ï»¿

Internet Engineering Task Force (IETF)
Request for Comments: 8979
Category: Standards Track
ISSN: 2070-1721

B. Sarikaya
D. von Hugo
Deutsche Telekom
M. Boucadair
Orange

February 2021

Subscriber and Performance Policy Identifier Context Headers in the Network Service Header (NSH)

Abstract

This document defines the Subscriber and Performance Policy Identifier Context Headers. These Variable-Length Context Headers can be carried in the Network Service Header (NSH) and are used to inform Service Functions (SFs) of subscriber- and performance-related information for the sake of policy enforcement and appropriate Service Function Chaining (SFC) operations. The structure of each Context Header and their use and processing by NSH-aware nodes are described.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in 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/rfc8979.

Copyright Notice

Copyright (c) 2021 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
2. Conventions and Terminology
3. Subscriber Identifier NSH Variable-Length Context Header
4. Performance Policy Identifier NSH Variable-Length Context Headers
5. MTU Considerations
6. IANA Considerations
7. Security Considerations
8. References
   8.1. Normative References
   8.2. Informative References

Acknowledgements
Authors' Addresses

1. Introduction

This document discusses how to inform Service Functions (SFs) [RFC7665] about subscriber and service policy information when required for the sake of policy enforcement within a single administrative domain. In particular, subscriber-related information may be required to enforce subscriber-specific SFC-based traffic policies. However, the information carried in packets may not be sufficient to unambiguously identify a subscriber. This document fills this void by specifying a new Network Service Header (NSH) [RFC8300] Context Header to convey and disseminate such information within the boundaries of a single administrative domain. As discussed in Section 3, the use of obfuscated and non-persistent identifiers is recommended.

Also, traffic steering by means of SFC may be driven, for example, by Quality of Service (QoS) considerations. Typically, QoS information may serve as an input for the computation, establishment, and selection of the Service Function Path (SFP). Furthermore, the dynamic structuring of Service Function Chains and their subsequent SFPs may be conditioned by QoS requirements that will affect the identification, location, and sequencing of SF instances. Hence, the need arises to provide downstream SFs with a performance policy identifier in order for them to appropriately meet the QoS requirements. This document also specifies a new NSH Context Header (Section 4) to convey such policy identifiers.

The context information defined in this document can be applicable in the context of mobile networks (particularly in the 3GPP-defined (S)Gi interface) [CASE-MOBILITY]. Typically, because of the widespread use of private IPv4 addresses in those networks, if the SFs to be invoked are located after a NAT function, the identification based on the internal IPv4 address is not possible once the NAT has been crossed. NAT functionality can reside in a distinct node. For a 4G 3GPP network, that node can be the Packet Data Network (PDN) Gateway (PGW) as specified in [TS23401]. For a 5G 3GPP network, it can be the User Plane Function (UPF) facing the external Data Network (DN) [TS23501]. As such, a mechanism to pass the internal information past the NAT boundary may optimize packet traversal within an SFC-enabled mobile network domain. Furthermore, some SFs that are not enabled on the PGW/UPF may require a subscriber identifier to properly operate (see, for example, those listed in [RFC8371]). It is outside the scope of this document to include a comprehensive list of deployments that may make use of the Context Headers defined in the document.

Since subscriber identifiers are distinct from those used to identify a performance policy and given that multiple policies may be associated with a single subscriber within a Service Function Chain, these identifiers are carried in distinct Context Headers rather than being multiplexed in one single Context Header. This approach avoids a requirement for additional internal structure in the Context Headers to specify whether an identifier refers to a subscriber or to a policy.

This document does not make any assumptions about the structure of subscriber or performance policy identifiers; each such identifier is treated as an opaque value. The semantics and validation of these identifiers are policies local to each SFC-enabled domain. This document focuses on the data plane behavior. Control plane considerations are out of the scope.

This document adheres to the SFC data plane architecture defined in [RFC7665]. This document assumes the reader is familiar with [RFC8300].

This document assumes the NSH is used exclusively within a single administrative domain. This document follows the recommendations in [RFC8300] for handling the Context Headers at both ingress and egress SFC boundary nodes (i.e., to strip the entire NSH, including Context

Headers). Revealing any subscriber-related information to parties outside the SFC-enabled domain is avoided by design. Accordingly, the scope for privacy breaches and user tracking is limited to within the SFC-enabled domain where the NSH is used. It is assumed that appropriate mechanisms to monitor and audit an SFC-enabled domain to detect misbehavior and to deter misuse are in place.

MTU considerations are discussed in Section 5.

2. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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

"SFC Control Element" refers to a logical entity that instructs one or more SFC data plane functional elements on how to process packets within an SFC-enabled domain.

3. Subscriber Identifier NSH Variable-Length Context Header

Subscriber Identifier is defined as an optional Variable-Length NSH Context Header. Its structure is shown in Figure 1.

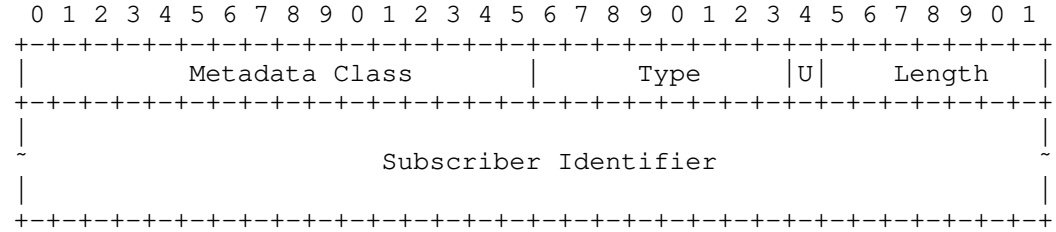

Figure 1: Subscriber Identifier Variable-Length Context Header

The fields are described as follows:

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

Type: 0x00 (see Section 6).

U bit: Unassigned bit (see Section 2.5.1 of [RFC8300]).

Length: Indicates the length of the Subscriber Identifier, in bytes (see Section 2.5.1 of [RFC8300]).

Subscriber Identifier: Carries an opaque local identifier that is assigned to a subscriber by a network operator.

While this document does not specify an internal structure for these identifiers, it also does not provide any cryptographic protection for them; any internal structure to the identifier values chosen will thus be visible on the wire if no secure transport encapsulation is used. Accordingly, in alignment with Section 8.2.2 of [RFC8300], identifier values SHOULD be obfuscated.

The Subscriber Identifier Context Header is used by SFs to enforce per-subscriber policies (e.g., resource quota, customized filtering profile, accounting). To that aim, network operators may rely on identifiers that are generated from those used in legacy deployments (e.g., Section 3.3 of [CASE-MOBILITY]). Alternatively, network operators may use identifiers that are associated with customized policy profiles that are preconfigured on SFs using an out-of-band mechanism. Such a mechanism can be used to rotate the identifiers, thus allowing for better unlinkability (Section 3.2 of [RFC6973]). Such alternative methods may be suboptimal (e.g., scalability issues

induced by maintaining and processing finer granular profiles) or inadequate for providing some per-subscriber policies. The assessment of whether a method for defining a subscriber identifier provides the required functionality and whether it is compatible with the capabilities of the SFs at the intended performance level is deployment specific.

The classifier and NSH-aware SFs MAY inject a Subscriber Identifier Context Header as a function of a local policy. This local policy should indicate the SFP(s) for which the Subscriber Identifier Context Header will be added. In order to prevent interoperability issues, the type and format of the identifiers to be injected in a Subscriber Identifier Context Header should be configured to nodes authorized to inject and consume such headers. For example, a node can be instructed to insert such data following a type/set scheme (e.g., node X should inject subscriber ID type Y). Other schemes may be envisaged.

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) might depend on the nature of the information in the Context Header. Parameters for sending alarms, such as frequency, thresholds, and content of the alarm, should be configurable.

The default behavior of intermediary NSH-aware nodes is to preserve Subscriber Identifier Context Headers (i.e., the information can be passed to next-hop NSH-aware nodes), but local policy may require an intermediary NSH-aware node to strip a Subscriber Identifier Context Header after processing it.

NSH-aware SFs MUST ignore Context Headers carrying unknown subscriber identifiers.

Local policies at NSH-aware SFs may require running additional validation checks on the content of these Context Headers (e.g., accepting only some lengths or types). These policies may also indicate the behavior to be followed by an NSH-aware SF if the validation checks fail (e.g., removing the Context Header from the packet). These additional validation checks are deployment specific. If validation checks fail on a Subscriber Identifier Context Header, an NSH-aware SF MUST ignore that Context Header. The event should be logged locally, while a notification alarm may be sent to a Control Element if the NSH-aware SF is instructed to do so. For example, an SF will discard Subscriber Identifier Context Headers conveying identifiers in all formats that are different from the one the SF is instructed to expect.

Multiple Subscriber Identifier Context Headers MAY be present in the NSH, each carrying a distinct opaque value but all pointing to the same subscriber. This may be required, e.g., by policy enforcement mechanisms in a mobile network where some SFs rely on IP addresses as subscriber identifiers, while others use non-IP-specific identifiers such as those listed in [RFC8371] and Section 3.3.2 of [CASE-MOBILITY]. When multiple Subscriber Identifier Context Headers are present and an SF is instructed to strip the Subscriber Identifier Context Header, that SF MUST remove all Subscriber Identifier Context Headers.

4. Performance Policy Identifier NSH Variable-Length Context Headers

Dedicated service-specific performance identifiers are defined to differentiate between services that require specific treatment in order to exhibit a performance characterized by, e.g., ultra-low latency (ULL) or ultra-high reliability (UHR). Other policies can be considered when instantiating a Service Function Chain within an SFC-enabled domain. They are conveyed in the Performance Policy Identifier Context Header.

The Performance Policy Identifier Context Header is inserted in an

NSH packet so that downstream NSH-aware nodes can make use of the performance information for proper selection of suitably distributed SFC paths, SF instances, or applicable policy at SFs. Note that the use of the performance policy identifier is not helpful if the path computation is centralized and a strict SFP is presented as local policy to SF Forwarders (SFFs).

The Performance Policy Identifier Context Header allows for the distributed enforcement of a per-service policy such as requiring an SFP to only include specific SF instances (e.g., SFs located within the same Data Center (DC) or those that are exposing the shortest delay from an SFF). Details of this process are implementation specific. For illustration purposes, an SFF may retrieve the details of usable SFs based upon the corresponding performance policy identifier. Typical criteria for instantiating specific SFs include location, performance, or proximity considerations. For the particular case of UHR services, the standby operation of backup capacity or the presence of SFs deployed in multiple instances may be requested.

In an environment characterized by frequent changes of link and path behavior (for example, due to variable load or availability caused by propagation conditions on a wireless link), the SFP may have to be adapted dynamically by on-the-move SFC path and SF instance selection.

Performance Policy Identifier is defined as an optional Variable-Length Context Header. Its structure is shown in Figure 2.

The default behavior of intermediary NSH-aware nodes is to preserve such Context Headers (i.e., the information can be passed to next-hop NSH-aware nodes), but local policy may require an intermediary NSH-aware node to strip one Context Header after processing it.

Multiple Performance Policy Identifier Context Headers MAY be present in the NSH, each carrying an opaque value for a distinct policy that needs to be enforced for a flow. Supplying conflicting policies may complicate the SFP computation and SF instance location. Corresponding rules to detect conflicting policies may be provided as a local policy to the NSH-aware nodes. When such conflict is detected by an NSH-aware node, the default behavior of the node is to discard the packet and send a notification alarm to a Control Element.

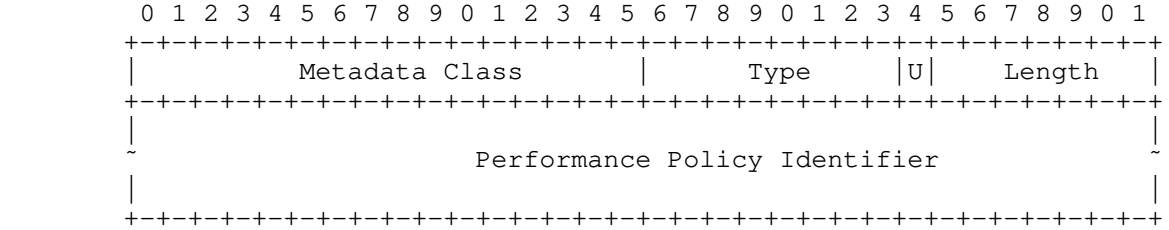

Figure 2: Performance Policy Identifier Variable-Length Context Header

The fields are described as follows:

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

Type: 0x01 (see Section 6).

U bit: Unassigned bit (see Section 2.5.1 of [RFC8300]).

Length: Indicates the length of the Performance Policy Identifier, in bytes (see Section 2.5.1 of [RFC8300]).

Performance Policy Identifier: Represents an opaque value pointing to a specific performance policy to be enforced. The structure and semantics of this field are deployment specific.

5. MTU Considerations

As discussed in Section 5.6 of [RFC7665], the SFC architecture prescribes that additional information be added to packets to:

* Identify SFPs. This is typically the NSH Base Header (Section 2.2 of [RFC8300]) and Service Path Header (Section 2.3 of [RFC8300]).

* Carry metadata such those defined in Sections 3 and 4.

* Steer the traffic along the SFPs: This is realized by means of transport encapsulation.

This added information increases the size of the packet to be carried along an SFP.

Aligned with Section 5 of [RFC8300], it is RECOMMENDED for network operators to increase the underlying MTU so that NSH traffic is forwarded within an SFC-enabled domain without fragmentation. The available underlying MTU should be taken into account by network operators when providing SFs with the required Context Headers to be injected per SFP and the size of the data to be carried in these Context Headers.

If the underlying MTU cannot be increased to accommodate the NSH overhead, network operators may rely upon a transport encapsulation protocol with the required fragmentation handling. The impact of activating such feature on SFFs should be carefully assessed by network operators (Section 5.6 of [RFC7665]).

When dealing with MTU issues, network operators should consider the limitations of various transport encapsulations such as those discussed in [INTAREA-TUNNELS].

6. IANA Considerations

IANA has assigned the following types from the "NSH IETF-Assigned Optional Variable-Length Metadata Types" subregistry (0x0000 IETF Base NSH MD Class) available at: <https://www.iana.org/assignments/nsh>.

Value	Description	Reference
0x00	Subscriber Identifier	[RFC8979]
0x01	Performance Policy Identifier	[RFC8979]

Table 1: NSH IETF-Assigned Optional Variable-Length Metadata Types Additions

7. Security Considerations

Data plane SFC-related security considerations, including privacy, are discussed in Section 6 of [RFC7665] and Section 8 of [RFC8300]. In particular, Section 8.2.2 of [RFC8300] states that attached metadata (i.e., Context Headers) should be limited to that necessary for correct operation of the SFP. Section 8.2.2 of [RFC8300] indicates that metadata considerations that operators can take into account when using NSH are discussed in [RFC8165].

As specified in [RFC8300], means to prevent leaking privacy-related information outside an SFC-enabled domain are natively supported by the NSH given that the last SFF of an SFP will systematically remove the NSH (and therefore the identifiers defined in this specification) before forwarding a packet exiting the SFP.

Nodes that are involved in an SFC-enabled domain are assumed to be trusted (Section 1.1 of [RFC8300]). Discussion of means to check that only authorized nodes are traversed when a packet is crossing an

SFC-enabled domain is out of scope of this document.

Both Subscriber Identifier and Performance Policy Identifier Context Headers carry opaque data. In particular, the Subscriber Identifier Context Header is locally assigned by a network provider and can be generated from some of the information that is already conveyed in the original packets from a host (e.g., internal IP address) or other information that is collected from various sources within an SFC-enabled domain (e.g., line identifier). The structure of the identifiers conveyed in these Context Headers is communicated only to entitled NSH-aware nodes. Nevertheless, some structures may be easily inferred from the headers if trivial structures are used (e.g., IP addresses). As persistent identifiers facilitate tracking over time, the use of indirect and non-persistent identification is thus RECOMMENDED.

Moreover, the presence of multiple Subscriber Identifier Context Headers in the same NSH allows a misbehaving node from within the SFC-enabled domain to bind these identifiers to the same subscriber. This can be used to track that subscriber more effectively. The use of non-persistent (e.g., regularly randomized) identifiers as well as the removal of the Subscriber Identifier Context Headers from the NSH by the last SF making use of such headers soften this issue (see "data minimization" discussed in Section 3 of [RFC8165]). Such behavior is especially strongly recommended in case no encryption is enabled.

A misbehaving node from within the SFC-enabled domain may alter the content of Subscriber Identifier and Performance Policy Identifier Context Headers, which may lead to service disruption. Such an attack is not unique to the Context Headers defined in this document; measures discussed in Section 8 of [RFC8300] are to be followed. A mechanism for NSH integrity is specified in [NSH-INTEGRITY].

If no secure transport encapsulation is enabled, the use of trivial subscriber identifier structures, together with the presence of specific SFs in a Service Function Chain, may reveal sensitive information to every on-path device. Also, operational staff in teams managing these devices could gain access to such user privacy-affecting data. Such disclosure can be a violation of legal requirements because such information should be available to very few network operator personnel. Furthermore, access to subscriber data usually requires specific access privilege levels. To maintain that protection, an SF keeping operational logs should not log the content of Subscriber and Performance Policy Identifier Context Headers unless the SF actually uses the content of these headers for its operation. As discussed in Section 8.2.2 of [RFC8300], subscriber-identifying information should be obfuscated, and, if an operator deems cryptographic integrity protection is needed, security features in the transport encapsulation protocol (such as IPsec) must be used. A mechanism for encrypting sensitive NSH data is specified in [NSH-INTEGRITY]. This mechanism can be considered by network operators when enhanced SF-to-SF security protection of NSH metadata is required (e.g., to protect against compromised SFFs).

Some events are logged locally with notification alerts sent by NSH-aware nodes to a Control Element. These events SHOULD be rate limited.

8. References

8.1. 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>.

[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>.

[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>.

8.2. Informative References

[CASE-MOBILITY] Haeffner, W., Napper, J., Stiemerling, M., Lopez, D. R., and J. Uttaro, "Service Function Chaining Use Cases in Mobile Networks", Work in Progress, Internet-Draft, draft-ietf-sfc-use-case-mobility-09, 1 January 2019, <https://tools.ietf.org/html/draft-ietf-sfc-use-case-mobility-09>.

[INTAREA-TUNNELS] Touch, J. and M. Townsley, "IP Tunnels in the Internet Architecture", Work in Progress, Internet-Draft, draft-ietf-intarea-tunnels-10, 12 September 2019, <https://tools.ietf.org/html/draft-ietf-intarea-tunnels-10>.

[NSH-INTEGRITY] Boucadair, M., Reddy.K, T., and D. Wing, "Integrity Protection for the Network Service Header (NSH) and Encryption of Sensitive Context Headers", Work in Progress, Internet-Draft, draft-ietf-sfc-nsh-integrity-03, 22 January 2021, <https://tools.ietf.org/html/draft-ietf-sfc-nsh-integrity-03>.

[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>.

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

[RFC8371] Perkins, C. and V. Devarapalli, "Mobile Node Identifier Types for MIPv6", RFC 8371, DOI 10.17487/RFC8371, July 2018, <https://www.rfc-editor.org/info/rfc8371>.

[TS23401] 3GPP, "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access, Release 16", Version 16.5.0, TS 23.401, December 2019.

[TS23501] 3GPP, "System architecture for the 5G System (5GS), Release 16", Version 16.3.0, TS 23.501, December 2019.

Acknowledgements

Comments from Joel Halpern on a previous draft version and from Carlos Bernardos are appreciated.

Contributions and review by Christian Jacquenet, Danny Lachos, Debashish Purkayastha, Christian Esteve Rothenberg, Kyle Larose, Donald Eastlake, Qin Wu, Shunsuke Homma, and Greg Mirsky are thankfully acknowledged.

Many thanks to Robert Sparks for the secdir review.

Thanks to Barry Leiba, Erik Kline, Ã\211ric Vyncke, Robert Wilton, and

Magnus Westerlund for the IESG review.

Special thanks to Benjamin Kaduk for the careful review and suggestions that enhanced this specification.

Authors' Addresses

Behcet Sarikaya

Email: sarikaya@ieee.org

Dirk von Hugo
Deutsche Telekom
Deutsche-Telekom-Allee 9
D-64295 Darmstadt
Germany

Email: Dirk.von-Hugo@telekom.de

Mohamed Boucadair
Orange
3500 Rennes
France

Email: mohamed.boucadair@orange.com