

Internet Engineering Task Force (IETF)
Request for Comments: 8329
Category: Informational
ISSN: 2070-1721

D. Lopez
Telefonica I+D
E. Lopez
Curveball Networks
L. Dunbar
J. Strassner
Huawei
R. Kumar
Juniper Networks
February 2018

Framework for Interface to Network Security Functions

Abstract

This document describes the framework for Interface to Network Security Functions (I2NSF) and defines a reference model (including major functional components) for I2NSF. Network Security Functions (NSFs) are packet-processing engines that inspect and optionally modify packets traversing networks, either directly or in the context of sessions to which the packet is associated.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8329>.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions Used in This Document	3
2.1. Acronyms	3
2.2. Definitions	4
3. I2NSF Reference Model	5
3.1. I2NSF Consumer-Facing Interface	6
3.2. I2NSF NSF-Facing Interface	6
3.3. I2NSF Registration Interface	7
4. Threats Associated with Externally Provided NSFs	8
5. Avoiding NSF Ossification	9
6. The Network Connecting I2NSF Components	10
6.1. Network Connecting I2NSF Users and the I2NSF Controller	10
6.2. Network Connecting the I2NSF Controller and NSFs	10
6.3. Interface to vNSFs	11
6.4. Consistency	12
7. I2NSF Flow Security Policy Structure	13
7.1. Customer-Facing Flow Security Policy Structure	13
7.2. NSF-Facing Flow Security Policy Structure	14
7.3. Differences from ACL Data Models	16
8. Capability Negotiation	16
9. Registration Considerations	17
9.1. Flow-Based NSF Capability Characterization	17
9.2. Registration Categories	18
10. Manageability Considerations	21
11. Security Considerations	22
12. IANA Considerations	22
13. References	22
13.1. Normative References	22
13.2. Informative References	23
Acknowledgements	24
Authors' Addresses	25

1. Introduction

This document describes the framework for Interface to Network Security Functions (I2NSF) and defines a reference model (including major functional components) for I2NSF. This includes an analysis of the threats implied by the deployment of Network Security Functions (NSFs) that are externally provided. It also describes how I2NSF facilitates implementing security functions in a technology- and vendor-independent manner in Software-Defined Networking (SDN) and Network Function Virtualization (NFV) environments, while avoiding potential constraints that could limit the capabilities of NSFs.

I2NSF use cases [RFC8192] call for standard interfaces for users of an I2NSF system (e.g., applications, overlay or cloud network management system, or enterprise network administrator or management system) to inform the I2NSF system which I2NSF functions should be applied to which traffic (or traffic patterns). The I2NSF system realizes this as a set of security rules for monitoring and controlling the behavior of different traffic. It also provides standard interfaces for users to monitor flow-based security functions hosted and managed by different administrative domains.

[RFC8192] also describes the motivation and the problem space for an Interface to Network Security Functions system.

2. Conventions Used in This Document

This memo does not propose a protocol standard, and the use of words such as "should" follow their ordinary English meaning and not that for normative languages defined in [RFC2119] [RFC8174].

2.1. Acronyms

The following acronyms are used in this document:

DOTS: Distributed Denial-of-Service Open Threat Signaling
IDS: Intrusion Detection System
IoT: Internet of Things
IPS: Intrusion Protection System
NSF: Network Security Function

2.2. Definitions

The following terms, which are used in this document, are defined in the I2NSF terminology document [I2NSF-TERMS]:

- Capability
- Controller
- Firewall
- I2NSF Consumer
- I2NSF NSF-Facing Interface
- I2NSF Policy Rule
- I2NSF Producer
- I2NSF Registration Interface
- I2NSF Registry
- Interface
- Interface Group
- Intrusion Detection System
- Intrusion Protection System
- Network Security Function
- Role

3. I2NSF Reference Model

Figure 1 shows a reference model (including major functional components and interfaces) for an I2NSF system. This figure is drawn from the point of view of the Network Operator Management System; hence, this view does not assume any particular management architecture for either the NSFs or how the NSFs are managed (on the developer's side). In particular, the Network Operator Management System does not participate in NSF data-plane activities.

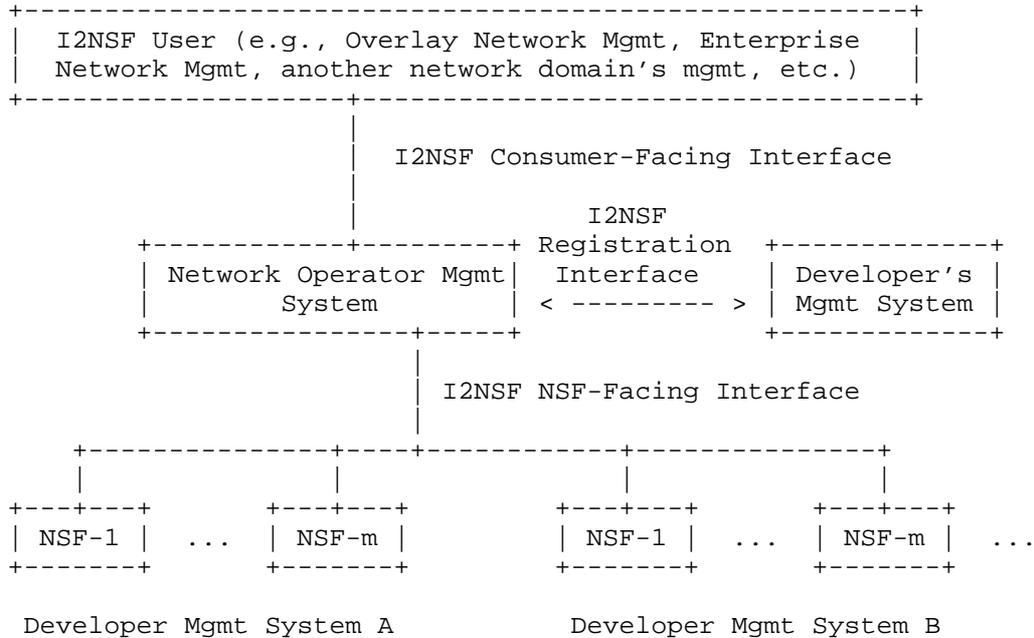


Figure 1: I2NSF Reference Model

When defining I2NSF Interfaces, this framework adheres to the following principles:

- o It is agnostic of network topology and NSF location in the network
- o It is agnostic of provider of the NSF (i.e., independent of the way that the provider makes an NSF available, as well as how the provider allows the NSF to be managed)
- o It is agnostic of any vendor-specific operational, administrative, and management implementation; hosting environment; and form factor (physical or virtual)

- o It is agnostic to NSF control-plane implementation (e.g., signaling capabilities)
- o It is agnostic to NSF data-plane implementation (e.g., encapsulation capabilities)

In general, all I2NSF Interfaces should require at least mutual authentication and authorization for their use. Other security and privacy considerations are specified in Section 11.

3.1. I2NSF Consumer-Facing Interface

The I2NSF Consumer-Facing Interface is used to enable different users of a given I2NSF system to define, manage, and monitor security policies for specific flows within an administrative domain. The location and implementation of I2NSF policies are irrelevant to the consumer of I2NSF policies.

Some examples of I2NSF Consumers include:

- o A video-conference network manager that needs to dynamically inform the underlay network to allow, rate-limit, or deny flows (some of which are encrypted) based on specific fields in the packets for a certain time span.
- o Enterprise network administrators and management systems that need to request their provider network to enforce specific I2NSF policies for particular flows.
- o An IoT management system sending requests to the underlay network to block flows that match a set of specific conditions.

3.2. I2NSF NSF-Facing Interface

The I2NSF NSF-Facing Interface (NSF-Facing Interface for short) is used to specify and monitor flow-based security policies enforced by one or more NSFs. Note that the I2NSF Management System does not need to use all features of a given NSF, nor does it need to use all available NSFs. Hence, this abstraction enables NSF features to be treated as building blocks by an NSF system; thus, developers are free to use the security functions defined by NSFs independent of vendor and technology.

Flow-based NSFs [RFC8192] inspect packets in the order that they are received. Note that all Interface Groups require the NSF to be registered using the Registration Interface. The interface to flow-based NSFs can be categorized as follows:

1. NSF Operational and Administrative Interface: an Interface Group used by the I2NSF Management System to program the operational state of the NSF; this also includes administrative control functions. I2NSF Policy Rules represent one way to change this Interface Group in a consistent manner. Since applications and I2NSF Components need to dynamically control the behavior of traffic that they send and receive, much of the I2NSF effort is focused on this Interface Group.
2. Monitoring Interface: an Interface Group used by the I2NSF Management System to obtain monitoring information from one or more selected NSFs. Each interface in this Interface Group could be a query- or a report-based interface. The difference is that a query-based interface is used by the I2NSF Management System to obtain information, whereas a report-based interface is used by the NSF to provide information. The functionality of this Interface Group may also be defined by other protocols, such as SYSLOG and DOTS. The I2NSF Management System may take one or more actions based on the receipt of information; this should be specified by an I2NSF Policy Rule. This Interface Group does NOT change the operational state of the NSF.

This document uses the flow-based paradigm to develop the NSF-Facing Interface. A common trait of flow-based NSFs is in the processing of packets based on the content (e.g., header/payload) and/or context (e.g., session state and authentication state) of the received packets. This feature is one of the requirements for defining the behavior of I2NSF.

3.3. I2NSF Registration Interface

NSFs provided by different vendors may have different capabilities. In order to automate the process of utilizing multiple types of security functions provided by different vendors, it is necessary to have a dedicated interface for vendors to define the capabilities of (i.e., register) their NSFs. This interface is called the I2NSF Registration Interface.

An NSF's capabilities can be either pre-configured or retrieved dynamically through the I2NSF Registration Interface. If a new function that is exposed to the consumer is added to an NSF, then the capabilities of that new function should be registered in the I2NSF

Registry via the I2NSF Registration Interface, so that interested management and control entities may be made aware of them.

4. Threats Associated with Externally Provided NSFs

While associated with a much higher flexibility, and in many cases a necessary approach given the deployment conditions, the usage of externally provided NSFs implies several additional concerns in security. The most relevant threats associated with a security platform of this nature are:

- o An unknown/unauthorized user can try to impersonate another user that can legitimately access external NSF services. This attack may lead to accessing the I2NSF Policy Rules and applications of the attacked user and/or generating network traffic outside the security functions with a falsified identity.
- o An authorized user may misuse assigned privileges to alter the network traffic processing of other users in the NSF underlay or platform.
- o A user may try to install malformed elements (e.g., I2NSF Policy Rules or configuration files) to directly take control of an NSF or the whole provider platform. For example, a user may exploit a vulnerability on one of the functions or may try to intercept or modify the traffic of other users in the same provider platform.
- o A malicious provider can modify the software (e.g., the operating system or the specific NSF implementation) to alter the behavior of one or more NSFs. This event has a high impact on all users accessing NSFs, since the provider has the highest level of privileges controlling the operation of the software.
- o A user that has physical access to the provider platform can modify the behavior of the hardware/software components or the components themselves. For example, the user can access a serial console (most devices offer this interface for maintenance reasons) to access the NSF software with the same level of privilege of the provider.

The use of authentication, authorization, accounting, and audit mechanisms is recommended for all users and applications to access the I2NSF environment. This can be further enhanced by requiring attestation to be used to detect changes to the I2NSF environment by authorized parties. The characteristics of these procedures will define the level of assurance of the I2NSF environment.

5. Avoiding NSF Ossification

A basic tenet in the introduction of I2NSF standards is that the standards should not make it easier for attackers to compromise the network. Therefore, in constructing standards for I2NSF Interfaces as well as I2NSF Policy Rules, it is equally important to allow support for specific functions, as this enables the introduction of NSFs that evolve to meet new threats. Proposed standards for I2NSF Interfaces to communicate with NSFs, as well as I2NSF Policy Rules to control NSF functionality, should not:

- o Narrowly define NSF categories, or their roles, when implemented within a network. Security is a constantly evolving discipline. The I2NSF framework relies on an object-oriented information model, which provides an extensible definition of NSF information elements and categories; it is recommended that implementations follow this model.
- o Attempt to impose functional requirements or constraints, either directly or indirectly, upon NSF developers. Implementations should be free to realize and apply NSFs in a way that best suits the needs of the applications and environment using them.
- o Be a limited lowest common denominator approach, where interfaces can only support a limited set of standardized functions, without allowing for developer-specific functions. NSFs, interfaces, and the data communicated should be extensible, so that they can evolve to protect against new threats.
- o Be seen as endorsing a best common practice for the implementation of NSFs; rather, this document describes the conceptual structure and reference model of I2NSF. The purpose of this reference model is to define a common set of concepts in order to facilitate the flexible implementation of an I2NSF system.

To prevent constraints on NSF developers' creativity and innovation, this document recommends flow-based NSF interfaces to be designed from the paradigm of processing packets in the network. Flow-based NSFs are ultimately packet-processing engines that inspect packets traversing networks, either directly or in the context of sessions in which the packet is associated. The goal is to create a workable interface to NSFs that aids in their integration within legacy, SDN, and/or NFV environments, while avoiding potential constraints that could limit their functional capabilities.

6. The Network Connecting I2NSF Components

6.1. Network Connecting I2NSF Users and the I2NSF Controller

As a general principle, in the I2NSF environment, users directly interact with the I2NSF Controller. Given the role of the I2NSF Controller, a mutual authentication of users and the I2NSF Controller is required. I2NSF does not mandate a specific authentication scheme; it is up to the users to choose available authentication schemes based on their needs.

Upon successful authentication, a trusted connection between the user and the I2NSF Controller (or an endpoint designated by it) will be established. This means that a direct, physical point-to-point connection, with physical access restricted according to access control, must be used. All traffic to and from the NSF environment will flow through this connection. The connection is intended not only to be secure but trusted in the sense that it should be bound to the mutual authentication between the user and the I2NSF Controller, as described in [I2NSF-ATTESTATION]. The only possible exception is when the required level of assurance is lower (see Section 4.1 of [I2NSF-ATTESTATION]), in which case the user must be made aware of this circumstance.

6.2. Network Connecting the I2NSF Controller and NSFs

Most likely, the NSFs are not directly attached to the I2NSF Controller; for example, NSFs can be distributed across the network. The network that connects the I2NSF Controller with the NSFs can be the same network that carries the data traffic, or it can be a dedicated network for management purposes only. In either case, packet loss could happen due to failure, congestion, or other reasons.

Therefore, the transport mechanism used to carry management data and information must be secure. It does not have to be a reliable transport; rather, a transport-independent reliable messaging mechanism is required, where communication can be performed reliably (e.g., by establishing end-to-end communication sessions and by introducing explicit acknowledgement of messages into the communication flow). Latency requirements for control message delivery must also be evaluated. Note that monitoring does not require reliable transport.

The network connection between the I2NSF Controller and NSFs can rely on either:

- o Open environments, where one or more NSFs can be hosted in one or more external administrative domains that are reached via secure external network connections. This requires more restrictive security control to be placed over the I2NSF Interface. The information over the I2NSF Interfaces shall be exchanged by using the trusted connection described in Section 6.1, or
- o Closed environments, where there is only one administrative domain. Such environments provide a more ****isolated**** environment but still communicate over the same set of I2NSF Interfaces present in open environments (see above). Hence, the security control and access requirements for closed environments are the same as those for open environments.

The network connection between the I2NSF Controller and NSFs will use the trusted connection mechanisms described in Section 6.1.

Following these mechanisms, the connections need to rely on the use of properly verified peer identities (e.g., through an Authentication, Authorization, and Accounting (AAA) framework). The implementations of identity management functions, as well as the AAA framework, are out of scope for I2NSF.

6.3. Interface to vNSFs

There are some unique characteristics in interfacing to virtual NSFs (vNSFs):

- o There could be multiple instantiations of one single NSF that has been distributed across a network. When different instantiations are visible to the I2NSF Controller, different policies may be applied to different instantiations of an individual NSF (e.g., to reflect the different roles that each vNSF is designated for). Therefore, it is recommended that Roles, in addition to the use of robust identities, be used to distinguish between different instantiations of the same vNSF. Note that this also applies to physical NSFs.
- o When multiple instantiations of one single NSF appear as one single entity to the I2NSF Controller, the I2NSF Controller may need to get assistance from other entities in the I2NSF Management System and/or delegate the provisioning of the multiple instantiations of the (single) NSF to other entities in the I2NSF Management System. This is shown in Figure 2 below.

- o Policies enforced by one vNSF instance may need to be retrieved and moved to another vNSF of the same type when user flows are moved from one vNSF to another.
- o Multiple vNSFs may share the same physical platform.
- o There may be scenarios where multiple vNSFs collectively perform the security policies needed.

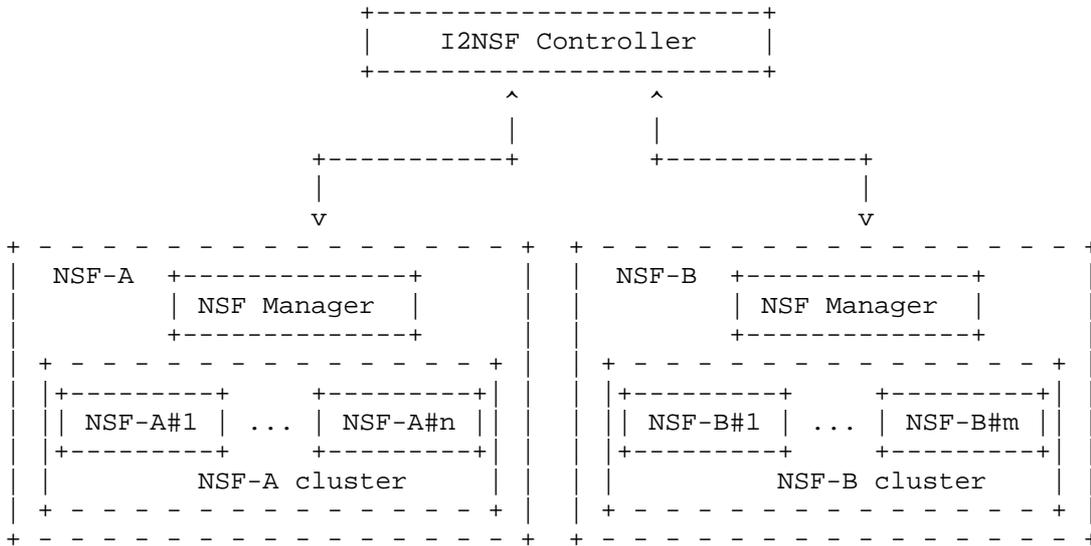


Figure 2: Cluster of NSF Instantiations Management

6.4. Consistency

There are three basic models of consistency:

- o centralized, which uses a single manager to impose behavior
- o decentralized, in which managers make decisions without being aware of each other (i.e., managers do not exchange information)
- o distributed, in which managers make explicit use of information exchange to arrive at a decision

This document does NOT make a recommendation on which of the above three models to use. I2NSF Policy Rules, coupled with an appropriate management strategy, is applicable to the design and integration of any of the above three consistency models.

7. I2NSF Flow Security Policy Structure

Even though security functions come in a variety of form factors and have different features, provisioning to flow-based NSF's can be standardized by using policy rules.

In this version of I2NSF, policy rules are limited to imperative paradigms. I2NSF is using an Event-Condition-Action (ECA) policy, where:

- o An Event clause is used to trigger the evaluation of the Condition clause of the I2NSF Policy Rule.
- o A Condition clause is used to determine whether or not the set of Actions in the I2NSF Policy Rule can be executed or not.
- o An Action clause defines the type of operations that may be performed on this packet or flow.

Each of the above three clauses are defined to be Boolean clauses. This means that each is a logical statement that evaluates to either TRUE or FALSE.

The above concepts are described in detail in [I2NSF-CAPABILITIES].

7.1. Customer-Facing Flow Security Policy Structure

This layer is for the user's network management system to express and monitor the needed flow security policies for their specific flows.

Some customers may not have the requisite security skills to express security requirements or policies that are precise enough to implement in an NSF. These customers may instead express expectations (e.g., goals or intent) of the functionality desired by their security policies. Customers may also express guidelines, such as which types of destinations are (or are not) allowed for certain users. As a result, there could be different levels of content and abstractions used in Service Layer policies. Here are some examples of more abstract security policies that can be developed based on the I2NSF-defined Customer-Facing Interface:

- o Enable Internet access for authenticated users
- o Any operation on a HighValueAsset must use the corporate network
- o The use of FTP from any user except the CxOGroup must be audited

- o Streaming media applications are prohibited on the corporate network during business hours
- o Scan email for malware detection; protect traffic to corporate network with integrity and confidentiality
- o Remove tracking data from Facebook [website = *.facebook.com]

One flow policy over the Customer-Facing Interface may need multiple NSF's at various locations to achieve the desired enforcement. Some flow security policies from users may not be granted because of resource constraints. [I2NSF-DEMO] describes an implementation of translating a set of 1) user policies to flow policies and 2) flow policies to individual NSF's.

I2NSF will first focus on user policies that can be modeled as closely as possible to the flow security policies used by individual NSF's. An I2NSF user flow policy should be similar in structure to the structure of an I2NSF Policy Rule, but with more of a user-oriented expression for the packet content, the context, and other parts of an ECA policy rule. This enables the user to construct an I2NSF Policy Rule without having to know the exact syntax of the desired content (e.g., actual tags or addresses) to match in the packets. For example, when used in the context of policy rules over the Client-Facing Interface:

- o An Event can be "the client has passed the AAA process"
- o A Condition can be matching the user identifier or from specific ingress or egress points
- o An Action can be establishing an IPsec tunnel

7.2. NSF-Facing Flow Security Policy Structure

The NSF-Facing Interface is to pass explicit rules to individual NSF's to treat packets, as well as methods to monitor the execution status of those functions.

Here are some examples of Events over the NSF-Facing Interface:

- o time == 08:00
- o notification that a NSF state changes from standby to active
- o user logon or logoff

Here are some examples of Conditions over the NSF-Facing Interface:

- o Packet content values that look for one or more packet headers, data from the packet payload, bits in the packet, or data that are derived from the packet.
- o Context values that are based on measured and/or inferred knowledge, which can be used to define the state and environment in which a managed entity exists or has existed. In addition to state data, this includes data from sessions, direction of the traffic, time, and geo-location information. State refers to the behavior of a managed entity at a particular point in time. Hence, it may refer to situations in which multiple pieces of information that are not available at the same time must be analyzed. For example, tracking established TCP connections (connections that have gone through the initial three-way handshake).

Actions to individual flow-based NSFs include:

- o Actions performed on ingress packets, such as pass, drop, rate limiting, and mirroring.
- o Actions performed on egress packets, such as invoke signaling, tunnel encapsulation, packet forwarding, and/or transformation.
- o Applying a specific functional profile or signature -- e.g., an IPS Profile, a signature file, an anti-virus file, or a URL filtering file. Many flow-based NSFs utilize profile and/or signature files to achieve more effective threat detection and prevention. It is not uncommon for an NSF to apply different profiles and/or signatures for different flows. Some profiles/signatures do not require any knowledge of past or future activities, while others are stateful and may need to maintain state for a specific length of time.

The functional profile or signature file is one of the key properties that determine the effectiveness of the NSF and is mostly NSF specific today. The rulesets and software interfaces of I2NSF aim to specify the format to pass profile and signature files while supporting specific functionalities of each.

Policy consistency among multiple security function instances is very critical because security policies are no longer maintained by one central security device; instead, they are enforced by multiple security functions instantiated at various locations.

7.3. Differences from ACL Data Models

Policy rules are very different from Access Control Lists (ACLs). An ACL is NOT a policy. Rather, policies are used to manage the construction and life cycle of an ACL.

[ACL-YANG] has defined rules for ACLs supported by most routers/switches that forward packets based on their L2, L3, or sometimes L4 headers. The actions for ACLs include Pass, Drop, or Redirect.

The functional profiles (or signatures) for NSFs are not present in [ACL-YANG] because the functional profiles are unique to specific NSFs. For example, most IPS/IDS implementations have their proprietary functions/profiles. One of the goals of I2NSF is to define a common envelope format for exchanging or sharing profiles among different organizations to achieve more effective protection against threats.

The "packet content matching" of the I2NSF policies should not only include the matching criteria specified by [ACL-YANG] but also the L4-L7 fields depending on the NSFs selected.

Some flow-based NSFs need matching criteria that include the context associated with the packets. This may also include metadata.

The I2NSF "actions" should extend the actions specified by [ACL-YANG] to include applying statistics functions, threat profiles, or signature files that clients provide.

8. Capability Negotiation

It is very possible that the underlay network (or provider network) does not have the capability or resources to enforce the flow security policies requested by the overlay network (or enterprise network). Therefore, it is required that the I2NSF system support dynamic discovery capabilities, as well as a query mechanism, so that the I2NSF system can expose appropriate security services using I2NSF capabilities. This may also be used to support negotiation between a user and the I2NSF system. Such dynamic negotiation facilitates the delivery of the required security service(s). The outcome of the negotiation would feed the I2NSF Management System, which would then dynamically allocate appropriate NSFs (along with any resources needed by the allocated NSFs) and configure the set of security services that meet the requirements of the user.

When an NSF cannot perform the desired provisioning (e.g., due to resource constraints), it must inform the I2NSF Management System. The protocol needed for this security function/capability negotiation

may be somewhat correlated to the dynamic service parameter negotiation procedure described in [RFC7297]. The Connectivity Provisioning Profile (CPP) template, even though currently covering only connectivity requirements, includes security clauses such as isolation requirements and non-via nodes. Hence, it could be extended as a basis for the negotiation procedure. Likewise, the companion Connectivity Provisioning Negotiation Protocol (CPNP) could be a candidate for the negotiation procedure.

"Security-as-a-Service" would be a typical example of the kind of (CPP-based) negotiation procedures that could take place between a corporate customer and a service provider. However, more security-specific parameters have to be considered.

[I2NSF-CAPABILITIES] describes the concepts of capabilities in detail.

9. Registration Considerations

9.1. Flow-Based NSF Capability Characterization

There are many types of flow-based NSFs. Firewall, IPS, and IDS are the commonly deployed flow-based NSFs. However, the differences among them are definitely blurring, due to more powerful technology, integration of platforms, and new threats. Basic types of flow-based NSFs include:

- o Firewall -- A device or a function that analyzes packet headers and enforces policy based on protocol type, source address, destination address, source port, destination port, and/or other attributes of the packet header. Packets that do not match policy are rejected. Note that additional functions, such as logging and notification of a system administrator, could optionally be enforced as well.
- o IDS (Intrusion Detection System) -- A device or function that analyzes packets, both header and payload, looking for known events. When a known event is detected, a log message is generated detailing the event. Note that additional functions, such as notification of a system administrator, could optionally be enforced as well.
- o IPS (Intrusion Prevention System) -- A device or function that analyzes packets, both header and payload, looking for known events. When a known event is detected, the packet is rejected. Note that additional functions, such as logging and notification of a system administrator, could optionally be enforced as well.

Flow-based NSF's differ in the depth of packet header or payload they can inspect, the various session/context states they can maintain, and the specific profiles and the actions they can apply. An example of a session is as follows: allowing outbound connection requests and only allowing return traffic from the external network.

9.2. Registration Categories

Developers can register their NSF's using packet content matching categories. The Inter-Domain Routing (IDR) Flow Specification [RFC5575] has specified 12 different packet header matching types.

IP Flow Information Export (IPFIX) data [IPFIX-D] defines IP flow information and mechanisms to transmit such information. This includes flow attributes as well as information about the metering and exporting processes. Such information may be stored in an IPFIX registry [IPFIX-R]. As such, IPFIX information should be considered when defining categories of registration information.

More packet content matching types have been proposed in the IDR WG. I2NSF should reuse the packet matching types being specified as much as possible. More matching types might be added for flow-based NSF's.

Figures 3-6 below list the applicable packet content categories that can be potentially used as packet matching types by flow-based NSF's:

Packet Content Matching Capability Index	
Layer 2 Header	Layer 2 header fields: Source Destination s-VID c-VID Ethertype
Layer 3 IPv4 Header	Layer 3 header fields: protocol dest port src port src address dest address dscp length flags ttl

IPv6 Header	<ul style="list-style-type: none"> protocol/nh src port dest port src address dest address length traffic class hop limit flow label dscp
<ul style="list-style-type: none"> Layer 4 TCP SCTP DCCP UDP 	<p>Layer 4 header fields:</p> <ul style="list-style-type: none"> Port syn ack fin rst ? psh ? urg ? window sockstress <p>Note: bitmap could be used to represent all the fields</p> <ul style="list-style-type: none"> flood abuse fragment abuse Port
HTTP layer	<ul style="list-style-type: none"> hash collision http - get flood http - post flood http - random/invalid url http - slowloris http - slow read http - r-u-dead-yet (rudy) http - malformed request http - xss https - ssl session exhaustion

IETF PCP	Configurable Ports
IETF TRAM	profile

Notes:

- DCCP: Datagram Congestion Control Protocol
- PCP: Port Control Protocol
- TRAM: TURN Revised and Modernized, where TURN stands for Traversal Using Relays around NAT

Figure 3: Packet Content Matching Capability Index

Context Matching Capability Index	
Session	Session State, Bidirectional State
Time	Time span Time occurrence
Events	Event URL, variables
Location	Text string, GPS coords, URL
Connection Type	Internet (unsecured), Internet (secured by VPN, etc.), Intranet, ...
Direction	Inbound, Outbound
State	Authentication State Authorization State Accounting State Session State

Note:

These fields are used to provide context information for I2NSF Policy Rules to make decisions on how to handle traffic. For example, GPS coordinates define the location of the traffic that is entering and exiting an I2NSF system; this enables the developer to apply different rules for ingress and egress traffic handling.

Figure 4: Context Matching Capability Index

Action Capability Index	
Ingress port	SFC header termination, VxLAN header termination
Actions	Pass Deny Mirror Simple Statistics: Count (X min; Day;...) Client-Specified Functions: URL
Egress	Encap SFC, VxLAN, or other header

Note:

SFC: Service Function Chaining

Figure 5: Action Capability Index

Functional Profile Index	
Profile types	name, type, or flexible
Signature	Profile/signature URL command for the I2NSF Controller to enable/disable

Figure 6: Functional Profile Index

10. Manageability Considerations

Management of NSF's include:

- o Life-cycle management and resource management of NSF's
- o Configuration of devices, such as address configuration, device internal attributes configuration, etc.
- o Signaling
- o Policy rules provisioning

Currently, I2NSF only focuses on the policy rule provisioning part.

11. Security Considerations

The configuration, control, and monitoring of NSFs provide access to and information about security functions that are critical for delivering network security and for protecting end-to-end traffic. Therefore, it is important that the messages that are exchanged within this architecture utilize a trustworthy, robust, and fully secure communication channel. The mechanisms adopted within the solution space must include proper secure communication channels that are carefully specified for carrying the controlling and monitoring information between the NSFs and their management entity or entities. The threats associated with remotely managed NSFs are discussed in Section 4, and solutions must address those concerns.

This framework is intended for enterprise users, with or without cloud service offerings. Privacy of users must be provided by using existing standard mechanisms, such as encryption; anonymization of data should also be done if possible (depending on the transport used). Such mechanisms require confidentiality and integrity.

12. IANA Considerations

This document has no IANA actions.

13. References

13.1. Normative References

- [IPFIX-D] "IP Flow Information Export (ipfix)",
<<https://datatracker.ietf.org/wg/ipfix/documents/>>.
- [IPFIX-R] IANA, "IP Flow Information Export (IPFIX) Entities",
<<https://www.iana.org/assignments/ipfix>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997,
<<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009,
<<https://www.rfc-editor.org/info/rfc5575>>.
- [RFC7297] Boucadair, M., Jacquenet, C., and N. Wang, "IP Connectivity Provisioning Profile (CPP)", RFC 7297, DOI 10.17487/RFC7297, July 2014,
<<https://www.rfc-editor.org/info/rfc7297>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

13.2. Informative References

[ACL-YANG]

Jethanandani, M., Huang, L., Agarwal, S., and D. Blair, "Network Access Control List (ACL) YANG Data Model", Work in Progress, draft-ietf-netmod-acl-model-15, January 2018.

[I2NSF-ATTESTATION]

Pastor, A., Lopez, D., and A. Shaw, "Remote Attestation Procedures for Network Security Functions (NSFs) through the I2NSF Security Controller", Work in Progress, draft-pastor-i2nsf-nsf-remote-attestation-02, September 2017.

[I2NSF-CAPABILITIES]

Xia, L., Strassner, J., Basile, C., and D. Lopez, "Information Model of NSFs Capabilities", Work in Progress, draft-i2nsf-capability-00, September 2017.

[I2NSF-DEMO]

Xie, Y., Xia, L., and J. Wu, "Interface to Network Security Functions Demo Outline Design", Work in Progress, draft-xie-i2nsf-demo-outline-design-00, April 2015.

[I2NSF-TERMS]

Hares, S., Strassner, J., Lopez, D., Xia, L., and H. Birkholz, "Interface to Network Security Functions (I2NSF) Terminology", Work in Progress, draft-ietf-i2nsf-terminology-05, January 2018.

[RFC8192] Hares, S., Lopez, D., Zarny, M., Jacquenet, C., Kumar, R., and J. Jeong, "Interface to Network Security Functions (I2NSF): Problem Statement and Use Cases", RFC 8192, DOI 10.17487/RFC8192, July 2017, <<https://www.rfc-editor.org/info/rfc8192>>.

Acknowledgements

This document includes significant contributions from Christian Jacquenet (Orange), Seetharama Rao Durbha (Cablelabs), Mohamed Boucadair (Orange), Ramki Krishnan (Dell), Anil Lohiya (Juniper Networks), Joe Parrott (BT), Frank Xialing (Huawei), and XiaoJun Zhuang (China Mobile).

Some of the results leading to this work have received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 611458.

Authors' Addresses

Diego R. Lopez
Telefonica I+D
Editor Jose Manuel Lara, 9
Seville, 41013
Spain

Email: diego.r.lopez@telefonica.com

Edward Lopez
Curveball Networks
Chantilly, Virginia
United States of America

Email: ed@curveballnetworks.com

Linda Dunbar
Huawei Technologies
United States of America

Email: Linda.Dunbar@huawei.com

John Strassner
Huawei Technologies
Santa Clara, CA
United States of America

Email: John.sc.Strassner@huawei.com

Rakesh Kumar
Juniper Networks
United States of America

Email: rakeshkumarcloud@gmail.com

