

RTG Working Group
Internet Draft
Intended status: Standard track
Expires: January 24, 2024

L. Dunbar
Futurewei
K. Majumdar
Microsoft
U. Chunduri
Intel
July 24, 2023

BGP Dissemination of FlowSpec for Transport Aware Mobility
draft-dmc-idr-flowspec-tn-aware-mobility-04

Abstract

This document defines 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].

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.....	2	
2. Conventions used in this document.....	3	
3. TN-Aware matching conditions.....	4	
4. Redirect a flow over an underlay tunnel.....	6	
5. FlowSpec Redirect to Indirection-ID Non-Transitive Extended nity.....	8	Commu
6. IANA Considerations.....	9	
7. Security Considerations.....	9	
8. Contributors.....	9	
9. References.....	9	
9.1. Normative References.....	9	
9.2. Informative References.....	10	
10. Acknowledgments.....	10	
Authors' Addresses.....	12	

1. Introduction

The [TN-AWARE-MOBILITY-EXT] describes a framework for extending the mobility-aware transport network characteristics through the Data Network outside the 5G UPFs.

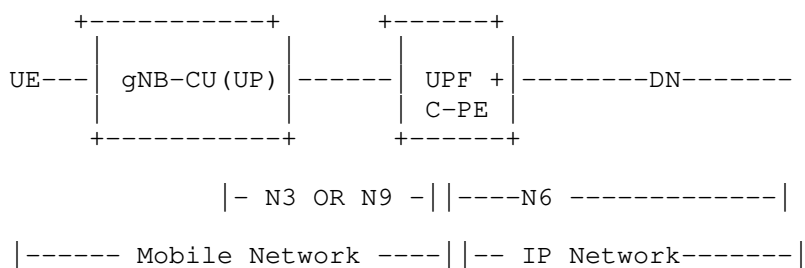


Figure 1: Mobile and IP Data Network for UE

The 5G UPF terminates the 5G GTP tunnels from gNB and passes the IP packets to the N6 Interface [3GPP] data networks, which deliver the packets over hybrid paths, like MPLS, SR paths, Private-IP, or public Internet to reach the packets' destinations.

This document specifies how to use FlowSpec 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].

Border Gateway Protocol (BGP) Flow Specification (FlowSpec) [RFC8955] and FlowSpec for IPv6 [RFC8956] leverage the BGP Control Plane to simplify the distribution of rules & policies for the specified flows. FlowSpec filter rules can be injected into all BGP peers simultaneously without changing router configuration.

2. Conventions used in this document

BSID - Binding SID
 DC - Data Center

DN	- Data Network (5G)
EMBB	- enhanced Mobile Broadband (5G)
gNB	- 5G NodeB
GTP-U	- GPRS Tunneling Protocol - Userplane (3GPP)
MIOT	- Massive IOT (5G)
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 (5G)
SR	- Segment Routing
SR-PCE	- SR Path Computation Element
UE	- User Equipment
UPF	- User Plane Function (5G)
URLLC	- Ultra reliable and low latency communications (5G)

3. TN-Aware matching conditions

[RFC8955] defines a BGP Network Layer Reachability Information (NLRI) format to distribute traffic flow specification rules. The NLRI for (AFI=1, SAFI=133) specifies IPv4 unicast filtering. The NLRI for (AFI=1, SAFI=134) specifies IPv4 BGP/MPLS VPN filtering [RFC7432]. The Flow Specification match part defined in [RFC8955]

includes L3/L4 information like IPv4 source/destination prefix, protocol, ports, etc., so traffic flows can be filtered based on L3/L4 information. [RFC8956] specifies the filtering to cover IPv6 (AFI=2) L3/L4.

The NLRI FlowSpec components described in RFC8955 and RFC8956 are adequate for specifying the UDP Source Port Range which is used to differentiate SLAs of flows from UPFs [EXT-TN-AWARE-Mobility].

The ingress PE, which can be a function integrated with a UPF or an edge router directly connected to a UPF, acting as BGP FlowSpec Receiver, is assumed to have a BGP FlowSpec session with the FlowSpec Controller. The mobility traffic destination would resolve in the BGP Peer Next Hop in the data network. The BGP FlowSpec Controller would be programmed with {5G UDP Src Port Range} to map different SSTs defined in [TN-AWARE-MOBILITY] to create an internal mapping table for {5G UDP Src Port Range} < -- > {BGP FlowSpec Generalized Indirection-ID}. The Mobility IP packets coming out of the UPF, i.e., the GTP header being decapsulated, carrying a specific UDP Source Port, can be classified based on the matching policies carried by the FlowSpec NLRI.

For example, to filter out flows with source UDP port number between [i, j], the following encoding can be used in the NLRI (SAFI=133 or SAFI 134):

Encoding

<Type = 6, [numeric_op1, i][numeric_op2, j]>

<Type = 2, [numeric_op3, Src-Prefix]>

<Type = 1, [numeric_op4, Dest-prefix]>

Numeric_Op1 is:

0	1	2	3	4	5	6	7
e	a	len	0	lt	gt	eq	
0	1	00	0	0	1	0	

Numeric_Op2 is:

0	1	2	3	4	5	6	7
e	a	len		0	lt	gt	eq
1	1	00		0	1	0	0

Where len ==0, which indicates two bytes of value [i] follow the Numeric_op1 and two bytes of value [j] follow the Numeric_op2.

The "numeric_op3" and "numeric_op4" are for comparing the source and destination addresses of the UE traffic.

4. Redirect a flow over an underlay tunnel

For the flows matching with the filter conditions carried by the FlowSpec NLRI, the policy for redirect path can indicate a set of underlay tunnels or one underlay tunnel.

As the BGP FlowSpec Receiver, i.e., the ingress PE, takes the action of redirecting traffic to specific underlay tunnels, a non-transitive Extended Community for Path Redirect [Flowspec-path-redirect] and [SRv6-flowspec-path-redirect] should be used.

[IANA Action: need a new type:

0x49 FlowSpec Redirect to Indirection-id Non-transitive Extended Community.

]

For hierarchical RR deployments where the FlowSpec rules need to be propagated via the RRs to the ingress PE, the Transitive Path Redirect Extended Community [FlowSpec-path-redirect] can be used.

The figure below depicts the overall topology, showing the mobility traffic from UPF being redirected to different paths per the BGP FlowSpec from the Controller:

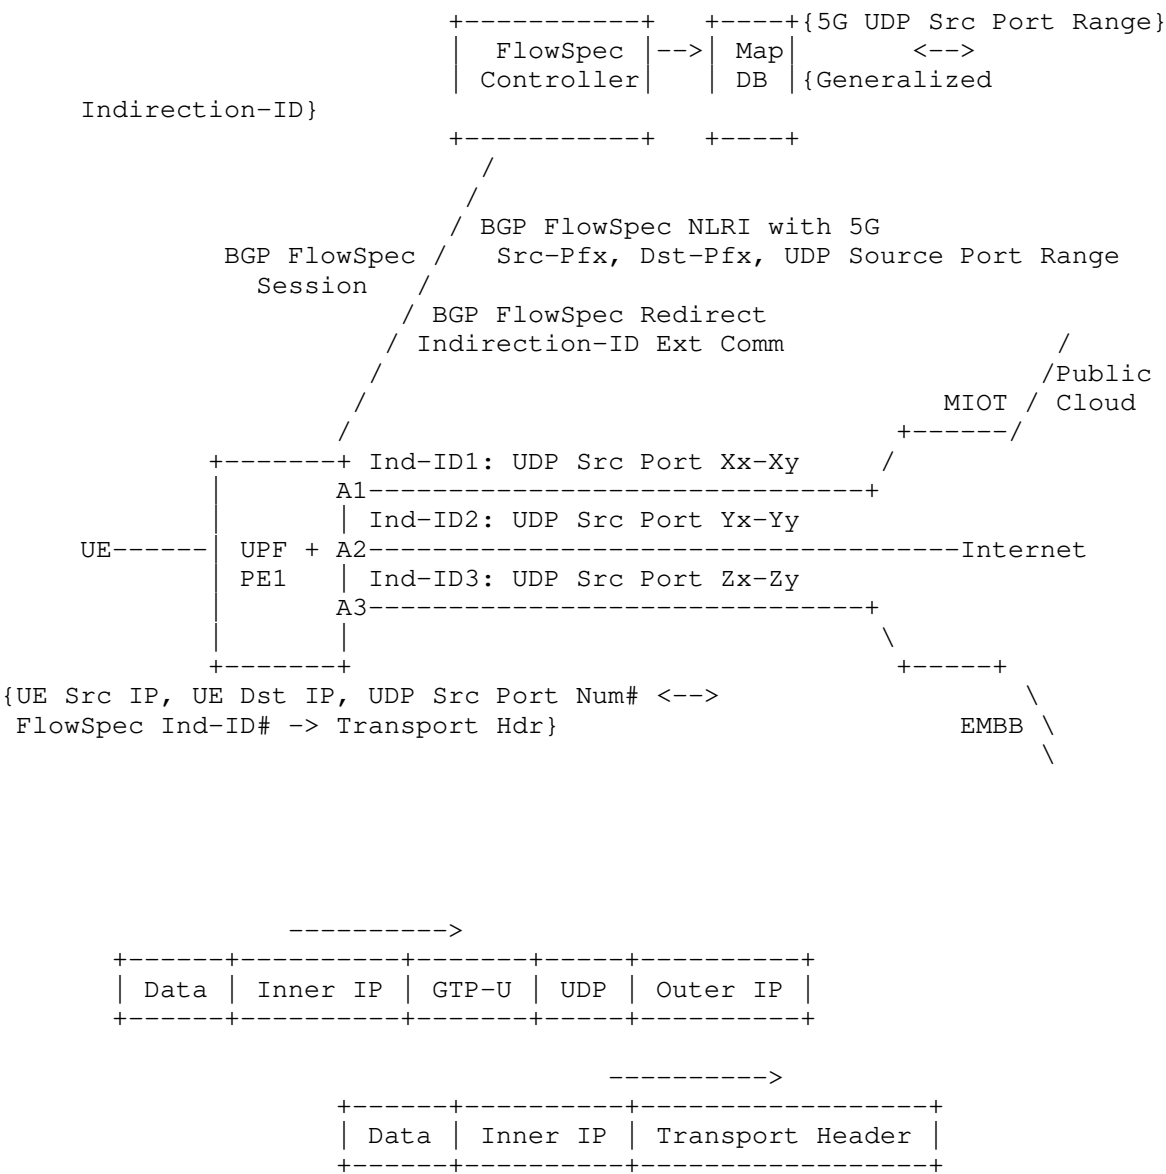


Figure 2: Mobility Traffic Mapping to Redirect Path

5. FlowSpec Redirect to Indirection-ID Non-Transitive Extended Community

This section defines "FlowSpec Redirect to Indirection-ID Non-Transitive Extended Community for IPsec Tunnel ID". The format of this extended community is shown below:




Figure 3: Redirect to Ind-ID Ext Community for IPsec Tunnel

Where

Type = 0x49 (to be assigned by IANA): Non-Transitive FlowSpec Redirect to Indirection-ID Extended Community for IPsec Tunnel ID.

[Note: Type = 0x09 for Transitive FlowSpec Redirect to Indirection-ID Extended Community can also be used for Hierarchical deployment, where the FlowSpec Update needs to be propagated]

IPsec SA Sub-Type: 1 octet, its value (TBD) will be assigned by IANA to indicate the ID carried by the Extended Community is IPsec SA ID. Assuming the IPsec SA is pre-established, its Security Association (SA) ID is within a single administrative domain a globally unique identifier. The allocation and establishment of the IPsec SA among peers is outside scope of the document.

Flags: Same as that defined in [Flowspec-path-redirect].

IPsec SA ID-Type: 1 octet value. Here are the new values needed for IPsec IPv4 tunnel (to be assigned by IANA)

v1 - Inner Encap type = IPSec+GRE

v2 - Inner Encap type = IPSec+Vxlan

6. IANA Considerations

This draft needs an IANA code point allocation for the Non-Transitive FlowSpec Redirect to Indirection-ID Extended Community.

Type: Non-Transitive FlowSpec Redirect to Indirection-ID Extended Community for IPSec Tunnel ID.

IPsec SA Sub-Type:

IPSec SA ID-Type:

v1 - Inner encap type = IPSec+GRE

v2 - Inner encap type = IPSec+Vxlan

7. Security Considerations

TBD.

8. Contributors

The following people have contributed to this document.

9. References

9.1. Normative References

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

[RFC8955] C. Loibl, et al, "Dissemination of Flow specification Rules", Dec 2020.

[RFC8956] C. Loibl, et al, "Dissemination of Flow Specification Rules for IPv6". Dec 2020.

9.2. Informative References

[RFC5440] JP. Vasseur, Ed., JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", March 2009

[Flowspec-path-redirect] G. Van De Velde, et al, "Flowspec Indirection-id Redirect", draft-ietf-idr-flowspec-path-redirect-11, March 2020

[SRv6-Flowspec-path-redirect] G. Van De Velde, et al, "Flowspec Indirection-id Redirect for SRv6", draft-ietf0-idr-srv6-flowspec-path-redirect-05, Jan. 2021

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

[TN-AWARE-MOBILITY-EXT] K. majumdar, et al, "Extension of Transport Aware Mobility in Data Network", draft-mcd-rtgwg-extension-tn-aware-mobility-06, July 2023

[BGP-SR-TE-POLICY] S. Previdi, et al, "Advertising Segment Routing Policies in BGP", draft-ietf-idr-segment-routing-te-policy-09, November 2020

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

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

10. Acknowledgments

TBD.

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

Authors' Addresses

Linda Dunbar
Futurewei
2330 Central Expressway
Santa Clara, CA 95050

Email: linda.dunbar@futurewei.com

Kausik Majumdar
Microsoft
Email: kmajumdar@microsoft.com

Uma Chunduri
Intel
2200 Mission College Blvd
Santa Clara, CA 95052

Email: umac.ietf@gmail.com

