Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2017

S. Hyun
J. Jeong
S. Woo
Y. Yeo
Sungkyunkwan University
J. Park
ETRI

March 13, 2017

NSF-triggered Traffic Steering Framework
draft-hyun-i2nsf-nsf-triggered-steering-02

Abstract

This document describes an architecture of the Interface to Network Security Functions (I2NSF) framework which enables traffic steering between Network Security Functions (NSFs) for security policy enforcement. Such traffic steering enables the composite inspection of network traffic by steering the traffic through multiple types of security functions according to the information model for the NSF-facing interface in the I2NSF framework. This document explains the additional components integrated into the I2NSF framework and their functionalities to achieve NSF-triggered traffic steering. It also describes representative use cases to address major benefits from the proposed architecture.

Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on September 14, 2017.

Copyright Notice

Table of Contents

1. Introduction 3
2. Objective 3
3. Terminology 4
4. Architecture 5
4.1. NSF Operation Manager 7
4.2. Developer's Management System 7
4.3. Network Security Function Forwarder (NSFF) 8
5. Information for Traffic Steering 8
5.1. Packet Forwarding Header 9
5.2. NSF Forwarding Information 9
5.2.1. Query of NSF forwarding information 10
5.2.2. Response of NSF forwarding information 10
6. Use Cases 11
6.1. Enforcing Different NSFs Depending on a Packet Source's Trust Level 11
6.2. Effective Load Balancing with Dynamic NSF Instantiation . 12
7. Security Considerations 13
8. Acknowledgements 13
9. References 13
9.1. Normative References 13
9.2. Informative References 14
Appendix A. Changes from draft-hyun-i2nsf-nsf-triggered-steering-01 15

1. Introduction

To effectively cope with emerging sophisticated network attacks, it is necessary that various security functions cooperatively analyze network traffic [sfc-ns-use-cases] [RFC7498] [i2nsf-problem] [capability-im]. In addition, depending on the characteristics of network traffic and their suspiciousness level, the different types of network traffic need to be analyzed through the different sets of security functions. [capability-im] proposes an information model for the interface between a Security Controller and Network Security Functions (NSFs) (called NSF-facing interface) in the Interface to Network Security Functions (I2NSF) framework [i2nsf-framework]. This NSF-facing interface enables an NSF to trigger further inspection by calling another NSF based on its own analysis results. However, the current design of the I2NSF framework does not consider network traffic steering fully in order to enable such consecutive inspections through multiple security functions.

In this document, we propose an architecture that integrates additional components for traffic steering over NSFs into the I2NSF framework. We extend the security controller's functionalities such that it can interpret a high-level policy of NSF-triggered traffic steering into a low-level policy and manage them. It also keeps track of the available network security function instances and their information (e.g., network information and workload), and makes a decision on which NSF instances to use for a given network security function. Based on the forwarding information provided by the security controller, the security function forwarder performs network traffic steering through required security functions. The security function forwarder is also responsible for interpreting inspection result from a network security function to enforce more advanced inspection. We define an additional packet header format to specify security inspection results and advanced inspection requests.

2. Objective

- Policy configuration for consecutive inspections: NSF-triggered traffic steering architecture allows policy configuration and management of network security function triggering. Based on the triggering policy, relevant network traffic can be analyzed through various security functions in a composite, cooperative manner.

- Network traffic steering for consecutive inspection: NSF-triggered traffic steering architecture allows network traffic to be steered through multiple required network security functions based on the triggering policy. Moreover, the I2NSF information model for NSF-facing interface [capability-im] requires a security function to

call another security function for further inspection based on its own inspection result. To meet this requirement, NSF-triggered traffic steering architecture also enables traffic forwarding from one security function to another security function.

- Load balancing over network security function instances: NSF-triggered traffic steering architecture provides load balancing of incoming traffic over available network security function instances by leveraging the flexible traffic steering mechanism. For this objective, it also performs dynamic instantiation of a security function when there are an excessive amount of requests for that network security function.

3. Terminology

This document uses the terminology described in [RFC7665], [RFC7665] [sfc-ns-use-cases] [i2nsf-terminology][ONF-SFC-Architecture].

- Network Security Function (NSF): A function that is responsible for specific treatment of received packets. A Network Security Function can act at various layers of a protocol stack (e.g., at the network layer or other OSI layers) [RFC7665]. Sample Network Security Service Functions are as follows: Firewall, Intrusion Prevention/Detection System (IPS/IDS), Deep Packet Inspection (DPI), Application Visibility and Control (AVC), network virus and malware scanning, sandbox, Data Loss Prevention (DLP), Distributed Denial of Service (DDoS) mitigation and TLS proxy.

- Advanced Inspection/Action: As like the I2NSF information model

for NSF-facing interface [capability-im], Advanced Inspection/ Action means that a security function calls another security function for further inspection based on its own inspection result.

- Network Security Function Profile (NSF Profile): NSF Profile represents NSF's inspection capabilities. Each NSF has its own NSF Profile to specify the type of security service it provides and its resource capacity etc.

- Network Security Function Operation Manager (NSF Operation Manager): NSF Operation Manager consistently manages information and state of NSF instances and provides NSF network access information to support advanced inspection request. For example, the information includes the supported transport protocols, IP addresses, and locations for the NSF instances. Also, NSF Operation Manager takes charge of dynamic management of a pool of NSF instances by consulting with Developer's Management System and load balancing over NSF instances.

- Packet Forwarding Header/Encapsulation: Packet Forwarding Header is used to forward a packet from one NSF to another for further inspection. The former NSF constructs a Packet Forwarding Header with the NSF profile of the latter NSF and transmits it to a NSFF. The required fields are the action code, the number of the metadata, and the metadata. In this context, the metadata is a part of NSF profile.

- Network Security Function Forwarder (NSFF): A security function forwarder is responsible for forwarding traffic to one or more connected network security functions according to the information carried in the packet forwarding encapsulation when the traffic comes back from an NSF. Additionally, an NSFF is responsible for transporting traffic to another NSFF (in the same or the different type of overlay), and terminating overlay inspection [RFC7665].

4. Architecture

This section describes an NSF-triggered traffic steering architecture and the basic operations of traffic steering. It also includes details about each component of the architecture.

Figure 1 describes the components of NSF-triggered traffic steering architecture. Our architecture enables support a composite inspection of packets in transit. According to the inspection result of each NSF, which is stored in the Packet Forwarding Header, the traffic packets could be steered to another NSF for futher detailed analysis. It is also possible to reflect a high-level advanced inspection policy and a configuration from I2NSF User which is a component of the original I2NSF framwork. Moreover, the proposed architecture provides load balancing, auto supplementary NSF instance generation, and the elimination of unused NSF instances. In order to achieve these design purposes, we integrate several components to the original I2NSF framwork. In the following sections, we explain the details of each component.

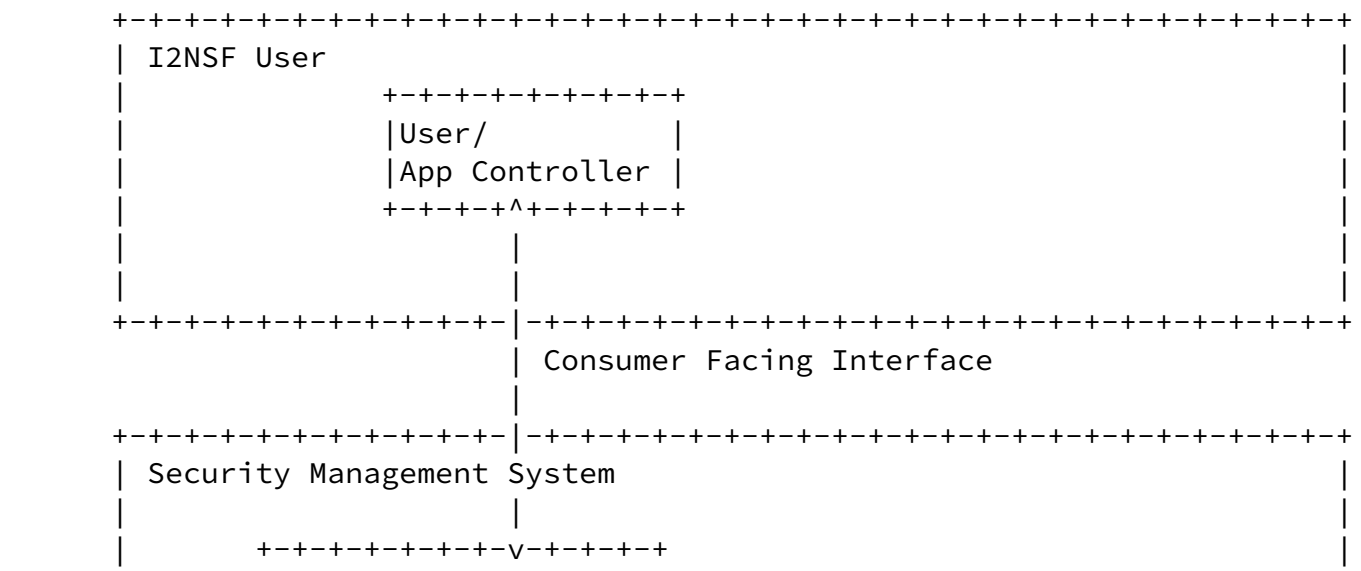

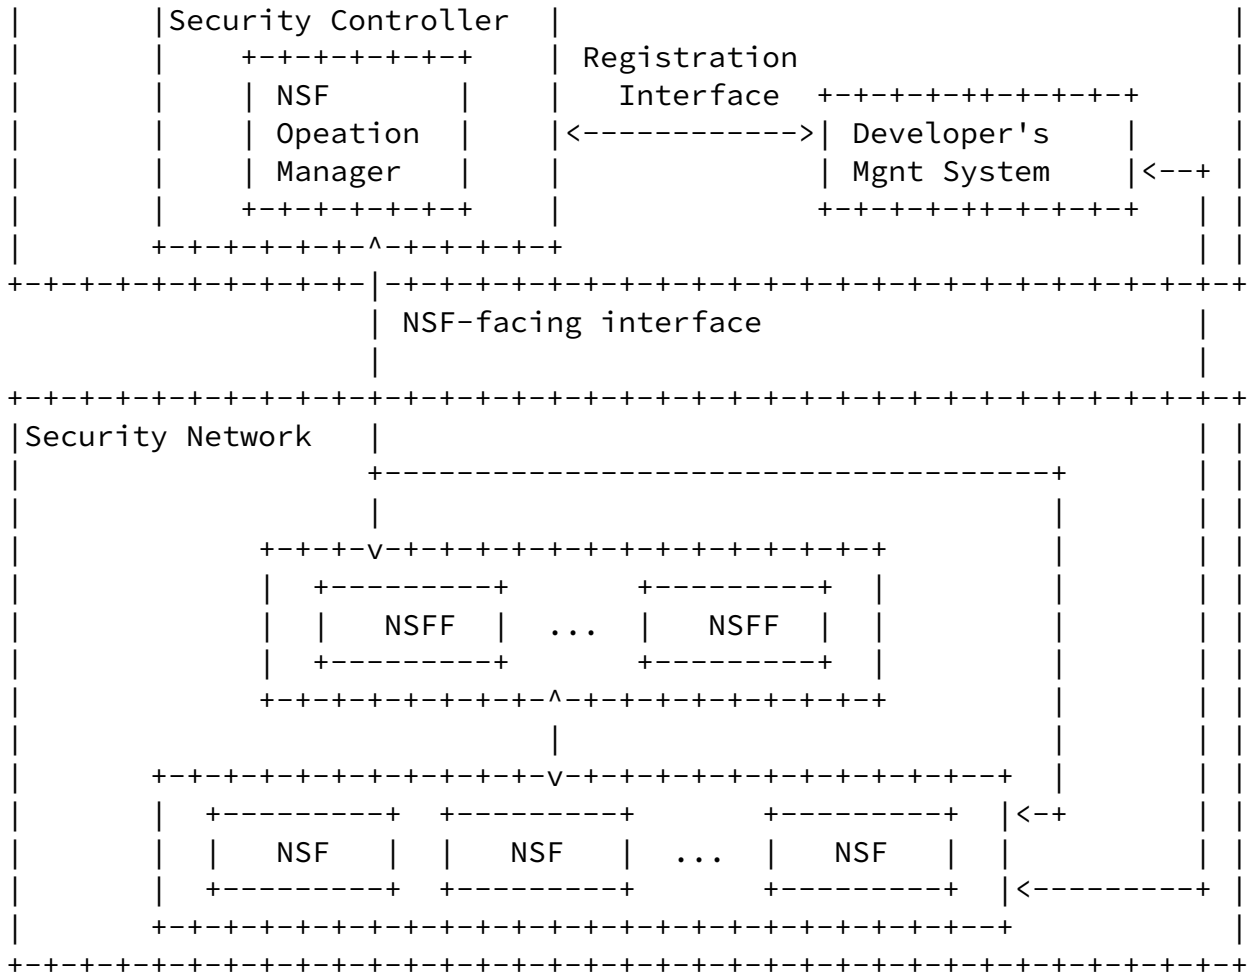

Figure 1: NSF-triggered Traffic Steering Architecture

4.1. NSF Operation Manager

NSF Operation Manager is a core component in our system. It is responsible for the following three things: (1) Maintaining the information of every available NSF instance such as IP address, supported transport protocol, NSF profile, and load status. (2) Responding the queries of available NSF instances from NSFF so as to help to conduct advanced inspection relevant to a given NSF profile.

(3) Requesting Developer's Management System for the dynamic instantiation of supplementary NSF instances to avoid service congestion or the elimination of an existing NSF instance to avoid resource waste. As Figure 1 describes, NSF Operation Manager is a sub-module of Security Controller.

Whenever a new NSF instance is registered, Developer's Management System passes the information of the registered NSF instance to NSF Operation Manager, so NSF Operation Manager maintains a list of the information of every available NSF instance. NSF Operation Manger will receive the request packet containing NSF profile for advanced inspection from NSFF. Once receiving a query of a certain NSF profile from NSFF, NSF Operation Manager searches for all the available NSF instances applicable for that NSF profile and then finds the best instance with selection criteria like location and load status. After finding the best instance, it returns the search result to NSFF.

In our system, each NSF instance periodically reports its load status to NSF Operation Manager. Based on such reports, NSF Operation Manager updates the information of the NSF instances and manages the pool of NSF instances by requesting Developer's Management System for the additional instantiation or elimination of the NSF instances. Consequently, NSF Operation Manager enables efficient resource utilization by avoiding congestion and resource waste.

4.2. Developer's Management System

We extend Developer's Management System for additional functionalities as follows. As mentioned above, NSF Operation Manager requests Developer's Management System to create additional NSF instances when the existing instances of that security function are congested. On the other hand, when there are an excessive number of instances for a certain security function, NSF Operation Manager requests Developer's Management System to eliminate some of the NSF instances. As a response to such requests, Developer's Management System creates and/or removes NSF instances. Once it creates a new NSF instance or removes an existing NSF instance, the changes must be notified to NSF Operation Manager.

4.3. Network Security Function Forwarder (NSFF)

It is responsible for the following two functionalities: (1) Initially forwarding the incoming traffic/packets to Network Security Sub-Module, as described in the I2NSF information model for NSF-facing interface [capability-im]. (2) Forwarding the traffic/packets to the matched NSF with the NSF profile which is specified in a Packet Forwarding Header.

An NSFF takes a gateway functionality, so it receives incoming traffic/packets first and attaches outer encapsulation in order to forward the traffic/packets to Network Sub-Module [capability-im]. The example of Network Sub-Moudle is a firewall which performs packet header inspection. This Network Security Sub-Module attaches a Packet Forwarding Header between the outer encapsulation and the original packet and specifies NSF Profile in that header so that it can be forwarded to Content Security Sub-Module or Mitigate Sub-Module for advanced inspection.

When receiving a packet attached with a packet forwarding header of a specific NSF profile, an NSFF searches for an available NSF instance which provides the network security service corresponding to (matching with) the NSF profile and forward the packet to the NSF instance. If an NSF decides that the packet requires further inspection via another type of network security function, it constructs a packet forwarding header specified with (including) the NSF profile of the advanced network security function, attaches the header to the packet, and then sends the resulting packet to the NSFF. Once receiving the packet, the NSFF checks the NSF profile specified in the packet forwarding header. Then it searches for an NSF instance matching with the NSF profile by consulting with NSF Operation Manager, and finally forwards the packet to the NSF instance.

5. Information for Traffic Steering

This section describes the details of Packet Forwarding Header and NSF Forwarding Information for traffic forwarding between NSFs in the I2NSF system

5.1. Packet Forwarding Header

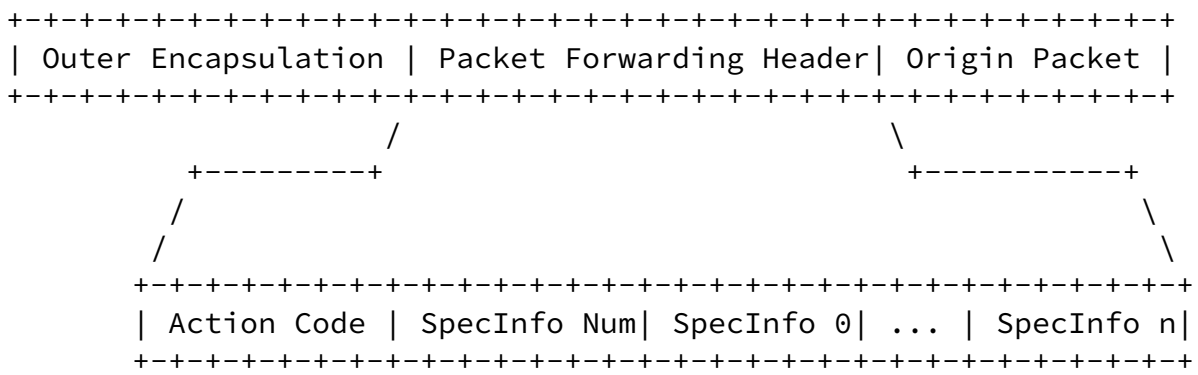

Figure 2: Packet Forwarding Header Format

An NSF uses Packet Forwarding Header to inform an NSFF(NSF Forwarder) of its inspection results and/or an advanced security inspection which is further required. As shown in Figure 2, Packet Forwarding Header consists of a single Action Code and a variable number of SpecInfo fields. The Action Code field has a value out of "allow", "deny", "advanced", and "mirror". SpecInfo Num field represents how many SpecInfos are included in this Packet Forwarding Header, and each SpecInfo includes a part of NSF Profile which describes the capabilities of an NSF required for an advanced security inspection. For instance, the value of SepcInfo can be "syn-flood-mitigate", "udp-flood-mitigate" or "content-matching-tcp" etc., which describes the service profile of an NSF.

5.2. NSF Forwarding Information

The NSF-facing interface takes charge of core functions for steering packets in the I2NSF system.

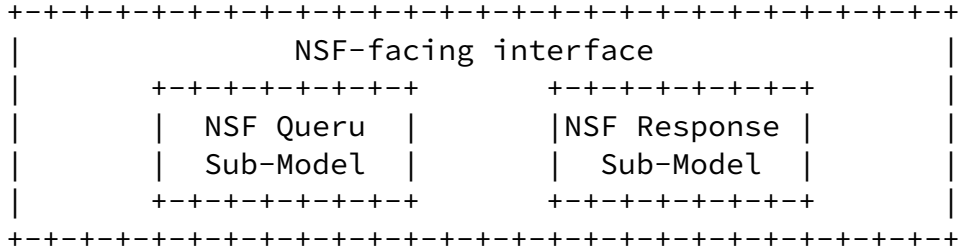

Figure 3: NSF-facing interface Model Design

5.2.1. Query of NSF forwarding information

An NSF can request an NSFF to forward packets to another NSF for more advanced security inspection of the packets. In this case, if the NSFF fails to find an NSF that can provide the security capabilities required for the advanced inspection in its forwarding information table, it sends a query to NSF Operation Manager through NSF-facing interface. The query contains an NSF profile which describes the security required for the advanced inspection capabilities. We share the definition of NSF profile in Section 4.4 of [i2nsf-reg-inf-im].

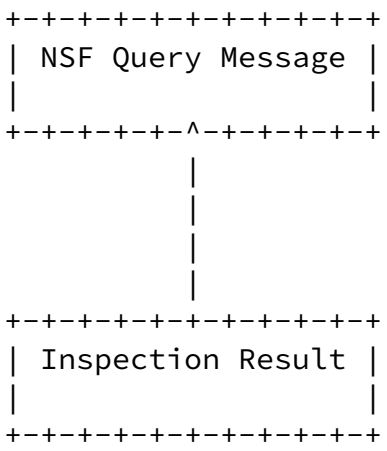

Figure 4: NSF Query Message

5.2.2. Response of NSF forwarding information

NSF Operation Manager maintains an information table of every NSF running in the system. If it receives the query from an NSFF, NSF Operation Manager searches the table for an NSF matching with the NSF profile included in the query. If there are multiple candidate NSFs, NSF Operation Manager could further consider the current workload levels of those NSFs. After choosing an NSF, it notifies the NSFF of the network forwarding information of the chosen NSF. The network forwarding information consists of IPv4 address, IPv6 address, supported transport protocols, and location information.

- IP address : As unique indetifier of an NSFs, IP address is the basic network information that allows forwarding packets to the NSF.

- Supported Transport Protocol : In order to forward packets to an NSF, it is essential to figure out which transport protocol(s) the NSF supports. Examples of the transport protocols are as follows: Virtual Extensible LAN (VXLAN) [RFC7348], Generic Protocol Extension for VXLAN (VXLAN-GPE) [nvo3-vxlan-gpe], Generic Route

Encapsulation (GRE), Ethernet). An NSFF performs encapsulation of the packets to forward as defined in the transport protocol.

- Location Information : NSFs in the system can be distributed over a wide physical region. Unlike IP address, location information specifies the physical location of an NSF. Thus, NSF Operation Manager can consider physical proximity as an additional factor to select an NSF.

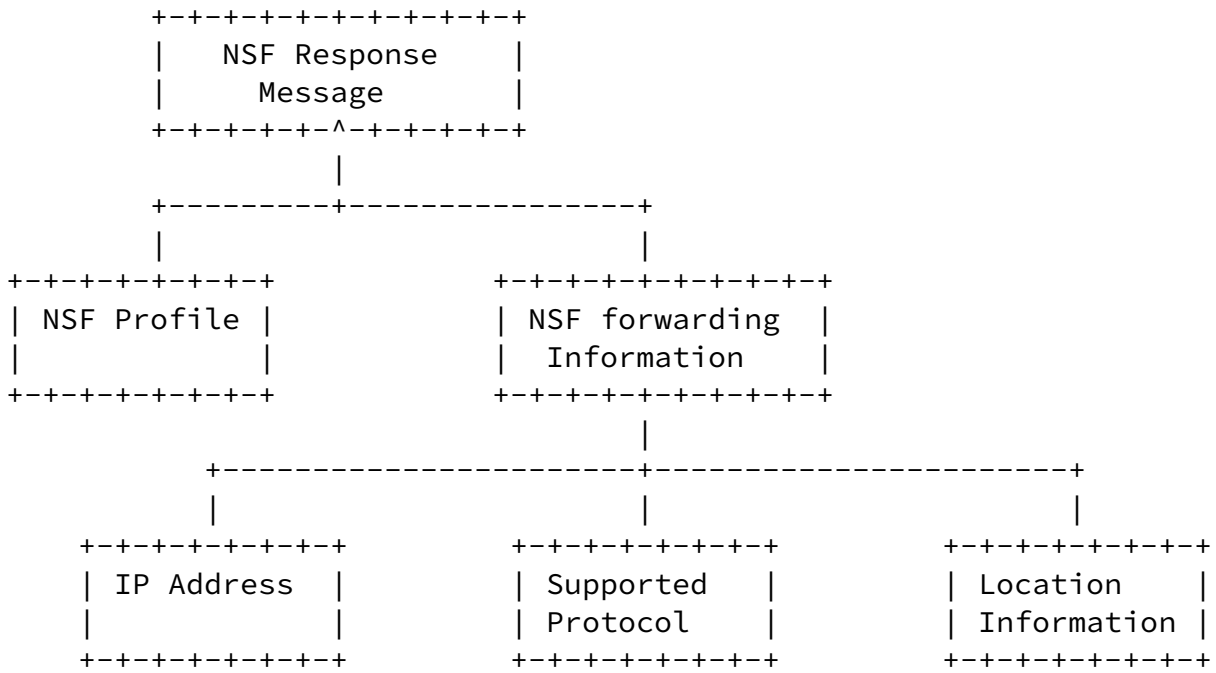

Figure 5: NSF Response Message

6. Use Cases

This section introduces two use cases for the NSF-triggered Traffic Steering Framework: (1) Enforcing Different NSFs Depending on a Packet Source's Trust Level, (2) Effective Load Balancing with Dynamic NSF Instantiation.

6.1. Enforcing Different NSFs Depending on a Packet Source's Trust Level

In the proposed architecture, all incoming packets initially arrive at the NSFF. We assume that the current security policy forces all incoming packets to be by default inspected by a firewall in this scenario. Thus the NSFF forwards the received packets to a firewall instance. Then the firewall identifies the source of the traffic and evaluates the trust level of the source. If the traffic comes from a

trusted source, it is likely to be benign. In this case, the traffic is just forwarded to the destination without further detailed inspection via different types of security functions as illustrated in Figure 6-(a). Otherwise if the traffic comes from an untrusted source, the firewall attaches a packet forwarding header including the NSF profile corresponding to DPI to the packet and returns the resulting packet to the NSFF. Once receiving the packet, the NSFF forwards the packet to the DPI instance which will perform detailed inspection for the packet payload. Figure 6-(b) illustrates this case.

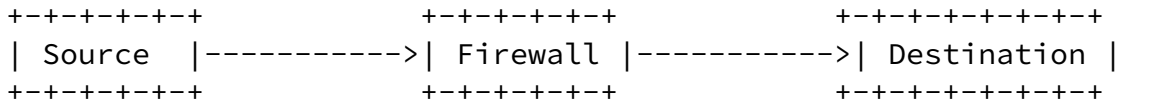

(a) Traffic flow of trusted source

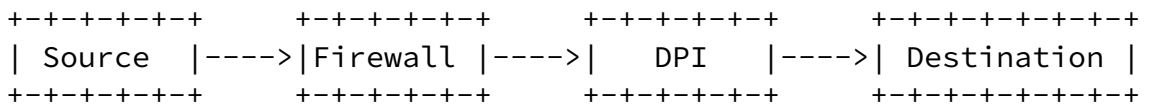

(b) Traffic flow of untrusted source

Figure 6: Different path allocation depending on source of traffic

6.2. Effective Load Balancing with Dynamic NSF Instantiation

In a large-scale network domain, there typically exist a large number of NSF instances that provide various security services. It is possible that a specific NSF instance experiences an excessive amount of traffic beyond its capacity. In this case, it is required to allocate some of the traffic to another available instance of the same security function. If there are no additional instances of the same security function available, we need to create a new NSF instance and then direct the subsequent traffic to the new instance. In this way, we can avoid service congestion and achieve more efficient resource utilization.

This process is commonly called load balancing. In our proposed architecture, NSF Operation Manager performs periodic monitoring of the load status of available NSF instances. In addition, it is possible to dynamically generate a new NSF instance through Developer's Management System. With these functionalities along with the flexible traffic steering mechanism, we can eventually provide load balancing service.

The following describes the detailed process of load balancing when congestion occurs at the firewall instance:

1. NSF Operation Manager detects that the firewall instance is receiving too much requests. Currently, there are no additional firewall instances available.

2. NSF Operation Manager requests Developer's Management System to create a new firewall instance.

3. Developer's Management System creates a new firewall instance and then registers the information of the new firewall instance to NSF Operation Manager.

4. NSF Operation Manager updates the SFC Information Table to reflect the new firewall instance, and notifies NSF and NSFF of this update.

5. According to the new forwarding information, the NSFF forwards the subsequent traffic to the new firewall instance. As a result, we can effectively alleviate the burden of the existing firewall instance.

7. Security Considerations

To enable security function chaining in the I2NSF framework, we adopt the additional components in the SFC architecture. Thus, this document shares the security considerations of the SFC architecture that are specified in [RFC7665] for the purpose of achieving secure communication among components in the proposed architecture.

8. Acknowledgements

This work was supported by Institute for Information and communications Technology Promotion(IITP) grant funded by the Korea government(MSIP) (No.R-20160222-002755, Cloud based Security Intelligence Technology Development for the Customized Security Service Provisioning).

9. References

9.1. Normative References

[RFC7665] Boucadair, M. and C. Jacquenet, "Software-Defined Networking: A Perspective from within a Service Provider Environment", RFC 7665, March 2014.

[RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, August 2014.

[sfc-ns-use-cases] Wang, E., Leung, K., Felix, J., and J. Iyer, "Service Function Chaining Use Cases for Network Security", draft-wang-sfc-ns-use-cases-02 , October

2016.

9.2. Informative References

[RFC7498] Quinn, P. and T. Nadeau, "Problem Statement for Service Function Chaining", RFC 7498, April 2015.

[capability-im] Xia, L., Strassner, J., Basile, C., and D. Lopez, "Information Model of NSFs Capabilities", draft-xibassnez-i2nsf-capability-01 (work in progress), March 2017.

[i2nsf-reg-inf-im] Hyun, S., Woo, S., Yeo, Y., Jeong, J., and J. Park, "Registration Interface Information Model", draft-hyun-i2nsf-registration-interface-im-01 (work in progress), March 2017.

[i2nsf-framework] Lopez, D., Lopez, E., Dunbar, L., Strassner, J., and R. Kumar, "Framework for Interface to Network Security Functions", draft-ietf-i2nsf-framework-04 (work in progress), October 2016.

[i2nsf-problem] Hares, S., Lopez, D., Zarny, M., Jacquenet, C., Kumar, R., and J. Jeong, "I2NSF Problem Statement and Use cases", draft-ietf-i2nsf-problem-and-use-cases-11 (work in progress), March 2017.

[i2nsf-terminology] Hares, S., Strassner, J., Lopez, D., Xia, L., and H. Birkholz, "Interface to Network Security Functions (I2NSF) Terminology", draft-ietf-i2nsf-terminology-03 (work in progress), March 2017.

[ONF-SFC-Architecture] ONF, "L4-L7 Service Function Chaining Solution Architecture", June 2015.

[nvo3-vxlan-gpe] Maino, F., Kreeger, L., and U. Elzur,

"Generic Protocol Extension for VXLAN", draft-ietf-nvo3-vxlan-gpe-03 (work in progress), October 2016.

Appendix A. Changes from draft-hyun-i2nsf-nsf-triggered-steering-01

The following changes are made from draft-hyun-i2nsf-nsf-triggered-steering-01:

- Section 5 is added to explain more details of Packet Forwarding Header and NSF Forwarding Information than the previous version.
- In Section 5.1, the explanation of Packet Forwarding Header is clarified more clearly.
- In Section 5.2, we specify the details of NSF Forwarding Information.

Authors' Addresses

Sangwon Hyun
Department of Software
Sungkyunkwan University
2066 Seobu-Ro, Jangan-Gu
Suwon, Gyeonggi-Do 16419
Republic of Korea

Phone: +82 31 290 7222
Fax: +82 31 299 6673
EMail: swhyun77@skku.edu
URI: http://imtl.skku.ac.kr/

Jaehoon Paul Jeong
Department of Software
Sungkyunkwan University
2066 Seobu-Ro, Jangan-Gu
Suwon, Gyeonggi-Do 16419
Republic of Korea

Phone: +82 31 299 4957
Fax: +82 31 290 7996
EMail: pauljeong@skku.edu
URI: http://iotlab.skku.edu/people-jaehoon-jeong.php

SangUk Woo
Department of Software
Sungkyunkwan University
2066 Seobu-Ro, Jangan-Gu
Suwon, Gyeonggi-Do 16419
Republic of Korea

Phone: +82 31 290 7222
Fax: +82 31 299 6673
EMail: suwoo@imtl.skku.ac.kr,
URI: http://imtl.skku.ac.kr/index.php?mid=member_student

YunSuk Yeo
Department of Software
Sungkyunkwan University
2066 Seobu-Ro, Jangan-Gu
Suwon, Gyeonggi-Do 16419
Republic of Korea

Phone: +82 31 290 7222
Fax: +82 31 299 6673
EMail: yunsuk@imtl.skku.ac.kr,
URI: http://imtl.skku.ac.kr/index.php?mid=member_student

Jung-Soo Park
Electronics and Telecommunications Research Institute
218 Gajeong-Ro, Yuseong-Gu
Daejeon 305-700
Republic of Korea

Phone: +82 42 860 6514
EMail: pjs@etri.re.kr