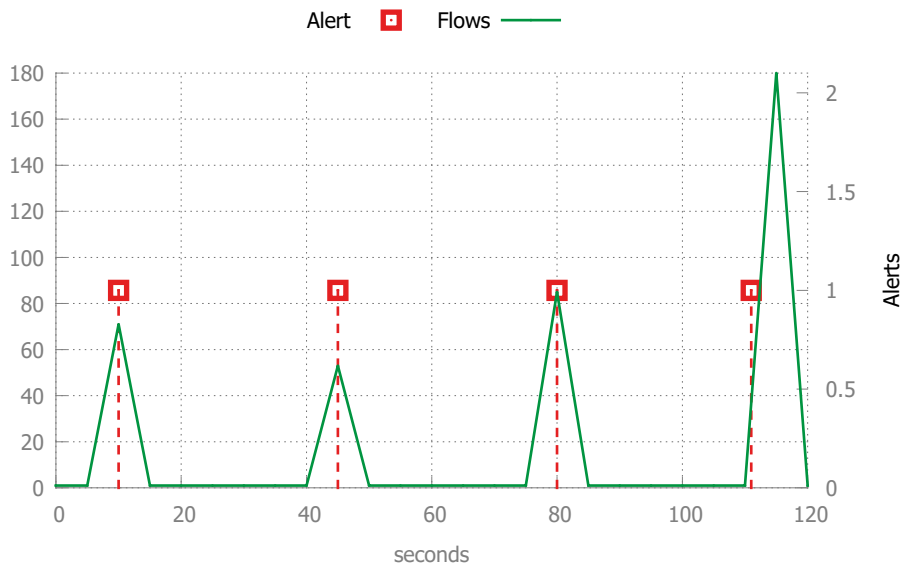


Figure 6.13.: Mitigation of a flow flooding attack

Afterwards the attack was launched again, while the mitigation strategy, presented above, was applied. Figure 6.13 shows that the generation of the flows was stopped before the flow table was even full. Only the last attack has been able to almost fill the flow table. But as the malicious flows are immediately removed again even if the attack manages to flood the table the controller functionality is restored.

6.4. Conclusion

A mitigation module was proposed that can handle various attacks based on alerts. In that context it was shown that flows serve as a powerful and versatile tool for the mitigation of attacks. Through the use of alerts and the Alert Handler, the module can be easily extended to handle alerts from various sources.

To install flows via the Northbound interface a powerful API was implemented that provides nearly the same functionality as if creating the flow instructions directly on the controller, by translating a python dictionary with a syntax closely related to the one used on the Southbound API into the respective OpenFlow commands.

The usability of the module has been shown by providing a mitigation strategy for a TCP SYN Flood, its distributed counterpart and a Flow Flooding attack. The results show that the mitigation module is able to quickly mitigate the attacks.

7. Conclusion and future work

An application stack for detection and mitigation of attacks in SDN was proposed and demonstrated. For the implementation of the application stack a test environment that can produce reproducible results was realized through a virtual machine based on a LXC architecture that is provided through vagrant. The test environment was then used to evaluate the proposed detection and mitigation mechanism.

A way of enhancing the testing environment can be to implement more caching capabilities. As a significant amount of the setup process is dependent on having access to the Internet, caching more downloaded packages can speed up the process and provide a faster to use setup. Additionally, the testing workflow could be enhanced by automating the data exported from the visualization instance. While sufficient graphs are available in Grafana to provide insight into the live status of the system, this could ease the workflow of generating high quality plots for further use. To move more into the direction of a production ready setup, security policies have to be enforced and components tested against real traffic. Parameter tuning has to be done and the network layout adapted to the new environment.

The monitoring instance used exemplifies the combination of using software already developed and originally running in a traditional network, and new one specifically developed for the analysis of SDN related data. It proves that already established ways of monitoring data can be used along with new methods to achieve detection of an attack. Both of these can provide valuable information on SDN specific and nonspecific attacks to be used together in order to defend against them.

Improving upon the monitoring component could be done in the future in two main directions. To enable it to detect more complicated and variable attacks, the analysis phase could be extended and already existing detection enhanced. Doing this can involve implementation of new logic to detect attack vectors or exchanging the stiff limit based implementation for a more flexible and adaptive one.

Though this would introduce more load to the instance, the second option could be to add an additional phase. This phase could be a lightweight component looking for any suspicious behavior and inserting a full monitoring instance into the path of traffic only after such behaviour is detected. This would improve resource usage by eliminating a lot of analysis time now spent on regular packets.

In the mitigation chapter a concept for mitigating attacks through the use of flows was introduced. This was achieved by creating flow rules from the information provided from the alerts raised by the monitoring module that block the attack. This was successfully proven on the example of three attacks.

A TCP SYN Flood attack was mitigated by installing a flow that blocks packets from the attacker on the way to the victim. While it was shown that the proposed mitigation blocks the attack, an attacker could spoof the IP Address of a system in the network to cause the mitigation to block the traffic of that system. To prevent this it might prove useful to implement source verification, such as VAVE [YBX11].

In the next step a Distributed TCP SYN Flood was tackled. Through the use of infor-

7. Conclusion and future work

mation provided by the flow statistics the ingress switch was identified and the malicious packets could be stopped, as soon as they entered the network. A known problem is that high traffic can be mistaken for a SYN flood attack. As the controller is written with support for up to OpenFlow version 1.4, matching the SYN flag in a TCP packet was not yet possible. As such supporting OpenFlow version 1.5 and matching on the SYN flag is a task that remains open.

Finally a Flow Flooding attack was mitigated. However more extensive research in regard to flow rules used in a production environment that are susceptible to this attack is to be done. This could also provide further insight on how the mitigation flows can be grained more finely to avoid blocking legitimate clients from accessing the network.

The goal set in chapter 1 to develop an SDN aware application stack has been reached and proven its capabilities while using the test setup, which was also implemented during the process. Further works can use this as a base to build upon.

A. Controller

```
1 # Copyright (C) 2013 Nippon Telegraph and Telephone Corporation.
2 #
3 # Licensed under the Apache License, Version 2.0 (the "License");
4 # you may not use this file except in compliance with the License.
5 # You may obtain a copy of the License at
6 #
7 # http://www.apache.org/licenses/LICENSE-2.0
8 #
9 # Unless required by applicable law or agreed to in writing, software
10 # distributed under the License is distributed on an "AS IS" BASIS,
11 # WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
12 # implied.
13 # See the License for the specific language governing permissions and
14 # limitations under the License.
15
16 import logging
17 import struct
18 import socket
19 import redis
20 import time
21
22 import json
23 from webob import Response
24
25 from threading import Timer
26
27 from ryu.app.wsgi import ControllerBase
28 from ryu.app.wsgi import WSGIApplication
29 from ryu.base import app_manager
30 from ryu.controller import dpset
31 from ryu.controller import ofp_event
32 from ryu.controller.handler import MAIN_DISPATCHER
33 from ryu.controller.handler import set_ev_cls
34 from ryu.exception import OFPUnknownVersion
35 from ryu.exception import RyuException
36 from ryu.lib import dpid as dpid_lib
37 from ryu.lib import hub
38 from ryu.lib import mac as mac_lib
39 from ryu.lib import addrconv
```

A. Controller

```
40 from ryu.lib.packet import packet
41 from ryu.lib.packet import arp
42 from ryu.lib.packet import ethernet
43 from ryu.lib.packet import ether_types
44 from ryu.ofproto import ether
45 from ryu.ofproto import ofproto_v1_3
46 from ryu.ofproto import ofproto_v1_4
47
48 UINT32_MAX = 0xffffffff
49 UINT64_MAX = 0xffffffffffffffff
50
51 ARP = arp.arp._name_
52
53 OFP_REPLY_TIMER = 1.0 # sec
54
55 SWITCHID_PATTERN = dpid_lib.DPID_PATTERN + r'|all'
56
57 COOKIE_DEFAULT_ID = 0
58
59 REST_COMMAND_RESULT = 'command_result'
60 REST_RESULT = 'result'
61 REST_DETAILS = 'details'
62 REST_OK = 'success'
63 REST_NG = 'failure'
64 REST_ALL = 'all'
65 REST_SWITCHID = 'switch_id'
66 REST_COMMAND = 'cmd'
67 REST_MATCH = 'match'
68 REST_COOKIE = 'cookie'
69 REST_COOKIEMASK = 'cookie_mask'
70 REST_ACTIONS = 'actions'
71 REST_ACTION = 'action'
72 REST_PARAMS = 'params'
73 REST_PRIORITY = 'priority'
74 REST_TABLEID = 'table_id'
75 REST_IDLETIMEOUT = 'idle_timeout'
76 REST_HARDTIMEOUT = 'hard_timeout'
77
78 IP_MONITORING = '10.10.10.5'
79 IP_SWITCHARP = '10.10.10.99'
80 REDIS_HOST = '10.20.0.6'
81 REDIS_PORT = 6379
82
83 DELAY_CONNECTION = 5.0 # sec
84
85 STATS_POLL_RATE = 5 # sec
86
```

```

87
88 class NotFoundError(RyuException):
89     message = 'Switch SW is not connected. : switch_id=%(switch_id)s'
90
91
92 class RestSwitchAPI(app_manager.RyuApp):
93     OFP_VERSIONS = [ofproto_v1_4.OFP_VERSION]
94
95     _CONTEXTS = {'dpset': dpset.DPSet,
96                 'wsgi': WSGIApplication}
97
98     def __init__(self, *args, **kwargs):
99         super(RestSwitchAPI, self).__init__(*args, **kwargs)
100
101         SwitchController.set_logger(self.logger)
102
103         wsgi = kwargs['wsgi']
104         self.waiters = {}
105         self.data = {'waiters': self.waiters}
106
107         mapper = wsgi.mapper
108         wsgi.registry['SwitchController'] = self.data
109         requirements = {'switch_id': SWITCHID_PATTERN}
110
111         # No vlan data
112         path = '/switch/{switch_id}'
113         mapper.connect('switch', path, controller=SwitchController,
114                       requirements=requirements,
115                       action='get_data',
116                       conditions=dict(method=['GET']))
117         mapper.connect('switch', path, controller=SwitchController,
118                       requirements=requirements,
119                       action='set_data',
120                       conditions=dict(method=['POST']))
121
122         mapper.connect('switch', path, controller=SwitchController,
123                       requirements=requirements,
124                       action='delete_data',
125                       conditions=dict(method=['DELETE']))
126
127     @set_ev_cls(dpset.EventDP, dpset.DPSET_EV_DISPATCHER)
128     def datapath_handler(self, ev):
129         if ev.enter:
130             SwitchController.register_switch(ev.dp, self.waiters)
131         else:
132             SwitchController.unregister_switch(ev.dp)
133

```

A. Controller

```
134 @set_ev_cls(ofp_event.EventOFPPacketIn, MAIN_DISPATCHER)
135 def packet_in_handler(self, ev):
136     SwitchController.packet_in_handler(ev.msg)
137
138 def _stats_reply_handler(self, ev):
139     msg = ev.msg
140     dp = msg.datapath
141
142     if dp.id not in self.waiters or msg.xid not in self.waiters[dp.id]:
143         return
144     event, msgs = self.waiters[dp.id][msg.xid]
145     msgs.append(msg.to_jsondict())
146     if ofproto_v1_3.OFP_VERSION == dp.ofproto.OFP_VERSION or ofproto_v1_4.OFP_VERSION
147     == dp.ofproto.OFP_VERSION:
148         more = dp.ofproto.OFPMPF_REPLY_MORE
149     else:
150         more = dp.ofproto.OFPSF_REPLY_MORE
151
152     if msg.flags & more:
153         return
154     del self.waiters[dp.id][msg.xid]
155     event.set()
156
157 # for OpenFlow version 1.0
158 @set_ev_cls(ofp_event.EventOFPPFlowStatsReply, MAIN_DISPATCHER)
159 def stats_reply_handler_v1_0(self, ev):
160     self._stats_reply_handler(ev)
161
162 # for OpenFlow version 1.2+
163 @set_ev_cls(ofp_event.EventOFPPStatsReply, MAIN_DISPATCHER)
164 def stats_reply_handler_v1_2(self, ev):
165     self._stats_reply_handler(ev)
166
167 @set_ev_cls(ofp_event.EventOFPPPortStatsReply, MAIN_DISPATCHER)
168 def port_stats_reply_handler(self, ev):
169     self._stats_reply_handler(ev)
170
171 @set_ev_cls(ofp_event.EventOFPPAggregateStatsReply, MAIN_DISPATCHER)
172 def aggregate_stats_reply_handler(self, ev):
173     self._stats_reply_handler(ev)
174
175 def rest_command(func):
176     def _rest_command(*args, **kwargs):
177         try:
178             msg = func(*args, **kwargs)
179             return Response(content_type='application/json',
```



```

180         body=json.dumps(msg))
181
182     except SyntaxError as e:
183         status = 400
184         details = e.msg
185     except (ValueError, NameError) as e:
186         status = 400
187         details = e.message
188
189     except NotFoundError as msg:
190         status = 404
191         details = str(msg)
192
193     msg = {REST_RESULT: REST_NG,
194           REST_DETAILS: details}
195     return Response(status=status, body=json.dumps(msg))
196
197     return _rest_command
198
199
200 class SwitchController(ControllerBase):
201     _SWITCH_LIST = {}
202     _LOGGER = None
203
204     def __init__(self, req, link, data, **config):
205         super(SwitchController, self).__init__(req, link, data, **config)
206         self.waiters = data['waiters']
207
208     @classmethod
209     def set_logger(cls, logger):
210         cls._LOGGER = logger
211         cls._LOGGER.propagate = False
212         hdlr = logging.StreamHandler()
213         fmt_str = '[RT][%(levelname)s] switch_id=%(sw_id)s: %(message)s'
214         hdlr.setFormatter(logging.Formatter(fmt_str))
215         cls._LOGGER.addHandler(hdlr)
216
217     @classmethod
218     def register_switch(cls, dp, waiters):
219         dpid = {'sw_id': dpid_lib.dpid_to_str(dp.id)}
220         try:
221             switch = Switch(dp, cls._LOGGER, waiters)
222         except OFPUnknownVersion as message:
223             cls._LOGGER.error(str(message), extra=dpid)
224         return
225         cls._SWITCH_LIST.setdefault(dp.id, switch)
226         cls._LOGGER.info('Join as switch.', extra=dpid)

```

A. Controller

```
227     # Try to setup connection to Monitoring after x seconds
228     t = Timer(DELAY_CONNECTION, switch.connect_monitoring)
229     t.start()
230
231     @classmethod
232     def unregister_switch(cls, dp):
233         if dp.id in cls._SWITCH_LIST:
234             cls._SWITCH_LIST[dp.id].delete()
235         del cls._SWITCH_LIST[dp.id]
236
237         dpid = {'sw_id': dpid.lib.dpid_to_str(dp.id)}
238         cls._LOGGER.info('Leave switch.', extra=dpid)
239
240     @classmethod
241     def packet_in_handler(cls, msg):
242         dp_id = msg.datapath.id
243         if dp_id in cls._SWITCH_LIST:
244             switch = cls._SWITCH_LIST[dp_id]
245             switch.packet_in_handler(msg)
246
247     # GET /switch/{switch_id}
248     @rest_command
249     def get_data(self, req, switch_id, **kwargs):
250         return self._access_switch(switch_id, 'get_data', req)
251
252     # POST /switch/{switch_id}
253     @rest_command
254     def set_data(self, req, switch_id, **kwargs):
255         return self._access_switch(switch_id, 'set_data', req)
256
257     # DELETE /switch/{switch_id}
258     @rest_command
259     def delete_data(self, req, switch_id, **kwargs):
260         return "Not implemented"
261
262     def _access_switch(self, switch_id, func, req):
263         rest_message = []
264         switches = self._get_switch(switch_id)
265         try:
266             param = req.json if req.body else {}
267         except ValueError:
268             raise SyntaxError('invalid syntax %s', req.body)
269         for switch in switches.values():
270             function = getattr(switch, func)
271             data = function(param)
272             rest_message.append(data)
273
```

```

274     return rest_message
275
276 def _get_switch(self, switch_id):
277     switches = {}
278
279     if switch_id == REST_ALL:
280         switches = self._SWITCH_LIST
281     else:
282         sw_id = dpid.lib.str_to_dpid(switch_id)
283         if sw_id in self._SWITCH_LIST:
284             switches = {sw_id: self._SWITCH_LIST[sw_id]}
285
286     if switches:
287         return switches
288     else:
289         raise NotFoundError(switch_id=switch_id)
290
291
292 class Switch(dict):
293     def __init__(self, dp, logger, waiters):
294         super(Switch, self).__init__()
295         self.dp = dp
296         self.dpid.str = dpid.lib.dpid_to_str(dp.id)
297         self.sw_id = {'sw_id': self.dpid.str}
298         self.logger = logger
299
300         self.monitoring_port = None
301
302         self.port_data = PortData(dp.ports)
303
304         self.history = {}
305
306         self.mac_to_port = {}
307
308         self.ofctl = OfCtl.factory(dp, logger)
309
310         self.setup_basic_flows()
311
312         self.redis_instance = self.setup_redis()
313
314         self.monitor_thread = hub.spawn(self._monitor, waiters)
315
316     def _monitor(self, waiters):
317         """Requesting statistic messages from switches at STATSPOLLRATE
318
319         :param waiters: the switches to request stats from
320         """

```

A. Controller

```
321
322 while True:
323     stats = self.request_stats(waiters)
324     message = {
325         'event_type': 'statistics',
326         'statistics': stats['statistics'],
327     }
328     self.redis_instance.publish('switch-stats', json.dumps(message))
329     hub.sleep(STATS_POLL_RATE)
330
331 def setup_redis(self, retry=0, max_retries=30):
332     try:
333         redis_instance = redis.StrictRedis(host=REDIS_HOST, port=REDIS_PORT, db=0)
334     except redis.exceptions.ConnectionError as e:
335         if retry < max_retries:
336             self.logger.warning("Retrying connections to database for the {}
337 time".format(retry + 1),
338                             extra=self.sw_id)
339             time.sleep(2)
340             redis_instance = self.setup_redis(retry=retry + 1)
341         else:
342             raise e
343     return redis_instance
344
345 def connect_monitoring(self):
346     # Send ARP Request to get MACAddress of Monitoring
347     self.send_arp_request(IP_SWITCHARP, IP_MONITORING)
348     self.logger.info('Send ARP request (flood)', extra=self.sw_id)
349
350 def delete_flows(self, table_id):
351     # Delete all flows in Table
352     # Used to delete old flows before monitoring was connected
353     ofproto = self.dp.ofproto
354     parser = self.dp.ofproto_parser
355     match = parser.OFPMatch()
356     inst = []
357     mod = parser.OFPFlowMod(self.dp, 0, 0, table_id, ofproto.OFPFC_DELETE, 0, 0, 1,
358                             ofproto.OFPCML_NO_BUFFER,
359                             ofproto.OFPP_ANY, ofproto.OFPG_ANY, 0, 0, match, inst)
360     self.dp.send_msg(mod)
361
362 def setup_basic_flows(self):
363     ofproto = self.dp.ofproto
364     parser = self.dp.ofproto_parser
365     # Table Miss Entry
```

```

366
367 match = parser.OFPMatch()
368 actions = [parser.OFPActionOutput(ofproto.OFPP_CONTROLLER,
369                                 ofproto.OFPCML_NO_BUFFER)]
370 self.ofctl.add_flow(0, 0, match, actions)
371
372 # Make sure our ARP replies are sent to the controller
373
374 match = parser.OFPMatch(arp_tpa=IP.SWITCHARP, eth_type=2054)
375 actions = [parser.OFPActionOutput(ofproto.OFPP_CONTROLLER,
376                                 ofproto.OFPCML_NO_BUFFER)]
377 self.ofctl.add_flow(3, 0, match, actions)
378
379 def delete(self):
380     self.logger.info('Stopped switch', extra=self.sw_id)
381
382 def get_data(self, data, waiters):
383     parser = self.dp.ofproto_parser
384     msg = "Not implemented" #TODO return load statistics
385     return {REST_SWITCHID: self.dpid_str,
386           REST_COMMAND_RESULT: msg}
387
388 def set_data(self, data):
389     parser = self.dp.ofproto_parser
390     ofp = self.dp.ofproto
391     msgs = []
392     try:
393         if REST_COMMAND in data:
394             action_list = []
395             cmd = data[REST_COMMAND]
396             command_obj = getattr(ofp, cmd)
397             if REST_MATCH in data:
398                 matches = data[REST_MATCH]
399                 match = parser.OFPMatch(**matches)
400             else:
401                 match = parser.OFPMatch()
402             if REST_ACTIONS in data:
403                 actions = data[REST_ACTIONS]
404                 if REST_ACTION in actions:
405                     for k, v in actions.items():
406                         if REST_COMMAND in v:
407                             cmd = v[REST_COMMAND]
408                             action_obj = getattr(parser, cmd)
409                         else:
410                             details = "No action Command found"
411                             raise ValueError(details)
412                 if REST_PARAMS in v:

```

A. Controller

```
413         params = v[REST_PARAMS]
414         action_list.append(action_obj(**params))
415     else:
416         details = "No Parameters provided for %s" % cmd
417         raise ValueError(details)
418     else:
419         details = "No action found in actions"
420         raise ValueError(details)
421     if REST_PRIORITY in data:
422         priority = data[REST_PRIORITY]
423     else:
424         # Default value for mitigation flows
425         priority = 32768
426     if REST_TABLEID in data:
427         table_id = data[REST_TABLEID]
428     else:
429         # Default value for mitigation flows
430         table_id = 0
431     if REST_COOKIE in data:
432         cookie = data[REST_COOKIE]
433     else:
434         cookie = 0
435     if REST_COOKIEMASK in data:
436         cookie_mask = data[REST_COOKIEMASK]
437     else:
438         cookie_mask = 0
439     if REST_IDLETIMEOUT in data:
440         idle_timeout = data[REST_IDLETIMEOUT]
441     else:
442         idle_timeout = 0
443     if REST_HARDTIMEOUT in data:
444         hard_timeout = data[REST_HARDTIMEOUT]
445     else:
446         hard_timeout = 0
447     if cmd == 'OPFPC_DELETE':
448         self.ofctl.delete_flow(priority=priority, table_id=table_id, match=match,
449                                command=command_obj,
450                                cookie=cookie, cookie_mask=cookie_mask)
451     else:
452         self.ofctl.add_flow(priority=priority, table_id=table_id, match=match,
453                             action_list=action_list,
454                             hard_timeout=hard_timeout, idle_timeout=idle_timeout,
455                             command=command_obj,
456                             cookie=cookie, cookie_mask=cookie_mask)
457     self.logger.info(
458         "Added mitigation flow: {}".format(priority, table_id, match,
459                                           action_list,
```

```

456             hard_timeout, idle_timeout), extra=self.sw_id)
457         details = 'Performed %s' % cmd
458     else:
459         raise ValueError("No command found")
460     msg = {REST_RESULT: REST_OK, REST_DETAILS: details}
461     except ValueError as err_msg:
462         msg = {REST_RESULT: REST_NG, REST_DETAILS: str(err_msg)}
463     msgs.append(msg)
464     return {REST_SWITCHID: self.dpid_str,
465           REST_COMMAND_RESULT: msgs}
466
467 def request_stats(self, waiters):
468     stats = {}
469     stat_replies = {}
470     stat_replies['port_aggregate'] = {}
471     stat_replies['port_aggregate'][self.dpid_str] = {}
472     for send_port in self.port_data.values():
473         send_port = send_port.port_no
474         stat_replies['port_aggregate'][self.dpid_str][send_port] =
475             self.ofctl.get_all_flow_aggregate(send_port,
476                                             waiters)
477     # flows = self.ofctl.get_all_flow(waiters)
478     # if flows:
479     #     # flows = self.ofctl.get_all_flow(waiters)[0]
480     #     # stat_replies['flow'] = {}
481     #     # stat_replies['flow'][self.dpid_str] = flows
482     portstats = {}
483     ports = self.ofctl.get_all_port(waiters)
484     if ports:
485         ports = self.ofctl.get_all_port(waiters)[0]
486         for msg in ports['OFPPortStatsReply']['body']:
487             msg = msg['OFPPortStats']
488             port_no = msg['port_no']
489             del msg['port_no']
490             # Calculate deltas
491             msg['rx_bytes_delta'] = self.get_port_delta(msg, port_no, 'rx_bytes')
492             msg['rx_dropped_delta'] = self.get_port_delta(msg, port_no, 'rx_dropped')
493             msg['rx_errors_delta'] = self.get_port_delta(msg, port_no, 'rx_errors')
494             msg['rx_packets_delta'] = self.get_port_delta(msg, port_no, 'rx_packets')
495             msg['tx_bytes_delta'] = self.get_port_delta(msg, port_no, 'tx_bytes')
496             msg['tx_dropped_delta'] = self.get_port_delta(msg, port_no, 'tx_dropped')
497             msg['tx_errors_delta'] = self.get_port_delta(msg, port_no, 'tx_errors')
498             msg['tx_packets_delta'] = self.get_port_delta(msg, port_no, 'tx_packets')
499             portstats[port_no] = msg
500     stat_replies['port'] = {}
501     stat_replies['port'][self.dpid_str] = portstats
502     stats['statistics'] = stat_replies

```

A. Controller

```
502     return stats
503
504 def get_port_delta(self, msg, port, field):
505     new_value = msg[field]
506     if port in self.history:
507         if field in self.history[port]:
508             old_value = self.history[port][field]
509             self.history[port][field] = new_value
510             return (new_value - old_value) / STATS_POLL_RATE
511         else:
512             self.history[port][field] = new_value
513             return new_value / STATS_POLL_RATE
514     else:
515         self.history[port] = {}
516         self.history[port][field] = new_value
517         return new_value / STATS_POLL_RATE
518
519 def packet_in_handler(self, msg):
520     pkt = packet.Packet(msg.data)
521     eth = pkt.get_protocol(ethernet.ethernet)
522
523     if not eth:
524         self.logger.warning("Non ethernet Packet received!", extra=self.sw_id)
525         return
526     arp_p = pkt.get_protocol(arp.arp)
527     if arp_p:
528         self._packetin_arp(msg, arp_p)
529     ofproto = self.dp.ofproto
530     parser = self.dp.ofproto_parser
531     in_port = msg.match['in_port']
532
533     if eth.ethertype == ether_types.ETH_TYPE_LLDP:
534         # ignore lldp packet
535         return
536     dst = eth.dst
537     src = eth.src
538
539     # self.logger.info("Packet in [src:%s] [dst:%s] [in_port:%s]", src, dst, in_port,
540     #                  extra=self.sw_id)
541
542     self.mac_to_port[src] = in_port
543
544     if dst in self.mac_to_port:
545         out_port = self.mac_to_port[dst]
546     else:
547         out_port = ofproto.OFPP_FLOOD
548
```



```

549     # self.logger.info("Additional [out_port:%s]", out_port,
550     #                 extra=self.sw_id)
551
552     # Check if Monitoring is connected and packet not already being forwarded to
    # Monitoring
553     if self.monitoring_port is not None and out_port != self.monitoring_port and \
554         in_port != self.monitoring_port and out_port != ofproto.OFPP_FLOOD:
555         match = parser.OFPMatch(in_port=in_port, eth_dst=dst)
556         if out_port == in_port:
557             # Only duplicate to monitoring but do not forward
558             actions = [parser.OFPActionGroup(group_id=0)]
559         else:
560             actions = [parser.OFPActionGroup(group_id=0), parser.OFPActionOutput(out_port)]
561
562         inst = [parser.OFPInstructionActions(ofproto.OFPIT_APPLY_ACTIONS,
563                                             actions)]
564
565         mod = parser.OFPFlowMod(datapath=self.dp, priority=1,
566                                 table_id=0, match=match, instructions=inst)
567         self.dp.send_msg(mod)
568     else:
569         if out_port == in_port:
570             # Do not forward packets to the in_port (only caused by duplicated packet
    forwarding?)
571             self.logger.info("Illegal rule?: in_port: [%s], mac_dst: [%s] -> out_port:
    [%s]",
572                             in_port, out_port, dst, extra=self.sw_id)
573             return
574         actions = [parser.OFPActionOutput(out_port)]
575
576         if out_port != ofproto.OFPP_FLOOD:
577             match = parser.OFPMatch(in_port=in_port, eth_dst=dst)
578             self.ofctl.add_flow(1, 0, match, actions)
579
580         data = None
581         if msg.buffer_id == ofproto.OFP_NO_BUFFER:
582             data = msg.data
583
584         out = parser.OFPPacketOut(datapath=self.dp, buffer_id=msg.buffer_id,
585                                 in_port=in_port, actions=actions, data=data)
586         self.dp.send_msg(out)
587
588     def _packetin_arp(self, msg, header):
589
590         if header.opcode == arp.ARP_REPLY:
591             src_ip = header.src_ip
592             dst_ip = header.dst_ip

```

A. Controller

```
593     src_mac = header.src_mac
594     srcip = ip_addr_ntoa(src_ip)
595     dstip = ip_addr_ntoa(dst_ip)
596     if dst_ip == IP_SWITCHARP and src_mac in self.mac_to_port:
597         out_port = self.mac_to_port[src_mac]
598         self.monitoring_port = out_port
599         # Delete old flows in table 0
600         self.delete_flows(0)
601         # Setup basic flows again
602         self.setup_basic_flows()
603         log_msg = 'Receive ARP request from [%s] to switch port [%s].' \
604                 ' Setting up duplicate packet flow to port [%s]'
605         self.logger.info(log_msg, srcip, dstip, out_port, extra=self.sw_id)
606         # Setup Group 0 for packet duplication to Monitoring port
607
608         parser = self.dp.ofproto_parser
609         ofproto = self.dp.ofproto
610
611         actions = [parser.OFPActionOutput(port=out_port)]
612
613         buckets = [parser.OFPBucket(actions=actions)]
614
615         req = parser.OFPGroupMod(datapath=self.dp, command=ofproto.OFPGC_ADD,
616                                type_=ofproto.OFPGT_ALL, group_id=0, buckets=buckets)
617
618         self.dp.send_msg(req)
619
620     def send_arp_request(self, src_ip, dst_ip, in_port=None):
621
622         # Send ARP request from all ports.
623         for send_port in self.port_data.values():
624             if in_port is None or in_port != send_port.port_no:
625                 src_mac = send_port.mac
626                 dst_mac = mac_lib.BROADCAST_STR
627                 arp_target_mac = mac_lib.DONTCARE_STR
628                 inport = self.ofctl.dp.ofproto.OFPP_CONTROLLER
629                 output = send_port.port_no
630                 self.ofctl.send_arp(arp.ARP_REQUEST, src_mac, dst_mac, src_ip, dst_ip,
631                                   arp_target_mac, inport, output)
632
633
634     class PortData(dict):
635     def __init__(self, ports):
636         super(PortData, self).__init__()
637         for port in ports.values():
638             data = Port(port.port_no, port.hw_addr)
639             self[port.port_no] = data
```

```

640
641
642 class Port(object):
643     def __init__(self, port_no, hw_addr):
644         super(Port, self).__init__()
645         self.port_no = port_no
646         self.mac = hw_addr
647
648
649 class OfCtl(object):
650     _OF_VERSIONS = {}
651
652     @staticmethod
653     def register_of_version(version):
654         def _register_of_version(cls):
655             OfCtl._OF_VERSIONS.setdefault(version, cls)
656             return cls
657
658         return _register_of_version
659
660     @staticmethod
661     def factory(dp, logger):
662         of_version = dp.ofproto.OFP_VERSION
663         if of_version in OfCtl._OF_VERSIONS:
664             ofctl = OfCtl._OF_VERSIONS[of_version](dp, logger)
665         else:
666             raise OFPUnknownVersion(version=of_version)
667
668         return ofctl
669
670     def __init__(self, dp, logger):
671         super(OfCtl, self).__init__()
672         self.dp = dp
673         self.sw_id = {'sw_id': dpid_lib.dpid_to_str(dp.id)}
674         self.logger = logger
675
676     def add_flow(self, priority, table_id, match, action_list=None, hard_timeout=0,
677                idle_timeout=0, command=0, cookie=0,
678                cookie_mask=0):
679         # Abstract method
680         raise NotImplementedError()
681
682     def send_arp(self, arp_opcode, src_mac, dst_mac,
683                src_ip, dst_ip, arp_target_mac, in_port, output):
684         # Generate ARP packet
685
686         ether_proto = ether.ETH_TYPE_ARP

```

A. Controller

```
686     hwtype = 1
687     arp_proto = ether.ETH_TYPE_IP
688     hlen = 6
689     plen = 4
690
691     pkt = packet.Packet()
692     e = ethernet.ethernet(dst_mac, src_mac, ether_proto)
693     a = arp.arp(hwtype, arp_proto, hlen, plen, arp_opcode,
694               src_mac, src_ip, arp_target_mac, dst_ip)
695     pkt.add_protocol(e)
696     pkt.add_protocol(a)
697     pkt.serialize()
698
699     # Send packet out
700     self.send_packet_out(in_port, output, pkt.data, data_str=str(pkt))
701
702     def send_packet_out(self, in_port, output, data, data_str=None):
703         actions = [self.dp.ofproto_parser.OFPActionOutput(output, 0)]
704         self.dp.send_packet_out(buffer_id=UINT32_MAX, in_port=in_port,
705                                actions=actions, data=data)
706         # TODO: ?Packet library convert to string
707         # if data_str is None:
708         # data_str = str(packet.Packet(data))
709         # self.logger.debug('Packet out = %s', data_str, extra=self.sw_id)
710
711     def send_stats_request(self, stats, waiters):
712         self.dp.set_xid(stats)
713         waiters_per_dp = waiters.setdefault(self.dp.id, {})
714         event = hub.Event()
715         msgs = []
716         waiters_per_dp[stats.xid] = (event, msgs)
717         self.dp.send_msg(stats)
718
719         try:
720             event.wait(timeout=OFP_REPLY_TIMER)
721         except hub.Timeout:
722             del waiters_per_dp[stats.xid]
723
724         return msgs
725
726
727 @OfCtl.register_of_version(ofproto_v1.4.OFP_VERSION)
728 class OfCtl_v1.4(OfCtl):
729     def __init__(self, dp, logger):
730         super(OfCtl_v1.4, self).__init__(dp, logger)
731
```

```

732 def add_flow(self, priority, table_id, match, action_list=None, hard_timeout=0,
733             idle_timeout=0, command=0, cookie=0,
734             cookie_mask=0):
735     datapath = self.dp
736     ofproto = datapath.ofproto
737     parser = datapath.ofproto_parser
738     if command == 0:
739         command = ofproto.OFPFC_ADD
740
741     inst = [parser.OFPInstructionActions(ofproto.OFPIT_APPLY_ACTIONS,
742                                       action_list)]
743
744     mod = parser.OFPFlowMod(datapath=datapath, priority=priority,
745                             table_id=table_id, match=match,
746                             instructions=inst, hard_timeout=hard_timeout,
747                             idle_timeout=idle_timeout, command=command, cookie=cookie,
748                             cookie_mask=cookie_mask)
749     datapath.send_msg(mod)
750
751 def delete_flow(self, priority, table_id, match, command, cookie, cookie_mask):
752     ofp = self.dp.ofproto
753     ofp_parser = self.dp.ofproto_parser
754
755     inst = []
756
757     flow_mod = ofp_parser.OFPFlowMod(self.dp, cookie, cookie_mask, table_id, command,
758                                     0, 0, priority, ofp.OFPCML_NO_BUFFER, ofp.OFPP_ANY,
759                                     ofp.OFPG_ANY, 0, 0, match, inst)
760     self.dp.send_msg(flow_mod)
761     self.logger.info('Delete flow [cookie=0x%x]', cookie, extra=self.sw_id)
762
763 def get_all_flow(self, waiters):
764     ofp = self.dp.ofproto
765     ofp_parser = self.dp.ofproto_parser
766
767     match = ofp_parser.OFPMatch()
768     stats = ofp_parser.OFPFlowStatsRequest(self.dp, 0, 0, ofp.OFPP_ANY,
769                                           ofp.OFPG_ANY, 0, 0, match)
770     return self.send_stats_request(stats, waiters)
771
772 def get_all_port(self, waiters):
773     ofp = self.dp.ofproto
774     ofp_parser = self.dp.ofproto_parser
775
776     stats = ofp_parser.OFPPortStatsRequest(self.dp, 0, ofp.OFPP_ANY)
777     return self.send_stats_request(stats, waiters)

```

A. Controller

```
777
778 def get_all_flow_aggregate(self, port, waiters):
779     ofp = self.dp.ofproto
780     ofp_parser = self.dp.ofproto_parser
781
782     match = ofp_parser.OFPMatch(in_port=port)
783
784     stats = ofp_parser.OFPAggregateStatsRequest(self.dp, 0, ofp.OFPTT_ALL,
785         ofp.OFPP_ANY, ofp.OFPG_ANY, 0, 0, match)
786     return self.send_stats_request(stats, waiters)
787
788 def ipv4_text_to_int(ip_text):
789     if ip_text == 0:
790         return ip_text
791     assert isinstance(ip_text, str)
792     return struct.unpack('!I', addrconv.ipv4.text_to_bin(ip_text))[0]
793
794
795 def ip_addr_ntoa(ip):
796     return socket.inet_ntoa(addrconv.ipv4.text_to_bin(ip))
797
798
799 def mask_ntob(mask, err_msg=None):
800     try:
801         return (UINT32_MAX <<< (32 - mask)) & UINT32_MAX
802     except ValueError:
803         msg = 'illegal netmask'
804         if err_msg is not None:
805             msg = '%s %s' % (err_msg, msg)
806         raise ValueError(msg)
```

B. Mitigation component

B.1. Documentation

Alert Handler Documentation

Release 1.0.0rc0

Alex Marczinek

Dec 11, 2016

CONTENTS:

1	Usage Documentation	1
1.1	Requirements	1
1.2	Standalone Usage	1
1.3	System Service	1
2	Welcome to Mitigation's documentation!	3
3	Usage Documentation	5
3.1	System Service	5
4	Documented Modules	7
4.1	AlertHandler	7
4.2	Mitigation	8
	Python Module Index	9

USAGE DOCUMENTATION

This module can be used in two ways. First, it can be run in as a standalone piece of software, described in *standalone*. Further usage can happen through installing the script as a system service, more in *system_service*.

This module connects to a Redis Database. As Redis uses multiple channels, there can be multiple instances of this module be running concurrently, listening to different channels. Messages received from the Redis channel are then parsed and forwarded to the Mitigation module.

1.1 Requirements

Required to run this script is at least Python 3.4 and the package 'redis' ('pip install redis').

1.2 Standalone Usage

The script can be run in standalone mode using the following options:

- -l : The loglevel for the application
- -r : The IP of the Redis instance to connect to
- -p : The port of the Redis instance
- -k : The key/channel to subscribe to
- -d : The database in Redis to subscribe to, usually 0

1.3 System Service

Usage as a system service can happen, a service file for usage with systemd is provided in *./miti-ids.service* and *./miti-stats.service*. Copying this file to */etc/systemd/system/* enables systemd to start the script. After a 'systemctl daemon-reload' it can be started with 'systemctl start <service-name>' and enabled at boot with 'systemctl enable <service-name>'.

The usage in the *.service* file is the same as described in *standalone*.

WELCOME TO MITIGATION'S DOCUMENTATION!

USAGE DOCUMENTATION

This module runs as a system service, as described in *system_service2*.

It creates an HTTP server for a RESTful API with the purpose of receiving the alerts from the AlertHandler

Based on these alerts a Mitigation strategy is chosen and implemented via the Northbound API of the controller.

3.1 System Service

Usage as a system service can happen, a service file for usage with systemd is provided in `./miti-mit.service`. Copying this file to `/etc/systemd/system/` enables systemd to start the script. After a `'systemctl daemon-reload'` it can be started with `'systemctl start miti-mit'` and enabled at boot with `'systemctl enable miti-mit'`.

DOCUMENTED MODULES

In the following sections, the modules contained within the AlertHandler and Mitigation project will be documented.

4.1 AlertHandler

`AlertHandler.setup_redis(ip, port, key, retry=0, max_retries=30)`

This will connect to a Redis instance and return the connection object.

Establishing a stable connection with a Redis instance is the goal. Therefore, the method will retry at most `max_retries` times.

Parameters

- **ip** – The IP of the Redis server to connect to.
- **port** – The port on which Redis is running.
- **key** – The redis channel/key to subscribe the a PubSub object to.
- **max_retries** – The maximum number of retries after which connection attempts will fail.
- **retry** – The current number of retries executed.

Returns A set of (redis_instance, pubsub)

`AlertHandler.parseargs()`

Parsing arguments passed

Using the built in argparse module, passed arguments are parsed. Simply calling this method is sufficient, the results will be returned.

Returns A namespace object containing all parsed arguments

`AlertHandler.parse_message(message)`

This will parse the messages received from Redis.

Parses the message received from Redis and sends alert messages to the mitigation module. The alerts are sent to <http://127.0.0.1:9000/store.json>

Parameters **message** – The message received from Redis.

`AlertHandler.main()`

The main method.

First, arguments passed to the system are parsed. Then, logging is set up. The connection to redis is established by calling `setup_redis()` and the main event loop started.

Returns The system exit code

4.2 Mitigation

`Mitigation.main()`

The main method.

Logging is setup and the HTTP server is started.

Returns The system exit code

class `Mitigation.MyServer` (*request, client_address, server*)

The HTTP Server receiving the alerts from Alert Handler

do_POST ()

Handles the POST messages and passes them to mitigation

mitigate (*parsed_alert: dict*)

Selects a mitigation strategy based on the Signature ID from the alert.

Depending on the strategy a flow instruction is created, that is then sent to the controller via a POST to 10.20.0.8:8080

Parameters `parsed_alert` – The alert from the Alert Handler as a python dictionary

Returns True if the attack was mitigated False if it could not mitigate the attack

PYTHON MODULE INDEX

a

AlertHandler, 7

m

Mitigation, 8

B.2. Source Code Alert Handler

```

1 import argparse
2 import json
3 import logging
4 import sys
5 import time
6
7 import redis
8
9 from src.tools import setup_logging, send_post
10
11
12 def parseargs():
13     """Parsing arguments passed
14
15     Using the built in argparse module, passed arguments are parsed. Simply calling this
16     method is sufficient,
17     the results will be returned.
18
19     :return: A namespace object containing all parsed arguments
20     """
21     loglevel = {
22         'debug': logging.DEBUG,
23         'info': logging.INFO,
24         'warning': logging.WARNING,
25         'error': logging.ERROR,
26         'critical': logging.CRITICAL,
27     }
28
29     parser = argparse.ArgumentParser()
30     parser.add_argument('-l', '--loglevel', help='Loglevels to output: debug, info,
31         warning, error, critical',
32         default='info', choices=loglevel.keys())
33     parser.add_argument('-r', '--redis', help='Redis instance to query for alerts',
34         required=True)
35     parser.add_argument('-k', '--key', help='Key to subscribe to in redis',
36         default='suricata')
37     parser.add_argument('-p', '--port', help='Port to bind to, default is 6379',
38         default=6379, type=int)
39     parser.add_argument('-d', '--database', help='Database number to subscribe to',
40         default=0, type=int)
41     args = parser.parse_args()
42
43     # Validating arguments

```

B. Mitigation component

```
39  args.loglevel = loglevel[args.loglevel]
40  return args
41
42
43  def setup_redis(ip, port, key, retry=0, max_retries=30):
44      """This will connect to a Redis instance and return the connection object.
45
46      Establishing a stable connection with a Redis instance is the goal. Therefore, the
47      method will retry at most
48
49      :param ip: The IP of the Redis server to connect to.
50      :param port: The port on which Redis is running.
51      :param key: The redis channel/key to subscribe the a PubSub object to.
52      :param max_retries: The maximum number of retries after which connection attempts
53      will fail.
54      :param retry: The current number of retries executed.
55      :return: A set of (redis_instance, pubsub)
56      """
57      try:
58          redInstance = redis.StrictRedis(host=ip, port=port, db=0)
59          pubsub = redInstance.pubsub()
60          pubsub.subscribe(key)
61      except redis.exceptions.ConnectionError as e:
62          if retry < max_retries:
63              logging.warning("Retrying connections to database for the {} time".format(retry
64              + 1))
65              time.sleep(2)
66              pubsub = setup_redis(ip, port, key, retry=retry + 1)
67          else:
68              raise e
69      return pubsub
70
71
72  def parse_message(message):
73      """This will parse the messages received from Redis.
74
75      Parses the message received from Redis and sends alert messages to the mitigation
76      module.
77
78      The alerts are sent to http://127.0.0.1:9000/store.json
79
80      :param message: The message received from Redis.
81      """
82      data = message['data']
83      if type(data) == int:
```

```

82     return
83
84     received_json = json.loads(data.decode())
85     if received_json['event_type'] == 'alert':
86         logging.info("Received alert : \n{}".format(received_json))
87         send_post(received_json, 'http://127.0.0.1:9000/store.json')
88
89
90 def main():
91     """The main method.
92
93     First, arguments passed to the system are parsed. Then, logging is set up. The
94     connection to redis is established
95     by calling setup_redis() and the main event loop started.
96
97     :return: The system exit code
98     """
99     args = parseargs()
100     setup_logging(args.loglevel)
101
102     pubsub = setup_redis(args.redis, args.port, args.key)
103     try:
104         while True:
105             message = pubsub.get_message()
106             if message is None:
107                 time.sleep(0.1)
108             else:
109                 parse_message(message)
110     except KeyboardInterrupt:
111         logging.info("Ctrl+C Pressed. Shutting down.")
112         pubsub.close()
113         return 0
114
115 if __name__ == "__main__":
116     sys.exit(main())

```

B.3. Source Code Mitigation

```

1 #!/usr/bin/env python
2 import datetime
3 import sys
4 import json
5 import logging
6 import signal

```

B. Mitigation component

```
7
8 from http.server import BaseHTTPRequestHandler, HTTPServer
9
10 from src.tools import setup_logging, send_post
11
12 ETH_TYPE_IPV4 = 2048
13
14 IP_PROTO_ICMP = 1
15 IP_PROTO_TCP = 6
16
17 hostName = "localhost"
18 hostPort = 9000
19
20
21 class MyServer(BaseHTTPRequestHandler):
22     """The HTTP Server receiving the alerts from Alert Handler"""
23     suricata_alert = {}
24     stats_alert = {}
25
26     def do_POST(self):
27         """Handles the POST messages and passes them to mitigation"""
28         if self.path == '/store.json':
29             length = self.headers['content-length']
30
31             data = self.rfile.read(int(length))
32             data = data.decode()
33             print("Got message: {}".format(data))
34             received_json = json.loads(data)
35
36             self.send_response(200, 'OK')
37             self.end_headers()
38
39             self.mitigate(received_json)
40
41     def mitigate(self, parsed_alert: dict):
42         """Selects a mitigation strategy based on the Signature ID from the alert.
43
44         Depending on the strategy a flow instruction is created, that is then sent to the
45         controller via
46         a POST to 10.20.0.8:8080
47
48         :param parsed_alert: The alert from the Alert Handler as a python dictionary
49         :return: True if the attack was mitigated False if it could not mitigate the attack
50         """
51         assert type(parsed_alert) == dict
52         alert = parsed_alert['alert']
```



```

53 sid = str(alert.get('signature_id', "0"))
54 try:
55     if sid == "10000001" or sid == "10000002":
56         #ICMP
57         return "ICMP: do nothing"
58     elif sid == "10000003":
59         logging.info('TCP SYN Flood Attack detected')
60         # Parameters:
61         hard_timeout = 20
62         #SYN Flood
63         src = parsed_alert["src_ip"]
64         dst = parsed_alert["dest_ip"]
65         match = dict()
66         match['ipv4_src'] = src
67         match['eth_type'] = ETH_TYPE_IPV4
68         match['ipv4_dst'] = dst
69         match['ip_proto'] = IP_PROTO_TCP
70         data = dict()
71         data['cmd'] = 'OFFPC_ADD'
72         data['hard_timeout'] = hard_timeout
73         data['match'] = match
74         post_result = send_post(data, 'http://10.20.0.8:8080/switch/all')
75         logging.info(
76             'Handled SYN Flood (Target: {dst}, Attacker: {src}): Drop packets for
77             {timeout}s. [{post}]'.format(
78                 dst=dst, src=src, timeout=hard_timeout, post=post_result))
79         return True
80     elif sid == "10000004":
81         if 'host' in parsed_alert and parsed_alert['host'] == "monitoring":
82             # Alert from Suricata
83             self.suricata.alert['timestamp'] = parsed_alert['timestamp']
84             self.suricata.alert['dest_ip'] = parsed_alert['dest_ip']
85             self.suricata.alert['dest_port'] = parsed_alert['dest_port']
86         elif 'sensor' in parsed_alert and parsed_alert['sensor'] == "statshandler":
87             # Alert from Statistical Monitoring
88             self.stats.alert['timestamp'] = parsed_alert['timestamp']
89             self.stats.alert['switch'] = parsed_alert['switch']
90             self.stats.alert['port'] = parsed_alert['port']
91         else:
92             logging.error('Unknown Alert Source')
93             return False
94     if 'timestamp' in self.suricata.alert and 'timestamp' in self.stats.alert:
95         logging.info('Found possible DDoS Syn Attack')
96         time_suricata = datetime.datetime.strptime(self.suricata.alert['timestamp'],
97             "%Y-%m-%dT%H:%M:%S.%fz")

```

B. Mitigation component

```
98         time_stats = datetime.datetime.strptime(self.stats_alert['timestamp'],
99         "%Y-%m-%dT%H:%M:%S.%f%z")
100         time_delta = abs((time_suricata - time_stats).total_seconds())
101         logging.debug("Having a timedelta of {} seconds".format(time_delta))
102         if time_delta < 5:
103             # Parameters:
104             hard_timeout = 20
105             # SYN Flood
106             switch = self.stats_alert['switch']
107             port = self.stats_alert['port']
108             dest_ip = self.suricata_alert['dest_ip']
109             dest_port = self.suricata_alert['dest_port']
110             match = dict()
111             match['tcp_dst'] = int(dest_port)
112             match['eth_type'] = ETH_TYPE_IPv4
113             match['ipv4_dst'] = dest_ip
114             match['ip_proto'] = IP_PROTO_TCP
115             match['in_port'] = int(port)
116             data = dict()
117             data['cmd'] = 'OFPPC_ADD'
118             data['hard_timeout'] = hard_timeout
119             data['match'] = match
120             post_result = send_post(data, 'http://10.20.0.8:8080/switch/' + switch)
121             logging.info('Handled Distributed SYN Flood (Target: {dst}, Port:
122             {dst_port}):' \
123             ' Drop packets on port {port} on switch{switch} for {timeout}s.
124             [{post}]' \
125             .format(dst=dest_ip, dst_port=dest_port, port=port, switch=switch,
126             timeout=hard_timeout, post=post_result))
127             return True
128             # DDoS Syn Flood
129             return True
130         elif sid == "10000005":
131             # Flow flood mitigation
132             logging.info('Flow Flood Attack detected')
133             # Parameters:
134             hard_timeout = 20
135             # SYN Flood
136             switch = parsed.alert['switch_id']
137             port = parsed_alert['bad_port']
138             match = dict()
139             match['in_port'] = int(port)
140             data = dict()
141             data['cmd'] = 'OFPPC_ADD'
142             data['hard_timeout'] = hard_timeout
143             data['match'] = match
144             data['cookie'] = 0x0000000000000001
```

```

142     data['cookie_mask'] = 0x0000000000000000f
143     post_result = send_post(data, 'http://10.20.0.8:8080/switch/' + switch)
144     # Delete malicious flow entries
145     data['cmd'] = 'OFPFC_DELETE'
146     data['cookie'] = 0x0000000000000000
147     data['cookie_mask'] = 0x0000000000000000f
148     post_result2 = send_post(data, 'http://10.20.0.8:8080/switch/' + switch)
149     logging.info('Handled Flow Flood:' \
150                 ' Drop packets {match} on port {port} on switch {switch} for
151                 {timeout}s. [{post}]' \
152                 .format(match=match, port=port, switch=switch,
153                           timeout=hard_timeout, post=post_result))
153     logging.info('Handled Flow Flood:' \
154                 ' Delete Flows on port {port} on switch {switch} for {timeout}s.
155                 [{post}]' \
156                 .format(port=port, switch=switch,
157                           timeout=hard_timeout, post=post_result2))
157     return True
158 elif sid == "20000001":
159     # Test for statistical monitoring module
160     logging.info("Statistical Test: do nothing")
161     return False
162 else:
163     logging.error("SignatureID: {} is not handled".format(sid))
164     return False
165 except LookupError as e:
166     return e.args
167
168
169 def main():
170     """The main method.
171
172     Logging is setup and the HTTP server is started.
173
174     :return: The system exit code
175     """
176     signal.signal(signal.SIGTERM, signal_term_handler)
177     setup_logging(logging.INFO)
178     myServer = HTTPServer((hostName, hostPort), MyServer)
179     logging.info('Server starting bound to {}:{}'.format(hostName, hostPort))
180     try:
181         myServer.serve_forever()
182     except KeyboardInterrupt:
183         shutdown()
184     return 0
185
186

```

B. Mitigation component

```
187 def shutdown():
188     myServer.server_close()
189     logging.info('Server shut down')
190     sys.exit(0)
191
192
193 def signal_term_handler(signal, frame):
194     shutdown()
195
196
197 if __name__ == '__main__':
198     sys.exit(main())
```

C. StatsHandler Documentation

Statshandler Documentation

Release 1.0.0rc0

Christoph Girstenbrei

Dec 11, 2016

CONTENTS:

1	Usage Documentation	1
1.1	Requirements	1
1.2	Standalone Usage	1
1.3	System Service	1
2	Development Guide	3
2.1	New Analyzer	3
3	Documented Modules	5
3.1	Statshandler	5
3.2	src.Monitoring	6
3.3	src.Analyzers	6
3.4	src.redis	8
	Python Module Index	11

USAGE DOCUMENTATION

This module can be used in two ways. First, it can be run in as a standalone piece of software, described in *standalone*. Further usage can happen through installing the script as a system service, more in *system_service*.

1.1 Requirements

Required to run this script is at least Python 3.4 and the package 'redis' ('pip install redis').

1.2 Standalone Usage

The script can be run in standalone mode using the following options:

- -r : The IP of the Redis instance to connect to
- -p : The port of the Redis instance
- -k : The key/channel to subscribe to
- -d : The database in Redis to subscribe to, usually 0
- -a : The critical value (int) when to alert if a port surpasses this in transmitting packages
- -e : Ports not to alert on, usefull for backbone ports transmitting lots of data
- -s : Switch, if statistics should be sent to Redis or not

If the statistics data is published to the Redis channel 'switch-stats', running the script can be as ease as running:
./StatsHandler.py -r <IP>

1.3 System Service

Usage as a system service can happen, a service file for usage with systemd is provided in *./statshandler.service*. Copying this file to */etc/systemd/system/statshandler.service* enables systemd to start the script. After a 'systemctl daemon-reload' it can be started with 'systemctl start statshandler' and enabled at boot with 'systemctl enable statshandler'.

The usage in the *.service* file is the same as described in *standalone*.

DEVELOPMENT GUIDE

This software is written with extensibility in mind. It tries to follow the Single Responsibility Principle (Robert C. Martin) and making it easy, to develop e.g. a new analyzer for a piece of statistics.

After finishing the setup process, there are three main phases to Statshandler. Receiving and preprocessing messages into available in-memory data. After that, the analysis phase takes place. At last, alerts generated are published. There is an optional forth phase, in which statistics are published to Redis, too.

2.1 New Analyzer

The main extension capability is in the analysis phase, providing a new analyzer to be able to find new attacks.

To develop a new analyzer, it can be added to the `src.Analyzer` module. It must inherit from `src.Analyzer.BaseAnalyzer` and build on that. There are three necessary steps in the subclass: - set an alerting function while calling the `__init__` method of the superclass - override the `analyze` method with a custom one - Adding the class to `Statshandler.initiate:to_initiate`

Setting the alerting function can be done like this: `super(MyAnalyzer, self).__init__(alerter=self._alert_limiter, limit=limit)` There are three different alerting functions implemented at the moment, each requiring their own set of arguments. These are specified in their method documentation.

All analyzers are initiated at startup. Therefore, the new analyzer has to be added to the list `Statshandler.initiate:to_initiate` with the following Syntax: `[(MyAnalyzer, (argument1, ...)), ...]`. Then, a analyzer will be instantiated during startup with the arguments provided in the argument set.

All analyzing logic has to start in the `analyze` method, as it is called automatically by the rest of the program while running in the mainloop.

DOCUMENTED MODULES

In the following sections, the modules contained within the Statshandler project will be documented.

3.1 Statshandler

`StatsHandler.setup_logging(loglevel, set_streamhandler=True)`

Set up basic logging functionality.

This method sets up default logging facilities using the built in logging module. A default loglevel can be set, and the decision whether to log to a file switched on and of.

Parameters

- **loglevel** – The loglevel that will end up in the outputed logs
- **set_streamhandler** – If set, a file will be created named `/var/log/statshandler/statshandler.log`

`StatsHandler.parseargs()`

Parsing arguments passed

Using the built in argparse module, passed arguments are parsed. Simply calling this method is sufficient, the results will be returned.

Returns A namespace object containing all parsed arguments

`StatsHandler.mainloop(pubsub, redis_instance, redis_channel, gather_stats, analyzers)`

The main event loop

Parameters

- **pubsub** – The redis publish-subscribe object from `setup_redis`.
- **redis_instance** – The redis instance itself.
- **redis_channel** – The redis channel/key to publish to
- **gather_stats** – A boolean whether to gather statistics or not
- **analyzers** – A list of instantiated analyzers being derivatives from `Analyzers.BaseAnalyzer`

`StatsHandler.main()`

The main method.

First, arguments passed to the system are parsed. Then, logging is set up. The connection to redis is established by calling `setup_redis()` and the main event loop started.

Returns The system exit code

3.2 src.Monitoring

`src.Monitoring.analyze` (*statistics, analyzers*)

Analyzing the provided stats to alert on.

Parameters

- **statistics** – The data in a dictionary to analyze
- **analyzers** – A list of tripples in the following format: (analyzer: callable, arguments: list, filter: str)

Returns A list of alerts to publish

`src.Monitoring.collect_port_stats` (*statistics*)

Gathering statistics about OVS ports.

Parameters **statistics** – A dictionary to fetch statistics from.

Returns A list of statistics gathered.

`src.Monitoring.collect_portaggregate_stats` (*statistics*)

Gathering statistics on flows used on ports.

Parameters **statistics** – A dictionary to fetch statistics from.

Returns A list of statistics gathered.

`src.Monitoring.collect_stats` (*statistics, collectors*)

Applying a list of callable collectors to the statistics message.

Parameters

- **statistics** – A dictionary of received statistics.
- **collectors** – A list of callables in the format [(callable, (argument1, ...),filter), ...].

Returns A list of statistics gathered.

3.3 src.Analyzers

All available analyzers are shown here.

First, a base class is provided to give other analyzers a standardize whay of checking for setting up an alerting mechanism and generating alert messages.

class `src.Analyzers.BaseAnalyzer` (*alerter, **alerter_setup*)

`__init__` (*alerter, **alerter_setup*)

A basis for other analyzers to inherit from.

Providing three different alerting methods, namely a limit, a range and a standard deviation, all sub-classes can use this reference implementation.

Parameters

- **alerter** – choose an alerting method like `self._alert_limiter, ...`. For a chosen alerter, additional values have to be set. These are specified in their method documentation.
- **alerter_setup** – additional arguments needed to set up the alerter. These differ for alerters and are specified in their method documentation.

`_alert_limiter` (*value*)

Generating an alert by comparing to a simple limit.

Necessary to use this alerter is to specify a limit when instantiating the class by providing the constructor with `limit=<n>`.

As called, an alert is generated if value is smaller than self.limit. Internally, this reverts to the range alerter, as a limit can be expressed with a bottom value of - infinity and a top value of limit.

Parameters *value* – The newest value.

Returns True, if an alert is necessary.

`_alert_range` (*value*)

Generating an alert by comparing value to a given range of valid data.

Necessary to use this alerter is to specify a top and bottom limit when instantiating the class. This is done by providing the constructor with `top=<n>` and `bottom=<m>`; n and m have to be floats.

This alerts, if value is `_not_` inside the range. If value is either supremum or infimum of the range, no alert is generated. For example, let `value = 5.0`. For `bottom = 4.0` and `top = 5.0`, no alert is generated. If value changes to 5.1, an alert is generated.

Parameters *value* (*float*) – The new value to compare against.

Returns True, if value not in [bottom, ..., top]

`_alert_standard_dev` (*new_value*)

Generating an alert by checking against a range based on the standard deviation.

Necessary to use this alerter is to specify `last_n` values and a multiplier. `Last_n` must be an integer and is used to determine, how many last values to keep. A range is then determined by computing the mean of this n last values. Adding and subtracting the standard deviation of the list of last values multiplied by the set multiplier gives a range of `mean +/- (m * standard_dev)`. A new value is only added to the list of last values, if it does not generate an alert.

Parameters *new_value* – The newest value to check against.

Returns True, if `new_value` is not in `mean +/- (m * standard_dev)`.

`analyze` (*new_statistics*)

The `analyze` method to be overridden by all subclasses.

This method is called by the program logic, to get a list of alerts back if any are present. It `_needs_` to be overridden in all subclasses.

Parameters *new_statistics* – The newest statistics gathered as a dictionary.

Returns A list of alerts.

`generate_alert_message` (*signature_id*, *additional_info*={})

Base method to generate alert messages

A signature is the minimum necessary to be able to identify an alert. Additional information can be passed as a dictionary to be merged into the alert message returned.

Parameters

- **`signature_id`** – An ID to identify the alert.
- **`additional_info`** – A dictionary of additional information.

Returns A combined alert message ready to be used in the output phase.

class `src.Analyzers.PortsAnalyzer` (*limit*, *excluded_ports*=[])

`__init__` (*limit*, *excluded_ports*=[])

This analyzer can handle single port statistics using a limit as alerter.

Parameters

- **`limit`** – The limit to alert on.
- **`excluded_ports`** – A list of ports not to produce alerts for.

analyze (*new_statistics*)

Single port statistics are analysed here.

For every switch and every port, statistics are gathered and alerts generated if a port surpasses the limit value with its `rx_packets_delta`. This is counting packets received by each port, even if no flow matched against them or a flow matched, but they were dropped.

Parameters `new_statistics` – The collected statistics on one Port.

`class src.Analyzers.PortAggregateAnalyzer` (*limit*)

__init__ (*limit*)

This analyzer can handler aggregated statistics combining port and flow information based on a limiter.

Parameters `limit` – The limit to alert on.

analyze (*new_statistics*)

Aggregated port statistics about used flows per port are analysed here.

For every switch and every port on them the limit value is checked against the `flow_count` usage of this port.

Parameters `new_statistics` – The dictionary containing raw statistics.

3.4 src.redis

`src.redis.setup_redis` (*ip, port, key, max_retries=30, retry=0*)

This will connect to a Redis instance and return the connection object.

Establishing a stable connection with a Redis instance is the goal. Therefore, the method will retry at most `max_retries` times.

Parameters

- **ip** – The IP of the Redis server to connect to.
- **port** – The port on which Redis is running.
- **key** – The redis channel/key to subscribe the a PubSub object to.
- **max_retries** – The maximum number of retries after which connection attempts will fail.
- **retry** – The current number of retries executed.

Returns A set of (`redis_instance`, `pubsub`)

`src.redis._fetch_message` (*pubsub*)

Fetching a single message from redis.

This fetches a single message and ensures, it has a valid data field.

Parameters `pubsub` – The pubsub object created by `setup_redis` to fetch the message from.

Returns A complete and data filled message.

Return type Python dict

`src.redis._parse_message` (*message*)

Parsing a message to get its valuable parts.

A message received from Redis contains information in json format and metadata from Redis, extracting the data is done here.

Parameters `message` – Raw message received from redis.

Returns A dictionary containing parsed data

`src.redis.get_statistics` (*pubsub*)

Returning a fully processed Redis message.

Parameters `pubsub` – The pubsub object created by `setup_redis` to fetch the message from.

Returns A received and parsed message.

`src.redis.publish_messages` (*messages, redis_instance, channel*)

Publish a list of messages via Redis.

Parameters

- **messages** – A list of json-dumpable messages to send.
- **redis_instance** – The instance to publish the messages to.
- **channel** – The Redis channel to publish to.

Returns

PYTHON MODULE INDEX

S

`src.Analyzers`, 6
`src.Monitoring`, 6
`src.redis`, 8
`StatsHandler`, 5

D. SynFlood Documentation

SynFlood Documentation

Release 1.0.0rc0

Christoph Girstenbrei

Dec 09, 2016

CONTENTS:

- 1 Usage Documentation** **1**
- 1.1 Requirements 1
- 1.2 Usage 1

- 2 Documented Modules** **3**
- 2.1 SynFlood 3

- Python Module Index** **5**

USAGE DOCUMENTATION

This script is intended to be used as a testing tool to perform SYN flood attacks against a known host.

1.1 Requirements

Required to run this script is at least Python 3.4.

1.2 Usage

The script can be run in standalone mode using the following options:

- `-l` : Setting the loglevel to one of {critical,warning,error,debug,info}
- `-s` : The source IP address to use or 'r' for random to spoof the address
- `-t` : The target IP address to attack
- `-p` : The source TCP port to use or 'r' for random to switch TCP port for every connection
- `-P` : The destination TCP port to use or 'r' for random to switch TCP port for every connection
- `-n` : The number of packages to send to the target
- `-w` : The wait period in between packages in milliseconds

To attack a <target> IP with 1000000 packages from different TCP source ports to TCP port 80, the following can be used: `./SynFlood.py -s r -t <target> -n 1000000 -p r -P 80`

DOCUMENTED MODULES

In the following sections, the modules contained within the SynFlood project will be documented.

2.1 SynFlood

This software is based on a script developed by Silver Moon. The following is the original copyright message:

Syn flood program in python using raw sockets (Linux)

Silver Moon (m00n.silv3r@gmail.com) source is <http://www.binarytides.com/python-syn-flood-program-raw-sockets-linux/>

SynFlood.**parseargs** ()

Parsing arguments passed

Using the built in argparse module, passed arguments are parsed. Simply calling this method is sufficient, the results will be returned.

Returns A namespace object containing all parsed arguments

PYTHON MODULE INDEX

S

SynFlood, 3

E. FlowFlood Documentation

FlowFlood Documentation

Release 1.0.0rc0

Christoph Girstenbrei

Dec 09, 2016

CONTENTS:

- 1 Usage Documentation** **1**
- 1.1 Requirements 1
- 1.2 Usage 1

- 2 Documented Modules** **3**
- 2.1 FlowFlood 3
- 2.2 src.netio 4

- Python Module Index** **5**

USAGE DOCUMENTATION

This script is intended to perform attacks in an SDN environment. It generates Flows in a MAC-learning SDN switch by performing a TCP handshake and sending a package back and forth.

1.1 Requirements

Required to run this script is at least Python 3.4.

1.2 Usage

The script can be run in standalone mode using the following options:

- `-l` : Setting a loglevel from one of {info,debug,error,critical,warning}
- `-m` : Running in master mode.

The simplest setup is having one instance running with `-m` as master and one without.

DOCUMENTED MODULES

In the following sections, the modules contained within the Statshandler project will be documented.

2.1 FlowFlood

`FlowFlood.setup_logging (loglevel, set_streamhandler=True)`

Set up basic logging functionality.

This method sets up default logging facilities using the built in logging module. A default loglevel can be set, and the decision whether to log to a file switched on and of.

Parameters

- **loglevel** – The loglevel that will end up in the outputed logs
- **set_streamhandler** – If set, a file will be created named `/var/log/statshandler/statshandler.log`

`FlowFlood.parseargs ()`

Parsing arguments passed

Using the built in argparse module, passed arguments are parsed. Simply calling this method is sufficient, the results will be returned.

Returns A namespace object containing all parsed arguments

`FlowFlood.mainloop (args)`

The main event loop.

If this is running as the master, the MAC address is changed, a socket set up and a connection is waited for. As a connection is established, one package is received and one sent.

If this is running in slave mode, connection to master is attempted. On success, a package is sent and one received.

Parameters **args** – Argumentes passed to the script.

`FlowFlood.main ()`

The main method.

First, arguments passed to the system are parsed. Then, logging is set up. The mainloop is then started, to run until keyboard interrupt.

Returns The system exit code

2.2 src.netio

`src.netio.setup_socket()`

Creating a TCP socket.

Returns A TCP socket.

`src.netio.rand_mac()`

Generating a random MAC address in private MAC address space.

Returns The random MAC as a string.

`src.netio.change_hwaddr(interface='eth1')`

Changing the MAC on interface.

Parameters `interface` – The one to change the MAC on.

`src.netio.setup_socket_master(src_ip)`

Setting up a server socket.

This sets up a server socket listening on address `src_ip` and a random port.

Parameters `src_ip` – The ip to bind to.

Returns A server socket.

`src.netio.yld_port()`

To generate a series of random ports, this returns a generator.

Returns A generator object returning ints from 1025 to 32768

`src.netio.next_port()`

Giving back the next port as an integer.

Returns Integer form 1025 to 32768.

`src.netio.stable_connect(sock, dst_ip, dst_port)`

Establishing a tcp connection to `dst_ip`.

This method tries to connect to a TCP socket, catching many exceptions possibly occurring during this operations.

Parameters

- `sock` – The socket to use to connect via.
- `dst_ip` – The IP to connect to.
- `dst_port` – The port to connect to.

`src.netio.clear_arp(ip)`

Clearing the arp cache of the host.

Parameters `ip` – IP that gets deleted

Returns The return code of the arp clearing command

PYTHON MODULE INDEX

f

`FlowFlood`, 3

s

`src.netio`, 4

F. VM Vagrant File

```
1 # -*- mode: ruby -*-
2 # vi: set ft=ruby :
3
4 # All Vagrant configuration is done below. The "2" in Vagrant.configure
5 # configures the configuration version (we support older styles for
6 # backwards compatibility). Please don't change it unless you know what
7 # you're doing.
8 Vagrant.configure(2) do |config|
9 # The most common configuration options are documented and commented below.
10 # For a complete reference, please see the online documentation at
11 # https://docs.vagrantup.com.
12
13 # Every Vagrant development environment requires a box. You can search for
14 # boxes at https://atlas.hashicorp.com/search.
15 config.vm.box = "bento/ubuntu-16.04"
16 config.vm.box_version = "2.2.9"
17 #config.vm.synced_folder ".", "/vagrant"
18
19 # Disable automatic box update checking. If you disable this, then
20 # boxes will only be checked for updates when the user runs
21 # 'vagrant box outdated'. This is not recommended.
22 config.vm.box_check_update = false
23
24 # Create a forwarded port mapping which allows access to a specific port
25 # within the machine from a port on the host machine. In the example below,
26 # accessing "localhost:8080" will access port 80 on the guest machine.
27 #config.vm.network "forwarded_port", guest: 80, host: 8080
28 #config.vm.network "forwarded_port", guest: 81, host: 8081
29
30 # Create a private network, which allows host-only access to the machine
31 # using a specific IP.
32 config.vm.network "private_network", ip: "192.168.33.10"
33
34 # Create a public network, which generally matched to bridged network.
35 # Bridged networks make the machine appear as another physical device on
36 # your network.
37 # config.vm.network "public_network"
38
39 # Share an additional folder to the guest VM. The first argument is
40 # the path on the host to the actual folder. The second argument is
41 # the path on the guest to mount the folder. And the optional third
42 # argument is a set of non-required options.
43 # config.vm.synced_folder "../data", "/vagrant_data"
44 config.vm.synced_folder ".", "/vagrant"
45
46 # Provider-specific configuration so you can fine-tune various
47 # backing providers for Vagrant. These expose provider-specific options.
48 # Example for VirtualBox:
49 #
50 config.vm.provider "virtualbox" do |vb|
51 # Display the VirtualBox GUI when booting the machine
52 # vb.gui = true
53 #
54 # Customize the amount of memory on the VM:
55 vb.cpus = 12
56 vb.memory = "8192"
```

F. VM Vagrant File

```
57 end
58 #
59 # View the documentation for the provider you are using for more
60 # information on available options.
61
62 # Define a Vagrant Push strategy for pushing to Atlas. Other push strategies
63 # such as FTP and Heroku are also available. See the documentation at
64 # https://docs.vagrantup.com/v2/push/atlas.html for more information.
65 # config.push.define "atlas" do |push|
66 #   push.app = "YOUR_ATLAS_USERNAME/YOUR_APPLICATION_NAME"
67 # end
68
69 # Enable provisioning with a shell script. Additional provisioners such as
70 # Puppet, Chef, Ansible, Salt, and Docker are also available. Please see the
71 # documentation for more information about their specific syntax and use.
72 config.vm.provision "shell", inline: <<-SHELL
73
74   PURPLE='\033[0;35m'
75   NC='\033[0m' # No Color
76   # Update system
77   printf "${PURPLE}|----- Updating -----|${NC}"
78
79   sudo timedatectl set-timezone Europe/Berlin
80   sudo cp /vagrant/sources_pre.list /etc/apt/sources.list
81   sudo apt-mark hold grub-pc
82   sudo apt-get update
83   sudo apt-get upgrade -y
84   sudo apt-get install -y -q tmux vim htop tofrodos curl
85   sudo apt-get install -y -q lxc
86
87   # Install OpenvSwitch
88   sudo apt-get install -y -q openvswitch-switch
89
90   # Install Vagrant
91   printf "${PURPLE}|----- Installing Vagrant -----|${NC}"
92   /vagrant/load_vagrant.sh
93   sudo dpkg -i vagrant_1.8.5_x86_64.deb
94   sudo vagrant plugin install vagrant-lxc
95
96   # Loading Cache
97   /vagrant/load_cache.sh
98
99   # Adding bridge to ovs
100  printf "${PURPLE}|----- Adding ovs bridge -----|${NC}"
101  sudo ovs-vsctl add-br switch0
102  sudo ovs-vsctl add-br switch1
103  sudo ovs-vsctl add-br switch-private
104
105  #Create Patch between switch0 and switch1
106  printf "${PURPLE}|----- Patching switches -----|${NC}"
107  sudo ovs-vsctl add-port switch0 patch0-1 -- set interface patch0-1 type=patch
108  options:peer=patch1-0
109  sudo ovs-vsctl add-port switch1 patch1-0 -- set interface patch1-0 type=patch
110  options:peer=patch0-1
111
112  # Adding veth-pair
113  printf "${PURPLE}|----- Adding veth-pair -----|${NC}"
114  ip link add ovs-acc type veth peer name ovs-ins
115  ip link add ovs-acc-priv type veth peer name ovs-ins-priv
116
117  # Setting up outside port
118  printf "${PURPLE}|----- Outside Port -----|${NC}"
119  ifconfig ovs-acc 10.10.10.10/16
120  ifconfig ovs-acc-priv 10.20.0.10/16
121  ifconfig ovs-acc up
122  ifconfig ovs-acc-priv up
```

```

121
122 # Setting up inside port
123 printf "${PURPLE}|----- Inside Port -----|${NC}"
124 ovs-vsctl add-port switch0 ovs-ins
125 ovs-vsctl add-port switch-private ovs-ins-priv
126 ifconfig ovs-ins up
127 ifconfig ovs-ins-priv up
128
129 # Allow access from internal network to external
130 printf "${PURPLE}|----- IpTables -> Inet -----|${NC}"
131 sudo iptables -t nat -A POSTROUTING -o enp0s3 -j MASQUERADE
132
133 #Connect eth0 to port and configure dhclient
134 printf "${PURPLE}|----- Set OVS Version -----|${NC}"
135 ovs-vsctl set bridge switch0
136     protocols=OpenFlow10,OpenFlow11,OpenFlow12,OpenFlow13,OpenFlow14,OpenFlow15
137 ovs-vsctl set bridge switch1
138     protocols=OpenFlow10,OpenFlow11,OpenFlow12,OpenFlow13,OpenFlow14,OpenFlow15
139 ovs-vsctl set bridge switch-private
140     protocols=OpenFlow10,OpenFlow11,OpenFlow12,OpenFlow13,OpenFlow14,OpenFlow15
141
142 # Copying the lxc-configs
143 printf "${PURPLE}|----- Copy lxc configs via unison -----|${NC}"
144 mkdir lxcs
145 cp -r /vagrant/lxcs/* lxcs
146
147 # Copy files to unix
148 printf "${PURPLE}|----- Chmod files -----|${NC}"
149 sudo su
150 cd /home/vagrant/lxcs/
151 chmod +x ifup ifdown lxc.conf ifup1 ifdown1
152 chmod +x /home/vagrant/lxcs/
153
154 printf "${PURPLE}|----- Converting files -----|${NC}"
155 fromdos -v ifup ifdown lxc.conf ifup1 ifdown1 ifup-private ifdown-private
156     test-main.sh main/resolv.conf
157 find . -name "*.sh" | xargs fromdos -v
158 find . -name "*.yaml" | xargs fromdos -v
159 find . -name "*.rules" | xargs fromdos -v
160 cp lxc.conf /etc/init/
161
162 # Starting Vagrant container
163 printf "${PURPLE}|----- Starting container -----|${NC}"
164 cd /home/vagrant/lxcs/main
165 sudo vagrant up dns redis
166 sudo vagrant up
167
168 # Forward Ports
169 printf "${PURPLE}|----- Forward Ports -----|${NC}"
170 ip='10.20.0.9'
171 # Grafana
172 iptables -t nat -A PREROUTING -i enp0s8 -p tcp --dport 80 -j DNAT --to $ip:3000
173 # Kibana
174 iptables -t nat -A PREROUTING -i enp0s8 -p tcp --dport 81 -j DNAT --to $ip:80
175 # Flooding Port
176 iptables -t nat -A PREROUTING -i enp0s8 -p tcp --dport 1234 -j DNAT --to
177     10.10.10.4:80
178
179 # Setting controller
180 printf "${PURPLE}|----- Setting OVS controller -----|${NC}"
181 sudo ovs-vsctl set-controller switch0 tcp:10.10.10.8:6633
182 sudo ovs-vsctl set-controller switch1 tcp:10.10.10.8:6633
183
184 # Install monitoring

```

F. VM Vagrant File

```
182 wget -q0 - https://artifacts.elastic.co/GPG-KEY-elasticsearch | sudo apt-key
    add -
183 sudo apt-get install -y apt-transport-https
184 echo "deb https://artifacts.elastic.co/packages/5.x/apt stable main" | sudo tee
    -a /etc/apt/sources.list.d/elastic-5.x.list
185 sudo apt-get update && sudo apt-get install metricbeat
186 cp /vagrant/lxcs/main/metricbeat.yml /etc/metricbeat/
187 systemctl start metricbeat
188 systemctl enable metricbeat
189
190 # Disable Standalone mode
191 #sudo ovs-vsctl set-fail-mode switch0 secure
192 #sudo ovs-vsctl set-fail-mode switch1 secure
193
194 # Limit the size of the flow table
195 sudo ovs-vsctl -- --id=@ft create Flow_Table flow_limit=200
    overflow_policy=refuse -- set Bridge switch0 flow_tables=0=@ft
196 sudo ovs-vsctl -- --id=@ft create Flow_Table flow_limit=200
    overflow_policy=refuse -- set Bridge switch1 flow_tables=0=@ft
197
198 # Deactivating default LXC-bridge
199 sudo ifconfig lxcbr0 down
200 #sudo brctl delbr lxcbr0
201
202 # Activating file sync
203 sudo chmod a+rx /home/vagrant/lxcs/sync-files.sh
204 sudo cp /home/vagrant/lxcs/sync-files.service /etc/systemd/system/
205 sudo systemctl daemon-reload
206 sudo systemctl enable sync-files
207 sudo systemctl start sync-files
208
209 # Deactivating syn cookies
210 sudo sysctl -w net.ipv4.tcp_syncookies=0
211 sudo systemctl restart networking
212
213 # Saving Cache
214 /vagrant/save_cache.sh
215
216 /home/vagrant/lxcs/test-main.sh
217
218 printf "Machine deployed."
219
220 SHELL
221 end
```

G. LXC's Vagrant File

```
1 # -*- mode: ruby -*-
2 # vi: set ft=ruby :
3
4 # IP-Table #
5 # DNS      | 10.0.3.61, 10.10.10.2
6 # attackone | 10.0.3.68, 10.10.10.3
7 # Manager   | 10.0.3.14, 10.10.10.8, 10.20.0.8
8 # Mitigation | 10.0.3.187, 10.20.0.7
9 # Monitoring | 10.0.3.119, 10.10.10.5, 10.20.0.5
10 # Redis     | 10.0.3.245, 10.20.0.6
11 # Target    | 10.0.3.162, 10.10.10.4
12 # ELK       | 10.20.0.9
13
14
15
16 Vagrant.configure("2") do |config|
17
18   config.vm.define "dns" do |dns|
19     dns.vm.hostname = "dns"
20     dns.vm.box = "developerinlondon/ubuntu_lxc_xenial_x64"
21
22     dns.vm.provider :lxc do |lxc, override|
23       lxc.container_name = "dns"
24       lxc.customize 'start.auto', '1'
25       lxc.customize 'start.delay', '3'
26       lxc.customize 'network.type', 'veth'
27       lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup1'
28       lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown1'
29       lxc.customize 'network.veth.pair', 'dns1eth1'
30       lxc.customize 'network.flags', 'up'
31       lxc.customize 'network.ipv4', '10.10.10.2/16'
32       lxc.customize 'group', 'onboot'
33       lxc.customize 'cgroup.cpuset.cpus', '0'
34     end
35
36     dns.vm.provision "shell", inline: <<-SHELL
37
38     # Custom configuration here!
39     sudo timedatectl set-timezone Europe/Berlin
40     sudo sed -i 's|nameserver 10.0.3.1|nameserver 8.8.8.8|' /etc/resolv.conf
41     ip route de default
42     ip route add default via 10.10.10.10
43     ping -c 2 8.8.8.8
44     ping -c 2 www.google.de
45     sudo apt-get update
46     sudo apt-get install -y bind9 dnsutils
47     sudo sed -i 's|nameserver 8.8.8.8|nameserver 10.10.10.2|' /etc/resolv.conf
48     sudo cp -f /vagrant/dns/named.conf.options /etc/bind
49     sudo cp -f /vagrant/dns/named.conf.local /etc/bind
50     sudo mkdir /etc/bind/zones
51     sudo cp /vagrant/dns/db.sdn.local /etc/bind/zones
52     sudo cp /vagrant/dns/db.10.10.10 /etc/bind/zones
53     sudo systemctl start bind9
54   SHELL
55   end
56 end
```

G. LXC's Vagrant File

```
57 config.vm.define "target" do |target|
58   target.vm.hostname = "target"
59   target.vm.box = "developerinlondon/ubuntu_lxc_xenial_x64"
60
61   target.vm.provider :lxc do |lxc, override|
62     lxc.container_name = "target"
63     lxc.customize 'start.auto', '1'
64     lxc.customize 'start.delay', '3'
65     lxc.customize 'network.type', 'veth'
66     lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup1'
67     lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown1'
68     lxc.customize 'network.veth.pair', 'target1eth1'
69     lxc.customize 'network.flags', 'up'
70     lxc.customize 'network.ipv4', '10.10.10.4/16'
71     lxc.customize 'group', 'onboot'
72     lxc.customize 'cgroup.cpuset.cpus', '1'
73   end
74
75   target.vm.provision "shell", inline: <<-SHELL
76
77     # Custom configuration here!
78     sudo timedatectl set-timezone Europe/Berlin
79     sudo sed -i 's|nameserver 10.0.3.1|nameserver 8.8.8.8|' /etc/resolv.conf
80     ip route de default
81     ip route add default via 10.10.10.10
82     ping -c 2 8.8.8.8
83     ping -c 2 www.google.de
84     sudo apt-get install -y -q tcpdump iperf
85     cp /vagrant/resolv.conf /etc
86     sudo apt-get update
87     sudo apt-get install -y nginx
88     sudo systemctl enable nginx
89     sudo systemctl start nginx
90     sudo cp /vagrant/target/BA.mp4 /var/www/html
91     sudo cp /vagrant/target/index.nginx-debian.html /var/www/html
92
93   SHELL
94 end
95
96 config.vm.define "attackone" do |attackone|
97   attackone.vm.hostname = "attackone"
98   attackone.vm.box = "developerinlondon/ubuntu_lxc_xenial_x64"
99
100  attackone.vm.provider :lxc do |lxc, override|
101    lxc.container_name = "attackone"
102    lxc.customize 'start.auto', '1'
103    lxc.customize 'start.delay', '3'
104    lxc.customize 'network.type', 'veth'
105    lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup'
106    lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown'
107    lxc.customize 'network.veth.pair', 'attackone1eth1'
108    lxc.customize 'network.flags', 'up'
109    lxc.customize 'network.ipv4', '10.10.10.3/16'
110    lxc.customize 'group', 'onboot'
111    lxc.customize 'cgroup.cpuset.cpus', '2'
112  end
113
114  attackone.vm.provision "shell", inline: <<-SHELL
115
116    # Custom configuration here!
117    sudo timedatectl set-timezone Europe/Berlin
118    sudo sed -i 's|nameserver 10.0.3.1|nameserver 8.8.8.8|' /etc/resolv.conf
119    ip route de default
120    ip route add default via 10.10.10.10
121    ping -c 2 8.8.8.8
122    ping -c 2 www.google.de
```

```

123     sudo apt-get install -y -q iperf iptables tcpdump ethtool
124     cp /vagrant/resolv.conf /etc
125
126     #sudo ethtool -K eth1 gro off
127     #sudo ifconfig eth1 promisc
128
129     SHELL
130 end
131
132 config.vm.define "attacktwo" do |attacktwo|
133     attacktwo.vm.hostname = "attacktwo"
134     attacktwo.vm.box = "developerinlondon/ubuntu_lxc_xenial_x64"
135
136     attacktwo.vm.provider :lxc do |lxc, override|
137         lxc.container_name = "attacktwo"
138         lxc.customize 'start.auto', '1'
139         lxc.customize 'start.delay', '3'
140         lxc.customize 'network.type', 'veth'
141         lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup'
142         lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown'
143         lxc.customize 'network.veth.pair', 'attacktwo1eth1'
144         lxc.customize 'network.flags', 'up'
145         lxc.customize 'network.ipv4', '10.10.10.6/16'
146         lxc.customize 'group', 'onboot'
147         lxc.customize 'cgroup.cpuset.cpus', '3'
148     end
149
150     attacktwo.vm.provision "shell", inline: <<-SHELL
151
152     # Custom configuration here!
153     sudo timedatectl set-timezone Europe/Berlin
154     sudo sed -i 's|nameserver 10.0.3.1|nameserver 8.8.8.8|' /etc/resolv.conf
155     ip route de default
156     ip route add default via 10.10.10.10
157     ping -c 2 8.8.8.8
158     ping -c 2 www.google.de
159     sudo apt-get install -y -q iperf iptables tcpdump ethtool
160     cp /vagrant/resolv.conf /etc
161
162     #sudo ethtool -K eth1 gro off
163     #sudo ifconfig eth1 promisc
164     SHELL
165 end
166
167 config.vm.define "monitoring" do |monitoring|
168     monitoring.vm.hostname = "monitoring"
169     monitoring.vm.box = "developerinlondon/ubuntu_lxc_xenial_x64"
170
171     monitoring.vm.provider :lxc do |lxc, override|
172         lxc.container_name = "monitoring"
173         lxc.customize 'start.auto', '1'
174         lxc.customize 'start.delay', '3'
175         lxc.customize 'network.type', 'veth'
176         lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup'
177         lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown'
178         lxc.customize 'network.veth.pair', 'monit1eth1'
179         lxc.customize 'network.flags', 'up'
180         lxc.customize 'network.ipv4', '10.10.10.5/16'
181         lxc.customize 'network.type', 'veth'
182         lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup-private'
183         lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown-private'
184         lxc.customize 'network.veth.pair', 'monit1eth2'
185         lxc.customize 'network.flags', 'up'
186         lxc.customize 'network.ipv4', '10.20.0.5/16'
187         lxc.customize 'group', 'onboot'
188         lxc.customize 'cgroup.cpuset.cpus', '4'

```

G. LXC's Vagrant File

```
189     end
190
191     monitoring.vm.provision "shell", inline: <<-SHELL
192         PURPLE='\033[0;35m'
193         NC='\033[0m' # No Color
194         # Standard Base Config
195         printf "${PURPLE}|----- Monitoring: Base setup -----|${NC}"
196         sudo timedatectl set-timezone Europe/Berlin
197         sudo sed -i 's|nameserver 10.0.3.1|nameserver 8.8.8.8|' /etc/resolv.conf
198         sudo ip link set down eth2 && sudo ip link set up eth2
199         ip route del default && ip route add default via 10.20.0.10
200         sudo apt-get install -y iperf
201         sudo apt-get update
202         sudo apt-get install -y ethtool
203         sudo apt-get install -y python3-pip
204         sudo pip3 install --upgrade pip
205         sudo pip3 install -r /vagrant/monitoring/statistical/requirements.txt
206
207         # Configure for packet sniffing
208         printf "${PURPLE}|----- Monitoring: Promisc -----|${NC}"
209         sudo ethtool -K eth1 gro off
210         sudo ifconfig eth1 promisc
211
212         # Install & configure Suricata
213         printf "${PURPLE}|----- Monitoring: Suricata -----|${NC}"
214         apt install -y software-properties-common
215         sudo add-apt-repository -y ppa:oisf/suricata-stable
216         sudo apt-get update
217         sudo apt-get -y install suricata
218         sudo systemctl stop suricata
219         sudo cp /vagrant/monitoring/local.rules /etc/suricata/rules/
220         sudo cp /vagrant/monitoring/suricata.yaml /etc/suricata/
221         sudo sed -i 's/IFACE=eth0/IFACE=eth1/' /etc/default/suricata
222         sudo sed -i 's/LISTENMODE=af-packet/LISTENMODE=pcap/' /etc/default/suricata
223
224         sleep 2
225         sudo systemctl enable suricata
226         sudo systemctl start suricata
227
228         sleep 2
229
230         systemctl restart suricata
231
232         cp /vagrant/monitoring/statistical/statshandler.service /etc/systemd/system/
233         systemctl daemon-reload
234         systemctl enable statshandler
235         systemctl start statshandler
236         cp /vagrant/resolv.conf /etc
237     SHELL
238     end
239
240     config.vm.define "redis" do |redis|
241         redis.vm.hostname = "redis"
242         redis.vm.box = "developerinlondon/ubuntu_lxc_xenial_x64"
243
244         redis.vm.provider :lxc do |lxc, override|
245             lxc.container_name = "redis"
246             lxc.customize 'start.auto', '1'
247             lxc.customize 'start.delay', '3'
248             lxc.customize 'network.type', 'veth'
249             lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup-private'
250             lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown-private'
251             lxc.customize 'network.veth.pair', 'redis1eth1'
252             lxc.customize 'network.flags', 'up'
253             lxc.customize 'network.ipv4', '10.20.0.6/16'
254             lxc.customize 'group', 'onboot'
```



```

255     lxc.customize 'cgroup.cpuset.cpus', '5'
256 end
257
258 redis.vm.provision "shell", inline: <<-SHELL
259     # Custom configuration here!
260     sudo timedatectl set-timezone Europe/Berlin
261     sudo sed -i 's|nameserver 10.0.3.1|nameserver 8.8.8.8|' /etc/resolv.conf
262     sudo ip link set down eth1 && sudo ip link set up eth1
263     ip route del default && ip route add default via 10.20.0.10
264
265     /vagrant/redis/setup-redis.sh
266
267
268 SHELL
269 end
270
271 config.vm.define "elk" do |elk|
272     elk.vm.hostname = "elk"
273     elk.vm.box = "developerinlondon/ubuntu_lxc_xenial_x64"
274
275     elk.vm.provider :lxc do |lxc, override|
276         lxc.container_name = "elk"
277         lxc.customize 'start.auto', '1'
278         lxc.customize 'start.delay', '3'
279         lxc.customize 'network.type', 'veth'
280         lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup-private'
281         lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown-private'
282         lxc.customize 'network.veth.pair', 'elk1eth1'
283         lxc.customize 'network.flags', 'up'
284         lxc.customize 'network.ipv4', '10.20.0.9/16'
285         lxc.customize 'group', 'onboot'
286         lxc.customize 'cgroup.cpuset.cpus', '6'
287     end
288
289     elk.vm.provision "shell", inline: <<-SHELL
290         # Custom configuration here!
291         sudo timedatectl set-timezone Europe/Berlin
292         sudo sed -i 's|nameserver 10.0.3.1|nameserver 8.8.8.8|' /etc/resolv.conf
293         sudo ip link set down eth1 && sudo ip link set up eth1
294         ip route del default && ip route add default via 10.20.0.10
295
296         printf "${PURPLE}|----- Monitoring: Copy elk setup -----|${NC}"
297         cd /home/vagrant/
298         cp /vagrant/elk/*.sh ./
299         # Make files executable
300         sudo chmod +x *.sh
301
302         # Install & configure elk stack
303         printf "${PURPLE}|----- Monitoring: ELK setup -----|${NC}"
304         sudo ./setup-elk.sh
305         sudo ./configure-grafana.sh
306
307     SHELL
308 end
309
310 config.vm.define "manager" do |manager|
311     manager.vm.hostname = "manager"
312     manager.vm.box = "developerinlondon/ubuntu_lxc_xenial_x64"
313
314     manager.vm.provider :lxc do |lxc, override|
315         lxc.container_name = "manager"
316         lxc.customize 'start.auto', '1'
317         lxc.customize 'start.delay', '3'
318         lxc.customize 'network.type', 'veth'
319         lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup'
320         lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown'

```

G. LXC's Vagrant File

```
321     lxc.customize 'network.veth.pair', 'mane1eth1'
322     lxc.customize 'network.flags', 'up'
323     lxc.customize 'network.ipv4', '10.10.10.8/16'
324     lxc.customize 'network.type', 'veth'
325     lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup-private'
326     lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown-private'
327     lxc.customize 'network.veth.pair', 'mane1eth2'
328     lxc.customize 'network.flags', 'up'
329     lxc.customize 'network.ipv4', '10.20.0.8/16'
330     lxc.customize 'group', 'onboot'
331     lxc.customize 'cgroup.cpuset.cpus', '7'
332 end
333
334 manager.vm.provision "shell", inline: <<-SHELL
335
336     # Custom configuration here!
337     sudo timedatectl set-timezone Europe/Berlin
338     sudo sed -i 's|nameserver 10.0.3.1|nameserver 8.8.8.8|' /etc/resolv.conf
339     sudo ip link set down eth1 && sudo ip link set up eth1
340     ip route del default && ip route add default via 10.20.0.10
341     echo "install iperf"
342     sudo apt-get install -y iperf
343     echo "Installing Ryu..."
344     LC_ALL=C
345     sudo apt-get update
346     sudo apt-get install -y python3-dev
347     sudo apt-get install -y python3-pip
348     sudo apt-get install -y python3-eventlet
349     sudo apt-get install -y python3-routes
350     sudo apt-get install -y python3-webob
351     sudo apt-get install -y python3-paramiko
352     sudo pip3 install --upgrade pip
353     sudo pip3 install --upgrade six
354     sudo pip3 install redis
355     sudo pip3 install hiredis
356     sudo pip3 install ryu
357     sudo pip3 install --upgrade tinyrpc
358     cp /vagrant/manager/controller.py ./
359
360     cp /vagrant/manager/controller.service /etc/systemd/system/
361     systemctl daemon-reload
362     systemctl enable controller
363     systemctl start controller
364     cp /vagrant/resolv.conf /etc
365     SHELL
366 end
367
368 config.vm.define "mitigation" do |mitigation|
369     mitigation.vm.hostname = "mitigation"
370     mitigation.vm.box = "developerinlondon/ubuntu_lxc_xenial_x64"
371
372     mitigation.vm.provider :lxc do |lxc, override|
373         lxc.container_name = "mitigation"
374         lxc.customize 'start.auto', '1'
375         lxc.customize 'start.delay', '3'
376         lxc.customize 'network.type', 'veth'
377         lxc.customize 'network.script.up', '/home/vagrant/lxcs/ifup-private'
378         lxc.customize 'network.script.down', '/home/vagrant/lxcs/ifdown-private'
379         lxc.customize 'network.veth.pair', 'mit1eth1'
380         lxc.customize 'network.flags', 'up'
381         lxc.customize 'network.ipv4', '10.20.0.7/16'
382         lxc.customize 'group', 'onboot'
383         lxc.customize 'cgroup.cpuset.cpus', '8'
384     end
385
386     mitigation.vm.provision "shell", inline: <<-SHELL
```

```
387     # Custom configuration here!
388     sudo timedatectl set-timezone Europe/Berlin
389     sudo sed -i 's|nameserver 10.0.3.1|nameserver 8.8.8.8|' /etc/resolv.conf
390     sudo ip link set down eth1 && sudo ip link set up eth1
391     ip route del default && ip route add default via 10.20.0.10
392
393     sudo apt-get update
394     sudo apt-get install -y iperf
395     sudo apt-get install -y python3-pip
396     sudo pip3 install --upgrade pip
397     sudo pip3 install --upgrade requests hiredis redis
398
399     cp /vagrant/mitigation/miti*.service /etc/systemd/system/
400     systemctl daemon-reload
401     systemctl enable miti-stats
402     systemctl start miti-stats
403     systemctl enable miti-ids
404     systemctl start miti-ids
405     systemctl enable miti-mit
406     systemctl start miti-mit
407     SHELL
408     end
409 end
```

Acronyms

CPU	Central Processing Unit
TLS	Transport Layer Security
DNS	Domain Name System
DoS	Denial Of Service
DDoS	Distributed Denial of Service
EVE	Extensible Event Format
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IDMEF	Intrusion Detection Message Exchange Format
IDS	Intrusion Detection System
IP	Internet Protocol
JSON	JavaScript Object Notation
LXC	Linux Container
MAC	Media-Access-Control
NIC	Network Interface Card
NOS	Network Operating System
OF	OpenFlow
OVS	OpenVSwitch
SDN	Software Defined Network
SNMP	Simple Network Management Protocol
SRP	Single Responsibility Principle
TCP	Transport Control Protocol
TES	Triple Exponential Smoothing
TLS	Transport Layer Security
VM	Virtual Machine
XML	Extensible Markup Language

List of Figures

2.1. Roles of the control, management, and data planes. From [GB14]	3
2.2. SDN architecture adapted from [Kre+15]	5
2.3. Flow Table and its match fields and actions. Based on [NG13]	7
2.4. Packet drop due to flow flooding at 100 Mbps. From [KKS13]	9
2.5. Components of a DDoS attack. Based on [Lau+00]	10
2.6. TCP three-way handshake.	11
4.1. Application stack concept	15
4.2. Information flow through the application stack	17
4.3. Test-setup of the virtual network	18
4.4. Flow of alert messages	19
5.1. Monitoring in the concept	29
5.2. Monitoring concept	30
5.3. Range by standard deviation	32
5.4. TES. Based on [Fil+13][6.4.3.5]	33
5.5. Range by fixed percentage	33
5.6. Monitoring implementation	35
5.7. DoS attack without any mitigation	38
5.8. Flow of alert messages	39
5.9. SYN Flood with <code>track by_dst</code>	41
5.10. SYN Flood with alerts	42
5.11. SYN Flood with Statshandler	43
5.12. Flow flood with Statshandler	44
6.1. Mitigation in the concept	47
6.2. Concept of the mitigation component	48
6.3. Process of a TCP SYN Flood Mitigation	53
6.4. Mitigation flow for a TCP SYN Flood	53
6.5. Network load during a TCP SYN flood with and without mitigation	54
6.6. Alerts during mitigated TCP SYN Flood	54
6.7. Process of a Distributed TCP SYN Flood Mitigation	55
6.8. Workflow for the mitigation of a Distributed TCP SYN Flood	56
6.9. Querying the webservice during an attack	56
6.10. Network load during a distributed TCP SYN flood with and without mitigation	57
6.11. Alerts during mitigated Distributed TCP SYN Flood	57
6.12. Response time during flow flooding attack	59
6.13. Mitigation of a flow flooding attack	59

List of Tables

4.1. Responsibilities of application stack components	16
4.2. Example changes to secure setup	27
5.1. Comparison of monitoring methods	32
5.2. Comparison of analyzing methods	34

Bibliography

- [Amb+15] Moreno Ambrosin et al. “Lineswitch: Efficiently managing switch flow in software-defined networking while effectively tackling dos attacks”. In: *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security*. ACM. 2015, pp. 639–644.
- [AX15] Izzat Alsmadi and Dianxiang Xu. “Security of software defined networks: A survey”. In: *computers & security* 53 (2015), pp. 79–108.
- [BAM09] Theophilus Benson, Aditya Akella, and David A Maltz. “Unraveling the Complexity of Network Management.” In: *NSDI*. 2009, pp. 335–348.
- [BBC14] BBC. *Sony Pictures computer system hacked in online attack*. 2014. URL: <http://www.bbc.com/news/technology-30189029> (visited on 11/21/2016).
- [BCS13] Kevin Benton, L Jean Camp, and Chris Small. “Openflow vulnerability assessment”. In: *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*. ACM. 2013, pp. 151–152.
- [BG13] Marcelo Bagnulo and Alberto Garcia-Martinez. “SAVI: The IETF standard in address validation”. In: *IEEE Communications Magazine* 51.4 (2013), pp. 66–73.
- [Cho+14] S. R. Chowdhury et al. “PayLess: A low cost network monitoring framework for Software Defined Networks”. In: *2014 IEEE Network Operations and Management Symposium (NOMS)*. May 2014, pp. 1–9. DOI: 10.1109/NOMS.2014.6838227.
- [Cui+16] Yunhe Cui et al. “SD-Anti-DDoS: Fast and efficient DDoS defense in software-defined networks”. In: *Journal of Network and Computer Applications* 68 (2016), pp. 65–79.
- [DBP05] Thomas Dubendorfer, Matthias Bossardt, and Bernhard Plattner. “Adaptive distributed traffic control service for DDoS attack mitigation”. In: *19th IEEE International Parallel and Distributed Processing Symposium*. IEEE. 2005, 8–pp.
- [Deb+07] H. Debar et al. *The Intrusion Detection Message Exchange Format*. Language. Online. RFC. Internet Engineering Steering Group, Mar. 2007. URL: <https://tools.ietf.org/html/rfc4765>.
- [Edd06] Wesley M Eddy. “Defenses against TCP SYN flooding attacks”. In: *The Internet Protocol Journal* 9.4 (2006), pp. 2–16.
- [Fil+13] James J. Filliben et al. *Engineering Statistics Handbook*. National Institute of Standards and Technology. Oct. 2013. URL: <http://www.itl.nist.gov/div898/handbook//index.htm>.
- [GB14] Paul Goransson and Chuck Black. *Software Defined Networks: A Comprehensive Approach*. Elsevier, 2014.

- [GDK] Tobias Guggemos, Vitalian Danciu, and Dieter Kranzlmüller. “Schichtung virtueller Maschinen zu Labor-und Lehrinfrastruktur”. In: *Gesellschaft für Informatik eV (GI) publishes this series in order to make available to a broad public recent findings in informatics (ie computer science and information systems), to document conferences that are organized in co-operation with GI and to publish the annual GI Award dissertation*. P. 35.
- [Gio+14] Kostas Giotis et al. “Combining OpenFlow and sFlow for an effective and scalable anomaly detection and mitigation mechanism on SDN environments”. In: *Computer Networks* 62 (2014), pp. 122–136.
- [Gua16] The Guardian. *Major cyber attack disrupts internet service across Europe and US*. 2016. URL: <https://www.theguardian.com/technology/2016/oct/21/ddos-attack-dyn-internet-denial-service> (visited on 11/21/2016).
- [Hab14] Itamar Haber. *Using stunnel to Secure Redis*. English. Mar. 2014. URL: <https://redislabs.com/blog/using-stunnel-to-secure-redis>.
- [Her+16] Andreas Herz et al. *Suricata User Guide*. OISF. 2016. URL: <https://suricata.readthedocs.io/en/latest/rules/index.html>.
- [Hon+15] Sungmin Hong et al. “Poisoning Network Visibility in Software-Defined Networks: New Attacks and Countermeasures.” In: *NDSS*. 2015.
- [Hu+14] Hongxin Hu et al. “FLOWGUARD: building robust firewalls for software-defined networks”. In: *Proceedings of the third workshop on Hot topics in software defined networking*. ACM. 2014, pp. 97–102.
- [ISO] ‘ISO/IEC’. *Information technology - Open Systems Interconnection - Basic Reference Model: The Basic Model*. ISO/IOC. URL: [http://standards.iso.org/ittf/PubliclyAvailableStandards/s020269_ISO_IEC_7498-1_1994\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/s020269_ISO_IEC_7498-1_1994(E).zip).
- [JM16] Damian Janowski and Michel Martens. *redis.io*. English. 2016. URL: <http://redis.io/>.
- [JMD14] Yosr Jarraya, Taous Madi, and Mourad Debbabi. “A survey and a layered taxonomy of software-defined networking”. In: *Communications Surveys & Tutorials, IEEE* 16.4 (2014), pp. 1955–1980.
- [KKS13] Rowan Klöti, Vasileios Kotronis, and Paul Smith. “Openflow: A security analysis”. In: *2013 21st IEEE International Conference on Network Protocols (ICNP)*. IEEE. 2013, pp. 1–6.
- [Kre+15] Diego Kreutz et al. “Software-defined networking: A comprehensive survey”. In: *Proceedings of the IEEE* 103.1 (2015), pp. 14–76.
- [Lau+00] Felix Lau et al. “Distributed denial of service attacks”. In: *Systems, Man, and Cybernetics, 2000 IEEE International Conference on*. Vol. 3. IEEE. 2000, pp. 2275–2280.
- [Lem+02] Jonathan Lemon et al. “Resisting SYN Flood DoS Attacks with a SYN Cache.” In: *BSDCon*. Vol. 2002. 2002, pp. 89–97.
- [LMK16] Wenjuan Li, Weizhi Meng, and Lam For Kwok. “A survey on OpenFlow-based Software Defined Networks: Security challenges and countermeasures”. In: *Journal of Network and Computer Applications* 68 (2016), pp. 126–139.

- [Mar03] Robert Cecil Martin. *Agile Software Development: Principles, Patterns, and Practices*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 2003. ISBN: 0135974445.
- [Mar16] Alex Marczinek. *Mitigation of attacks in the environment of Software Defined Networks*. PDF. Dec. 2016.
- [MBR16] Louis Marinos, Adrian Belmonte, and Evangelos Rekleitis. *ENISA Threat Landscape 2015*. 2016.
- [McK+08] Nick McKeown et al. “OpenFlow: enabling innovation in campus networks”. In: *ACM SIGCOMM Computer Communication Review* 38.2 (2008), pp. 69–74.
- [MR04] Jelena Mirkovic and Peter Reiher. “A taxonomy of DDoS attack and DDoS defense mechanisms”. In: *ACM SIGCOMM Computer Communication Review* 34.2 (2004), pp. 39–53.
- [NG13] Thomas D Nadeau and Ken Gray. *SDN: software defined networks.* ” O’Reilly Media, Inc.”, 2013.
- [Nyg+14] Anders Nygren et al. *OpenFlow Switch Specification Version 1.5.0 (Protocol version 0x06)*. PDF. Dec. 2014. URL: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-switch-v1.5.0.noipr.pdf>.
- [ONF14] ONF. *SDN Architecture Overview*. Tech. rep. Open Networking Foundation (ONF), Nov. 2014.
- [Sec16] Calyptix Security. *Top 7 Network Attack Types in 2016*. 2016. URL: <http://www.calyptix.com/top-threats/top-7-network-attack-types-2016/> (visited on 08/10/2016).
- [She+12] Justine Sherry et al. “Making middleboxes someone else’s problem: network processing as a cloud service”. In: *ACM SIGCOMM Computer Communication Review* 42.4 (2012), pp. 13–24.
- [Shi+13] Seungwon Shin et al. “AVANT-GUARD: scalable and vigilant switch flow management in software-defined networks”. In: *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM. 2013, pp. 413–424.
- [SMC74] W. P. Stevens, G. J. Myers, and L. L. Constantine. “Structured design”. In: *IBM Systems Journal* 13.2 (1974), pp. 115–139. ISSN: 0018-8670. DOI: 10.1147/sj.132.0115.
- [Sys15] CS Communication & Systems. *Prelude OSS — The Open Source Reference*. Online. 2015. URL: <http://www.prelude-siem.com/en/products/prelude-oss/>.
- [Tan03] Matthew Tanase. “IP spoofing: an introduction”. In: *Security Focus* 11 (2003).
- [Tea00] Nagios Plugins Team. *Nagios Plugin Development Guidelines*. Nagios Enterprises. 2000. URL: <https://nagios-plugins.org/doc/guidelines.html>.
- [VE06] Randal Vaughn and Gadi Evron. “DNS amplification attacks”. In: *Go online to http://www.isotf.org/news/DNS-Amplification-Attacks.pdf* (2006).

Bibliography

- [Ver16] Verisign. *Q3 2016 DDoS trends report: UDP Flood Attacks make up 49 percent of attacks*. 2016. URL: <https://blog.verisign.com/security/q3-2016-ddos-trends-report-udp-flood-attacks-make-up-49-percent-of-attacks/> (visited on 11/26/2016).
- [YBX11] Guang Yao, Jun Bi, and Peiyao Xiao. “Source address validation solution with OpenFlow/NOX architecture”. In: *2011 19th IEEE International Conference on Network Protocols*. IEEE. 2011, pp. 7–12.
- [ZJT13] Saman Taghavi Zargar, James Joshi, and David Tipper. “A survey of defense mechanisms against distributed denial of service (DDoS) flooding attacks”. In: *IEEE Communications Surveys & Tutorials* 15.4 (2013), pp. 2046–2069.