04-09-2017

Deliverable D8.3
Distributed Denial of Service Mitigation v1.0 Pilot

Deliverable D8.3

Contractual Date:	31-07-2017
Actual Date:	04-09-2017
Grant Agreement No.:	731122
Work Package/Activity:	8/JRA2
Task Item:	Task 6
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	DFN (LRZ)
Document ID:	GN4-2-17-9F2F2
Authors:	David Schmitz (DFN/LRZ), Tomáš Čejka (CESNET), Václav Bartoš (CESNET), Ivana Golub (PSNC), Evangelos Spatharas (GÉANT)

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 731122 (GN4-2).

Abstract

This document describes the pilot of the DDoS detection and mitigation tool. It starts with the explanation of the DDoS detection and mitigation architecture that includes Firewall-On-Demand (FoD), Network Security Handling and Response Process (NSHaRP) and Reputation Shield (RepShield), and continues focusing on the DDoS detection and mitigation pilot that is planned for the second reporting period of the GN4-2 project within Network Services Development Joint Research Activity (JRA2) Network Security Task (Task 6).

Table of Contents

Table of Figures

Executive Summary

The Distributed Denial of Service (DDoS) Detection and Mitigation pilot is based on GÉANT DDoS detection and mitigation architecture (detailed in Section 2), which includes: Firewall-On-Demand (FoD) [FLOWSPY], Network Security Handling and Response Process (NSHaRP) [GNNSHARP], Firewall Rule Updater (FRU) and Reputation Shield - RepShield [REPSHIELD]. The DDoS detection and mitigation pilot (the pilot) is planned for the first quarter of 2018. The development work and the pilot are conducted within the GN4-2 project Network Services Development Joint Research Activity (JRA2) Network Security Task (Task 6).

The main component of the detection and mitigation architecture is FoD (current versions: version 1.1, in production and version 1.5, in development). A progress report on FoD v1.5 development is presented in the GN4-2 Deliverable D8.2 [D8.2]. It is expected that FoD v1.5 will become a production version at the time of the pilot.

Therefore, this document presents the architecture and objectives of the upcoming pilot, based on FoD v1.6, which will enhance FoD v1.5 with an automated BGP FlowSpec rule proposal for DDoS mitigation, where the NREN NOC FoD user is only expected to check and apply proposed mitigation rules instead of having to manually enter them. The rules are proposed according to detected security events, (such as via NSHaRP) and information about network entities retrieved from several other publicly available sources, such as Whois and geolocation databases and blacklists related to SPAM senders, malware infection and botnet C&C servers. Detected DDoS attacks, which are the main concern of FoD, will be correlated and quality checked by the Reputation Shield (RepShield) tool, also developed in JRA2-T6, in cooperation with CESNET organisation. Based on the input from NSHaRP and RepShield, the Firewall Rule Updater will prepare rule proposals for FoD, and users will be prompted to apply the rules, thus speeding the overall DDoS detection and mitigation process.

The pilot of v1.6 will provide a new automated rule proposal functionality to a selected set of FoD test users. FoD and DDoS detection and mitigation users are GÉANT community organisations (with their own AS) and their NOC and CERT operators. A few organisations will participate in the pilot and the final list of participants will be defined before the beginning of the pilot. The users will test and check the tool based on defined pilot acceptance criteria. At the end of the pilot, the test users will share their experience via a survey, the results of which will also be used to validate the pilot.

The final results of the DDoS detection and mitigation pilot will be visible through the further development of DDoS detection and mitigation system and its components, especially through further versions of Firewall on Demand. The progress on the FoD development will be reported in the GN4-2 Quarterly Management Reports, as well as in the Deliverable D1.12 Assessment of GÉANT Service Catalogue.

1 Introduction

With the proliferation of Distributed Denial of Service (DDoS) attacks, an increased effort is placed on the design and construction of a successful detection and mitigation system. Such endeavours have been nourished through several generations of GÉANT projects, yielding a Firewall on Demand (FoD) tool, which is now complemented with other tools and solutions to form a DDoS detection and mitigation system. The further development of FoD and the DDoS detection and mitigation system is a part of Network Services Development Joint Research Activity (JRA2) in Network Security Task (Task 6) of the GN4-2 project [GN4-2].

FoD [FLOWSPY] v1.6 adds support for (semi-)automated rule proposal to the previous FoD v1.5. In addition to the existing functionality of manually entering mitigation rules in the web interface of FoD, users of FoD – i.e. NREN NOC administrators – will be able to access automatically generated proposals for mitigation rules created out of Network Security Handling and Response Process (NSHaRP) [GNNSHARP] events.

Through the DDoS detection and mitigation system the users are informed about the proposed rules, e.g. via e-mail, and they subsequently decide whether to apply these rules in the FoD web GUI. In FoD v1.6, the users are always required to manually apply the proposed rules because a rule's initial status is inactive by default. This new feature of FoD v1.6 makes the process of entering and applying the rules faster and easier, however, the final decision regarding what rules to apply is still up to the users.

In order to ensure a certain quality regarding rule granularity, as well as to improve the avoidance of false-positives, the RepShield [REPSHIELD] will be used in between the NSHaRP events and Firewall Rule Updater, which will prepare the rule proposal for FoD. RepShield is developed in JRA2-T6, in cooperation with CESNET.

The proposed solution is designed as multi-domain, taking into consideration all GÉANT project participating organisations, each with at least one network Autonomous Systems Number (ASN) and approximately 50 million end users. Alternative solutions were considered instead of, and together with, the proposed DDoS detection and mitigation system.

Commercial DDoS mitigation solutions carry several disadvantages: primarily single-domain, they are normally expensive for licensing or support for the GÉANT range of ASNs and/or end-users; they often do not support open standards such as BGP FlowSpec; they usually come only with a vendor-specific UI, which is normally not multi-tenant. Commercial solutions can hardly, if at all, be extended and/or customised in sufficient way and degree, especially regarding automation and integration with other tools or processes in the organisation.

Another option for DDoS mitigation is the usage access control lists (ACL), typically manually managed through a router command-line user interface (CLI). This, frequently used and effective solution is much slower, error-prone and requires interaction with GÉANT core NOC engineers (thus increasing their workload). The drawback of such a solution can be that it can create inconsistencies in configuration setup or removal in case of multi-homed networks. In the case of FoD, users can administrate the rules in multi-tenant-manner, efficiently controlling and directly managing their rules without interaction with the GÉANT NOC. Rules entered at one place (FoD UI) are automatically propagated to all routers in the GÉANT core network, so inconsistencies between routers will no longer occur. Users can get an overview of all of their active rules, so they can find, review and remove any unchecked active rules that remain.

The DDoS detection and mitigation system presented in this document is based on open source software, so its operation does not bring any high costs for software licenses. As this pilot will demonstrate, it is multi-tenant and easily extensible, especially regarding automation and integration with existing tools and processes. In addition, FoD can be run not only in GÉANT core, but also in NRENs or even in connected institutions in a multi-domain manner. Any of these instances could exchange and share BGP FlowSpec rules via eBGP [rfc5575] [rfc7674].

With these considerations, it has been concluded that the pilot of FoD v1.6 presented here has multiple advantages to alternative DDoS mitigation techniques and solutions.

The rest of this document is structured as follows: Section 2 describes the elements of the DDoS detection and mitigation architecture, including FoD, NSHaRP, Firewall Rule Updater and RepShield. Section 3 describes the DDoS detection and mitigation pilot. Section 4 presents the future work and the conclusions are summarised in Section 5.

2 DDoS Detection and Mitigation Architecture

Architecture for the detection and mitigation of DDoS attacks in GÉANT network is shown in Figure 2.1.

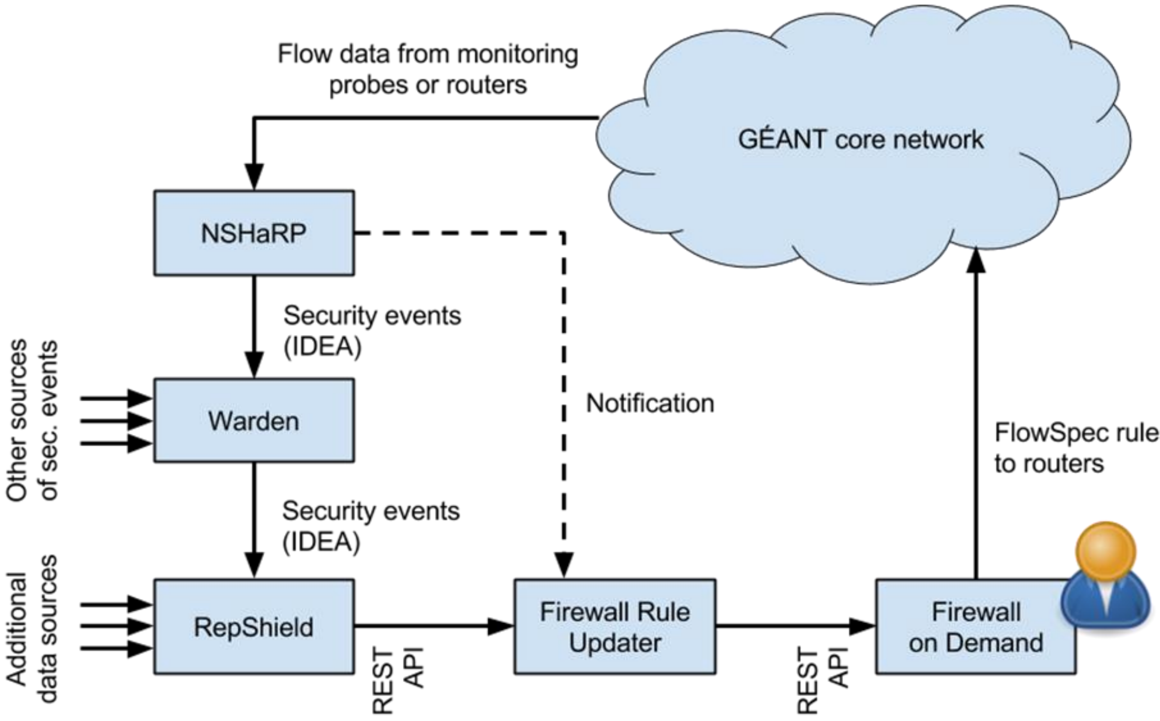

Figure 2.1: Automatic proposals of mitigation rules created by RepShield for FoD v1.6

It consists of the central component – Firewall on Demand (FoD), which is attached to GÉANT core network. The DDoS detection architecture also includes Network Security Handling and Response Process (NSHaRP), a newly developed tool – Reputation Shield (RepShield) and the Firewall Rule Updater (FRU). The four components will be presented in the remainder of this document. In addition to those elements, the architecture also includes Warden, which is an alert-sharing system that feeds the information about the security events into the Reputation Shield.

2.1 Firewall on Demand

Firewall on Demand (FoD) is a solution for a mitigation of large-scale network attacks, such as Distributed Denial of Service attacks. FoD is based on BGP-FlowSpec to allow normally routed GÉANT IP traffic to be filtered, based on the administered BGP FlowSpec rules [rfc5575] [rfc7674], thus preventing the attack to the protected network.

Current provision within the GÉANT core network allows users (NREN NOC administrators) to administer BGP FlowSpec rules via a web-based interface. The web interface is capable of crafting, disseminating or withdrawing FlowSpec rules 'on demand'. FlowSpec is preferred rather than older DDoS detection and mitigation tools such as access control lists (ACLs) and remotely triggered black hole (RTBH) filtering, because it is faster and enables bigger granularity of firewall rules.

More about Firewall on Demand, its features and development roadmap is provided in the Deliverable D8.2 Firewall-on-Demand Progress Report [D8.2].

2.2 Network Security Handling and Response Process - NSHaRP

The NSHaRP provides a mechanism to quickly and effectively inform affected users and to manage the attack mitigation process. It allows Computer Emergency Response Teams (CERTs) to tailor how and for what type of incidents they want their notifications to be triggered. The system serves as an extension to the national research and education networking (NREN) organisations' CERTs, if they do not have either the available human or the technical resources to monitor for security incidents affecting their users. Users can configure – depending on how important a type of attack is for the user – for each corresponding type of NSHaRP event, e.g. SSH dictionary attack or DDoS attack, the severity/criticality, from either critical (single 5-min mail notification), high (aggregated mail every 6 hours), medium (12 hours), low (every day).

NSHaRP extends the NRENs' detection and mitigation capability across GÉANT network and to its borders with other networks, therefore enabling the attack to be mitigated before it transits the GÉANT network. This is a highly innovative and unique security service in that it caters for different requirements from each NREN, by enabling the customisation of their NREN specific alerts in their hands.

2.3 Reputation Shield – RepShield

Rep(utation)Shield [REPSHIELD] is a component designed and currently developed by CESNET within JRA2 T6. RepShield analyses security alerts/events and correlates them with other various information sources to estimate a so-called reputation score for network entities related to the events.

A network entity can be an IP address, Network (IP prefix), Autonomous System (AS) or Domain name. Additional information sources might be various detectors (honeypots, flow analysers, IDS, …), blacklists, lists of open resolvers, sharing systems such as MISP or AlienVault OTX, as well as data acquired by DNS, geolocation, Whois, etc. Information about the network security events are gathered via a Warden alert-sharing system, also developed by CESNET or via NSHaRP.

Reputation score is formally defined as the probability and severity of future attacks originating from the network entity. This is estimated by machine-learning techniques using past behaviour of the entity, similar entities and other information related to the entity as input. The current implementation is based on neural networks. The best models so far are able to correctly predict detection of an IP address in next 24 hours with 74% accuracy for port scanning data. Accuracy for DDoS attacks is currently unknown, since more DDoS related data are needed in order to provide such information. However, the estimation of reputation score is still a topic of ongoing research.

The quality of computed reputation score highly depends on quantity and quality of data from external sources. Inputs checked by RepShield were extended by several new publicly available blacklists sources related to malware infection, SPAM senders, etc. Each blacklist is used to check the reported suspicious entities. To solve the issue of ageing information about entities, an automatic update feature was implemented in the core of RepShield. Regular updates fetch new information about entities and trigger computation of the new value of reputation score for every entity in the database.

Since the automated creation of the FoD rules considers just a subset of event types related to high-severity DDoS attacks, a classification algorithm was designed and implemented in RepShield to assign labels to every entity. This classification is used to recognise major characteristics of malicious activity and the assignment of labels allows for improved filtering of the entities in RepShield. As a result, a list of frequent sources of DDoS attacks with bad reputation scores can be obtained from RepShield.

RepShield has a graphical user interface that can be used to manually check and investigate information about entities stored in the database. Figure 2.2 shows the list of entities observed in reported events. Each entity is queried in several external sources of information and looked up information are presented to the user. The detailed information about a selected entity is shown in Figure 2.3.

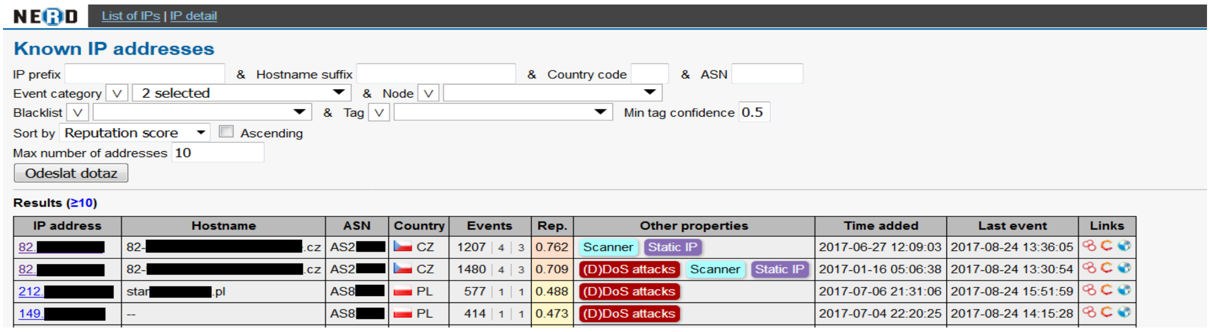

Known IP addresses

IP prefix & Hostname suffix & Country code & ASN
Event category V 2 selected & Node V
Blacklist V & Tag V Min tag confidence 0.5
Sort by Reputation score □ Ascending
Max number of addresses 10

Results (≥10)

IP address	Hostname	ASN	Country	Events	Rep.	Other properties	Time added	Last event	Links
82.[redacted]	82-[redacted].cz	AS2[redacted]	CZ	1207 | 4 | 3	0.762	Scanner Static IP	2017-06-27 12:09:03	2017-08-24 13:36:05	
82.[redacted]	82-[redacted].cz	AS2[redacted]	CZ	1480 | 4 | 3	0.709	(D)DoS attacks Scanner Static IP	2017-01-16 05:06:38	2017-08-24 13:30:54	
212.[redacted]	star[redacted].pl	AS8[redacted]	PL	577 | 1 | 1	0.488	(D)DoS attacks	2017-07-06 21:31:06	2017-08-24 15:51:59	
149.[redacted]	--	AS8[redacted]	PL	414 | 1 | 1	0.473	(D)DoS attacks	2017-07-04 22:20:25	2017-08-24 14:15:28	

Figure 2.2: RepShield records in the database

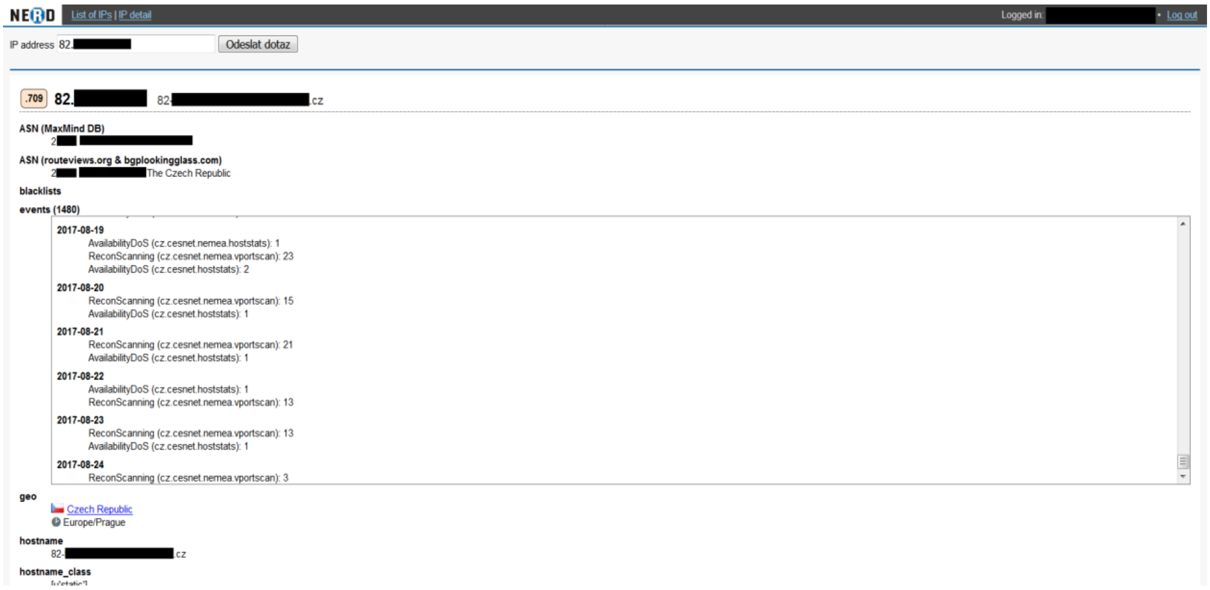

IP address 82.[redacted]

.709 82.[redacted] 82-[redacted].cz

ASN (MaxMind DB)
2[redacted] [redacted]

ASN (routeviews.org & bgplookingglass.com)
2[redacted] [redacted]The Czech Republic

blacklists

events (1480)

```
2017-08-19
        AvailabilityDoS (cz.cesnet.nemea.hoststats): 1
        ReconScanning (cz.cesnet.nemea.vportscan): 23
        AvailabilityDoS (cz.cesnet.hoststats): 2
2017-08-20
        ReconScanning (cz.cesnet.nemea.vportscan): 15
        AvailabilityDoS (cz.cesnet.hoststats): 1
2017-08-21
        ReconScanning (cz.cesnet.nemea.vportscan): 21
        AvailabilityDoS (cz.cesnet.hoststats): 1
2017-08-22
        AvailabilityDoS (cz.cesnet.hoststats): 1
        ReconScanning (cz.cesnet.nemea.vportscan): 13
2017-08-23
        ReconScanning (cz.cesnet.nemea.vportscan): 13
        AvailabilityDoS (cz.cesnet.hoststats): 1
2017-08-24
        ReconScanning (cz.cesnet.nemea.vportscan): 3
```

geo
Czech Republic
Europe/Prague

hostname
82-[redacted].cz

hostname_class
[u'static']

Figure 2.3: RepShield graphical user interface: details about entity

For the Network Security Task (JRA2 T6), the main interest for RepShield is to use the provided Reputation Scores and prepare the data for the Firewall Rule Updater, which will derive automated FoD rules that can be proposed to the users in order to make the system easier to use.

The arrow leading from NSHaRP into Firewall Rule Updater in Figure 2.1 presents the trigger mechanism for the preparation of the list of rules that is then pushed into FoD. Since RepShield is able to get data from different sources not just NSHaRP, it is necessary to consider only events related to GÉANT infrastructure. However, other sources of information can increase precision and reliability of the proposed rules. Theoretically, RepShield can be improved in future in order to get rid of the direct connection from NSHaRP to Firewall Rule Updater. This solution requires some kind of labelling of the events in RepShield to be able to distinguish those that are meant to generate the list of rules.

2.4 Firewall Rule Updater (FRU)

Firewall Rule Updater FRU is an intermediate firewall rule processing component that interconnects several architectural components: NSHaRP, RepShield and FoD.

Based on an NSHaRP event and finished update of RepShield computation that act as a starting point, FRU gathers data about the potential incident and use records from RepShield to prepare a rule proposal for FoD. The rule proposal is sent to a user that then decides whether to apply the rule to its network via a FoD GUI or decline it. FRU uses both FoD and RepShield APIs to fetch the information and to pass the information with a lists of rules and insert them into FoD for approval by a user.

3 DDoS Detection and Mitigation Pilot

The GÉANT network DDoS detection and mitigation pilot will be based on FoD v1.6, with the RepShield automated rule proposal. The pilot will include up to three NRENs and the exact NREN list will be defined before the pilot. The pilot is scheduled for Q1 2018, after which it is expected to transition into production and replace the FoD service in production at that time.

3.1 DDoS Detection and Mitigation Pilot Infrastructure

The infrastructure that will be used for the pilot is presented in Figure 2.1 in the previous section. It resembles the infrastructure prepared for existing FoD user acceptance testing (UAT) and it consists of one FoD VM that communicates with one GÉANT router via NETCONF to synchronise the FlowSpec rules. FoD communicates with the Firewall Rule Updater which is linked with RepShield installed on additional VM. RepShield on the other hand gets information about the security events from Warden and NSHaRP.

Warden is an incident sharing system that can be used to capture NSHaRP events which are converted into Intrusion Detection Extensible Alert (IDEA) format [IDEA]. The pilot will consider subscribing this system to also receive events from the production Warden system deployed in CESNET, as presented in Figure 3.1: .

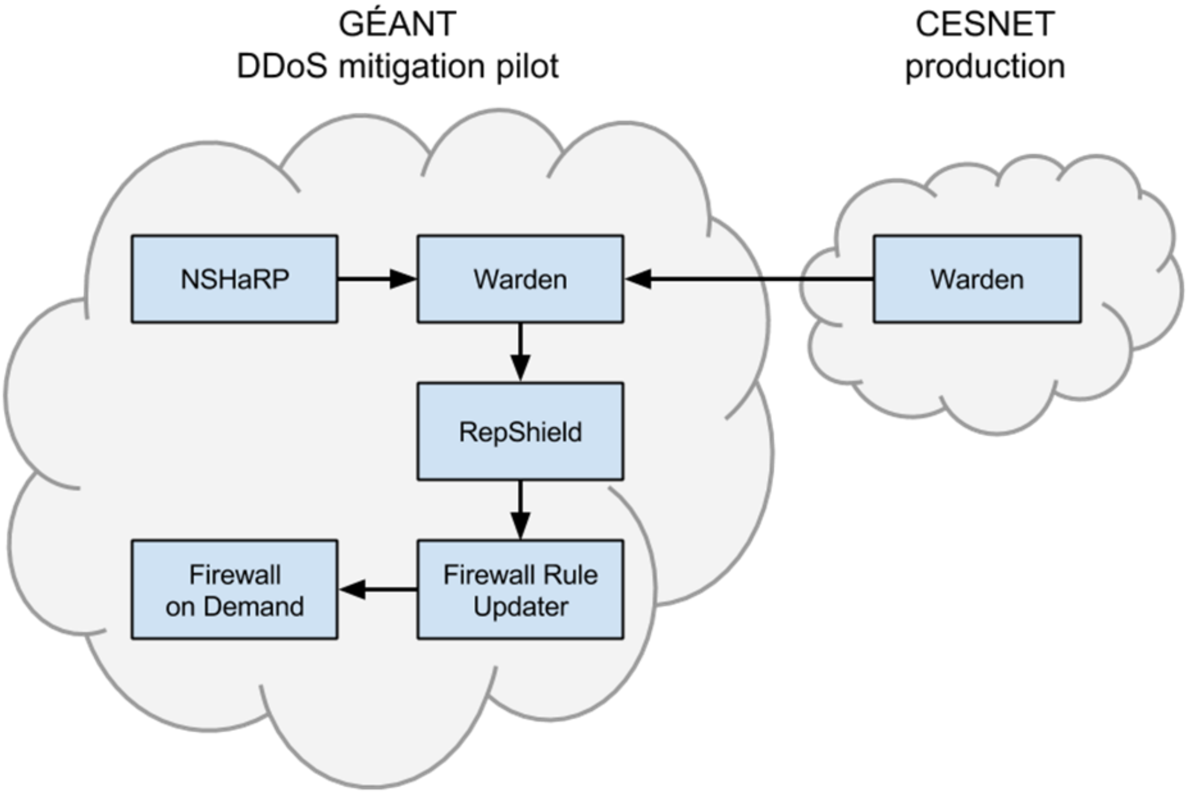

Figure 3.1: GÉANT pilot infrastructure and CESNET alert-sharing production system.

The FRU will serve as an intermediate component between NSHaRP input events exported in IDEA format as an input to Warden, RepShield REST API and FoD rule control API. RepShield REST API is described in deliverable D7.1 of the GN4-1 project.

FoD was designed to fully support BGP FlowSpec mitigation rules for both IPv4 and IPv6 protocols. Currently, the pilot is planned for IPv4 support only. For full IPv6 support, additional work will be required on DDoS detection and mitigation system including: some extension of input field checks in FoD, adaptation of the user registration (currently, automated discovery is in place only for IPv4 user prefixes), enabled and tested IPv6 support in core routers, as well as NShaRP support for IPv6 events. This is planned for future work on the development on DDoS detection and mitigation system.

3.2 DDoS Detection and Mitigation Pilot Description

The pilot will be created for the following use case: a protected network gets under a DDoS attack, and the DDoS detection and mitigation system reacts to this attack with a preparation of a rule proposal for the FoD system. The pilot infrastructure is set up as shown in Figure 2.1, with the correlation of Warden systems as per Figure 3.1: The defence process is presented at Figure 3.2: .

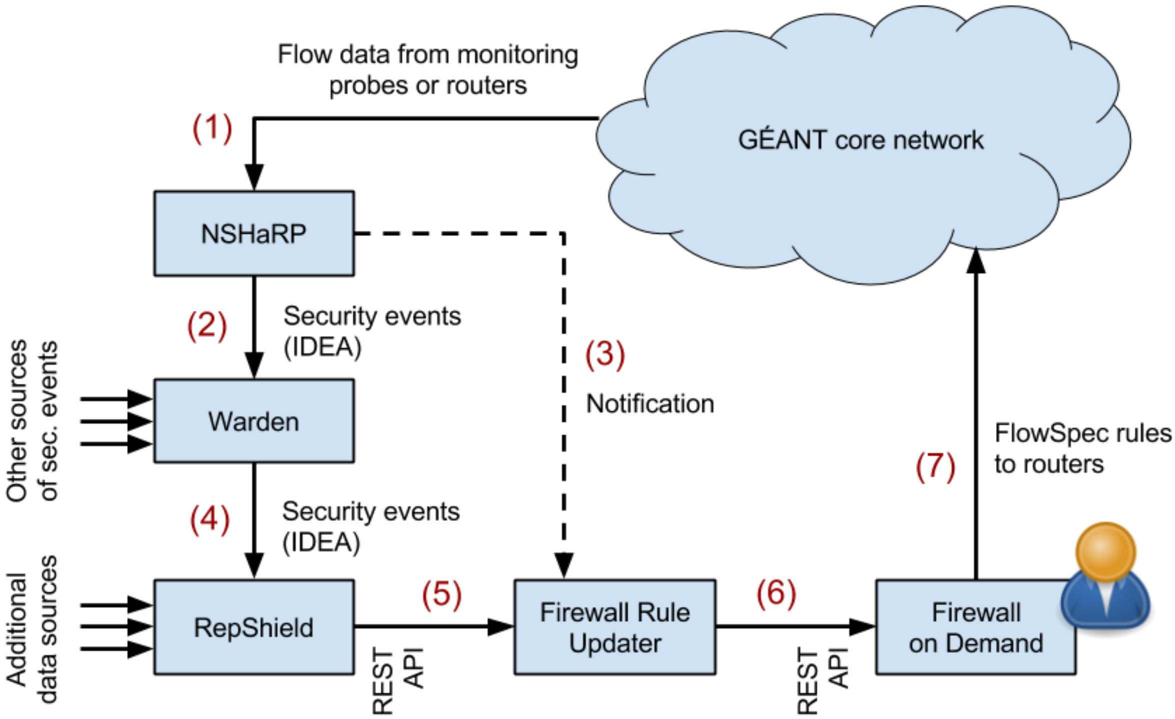

Figure 3.2: DDoS mitigation process

A security incident that is coming into a protected network is recorded in NSHaRP (1). This incident triggers NSHaRP data export about the event (in IDEA format) into Warden (2) which then populates the RepShield database. This action triggers reputation score re-calculation for this network entity in the RepShield. An NSHaRP event will also trigger Firewall Rule Updater (3) which will query for correlated information about all DDoS events from local RepShield. RepShield will send to FRU (5) information about the correct NREN and contact data (which are already included n NSHaRP). FRU will map this to create inactive FoD rules via FoD rule control API (6) for the respective particular FoD user (7). The user should then manually accept or decline the rule.

The description of that rule includes: an explanation that it is an automated rule, a summary about correlated NSHaRP event and a link to local RepShield query for that event. FRU then sends an e-mail with this rule to the contact e-mail address of the corresponding NREN's representative NOC engineer.

This notifies the user about the correlated NShaRP event and includes links to the FoD GUI, which will guide the user directly to rules to decide about whether to activate the proposed rules or not. Moreover, the notification also includes an appropriate link to local RepShield with a query for the correlated NSHaRP event. The user can get a good overview as well as details about the respective detected attack by just looking into RepShield. If the user finally decides that the proposed rules should be activated, s/he can easily do it through a link provided in the notification email for FoD.

It should be noted that while automated mitigation is the main goal, applying such filters – especially the false positive ones – can have a substantial negative impact on an NREN or campus. Therefore, the pilot still requires some review and approval by the “human in the middle”. However, performance results, and results of the planned survey can be used as feedback for improvement of the automatic creation of the proposed rules. That means additional heuristics may be employed to learn human decisions about the proposed rules and automate the rule activation in the future for those routinely reviewed and approved.

Users can control mitigation rules status at any time. This possibility exists already in FoD v1.1,and is kept in any subsequent FoD versions. That means that a user can easily switch status of any rule from active to inactive or vice versa, either for manually entered rules in previous versions or for the rules that are automatically proposed via RepShield in FoD version 1.6. This will automatically switch the rule status to inactive state, for example, based on NSHaRP DDoS stop events, by examining the rule mitigation statistics, or based on a predefined period of time. However, this will not be part of the pilot. In this pilot, the rules will need to be deactivated manually by the NOC/security operator.

3.3 DDoS Detection and Mitigation Pilot Validation

In order to validate the pilot test and the results, a pre-defined list of acceptance criteria will be prepared to include the user view, as well as the operator/administrator/management view about the new rule proposal functionality. The list will evolve from the UATs previously used for FoD v1.1 and FoD v1.5 testing, and will include additional criteria specific to the rule proposal mechanism.

It can include, for example, evaluation whether all occurred NSHaRP DDoS events have been correctly converted into FoD rules (e.g. by comparing proposed FoD rules with all respective e-mails sent by NSHaRP or respective logs), as well as the evaluation for the usefulness of these proposals.

At the end of the pilot, a survey will be completed by the pilot users to assess and validate the proposed solution.

4 Future Work

After the pilot is completed, future work will be based on the pilot results. In the case of (anticipated) positive feedback, steps will be taken in order to transition the pilot into production and offer it to broader range of users. The results of this work will be reported in GN4-2 Quarterly Management Reports, as well as in the Deliverable D1.12 Assessment of GÉANT Service Catalogue.

Further development of the DDoS detection and mitigation infrastructure will be guided by the need for process automatisation in order to help speed up the network and security services operation. One possible aspect for improvement is to enable automatic application and retrieval of pre-prepared rule proposals with a timing option. With this, the system will automatically prepare the rule proposal ready for the deployment in the system, which will delay the automatic deployment so that an operator can decline the rule deployment. If the operator does not react, the rule will be automatically applied, in line with the operator's tacit agreement. The rules could similarly be automatically retrieved, after the threat fades away. With stronger trust in the DDoS mitigation system, such an approach can help the operators to avoid false positives, speed up the process, alleviate the stress of needing a fast reaction without mistake, and result in stronger infrastructure protection.

The second area of future work might be to introduce the scrubbing centre in the current architecture. Figure 4.1: shows the architecture of FoD connected to a scrubbing centre.

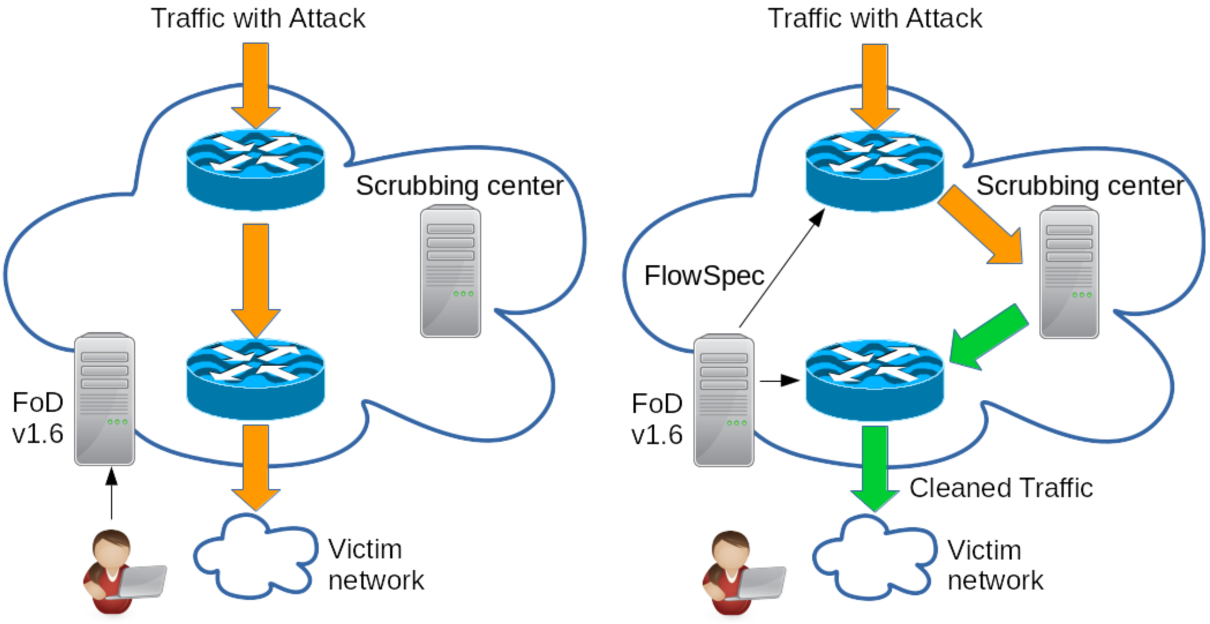

Figure 4.1: Usage of scrubbing centre with FoD v1.6

The scrubbing centre would accept all incoming traffic, examine it from the perspective of potential threats, clean the traffic and forward it to the user's desired destination. The information about the malicious traffic, attackers and threats from the scrubbing centre should be fed into the RepShield for reputation score calculation.

The third area of future work is to consider replacing FRU with direct communication between RepShield and Firewall on Demand through a RepShield notification API. However, this work is not planned for the duration of the pilot.

The fourth are might consider adjustment of the proposed DDoS detection and mitigation system to fully support IPv6. Although it is supported by design, this could be achieved only if all components - from the core of GÉANT network through all components - will enable the needed support.

Finally, the fifth area of future work might include the analysis of the proposed and performed actions. Such an analysis might help to detect potential false alarms or aid the process for the continual improvement of the DDoS detection and mitigation.

5 Conclusion

This document describes the architecture of the DDoS detection and mitigation architecture, as created in JRA2, Task 6. The architecture includes the Firewall on Demand tool for DDoS mitigation, enhanced with NSHaRP tool for attack detection, and the Firewall Rule Updater and RepShield tool for automatic rule proposal.

FoD has undergone years of improvement and development within the GÉANT project. With the information about the reputation scores for each recognised network entity, RepShield can help FoD create better mitigation rule proposals. The rule proposals can be based on multiple sources, including various detectors (honeypots, flow analysers, IDS), blacklists and many open sources that can act in addition to the data sources recognised by NSHaRP and Warden. The integration of FoD, NSHaRP and RepShield is planned in a DDoS detection and mitigation pilot, as described in this document.

Future work to be considered includes: the development of a DDoS detection and mitigation system that further automates rule proposal, application and retrieval, addition of a scrubbing centre for more efficient cleaning of the NREN's incoming traffic, FRU replacement with the direct integration of RepShield and Firewall on Demand, full IPv6 support, as well as the system learning analytics for service optimisation and continuous improvement.

References

[D8.2] https://www.geant.org/Projects/GEANT_Project_GN4/deliverables/D8.2_Firewall on-Demand-Progress-Report.pdf

[FLOWSPY] GitHub grnet/flowspy: GRNET Firewall on Demand platform. Powers, Aug 2016, https://github.com/grnet/flowspy

[GN4-2] (GÉANT Network 4, Phase 2) project part-funded from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No.731122

[GNNSHARP] GEANT Limited, NSHaRP, Aug 2016, http://geant3.archive.geant.net/Network/NetworkOperations/Pages/NSHaRP-NetworkSecurity.aspx

[IDEA] Intrusion Detection Extensible Alert https://idea.cesnet.cz/en/index

[REPSHIELD] CESNET z.s.p.o., , Aug 2016, https://www.cesnet.cz/wp-content/uploads/2015/12/Reputation-Shield-BARTOS.pdf

[rfc7674] Clarification of the Flowspec Redirect (updates rfc5575) https://tools.ietf.org/html/rfc7674

[rfc5575] GP-FlowSpec, https://tools.ietf.org/html/rfc5575

Glossary

ACL	Access Control Lists
AS	Autonomous System
BGP	Border Gateway Protocol
CERT	Computer Emergency Response Team
CLI	Command Line user Interface
DDoS	Distributed Denial of Service
DNS	Domain Name System
FoD	Firewall-On-Demand
FRU	Firewall Rule Updater
GN4-2	GÉANT Network 4, Phase 2 project
GUI	Graphical User Interface
IDEA	Intrusion Detection Extensible Alert
IDS	Intrusion Detection System
JRA2	Joint Research Activity 2
NOC	Network Operations Centre
NREN	National Research and Education Networking
NSHaRP	Network Security Handling and Response Process
RepShield	Reputation Shield
REST API	Representational state transfer Application Programming Interface
RTBH	Remotely Triggered Black Hole
UAT	User Acceptance Testing
UI	User Interface
VM	Virtual Machine