

Service Function Chaining
Internet-Draft
Intended status: Informational
Expires: November 6, 2014

W. Haeffner
Vodafone
J. Napper
Cisco Systems
M. Stiemerling
NEC
D. Lopez
Telefonica I+D
J. Uttaro
AT&T
May 5, 2014

Service Function Chaining Use Cases in Mobile Networks
draft-haeffner-sfc-use-case-mobility-02

Abstract

This document provides some exemplary use cases for service function chaining in mobile service provider networks. The objective of this draft is not to cover all conceivable service chains in detail. Rather, the intention is to localize and explain the application domain of service chaining within mobile networks as far as it is required to complement the problem statement and framework statements of the working group.

Service function chains typically reside in a LAN segment which links the mobile access network to the actual application platforms located in the carrier's datacenters or somewhere else in the Internet. Service function chains ensure a fair distribution of network resources according to agreed service policies, enhance the performance of service delivery, take care of security and privacy or support application and business support platforms. General considerations and specific use cases are presented in this document to demonstrate the different technical requirements of these goals for service function chaining in mobile service provider networks.

The specification of service function chaining for mobile networks must take into account an interaction between service function chains and the 3GPP Policy and Charging Control (PCC) environment.

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 November 6, 2014.

Copyright Notice

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

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Terminology and abbreviations	3
1.2.	General scope of mobile service chains	4
1.3.	General structure of end-to-end carrier networks	5
2.	Mobile network overview	6
2.1.	Building blocks of 3GPP mobile networks	6
2.2.	Overview of mobile service chains	7
2.3.	The most common classification scheme	9
2.4.	More sophisticated classification schemes	10
3.	Example use cases specific to mobile networks	11
3.1.	Service chain model for Internet HTTP services	11
3.1.1.	Weaknesses of current approaches	15
3.2.	Service chain for TCP optimization	15
3.2.1.	Weaknesses of current approaches	16

3.3.	HTTP header enrichment in mobile networks	16
4.	Remarks on QoS in mobile networks	17
5.	Weaknesses of current implementations	18
5.1.	Granularity of the classification scheme	18
5.2.	Service chain implementations	18
6.	Operator requirements	19
6.1.	Simplicity of service function chain instantiation . . .	19
6.2.	Extensions	20
6.3.	Delimitations	21
7.	Security Considerations	21
8.	IANA Considerations	21
9.	Acknowledgments	21
10.	References	21
10.1.	Normative References	21
10.2.	Informative References	21
	Authors' Addresses	22

[1.](#) Introduction

[1.1.](#) Terminology and abbreviations

Much of the terminology used in this document has been defined by the 3rd Generation Partnership Project (3GPP), which defines standards for mobile service provider networks. Although a few terms are defined here for convenience, further terms can be found in [\[RFC6459\]](#).

UE User equipment like tablets or smartphones

eNB enhanced NodeB, radio access part of the LTE system

S-GW Serving Gateway, primary function is user plane mobility

P-GW Packet Gateway, actual service creation point, terminates 3GPP mobile network, interface to Packet Data Networks (PDN)

HSS Home Subscriber System (control plane element)

MME Mobility Management Entity (control plane element)

GTP GPRS (General Packet Radio Service) Tunnel Protocol

S-IP Source IP address

D-IP Destination IP address

IMSI The International Mobile Subscriber Identity that identifies a mobile subscriber

(S)Gi Egress termination point of the mobile network (SGi in case of LTE, Gi in case of UMTS/HSPA). The internal data structure of this interface is not standardized by 3GPP

PCRF 3GPP standardized Policy and Charging Rules Function

1.2. General scope of mobile service chains

Mobile access networks terminate at a mobile service creation point (Packet Gateway) typically located at the edge of an operator IP backbone. From the user equipment (UE) up to the Packet Gateway (P-GW) everything is fully standardized by the 3rd Generation Partnership Project (3GPP) e.g., in [[TS.23.401](#)]. Within the mobile network, the user payload is encapsulated in 3GPP specific tunnels terminating eventually at the P-GW. In many cases application-specific IP traffic is not directly exchanged between the original mobile network, more specific the P-GW, and an application platform, but will be forced to pass a set of service functions. Those application platforms are, for instance, a web server environment, a video platform, a social networking platform or some other multimedia platform. Network operators use these service functions to differentiate their services to their subscribers. Service function chaining is thus integral to the business model of operators.

Important use cases classes for service function chains generally include:

1. functions to protect the carrier network and the privacy of its users(IDS, FW, ACL, encryption, decryption, etc.),
2. functions that ensure the contracted quality of experience using e.g., performance enhancement proxies (PEP) like video optimizers, TCP optimizers or functions guaranteeing fair service delivery based on policy based QoS mechanisms,
3. functions like HTTP header enrichment that may be used to identify and charge subscribers real time,
4. functions like CG-NAT/PAT, which are required solely for technical reasons, and
5. functions like parental control or malware detection that may be a cost option of a service offer.

1.3. General structure of end-to-end carrier networks

Although this memo is focused on the Service Function Chaining use cases for mobile carrier networks, such as 3GPP- based ones, a number of other, different carrier networks exists that share similarities in the structure of the access networks and the service functions with mobile networks.

Figure 1 shows 4 different carrier networks as examples to show similarities with respect to Service Functions and their Chaining.

The service networks consist of access-specific user equipment, a dedicated access network, a related service creation point and finally a (LAN) infrastructure hosting Service Functions which finally interconnect to application platforms in the Internet or in the carrier's own datacenter (DC).

From top to down, there is a 3GPP mobile network terminating at the P-GW, a xDSL network with its PPP tunnels terminating at a BNG (Broadband Network Gateway), a FTTH network terminating at an OLT (optical line terminal) and finally a cable TV network terminating at a CMTS (cable modem termination system).

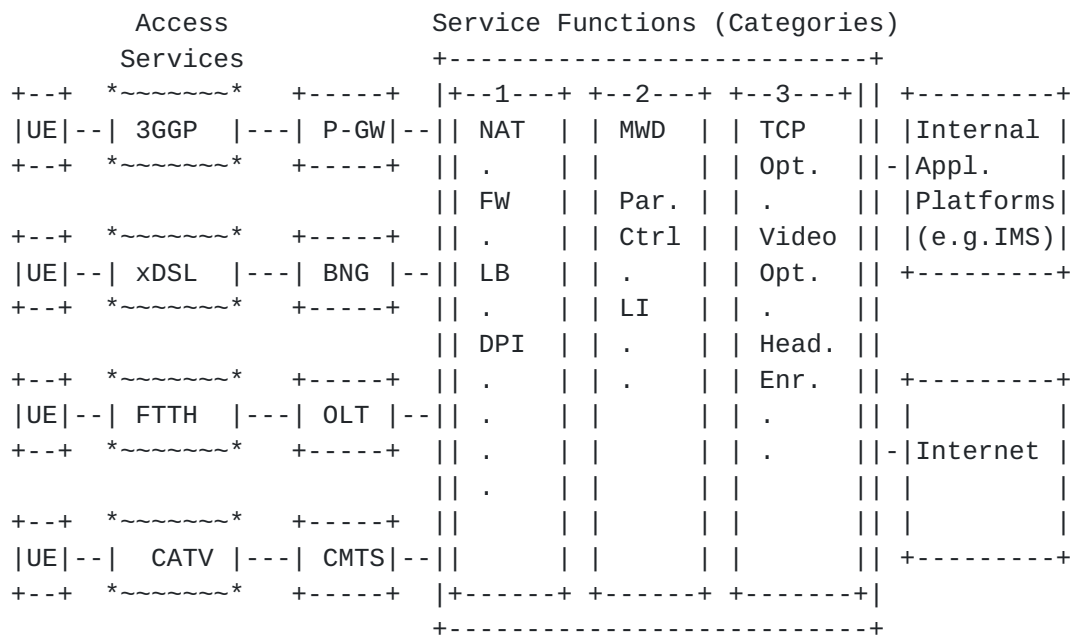


Figure 1: Various end-to-end carrier networks and service functions.

Category 1 of service functions like NAT or DPI may be used by all of these service networks mainly just (but not exclusively) for technical reasons. The same is true for category 2, value added services (VAS) like parental control, malware detection and

elimination (MWD) or legal intercept (LI). TCP optimization is basically seen in mobile networks only, the same may be true for video optimizers or HTTP header enrichment, i.e., category 3 as a rule mainly belongs to mobile networks only.

In our view, 3GPP-based mobile networks seem to have the largest demand for service functions and service function chains. Service Function Chains used in other access networks are very likely a subset of what one can expect in 3GPP-based mobile networks.

2. Mobile network overview

For simplicity we only describe service function chaining in the context of LTE (Long Term Evolution) networks. But indeed our service chaining model also applies to earlier generations of mobile networks, such as purely UMTS-based mobile networks.

2.1. Building blocks of 3GPP mobile networks

The major functional components of a LTE network are shown in Figure 2 and include user equipment (UE) like smartphones or tablets, the LTE radio unit named enhanced NodeB (eNB), the serving gateway (S-GW) which together with the mobility management entity (MME) takes care of mobility and the packet gateway (P-GW), which finally terminates the actual mobile service. These elements are described in detail in [\[TS.23.401\]](#). Other important components are the home subscriber system (HSS) and the policy and charging rule function (PCRF), which are described in [\[TS.23.203\]](#). The P-GW interface towards the SGi-LAN is called the SGi-interface, which is described in [\[TS.29.061\]](#) and finally the SGi-LAN is the home of service function chains (SFC), which are not standardized by 3GPP.

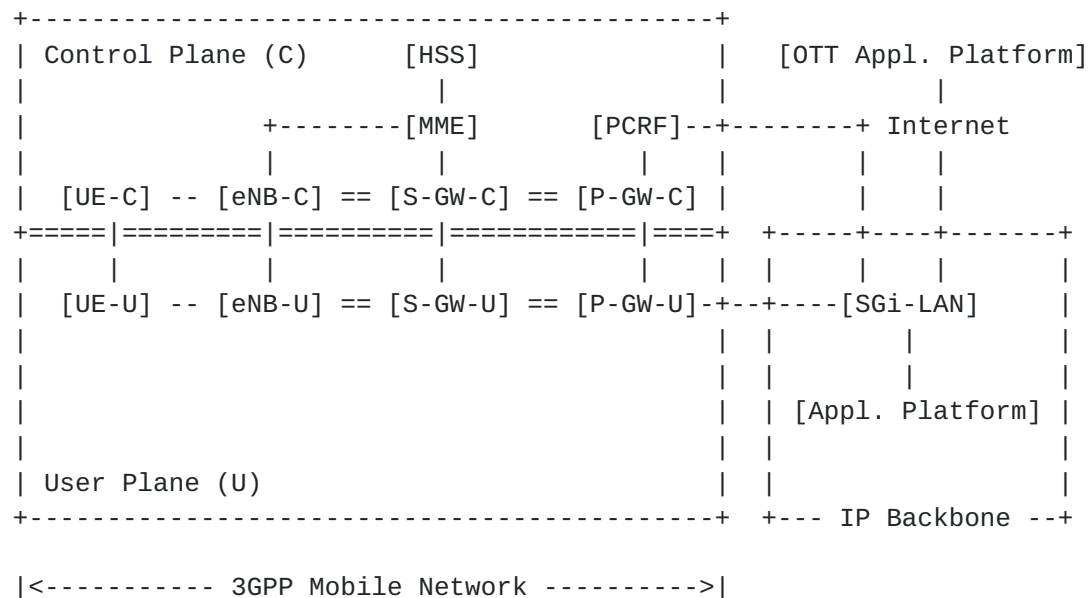


Figure 2 shows the end to end context including all major components of a LTE network. The actual 3GPP mobile network includes the elements from the user equipment [UE] to the packet gateway [P-GW].

Figure 2: End to end context including all major components of a LTE network. The actual 3GPP mobile network includes the elements from the user equipment [UE] to the packet gateway [P-GW].

The radio-based IP traffic between the UE and the eNB is encrypted according 3GPP standards. Between eNB, S-GW, P-GW user IP packets as well as control plane packets are encapsulated in 3GPP-specific tunnels. In some mobile carrier networks the 3GPP specific tunnels between eNB and S-GW are even additionally IPsec-encrypted. For more details see [[TS.29.281](#)] and [[TS.29.274](#)].

Service function chains act on user plane IP traffic only. But the way these act on user traffic may depend on subscriber, service or network specific control plane metadata.

2.2. Overview of mobile service chains

The original user IP packet, including the Source-IP-Address (S-IP) of the UE and the Destination-IP-Address (D-IP) of the addressed application platform, leaves the Packet Gateway from the mobile network via the so-called Gi-interface (3G service, e.g., UMTS) respectively SGi-interface (4G service, e.g., LTE). Between this (S)Gi-interface and the actual application platform the user generated upstream IP packets and the corresponding downstream IP packets are typically forced to pass an ordered set of service functions, loosely called a service function chain (SFC).

Figure 3: Typical service chain topology.

Service functions typically observe, alter or even terminate and reestablish application session flows between mobile user equipment and application platforms. Control plane metadata supporting policy based traffic handling may be linked to individual service functions SFn. Because in Figure 3 the P-GW classifies service chains, we consider the P-GW as a component of the service chaining environment.

2.3. The most common classification scheme

Mobile user equipment like smartphones, tablets or other mobile devices address use Access Point Names (APNs) to address a service network or service platform. APNs are DNS host names and comparable to FQDN host names. While a FQDN refers to an Internet IP address, an APN (loosely speaking) specifies a P-GW IP address. These APNs are used to distinguish certain user groups and their traffic, e.g., there can be an APN for a mobile service offered to the general public while enterprise customers get their own APN. For APN details see [TS.23.003].

Operators often associate a designated VLAN-ID with an APN. A VLAN-ID n then may classify the service function chain n (SFC n) related to an application platform n (Appl. n), as shown in the following Figure 4.




Figure 4: Association of a service chain to an application platform.

Examples for an APN are, e.g.:

APN:	web.vodafone.de
User Name:	not required
Password:	not required

Table 1: Example APN for Vodafone Germany

+-----+-----+	
APN:	internet.telekom
User Name:	t-mobile
Password:	tm
+-----+-----+	

Table 2: Example APN for Deutsche Telekom/T-Mobile

2.4. More sophisticated classification schemes

More sophisticated classifications are feasible using metadata that is UE related, subscriber and service related, as well as network related metadata. Typical metadata and their sources are:

UE: terminal type (e.g., HTC one); IMSI (country, carrier, user);

GTP tunnel endpoint: eNB-Identifier; time;

PCRF: subscriber info; APN (service name); QoS; policy rules.

Mobile operator defined subscriber, service or network specific policies are typically encoded in the 3GPP-based "policy and charging rules function" (PCRF), see [TS.23.203]. For instance, a PCRF may encode the rules that apply to pre-paid and post-paid users, users with a classification of gold, silver, or bronze, or even as detailed as describing rules that apply to "gold users, wishing to download a video file, while these subscribers are subjected to a fair-usage policy". It is up to the mobile service providers to encode the precise mappings between its subscriber classes and the associated service chains.

The Traffic Detection Function (TDF) is part of the 3GPP PCC (Policy and Charging Architecture, [TS.23.203]). Such a TDF inspects the user traffic after leaving the PGW (see Figure 4). The TDF can be used to classify traffic originating from an APN into more detailed services. This could be used to classify traffic into different Service Functions.

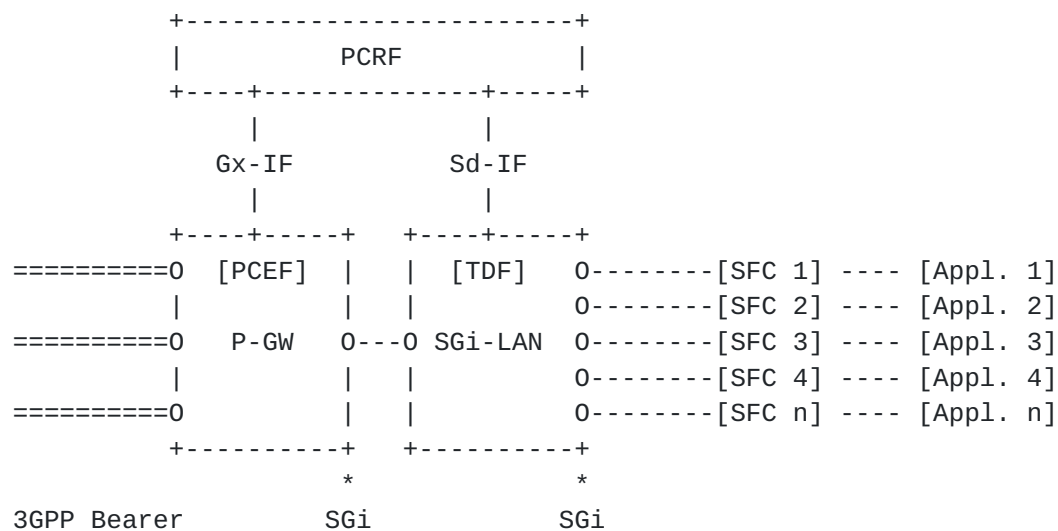


Figure 5: 3GPP Traffic Detection Function (TDF) for classification.

The TDF will typically observe the traffic on all layers. On application start and stop the TDF provides feedback to the PCRF for further actions to be taken on a particular flow. The PCRF can request that the TDF apply application and detection controls to application flows including charging and usage monitoring. The TDF can also act without any interaction with the PCRF taking care of gating (firewalling), traffic redirection, bandwidth management or charging.

3. Example use cases specific to mobile networks

Because HTTP via TCP port 80 (or TCP port 443 for HTTPS) is by far the most common Internet traffic class, we discuss two typical examples of an associated service function chaining model in some more detail.

The models presented below are simplified compared to real life service function chain implementations because we do not discuss differentiated traffic handling based on different subscriber-specific service level agreements and price plans or even actual network load conditions.

3.1. Service chain model for Internet HTTP services

With the increase of Internet traffic in mobile networks mobile operators have started to introduce Performance Enhancement Proxies (PEPs) to optimize network resource utilization. PEPs are more or less integrated platforms that ensure the best possible Quality of Experience (QoE). Their service functions include but are not limited to Deep Packet Inspection (DPI), web and video optimizations,

subscriber and service policy controlled dynamic network adaption, analytics and management support.

A simple service function chain model for mobile Internet upstream and downstream traffic is shown in Figure 6 below. The function chain includes Load Balancers (LB), which split HTTP over TCP port 80 away from the rest of the internet traffic. Beside basic web content, this traffic class includes more and more video. To act on this traffic type we force this traffic to pass Performance Enhancement Proxies. The firewall function (FW) protects the carrier network from the outside and Network Address Translation (NAT) maps the private IP address space dedicated to user equipment to a public IP address.




Figure 6: Service function chain for HTTP traffic over TCP port 80.

The first application in the service chain caches web content to help reduce Round Trip Times (RTT) and therefore contribute to improved web page load times. This is generally more important for mobile service providers than reducing Internet peering costs. Similar arguments hold for cached video content. We avoid potential large jitter imported from the Internet.

An example for non HTTP:80 traffic in Figure 6 is UDP-encapsulated IPsec traffic, which is dedicated to port 10000. Note that in a real environment not only port 80 but for example additional traffic via port 8080, 25 for SMTP, 110 for POP3 or 143 for IMAP may be forced to pass a service chain.

A second application is video optimization. Video content from the Internet may not fit in the size of mobile device displays or simply would overload the mobile network when used natively. Therefore mobile operators adapt internet-based video content to ensure the best Quality of Experience.

Video content optimization very often is also an additional premium-related component in service offers and price plans.

Our PEP environment for video optimization consists of three basic service functions shown in Figure 7: a steering proxy which is able to redirect HTTP traffic, a DPI-based controller which checks for video content and an optimizer which transcodes videos to an appropriate format on the fly in real time.

[PEP for video] ==>> [Steering Proxy]---[DPI Contr.]---[Optimizer]

Figure 7: Service functions required for video optimization.

In Figure 8 we show the detailed HTTP flows and their redirection in some more detail. The intention here is to show every elementary functional step in a chain as a separate physical or virtualized item, but this state diagram must not necessarily reflect every existing vendor-specific proprietary implementation.

```
[UE]----[Steering Proxy]----[DPI Contr.]----[Optimizer]----[Content]

|-- HTTP GET ->|----- HTTP GET ----->|
                |<----- HTTP Response -----|
                |-- Is it Video? ->|
                |<-- Video found --|

|<--- HTTP ----|
      Redirect

|-- HTTP GET ->|-----HTTP GET ----->|
                                     |-- HTTP GET -->|
                                     Video
                                     |<--- HTTP ----|
                                     Response
                                     Orig. Video

                                     {Optimize}

                |<----- HTTP Response -----|
                Transcoded Video

                |-- Is it Video? ->|
                |<-- Video found --|

|<--- HTTP ----|
      Response
Transcoded Video
```

Figure 8: Flow diagram between UE and video optimization PEP.

In such an application scenario one can have reclassification or off-loading on the fly.

Assume a video streams within a 4G LTE radio cell. The video optimizer would then apply a transcoding scheme appropriate to the

abilities of the 4G network. If one is now leaving the 4G cell and entering a 3G radio cell, the network conditions will most probably become different and the video optimizer has to use another transcoding scheme to keep a certain QoE. This requires that the video optimizer service function is aware of the Radio Access Technology (RAT) in use. One may transfer RAT type from the P-GW (or GGSN in case of 3G traffic) via an AAA Proxy to the service function chain. The RAT information will then be embedded in an appropriate Radius message. Other 3GPP steering mechanisms may apply as well.

If for example the 4G network has sufficient bandwidth, one could also think of another, different use case. The rule could be that only 3G video streams are forced to pass the video optimizer while all 4G video traffic will be bypassed.

Additionally, network utilization information can be used to trigger the behavior of the service function. The degree of video compression applied could depend on the actual current network load.

Last not least the behavior of the video optimizer service function (or any other service function) could additionally depend on the user-specific service contract (price plan, gold, silver, bronze) or on individual on demand requests.

3.1.1. Weaknesses of current approaches

This use case model highlights the weakness of current service deployments in the areas of traffic selection, reclassification, and multi-vendor support. Traffic in this example is classified after the P-GW and separated into separate flows based on whether it is (in this example) TCP traffic destined to port 80. This classification could be done by the load balancer if it initiates the service chain selection, or the traffic can be reclassified at the load balancer if it the traffic already embedded in a service chain (e.g., when combined with other functions such as the TCP optimization in the following use case). Multi-vendor support is needed because every element in the use case can be provided by a different vendor: load-balancer, http proxy, firewall, and NAT.

3.2. Service chain for TCP optimization

The essential parameters influencing TCP behavior are latency, packet loss and available throughput.

Content servers are mostly attached to fixed networks. These are characterized by high bandwidth and low latency. Therefore content servers often experience large TCP window sizes. In fixed networks,

end-to-end TCP window size mismatches do not have that much negative impact on data flows.

On the other hand, mobile networks show a different behavior. While the (S)Gi-side of the network typically exhibits low latency, low packet loss, high bandwidth and minimal congestion, the Radio Access Network (RAN) tends to have higher latency, packet loss, and congestion. Therefore mobile devices normally experience much smaller TCP window sizes.

One way to mitigate these different environments, i.e., the LAN and the mobile wireless part, is to use split TCP. However, this leads to the case that the mobile wireless part can experience a different TCP window size than the fixed LAN part.

In mobile networks, these TCP window size mismatches may result in poor end-to-end throughput and bad user experience.

Therefore mobile operators very often use TCP optimization proxies in the data path. These proxies monitor latency and throughput real-time and dynamically optimize TCP parameters for each TCP connection to ensure a better transmission behavior.

A rudimentary service chain setup for TCP optimization is shown in Figure 9. Examples of non TCP flows are UDP/RTP voice or video traffic.




Figure 9: Optimizing TCP parameterization in a mobile network.

3.2.1. Weaknesses of current approaches

This use case highlights weaknesses of current service deployments in the areas of traffic selection, reclassification, and multi-vendor support as in the previous use case presented in [Section 3.1](#).

3.3. HTTP header enrichment in mobile networks

In legacy mobile networks WAP (Wireless Application Protocol) gateways mediated between traditional mobile phones and the Internet translating HTML web content into a WML (Wireless Markup Language) and vice versa. By functionality, WAP-GWs fit also in the SFC category.

Traditionally WAP-GWs use HTTP header enrichment to insert subscriber related data into WAP and HTTP request headers in real time. These data were (are) used to identify and charge subscribers on third party web sites.

Of today 3G and 4G mobile networks HTTP header enrichment is done by the GGSN/P-GW or a dedicated transparent HTTP optimizer as most of the data traffic on a mobile network no longer passes a WAP-GW.

Information typically added to the header includes:

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

4. Remarks on QoS in mobile networks

As indicated in Figure 3, service functions may be linked to the control plane to take care of additional subscriber or service related metadata. In many cases the source of metadata would be the PCRF and the link would be a Gx or Diameter-based Sd interface. Diameter is specified in [[RFC6733](#)] and Gx in [3GPP].

Service functions within the SGi/Gi-LAN are less focused on the explicit QoS steering of the actual mobile wireless network. QoS in mobile networks is based on the 3GPP "Bearer" concept. A Bearer is the essential traffic separation element enabling traffic separation according different QoS settings and represents the logical transmission path between the User Equipment (UE) and the Packet Gateway (P-GW).

5. Weaknesses of current implementations

In many operator environments every new service introduction can result in a further dedicated (S)Gi-LAN service chain, because service chaining has been deployed historically in an ad hoc manner. It typically requires placement of new functions in the operator's data center, changing the actual wiring to include any new or change service function, configuration of the functions and network equipment, and finally testing of the new configuration to ensure that everything has been properly setup.

5.1. Granularity of the classification scheme

Often the coarse grained classification according to APNs is not fine enough to uniquely select a service function chain or a processing scheme within a service function chain required to support the typical user-, service- or network- related policies which the operator likes to apply to a specific user plane flow.

It is very likely that an APN, such as shown in [Section 2.3](#), is carrying an extremely diverse set of user traffic. This can range from HTTP web traffic to real-time traffic.

5.2. Service chain implementations

In many carrier networks service chain LANs grow incrementally according product introductions or product modifications. This very often ends in a mix of static IP links, policy based routing or individual VRF implementations etc. to enforce the required sequence of service functions. Major weak points seen in many carrier networks are:

- o Very static service chain instances, hard-wired on the network layer leads to no flexibility with respect to reusing, adding, and removing service nodes and reprogramming service chains.
- o Evolutionary grown "handcrafted" connectivity models require high complexity to manage or maintain.
- o Basic implementation paradigm is based on APNs (that is service types) only, which requires individual injections of context-related metadata to obtain granularity down to user/service level.
- o There is currently no natural (or standardized) information exchange on network status between services and the network, complicating management of network resources based on subscriber profiles.

- o It is currently practically impossible to implement an automated service provisioning and delivery platform.
- o Scaling up flows or service chains with service or subscriber related metadata is extremely difficulty.

6. Operator requirements

Mobile operators use service function chains to enable and optimize service delivery, offer network related customer services, optimize network behavior or protect networks against attacks and ensure privacy. Service function chains are essential to their business. Without these, mobile operators are not able to deliver the necessary and contracted Quality of Experience or even certain products to their customers.

6.1. Simplicity of service function chain instantiation

Because product development cycles are very fast in mobile markets, mobile operators are asking for service chaining environments which allow them to instantly create or modify service chains in a very flexible and very simple way. The creation of service chains should be embedded in the business and operation support layers of the company and on an abstract functional level, independent of any network underlay. No knowledge about internetworking technology should be required at all. The mapping of the functional model of a service chain onto the actual underlay network should be done by a provisioning software package similar to that shown in Figure 10. Details of the architecture and design are the subject of forthcoming standards and proprietary implementation details.

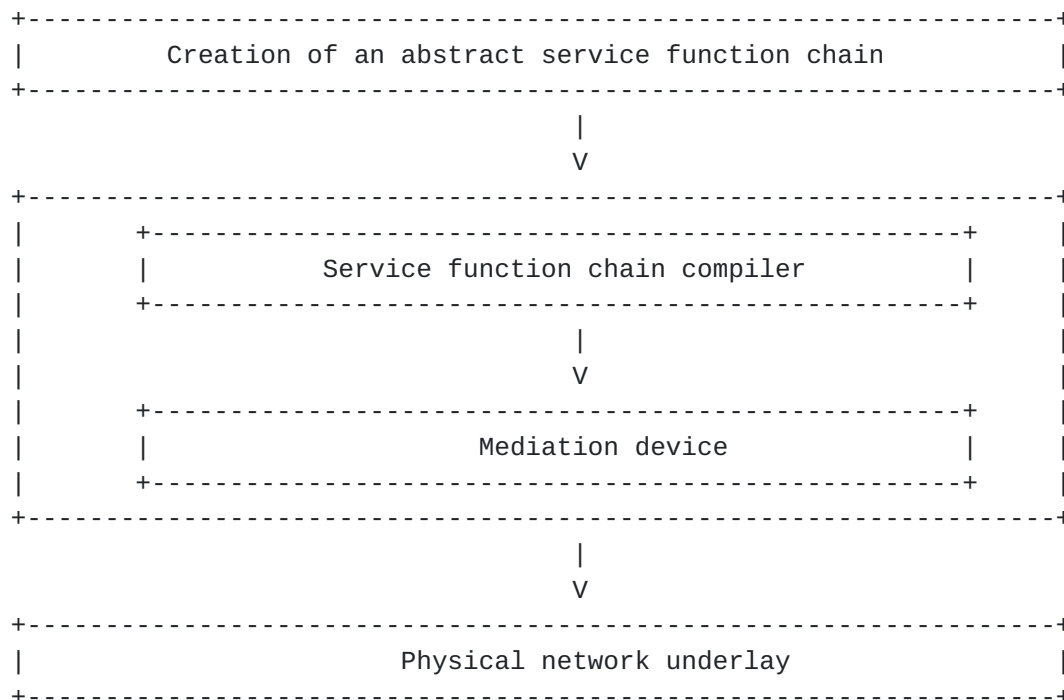


Figure 10: Creation and provisioning system for service function chains.

Service functions can be physical or virtualized. In the near future the majority of mobile service functions will typically reside in the local cloud computing environment of a mobile core location.

Nevertheless, the architecture and design should allow and support also remote service functions if applicable.

6.2. Extensions

A service function chain should be generalized by a directed graph where the vertices (nodes) represent an elementary service function. This model allows branching conditions at the vertices. Branching in a graph could then be triggered by typical 3GPP specified mobile metadata (see [Section 2.3](#)) and allow for more sophisticated steering methods in a service chain. Typically this data will be injected by the mobile control plane, commonly by the PCRF via a Diameter-based 3GPP Sd interface.

Service chain behavior and output should additionally depend on actual network conditions. For example, the selection of a video compression format could depend on the actual load of the mobile network segment a mobile user is attached to. That is to say that classification of flows may allow very dynamic inputs while specification of such inputs from a 3GPP network will need to be done by 3GPP or another standards body.

Although necessary metadata can be obtained from the PCRF, to avoid Diameter signaling storms in the (S)Gi-LAN, individual service functions should probably not be attached individually to the control plane. A mechanism where such metadata is carried by a metadata header can reduce requests to the PCRF, provided these extensions do not increase the original payload size too much.

6.3. Delimitations

A clear separation between service chaining functionality and 3GPP bearer handling is required. This may be subject of forthcoming studies.

7. Security Considerations

TBD.

8. IANA Considerations

This document has no actions for IANA.

9. Acknowledgments

We thank Peter Bosch, Carlos Correia, Dave Dolson, Linda Dunbar, Alla Goldner, Wim Hendericks, Konstantin Livanos, Praveen Muley, Ron Parker, and Nirav Salot for valuable discussions and contributions.

10. References

10.1. Normative References

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

10.2. Informative References

[RFC6459] Korhonen, J., Soininen, J., Patil, B., Savolainen, T., Bajko, G., and K. Iisakkila, "IPv6 in 3rd Generation Partnership Project (3GPP) Evolved Packet System (EPS)", [RFC 6459](#), January 2012.

[RFC6733] Fajardo, V., Arkko, J., Loughney, J., and G. Zorn, "Diameter Base Protocol", [RFC 6733](#), October 2012.

[TS.23.003] "Numbering, addressing and identification", 3GPP TS 23.003 12.1.0, December 2013.

[TS.23.203]

"Policy and charging control architecture", 3GPP TS 23.203 12.3.0, December 2013.

[TS.23.401]

"General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access", 3GPP TS 23.401 12.3.0, December 2013.

[TS.29.061]

"Interworking between the Public Land Mobile Network (PLMN) supporting packet based services and Packet Data Networks (PDN)", 3GPP TS 29.061 12.4.0, December 2013.

[TS.29.274]

"3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C); Stage 3", 3GPP TS 29.274 12.3.0, December 2013.

[TS.29.281]

"General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U)", 3GPP TS 29.281 11.6.0, March 2013.

Authors' Addresses

Walter Haeffner
Vodafone
Vodafone D2 GmbH
Ferdinand-Braun-Platz 1
Duesseldorf 40549
DE

Phone: +49 (0)172 663 7184
Email: walter.haeffner@vodafone.com

Jeffrey Napper
Cisco Systems
Cisco Systems, Inc.
Haarlerbergweg 13-19
Amsterdam 1101 CH
NL

Email: jenapper@cisco.com

Martin Stiernerling
NEC
NEC Europe Ltd.
Kurfuersten-Anlage 36
Heidelberg 69181
DE

Email: mls.ietf@gmail.com

Diego R. Lopez
Telefonica I+D
Don Ramon de la Cruz, 82
Madrid 28006
ES

Phone: +34 913 129 041
Email: diego@tid.es

Jim Uttaro
AT&T
200 South Laurel Ave
Middletown, NJ 07748
US

Email: uttaro@att.com

