

RTG Working Group
Internet Draft
Intended status: Informational
Expires: October 12, 2024

K. Majumdar
Microsoft
U. Chunduri
Intel
L. Dunbar
Futurewei
September 12, 2023

Extension of Transport Aware Mobility in Data Network
draft-mcd-rtgwg-extension-tn-aware-mobility-08

Abstract

The existing Transport Network Aware Mobility for 5G (TN-AWARE-MOBILITY) specifies a framework for mapping the 5G mobile systems' Slice and Service Types (SSTs) to corresponding underlying network paths to ensure the desired QoS for the services. The focus of (TN-AWARE-MOBILITY) is limited to the 3GPP domain and doesn't cover the network for user plane traffic between the UPFs and the services hosted in data centers.

This document describes the mechanisms to achieve the same QoS for the mobility traffic from the 3GPP domain to the service destinations over the IP Data Networks. Extending the QoS guarantee for the SSTs to the data centers where the services are hosted becomes necessary to maintain the end-to-end QoS for data plane traffic between UEs and their corresponding service destinations.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on April 23, 2021.

Copyright Notice

Copyright (c) 2023 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
2. Conventions used in this document.....	3
3. Extending 5G SSTs QoS to Data Network.....	4
4. Mobility Packets Transition to the Data Network.....	5
5. Transport Network Characteristics Mapping to TE Paths.....	8
6. Extend TN Aware Mobility via BGP SR-TE Policy.....	9
6.1. TN Aware Mobility Policy Encoding.....	11
6.2. Operations of Applying the TN Aware Mobility Policy.....	12
7. TN Aware Mobility Policies Distribution via PCE Controller....	12
8. TN Aware Mobility Policies via NETCONF/RestConf/gRPC.....	13
9. Extend BGP FlowSpec for TN Aware Mobility.....	14
10. Mapping of TN Characteristics on SD-WAN Edge Node.....	14
11. IANA Considerations.....	16
12. Security Considerations.....	16
13. Contributors.....	16
14. References.....	16
14.1. Normative References.....	16

14.2. Informative References.....	17
15. Acknowledgments.....	18
Authors' Addresses.....	19

1. Introduction

The [TN-AWARE-MOBILITY] describes the mapping of 5G slices to IP or Layer 2 transport network for both control plane and data plane traffic across back-haul, mid-haul, and front-haul segments within the 3GPP domain [TS.23.501-3GPP]. While the 5G control plane traffic is end-to-end from UEs, gNB, to the 3GPP specified control functions (such as Session Management function, Access and Mobility Management Function, etc.), the data plane traffic management limited to UEs<-> gNB <-> UPFs (user plane functions) is not end-to-end.

As 5G services are processed by their corresponding servers in data centers through the IP network, achieving the desired QoS for the data plane traffic between gNBs <-> UPFs by mapping to the appropriate underlay IP networks is insufficient to ensure the end-to-end QoS. This draft describes the mechanisms for extending the 5G Service Slices mapping enforced within the 3GPP domain to the appropriate IP underlays from UPFs to 5G service destinations, including RSVP-TE[RFC3209], SRv6[RFC8986], MPLS-TE[RFC3812], and Preferred Path Routing (PPR) [RTG-PPR]. The goal is to ensure the end-to-end QoS guarantee for the 5G service slices.

Note: even though the 3GPP domain has underlay data networks connecting various 3GPP specified functions, the term Data Network in this document represents the IP data network outside the 3GPP domain, a.k.a. the N6 interface described in [UPlane-5G-ANA].

2. Conventions used in this document

BSID	- Binding SID
DC	- Data Center
DN	- Data Network [TS.23.501-3GPP]
EMBB	- enhanced Mobile Broadband [TS.23.501-3GPP]

gNB	- Next Generation NodeB [TS.23.501-3GPP]
GTP-U	- GPRS Tunneling Protocol - User Plane [TS.23.501-3GPP]
MIOT	- Massive IOT [TS.23.501-3GPP]
PECP	- Path Computation Element (PCE) Communication Protocol
SD-WAN	- Software-Defined Wide Area Network
SID	- Segment Identifier
SLA	- Service Layer Agreement
SST	- Slice and Service Types [TS.23.501-3GPP]
SR	- Segment Routing
SR-PCE	- SR Path Computation Element
UE	- User Equipment
UPF	- User Plane Function [TS.23.501-3GPP]
URLLC	- Ultra reliable and low latency communications [TS.23.501-3GPP]

3. Extending 5G SSTs QoS to Data Network

In actual deployment, a UPF can be connected to a C-PE node that ingresses to the IP network (i.e., the N6 Interface of the 3GPP domain) or embedded with the C-PE function. The C-PE node can be an SD-WAN [NET2CLOUD] edge node, L2/L3 VPN [RFC4364] edge node, or SR-TE [RFC8402] edge node. For the enterprise 5G scenario, the L3 aggregator function can be integrated with the UPFs for all the enterprise's 5G mobility connections.

The proposed mechanisms to extend the 5G SSTs QoS in the Data Network focuses on the following areas:

- a) The Mobility packets transition in and out from the UPFs to the C-PEs node maintaining the same transport path characteristics.
- b) On a C-PE node, based on the transport characteristics, use different methods of fetching TE (traffic engineered) path segments from the IP network controller and map the TE path segments to the packets from UEs.
- c) On an SD-WAN CE Node, based on the transport characteristics, mapping of UP packets to the secure and un-secure tunnel paths based on 5G service policies.

Figure 1 under Section 4 depicts data from UEs through the UPFs and C-PEs to the service destinations. In some cases, UPF is integrated with the C-PE functions.

4. Mobility Packets Transition to the Data Network

As a mobility packet transitioning from the UPF to the C-PE node, the GTP header is decapsulated, i.e., the UDP port number carried by the GTP header as specified by the [TN-AWARE-MOBILITY] is no longer present in the packet. Here are some approaches to classify the mobility packets when transitioning to the Data Network so that the appropriate policies can be applied to achieve the same QoS as the 3GPP domain:

A) When the UPF is connected to a C-PE via a direct link (Figure 1 below), the respective application controller needs to pass the desired QoS policy information to the Data Network's management system (DN Mgnt) and the 5G Management Plane (NSSMF) [TS.23.501-3GPP]. In Some deployments, the application controller, 5G NSSMF, and DN Mgnt belong to one operator. In some deployments, they are separate entities, thus requiring contract binding for the desired QoS.

The C-PE can encapsulate the mobility packet from the UPF with another outer IP header, with the destination address of the outer header set to the service destination hosted in a remote data center and the source UDP port number set as [TN-AWARE-MOBILITY]. Here is the packet illustration:

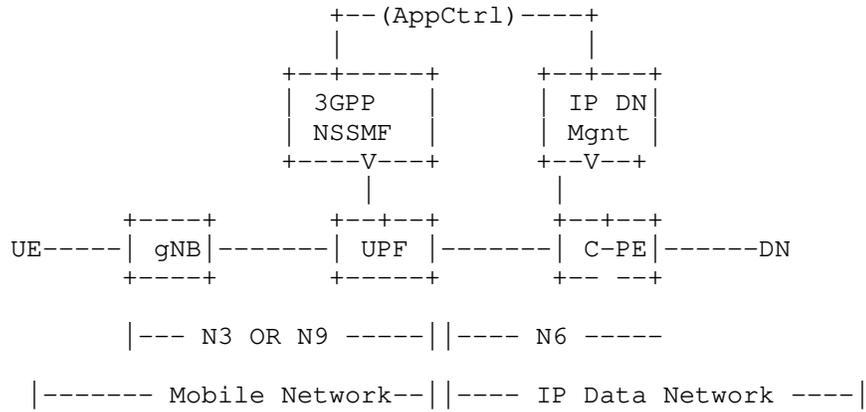
- From the UPF to the C-PE Node:
Inner IP Hdr (UE Packet) + Transport Hdr (Carrying UDP Src Port)
+ Outer IP (C-PE Node Address)

- From the C-PE to the UPF Node:

Outer IP (UPF Node Address) + Transport Hdr (Carrying UDP Src Port) + Inner IP Hdr (UE Packet)

B) When the 5G Core and the IP Data Network (N6 interface) are managed by one operator, a UPF can have the embedded C-PE function (Scenario B of Figure 1). Under this scenario, the original UDP Source Port information can be used to select the underlays for reaching the service destinations based on the local policy.

Scenario A:



Scenario B:

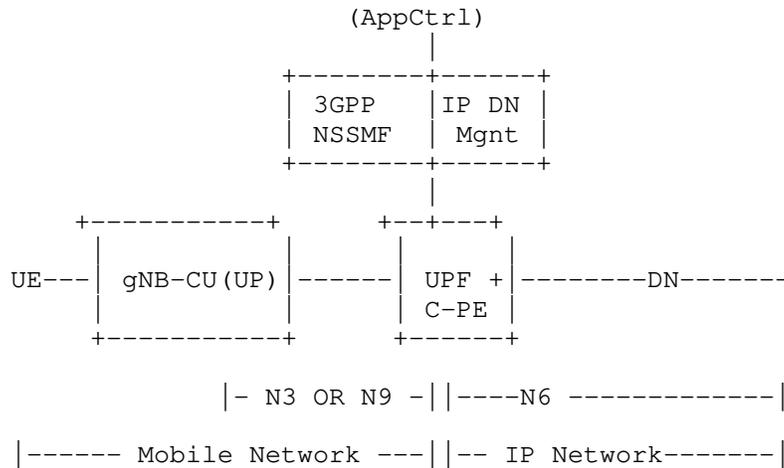
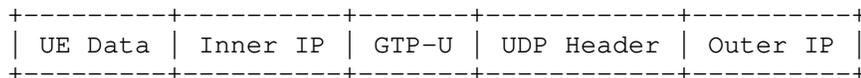


Figure 1: Mobile and IP Data Network for UE

Figure 2 illustrates the TN Aware Mobility packet format under scenario A.

- UE Packet format in the Mobile Network from gNB to the UPF:



- UE Packet format in the IP Network to the Ingress C-PE Node:

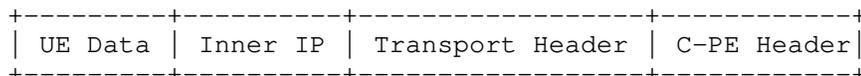


Figure 2: UE Packet Transition from Mobile to IP Network

The source port in the original UDP header indicates the TN Aware Mobility SST type.

5. Transport Network Characteristics Mapping to TE Paths

Within the 5G Core [TS.23.501-3GPP], UE traffic is carried by the GTP tunnels terminated by the UPFs. [TN-AWARE-MOBILITY] specifies using the GTP header's source UDP port number to indicate the desired transport network characteristics. The same transport network characteristics should be carried through the IP data network to the destinations to maintain the end-to-end SLA for the UE traffic.

3GPP specifies that the User Plane Function (UPF) decapsulates the GTP header and hands the inner payload to the N6 interface. As more and more deployments have integrated UPF with the ingress routing functions to the IP data network, it is relatively easy to carry the UDP port information from the GTP header into the IP data network, e.g.,

- Carry the UDP port number in an extra IP encapsulation (VxLAN, GENEVE, etc.) to the IP network, or,
- Assign an identifier in the header to the IP networks based on the UDP port number in the GTP header.

Here are some options in passing the corresponding policies:

Option 1: Assume the ingress C-PE is connected to a BGP SR-TE Controller through a BGP SR-TE Policy SAFI Session. Section 6 describes the mechanism to map mobility traffic to the BGP SR-TE Underlay Path Segments based on the mobility transport characteristics:

Option 2: Assume the ingress C-PE is connected to the PCE (Path Compute Element) Controller through a PCEP Session, which can pass the policy to map mobility traffic to the TE underlay paths [RFC9168] based on the mobility transport characteristics.

Option 3: Assume the ingress C-PE is connected to the TE Controller over NETCONF/Restconf/ Netconf or gRPC session. The existing mechanism would be used to download the TE Underlay Path Segments to the C-PE node based on the mobility transport characteristics.

Option 4: Assume the ingress C-PE is connected to the BGP SR FlowSpec Controller through a BGP FlowSpec session. The FlowSpec Redirect to indirection-id Extended community [FLOWSPEC-PATH-REDIRECT] can be used to map the Mobility Transport Path Characteristics to SR Traffic rules. [FLOWSPEC-TN-MOBILITY] specifies BGP FlowSpec extension to disseminate the policies from 5G mobile networks so that the 5G mobile systems slices and Service Types (SSTs) can be mapped to optimal underlying network paths in the data network.

6. Extend TN Aware Mobility via BGP SR-TE Policy

To integrate the 5G Core's transport aware mobility policies with BGP SR-TE Policy at the Ingress C-PE UPF, a class map needs to be defined to classify the incoming mobility traffic with different transport path characteristics.

Assume that the ingress C-PE (embedded in UPF) support the BGP SR-TE Policy SAFI [BGP-SR-TE-POLICY] and the BGP SR-TE Controller can compute the SR segments that can satisfy the SLA for the transport characteristics identified by the 5G UDP SRC Port Range based on the [TN-AWARE-MOBILITY] draft.

The figure below captures the overall topology and how to map the mobility traffic on the Ingress C-PE, which has the BGP SR-Policy SAFI connection with the BGP SR-TE Controller:

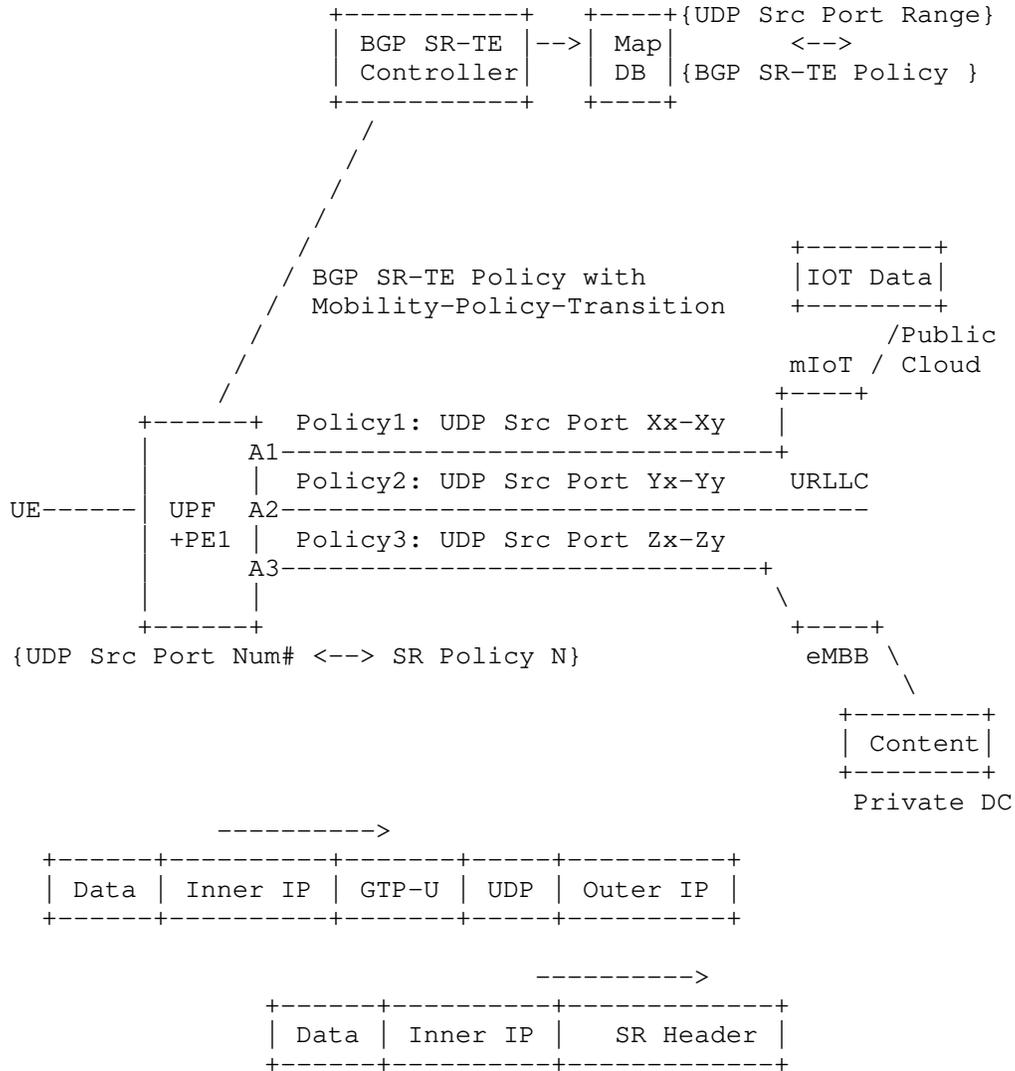


Figure 3: TN Aware Mobility Traffic Mapping to BGP SR-TE Policy Path

Note that, the SR Header in the figure above is shown as an illustrative purpose and the actual outgoing packet format is based on the BGP SR-TE mechanism (SR-MPLS or SRv6). If the BSID is not present with the BGP SR-TE Policy, the local ingress C-PE will map the incoming traffic to the best effort policy map path in the underlay.

6.1. TN Aware Mobility Policy Encoding

Mobility-Policy-Transition Sub-TLV specified here is to maintain the policy in the SR-TE network consistent with the 5G Core for the selective mobility traffic. The Mobility-Policy-Transition Sub-TLV is one of the Sub-TLVs under the Tunnel Type = 15 (SR Policy) under the SR-TE Policy NLRI [BGP-SR-TE-POLICY] sent to the relevant ingress C-PEs.

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>

Attributes:

Tunnel Encapsulation Attribute (23)

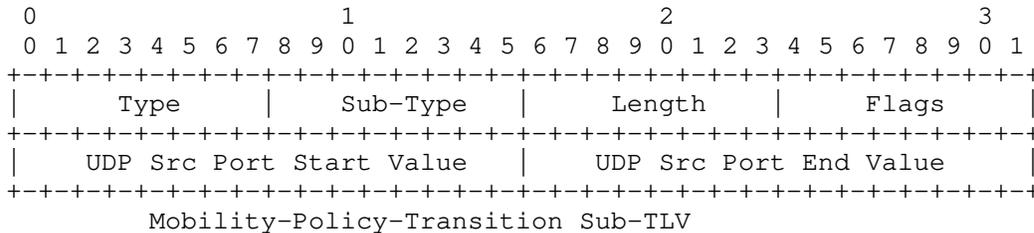
Tunnel Type: SR Policy (15)

{other SR policies Sub-TLVs specified in [BGP-SR-TE-POLICY]}

Mobility-Policy-Transition Sub-TLV

Suppose a mobility service has multiple instances reachable from different SR egress routers. In that case, the SR-TE controller might need to send multiple SR-TE policies, each with one of those egress addresses in the "Endpoint" field.

The Mobility-Policy-Transition Sub-TLV is optional, only applicable to the mobility traffic whose destinations match with the prefixes in the SR Policy NLRI. It MUST NOT appear more than once in the SR-TE Policy.



Mobility-Policy-Transition Sub-TLV

where:

- Type = Mobility-Policy-Transition: TBD1 by IANA.
- Sub-Type: This field has one of the following values:
 - 0: Reserved.
 - 1: UDP Source Port Range.
 - 2 - 255: Reserved for future use.
- Length: 6 octets.

- Flags: 1 octet of flags. None are defined at this stage. Flags SHOULD be set to zero on transmission and MUST be ignored on receipt.
- UDP Src Port Start Value: 2 octets value to define the starting of the value of the UDP Src Port range.
- UDP Src Port End Value: 2 octets value to define the end value of the UDP Src Port range.

6.2. Operations of Applying the TN Aware Mobility Policy

The TN Aware Mobility Policy might be based on something other than the destination or source addresses of the packets from/to UEs because the 5G's Session Controller manages the UE traffic on a finer granularity than the IP addresses. In this case, the BGP UPDATE with the SR-TE Policy SAFI doesn't include the top-level NLRI [RFC4271], which means that the policy specified by the Mobility-Policy-Transition Sub-TLV applies to all the mobility traffic from the 5G N6 Interface (i.e., from the UPFs).

When the TN Aware Mobility Policy is based on UE traffic's destination addresses, the BGP UPDATE with the SR-TE Policy SAFI includes the top-level NLRI [RFC4271], which means that the policy specified by the Mobility-Policy-Transition Sub-TLV only applies to the mobility traffic with destinations specified in the top-level NLRI [RFC4271] from the 5G N6 Interface (i.e., from the UPFs).

The ingress nodes compare all the BGP UPDATE messages for the routes and the SR-TE policies sent from the controller per procedure described by [RFC9256].

7. TN Aware Mobility Policies Distribution via PCE Controller

To achieve the desired QoS for the mobility traffic via the PCE Controller, the Ingress C-PE, standalone or embedded within the UPF, is assumed to support the PCEP communication (PCEP) protocol with the PCE Controller.

The PCE controller must have the mapping of the mobility traffic from the GTP tunnels [TS.23.501-3GPP] with the desired transport path characteristics.

[RFC9168] specifies using the BGP's Flow Specification Rules for IPv4 [RFC8955] and IPv6 [RFC8956] in PCE Communication Protocol (PCEP). Section 9 of this document specifies the mechanism of using UDP Source port filtering in the BGP FlowSpec.

8. TN Aware Mobility Policies via NETCONF/RestConf/gRPC

Assume that the Ingress C-PE UPF has Restconf or gRPC connection with the TE Controller, the source UDP port YANG Data Model of the [RFC8519] can be utilized to pass the allowed source UDP ports.

Here is an example using the [RFC8519] specified YANG Data Model:

```
<?xml version="1.0" encoding="UTF-8"?>
  <config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
    <acls
      xmlns="urn:ietf:params:xml:ns:yang:ietf-access-control-list">
      <acl>
        <name>sample-port-acl</name>
        <type>ipv4-acl-type</type>
        <aces>
          <ace>
            <name>rule1</name>
            <matches>
              <udp>
                <source-port>
                  <lower-port>16384</lower-port>
                  <upper-port>16387</upper-port>
                </source-port>
              </udp>
            </matches>
            <actions>
              <forwarding>[TE-Policy]</forwarding>
            </actions>
          </ace>
        </aces>
      </acl>
    </acls>
  </config>
```

Note: Need to search TE YANG modules to find one that includes the ACTION of selecting the specified "TE-Policy".

9. Extend BGP FlowSpec for TN Aware Mobility

[FlowSpec-TN-Aware] specifies a BGP Flow Specification (FlowSpec) extension to disseminate the policies from 5G mobile networks so that the 5G mobile systems slices and Service Types (SSTs) can be mapped to optimal underlying network paths in the data network outside the 5G UPFs which is the N6 interface in 3GPP 5G Architecture [3GPP TR 23.501].

10. Mapping of TN Characteristics on SD-WAN Edge Node

This section specifies a generic approach for SD-WAN Edge Node to map the mobility traffic to the appropriate IP data network paths based on the policies passed from the 5G Core. The existing [TN-AWARE-MOBILITY] draft is extended to support new Transport Path Characteristics "Security" for the mobile traffic where security is important for specific mobile traffic.

Based on the UDP Src Port characteristics from the mobile network, the SD-WAN edge node can determine which traffic needs to be carried by a secure tunnel or an un-secure tunnel where low latency is more important than security.

The below figure tries to capture the overall topology, and how to map the mobility traffic for Enterprise 5G traffic:

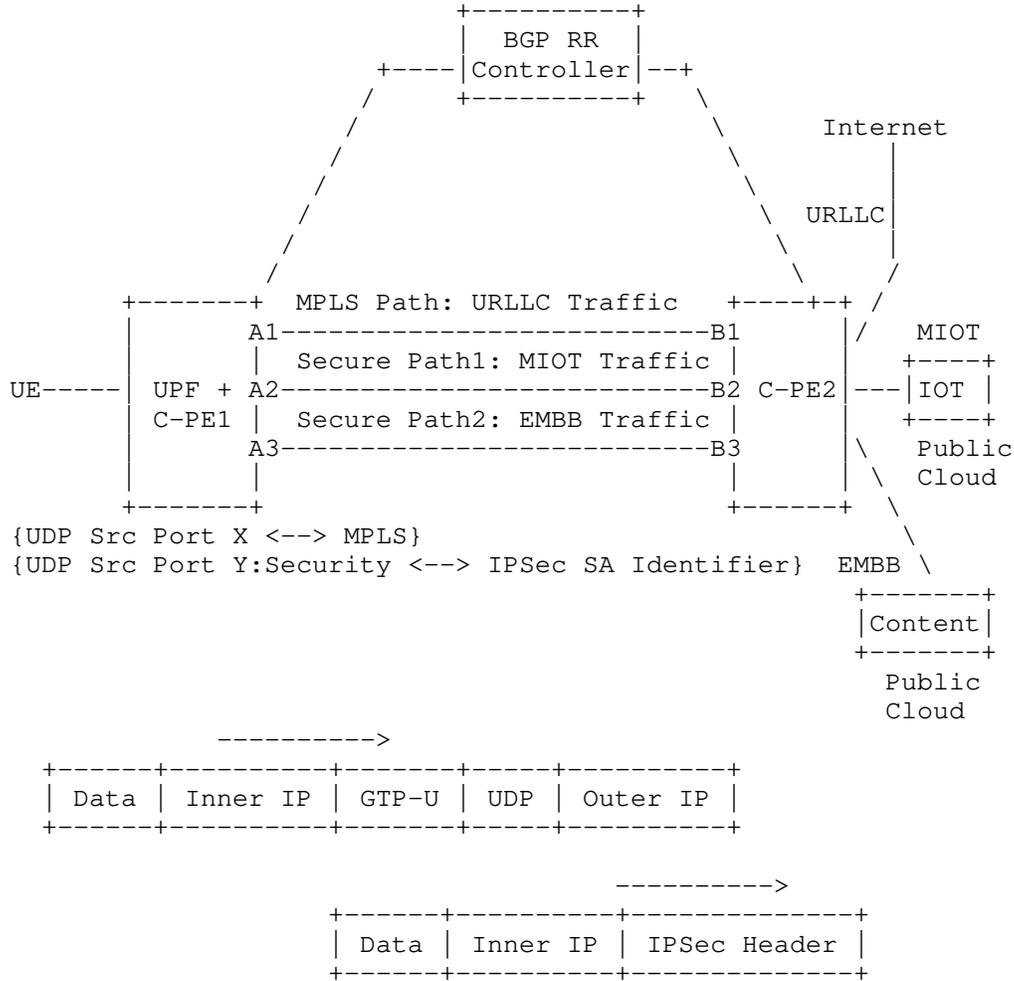


Figure 7: Secure TN Aware Mobility Traffic Mapping on SD-WAN Edge

The SD-WAN Edge Node can map the URLLC traffic without any security characteristics to the underlay MPLS path, whereas MIOT, and EMBB traffic with security characteristics gets mapped to the underlay IPsec Tunnel path. The mapping between SD-WAN overlay and underlay routes are described in the [BGP-SDWAN-Discovery] draft.

The SD-WAN Edge identifies the incoming mobility traffic characteristics using the class-map defined in Section 5.1 and selects the appropriate SD-WAN underlay tunnel.

11. IANA Considerations

No IANA action is needed.

12. Security Considerations

This document does not introduce any new security issues.

13. Contributors

The following people have contributed to this document.

Dhruv Dhody
Huawei Technologies

Email: dhruv.ietf@gmail.com

14. References

14.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3209] D. Awduche, et al, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC3209, Dec. 2001.
- [RFC5440] JP. Vasseur, Ed., JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", March 2009.
- [RFC9012] K. Patel, et al, "The BGP Tunnel Encapsulation Attribute", April 2021.

[RFC8986] c. Filsfils, et al, "Segment Routing over IPv6 (SRv6) Network Programming", RFC8986, Feb. 2021.

[RFC9168] D. Dhody, A. Farrel, Z. Li, "Path Computation Element Communication Protocol (PCEP) Extension for Flow Specification", RFC9168, Jan. 2022.

14.2. Informative References

[FlowSpec-TN-Aware] L. Dunbar, K. Majumdar, U. Chunduri, "BGP Dissemination of FlowSpec for Transport Aware Mobility", draft-dmc-idr-flowspec-tn-aware-mobility-04, July 2023.

[TN-AWARE-MOBILITY] U. Chunduri, et al, "Transport Network aware Mobility for 5G", draft-ietf-dmm-tn-aware-mobility-07, July 2023

[BGP-SR-TE-POLICY] S. Previdi, et al, "Advertising Segment Routing Policies in BGP", draft-ietf-idr-segment-routing-te-policy-23, July 2023

[SDWAN-BGP-USAGE] L. Dunbar, et al, "BGP Usage for SDWAN Overlay Networks", draft-ietf-bess-bgp-sdwan-usage-14, July 2023

[BGP-SDWAN-Discovery] L. Dunbar, et al, "BGP UPDATE for SDWAN Edge Discovery", draft-ietf-idr-sdwan-edge-discovery-10, June 2023

[FLOWSPEC-PATH-REDIRECT] G. Van de Velde, K. Patel, Z.Li, "Flowspec Indirection-id Redirect", draft-ietf-idr-flowspec-path-redirect-12, Nov. 2022.

[FLOWSPEC-TN-MOBILITY] L. Dunbar, K. Majumdar, U. Chunduri, "BGP Dissemination of FlowSpec for Transport Aware Mobility", draft-dmc-idr-flowspec-tn-aware-mobility-04. July 2023,

[RTG-PPR] U. Chunduri, L. Contreras, M. Bhaskaran, J. Tantsura, and P. Muley, "Preferred Path Routing Framework", Work in Progress, draft-chunduri-rtgwg-preferred-path-routing-03, Nov 2022.

[NET2CLOUD] L. Dunbar, et al, "Dynamic Networks to Hybrid Cloud DCs: Problem Statement and Mitigation Practices", draft-ietf-rtgwg-net2cloud-problem-statement-29, Aug. 2023.

[UPlane-Ana-5G] S. Homma, T. Miyasaka, S. Matsushima, and D. Voyer, "User Plane Protocol and Architectural Analysis on 3GPP 5G System", Work in Progress, draft-ietf-dmm-5g-uplaneanalysis-04, 2 November 2020.

15. Acknowledgments

TBD.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Kausik Majumdar
Microsoft

Email: kmajumdar@microsoft.com

Uma Chunduri
Intel Corporation

Email: umac.ietf@gmail.com

Linda Dunbar
Futurewei

Email: linda.dunbar@futurewei.com

