30-11-2018

Deliverable D8.12
Distributed Denial of Service Mitigation v1.0 Pilot Follow-up

Deliverable D8.12

Contractual Date:	30-11-2018
Actual Date:	30-11-2018
Grant Agreement No.:	731122
Work Package/Activity:	JRA2
Task Item:	Task 6
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	DFN (LRZ)>
Document ID:	GN4-2-18-5786D6
Authors:	David Schmitz (DFN (LRZ)), Ivana Golub (PSNC), Václav Bartoš (CESNET), Evangelos Spatharas (GÉANT Association), Tomáš Čejka (CESNET)

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 results of the DDoS detection and mitigation tool that is developed within Network Services Development Joint Research Activity (JRA2) Network Security Task (Task 6). The previous Deliverable D8.3 presented the DDoS detection and mitigation architecture that includes Firewall on Demand (FoD), Network Security Handling and Response Process (NSHaRP) and Reputation Shield (RepShield), and explained the DDoS detection and mitigation pilot. This document presents the pilot results and the next steps.

Table of Contents

Table of Figures

Executive Summary

The Distributed Denial of Service (DDoS) detection and mitigation solution is developed within the GN4-2 project's Network Services Development Joint Research Activity (JRA2) by the Network Security Task (Task 6). DDoS detection and mitigation solution users are GÉANT community organisations that have their own Autonomous System (AS) number, and NOC and CERT operators.

The DDoS Detection and Mitigation pilot is based on the GÉANT DDoS detection and mitigation architecture, which includes:

- Firewall on Demand (FoD) [FLOWSPY]
- Network Security Handling and Response Process (NSHaRP) [GNNSHARP]
- Firewall Rule Updater (FRU)
- Reputation Shield (RepShield) [REPSHIELD].

The architecture and components of the GÉANT DDoS detection and mitigation solution are detailed further in *Deliverable D8.3 Distributed Denial of Service Mitigation v1.0 Pilot* [D8.3].

The main component of the detection and mitigation architecture is FoD version 1.6. More details of the FoD architecture and features are presented in *Deliverable D8.2 Firewall on Demand Progress Report* [D8.2].

This document presents the results of the DDoS detection and mitigation solution pilot that was presented in *Deliverable D8.3 Distributed Denial of Service Mitigation v1.0 Pilot* [D8.3].

The pilot was successful in confirming that it is possible to automatically propose BGP FlowSpec rules for DDoS mitigation. It demonstrated that an NREN Network Operations Centre FoD user only needed to check and apply the proposed mitigation rules, rather than having to enter them manually.

Rules are proposed according to detected security events (via NSHaRP) and information on network entities retrieved from several other publicly available sources, such as Whois and geolocation databases, blacklists related to SPAM senders, and malware infection and botnet C&C servers.

Detected DDoS attacks, which are the main concern of the FoD, can be correlated and quality checked by the Reputation Shield (RepShield) tool. Based on the input from NSHaRP and RepShield, the Firewall Rule Updater prepares rule proposals for the FoD, and users are prompted to apply the rules, accelerating the overall DDoS detection and mitigation process.

1 Introduction

A Distributed Denial of Service (DDoS) detection and mitigation system is being developed by the Network Security Task (Task 6) as a part of the Network Services Development Joint Research Activity (JRA2) of the GN4-2 project [GN4-2].

DDoS detection is performed by the Network Security Handling and Response Process (NSHaRP) [GNNSHARP] and RepShield [REPSHIELD], which provide information about potential security threats, and suggest mitigation rules that are then implemented by the DDoS mitigation tool via the Firewall Rule Updater (FRU).

The mitigation of detected DDoS attacks is carried out by the Firewall on Demand (FoD) [FLOWSPY] tool. FoD development began in GN3plus, and continues in the GN4-2 project. The DDoS detection and mitigation system hosts FoD version 1.6., which enables semi-automated rule proposal. FoD users are NOC and SOC administrators from organisations that have their own autonomous system number (ASN).

The system is designed as multi-domain to better suit the GÉANT community. Rules entered in the FoD user interface (UI) are automatically propagated to all routers in the GÉANT core network, to avoid inconsistencies in the routers' configurations. Users can view their active rules in order to find, review and remove any remaining unchecked rules.

The DDoS detection and mitigation system presented in this document has a number of benefits. It is based on open source software, so its operation does not entail licensing costs. It is multi-tenant and can be easily extended to facilitate automation and integration with existing tools and processes. In addition to running in the GÉANT core, it can also run in NRENs or even connected institutions in a multi-domain manner. Any of these instances can exchange and share BGP FlowSpec rules via eBGP [rfc5575] [rfc7674].

The architecture of the GÉANT DDoS detection and mitigation system is presented in *Deliverable D8.2 Firewall on Demand Progress Report* [D8.2]. The pilot in which the DDoS detection and mitigation system was tested is described in *Deliverable D8.3 Distributed Denial of Service Mitigation v1.0 Pilot* [D8.3].

This document presents the pilot results in the following sections. Section 2 describes the DDoS detection and mitigation pilot, Section 3 describes the DDoS detection and mitigation pilot results and Section 4 presents conclusions on the pilot results.

2 DDoS Detection and Mitigation Pilot

The architecture of the DDoS detection and mitigation pilot described in *Deliverable D8.3 Distributed Denial of Service Mitigation v1.0 Pilot* [D8.3] is shown in Figure 2.1.

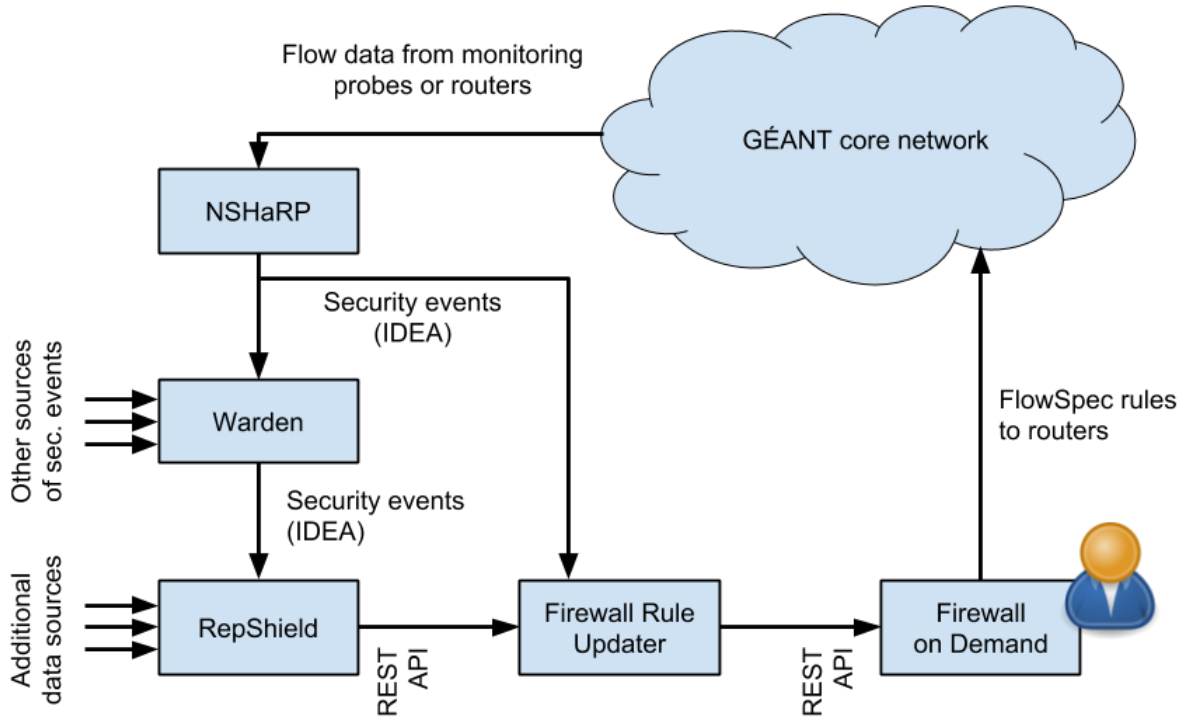

Figure 2.1: Automatic proposals of mitigation rules created by RepShield for FoD v1.6

The Firewall on Demand (FoD) is a central component attached to the GÉANT core network. It receives firewall rules proposed by the Firewall Rule Updater (FRU) which in turn, receives information about potential security threats from the Network Security Handling and Response Process (NSHaRP), the Reputation Shield (RepShield) and the alert-sharing system Warden [REPSHIELD]. The characteristics and functionalities of these components are presented in more detail in Deliverable D8.3 *Distributed Denial of Service Mitigation v1.0 Pilot* [D8.3].

The pilot tested a use case where a protected network is under a DDoS attack, and the DDoS detection and mitigation system reacts to this attack by proposing a rule to the FoD system. The defence process is presented in Figure 2.2.

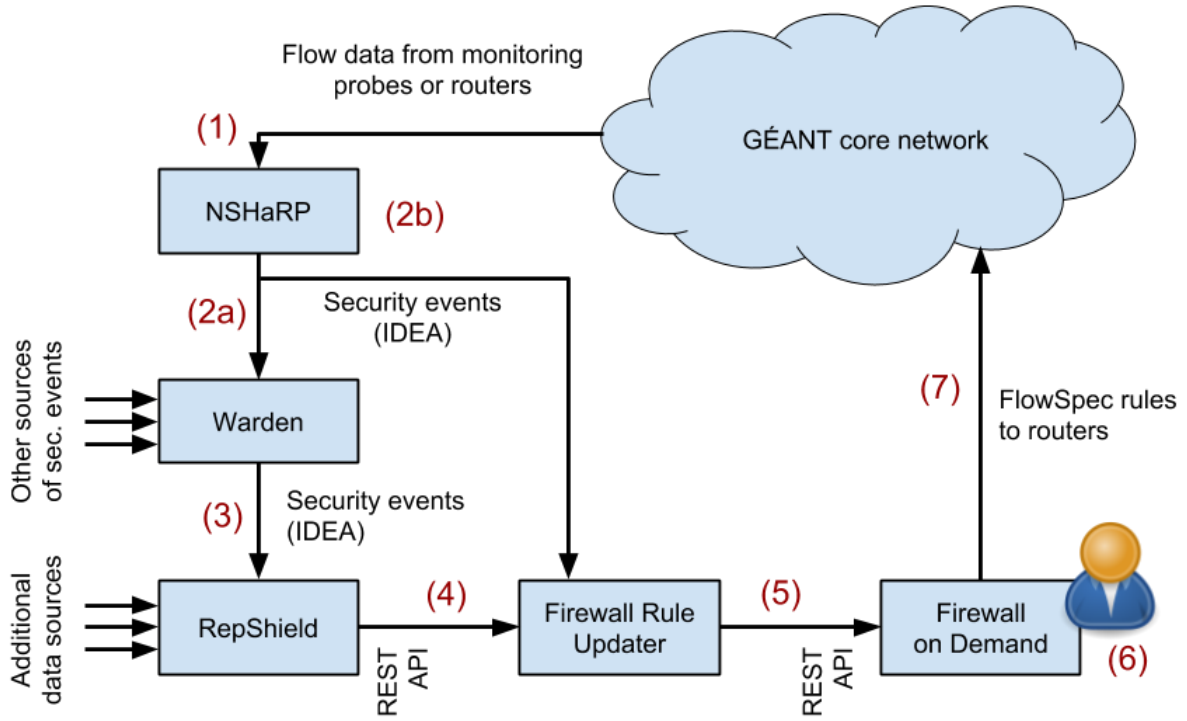

Figure 2.2: DDoS mitigation process

Figure 2.2 illustrates the sequence of actions in the DDoS mitigation process:

A security incident entering a protected network is recorded in NSHaRP (1). The incident triggers NSHaRP to export event data (in IDEA format) to Warden (2a), which then updates the RepShield database (3). This action triggers a re-calculation of the reputation score for this network entity in the RepShield. The NSHaRP event also triggers the Firewall Rule Updater (2b), which queries the local RepShield for reputational information on all IP addresses in this DDoS event (4). The FRU creates the proposed mitigation rules in an inactive state via the FoD's rule control REST API with admin rights (5). The user can then accept, modify or decline any of the proposed rules (6). The FoD installs the accepted rules via NETCONF on the routers which synchronise these rules among them via IBGP (7).

The description of a proposed rule includes:

- The information that the rule was automatically proposed by the FRU.
- A summary of the triggering NSHaRP event.
- A summary of the reputational information RepShield provided regarding the attacker IP addresses and prefixes of the event.
- Links to more detailed local RepShield queries for each attacker IP prefix and address in this event.

The FoD sends the information to all respective FoD users via email. The FoD users are determined based on the destination IP prefixes in the rules proposed by the FRU and notified via the configured notification email addresses.

The email notifies FoD users of the created rule and includes the information that the rule was automatically proposed by the FRU. A link to the FoD GUI rules is included to allow the users to decide whether the proposed rules should be activated. Also included is a summary of the triggering NSHaRP event, a summary of the reputational information from RepShield, and links to RepShield for executing more detailed queries about the reputational information regarding attacker IP addresses and prefixes in the NSHaRP event.

The user can get an overview as well as details about a detected attack by simply looking into RepShield. If the user decides that a proposed rule should be activated, they can easily do this using the link to the FoD, which is provided in the notification email.

Currently, the FRU proposes four different mitigation rules for each received NSHaRP DDoS event, which the user can choose from. Two of them do not take information from RepShield into consideration:

- Type 1 – a rule that prevents the attack from the IP addresses reported in the event.
- Type 2 – a rule which mitigates any traffic towards the attacked Layer 4 ports regardless of the source address.

Two other rules are created based on RepShield information:

- Type 3 – a rule which calculates a list of IP prefixes comprising the reported attacker IP addresses based on the reputational information on IP prefixes.
- Type 4 – a rule which uses the list of attacker IP addresses with the worst reputation score currently known to RepShield reported from the original DDoS event.

The user can choose any of these four proposed rules, edit the chosen rule if required, and then activate or just decline it.

It should be noted that while automated mitigation is the main goal, applying such filters, especially false positive ones, without review can have a significant, negative impact on an NREN or campus. Therefore, the solution still requires some human review and approval. However, performance and usability recommendation results provide feedback on the improvement of the automatic creation of proposed rules. That means additional heuristics may be employed to learn human decisions about proposed rules, and automate the activation of rules that are routinely reviewed and approved in the future.

Users can control the status of mitigation rules at any time. A user can easily switch the status of a rule from active to inactive or vice versa, both for rules that were entered manually in previous versions and for rules that are automatically proposed via RepShield in FoD version 1.6.

In the future, it could be useful to have the possibility to deactivate the rules in FoD automatically, for example, based on NSHaRP DDoS stop events, rule mitigation statistics, or a predefined period of time. However, this might only be part of some future release. In this FoD version and in this pilot, the rules need to be deactivated manually by the NOC/security operator.

3 DDoS Detection and Mitigation Pilot Results

The DDoS detection and mitigation solution presented in the previous chapters was tested to validate whether it is fit for use and fit for purpose, and whether it provides acceptable performance and accuracy. Three organisations participated in the testing and evaluation of the solution – GÉANT (as the organisation that will be running the solution in production), CESNET and DFN/LRZ. They concluded that the added new features work as designed and improve the mitigation and defence of their networks by making the workflow faster and more efficient.

The first test evaluated the detection and mitigation process. All components performed as expected:

- A security incident was recorded in NSHaRP which triggered an NSHaRP data export to Warden.
- Warden then updated the RepShield database which triggered a re-calculation of the reputational score for this network entity in the RepShield.
- An NSHaRP event also triggered the FRU which then queried the local RepShield for correlated reputational information about attacker IP addresses in this DDoS event.
- The FRU created mitigation rules in inactive state via the FoD's rule control REST API with admin rights.
- The FoD sent an email to the user, prompting them to accept, modify or decline the proposed rules.
- The FoD installed the rules accepted by the user in the routers via NETCONF, and the routers where then synchronised via IBGP.

During the tested period, only minor incidents were reported, and unfortunately for the pilot, but fortunately for the network, none were a huge DDoS attack. The attacks that were tested were those reported by either NSHaRP or Warden, and these were used to assess the DDoS detection and mitigation solution workflow, the functionalities of each of the components, and the system as a whole.

The accuracy of the proposed rules depends on the accuracy of NSHaRP and Warden. During the pilot, no false positives were recorded.

The second test verified the content of the email that is sent to the FoD user. It included all the expected information:

- Each proposed rule included all the information.
- Every user was able to get additional information about an event via the appropriate link to RepShield.

- Users were also able to activate or deactivate any proposed rule. If a rule was accepted, it was installed from the FoD to the routing configuration as an access control list. If a rule was declined, the router's configuration was not modified by the rule.
- Users were also able to easily switch the status of any rule from active to inactive or vice versa, both for rules that were manually entered in previous versions and for rules proposed by the FRU.

Per design, FoD is connected to one router directly, so all rules that are proposed for FoD by FRU (based on the information received from NSHaRP or Warden) directly impact the first router. The information is then propagated to other routes via BGP and in this pilot it included three more routers.

The third test validated that the FRU can create four types of rules, and that the user is able to choose, modify, approve or decline any rule. All four rules worked as expected, and testing proved that all four rules are useful:

- Type 1 rules enable fast reaction to prevent attacks, when an event's exact IP address and port data are used.
- Type 2 rules proved useful when multiple IP addresses attack a specified port, by disabling all access to this Layer 4 port.
- Rules Type 3 and Type 4 were using the information from RepShield. The type 3 rule contained a list of IP prefixes that RepShield recognised as attacker IP addresses and port numbers.
- Type 4 rule was seen as useful also, as it provided enhanced network protection through the use of additional intelligence.

During the testing, NSHaRP generated an event which states that attackers 10.97.92.182 and 10.97.92.183 attacked 192.168.111.139 on TCP ports 38629, 54899, 23268, 10792 and 50730. The individual rule types that were proposed, produced the following effect:

Type 1 (with two elements, one for each source address):

```
deny 10.97.92.182/32 -> 192.168.111.139/32 (tcp) 38629,54899,23268,10792,50730

deny 10.97.92.183/32 -> 192.168.111.139/32 (tcp) 38629,54899,23268,10792,50730
```

Type 2:

```
deny 0.0.0.0/0 -> 192.168.111.139/32 (tcp) 38629,54899,23268,10792,50730
```

Type 3:

```
deny 10.97.0.0/16 -> 192.168.111.139/32 (tcp) 38629,54899,23268,10792,5073
```

Type 4:

```
deny 172.16.20.0/24, 172.18.89.203/32,... , 172.17.29.0/24 -> 192.168.111.139/32
(tcp) 38629,54899,23268,10792,50730
```

The Type 1 rule replied only to the specific NSHaRP event by using the event's socket parameters. All other types discarded much more traffic, denying the service to the user that initiated the potentially malicious traffic, but also preventing other sources from connecting to the network. Types 3 and 4 additionally checked the reputation of selected source IP addresses which should minimise the risk that the blocked addresses might still represent valid traffic. This proves that using RepShield in the architecture provides additional benefits to network administrators, the users of the DDoS detection and the mitigation system.

However, it was also noted that all rules should be used with caution, in order to avoid unnecessary denial of network services. Multiple reactions to an event that NSHaRP recognises as a DDoS attack are possible:

- Only one IP address and/or selected Layer 4 port number can be denied.
- A range of IP addresses / port numbers can be forbidden.
- All traffic that is directed at a specific destination IP address, prefix or port number can be discarded.

Therefore, administrators need to be aware of the traffic, the risks and the consequences to avoid limiting genuine traffic. Every administrator needs to make decisions based on the requirements of their network, organisation and risk assessment, as discarding traffic that is not malicious presents a denial of service in itself. The DDoS detection and mitigation system's feature that enables administrators to review, manually edit, and confirm or decline a rule serves exactly this purpose.

Based on the successful pilot testing, the development team will intensify the collaboration with the operations team to prepare the DDoS detection and mitigation solution for the transition to production.

Future work on the development of the DDoS detection and mitigation system may include employing additional heuristics to learn human decisions about rule application, in order to automate the activation of any further rules.

4 Conclusions

This document described the results of the DDoS detection and mitigation system pilot, as performed in JRA2, Task 6. The architecture includes the Firewall on Demand tool for DDoS mitigation, enhanced by the NSHaRP tool for attack detection, and the Firewall Rule Updater (FRU) and RepShield tool for automatic rule proposal.

The pilot tested the detection and mitigation process, performance, accuracy, utility and warranty of the solution. Three institution representatives who tested the solution concluded that the system works as designed and planned, and that, based on the events detected by NSHaRP, the FRU generates four types of rules that network administrators can select, modify, apply or reject. That means, that via the FRU, malicious IP addresses and/or network prefixes recognised by either NSHaRP or the RepShield tool can be semi-automatically blocked in the FoD.

RepShield allows for future rule proposals to be based on multiple sources, including various detectors (such as honeypots, flow analysers and IDS), blacklists and many open sources that can be used in addition to the data sources recognised by NSHaRP and Warden. This functionality, as well as the information and statistics that can be provided by RepShield and the FoD were tested as a part of the DDoS detection and mitigation system.

The ability of a network administrator to manually control proposed rules was seen as a benefit of the solution, as the automatic application of filters, especially false positives, can have a negative impact on an NREN or campus. During the testing period, none of the incidents reported by NSHaRP or Warden were false positives, but as this can happen as part of a learning curve, the operators have to be aware and able to address such situations.

Future work will include closer collaboration with operations and PLM teams to prepare the DDoS detection and mitigation solution for transition to production. A future roadmap for the DDoS detection and mitigation system may include the implementation of heuristics for additional automation, optimisation of the generation and application of the rules within the system, and the definition of further use-cases that the system might address.

References

[D8.2]	https://www.geant.org/Projects/GEANT_Project_GN4/deliverables/D8.2_Firewall-on-Demand-Progress-Report.pdf
[D8.3]	https://www.geant.org/Projects/GEANT_Project_GN4/deliverables/D8.3_Distributed-Denial-of-Service-Mitigation-v1.0-Pilot.pdf
[FLOWSPY]	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]	http://geant3.archive.geant.net/Network/NetworkOperations/Pages/NSHaRP-NetworkSecurity.aspx
[IDEA]	https://idea.cesnet.cz/en/index
[REPSHIELD]	https://www.cesnet.cz/wp-content/uploads/2015/12/Reputation-Shield-BARTOS.pdf
[rfc5575]	https://tools.ietf.org/html/rfc5575
[rfc7674]	https://tools.ietf.org/html/rfc7674

Glossary

ACL	Access Control Lists
AS	Autonomous System
ASN	Autonomous System Number
BGP	Border Gateway Protocol
C&C	Command and Control
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
IBGP	Internal Border Gateway Protocol
IDEA	Intrusion Detection Extensible Alert
IDS	Intrusion Detection System
JRA2	Network Services Development Joint Research Activity
NOC	Network Operations Centre
NREN	National Research and Education Networking
NSHaRP	Network Security Handling and Response Process
PLM	Product Lifecycle Management
SOC	Security Operations Centre
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