

SFC WG
Internet-Draft
Intended status: Informational
Expires: May 4, 2017

CJ. Bernardos
UC3M
A. Rahman
A. Mourad
InterDigital
October 31, 2016

Service Function Chaining Use Cases in Fog RAN
draft-bernardos-sfc-fog-ran-00

Abstract

Fog Radio Access Networks (Fog RAN) refers to the part of the RAN that is virtualized at the very edge of the network, even at the end-user device. Fog RAN support is considered critical for the 5G mobile network architectures currently being developed in various research, standardization and industry forums. Since fog RAN builds on top of virtualization and can involve several virtual functions running on different virtualized resources, Service function chaining (SFC) support for the fog RAN will be critical. This document describes the overall fog RAN approach and also gives some use cases. Finally it proposes some requirements to be considered in the development of the SFC architecture and related protocols.

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 <http://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 May 4, 2017.

Copyright Notice

Copyright (c) 2016 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	2
2. Terminology	3
3. Fog RAN Overview	4
4. Applicability of SFC to Fog RAN	8
5. Fog RAN requirements	11
6. IANA Considerations	12
7. Security Considerations	12
8. Acknowledgments	12
9. References	12
9.1. Normative References	12
9.2. Informative References	12
Appendix A. 4G (LTE)	13
Appendix B. 5G	15
Authors' Addresses	15

1. Introduction

The telecommunications sector is experiencing a major revolution that will shape the way networks and services are designed and deployed for the next decade. We are witnessing an explosion in the number of applications and services demanded by users, which are now really capable of accessing them on the move. In order to cope with such a demand, some network operators are looking at the cloud computing paradigm, which enables a potential reduction of the overall costs by outsourcing communication services from specific hardware in the operator's core to server farms scattered in data centers. These services have different characteristics if compared with conventional IT services that have to be taken into account in this cloudification process. Also the transport network is affected in that it is evolving to a more sophisticated form of IP architecture with trends like separation of control and data plane traffic, and more fine-grained forwarding of packets (beyond looking at the destination IP address) in the network to fulfill new business and service goals.

Virtualization of functions also provides operators with tools to deploy new services much faster, as compared to the traditional use

of monolithic and tightly integrated dedicated machinery. As a natural next step, mobile network operators need to re-think how to evolve their existing network infrastructures and how to deploy new ones to address the challenges posed by the increasing customers' demands, as well as by the huge competition among operators. All these changes are triggering the need for a modification in the way operators and infrastructure providers operate their networks, as they need to significantly reduce the costs incurred in deploying a new service and operating it. Some of the mechanisms that are being considered and already adopted by operators include: sharing of network infrastructure to reduce costs, virtualization of core servers running in data centers as a way of supporting their load-aware elastic dimensioning, and dynamic energy policies to reduce the monthly electricity bill. However, this has proved to be tough to put in practice, and not enough. Indeed, it is not easy to deploy new mechanisms in a running operational network due to the high dependency on proprietary (and sometime obscure) protocols and interfaces, which are complex to manage and often require configuring multiple devices in a decentralized way.

Network function virtualization (NFV) [etsi_nfv_whitepaper] and software defined networking (SDN) [onf_sdn_architecture] are changing the way the telecommunications sector will deploy, extend and operate their networks. In this document we focus on one particular application of network softwarization: the radio access network (RAN), and in particular, how to run RAN functions on dynamic virtual resources very close to the users, the so-called Fog. We analyze the applicability of the SFC architecture to the Fog RAN use case.

2. 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 [RFC2119].

While [RFC2119] describes interpretations of these key words in terms of protocol specifications and implementations, they are used in this document to describe requirements for the SFC mechanisms to efficiently enable fog RAN.

The following terms used in this document are defined by the ETSI NFV ISG, the ONF and the IETF:

NFV Infrastructure (NFVI): totality of all hardware and software components which build up the environment in which VNFs are deployed

NFV Management and Orchestration (NFV-MANO): functions collectively provided by NFVO, VNFM, and VIM.

NFV Orchestrator (NFVO): functional block that manages the Network Service (NS) lifecycle and coordinates the management of NS lifecycle, VNF lifecycle (supported by the VNFM) and NFVI resources (supported by the VIM) to ensure an optimized allocation of the necessary resources and connectivity.

OpenFlow protocol (OFP): allowing vendor independent programming of control functions in network nodes.

Service Function (SF): a function that is responsible for specific treatment of received packets (e.g., firewall, load balancer).

Service Function Chain (SFC): for a given service, the abstracted view of the required service functions and the order in which they are to be applied. This is somehow equivalent to the Network Function Forwarding Graph (NF-FG) at ETSI.

Service Function Path (SFP): the selection of specific service function instances on specific network nodes to form a service graph through which an SFC is instantiated.

Virtualized Infrastructure Manager (VIM): functional block that is responsible for controlling and managing the NFVI compute, storage and network resources, usually within one operator's Infrastructure Domain.

Virtualized Network Function (VNF): implementation of a Network Function that can be deployed on a Network Function Virtualisation Infrastructure (NFVI).

Virtualized Network Function Manager (VNFM): functional block that is responsible for the lifecycle management of VNF.

3. Fog RAN Overview

Virtualization is invading all domains of the E2E 5G network, including the access, as a mean to achieve the necessary flexibility in support of the E2E slicing concept. The ETSI NFV framework is the cornerstone for making virtualization such a promising technology that can be matured in time for 5G. Typically, virtualization has been mostly envisaged in the core network, where sophisticated data centers and clouds provided the right substrate. And mostly, the framework focused on virtualizing network functions, so called VNFs (virtualized network functions), which were somewhat limited to functions that are delay tolerant, typically from the core and

aggregation transport. Yet in the access, virtualization of RAN functions could still be envisaged as one step forward for the Cloud-RAN concept, assuming an extremely low latency fronthaul transport (typically over fiber) and again limiting to those RAN functions which delay requirements match for execution in a virtualized environment.

As the community has recently been developing the 5G applications and their technical requirements, it has become clear that certain applications would require very low latency which is extremely challenging and stressing for the network to deliver through a pure centralized architecture. The need to provide networking, computing, and storage capabilities closer to the users has therefore emerged, leading to what is known today as the concept of intelligent edge. ETSI has been the first to address this need recently by developing the framework of mobile edge computing (MEC).

Such an intelligent edge could not be envisaged without virtualization. Beyond applications, it raises a clear opportunity for networking functions to execute at the edge benefiting from inherent low latencies. Being in close proximity to the access, the edge becomes thus an attractive place for hosting C-RAN functions in particular, which could also be envisaged virtualized thanks to the virtualization capabilities now available at the edge. Transport and core networking functions could also take advantage of such a hosting environment, thus saving bandwidth in their respective domains and offering local breakout options where required. Furthermore, a rich set of context information from the RAN but also from other network domains could be offered as services through the edge for applications to consume and hence optimize their behavior or key performance indicators (KPIs).

Whilst it is appreciated the particular challenge for the intelligent edge concept in dealing with mobile users, the edge virtualization substrate has been largely assumed to be fixed or stationary. Although little developed, the intelligent edge concept is being extended further to scenarios where for example the edge computing substrate is on the move, e.g., on-board a car or a train, or that it is distributed further down the edge, even integrating resources from different stakeholders, into what is known as the fog. The challenges and opportunities for such extensions of the intelligent edge remain an exciting area of future research.

The computing and virtualization capabilities available down into the fog are of particular advantage to the virtualized-RAN/cloud-RAN, leading to what we refer to as "Fog RAN". Virtual networking functions (VNFs) related to the RAN may execute in the fog. The close proximity of the fog to the end user devices opens new

possibilities for extending the scope of virtual RAN (VRAN) functions from just access nodes so far to end user devices, so that functions traditionally running on the end user devices may be moved to the fog. In addition, an additional tier of intelligent VRAN functions could be envisaged to run across other functions of various nodes or devices and from different radio access technologies (RATs), supporting tight coordination between these nodes and devices, and enabling convergence between the RATs. VNFs from the transport and core could also be hosted in the fog so as to save bandwidth in their respective domains.

Figure 1 shows a diagram representing the fog RAN concept. The fog is composed by virtual resources on top of heterogeneous resources available at the edge and even further in the RAN and end-user devices. These resources are therefore owned by different stakeholders who collaboratively form a single hosting environment for the VNFs to run. As an example, virtual resources provided to the fog might be running on eNBs, APs, at micro data centers deployed in shopping malls, cars, trains, etc. The fog is connected to data centers deeper into the network architecture (at the edge or the core). On the top part of the figure, an example of user and control plane VNFs is shown. User plane VNFs are represented as "fx", and control ones as "ctrlx". Depending on the functionality implemented by these VNFs and the service requirements, these VNFs would be mapped (i.e., instantiated) differently to the physical resources (as described in [I-D.aranda-sfc-dp-mobile]).




Figure 1: Fog RAN

The fog is also well suited to offer a rich set of context information noticeably (but not only) from the RAN that could be collected from the various RATs co-existing in the same service area. This information can be obtained either from networking resources (e.g., nodes and devices) or functions, some of which might be hosted in the fog. Such information may be exposed through services running in the fog or in the edge, for applications on top to consume and hence optimize their performance or behavior. The fog or edge will therefore be in charge of collecting and publishing context information towards network applications as well as end user/3rd party applications. This is accomplished by the so-called "services" which follow the concept of ETSI MEC ISG services and may run in the fog. Applications may also run directly in the fog and subscribe to one or more services provided in the fog. These fog applications

could also offer services to other applications residing inside the fog or outside, e.g., in the edge or central cloud. This enables different kinds of Over-The-Top (OTT) applications to utilize the RAN context information available at the fog.

4. Applicability of SFC to Fog RAN

For a given service, the abstracted view of the required service functions and the order in which they are to be applied is called a service function chain (SFC), which is called network function forwarding graph (NF-FG) in ETSI. An SFC is instantiated through selection of specific service function instances on specific network nodes to form a service graph: this is called a service function path (SFP). The service functions may be applied at any layer within the network protocol stack (network layer, transport layer, application layer, etc.).

[RFC7665] describes the architecture for the specification, creation, and ongoing maintenance of service function chains (SFCs) in a network. The SFC architecture is composed of the following logical components: classifiers, service function forwarders (SFFs), service functions (SFs), and SFC proxies. These components are interconnected using the SFC encapsulation, as shown in Figure 2.




Figure 2: SFC architecture

An overview of how the SFC logical components interact after the initial classification is also shown in Figure 3 (reproduced from [RFC7665] for completeness).




Figure 3: SFC architecture components

There are different use cases for service function chaining in mobile networks. [I-D.ietf-sfc-use-case-mobility] describes in general how to use SFC for mobile networks focusing on the core functions, while [I-D.aranda-sfc-dp-mobile] looks more at RAN aspects.

This document focuses on the specific use case of applying SFC concepts to the fog RAN environment, which is characterized mainly by:

- o The VNFs being chained implement mostly RAN functionality. The Cloud RAN (C-RAN) approach centralizes baseband processing and network resource allocation (which are today functions executed in a distributed way at the access nodes, such as eNBs). The strict latency and bandwidth requirements imposed by C-RAN triggered the evolution of the virtualization of the access infrastructure to the concept of RAN as a Service, where the centralization (i.e.

function virtualization in a data center) of processing and management in mobile networks is flexible and adapted to the actual service requirements and the network conditions. This is also referred to as flexible functional split of the radio protocol stack. Up to now, RAN virtualization approaches have only considered offloading of computation tasks in central clouds or regional data centers close to the network aggregation points. With fog RAN, virtualization resources are placed closer to users, from edge computing to fog computing at user devices, leveraging micro data centers at user premises, thus reducing end-to-end latency.

- o Virtualizing RAN functions at the fog could also enable new optimizations when information of (or available at) the access is used. Examples of this information include: RAN measured parameters, location information, etc. This information can be used for example to jointly optimize multiple RATs.
- o The fog computing environment can also be used to virtualize functions from the end-user devices, not only from the RAN. This can help facilitating a better convergence of multiple access technologies.

Fog RAN implementations will benefit from applying SFC concepts as virtual RAN functions will be executed on virtual resources from different stakeholders, meaning different data centers managed by different entities. In this environment, SFC encapsulation can be used to ensure proper data processing. Figure 4 shows an example of scenario of application of SFC in the fog RAN. On the top part we show a UE connected to the network. The eNB functionality is actually split, so part of it (e.g., RLC, PDCP and RRC) runs virtualized in the fog, while the rest (e.g., MAC and PHY) stay at the remote radio head (RRH). Other mobile network RAN functionality can also be running virtualized in the fog, such as a local virtual MME, multi-RAT authentication and RAT selection mechanisms, offloading solutions (such as local vEPCs), etc. The rest of the mobile network functions (see Appendix A for a short overview) runs in the core. On the bottom part of the figure we show a potential mapping of the virtual functions to the physical resources that compose the fog. SFC mechanisms would be used to interconnect the different functions in the right order and meeting the requirements (latency, compute, etc) needed.

REQ-R3: SFC MUST support Service Functions (e.g., MEC) that are located at the edge of the network and that perform L7-Application processing very early in the forwarding path.

REQ-R4: SFC MUST support metadata used to exchange information about the network and virtual resources status, so it can be used to decide about updates of the service function path.

REQ-RX: More TBD...

6. IANA Considerations

N/A.

7. Security Considerations

TBD.

8. Acknowledgments

TBD.

9. References

9.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.

9.2. Informative References

[etsi_nvf_whitepaper]
"Network Functions Virtualisation (NFV). White Paper 2", October 2014.

[I-D.aranda-sfc-dp-mobile]
Aranda, P., Gramaglia, M., Lopez, D., and W. Haeffner, "Service Function Chaining Dataplane Elements in Mobile Networks", draft-aranda-sfc-dp-mobile-02 (work in progress), July 2016.

[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-07 (work in progress), October 2016.

[onf_sdn_architecture]

"SDN Architecture (Issue 1.1), ONF TR-521", February 2016.

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

Appendix A. 4G (LTE)

This annex includes a high level summary of the 3GPP EPS architecture, commonly referred to as 4G (LTE), detailing both the EPC (core) and the RAN (access) parts.

The EPS architecture and some of its standardized interfaces are depicted in Figure 5. The EPS provides IP connectivity to user equipment (UE) (i.e., mobile nodes) and access to operator services, such as global Internet access and voice communications. The EPS comprises the core network -- called Evolved Packet Core (EPC) -- and different radio access networks: the 3GPP Access Network (AN), the Untrusted non-3GPP AN and the Trusted non-3GPP AN. There are different types of 3GPP ANs, with the evolved UMTS Terrestrial Radio Access Network (E-UTRAN) as the most advanced one. QoS is supported through an EPS bearer concept, providing bindings to resource reservation within the network.

The evolved NodeB (eNB), the Long Term Evolution (LTE) base station, is part of the access network that provides radio resource management, header compression, security and connectivity to the core network through the S1 interface. In an LTE network, the control plane signaling traffic and the data traffic are handled separately. The eNBs transmit the control traffic and data traffic separately via two logically separate interfaces.

The Home Subscriber Server, HSS, is a database that contains user subscriptions and QoS profiles. The Mobility Management Entity, MME, is responsible for mobility management, user authentication, bearer establishment and modification and maintenance of the UE context.

The Serving gateway, S-GW, is the mobility anchor and manages the user plane data tunnels during the inter-eNB handovers. It tunnels all user data packets and buffers downlink IP packets destined for UEs that happen to be in idle mode.

The Packet Data Network (PDN) Gateway, P-GW, is responsible for IP address allocation to the UE and is a tunnel endpoint for user and control plane protocols. It is also responsible for charging, packet

filtering, and policy-based control of flows. It interconnects the mobile network to external IP networks, e.g. the Internet.

In this architecture, data packets are not sent directly on an IP network between the eNB and the gateways. Instead, every packet is tunneled over a tunneling protocol - the GPRS Tunneling Protocol (GTP over UDP/IP). A GTP path is identified in each node with the IP address and a UDP port number on the eNB/gateways. The GTP protocol carries both the data traffic (GTP-U tunnels) and the control traffic (GTP-C tunnels). Alternatively Proxy Mobile IP (PMIPv6) is used on the S5 interface between S-GW and P-GW.

In addition to the above basic functions and entities, there are also additional features being discussed by the 3GPP that are relevant from a network virtualization viewpoint. One example is the Traffic Detection Function (TDF), which can be used by the P-GW, and in general by the whole transport network, to decide how to forward the traffic. In a virtualized infrastructure, this kind of information can be used to elastic and dynamically adapt the network capabilities to the traffic nature and volume.




Figure 5: EPS (non-roaming) architecture overview

Appendix B. 5G

TBD. This section will include a description of main 5G building blocks.

Authors' Addresses

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: <http://www.it.uc3m.es/cjbc/>

Akbar Rahman
InterDigital Communications, LLC
1000 Sherbrooke Street West, 10th floor
Montreal, Quebec H3A 3G4
Canada

Email: Akbar.Rahman@InterDigital.com
URI: <http://www.InterDigital.com/>

Alain Mourad
InterDigital Europe

Email: Alain.Mourad@InterDigital.com
URI: <http://www.InterDigital.com/>

Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 16, 2017

T. Mizrahi
I. Yerushalmi
D. Melman
Marvell
R. Browne
Intel
January 12, 2017

Network Service Header (NSH) Context Header Allocation: Timestamp
draft-mymb-sfc-nsh-allocation-timestamp-00

Abstract

This memo defines an allocation for the Context Headers of the Network Service Header (NSH), which incorporates the packet's ingress timestamp, a sequence number, and a source interface identifier.

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 <http://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 July 16, 2017.

Copyright Notice

Copyright (c) 2017 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	2
2. Terminology	3
2.1. Requirements Language	3
2.2. Abbreviations	3
3. NSH Context Header Allocation Allocation	3
4. Ingress Timestamping Use Cases	4
4.1. Network Analytics	4
4.2. Alternate Marking	5
4.3. Consistent Updates	5
5. Synchronization Considerations	5
6. IANA Considerations	5
7. Security Considerations	5
8. References	6
8.1. Normative References	6
8.2. Informative References	6
Authors' Addresses	7

1. Introduction

The Network Service Header (NSH), defined in [I-D.ietf-sfc-nsh], is an encapsulation header that is used in Service Function Chains (SFC).

The NSH specification [I-D.ietf-sfc-nsh] supports two possible methods of including metadata in the NSH; MD Type 0x1 and MD Type 0x2. When using MD Type 0x1 the NSH includes 16 octets of Context Header fields. The current memo proposes an allocation for the MD Type 0x1 Context Headers, which incorporates the ingress timestamp of the packet, a sequence number, and a source interface identifier.

In a nutshell, packets that enter the SFC-Enabled Domain are timestamped. The ingress timestamp is measured by the Classifier [RFC7665], and incorporated in the NSH. The ingress timestamp may be used for various different purposes, including delay measurement, packet marking for passive performance monitoring, and timestamp-based policies. Notably, the timestamp does not increase the packet length, since it is incorporated in the MD Type 0x1 Mandatory Context Headers.

The source interface identifier indicates the interface through which the packet was received at the classifier. This identifier may specify a physical or a virtual interface. The sequence numbers can be used by Service Functions (SFs) to detect out-of-order delivery or

duplicate transmissions. The sequence number is maintained on a per-source-interface basis.

KPI-stamping [I-D.browne-sfc-nsh-kpi-stamp] defines an NSH timestamping mechanism that uses the MD Type 0x2 format. The current memo defines a compact MD Type 0x1 Context Header that does not require the packet to be extended beyond the NSH header. Furthermore, the two timestamping mechanisms can be used in concert, as further discussed below.

2. Terminology

2.1. Requirements Language

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].

2.2. Abbreviations

KPI	Key Performance Indicators [I-D.browne-sfc-nsh-kpi-stamp]
NSH	Network Service Header [I-D.ietf-sfc-nsh]
MD	Metadata [I-D.ietf-sfc-nsh]
SF	Service Function [RFC7665]
SFC	Service Function Chaining [RFC7665]

3. NSH Context Header Allocation Allocation

This memo defines the following Context Header allocation, as presented in Figure 1.

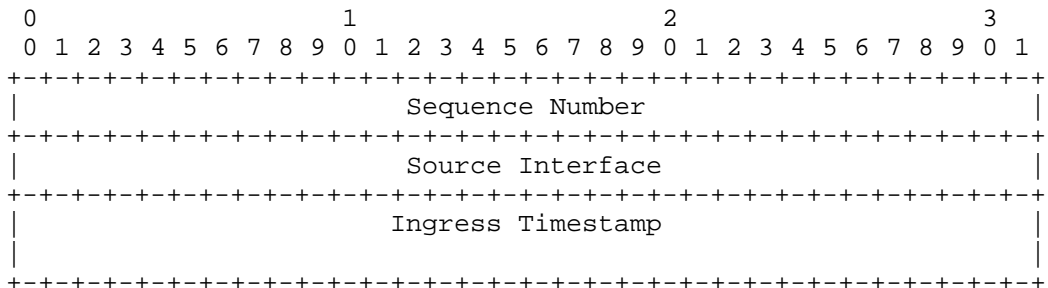


Figure 1: NSH Ingress Timestamp Allocation.

The NSH Timestamp Allocation includes the following fields:

- o Sequence Number - a 32-bit sequence number. The sequence number is maintained on a per-source-interface basis. The sequence numbers can be used by SFs to detect out-of-order delivery, or duplicate transmissions.
- o Source Interface - a 32-bit source interface identifier that is assigned by the Classifier.
- o Ingress Timestamp - this field is 8 octets long, and specifies the time at which the packet was received by the Classifier. The format of this field uses the 64 least significant bits of the IEEE 1588-2008 Precision Time Protocol format [IEEE1588]. This truncated format consists of a 32-bit seconds field followed by a 32-bit nanoseconds field. As defined in [IEEE1588], the timestamp specifies the number of seconds elapsed since 1 January 1970 00:00:00 according to the International Atomic Time (TAI).




Figure 2: IEEE 1588 Truncated Timestamp Format [IEEE1588].

4. Ingress Timestamping Use Cases

4.1. Network Analytics

Per-packet timestamping enables coarse-grained monitoring of the network delay along the Service Function Chain. Once a potential problem or bottleneck is detected, for example when the delay exceeds a certain policy, a highly-granular hop-by-hop monitoring mechanism, such as [I-D.browne-sfc-nsh-kpi-stamp] or [I-D.brockners-inband-oam-data], can be triggered, allowing to analyze and localize the problem.

Timestamping is also useful for logging and for flow analytics. It is often useful to maintain the timestamp of the first and last packet of the flow. Furthermore, traffic mirroring and sampling often requires a timestamp to be attached to analyzed packets. Attaching the ingress timestamp to the NSH Context Header provides an

in-band common time reference that can be used for various network analytics applications.

4.2. Alternate Marking

A possible approach for passive performance monitoring is to use an alternate marking method [I-D.ietf-ippm-alt-mark]. This method requires data packets to carry a field that marks (colors) the traffic, and enables passive measurement of packet loss, delay, and delay variation. The value of this marking field is periodically toggled between two values.

When the ingress timestamp is incorporated in the NSH Context Header, it can natively be used for alternate marking. For example, the least significant bit of the timestamp Seconds field can be used for this purpose, since the value of this bit is inherently toggled every second.

4.3. Consistent Updates

The timestamp can be used for taking policy decisions such as 'Perform action A if timestamp>=T_0'. This can be used for enforcing time-of-day policies or periodic policies in service functions. Furthermore, timestamp-based policies can be used for enforcing consistent network updates, as discussed in [DPT].

5. Synchronization Considerations

Some of the applications that make use of the ingress timestamp require the Classifier and SFs to be synchronized to a common time reference, for example using the Network Time Protocol [RFC5905], or the Precision Time Protocol [IEEE1588].

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

The security considerations of NSH in general are discussed in [I-D.ietf-sfc-nsh]. The security considerations of in-band timestamping in the context of NSH is discussed in [I-D.browne-sfc-nsh-kpi-stamp], and the current section is based on that discussion.

The use of in-band timestamping, as defined in this document, can be used as a means for network reconnaissance. By passively eavesdropping to timestamped traffic, an attacker can gather

information about network delays and performance bottlenecks. A man-in-the-middle attacker can maliciously modify timestamps in order to attack applications that use the timestamp values, such as performance monitoring applications.

Since the timestamping mechanism relies on an underlying time synchronization protocol, by attacking the time protocol an attack can potentially compromise the integrity of the NSH timestamp. A detailed discussion about the threats against time protocols and how to mitigate them is presented in [RFC7384].

8. References

8.1. Normative References

- [I-D.ietf-sfc-nsh]
Quinn, P. and U. Elzur, "Network Service Header", draft-ietf-sfc-nsh-10 (work in progress), September 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

8.2. Informative References

- [DPT] Mizrahi, T., Moses, Y., "The Case for Data Plane Timestamping in SDN", IEEE INFOCOM Workshop on Software-Driven Flexible and Agile Networking (SWFAN), 2016.
- [I-D.brockners-inband-oam-data]
Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., and R. <>, "Data Formats for In-situ OAM", draft-brockners-inband-oam-data-02 (work in progress), October 2016.
- [I-D.browne-sfc-nsh-kpi-stamp]
Browne, R., Chilikin, A., and T. Mizrahi, "Network Service Header KPI Stamping", draft-browne-sfc-nsh-kpi-stamp-00 (work in progress), October 2016.
- [I-D.ietf-ippm-alt-mark]
Fioccola, G., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate Marking method for passive performance monitoring", draft-ietf-ippm-alt-mark-02 (work in progress), October 2016.

- [IEEE1588] IEEE, "IEEE 1588 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems Version 2", 2008.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<http://www.rfc-editor.org/info/rfc5905>>.
- [RFC7384] Mizrahi, T., "Security Requirements of Time Protocols in Packet Switched Networks", RFC 7384, DOI 10.17487/RFC7384, October 2014, <<http://www.rfc-editor.org/info/rfc7384>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <<http://www.rfc-editor.org/info/rfc7665>>.

Authors' Addresses

Tal Mizrahi
Marvell
6 Hamada
Yokneam 2066721
Israel

Email: talmi@marvell.com

Ilan Yerushalmi
Marvell
6 Hamada
Yokneam 2066721
Israel

Email: yilan@marvell.com

David Melman
Marvell
6 Hamada
Yokneam 2066721
Israel

Email: davidme@marvell.com

Rory Browne
Intel
Dromore House
Shannon, Co.Clare
Ireland

Email: rory.browne@intel.com

Service Function Chaining
Internet-Draft
Intended status: Standards Track
Expires: July 20, 2017

J. Napper
S. Kumar
Cisco Systems, Inc.
P. Muley
W. Hendericks
Nokia
M. Boucadair
Orange
January 16, 2017

NSH Context Header Allocation -- Broadband
draft-napper-sfc-nsh-broadband-allocation-02

Abstract

This document provides a recommended allocation of context headers for a Network Service Header (NSH) within the broadband service provider network context. NSH is described in detail in [ietf-sfc-nsh]. This allocation is intended to support uses cases as defined in [ietf-sfc-use-case-mobility].

Requirements Language

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 [RFC2119].

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 <http://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 July 20, 2017.

Copyright Notice

Copyright (c) 2017 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	2
2. Definition Of Terms	3
3. Network Service Header (NSH) Context Headers	3
4. Recommended Context Allocation	4
4.1. MD Type 0x01 Allocation Specifics	4
4.2. MD Type 0x02 Allocation Specifics	6
5. Context Allocation and Control Plane Considerations	7
6. Security Considerations	7
7. IANA Considerations	7
8. Acknowledgments	7
9. References	8
9.1. Normative References	8
9.2. Informative References	8
Authors' Addresses	9

1. Introduction

Service function chaining provides a mechanism for network traffic to be steered through multiple service functions in a sequence. Metadata can be useful to service functions. The Network Service Header (NSH) provides support for carrying shared metadata between service functions (and devices) either using 4 fixed-length 32-bit context headers or as optional TLVs as defined in [ietf-sfc-nsh]. NSH is then encapsulated within an outer header for transport.

This document provides a recommended default allocation scheme for the fixed-length context headers and for TLV types in the context of service chaining within fixed and mobile broadband service provider networks. Supporting use cases describing the need for a metadata header in these contexts are described in

[ietf-sfc-use-case-mobility]. This draft does not address control plane mechanisms.

2. Definition Of Terms

This document uses the terms as defined in [RFC7498] and [RFC7665].

3. Network Service Header (NSH) Context Headers

In Service Function Chaining, the Network Service Header is composed of a 4-byte base header (BH1), a 4-byte service path header (SH1) and four mandatory 4-byte context headers (CH1-CH4) in the case of MD Type 0x01 and optional TLVs in the case of MD Type 0x02 as described in [ietf-sfc-nsh].

The following Figure 1 shows the MD Type 0x01 mandatory context headers.

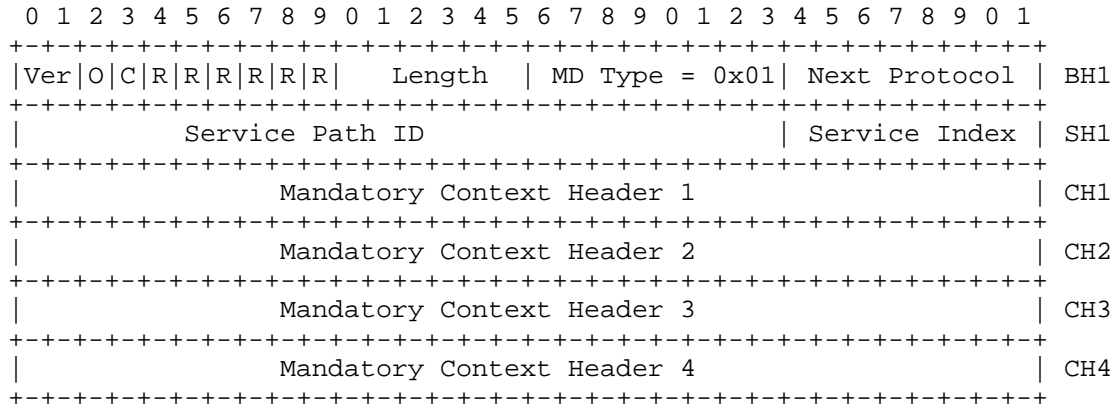


Figure 1: Network Service Header - MD Type 0x01

The following Figure 2 shows the MD Type 0x02 optional TLV header format.




Figure 2: Network Service Header - MD Type 0x02

4. Recommended Context Allocation

The following header allocations provide information to support service function chaining in a service provider network, for example as described for mobility in [ietf-sfc-use-case-mobility].

The set of metadata headers can be delivered to service functions that can use the metadata within to enforce policy, communicate between service functions, provide subscriber information and other functionality. Several of the headers are typed allowing for different metadata to be provided to different service functions or even to the same service function but on different packets within a flow. Which metadata are sent to which service functions is decided in the SFC control plane and is thus out of the scope of this document.

4.1. MD Type 0x01 Allocation Specifics

The following Figure 3 provides a high-level description of the fields in the recommended allocation of the fixed context headers for a mobility context.




Figure 3: NSH Context Allocation

The intended use for each of the context header allocations is as follows:

R - Reserved.

Sub - Sub/Endpoint ID type field. These bits determine the type of the 64-bit Sub/Endpoint ID field that spans CH2 and CH3.

000 - If the Sub field is not set, then the 64-bit Sub/Endpoint ID field is an opaque field that can be used or ignored by service functions as determined by the control plane.

001 - The Sub/Endpoint ID field contains an IMSI [itu-e-164].

010 - The Sub/Endpoint ID field contains an MSISDN (8-15 digit) [itu-e-164].

011 - The Sub/Endpoint ID field contains a 64-bit identifier that can be used to group flows (e.g., in Machine-to-Machine, M2M).

100 - The Sub/Endpoint IP field contains a wireline subscriber ID in CH2, and CH3 contains the home identifier.

101-111 - Reserved.

Tag - The Tag field indicates the type of the Service Information field in CH4. Some types for this field are specified by the Tag field as follows:

000 - If the Tag field is not set, then the Service Information field in CH4 is an opaque field that can be used or ignored by service functions as determined by the control plane.

001 - The Service Information field in CH4 contains information related to the Access Network (AN) for the subscriber. This is shown in Figure 4 for a 3GPP Radio Access Network (RAN). Note that these values should correspond to those that can be obtained for the flow from the corresponding 3GPP PCRF (Policy and Charging Rules Function) component using Diameter as described in [TS.29.230].




Figure 4: Service Information RAN Allocation

The intended use of the header is for TLVs associated with 3GPP Radio Access Networks as described in [TS.29.230]. This TLV can be used by 3GPP to extend the metadata as per use cases. Having this TLV helps to carry more information that does not fit within the MD Type 0x01.

The Len field carries the total length. Type = 0x01 is reserved. If set to 0x01, the TLV carries the 4 context headers as defined in Section 4.1.

5. Context Allocation and Control Plane Considerations

This document describes an allocation scheme for both the mandatory context headers and optional TLV headers in the context of broadband service providers. This suggested allocation of headers should be considered as a guideline and may vary depending on the use case. The control plane aspects of specifying and distributing the allocation scheme among different service functions within the Service Function Chaining environment to guarantee consistent semantics for the metadata is beyond the scope of this document.

6. Security Considerations

The header allocation recommended by this document includes numbers that must be distributed consistently across a Service Function Chaining environment. Protocols for distributing these numbers securely are required in the control plane, but are out of scope of this document.

Furthermore, some of the metadata carried in the headers require secure methods to prevent spoofing or modification by service function elements that may themselves be exposed to subscriber traffic and thus might be compromised. This document does not address such security concerns.

7. IANA Considerations

This document requests IANA to assign a TLV class for 3GPP to be used for its use cases.

8. Acknowledgments

The authors would like to thank Jim Guichard for his assistance structuring the document.

9. References

9.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.

9.2. Informative References

- [ietf-sfc-nsh] Quinn, P. and U. Elzur, "Network Service Header", I-D draft-ietf-sfc-nsh-10 (work in progress), September 2016.
- [ietf-sfc-use-case-mobility] Haeffner, W., Napper, J., Stiemerling, M., Lopez, D., and J. Uttaro, "Service Function Chaining Use Cases in Mobile Networks", I-D draft-ietf-sfc-use-case-mobility-07 (work in progress), October 2016.
- [itu-e-164] "The international public telecommunication numbering plan", ITU-T E.164, November 2010.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<http://www.rfc-editor.org/info/rfc2474>>.
- [RFC7498] Quinn, P., Ed. and T. Nadeau, Ed., "Problem Statement for Service Function Chaining", RFC 7498, DOI 10.17487/RFC7498, April 2015, <<http://www.rfc-editor.org/info/rfc7498>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <<http://www.rfc-editor.org/info/rfc7665>>.
- [TS.29.230] "Diameter applications; 3GPP specific codes and identifiers", 3GPP TS 29.230 14.3.0, December 2016.

Authors' Addresses

Jeffrey Napper
Cisco Systems, Inc.

Email: jenapper@cisco.com

Surendra Kumar
Cisco Systems, Inc.

Email: smkumar@cisco.com

Praveen Muley
Nokia

Email: praveen.muley@nokia.com

Wim Hendericks
Nokia

Email: Wim.Henderickx@nokia.com

Mohamed Boucadair
Orange

Email: mohamed.boucadair@orange.com