

Network Working Group  
Internet Draft  
Intended Status: Informational

M. Zarny  
Goldman Sachs  
S. Magee  
F5  
N. Leymann  
Deutsche Telecom  
L. Dunbar  
Huawei

Expires: April 28, 2015

October 25, 2014

I2NSF Data Center Use Cases  
draft-zarny-i2nsf-data-center-use-cases-00

## Abstract

This document describes data center use cases and their requirements that a common Interface to Network Security Functions (I2NSF) needs to take into account.

## 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/lid-abstracts.html>

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

## Copyright and License Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

I2NSF Data Center Use Cases

October 25, 2014

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://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

<a href="#">1. Introduction</a>	<a href="#">3</a>
<a href="#">2. Conventions used in this document</a>	<a href="#">3</a>
<a href="#">3. Terminology</a>	<a href="#">4</a>
<a href="#">4. On-demand, elastic deployment of firewalls</a>	<a href="#">4</a>
4.1 On demand virtual firewall deployment in cloud data centers	<a href="#">5</a>
<a href="#">5. Firewall policy deployment automation</a>	<a href="#">6</a>
<a href="#">5.1 Client-specific security policy in cloud VPNs</a>	<a href="#">6</a>
<a href="#">6. Key requirements for the use cases</a>	<a href="#">6</a>
<a href="#">7. Conclusion</a>	<a href="#">7</a>
<a href="#">8. Security considerations</a>	<a href="#">8</a>
<a href="#">9. IANA considerations</a>	<a href="#">8</a>
<a href="#">10. References</a>	<a href="#">8</a>
<a href="#">10.1 Normative references</a>	<a href="#">8</a>
<a href="#">10.2 Informative references</a>	<a href="#">8</a>
<a href="#">11. Acknowledgments</a>	<a href="#">8</a>
<a href="#">12. Authors' addresses</a>	<a href="#">8</a>

Internet-Draft

I2NSF Data Center Use Cases

October 25, 2014

## [1.](#) Introduction

Enterprises today increasingly consume cloud-based network security functions. The reasons are the same as those for the move toward cloud computing: greater economies of scale; faster service delivery; greater flexibility to respond to changing requirements; faster deployment of more sophisticated solutions; among others.

The cloud security services can in theory be offered in a number of ways. They can be operated by service providers or enterprises themselves; they can be run on shared or dedicated infrastructure; they can be deployed off- or on-premises; or any combination thereof. In practice, however, since most firms today possess neither the expertise nor resources to build and manage clouds, most firms that consume cloud-based security services do so on off-premise provider-managed clouds.

In response, providers and security vendors offer cloud-based models to deliver security solutions. Providers in particular are striving to standardize the offering methodologies through efforts like Network Functions Virtualization (NFV).

I2NSF is an IETF effort to standardize the interface for network security functions offered on any kind of cloud regardless of its location or operator. Since the term "network security service" can mean many things, we will limit the term to include only the following services in this draft.

- \* Firewall
- \* DDOS/Anti-DOS (Distributed Denial-of-Service/Anti-Denial-of-Service)
- \* AAA (Authentication, Authorization, Accounting)
- \* Remote identity management

- \* Secure key management

- \* IDS/IPS (Intrusion Detection System/Intrusion Prevention System)

## [2.](#) Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Zarny, et al

Expires April 28, 2015

[Page 3]

---

Internet-Draft

I2NSF Data Center Use Cases

October 25, 2014

## [3.](#) Terminology

**Cloud-scale network resources:** Networked resources which provide various functions (related to infrastructure, platform, software, etc.) in a scalable, automated and secure fashion. Resources in the cloud may be fully owned, operated, and used by a single organization; dedicated to a single client and managed by a provider; shared amongst several clients; hosted on-premises or off-premises of an organization; or a combination thereof. In the context of this draft, a cloud offers network security services.

**DC:** Data Center

**Domain relationships:** The term "Domain" in this draft has different connotations in different scenarios:

Client <-> Provider relationship, i.e. a client requesting network service functions from its provider;

Domain A <-> Domain B relationship, i.e. one operator domain requesting network service functions from another operator domain; or

Applications <-> Network relationship, an application (e.g., cluster of servers) requesting some functions from network.

**Network function:** In the context of I2NSF, the term "network function" describes services that provide network functions including L4-L7 functions. The network service functions may not

necessarily be owned or hosted by consumers of those functions. Furthermore, the network functions may be hosted on physical appliances, inside containers, or inside VMs instantiated on common compute servers (e.g., the ETSI NFV defined Virtualized Network Functions).

Virtual Security Function: A security function that can be requested by one domain but may be owned or managed by another domain.

Cloud-based security functions: Used interchangeably with the "Virtual Security Functions" in this draft.

#### [4.](#) On-demand, elastic deployment of firewalls

Network security devices such as firewalls may need to be added or removed dynamically for a number of reasons. It may have been explicitly requested by the user, or triggered by a pre-agreed-upon service level agreement (SLA) between the user and the provider of

the service. For example, the service provider may be required to add more firewall capacity within a set timeframe whenever the bandwidth utilization hits a certain threshold for a specified period. This capacity expansion could result in adding new instances of firewalls. Likewise, a service provider may need to provision a new firewall instance in a completely new environment due to a new requirement.


The on-demand, dynamic nature of deployment essentially requires that the network security "devices" be in software or virtual form factors, rather than in a physical appliance form. (This is a provider-side concern. Users of the firewall service are agnostic, as they should, as to whether or not the firewall service is run on a VM or any other form factor. Indeed, they may not even be aware that their traffic traverses firewalls.)

Furthermore, new firewall instances need to be placed in the "right zone" (domain). The issue applies not only to multi-tenant environments where getting the tenant right is of paramount importance but also to environments owned and operated by a single organization with its own service segregation policies. For example, an enterprise may mandate that firewalls serving Internet traffic and business-to-business (B2B) traffic be separate; or that IPS/IDS

services for investment banking and non-banking traffic be separate for regulatory reasons.

#### [4.1](#) On demand virtual firewall deployment in cloud data centers

A service provider operated cloud data center could serve tens of thousands of clients. Clients' compute servers are typically hosted on virtual machines (VMs), which could be deployed across different server racks located in different parts of the data center. It is often not technically and/or financially feasible to deploy dedicated physical firewalls to suit each client's myriad security policy requirements. What is needed is the ability to dynamically deploy virtual firewalls for each client's set of servers based on established security policies and underlying network topologies.



### [5](#). Firewall policy deployment automation

Firewall configuration today is a highly complex process that involves consulting established security policies, translating those policies into firewall rules, further translating those rules into vendor-specific configuration sets, identifying all the firewalls, and pushing configurations to those firewalls.

This is often a time consuming, complex and error-prone process even within a single organization/enterprise framework. It becomes far more complex in provider-owned cloud networks that serve myriad customers.

Automation can help address many of these issues. Automation works best when it can leverage a common set of standards that will work

across multiple entities.

### [5.1](#) Client-specific security policy in cloud VPNs

Clients of service provider operated cloud data centers need not only secure virtual private networks (VPNs) but also virtual security functions that enforce the clients' security policies. The security policies may govern communications within the clients' own virtual networks and those with external networks. For example, VPN service providers may need to provide firewall and other security services to their VPN clients. Today, it is generally not possible for clients to dynamically view, much less change, what, where and how security policies are implemented on their provider-operated clouds. Indeed, no standards-based framework that allows clients to retrieve/manage security policies in a consistent manner across different providers exists.

## [6.](#) Key requirements for the use cases

The I2NSF framework should provide a set of standard interfaces that facilitate:

- \* Dynamic creation, enablement, disablement, and removal of network security applications;
- \* Policy-driven placement of new service instances in the right administrative domain;
- \* Attachment of appropriate security and traffic policies to the service instances

- \* Management of deployed instances in terms of fault monitoring, utilization monitoring, event logging, inventory, etc.

Moreover, an I2NSF must support different deployment scenarios:

- \* Single and multi-tenant environments: The term multi-tenant does not mean just different companies subscribing to a provider's

cloud offering. It can for instance cover administrative domains/departments within a single firm that require different security and traffic policies.

- \* **Premise-agnostic:** Said network security services may be deployed on premises or off premises of an organization.

The I2NSF framework should provide a standard set of interfaces that enable:

- \* Translation of security policies into functional tasks. Security policies may be carried out by one or more security service functions. For example, a security policy may be translated into an IDS/IPS policy and a firewall policy for a given application type.

- \* Translation of functional tasks into vendor-specific configuration sets. For example, a firewall policy needs to be converted to vendor-specific configurations.

- \* Retrieval of information such as configuration, utilization, status, etc. Such information may be used for monitoring, auditing, troubleshooting purposes. The above functions should be available in single- or multi-tenant environments as well as on-premise or off-premise clouds.

## [7.](#) Conclusion

The need for common interfaces to network service functions goes beyond network security functions described here. Efforts like NFV will drive efforts to address this broad need. This draft covers common network security functions deployed in data centers as a way to scope the problem set. The use cases here are relevant to service provider and large enterprise networks, and they can all benefit significantly from an I2NSF.

We recommend the IETF to start a program to establish a common framework for network security functions that will address the issues raised here.

## [8.](#) Security considerations



TBD.

## 9. IANA considerations

This document requires no IANA actions. RFC Editor: Please remove this section before publication.

## 10. References

### 10.1 Normative references

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC7297] Boucadair, M., "IP Connectivity Provisioning Profile", [RFC7297](#), April 2014.

### 10.2 Informative references

[PS] Dunbar, et al, "Dynamic Network Security as a Service Problem Statement", <[draft-dunbar-nsaas-problem-statement-00](#)>, July 2014.

[GS-NFV] ETSI NFV Group Specification, Network Functions Virtualization (NFV) Use Cases. ETSI GS NFV 001v1.1.1, 2013.

[Boucadair-framework] Boucadair, M., et al, "Differentiated Service Function Chaining Framework", <[draft-boucadair-service-chaining-framework-00](#)>; Aug 2013

[SC-MobileNetwork] Haeffner, W. and N. Leymann, "Network Based Services in Mobile Network", IETF87 Berlin, July 29, 2013

[Application-SDN] Giacomonni, J., "Application Layer SDN", Layer 123 ONF Presentation, Singapore, June 2013

## 11. Acknowledgments

We would like to acknowledge Andrew Malis for his review and contribution.

## 12. Authors' addresses

Myo Zarny  
Goldman Sachs  
Email: [myo.zarny@gs.com](mailto:myo.zarny@gs.com)

Sumandra Majee  
F5 Netowrks  
Email: lal2ghar@gmail.com

Nic Leymann  
Deutsche Telekom  
Email: n.leymann@telekom.de

Linda Dunbar  
Huawei  
Email: linda.dunbar@huawei.com

