

SFC Working Group
Internet Draft
Category: Informational

R. Krishnan
Brocade Communications
A. Ghanwani
Dell
J. Halpern
S. Kini
Ericsson
D. R. Lopez
Telefonica I+D

Expires: August 7, 2015

February 6, 2015

SFC Long-lived Flow Use Cases

[draft-ietf-sfc-long-lived-flow-use-cases-03](#)

Abstract

Long-lived flows such as file transfers, video streams are common in today's networks. In the context of service function chaining, this draft suggests use cases for dynamic bypass of certain service functions for such flows. The benefit of this approach would be to avoid expensive Layer 7 service function processing for such flows based on dynamic decisions and thus improve overall performance.

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/lid-abstracts.txt>.

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

Internet-Draft

SFC Long Lived Flow Use Case

February 2015

This Internet-Draft will expire on August, 2015.

Copyright Notice

Copyright (c) 2015 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 (<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

1.	Introduction.....	3
1.1.	Acronyms.....	4
2.	Transparent Firewall Use Case.....	4
2.1.	Event Sequence.....	5
3.	Long-tail Content CDN Use Case.....	6
3.1.	Event Sequence.....	6
4.	IPsec Management in Mobile Environments.....	7
4.1.	Event Sequence.....	8
5.	Operational Considerations.....	8
6.	IANA Considerations.....	8
7.	Security Considerations.....	9
8.	Acknowledgements.....	9
9.	References.....	9
9.1.	Normative References.....	9
9.2.	Informative References.....	9
	Authors' Addresses.....	9

1. Introduction

In the context of service function chaining (SFC), this draft suggests use cases for dynamic bypass of certain service functions for long-lived flows such as file transfers and video streams. The benefit of this approach would be to avoid expensive Layer 4-7 service function processing for such flows and improve overall performance. The focus would be only on long-lived flows which are observable and controllable from a control plane perspective; attempting dynamic bypass for short-lived flows would cause excessive control plane chattiness without any significant performance benefit.

For long-lived flows, in order to dynamically bypass certain service functions in the service function chain, the key is to make sure that the Layer 7 flow can be identified using Layer 2/3/4 fields in the packet. Examples of such flows are file transfers (which typically use FTP/SFTP) and video streams (which typically use HTTP/HTTPS) which can be mapped to a unique IP 5 tuple (IP source address, IP destination address, IP protocol, transport protocol source port, transport protocol destination port). We note that it may not always be possible to identify a Layer 7 flow based on L2/L3/L4 fields in the packet header. An example of this could be file transfers under persistent HTTP sessions where multiple files may be transferred using the same values for these fields in the packet headers.

There are cases where the transfer of large content may be split across multiple transport protocol connections. In such cases, what may be considered a large flow at the application layer is dealt with using multiple small flows at the network layer. Such cases would not fit the ones described in this document.

The definition of long-lived flow in this context can reuse the definition in [[I2RS-large-flow](#)] and [[OPSAWG-large-flow](#)], where flows are categorized into 4 types - short-lived small flows, short-lived large flows, long-lived small flows and long-lived large flows. In this draft we are concerned with the last 2 types -- long-lived

small flows and long-lived large flows -- and we refer to these as long-lived flows. This identification of long-lived flows is based on L2/L3/L4 fields in the packet header that is consistent with that the definition of a flow in IPFIX [[RFC 7011](#)].

The criteria used by the service function for identifying a long-lived Layer 4-7 flow can use similar criteria, with appropriate modification to account for long-lived small flows, as the

techniques described in [[OPSAWG-large-flow](#)] for large flow identification. The mechanics of dynamic bypass are quite different for different service functions and are described in the following sections.

For the mechanisms in this draft, our focus is on the following SFC components:

- . An SFC Control Plane Application which is responsible for implementing the control plane functionality and programming the data plane for SFC.
- . An SFC edge, which is a switch/router responsible for adding/removing the service chain header to the packets.

[1.1](#). Acronyms

CDN: Content Delivery Network

DNS: Domain Name Service

DPI: Deep Packet Inspection

eNodeB: Evolved Node B

LTE: Long-Term Evolution

NAPT: Network Address Port Translation

SecGW: Security Gateway

SFC: Service Function Chaining

[2](#). Transparent Firewall Use Case

A transparent firewall may be able to determine that a long-lived flow (e.g. video stream, file transfer) has no security issues. It is desirable to have such a long-lived flow dynamically bypass the firewall service function but continue to execute the other service functions in the chain (e.g. NAT). The key benefit is overall performance improvement. The event sequence for this use case is detailed below. For this use case, it is assumed that the firewall is transparent and does not perform any packet modification.

Note that the firewall functionality is applicable only to Layer 2/3/4 headers.

[2.1](#). Event Sequence

1. The firewall examines packets of a flow and deems that it is benign. While the criteria for how the firewall determines that the flow is benign are beyond the scope of this document, some of the criteria that may be used for this include:
 - a. The packets which are encrypted at the application layer using protocols such as HTTPS [[RFC 2660](#)] cannot be decrypted and examined further.
 - b. The packets are from a trusted source.
 - c. The packets are from a trusted application.
2. The firewall determines that the flow can be identified using a Layer 2/3/4 rule in the fast path. The firewall moves the flow from the internal slow path (which inspects every packet) to the fast path (which does only switching and skips the detailed inspection of every packet).
3. Based on the above criteria and also having identified the flow as a long-lived flow, the firewall determines that the flow is a benign one and does not need to be processed by the firewall any more.
4. The firewall signals this information to the SFC Control Plane

Application.

5. The SFC Control Plane Application assigns the flow to a different service function chain that excludes the firewall.
6. The flow continues to be monitored by the SFC edge switch/router for activity.
7. Once the flow is detected as having become inactive, the flow is aged out by the SFC edge switch/router.
8. The SFC edge switch/router signals a flow age event to the SFC Control Plane Application.
9. The SFC Control Plane Application removes the dynamic service chain association created for the flow.

3. Long-tail Content CDN Use Case

Most popular content is of interest to a number of users; typical examples are newly released movies, latest television episodes, etc. Such content is very amenable to caching. A single copy of the content is delivered to the cache; the content is delivered to multiple users from the cache.

Long-tail personalized content is of interest to only a few users; typical examples are documentaries, older movies etc. Long-tail personalized content is typically not shared by many users and is not amenable to caching [[CDNI-long-tail](#)]. Caching of such content could cause excessive thrashing of the cache.

The idea is to improve performance by identifying such long-tail content and bypassing the CDN cache in the service chain for such content. This would be dynamic in nature, since content that is not popular can later become popular and vice versa. The focus will be on long-lived content such as movies, catch up episodes which generate long-lived flows. The key benefit is overall performance improvement. The event sequence for this use case is detailed below.

For the purpose of this draft, our focus is on the following components in the CDN:

- . CDN Monitoring System: The CDN Monitoring System monitors various aspects of the content such as
 - o Dynamic Content Usage: Number of users simultaneously viewing the same content.
 - o Content Life: If the content is long-lived or short-lived. Examples of long-lived content are movies, catch up episodes, etc., while examples of short-lived content are video clips, advertisements, etc.
- . CDN Cache: This is the node in the network where the content is cached.

For a general overview of CDNs, see [[CDN-overview](#)].

[3.1](#). Event Sequence

1. The CDN Monitoring System monitors the numbers of users and type of content being accessed. By default, we assume the CDN Cache is bypassed.

2. If the number of users viewing the same content exceeds a pre-programmed up-threshold and the content is long-lived, the CDN Monitoring System instructs the SFC Control Plane Application to dynamically switch the any new flows from the existing service chain A to another service chain B which includes a CDN cache for caching the content in addition to all of the service functions in service chain A in the same order. This is done by installing a rule for the flows corresponding to the content in the SFC edge switch/router.
3. If the number of users viewing the same content falls below another pre-programmed down-threshold and the content is long-lived, the monitoring server instructs the SFC Control Plane Application to dynamically switch any new flows from the existing service chain B (the one that was used in the previous event) to service chain A (again, which includes all of the service functions in service chain B other than the CDN cache). This is done by removing the previously installed rule for the flows corresponding to the content from the SFC edge switch/router.

Note that the CDN use case applies only to new flows; existing flows follow the service chain that they were originally assigned to. Additionally, this use case assumes that the caching is transparent wherein the user does not address the cache explicitly. In other words, the decision of whether or not to retrieve content from the cache is not based on DNS, rather it is accomplished using SFC. The mechanisms described here will apply to encrypted traffic as long the encryption is at the application layer.

[4.](#) IPsec Management in Mobile Environments

Existing security procedures for flow protection in LTE are based on the use of IPsec tunnels between the radio base stations (eNodeBs) and some central node in the core, where a security gateway (SecGW) is deployed. The eNodeB device located on the cell site initiates the IPsec tunnel through the backhaul network to the SecGW, where the tunnel is terminated and the traffic is forwarded towards its final destination. IPsec ESP is the method that LTE standards use for achieving the required levels of security [[TS33.401](#)].

To avoid traffic bottlenecks and in order to guarantee a high level of service availability, a recommended practice is the concurrent use of several SecGW devices. The one that is to be used for a given traffic flow may be determined by several criteria such as the origin of the traffic (user traffic vs network control), flows with well-known characteristics, e.g. security properties (HTTPS, secure VPNs), etc. In this way, more critical traffic can be prioritized,

and different levels of security can be applied depending of payload characteristics.

Such an optimization could be applied as well to long-lived flows in a dynamic way, relaxing security procedures for non-sensitive ones, e.g. it may not be necessary to secure a well-known video stream that is openly available, applying differentiated policies to avoid congestion, or even hardening the security procedures according to the user's data profile.

[4.1.](#) Event Sequence

1. A monitoring element analyzes the new flows arriving at the default SecGW device used by a given eNodeB device according to

criteria such as:

- . Security payload protection;
 - . Application and transport protocol(s) in use including encryption schemes;
 - . Relevant parameters in those protocols (URL, content-transfer declarations, etc.).
2. If the monitoring element identifies a long-lived flow that matches its differentiating criteria, it signals the flow to the SFC Control Plane Application.
 3. The SFC Control Plane Application assigns the flow to a different service function chain that makes the eNodeB device use a different SecGW device.
 4. Once the flow becomes inactive, it is aged out by the eNodeB device and signaled as such to the SFC Control Plane Application.
 5. The SFC Control Plane Application removes the dynamic service chain association that was created for the flow.

[5. Operational Considerations](#)

Any modification to the SFC path (due to insertion or removal of a service function) could result in temporary mis-ordering in the delivery of packets.

[6. IANA Considerations](#)

None.

[7. Security Considerations](#)

This draft specifies a use case for SFC and does not introduce any new security requirements beyond those already under consideration for SFC.

[8. Acknowledgements](#)

The authors would like to thank Qin Wu, Myo Zarny, Reinaldo Penno,

and Tiru Reddy for their comments on the document.

[9. References](#)

[9.1. Normative References](#)

[9.2. Informative References](#)

[OPSAWG-large-flow] Krishnan, R. et al., "Mechanisms for Optimal LAG/ECMP Component Link Utilization in Networks," February 2014.

[I2RS-large-flow] Krishnan, R. et al., "I2RS Large Flow Use Case," November 2013.

[CDNI-long-tail] Krishnan, R. et al., "Best practices and Requirements for delivering Long Tail personalized content delivery over CDN Interconnections," work in progress, May 2013.

[CDN-overview] Dilley, J. et al., "Globally distributed content delivery," IEEE Internet Computing, September-October 2002.

[RFC 7011] Claise, B., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information," September 2013.

[TS33.401] 3GPP Technical Specification 33.401, "Security Architecture," December 2013.

[RFC 2660] Eric Rescorla et al., "The Secure HyperText Transfer Protocol," August 1999.

Authors' Addresses

Ram Krishnan
Brocade Communications
ramk@brocade.com

Krishnan et al.

Expires August 2015

[Page 9]

Internet-Draft

SFC Long Lived Flow Use Case

February 2015

Anoop Ghanwani
Dell
anoop@alumni.duke.edu

Joel Halpern
Ericsson
joel.halpern@ericsson.com

Sriganesh Kini
Ericsson
Sriganesh.kini@ericsson.com

Diego Lopez
Telefonica I+D
Don Ramon de la Cruz, 82 Street
Madrid, 28006, Spain
+34 913 129 041
diego.r.lopez@telefonica.com