

SPRING and DMM
Internet-Draft
Intended status: Standards Track
Expires: February 16, 2020

P. Camarillo, Ed.
C. Filsfils
Cisco Systems, Inc.
H. Elmalaky, Ed.
Individual
S. Matsushima
SoftBank
D. Voyer
Bell Canada
A. Cui
AT&T
B. Peirens
Proximus
August 15, 2019

SRv6 Mobility Use-Cases
draft-camarilloelmalaky-springdmm-srv6-mob-usecases-02

Abstract

This document describes the SRv6 use-cases in the mobile network in association with different mobile generations (3G, 4G, and 5G). It also highlights potential interworking with SR-MPLS in relevant use-cases.

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](#)].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on February 16, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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	3
2.	Terminology	4
3.	Use-cases	5
3.1.	SP Network Simplification use-cases	5
3.1.1.	Radio-core Handoff	5
3.1.1.1.	Radio-transport programmability	5
3.1.1.2.	User-plane state transfer, offload, and mutation	6
3.1.1.3.	Rip-n-replace of GTP with SRv6	8
3.1.2.	End-to-end network slicing [N3, N9, N6 and transport]	9
3.1.3.	GiLAN Service Programming [N6 and N9]	9
3.1.3.1.	Service Programming on Gi-LAN for 3G/4G [SGi] . .	10
3.1.3.2.	Service Programming for 5G [N6 and N9]	10
3.1.4.	ID-Location Isolation at anchors	10
3.2.	New mobility use-cases	10
3.2.1.	eMBB (Enhanced Mobile Broadband)	10
3.2.1.1.	Fixed/Mobile Convergence (HA, FWA & WA)	10
3.2.1.2.	Mobile Enforced SD-WAN	11
3.2.2.	mMTC (massive Machine Type Communications)	11
3.2.2.1.	Stationary IoT Devices (industrial applications)	11
3.2.3.	URLLC (Ultra Reliable Low Latency Communications) . .	12
4.	Work in progress	12
5.	Acknowledgements	12
6.	References	12
6.1.	Normative References	12
6.2.	Informative References	12
	Authors' Addresses	13

1. Introduction

4G/LTE mobile networks are complex and the use cases that 5G has been architected to address, introduce new requirements and additional complexity to both the RAN and the mobile core. The current architecture employs the GPRS tunneling protocol (GTP) as the primary vehicle for user plane interconnect in the RAN and 5GC. GTP is currently used in two contexts, from the RAN to the first anchor point; the S-PGW/UPF (S1-U/N3 interface) and for inter S-PGW/UPF connectivity (S5-S8/N9 interface). While the tunnels themselves do not impose significant state beyond that needed, they do have a significant control plane setup component and are a potential target for network layering.

Segment Routing [[I-D.ietf-spring-segment-routing](#)] is a network architecture that simplifies networks by removing state from the network infrastructure, creating a scalable SDN architecture for overlays (VPNs), underlay (SLA, Traffic Engineering, FRR) and service programming (GiLAN). The IPv6 instantiation -also known as SRv6 [[I-D.ietf-spring-srv6-network-programming](#)]- takes this even further with the introduction of the Network Programming concept, allowing to bind segments to any kind of VNF anywhere in the network -from private DCs to public cloud services.

Segment routing embodies a number of potentially useful properties for consideration in a 4G/5G mobile networking context:

1.- Direct manipulation of path routing by the head-end

SRv6 provides the ability to direct traffic through an arbitrary path without the imposition of path state in the network or requiring a separate signaling system. It does this without signaling by encoding the path state in the packet header. This means the path head-end can instantly fulfill changes to a path by simply changing the header encoded information.

This capability has numerous applications as far as networking in general (traffic engineering, policy routing etc.), but has additional applicability to mobile networks:

- o The ability of a path head-end to manipulate the intermediate hops in a path can be exploited for end system mobility, the penultimate hop simply becomes the "care of" address.
- o The ability of path head-end to imply an asymmetric return path for a specific forwarding equivalence class (FEC).

- o Densification in the radio topology embodied in concepts like coordinated multipoint and multi-connectivity require the instantaneous redirection of traffic from the coordinating radio controller to any of several base stations. This is critical to exploit ephemeral "rich paths" that 4G & 5G radio technologies depend upon to achieve high rates of information transfer.

2.- Network programmability

The ability to bind segments to network functions provides an increased level of abstraction in service delivery combined with a practical realization. This would have applications in the GiLAN/N6, combined with the ability to specify a path from the head-end as applications in the GiLAN/N6.

3.- Overall simplification of the control plane

As noted previously SRv6 dispenses with a signaling system. This has obvious benefits as a simplification to overall network operation, but may have additional benefits in the "signaling rich" environment of mobile networks.

This memo serves to critically explore the applicability of SRv6 to 4G/5G mobile networks. It does that via an exploration of how SRv6 can simplify current mobile network architecture to improve the status quo of eMBB operation, and then delves into the new use cases that 5G is targeted towards.

2. Terminology

This document focuses on the use-cases, and it's associated terminologies. The full list of terminologies exists in [\[I-D.ietf-spring-srv6-network-programming\]](#).

In this document we focus on the 5G systems architecture, as specified in [\[TS.23501\]](#). This document also refers to 3G and 4G networks as specified in [\[TS.23002\]](#).

The uplink/upstream traffic is the traffic originated at the UE, while the downlink/downstream traffic is traffic destined towards the UE.

3. Use-cases

Use-cases have been classified into multiple categories depending on their fit into the mobile-network domain (Radio, Transport, Core) or mobile network generation (3G/4G, or 5G).

3.1. SP Network Simplification use-cases

3.1.1. Radio-core Handoff

3.1.1.1. Radio-transport programmability

Advances in radio technology, the deployment of new spectrum for 5G and the quest for ever increasing spectral efficiency results in increasingly complex RAN and air interface topologies. The result is that the RAN end of a GTP tunnel may appear as a single end point to the core network, but the actual realization is substantially more complex.

Modern radio scheduling is increasingly focused on using techniques such as MIMO to multiply the instantaneous bandwidth available for information transfer for a given unit of spectrum. A "rich path" can be very ephemeral so any latency between path measurement and initiating data transfer to a UE can be parasitic in the overall system efficiency.

3.1.1.1.1. Multi-connectivity and coordinated multi-point

There are multiple scenarios where a UE can be associated with more than one antenna and the associated spectrum:

Coordinated multipoint (CoMP) involves a UE associated with multiple geographically distributed antennas serving a common block of spectrum, and the radio controller selecting the best antenna at any given time. The other antennas being quiescent in the sector occupied by the UE at the time of transmission to avoid overlap. This can be in the context of an RRC/RLC split (F1-U interface) or a Phy Hi/Phy Lo split (F2-U interface).

Multi-connectivity can see a UE associated with multiple antennas each serving different spectral allocations. Applications include offload from a macro cell to a small cell. The possibility of simultaneous transfer from multiple antennas also exists. And again this can be on the basis of an F1 or F2 split in the RAN.

In both cases the radio controller is required to be able to instantly redirect traffic based on current radio measurements to any of a constellation of antennas serving a given UE.

3.1.1.1.2. Fronthaul

Modern radio systems have been deconstructed in order to drive efficiency across a variety of metrics. In essence various stages of waveform construction have been abstracted and exposed on interfaces as part of the 5G RAN architecture. In the most simplest form it allows putting functionality where it is easy to service, such as the equipment at the bottom of the tower rather than the top. In a rich radio connectivity context it permits co-location of radio scheduling and waveform generation which drives spectral efficiency, but where applied also results in significant multiples of bandwidth, and very tight jitter and delay requirements. The current specification for the F2-U or e(CPRI) packet interface has a maximum latency of 75 usec and correspondingly tight jitter requirements.

3.1.1.2. User-plane state transfer, offload, and mutation

A proper session handoff between radio, transport, and mobile-core requires storing/recalling user-plane session state on multiple levels. The use of SRv6 reduces the number of states needed in the network nodes by mapping the UE session state into IPv6 SID (Segment IDs) in SRH. Furthermore, mutation of SID-lists shall enable SMF to program data-paths (handling-state) and policies (serving-state) on per-subscriber / per-application level.

That session state can be broken down into two categories:

1.- Handling state: Who is the session handler?

- o A 1-to-1 mapping between GTP tunnel (TEID) and S-PGW/UPF
- o Usually stored at load-balancers deployed ahead of S-PGW/UPF instances or embedded inside the S-PGW/UPF system.

2.- Serving state: What is the serving-policy associated with this session?

- o A 1-to-1 mapping between the UE and a specific policy to be enforced on the subscriber traffic.
- o The policy may include (but not limited to) the authorization & accounting profile, one or multiple QoS profiles, one or multiple service-chaining/programming profiles.
- o A 1-to-1 mapping between a UE session and it's stats-registers at the S-PGW/UPF.

- o It is typical that S-PGW/UPF may break down the service-state into sub-states reflecting groups of 5-tuple flows, or employ other techniques (ex. DPI, deep packet inspection) to break down the serving-state even further within the same 5-tuple flow.

The ability to transfer, offload, or mutate the user-plane state with no/minimum disruption to end-users is one of the most significant challenges facing the mobile network's scalability towards mMTC use-cases (The current GTP-U mandates a per-session tunnel creation & handling). Moreover, the direct 1-to-1 binding between UE session ID and Location affects the optimal-path selection, which is one of the most significant challenges facing URLLC use-cases in 5G.

The use of SRv6 shall simplify the state storage dramatically where a single SID-list embedded in the UE session packet can store the handling-state and the majority of the serving-state. SRv6 programmability and traffic-engineering shall allow an easy way to transfer, offload, or mutate that state.

3.1.1.2.1. State-offload:

Upstream state-offload:

the use of SRv6 shall allow the S-PGW/UPF anchor(s) to offload the load-balancing function from a dedicated load-balancer in mobile-core to be a standard function in packet-forwarding in transport network where any SR-aware node on the path between eNB and S-PGW/UPF can forward the UE session to the proper S-PGW/UPF handling instant by relying on the handling-state stored in the SID-list in each packet.

Downstream state-offload:

The L3 anchor (PGW/UPF) is the first node that handles the subscriber traffic in the downstream direction, depending on the policy associated with the subscriber traffic. The PGW/UPF may decide to hairpin the traffic through multiple application (service chain) before sending it towards the radio-network. This implies double packet-processing on PGW/UPF instant (50% penalty on the VNF useful throughput).

The use of SRv6 shall allow the PGW/UPF to impose a specific data-path on a group of 5-tuple flows without the need for hairpins all the traffic through PGW/UPF. Which means the PGW/UPF can offload the first packet processing towards another none-SR- node earlier in the downstream path (ex. Service-proxy, or packet inspector) as per specific service-pipeline policy.

Moreover, that offload-service can be programmed once the S-PGW/UPF terminate the subs-session on the upstream direction. Alternatively,

the offload-service can be programmed on-demand after the first few packets been hair-pinned through the PGW/UPF on the downstream path.

3.1.1.2.2. State-transfer:

Handling-state:

SRv6 shall enable the handling-state to be embedded in the data-flow as metadata (in a form of SID-list). This means that all load-balancing operations can be performed by any of the SR-aware intermediate nodes in a stateless fashion with a zero-state transfer at failure scenarios.

Serving-state:

Depending on the applied policy, a significant portion of the serving-state can be embedded in the data-flow as metadata (in a form of SID-list). This means that serving nodes (S-PGW/UPF) have a smaller amount of data to store/recall to serve the UE session.

3.1.1.2.3. State-mutation:

SRv6 provides a more natural way to mutate the handling-state and serving-state to follow the optimal data path or fulfill traffic-engineering constrain(s).

In contrast to the current limitation of mutating the state only at SGW (session L2-anchor point) or PGW (session L3-anchor point). SRv6 shall allow the state mutation on any authorized SR-aware node between radio and mobile core.

3.1.1.3. Rip-n-replace of GTP with SRv6

A possible mechanism to do an early-deployment of SRv6 is to keep the tunnel-nature of GTP but do a simple data-plane replacement of IP/UDP/GTP-U with SRv6 for specific PDU sessions. In this case, there is no session aggregation, and the SRv6 segment corresponding to the overlay creation now carries the TEID, QFI and RQI as part of the SID arguments.

In this use-case there is no subscriber-traffic integration with the underlay or service programming. There could be some integration but it is based on static policies and not configured via the currently existing mobility management.

This is an interworking mechanism that shall used for an early stage implementation with no changes to the N4 interface.

3.1.2. End-to-end network slicing [N3, N9, N6 and transport]

One of operator's main challenges is providing end-to-end network slicing, taking into consideration the RAN, the S-PGW/UPF and the VNFs in the GiLAN; but more importantly taking also into consideration the transport network.

SRv6 can help bridging the gap in between all of these since it integrates the overlay, underlay and service programming into a single protocol. End-to-end SR policies can be defined that span across the RAN, S-PGW/UPF and transport network, without requiring any stitching configuration at the domain boundaries. From an overlay perspective, it is clear that SRv6 can provide -if desired- isolation among different RAN or S-PGW/UPF nodes.

In the transport network, the SRv6 overlay can integrate with an existing SRv6 or SR-MPLS transport network to provide traffic engineering in the underlay network infrastructure. SR provides operators with a stateless mechanism to build network slices with different optimization objectives or constrains i.e. low-latency (uRLLC), resource isolation (disjointness), etc...

Also, SR provides mechanisms for in-band performance monitoring. This implies that the end-to-end network slice can react upon topology changes -that for example might change the low-latency path-.

3.1.3. GiLAN Service Programming [N6 and N9]

Service Programming, in coordination with SRv6 can be used for optimal placement of VNFs in the Gi-LAN of mobile operators for flawless VNF management and placement -DC resource utilization-.

SRv6 transparently integrates VNFs [[I-D.xuclad-spring-sr-service-programming](#)], in the same SR policy used for overlay creation and underlay control. The VNFs are cloud-infrastructure agnostic -can be hosted on a private DC or public cloud-, and there is no state per-flow or per-chain in the network infrastructure. This implies a huge flexibility for mobile operators. Note that VMs can be distributed in different tenants, or can be migrated while there is live traffic without any major manageability complexity, state to update in the network infrastructure or packet loss. Note also that in the case of network slicing, the VNFs can be shared across multiple slices or can be restricted to only a particular slice. This can be chosen on a per-VNF granularity.

In addition, SRv6 offers mechanisms to do VNF load-balancing and to convey additional flow information to stateless VNFs using the SRv6 SID arguments, by leveraging the network programming concept.

3.1.3.1. Service Programming on Gi-LAN for 3G/4G [SGi]

SRv6-based NFV provides an approach to optimally steer traffic through Gi-LAN network functions in 3G/4G networks.

The PGW can steer uplink traffic into a specific SR policy that contains as many segments as VNFs that the packet must traverse. The packet follows the path specified in the SR policy, traversing the set of VNFs before getting delivered to the external PDN -i.e. internet-.

3.1.3.2. Service Programming for 5G [N6 and N9]

In 5G networks SRv6 can offer NFV control, as done in the Gi-LAN for 3G-4G networks (N6 interface), but can also integrate the VNFs within the N9 interface. This means that we can have more flexibility regarding the distribution and association of the functions/VNFs/micro-services, and bring applications closer to the user, where they might be better located for the operator and improve the overall customer experience.

3.1.4. ID-Location Isolation at anchors

TBD

3.2. New mobility use-cases

3.2.1. eMBB (Enhanced Mobile Broadband)

3.2.1.1. Fixed/Mobile Convergence (HA, FWA & WA)

The end users of different access networks under control of the same service provider would obtain significant benefit if there is a tight integration for service delivery in between the mobile access network and the fixed network.

This is the example of a residential user that is accessing content from his mobile phone, and once he arrives home his phone automatically connects to his home wireless network provided whose connectivity is provided by the same operator. As per today, these networks have different architectures, with different control-planes and data-planes, and with different policy control and service management.

SRv6 helps uniting the gap in between different access networks by optimizing the data path in between hierarchical networks and directly adding an SR policy that spans from the mobile packet core up to the broadband network BNG. Such capability will simplify the delivery of fixed-services on top of wireless infrastructure. It will also enable the simultaneous use of wireless and fixed connections towards end-user.

3.2.1.2. Mobile Enforced SD-WAN

TBD

3.2.2. mMTC (massive Machine Type Communications)

3.2.2.1. Stationary IoT Devices (industrial applications)

There are many types of IoT devices, ranging from connected cars to massive machine type devices like meter readers, which are stationary. One of these examples is electricity meters. These devices are static and might only attach to other gNBs due to changing RF conditions.

Massive machine type devices is projected to grow to 10's of billions in operator networks in the next few years. However, the traditional 3GPP GTP tunnel/bearer based connection-oriented architecture does not scale for billions of IoT devices due to the amount of signaling overhead associated with GTP tunnel setup/tear-down and the UE context information maintained at various parts of the mobile network.

Unlike smart devices, electric meters never move and each generates low RPU for carriers. For this reason, to efficiently support the massive machine type of stationary IoT devices, a simpler and more scalable control and user plane architecture is needed that can reduce the amount of signaling overhead and the UE context information kept in the network. This new architecture will need to work across all types of access technologies to improve adaptability to future RAT networks.

SRv6 can help improve scalability in the RAN, transport, and packet core networks significantly by removing GTP tunnels for each individual stationary IoT device, and replacing by the aggregated SRv6 route information for all the similar stationary IoT devices. For instance, at the eNB/gNB, only the first electric meter device for an electric company needs the SRv6 route set up procedure, which has one SRv6 look up table entry associated with it. No subsequent SRv6 route set up procedures and no additional SRv6 table entries for the succeeding electric meters are needed at the same eNB/gNB. This

effectively reduces the signaling overhead and UE context overhead by $(1-1/N)\%$ (where N is the number of the electric company meter readers in the same eNB/gNB). In the case of RAN virtualization with an aggregated vBBU for many cell sites, the reduction of the signaling and UE context overhead will be greater since N is a much bigger number.

The significant reduction of the signalling overhead and UE context overhead can be translated to the cost reduction of running operators' wireless network. In addition, this new architecture using SRv6 allows flexible service edge treatment, service chaining, such as billing, TE or other capabilities.

3.2.3. URLLC (Ultra Reliable Low Latency Communications)

TBD

4. Work in progress

- o Use of SRv6 in optimizing interface (reference N4 as defined by 3GPP xxx r16) between control-plane and user-plane.
- o Security implications & benefits of SRv6 in mobile networks.

5. Acknowledgements

We would like to thank Francois Clad, Darren Dukes, Zafar Ali, Peter Bosch, Simon Spraggs and Tom Anschutz for their help.

6. References

6.1. 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>>.

[TS.23501] 3GPP, "System Architecture for the 5G System", 3GPP TS 23.501 15.2.0, June 2018.

6.2. Informative References

[I-D.ietf-spring-segment-routing]

Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-15](#) (work in progress), January 2018.

[I-D.ietf-spring-srv6-network-programming]

Filsfils, C., Camarillo, P., Leddy, J., daniel.voyer@bell.ca, d., Matsushima, S., and Z. Li, "SRv6 Network Programming", [draft-ietf-spring-srv6-network-programming-01](#) (work in progress), July 2019.

[I-D.xuclad-spring-sr-service-programming]

Clad, F., Xu, X., Filsfils, C., daniel.bernier@bell.ca, d., Li, C., Decraene, B., Ma, S., Yadlapalli, C., Henderickx, W., and S. Salsano, "Service Programming with Segment Routing", [draft-xuclad-spring-sr-service-programming-02](#) (work in progress), April 2019.

[TS.23002]

3GPP, "Network Architecture", 3GPP TS 23.23002 15.0.0, March 2018.

Authors' Addresses

Pablo Camarillo Garvia (editor)
Cisco Systems, Inc.
Spain

Email: pcamaril@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Belgium

Email: cf@cisco.com

Hani Elmalky (editor)
Individual
United States of America

Email: hani.elmalky@gmail.com

Satoru Matsushima
SoftBank
1-9-1, Higashi-Shimbashi, Minato-Ku
Tokyo 105-7322
Japan

Email: satoru.matsushima@g.softbank.co.jp

Daniel Voyer
Bell Canada
Canada

Email: daniel.voyer@bell.ca

Anna Cui
AT&T
United States of America

Email: zc1294@att.com

Bart Peirens
Proximus
Belgium

Email: bart.peirens@proximus.com

