SFC
Internet-Draft
Intended status: Standards Track
Expires: December 23, 2018

B. Sarikaya
Denpel Informatique
M. Boucadair
Orange
D. von Hugo
Deutsche Telekom
June 21, 2018

Service Function Chaining: Subscriber and Service Identification Use Cases and Variable-Length NSH Context Headers
draft-sarikaya-sfc-hostid-serviceheader-07

Abstract

This document discusses how to inform Service Functions about service- and subscriber-related information for the sake of policy enforcement and appropriate SFC-inferred forwarding. Once the information is consumed by SFC-aware elements of an SFC-enabled domain, it is stripped from packets when they leave the SFC-enabled domain. Thus privacy-sensitive information is not leaked outside the domain.

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on December 23, 2018.

Copyright Notice

Table of Contents

1. Introduction

This document discusses how to inform Service Functions about service- and subscriber-related information when required for the sake of policy enforcement. Indeed, subscriber-related information may be required to enforce subscriber-specific, SFC-based traffic forwarding policies, since the information carried in packets may not be sufficient.

The enforcement of SFC-based differentiated traffic forwarding policies may also be inferred by QoS considerations. QoS information may serve as an input to classification of SFP for path computation and establishment.

The dynamic structuring of service function chains and their subsequent enforcement may be conditioned by QoS requirements that will affect SF instance identification, location and sequencing.

We refer here to the definition of a logical network slice as a sub-network being isolated from other sub-networks using the same physical infrastructure. Each of these slices are constructed to provide service specific QoS requirements (such as low latency, high availability, or high reliability) efficiently.

SFs and SF Forwarders (SFFs) involved in an SFC have to contribute to the respective QoS requirements characterized by low transmission delay between each other, by exposing a high availability of resources to process function tasks, or by redundancy provided by stand-by machines for seamless execution continuation in case of failures. These requirements may be satisfied by means of control protocols, but in some contexts, carrying QoS-related information in packets may improve the overall SFC operation instead of relying upon the potential complexity of SFC control plane features.

This document adheres to the architecture defined in [RFC7665]. This document assumes the reader is familiar with [RFC7665] and [I-D.ietf-sfc-hierarchical].

Subscriber-related information may be required to implement services such as, but not limited to, traffic policy control, parental control, traffic offload. Such features are often provided by operators as part of their service portfolio.

Another example is the applicability of service chaining in the context of mobile networks (typically, in the 3GPP defined (S)Gi Interface) [I-D.ietf-sfc-use-case-mobility]. Because of the widespread use of private addressing in those networks, if advanced SFs to be invoked are located after a NAT device (that can reside in the Packet Data Network (PDN) Gateway (PGW) or in a distinct operator-specific node), the identification based on the internal IP address is not anymore possible once the NAT has been crossed. As such, means to allow passing the internal information may optimise packet traversal within an SFC-enabled mobile network domain. Furthermore, some SFs that are not enabled on the PGW may require a subscriber identifier e.g., International Mobile Subscriber Identity (IMSI), to properly operate. Other use cases that suffer from identification problems further are discussed in [RFC7620].

Subscriber-specific information can be useful for optimized SFC design and SF placement/invocation, let alone slice/VPN design and operation.

To ensure a service specific quality and performance per use case logically separated network slices will be deployed. Each one is flagged by a corresponding service-related information in terms of a service identifier. Examples are ’tactile internet’, ’eHealth’ or

'industry control' requiring, e.g. granted low latency or extreme high reliability.

This document does not make any assumption about the structure of service identifiers; each such service-related information is treated as an opaque value by the SFC operations and protocols. The semantics and validation of these identifiers are up to the control plane used for SFC. Expectations to SFC control plane protocols are laid down in [I-D.ietf-sfc-hierarchical].

Subscriber- and service-related information is stripped from packets exiting an SFC-enabled domain for the sake of privacy protection in particular. See Section 8 for more discussion on privacy.

The use cases discussed in this document assume the NSH is used exclusively within a single administrative domain.

2. Conventions and Terminology

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].

The reader should be familiar with the terms defined in [RFC7665].

3. Problem Space and Sample Use Cases

Enforcing Policies based on an internal IP address:

> Because of the address sharing, implicit CPE/UE identification that relies on the source IP address cannot be implemented within the administrative domain because the same global IPv4 address is shared by various connected devices (CPE for the fixed case or UE for the mobile case). In the meantime, policies are something provisioned based on the internal IP address assigned to those devices. Means to pass the internal IP address beyond an address sharing device for the sake of per-subscriber policy enforcement is needed in some SFC deployments.
>
> Also, identifiers like a MAC address, or an IMSI may be required to optimize the corresponding SFC operation.

Enforcing Policies based on a subscriber identifier:

> In case some deployments may require per-subscriber policies, carrying subscriber ID information may be required for the sake of proper SFC operation..

Enforcing Policies based on a service specific identifier:

SFCs can be structured according to QoS/QoE requirements that may be shared by different services. In that case, service identification information is required to be accessible across the SFC for the sake of proper SFC operation.

Below we present some use cases where problems related to enforcing policies based on subscriber identifiers and those based on service and/or slice identifiers cannot be addressed by service function chaining. It is important to note that subscriber identification issues raised by address sharing environments are not specific to service function chaining.

3.1. Parental Control Use Case

Parental control service function searches each packet for certain content. Parental control function should have permanent access to corresponding specific information (URL and source IP address), e.g. in a cache, so that all packets of the corresponding flow(s) can be filtered [WT317].

Parental control function receives next packet from the recorded URL. Enforcing the parental control policies may depend on the internal IP address, i.e., the address of the subscriber’s host that is being subject to the parental control. Parental control function must be able to identify incoming traffic to be filtered, e.g., specific URL information. All other traffic is not subject to parental control filtering. Parental control function filters all traffic coming from the indicated URL only for the specific subscriber’s hosts identified by the service logic.

For the virtual CPE case, the access node will receive privately-addressed packets. Because private IPv4 addresses are likely to overlap between several subscribers, the internal private IPv4 address will need to be copied into a dedicated header in the NSH packet so that SFs responsible for parental control can process the packets appropriately. Furthermore, the subscriber identifier may also be required for authorization purposes.

3.2. Traffic Offload Use Case

A traffic offload service function is invoked for each flow/service originated from a mobile terminal and this SF decides whether traffic should be offloaded to the broadband network or sent back to the mobile network. In this use case, policy enforcement is based on the subscriber identifier. The broadband network must obtain the subscription profile from the mobile network and decide if the

traffic coming from this subscriber needs to be offloaded or not. If offloading is needed, this usually means that the subscriber identifier needs to be known by SFFs.

3.3. Mobile Network Use Cases

Many SFs can be executed in different combinations in a mobile network [I-D.ietf-sfc-use-case-mobility]. In particular, placement of NAT function (if used) plays an important role.

If a NAT function is collocated with P-GW as in [TR23.975] or right after the P-GW (i.e. between P-GW and (S)Gi-LAN) then all service functions located upstream can only see the translated IPv4 address as the source address from all User Equipments (UEs). Internal IPv4 address-related part of their policy set won't be able to execute their service logic. As a consequence, means to inform the various SFs of a given chain about the IPv4 address assigned to the UE and which will be translated into a global IPv4 address may be needed.

Note that the same problem occurs in case IPv6 is being used by UEs, whenever such UEs communicate with an IPv4-only web site. In that case, a NAT64 function is deployed at the P-GW. So in the case of chaining NAT64 SF needs to be invoked as part of a given chain, the IPv6 address used by the UE may be required for the service function chain to work properly.

[I-D.ietf-sfc-use-case-mobility] identifies the following information:

- Charging ID
- Subscriber ID
- GGSN or PGW IP address
- Serving Gateway Support Node (SGSN) or SGW IP address
- International Mobile Equipment Identity (IMEI)
- International Mobile Subscriber Identity (IMSI)
- Mobile Subscriber ISDN Number (MSISDN)
- UE IP address

Several other use cases where support of traffic classification with respect to service chain selection to achieve efficient and flexible

mobile service steering are described in [TR22.808]. A set of potential solutions are proposed and discussed in [TR23.718].

3.4. Extreme Low Latency Service Use Cases

Extreme or ultra-low latency requirements may be addressed by specific architectural and protocol characteristics to allow for rapid execution and low transmission delay of packets. Candidate services for such requirements include e-health or vehicular applications. This can be granted by forwarding all packets via the shortest paths only and/or via the service function instances with the lowest processing delay, possibly as a function of the location of the user.

The corresponding service function chain should be configured based on the service demanding for the performance, but policies are also tightly related to the subscriber, i.e. whether being entitled to request the specific service.

3.5. High Reliability Applications Use Cases

Another set of use cases that require very (or ultra-) high reliability of services assume committed QoS parameter values and its possibility to act upon an expected change of the network fulfillment of the QoS targets [TR22.862]. That means: the QoS fulfillment is controlled such that in case of expected or predicted deviation a countermeasure by the network is invoked, e.g. either resources for that session are increased or a backup path is assigned in case no improvement is possible at least the application is informed on the current performance to react upon. This can be granted by forwarding all packets via the most reliable and secure paths only.

4. Subscriber Identification NSH Variable-Length Context Header

Subscriber Identifier is defined as an optional variable-length NSH context header. Its structure is shown in Figure 1.

The subscriber identifier is used to convey an identifier already assigned by the service provider to uniquely identify a subscriber or an information that is required to enforce per-subscriber policies, the structure of the identifier being deployment-specific. Typically, this header may convey the IMSI, an opaque subscriber Identifier, an IP address, etc.

The classifier and SFC-aware Service Functions MAY be instructed via a control interface to inject or strip a subscriber identifier context header. Also, the data to be injected in such header SHOULD be configured to nodes authorized to inject such headers. Failures

to inject such headers SHOULD be logged locally while a notification alarm MAY be sent to a Control Element. The details of sending notification alarms (i.e. the parameters affecting the transmission of the notification alarms depend on the information in the context header such as frequency, thresholds, and content in the alarm: full header, header ID, timestamp etc.) SHOULD be configurable by the control plane.

This document adheres to the recommendations in [RFC8300] for handling the context headers at both ingress and egress SFC boundary nodes. That is, to strip such context headers.

SFC-aware SFs and proxies MAY be instructed to strip a subscriber identifier from the packet or to pass the data to the next SF in the chain after processing the content of the headers. If no instruction is provided, the default behavior is to maintain such context headers so that the information can be passed to next SFC-aware hops.

SFC-aware functions MAY be instructed via the control plane about the validation checks to run on the content of these context headers (e.g., accept only some lengths, accept some subtypes) and the behavior to adopt. For example, SFC-aware nodes may be instructed to ignore the context header, to remove the context header from the packet, etc. Nevertheless, this specification does not require nor preclude such additional validation checks. These validation checks are deployment-specific. If validation checks fail on a context header, an SFC-aware node ignores that context header. The event SHOULD be logged locally while a notification alarm MAY be sent to a control element if the SFC-aware node is instructed to do so.

Multiple subscriber Identifier context TLVs MAY be present in the NSH each carrying a distinct sub-type.

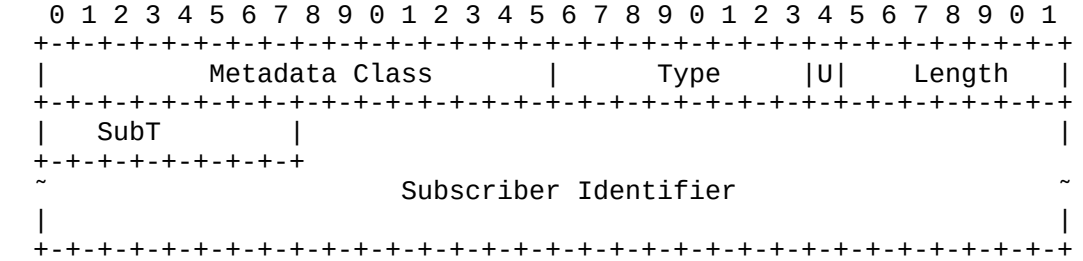

Figure 1: Subscriber Identifier Variable-Length Context Header

The description of the fields is as follows:

o Metadata Class: MUST be set to 0x0 [RFC8300].

- Type: TBD1 (See Section 6)
- SubT field of 8 bits indicates the sub-type of the information conveyed in the "Subscriber Identifier" field. The following values are defined:
  - 0x00: Opaque value
  - 0x01: Charging ID. The structure of this ID is deployment-specific.
  - 0x02: Subscriber ID. The structure of this ID is deployment-specific.
  - 0x03: GGSN or PGW IP address/prefix
  - 0x04: Serving Gateway Support Node (SGSN) or SGW IP address/ prefix
  - 0x05: International Mobile Equipment Identity (IMEI)
  - 0x06: International Mobile Subscriber Identity (IMSI)
  - 0x07: Mobile Subscriber ISDN Number (MSISDN)
  - 0x08: UE IP address
- Subscriber Identifier: Carries an opaque subscriber identifier or an identifier that corresponds to the sub-type.

5. Slice and Service Identification NSH Variable-Length Context Headers

Dedicated service- and slice-specific performance identifiers are defined to differentiate between services requiring specific treatment to exhibit a performance characterized by, e.g., ultra-low latency (ULL) or ultra-high reliability (UHR). These parameters are related to slice and service identifiers, among others. They are contained in the service Identifier. The service Identifier thus allows for the enforcement of a per service policy such as a service classification function to only consider specific Service Function instances during service function path establishment. Details of this process are implementation- specific. For illustration purposes, the classifier may retrieve the details of usable service functions based upon the corresponding service or slice ID. Typical criteria for instantiating specific service functions include location, performance or proximity considerations. For UHR services, the stand-by operation of back-up capacity or the deployment of multiple service function instances may be requested.

In other words, the classifier uses this kind of information to decide about the set of SFFs to invoke to honor the latency or reliability requirement (e.g., compute an Rendered Service Path, RSP, or insert a pointer to be shared with involved SFFs).

Slice and Service Identifiers are defined as optional variable length context headers. Their structure is shown in Figure 2 and Figure 3, respectively.

Service/Slice Identifier context header MAY convey a user or service provider defined unique identity which can be described by an opaque value.

The service requirements in terms of, e.g., maximum latency or minimum outage probability are specified by service providers and are out of the scope of this document.

Only one Slice Identifier context header (as described in Section 1) MUST be present in the NSH.

Multiple Service Identifier context headers MAY be present in the NSH; each carrying a distinct sub-type.

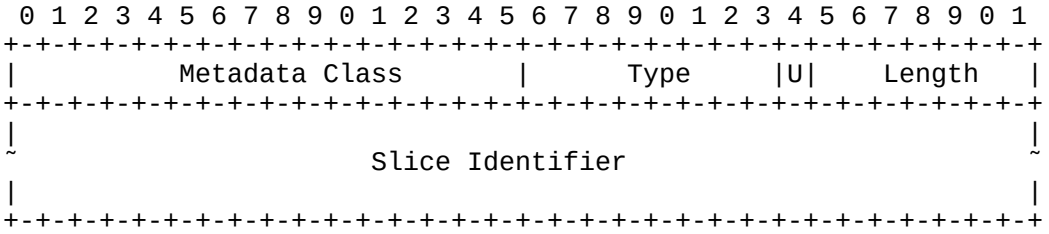

Figure 2: Slice Identifier Variable-Length Context Header

The description of the fields is as follows:

- Metadata Class: MUST be set to 0x0 [RFC8300].
- Type: TBD2 (See Section 6)
- Slice Identifier: The structure of the identifier is deployment-specific. This field carries an identifier that uniquely identifies a slice within a network, e.g. it could be an opaque value with an arbitrary number of characters.

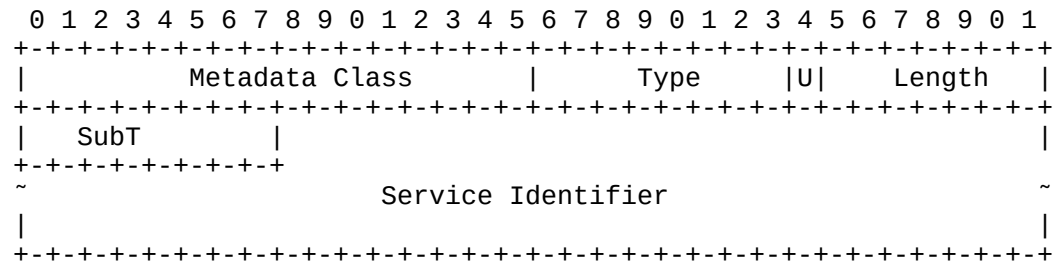

Figure 3: Service Identifier Variable-Length Context Header

The description of the fields is as follows:

- Metadata Class: MUST be set to 0x0 [RFC8300].

- Type: TBD3 (See Section 6)

- SubT: 8-bit field that carries the sub-type of the information conveyed in the "Service Identifier" field. The following values are defined:

  * 0x00: Opaque value

  * 0x01: Ultra-low latency ID. The structure of this ID is service deployment-specific.

  * 0x02: Ultra-high reliability ID. The structure of this ID is service deployment-specific.

  * 0x03: Slice Identifier. The structure of this ID is service deployment-specific.

  * 0x04 - 0x08: reserved

- Service Identifier: Represents a specific service performance characteristic reflected in the SubT field, but also denotes a default basic (best effort) service without specifically defined requirements. It MAY also be an opaque value which semantic is defined by the operator.

6. IANA Considerations

This document requests IANA to assign the followiing types from the "NSH IETF- Assigned Optional Variable-Length Metadata Types" (0x0000 IETF Base NSH MD Class) registry available at: https://www.iana.org/assignments/nsh/nsh.xhtml#optional-variable-length-metadata-types.

C1	C2	
TBD1	Subscriber Identifier	[ThisDocument]
TBD2	Slice Identifier	[ThisDocument]
TBD3	Service Identifier	[ThisDocument]

7. Security Considerations

Data plane SFC-related security considerations are discussed in [RFC7665] and [RFC8300].

A misbehaving node can inject subscriber Identifiers to disturb the service offered to some subscribers. Also, a misbehaving node can inject subscriber identifiers as an attempt to be granted access to some services. To prevent such misbehavior, only trusted nodes MUST be able to inject such context headers. Nodes that are involved in a SFC-enabled domain are assumed to be trusted ([RFC8300]). Means to check that only authorized nodes are solicited when a packet is crossing an SFC-enabled domain.

8. Privacy Considerations

The metadata defined in this document for subscriber identifiers may reveal private information about the subscriber. Some privacy-related considerations for Internet Protocols are discussed in [RFC6973] and [RFC6967]. In the light of these privacy considerations, it is important to state that the subscriber metadata MUST NOT be exposed outside the operator’s domain. This requirement is already supported by the NSH [RFC8300]. That is, NSH is stripped systematically at the egress of a service chain.

The information conveyed in subscriber identifiers is already known to an administrative entity managing an SFC-enabled domain. Some of that information is already conveyed in the original packets from a host (e.g., internal IP address) while other information is collected from various sources (e.g., GTP tunnel, line identifier, etc.). Conveying such sensitive information in packets may expose subscribers’ sensitive data to entities that are not allowed to receive such information. Misbehaving SFC egress nodes is a threat that may have negative impacts on privacy (e.g., some operational networks leak the MSISDN outside). Operators MUST ensure their SFC-enabled domain is appropriately conforming to the NSH specification so that any privacy-related information is not exposed outside the SFC-enabled domain.

Some use cases that rely upon the solution defined in this document may disclose some additional privacy-related information (e.g., a host identifier of a terminal within a customer premises for the parental control case). It is assumed that this information is provided upon approval from a subscriber [RFC8165]. For example, a customer may provide the information as part of its service management interface or as part of explicit subscription form.

9. Acknowledgements

Comments from Joel Halpern on a previous version and by Carlos Bernardos are appreciated. Contributions by Christian Jacquenet are thankfully acknowledged.

This work has been partially performed in the framework of the EU-funded H2020-ICT-2014-2 project 5G NORMA. Contributions of the project partners are gratefully acknowledged. The project consortium is not liable for any use that may be made of any of the information contained therein.

10. References

10.1. Normative References

[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>.

[RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <https://www.rfc-editor.org/info/rfc7665>.

[RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <https://www.rfc-editor.org/info/rfc8300>.

10.2. Informative References

[I-D.ietf-sfc-hierarchical] Dolson, D., Homma, S., Lopez, D., and M. Boucadair, "Hierarchical Service Function Chaining (hSFC)", draft-ietf-sfc-hierarchical-09 (work in progress), June 2018.

[I-D.ietf-sfc-use-case-mobility]
Haeffner, W., Napper, J., Stiemerling, M., Lopez, D., and J. Uttaro, "Service Function Chaining Use Cases in Mobile Networks", draft-ietf-sfc-use-case-mobility-08 (work in progress), May 2018.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <https://www.rfc-editor.org/info/rfc6146>.

[RFC6967] Boucadair, M., Touch, J., Levis, P., and R. Penno, "Analysis of Potential Solutions for Revealing a Host Identifier (HOST_ID) in Shared Address Deployments", RFC 6967, DOI 10.17487/RFC6967, June 2013, <https://www.rfc-editor.org/info/rfc6967>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <https://www.rfc-editor.org/info/rfc6973>.

[RFC7620] Boucadair, M., Ed., Chatras, B., Reddy, T., Williams, B., and B. Sarikaya, "Scenarios with Host Identification Complications", RFC 7620, DOI 10.17487/RFC7620, August 2015, <https://www.rfc-editor.org/info/rfc7620>.

[RFC8165] Hardie, T., "Design Considerations for Metadata Insertion", RFC 8165, DOI 10.17487/RFC8165, May 2017, <https://www.rfc-editor.org/info/rfc8165>.

[TR22.808]
"3GPP TR22.808, Technical Specification Group Services and System Aspects; Study on flexible mobile service steering", 2015.

[TR22.862]
"3GPP TR22.862, Feasibility Study on New Markets and Technology Enablers - Critical Communications; Stage 1 (Release 14)", 2015.

[TR23.718]
"3GPP TR23.718, Technical Specification Group Services and System Aspects; Architecture enhancement for flexible mobile service steering", 2015.

[TR23.975]
"3GPP TR23.975, IPv6 Migration Guidelines", June 2011.

[TS23.003]
"3GPP TS23.003, Technical Specification Group Core Network and Terminals; Numbering, addressing and identification", 2015.

[TS29.212]
"3GPP TS29.212, Policy and Charging Control (PCC) over Gx/ Sd reference point", December 2011.

[WT317] BBF, "Network Enhanced Residential Gateway", August 2015.

Authors' Addresses

Behcet Sarikaya
Denpel Informatique

Email: sarikaya@ieee.org

Mohamed Boucadair
Orange
Rennes 3500, France

Email: mohamed.boucadair@orange.com

Dirk von Hugo
Telekom Innovation Laboratories
Deutsche-Telekom-Allee 7
D-64295 Darmstadt
Germany

Email: Dirk.von-Hugo@telekom.de