
Stream: Internet Engineering Task Force (IETF)

RFC: [9889](#)

Category: Informational

Published: November 2025

ISSN: 2070-1721

Authors:

K. Szarkowicz, Ed.
HPE

R. Roberts, Ed.
Nokia

J. Lucek
HPE

M. Boucadair, Ed.
Orange

LM. Contreras
Telefonica

RFC 9889

A Realization of Network Slices for 5G Networks Using Current IP/MPLS Technologies

Abstract

Network slicing is a feature that was introduced by the 3rd Generation Partnership Project (3GPP) in Mobile Networks. Realization of 5G slicing implies requirements for all mobile domains, including the Radio Access Network (RAN), Core Network (CN), and Transport Network (TN).

This document describes a network slice realization model for IP/MPLS networks with a focus on the Transport Network fulfilling the service objectives for 5G slicing connectivity. The realization model reuses many building blocks commonly used in service provider networks at the current time.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9889>.

Copyright Notice

Copyright (c) 2025 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 (<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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Terminology	5
2.1. Definitions	5
2.2. Abbreviations	5
3. 5G Network Slicing Integration in Transport Networks	8
3.1. Scope of the Transport Network	8
3.2. 5G Network Slicing Versus Transport Network Slicing	9
3.3. Transport Network Reference Design	10
3.4. Orchestration Overview	14
3.5. Mapping 5G Network Slices to Transport Network Slices	18
3.6. First 5G Network Slice Versus Subsequent Slices	20
3.7. Overview of the Transport Network Realization Model	23
4. Handoff Between Domains	25
4.1. VLAN Handoff	25
4.2. IP Handoff	26
4.3. MPLS Label Handoff	27
5. QoS Mapping Realization Models	32
5.1. QoS Layers	32
5.2. QoS Realization Models	33
5.3. Transit Resource Control	46
6. PE Underlay Transport Mapping Models	47
6.1. 5QI-Unaware Model	49
6.2. 5QI-Aware Model	50

7. Capacity Planning/Management	51
7.1. Bandwidth Requirements	51
7.2. Bandwidth Models	54
8. Network Slicing OAM	56
9. Scalability Implications	57
10. IANA Considerations	57
11. Security Considerations	57
12. References	58
12.1. Normative References	58
12.2. Informative References	59
Appendix A. Example of Local IPv6 Addressing Plan for Network Functions	64
Acknowledgments	66
Contributors	66
Authors' Addresses	67

1. Introduction

This document focuses on network slicing for 5G networks, covering the connectivity between Network Functions (NFs) across multiple domains such as edge clouds, data centers, and the Wide Area Network (WAN). The document describes a network slice realization approach that fulfills 5G slicing requirements by using existing IP/MPLS technologies (at the time of publication of this document) to optimally control connectivity Service Level Agreements (SLAs) offered for 5G Network Slices. To that aim, this document describes the scope of the Transport Network in 5G architectures ([Section 3.1](#)), disambiguates 5G Network Slicing versus Transport Network Slicing ([Section 3.2](#)), draws the perimeter of the various orchestration domains to realize slices ([Section 3.4](#)), and identifies the required coordination between these orchestration domains for adequate setup of Attachment Circuits (ACs) ([Section 3.4.2](#)).

This work is compatible with the framework defined in [[RFC9543](#)], which describes network slicing in the context of networks built from IETF technologies. Specifically, this document describes an approach to how RFC 9543 Network Slices are realized within provider networks and how such slices are stitched to Transport Network resources in a customer site in the context of Transport Network Slices ([Figure 1](#)). The realization of an RFC 9543 Network Slice (i.e., connectivity with performance commitments) involves the provider network and partially the

AC (the Provider Edge (PE) side of the AC). This document assumes that the customer site infrastructure is over-provisioned and involves short distances (low latency) where basic QoS/scheduling logic is sufficient to comply with the Service Level Objectives (SLOs).

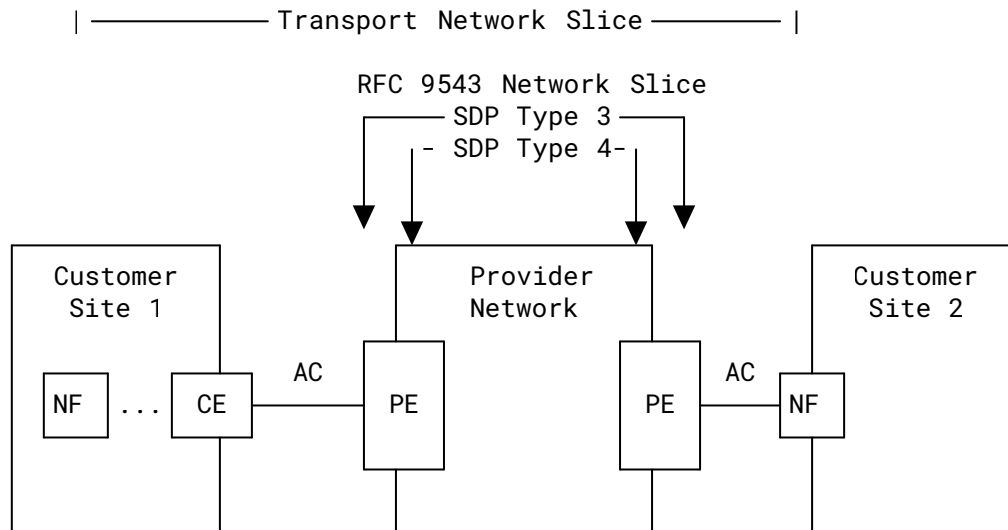


Figure 1: Transport Network Slice and RFC 9543 Network Slice Scopes

This document focuses on RFC 9543 Network Slice deployments where the Service Demarcation Points (SDPs) are located per Types 3 and 4 in Figure 1 of [RFC9543].

The realization approach described in this document is typically triggered by Network Slice Service requests. How a Network Slice Service request is placed for realization, including how it is derived from a 5G Slice Service request, is out of scope. Mapping considerations between 3GPP and RFC 9543 Network Slice Service (e.g., mapping of service parameters) are discussed in documents such as [NS-APP].

The 5G control plane uses the Single Network Slice Selection Assistance Information (S-NSSAI) for slice identification [TS-23.501]. Because S-NSSAIs are not visible to the transport domain, 5G domains can expose the 5G Network Slices to the transport domain by mapping to explicit data plane identifiers (e.g., Layer 2, Layer 3, or Layer 4). Passing information between customer sites and provider networks is referred to as the "handoff". Section 4 lists a set of handoff methods for slice mapping purposes.

Unlike approaches that require new protocol extensions (e.g., [NS-IP-MPLS]), the realization model described in this document uses a set of building blocks commonly used in service provider networks (at the time of publication of this document). The model uses (1) L2VPN [RFC4664] and/or L3VPN [RFC4364] service instances for logical separation, (2) fine-grained resource control at the PEs, (3) coarse-grained resource control within the provider network, and (4) capacity planning and management. More details are provided in Sections 3.7, 5, 6, and 7.

This realization model uses a single Network Resource Partition (NRP) ([Section 7.1](#) of [\[RFC9543\]](#)). The applicability to multiple NRPs is out of scope.

Although this document focuses on 5G, the realizations are not fundamentally constrained by the 5G use case. The document is not intended to be a BCP and does not claim to specify mandatory mechanisms to realize network slices. Rather, a key goal of the document is to provide pragmatic implementation approaches by leveraging existing techniques that are readily available and widely deployed. The document is also intended to align the mobile and the IETF perspectives of slicing from a realization perspective.

For a definitive description of 3GPP network architectures, the reader should refer to [\[TS-23.501\]](#). More details can be found in [\[Book-5G\]](#).

2. Terminology

2.1. Definitions

The document uses the terms defined in [\[RFC9543\]](#). Specifically, the use of "Customer" is consistent with [\[RFC9543\]](#) but with the following contextualization (see also [Section 3.3](#)):

Customer: An entity that is responsible for managing and orchestrating the end-to-end 5G Mobile Network, notably the Radio Access Network (RAN) and Core Network (CN).

This entity is distinct from the customer of a 5G Network Slice Service.

This document makes use of the following terms:

Customer site: A customer manages and deploys 5G NFs (e.g., gNodeB (gNB) and 5G Core (5GC)) in customer sites. A customer site can be either a physical or a virtual location. A provider is responsible for interconnecting customer sites.

Examples of customer sites are a customer private locations (e.g., Point of Presence (PoP) and Data Center (DC)), a Virtual Private Cloud (VPC), or servers hosted within the provider network or colocation service.

Resource Control: In the context of this document, resource control is used mainly to refer to buffer management and relevant Quality of Service (QoS) functions.

"5G Network Slicing" and "5G Network Slice": Refer to "Network Slicing" and "Network Slice" as defined in [\[TS-28.530\]](#).

2.2. Abbreviations

3GPP: 3rd Generation Partnership Project

5GC: 5G Core

5QI: 5G QoS Indicator

A2A: Any-to-Any

AC: Attachment Circuit

AMF: Access and Mobility Management Function

CE: Customer Edge

CIR: Committed Information Rate

CS: Customer Site

CN: Core Network

CoS: Class of Service

CP: Control Plane

CU: Centralized Unit

CU-CP: Centralized Unit Control Plane

CU-UP: Centralized Unit User Plane

DC: Data Center

DDoS: Distributed Denial of Service

DM: Data Model

DSCP: Differentiated Services Code Point

eCPRI: enhanced Common Public Radio Interface

FIB: Forwarding Information Base

GPRS: General Packet Radio Service

gNB: gNodeB

GTP: GPRS Tunneling Protocol

GTP-U: GPRS Tunneling Protocol User Plane

IGP: Interior Gateway Protocol

L2VPN: Layer 2 Virtual Private Network

L3VPN: Layer 3 Virtual Private Network

LSP: Label Switched Path

MACsec: Media Access Control Security

MIoT: Massive Internet of Things

MNO: Mobile Network Operator

MPLS: Multiprotocol Label Switching

NF: Network Function

NS: Network Slice

NRP: Network Resource Partition

NSC: Network Slice Controller

PE: Provider Edge

PIR: Peak Information Rate

QoS: Quality of Service

RAN: Radio Access Network

RIB: Routing Information Base

RSVP: Resource Reservation Protocol

SD: Slice Differentiator

SDP: Service Demarcation Point

SLA: Service Level Agreement

SLO: Service Level Objective

S-NSSAI: Single Network Slice Selection Assistance Information

SST: Slice/Service Type

SR: Segment Routing

SRv6: Segment Routing version 6

TC: Traffic Class

TE: Traffic Engineering

TN: Transport Network

UP: User Plane

UPF: User Plane Function

URLLC: Ultra-Reliable Low-Latency Communication

VLAN: Virtual Local Area Network

VPN: Virtual Private Network

VRF: Virtual Routing and Forwarding

VXLAN: Virtual Extensible Local Area Network

3. 5G Network Slicing Integration in Transport Networks

3.1. Scope of the Transport Network

The main 5G network building blocks are the Radio Access Network (RAN), Core Network (CN), and Transport Network (TN). The Transport Network is defined by the 3GPP in Section 1 of [TS-28.530]:

part supporting connectivity within and between CN and RAN parts.

The 3GPP management system does not directly control the Transport Network; it is considered a non-3GPP managed system. This is discussed in Section 4.4.1 of [TS-28.530]:

The non-3GPP part includes TN parts. The 3GPP management system provides the network slice requirements to the corresponding management systems of those non-3GPP parts, e.g. the TN part supports connectivity within and between CN and AN parts.

In practice, the TN may not map to a monolithic architecture and management domain. It is frequently segmented, non-uniform, and managed by different entities. For example, [Figure 2](#) depicts an NF instance that is deployed in an edge data center (DC) connected to an NF located in a Public Cloud via a WAN (e.g., MPLS-VPN service). In this example, the TN can be seen as an abstraction representing an end-to-end connectivity based upon three distinct domains: DC, WAN, and Public Cloud. A model for the Transport Network based on orchestration domains is introduced in [Section 3.4](#).




Figure 2: Example of Transport Network Decomposition

3.2. 5G Network Slicing Versus Transport Network Slicing

Network slicing has a different meaning in the 3GPP mobile world and transport world. This difference can be seen from the descriptions below that set out the objectives of 5G Network Slicing (Section 3.2.1) and Transport Network Slicing (Section 3.2.2). These descriptions are not intended to be exhaustive.

3.2.1. 5G Network Slicing

In [TS-28.530], the 3GPP defines 5G Network Slicing as an approach:

where logical networks/partitions are created, with appropriate isolation, resources and optimized topology to serve a purpose or service category (e.g. use case/traffic category, or for MNO internal reasons) or customers (logical system created "on demand").

These resources are from the TN, RAN, and CN domains together with the underlying infrastructure.

Section 3.1 of [TS-28.530] defines a 5G Network Slice as:

a logical network that provides specific network capabilities and network characteristics, supporting various service properties for network slice customers.

3.2.2. Transport Network Slicing

The term "Transport Network Slice" refers to a slice in the Transport Network domain of the 5G architecture. This section elaborates on how Transport Network Slicing is defined in the context of this document. It draws on the 3GPP definitions of "Transport Network" and "Network Slicing" in [TS-28.530].

The objective of Transport Network Slicing is to isolate, guarantee, or prioritize Transport Network resources for Slice Services. Examples of such resources include buffers, link capacity, or even Routing Information Base (RIB) and Forwarding Information Base (FIB).

Transport Network Slicing provides various degrees of sharing of resources between slices (Section 8 of [RFC9543]). For example, the network capacity can be shared by all slices, usually with a guaranteed minimum per slice, or each individual slice can be allocated dedicated network capacity. Parts of a given network may use the former, while others use the latter. For example, in order to satisfy local engineering guidelines and specific service requirements, shared TN resources could be provided in the backhaul (or midhaul), and dedicated TN resources could be provided in the midhaul (or backhaul). The capacity partitioning strategy is deployment specific.

There are different components to implement Transport Network Slices based upon mechanisms such as Virtual Routing and Forwarding (VRF) instances for logical separation, QoS, and Traffic Engineering (TE). Whether all or a subset of these components are enabled is a deployment choice.

3.3. Transport Network Reference Design

Figure 3 depicts the reference design used in this document for modeling the Transport Network based on management perimeters (customer vs. provider).

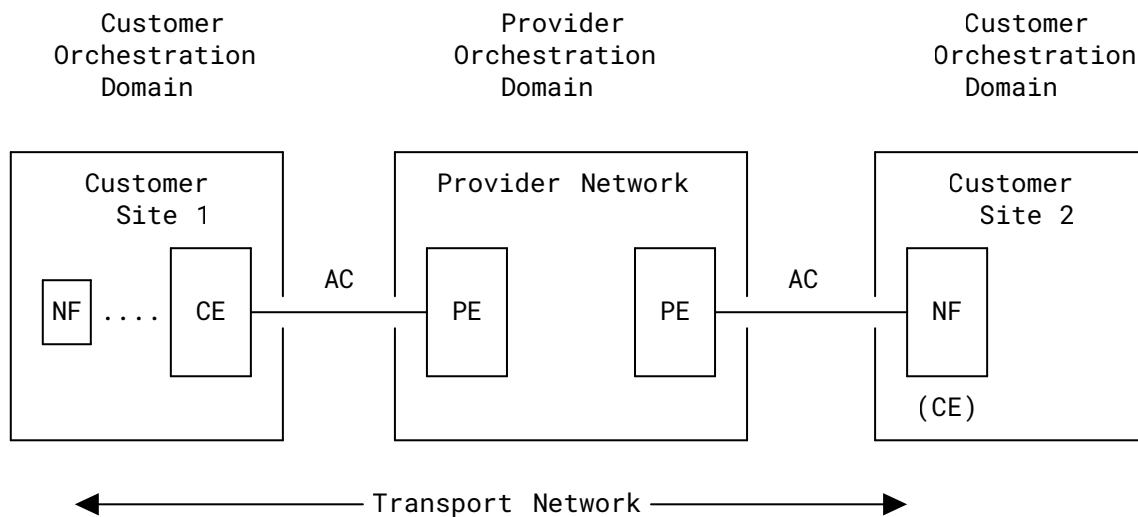


Figure 3: Reference Design with Customer Sites and Provider Network

The description of the main components shown in Figure 3 is provided in the following subsections.

3.3.1. Customer Site (CS)

On top of 5G NFs, a customer may manage additional TN elements (e.g., servers, routers, and switches) within a customer site.

NFs may be hosted on a CE, directly connected to a CE, or located multiple IP hops from a CE.

In some contexts, the connectivity between NFs that belong to the same site can be achieved via the provider network.

The orchestration of the TN within a customer site involves a set of controllers for automation purposes (e.g., Network Function Virtualization Infrastructure (NFVI), Container Network Interface (CNI), Fabric Managers, or Public Cloud APIs). Documenting how these controllers are implemented is out of scope for this document.

3.3.2. Customer Edge (CE)

A CE is a function that provides logical connectivity of a customer site (Section 3.3.1) to the provider network (Section 3.3.3). The logical connectivity is enforced at Layer 2 and/or Layer 3 and is denominated an Attachment Circuit (AC) (Section 3.3.5). Examples of CEs include TN components (e.g., router, switch, and firewalls) and also 5G NFs (i.e., an element of the 5G domain such as Centralized Unit (CU), Distributed Unit (DU), or User Plane Function (UPF)).

A CE is typically managed by the customer, but it can also be co-managed with the provider. A co-managed CE is orchestrated by both the customer and the provider. In this case, the customer and provider usually have control on distinct device configuration perimeters. A co-managed CE

has both PE and CE functions, and there is no strict AC connection, although one may consider that the AC stitching logic happens internally within the CE itself. The provider manages the AC between the CE and the PE.

This document generalizes the definition of a CE with the introduction of "distributed CE"; that is, the logical connectivity is realized by configuring multiple devices in the customer domain. The CE function is distributed. An example of distributed CE is the realization of an interconnection using an L3VPN service based on a distributed CE composed of a switch (SW) in Layer 2 and a router (RTR) in Layer 3, as shown in [Figure 4](#). Another example of distributed CE is shown in [Figure 5](#).

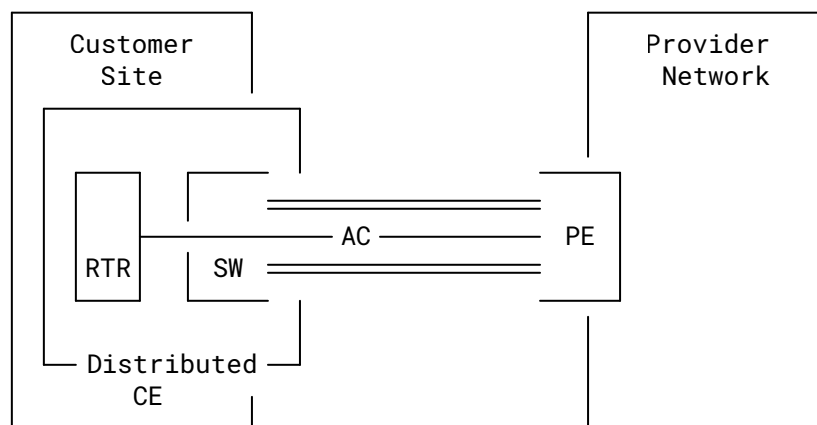


Figure 4: Example of Distributed CE

In most cases, CEs connect to PEs using IP (e.g., via Layer 3 VLAN subinterfaces), but a CE may also connect to the provider network using other technologies such as MPLS (potentially over IP tunnels) or Segment Routing over IPv6 (SRv6) [[RFC8986](#)]. Thus, the CE has awareness of provider service configuration (e.g., control plane identifiers such as Route Targets (RTs) and Route Distinguishers (RDs)). However, the CE is still managed by the customer, and the AC is based on MPLS or SRv6 data plane technologies. The complete termination of the AC within the provider network may happen on distinct routers; this is another example of distributed PE. Service-aware CEs are used, for example, in the deployments discussed in [Sections 4.3.2 and 4.3.3](#).

3.3.3. Provider Network

A provider uses a provider network to interconnect customer sites. This document assumes that the provider network is based on IP, MPLS, or both.

3.3.4. Provider Edge (PE)

A PE is a device managed by a provider that is connected to a CE. The connectivity between a CE and a PE is achieved using one or multiple ACs ([Section 3.3.5](#)).

This document generalizes the PE definition with the introduction of "distributed PE"; that is, the logical connectivity is realized by configuring multiple devices in the provider network (i.e., the Provider Orchestration Domain). The PE function is distributed.

An example of a distributed PE is the "managed CE service". For example, a provider delivers VPN services using CEs and PEs that are both managed by the provider (example (i) in Figure 5). The managed CE can also be a Data Center Gateway as depicted in example (ii) of Figure 5. A provider-managed CE may attach to CEs of multiple customers. However, this device is part of the provider network.

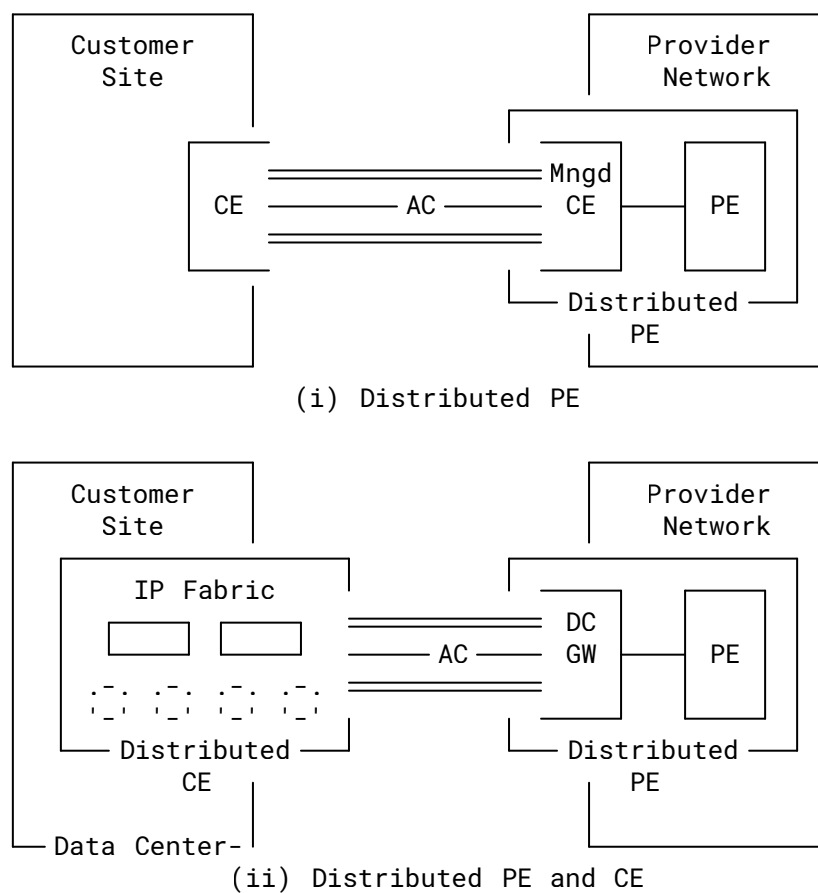


Figure 5: Examples of Distributed PE

In subsequent sections of this document, the terms "CE" and "PE" are used for both single and distributed devices.

3.3.5. Attachment Circuit (AC)

The AC is the logical connection that attaches a CE (Section 3.3.2) to a PE (Section 3.3.4). A CE is connected to a PE via one or multiple ACs.

This document uses the concept of distributed CE and PE (Sections 3.3.2 and 3.3.4) to consolidate a CE/AC/PE definition that is consistent with the orchestration perimeters (Section 3.4). The CEs and PEs delimit the customer and provider orchestration domains, respectively, while an AC interconnects these domains.

For consistency with the terminology used in AC data models (e.g., the data models defined in [RFC9834] and [RFC9835]), this document assumes that an AC is configured on a "bearer", which represents the underlying connectivity. For example, the bearer is illustrated with "=" in Figures 4 and 5.

An AC is technology specific. Examples of ACs are VLAN ACs configured on a physical interface (bearer) or Overlay VXLAN EVI ACs configured on an IP underlay (bearer).

Deployment cases where the AC is also managed by the provider are not discussed in this document because the setup of such an AC does not require any coordination between the customer and provider orchestration domains.

Note: In order to keep the figures simple, only one AC and single-homed CEs are represented. Also, the underlying bearers are not represented in most of the figures. However, this document does not exclude the instantiation of multiple ACs between a CE and a PE nor the presence of CEs that are attached to more than one PE.

3.4. Orchestration Overview

3.4.1. 5G End-to-End Slice Orchestration Architecture

This section introduces a global framework for the orchestration of a 5G end-to-end slice (a.k.a. 5G Network Slice) with a zoom on TN parts. This framework helps to delimit the realization scope of RFC 9543 Network Slices and identify interactions that are required for the realization of such slices.

This framework is consistent with the management coordination example shown in Figure 4.7.1 of [TS-28.530].

In Figure 6, a 5G End-to-End Network Slice Orchestrator (5G NSO) is responsible for orchestrating 5G Network Slices end-to-end. The details of the 5G NSO are out of the scope of this document. The realization of the 5G Network Slices spans RAN, CN, and TN. As mentioned in Section 3.1, the RAN and CN are under the responsibility of the 3GPP management system, while the TN is not. The orchestration of the TN is split into two subdomains in conformance with the reference design in Section 3.3:

Provider Network Orchestration Domain (simply referred to as "Provider Orchestration Domain"):

As defined in [RFC9543], the provider relies on a Network Slice Controller (NSC) to manage and orchestrate RFC 9543 Network Slices in the provider network. This framework allows for managing connectivity with SLOs.

Customer Site Orchestration Domain (simply referred to as "Customer Orchestration Domain"):

The orchestration of TN elements of the customer sites relies upon a variety of controllers (e.g., Fabric Manager, Element Management System, or Virtualized Infrastructure Manager (VIM)).

A Transport Network Slice relies upon resources that can involve both the provider and customer TN domains. More details are provided in [Section 3.4.2](#).

A Transport Network Slice might be considered as a variant of horizontal composition of network slices mentioned in [Appendix A.6](#) of [RFC9543].

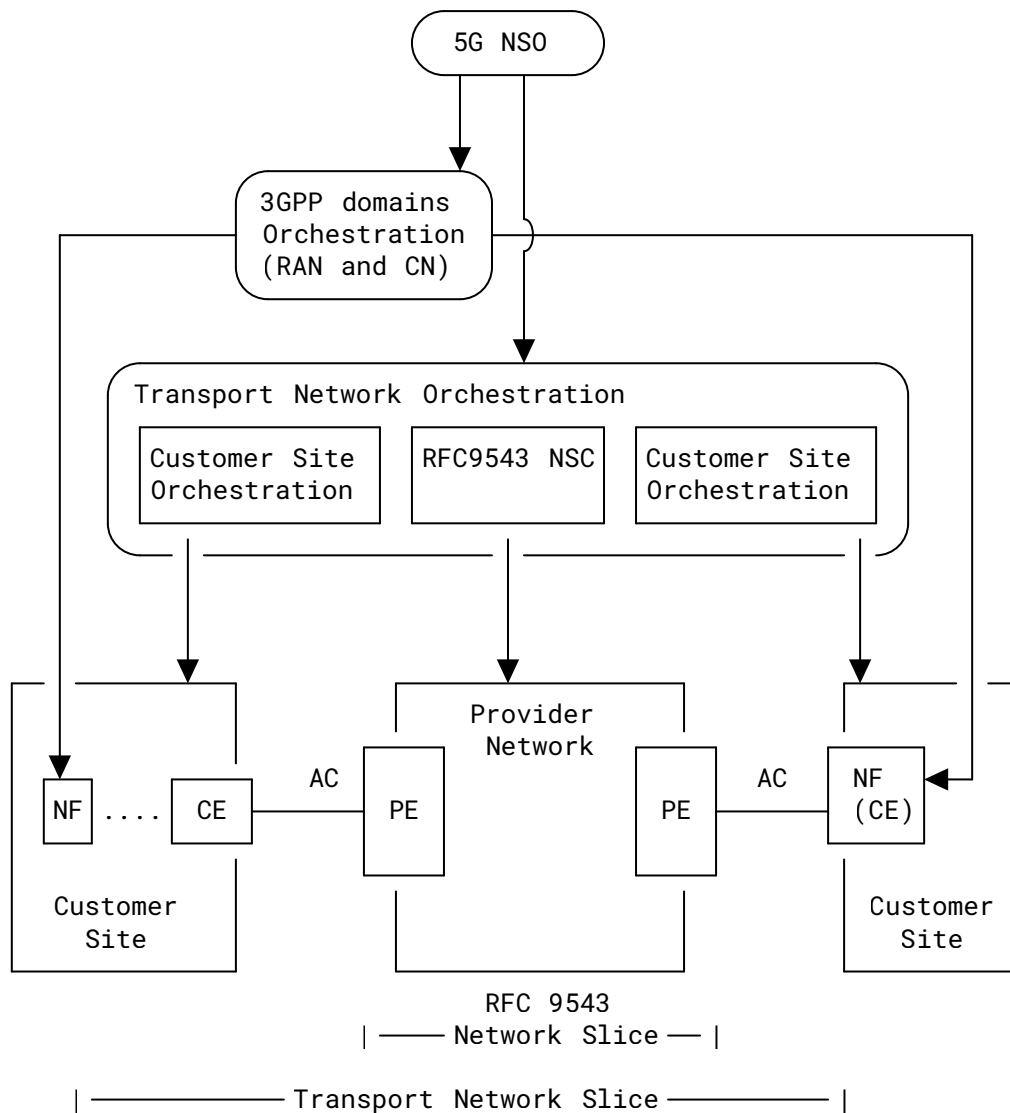


Figure 6: 5G End-to-End Slice Orchestration with TN

The various orchestrations depicted in Figure 6 encompass the 3GPP's Network Slice Subnet Management Function (NSSMF) mentioned, for instance, in Figure 5 of [NS-APP].

3.4.2. Transport Network Segments and Network Slice Instantiation

The concept of distributed PE (Section 3.3.4) assimilates the CE-based SDPs defined in Section 5.2 of [RFC9543] (i.e., Types 1 and 2) as SDP Types 3 or 4 in this document.

In the architecture depicted in Section 3.4.1, the connectivity between NFs can be decomposed into three main segment types:

Customer Site: Either connects NFs located in the same customer site or connects an NF to a CE.

This segment may not be present if the NF is the CE. In this case, the AC connects the NF to a PE.

The realization of this segment is driven by the 5G Network Orchestration (e.g., NF instantiation) and the Customer Site Orchestration for the TN part.

Provider Network: Represents the connectivity between two PEs. The realization of this segment is controlled by an NSC ([Section 6.3](#) of [\[RFC9543\]](#)).

Attachment Circuit: The orchestration of this segment relies partially upon an NSC for the configuration of the AC on the PE customer-facing interfaces and the Customer Site Orchestration for the configuration of the AC on the CE.

PEs and CEs that are connected via an AC need to be provisioned with consistent data plane and control plane information (VLAN IDs, IP addresses/subnets, BGP Autonomous System Number (ASN), etc.). Hence, the realization of this interconnection is technology specific and requires coordination between the Customer Site Orchestration and an NSC. Automating the provisioning and management of the AC is thus key to automate the overall service provisioning. Aligned with [\[RFC8969\]](#), this document assumes that this coordination is based upon standard YANG data models and APIs.

The provisioning of an RFC 9543 Network Slice may rely on new or existing ACs.

[Figure 7](#) is a basic example of a Layer 3 CE-PE link realization with shared network resources (such as VLAN IDs and IP prefixes), which are passed between orchestrators via a dedicated interface, e.g., the Network Slice Service Model (NSSM) [\[NSSM\]](#) or Attachment Circuits as a Service (ACaaS) [\[RFC9834\]](#).

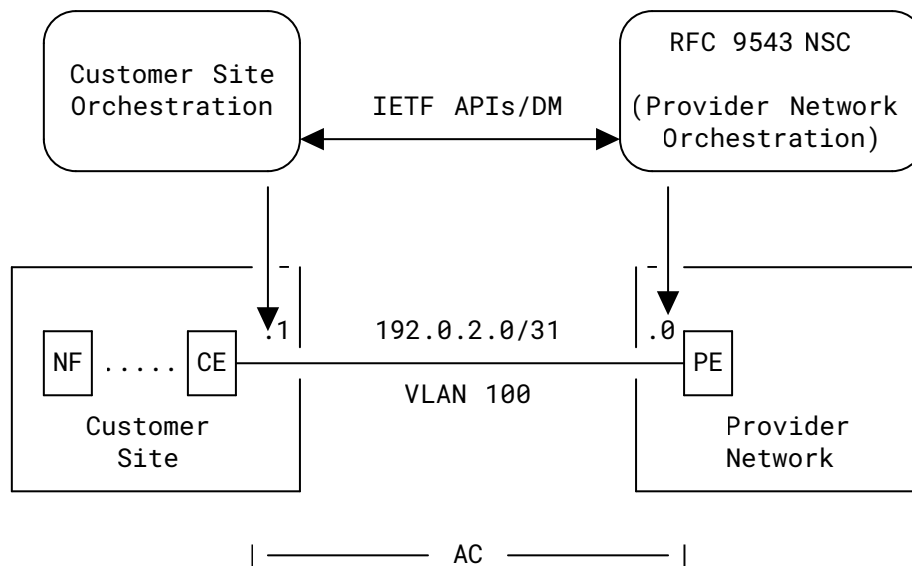


Figure 7: Coordination of Transport Network Resources for AC Provisioning

3.5. Mapping 5G Network Slices to Transport Network Slices

There are multiple options for mapping 5G Network Slices to Transport Network Slices:

1-to-N mapping: A single 5G Network Slice can be mapped to multiple Transport Network Slices. For instance, consider the scenario depicted in Figure 8, which illustrates the separation of the 5G control plane and user plane in Transport Network Slices for a single 5G Enhanced Mobile Broadband (eMBB) network slice. It is important to note that this mapping can serve as an interim step to M-to-N mapping. Further details about this scheme are described in Section 3.6.

M-to-1 mapping: Multiple 5G Network Slices may rely upon the same Transport Network Slice. In such a case, the Service Level Agreement (SLA) differentiation of slices would be entirely controlled at the 5G control plane, for example, with appropriate placement strategies. This use case is illustrated in Figure 9, where a UPF for the Ultra-Reliable Low-Latency Communication (URLLC) slice is instantiated at the edge cloud, close to the gNB CU-UP, to improve latency and jitter control. The 5G control plane and the UPF for the eMBB slice are instantiated in the regional cloud.

M-to-N mapping: The mapping of 5G to Transport Network Slice combines both approaches with a mix of shared and dedicated associations.

In this scenario, a subset of the Transport Network Slices can be intended for sharing by multiple 5G Network Slices (e.g., the control plane Transport Network Slice is shared by multiple 5G Network Slices).

In practice, for operational and scaling reasons, M-to-N mapping would typically be used, with M much greater than N.




Figure 8: 1-to-N Mapping (Single 5G Network Slice to Multiple Transport Network Slices)

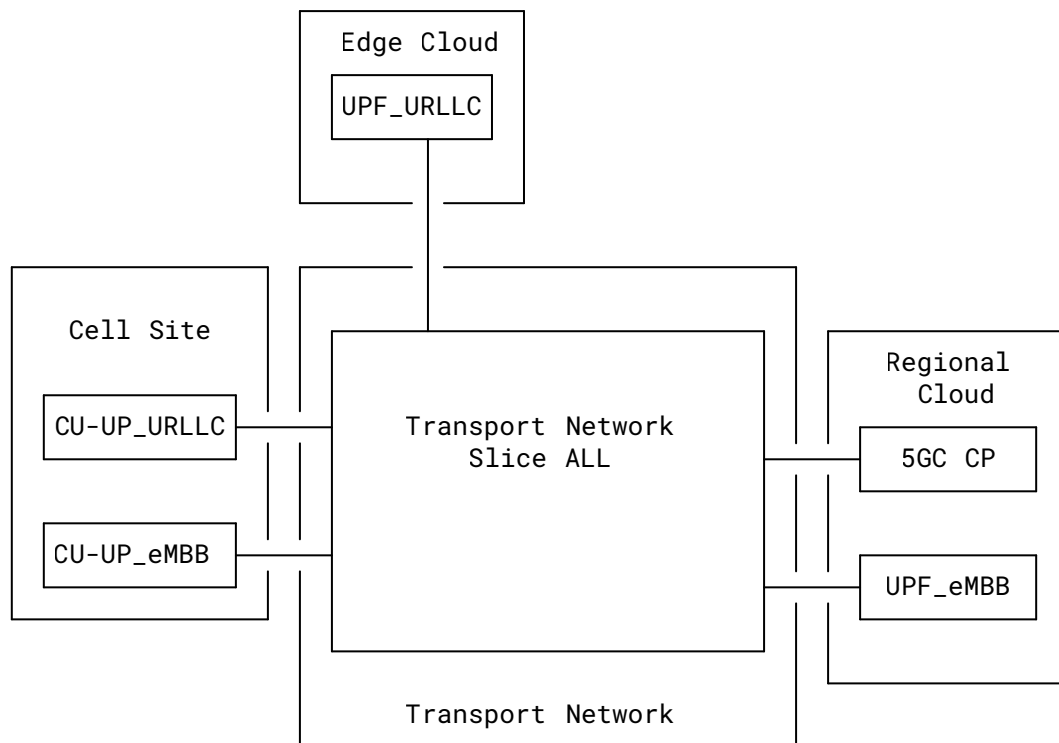


Figure 9: M-to-1 Mapping (Multiple 5G Network Slices to Single Transport Network Slice)

Note that the actual realization of the mapping depends on several factors, such as the actual business cases, the NF vendor capabilities, the NF vendor reference designs, as well as service provider or even legal requirements.

Mapping approaches that preserve the 5G Network Slice identification in the TN (e.g., the approach in [Section 4.2](#)) may simplify required operations to map Transport Network Slices back to 5G Network Slices. However, such considerations are not detailed in this document because these are under the responsibility of the 3GPP orchestration.

3.6. First 5G Network Slice Versus Subsequent Slices

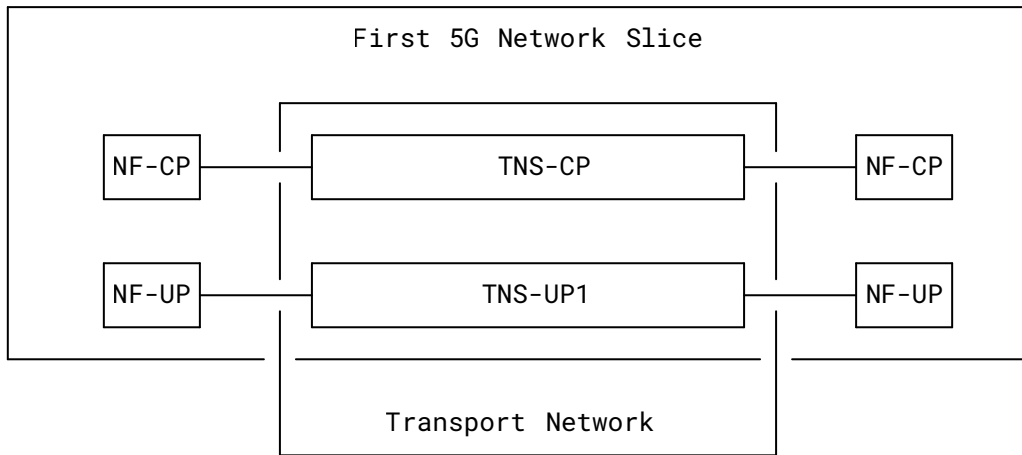
An operational 5G Network Slice incorporates both 5G control plane and user plane capabilities. For instance, in some deployments, in the case of a slice based on split CU in the RAN, both CU-UP and CU-CP may need to be deployed along with the associated interfaces E1, F1-c, F1-u, N2, and N3, which are conveyed in the TN. In this regard, the creation of the "first slice" can be subject to a specific logic that does not apply to subsequent slices. Let us consider the example depicted in [Figure 10](#) to illustrate this deployment. In this example, the first 5G Network Slice relies on the deployment of NF-CP and NF-UP functions together with two Transport Network Slices for the control and user planes (TNS-CP and TNS-UP1). Next, in many cases, the deployment of a second slice relies solely on the instantiation of a UPF (NF-UP2) together with a

dedicated Transport Network Slice for the user plane (TNS-UP2). The control plane of the first 5G Network Slice is also updated to integrate the second slice; the Transport Network Slice (TNS-CP) and Network Functions (NF-CP) are shared.

The model described here, in which the control plane is shared among multiple slices, is likely to be common; it is not mandatory, though. Deployment models with a separate control plane for each slice are also possible.

Section 6.1.2 of [NG.113] specifies that the eMBB slice (SST=1 and no Slice Differentiator (SD)) should be supported globally. This 5G Network Slice would be the first slice in any 5G deployment.

(1) Deployment of first 5G Network Slice



(2) Deployment of additional 5G Network Slice with shared Control Plane

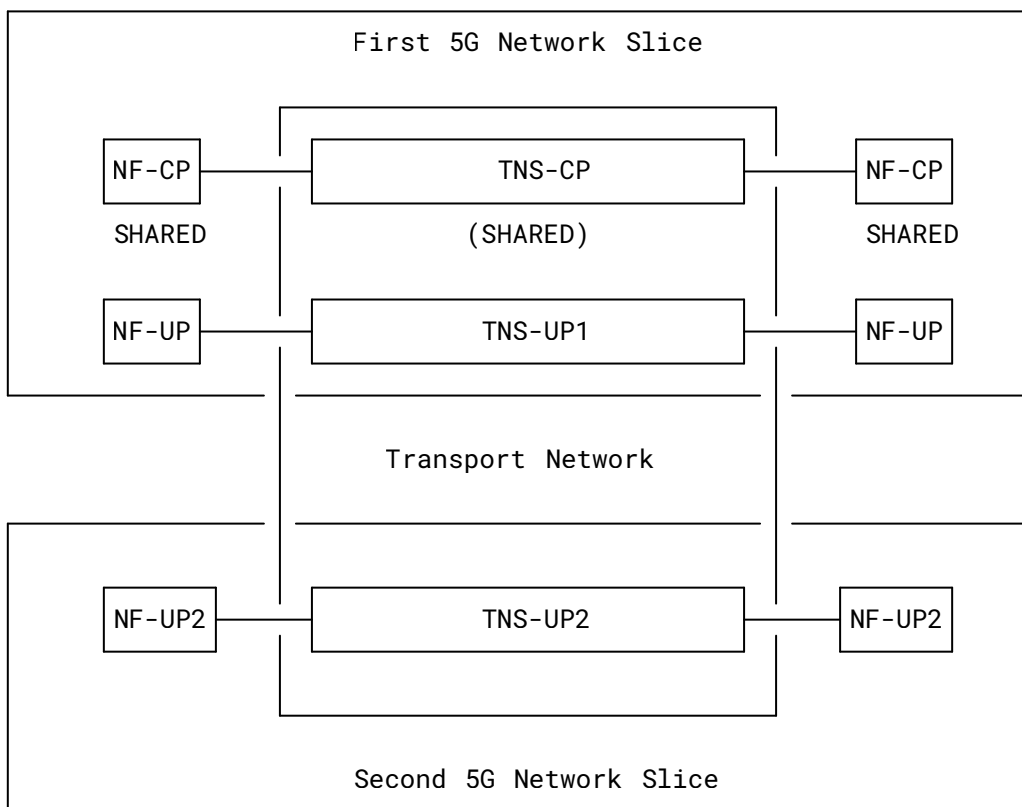


Figure 10: First and Subsequent Slice Deployment

Transport Network Slice mapping policies can be enforced by an operator (e.g., provided to a TN Orchestration or 5G NSO) to determine whether existing Transport Network Slices can be reused for handling a new Slice Service creation request. Providing such a policy is meant to better automate the realization of 5G Network Slices and minimize the realization delay that might be induced by extra cycles to seek for operator validation.

3.7. Overview of the Transport Network Realization Model

The realization model described in this document is depicted in [Figure 11](#). The following building blocks are used:

- L2VPN [[RFC4664](#)] and/or L3VPN [[RFC4364](#)] service instances for logical separation:

This realization model of transport for 5G Network Slices assumes Layer 3 delivery for midhaul and backhaul transport connections and a Layer 2 or Layer 3 delivery for fronthaul connections. Enhanced Common Public Radio Interface (eCPRI) [[ECPRI](#)] supports both delivery models. L2VPN/L3VPN service instances might be used as a basic form of logical slice separation. Furthermore, using service instances results in an additional outer header (as packets are encapsulated/decapsulated at the nodes hosting service instances), providing clean discrimination between 5G QoS and TN QoS, as explained in [Section 5](#).

Note that a variety of L2VPN mechanisms can be considered for slice realization. A non-comprehensive list is provided below:

- Virtual Private LAN Service (VPLS) [[RFC4761](#)] [[RFC4762](#)]
- Virtual Private Wire Service (VPWS) ([Section 3.1.1](#) of [[RFC4664](#)])
- Various flavors of EVPNs:
 - VPWS EVPN [[RFC8214](#)],
 - Provider Backbone Bridging combined with EVPN (PBB-EVPN) [[RFC7623](#)],
 - EVPN over MPLS [[RFC7432](#)], and
 - EVPN over Virtual Extensible LAN (VXLAN) [[RFC8365](#)].

The use of VPNs for realizing network slices is briefly described in [Appendix A.4](#) of [[RFC9543](#)].

- Fine-grained resource control at the PE:

This is sometimes called "admission control" or "traffic conditioning". The main purpose is the enforcement of the bandwidth contract for the slice right at the edge of the provider network where the traffic is handed off between the customer site and the provider network.

The method used here is granular ingress policing (rate limiting) to enforce contracted bandwidths per slice and, potentially, per traffic class within the slice. Traffic above the enforced rate might be immediately dropped or marked as high drop-probability traffic, which is more likely to be dropped somewhere inside the provider network if congestion occurs. In the egress direction at the PE node, hierarchical schedulers/shapers can be deployed, providing guaranteed rates per slice, as well as guarantees per traffic class within each slice.

For managed CEs, edge admission control can be distributed between CEs and PEs, where part of the admission control is implemented on the CE and the other part on the PE.

- Coarse-grained resource control at the transit links (non-attachment circuits) in the provider network, using a single NRP (called "base NRP" in Figure 11), spanning the entire provider network. Transit nodes in the provider network do not maintain any state of individual slices. Instead, only a flat (non-hierarchical) QoS model is used on transit links in the provider network, with up to 8 traffic classes. At the PE, traffic flows from multiple Slice Services are mapped to the limited number of traffic classes used on transit links in the provider network.
- Capacity planning/management for efficient usage of provider network resources:

The role of capacity planning/management is to ensure the provider network capacity can be utilized without causing any bottlenecks. The methods used here can range from careful network planning that ensures a more or less equal traffic distribution (i.e., equal-cost load balancing) to advanced TE techniques, with or without bandwidth reservations, that force more consistent load distribution, even in non-ECMP-friendly network topologies. See also Section 8 of [RFC9522].

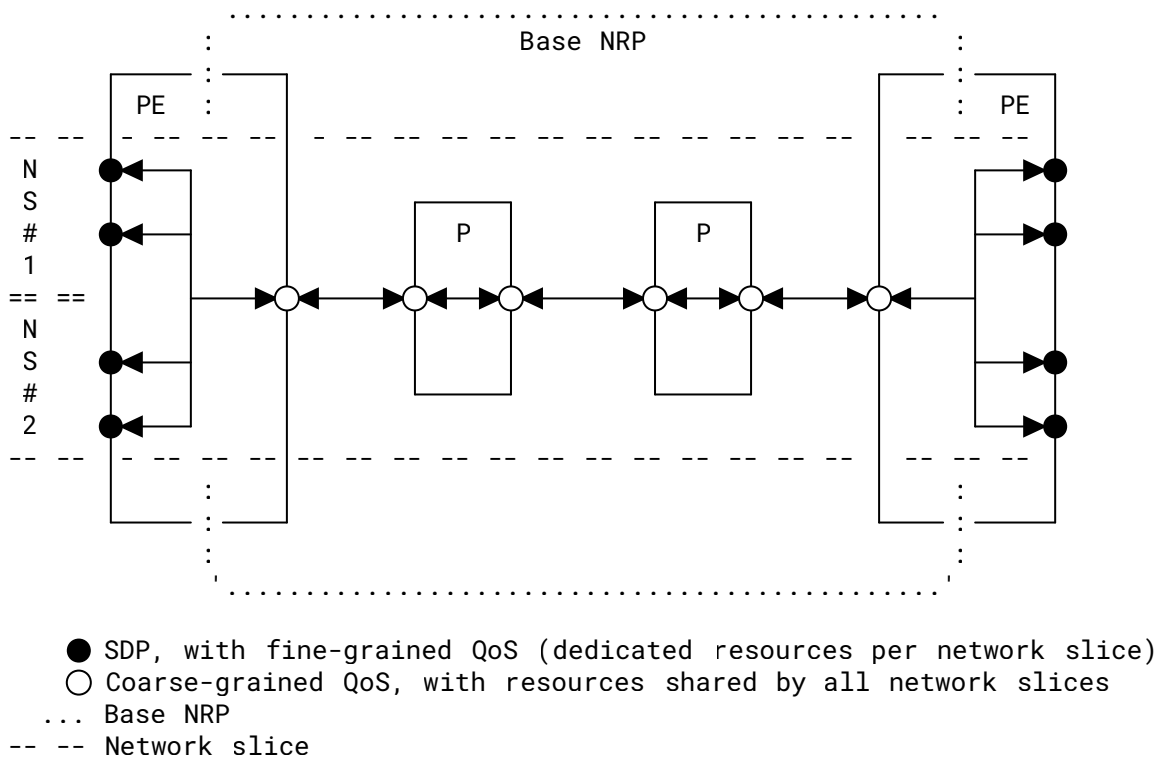


Figure 11: Resource Allocation Slicing Model with a Single NRP

The P nodes shown in Figure 11 are routers that do not interface with customer devices. See Section 5.3.1 of [RFC4026].

(with the exception of the F2 fronthaul interface, where eCPRI with Ethernet encapsulation is used), this VLAN is typically not transported across the provider network. Typically, it has only local significance at a particular SDP. For simplification, a deployment may rely on the same VLAN identifier for all ACs. However, that may not be always possible. As such, SDPs for the same slice at different locations may use different VLAN values. Therefore, a table mapping VLANs to RFC 9543 Network Slices is maintained for each AC, and the VLAN allocation is coordinated between customer orchestration and provider orchestration.

While VLAN handoff is simple for NFs, it adds complexity at the provider network because of the requirement of maintaining mapping tables for each SDP and performing a configuration task for new VLANs and IP subnet for every slice on every AC.

4.2. IP Handoff

In this option, an explicit mapping between source/destination IP addresses and a slice's specific S-NSSAI is used. The mapping can have either local (e.g., pertaining to a single NF attachment) or global TN significance. The mapping can be realized in multiple ways, including (but not limited to):

- S-NSSAI to a dedicated IP address for each NF
- S-NSSAI to a pool of IP addresses for global TN deployment
- S-NSSAI to a subset of bits of an IP address
- S-NSSAI to a DSCP value
- S-NSSAI to SRv6 Locators or Segment Identifiers (SIDs) [[RFC8986](#)]
- Use of a deterministic algorithm to map S-NSSAI to an IP subnet, prefix, or pools. For example, adaptations to the algorithm defined in [[RFC7422](#)] may be considered.

Mapping S-NSSAIs to IP addresses makes IP addresses an identifier for slice-related policy enforcement in the Transport Network (e.g., differentiated services, traffic steering, bandwidth allocation, security policies, and monitoring).

One example of the IP handoff realization is the arrangement in which the slices in the TN domain are instantiated using IP tunnels (e.g., IPsec or GTP-U tunnels) established between NFs, as depicted in [Figure 13](#). The transport for a single 5G Network Slice might be constructed with multiple such tunnels, since a typical 5G Network Slice contains many NFs, especially DUs and CUs. If a shared NF (i.e., an NF that serves multiple slices, such as a shared DU) is deployed, multiple tunnels from the shared NF are established, each tunnel representing a single slice.

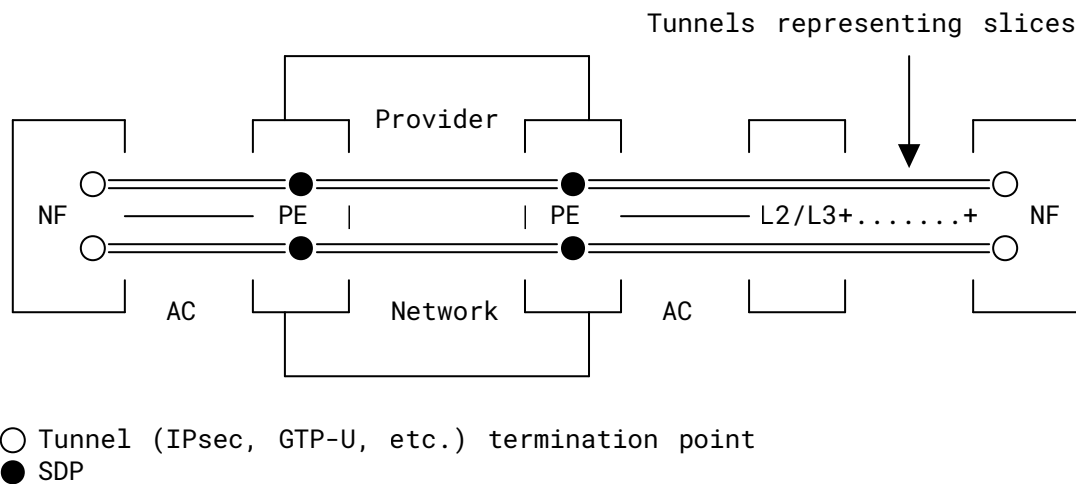


Figure 13: Example of 5G Network Slice with IP Handoff Providing End-to-End Connectivity

As opposed to the VLAN handoff case (Section 4.1), there is no logical interface representing a slice on the PE; hence, all slices are handled within a single service instance. The IP and VLAN handoffs are not mutually exclusive but instead could be used concurrently. Since the TN doesn't recognize S-NSSAIs, a mapping table similar to the VLAN handoff solution is needed (Section 4.1).

The mapping table can be simplified if, for example, IPv6 addressing is used to address NFs. An IPv6 address is a 128-bit field, while the S-NSSAI is a 32-bit field: The Slice/Service Type (SST) is 8 bits, and the Slice Differentiator (SD) is 24 bits. Out of the 128 bits of the IPv6 address, 32 bits may be used to encode the S-NSSAI, which makes an IP-to-slice mapping table unnecessary.

The S-NSSAI/IPv6 mapping is a local IPv6 address allocation method to NFs not disclosed to on-path nodes. IP forwarding is not altered by this method and is still achieved following BCP 198 [RFC7608]. Intermediary TN nodes are not required to associate any additional semantic with the IPv6 address.

However, operators using such mapping methods should be aware of the implications of any change of S-NSSAI on the IPv6 addressing plans. For example, modifications of the S-NSSAIs in use will require updating the IP addresses used by NFs involved in the associated slices.

An example of a local IPv6 addressing plan for NFs is provided in Appendix A.

4.3. MPLS Label Handoff

In this option, the service instances representing different slices are created directly on the NF, or within the customer site hosting the NF, and attached to the provider network. Therefore, the packet is encapsulated outside the provider network with MPLS encapsulation or MPLS-in-UDP encapsulation [RFC7510], depending on the capability of the customer site, with the service label depicting the slice.

There are three major methods (based upon [Section 10](#) of [[RFC4364](#)]) for interconnecting MPLS services over multiple service domains:

Option A ([Section 4.3.1](#)): VRF-to-VRF connections.

Option B ([Section 4.3.2](#)): Redistribution of labeled VPN routes with next-hop change at domain boundaries.

Option C ([Section 4.3.3](#)): Redistribution of labeled VPN routes without next-hop change and redistribution of labeled transport routes with next-hop change at domain boundaries.

[Figure 14](#) illustrates the use of service-aware CE ([Section 3.3.2](#)) for the deployment discussed in [Sections 4.3.2](#) and [4.3.3](#).

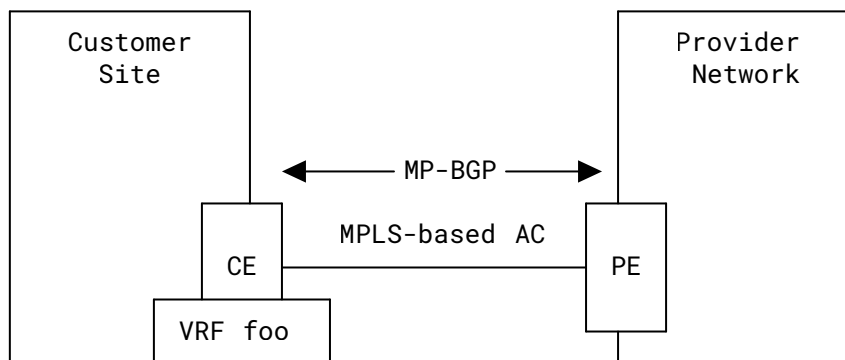


Figure 14: Example of MPLS-Based Attachment Circuit

4.3.1. Option A

This option is based on the VLAN handoff, described in [Section 4.1](#); it is not based on the MPLS label handoff.

4.3.2. Option B

In this option, L3VPN service instances are instantiated outside the provider network. These L3VPN service instances are instantiated in the customer site, which could be, for example, either on the compute that hosts mobile NFs ([Figure 15](#), left-hand side) or within the DC/cloud infrastructure itself (e.g., on the top of the rack or leaf switch within cloud IP fabric ([Figure 15](#), right-hand side)). On the AC connected to a PE, packets are already MPLS encapsulated (or MPLS-in-UDP/MPLS-in-IP encapsulated, if cloud or compute infrastructure don't support MPLS encapsulation). Therefore, the PE uses neither a VLAN nor an IP address for slice identification at the SDP but instead uses the MPLS label.

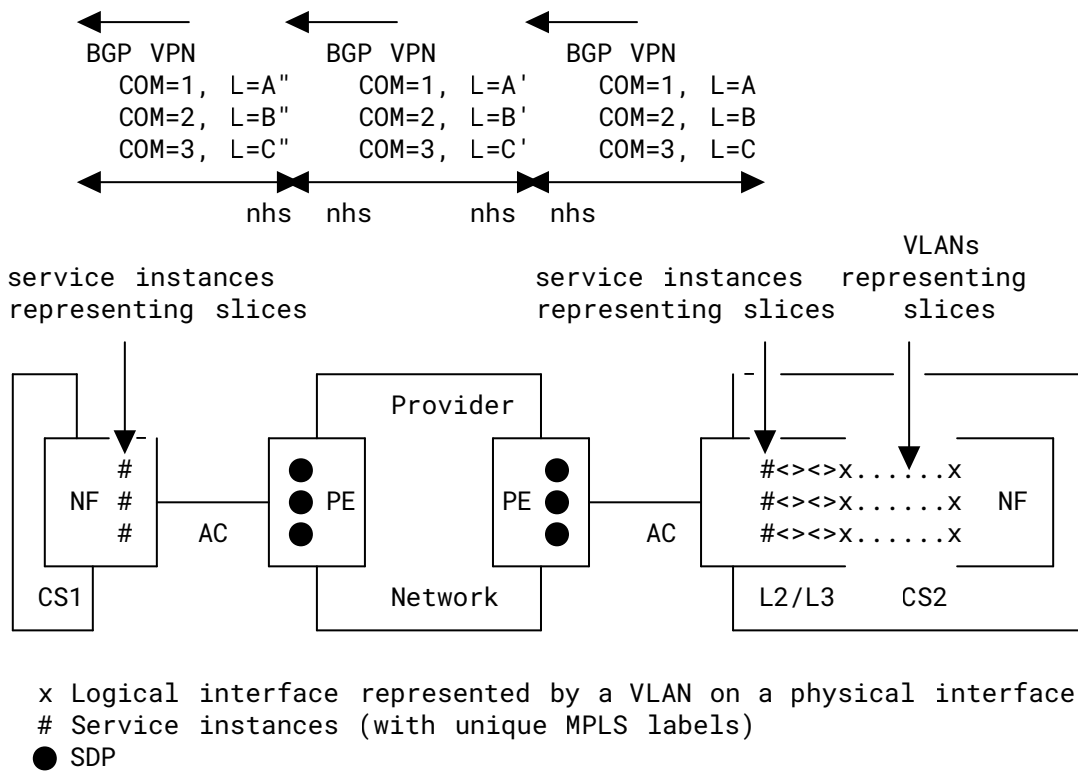


Figure 15: Example of MPLS Handoff with Option B

MPLS labels are allocated dynamically in Option B deployments, where, at the domain boundaries, service prefixes are reflected with next-hop self (nhs), and a new label is dynamically allocated, as shown in Figure 15 (e.g., labels A, A', and A" for the first depicted slice). Therefore, for any slice-specific per-hop behavior at the provider network edge, the PE needs to determine which label represents which slice. In the BGP control plane, when exchanging service prefixes over an AC, each slice might be represented by a unique BGP community, so tracking label assignment to the slice might be possible. For example, in Figure 15, for the slice identified with COM=1, the PE advertises a dynamically allocated label A". Since, based on the community, the label-to-slice association is known, the PE can use this dynamically allocated label A" to identify incoming packets as belonging to "slice 1" and execute appropriate edge per-hop behavior.

It is worth noting that slice identification in the BGP control plane might be with per-prefix granularity. In the extreme case, each prefix can have a different community representing a different slice. Depending on the business requirements, each slice could be represented by a different service instance as outlined in Figure 15. In that case, the route target extended community (Section 4 of [RFC4360]) might be used as a slice differentiator. In other deployments, all prefixes (representing different slices) might be handled by a single "mobile" service instance, and some other BGP attribute (e.g., a standard community [RFC1997]) might be

used for slice differentiation. There could also be a deployment option that groups multiple slices together into a single service instance, resulting in a handful of service instances. In any case, fine-grained per-hop behavior at the edge of provider network is possible.

4.3.3. Option C

Option B relies upon exchanging service prefixes between customer sites and the provider network. This may lead to scaling challenges in large-scale 5G deployments as the PE node needs to carry all service prefixes. To alleviate this scaling challenge, in Option C, service prefixes are exchanged between customer sites only. In doing so, the provider network is offloaded from carrying, propagating, and programming appropriate forwarding entries for service prefixes.

Option C relies upon exchanging service prefixes via multi-hop BGP sessions between customer sites, without changing the NEXT_HOP BGP attribute. Additionally, IPv4/IPv6 labeled unicast (SAFI-4) host routes, used as NEXT_HOP for service prefixes, are exchanged via direct single-hop BGP sessions between adjacent nodes in a customer site and a provider network, as depicted in [Figure 16](#). As a result, a node in a customer site performs hierarchical next-hop resolution.

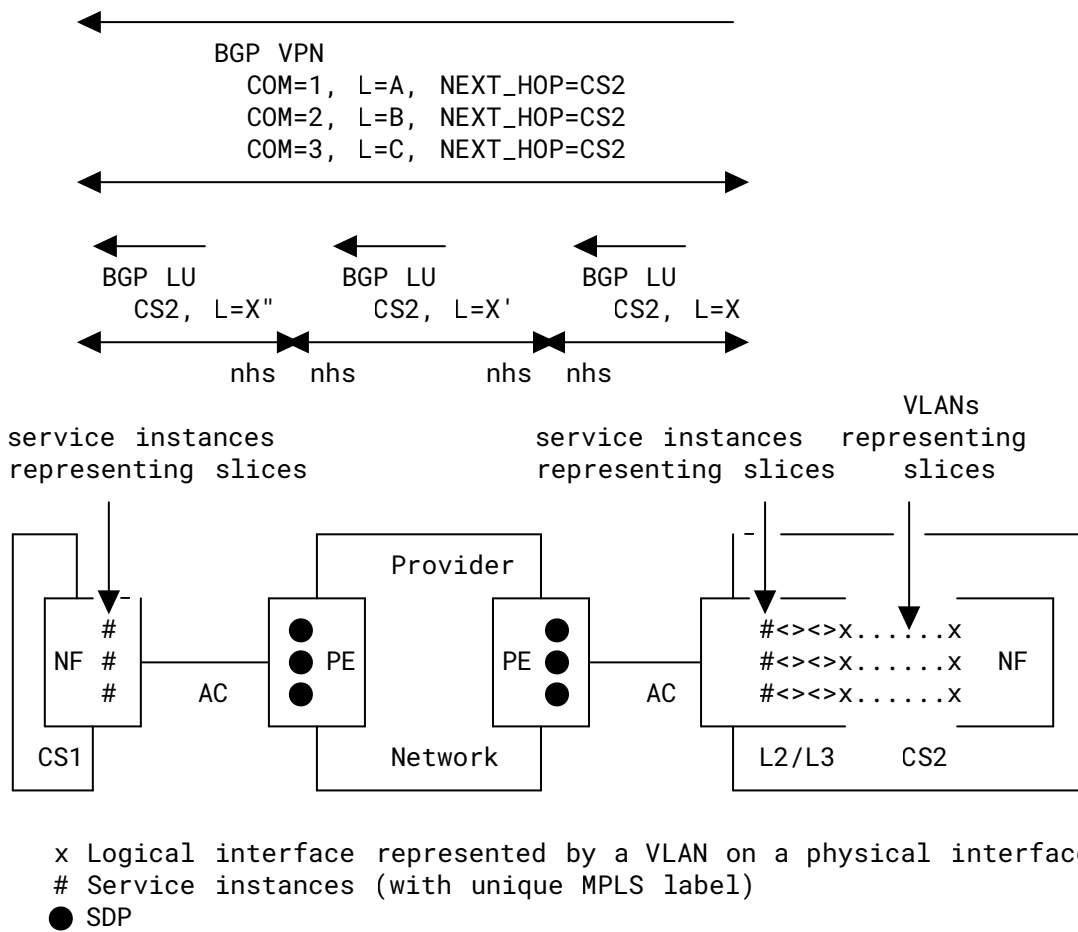


Figure 16: Example of MPLS Handoff with Option C

This architecture requires an end-to-end Label Switched Path (LSP) leading from a packet's ingress node inside one customer site to its egress inside another customer site, through a provider network. Hence, at the domain (customer site and provider network) boundaries, the NEXT_HOP attribute for IPv4/IPv6 labeled unicast needs to be modified to next-hop self (nhs), which results in a new IPv4/IPv6 labeled unicast label allocation. Appropriate forwarding entries for label swaps for IPv4/IPv6 labeled unicast labels are programmed in the data plane. There is no additional "labeled transport" protocol on the AC (e.g., no LDP, RSVP, or SR).

Packets are transmitted over the AC with the IPv4/IPv6 labeled unicast as the top label, with the service label deeper in the label stack. In Option C, the service label is not used for forwarding lookup on the PE. This significantly lowers the scaling pressure on PEs, as PEs need to program forwarding entries only for IPv4/IPv6 labeled unicast host routes, which are used as NEXT_HOP for service prefixes. Also, since one IPv4/IPv6 labeled unicast host route represents one customer site, regardless of the number of slices in the customer site, the number of forwarding entries on a PE is considerably reduced.

For any slice-specific per-hop behavior at the provider network edge, as described in detail in [Section 3.7](#), the PE needs to determine which label in the packet represents which slice. This can be achieved, for example, by allocating non-overlapping service label ranges for each slice and using those ranges for slice identification purposes on the PE.

5. QoS Mapping Realization Models

5.1. QoS Layers

The resources are managed via various QoS policies deployed in the network. QoS mapping models to support 5G slicing connectivity implemented over a packet switched provider network use two layers of QoS, which are discussed in the following subsections.

5.1.1. 5G QoS Layer

QoS treatment is indicated in the 5G QoS layer by the 5G QoS Indicator (5QI), as defined in [\[TS-23.501\]](#). The 5QI is an identifier that is used as a reference to 5G QoS characteristics (e.g., scheduling weights, admission thresholds, queue management thresholds, and link-layer protocol configuration) in the RAN domain. Given that 5QI applies to the RAN domain, it is not visible to the provider network. Therefore, if 5QI-aware treatment is desired in the provider network, 5G network functions might set DSCP with a value representing 5QI so that differentiated treatment can be implemented in the provider network as well. Based on these DSCP values, very granular QoS enforcement might be implemented at the SDP of each provider network segment used to construct transport for given 5G Network Slice.

The exact mapping between 5QI and DSCP is out of scope for this document. Mapping recommendations are documented, e.g., in [\[MAPPING\]](#).

Each Slice Service might have flows with multiple 5QIs. 5QIs (or, more precisely, corresponding DSCP values) are visible to the provider network at SDPs (i.e., at the edge of the provider network).

In this document, this layer of QoS is referred to as "5G QoS Class" ("5G QoS" in short) or "5G DSCP".

5.1.2. Transport Network (TN) QoS Layer

Control of the TN resources and traffic scheduling/prioritization on provider network transit links are based on a flat (non-hierarchical) QoS model in this network slice realization. That is, RFC 9543 Network Slices are assigned dedicated resources (e.g., QoS queues) at the edge of the provider network (at SDPs), while all RFC 9543 Network Slices are sharing resources (sharing QoS queues) on the transit links of the provider network. Typical router hardware can support up to 8 traffic queues per port; therefore, this document assumes support for 8 traffic queues per port in general.

At this layer, QoS treatment is indicated by a QoS indicator specific to the encapsulation used in the provider network. Such an indicator may be a DSCP or MPLS Traffic Class (TC). This layer of QoS is referred to as "TN QoS Class" ("TN QoS" for short) in this document.

5.2. QoS Realization Models

While 5QI might be exposed to the provider network via the DSCP value (corresponding to a specific 5QI value) set in the IP packet generated by NFs, some 5G deployments might use 5QI in the RAN domain only, without requesting per-5QI differentiated treatment from the provider network. This might be due to an NF limitation (e.g., no capability to set DSCP), or it might simply depend on the overall slicing deployment model. The O-RAN Alliance, for example, defines a phased approach to the slicing, with initial phases utilizing only per-slice, but not per-5QI, differentiated treatment in the TN domain (see Annex F of [O-RAN.WG9.XPSAAS]).

Therefore, from a QoS perspective, the 5G slicing connectivity realization defines two high-level realization models for slicing in the TN domain: a 5QI-unaware model and a 5QI-aware model. Both slicing models in the TN domain could be used concurrently within the same 5G Network Slice. For example, the TN segment for 5G midhaul (F1-U interface) might be 5QI-aware, while at the same time, the TN segment for 5G backhaul (N3 interface) might follow the 5QI-unaware model.

These models are further elaborated in the following two subsections.

5.2.1. 5QI-Unaware Model

In the 5QI-unaware model, the DSCP values in the packets received from NF at SDP are ignored. In the provider network, there is no QoS differentiation at the 5G QoS Class level. The entire RFC 9543 Network Slice is mapped to a single TN QoS Class and therefore to a single QoS queue on the routers in the provider network. With a low number of deployed 5G Network Slices (for example, only two 5G Network Slices: eMBB and MIIoT), it is possible to dedicate a separate QoS queue for each slice on transit routers in the provider network. However, with the introduction of private/enterprises slices, as the number of 5G Network Slices (and thus the corresponding RFC 9543 Network Slices) increases, a single QoS queue on transit links in the provider network serves multiple slices with similar characteristics. QoS enforcement on transit links is fully coarse-grained (single NRP, sharing resources among all RFC 9543 Network Slices), as displayed in [Figure 17](#).

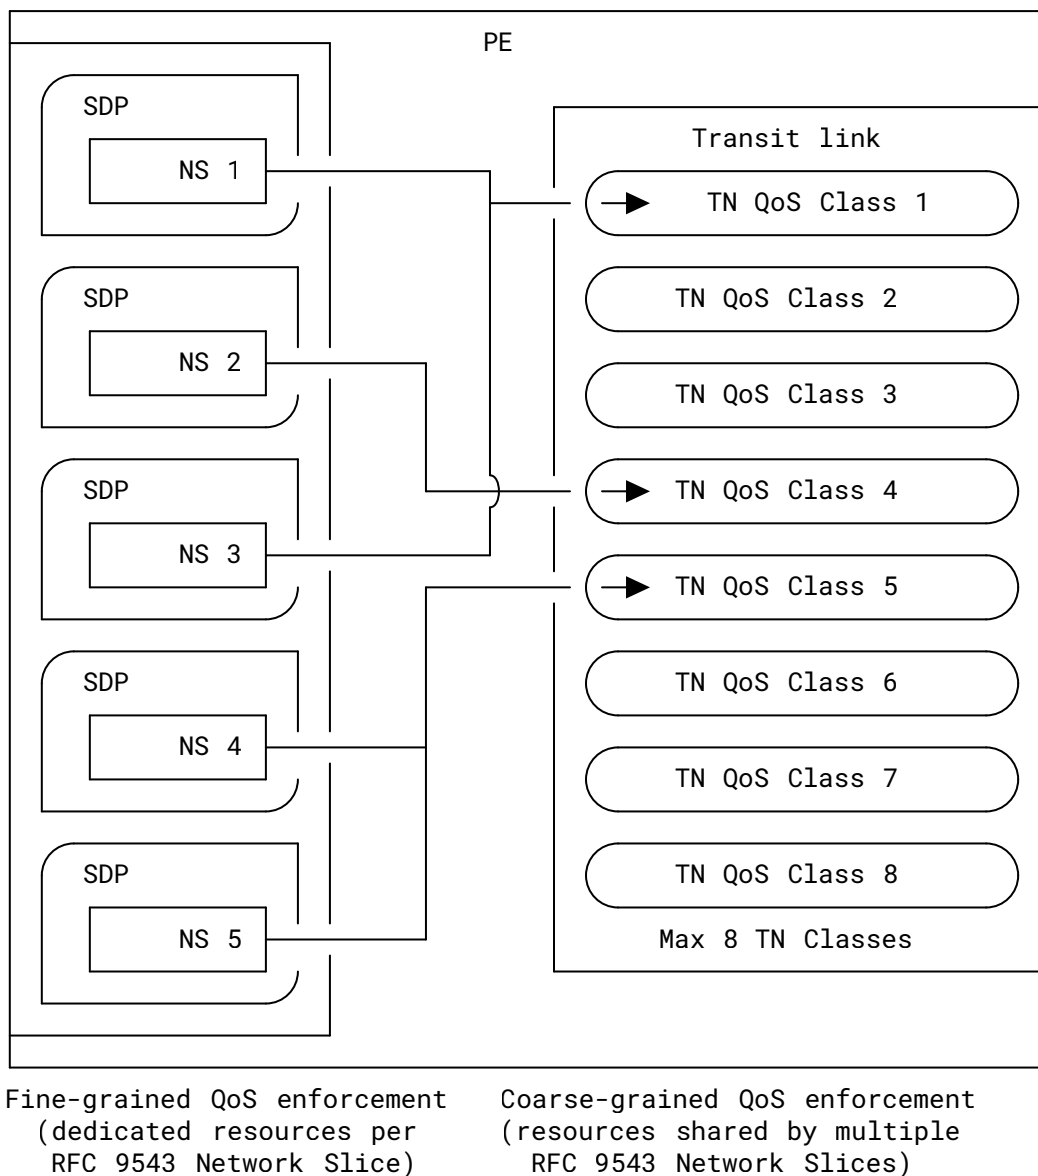


Figure 17: Mapping of Slice to TN QoS (5QI-Unaware Model)

When the IP traffic is handed over at the SDP from the AC to the provider network, the PE encapsulates the traffic into MPLS (if MPLS transport is used in the provider network) or IPv6, optionally with some additional headers (if SRv6 transport is used in the provider network), and sends out the packets on the provider network transit link.

The original IP header retains the DSCP marking (which is ignored in the 5QI-unaware model), while the new header (MPLS or IPv6) carries the QoS marking (MPLS Traffic Class bits for MPLS encapsulation or DSCP for SRv6/IPv6 encapsulation) related to the TN Class of Service (CoS). Based on the TN CoS marking, per-hop behavior for all RFC 9543 Network Slices is executed on

provider network transit links. Provider network transit routers do not evaluate the original IP header for QoS-related decisions. This model is outlined in [Figure 18](#) for MPLS encapsulation and in [Figure 19](#) for SRv6 encapsulation.

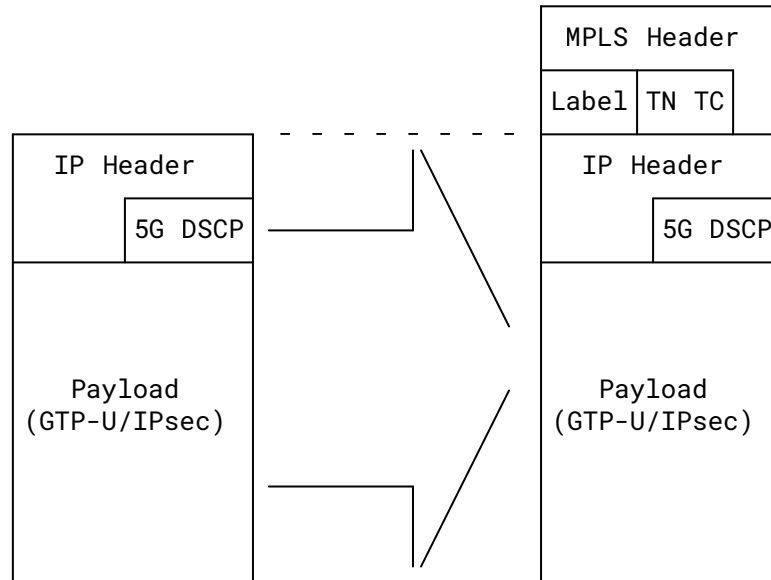


Figure 18: QoS with MPLS Encapsulation

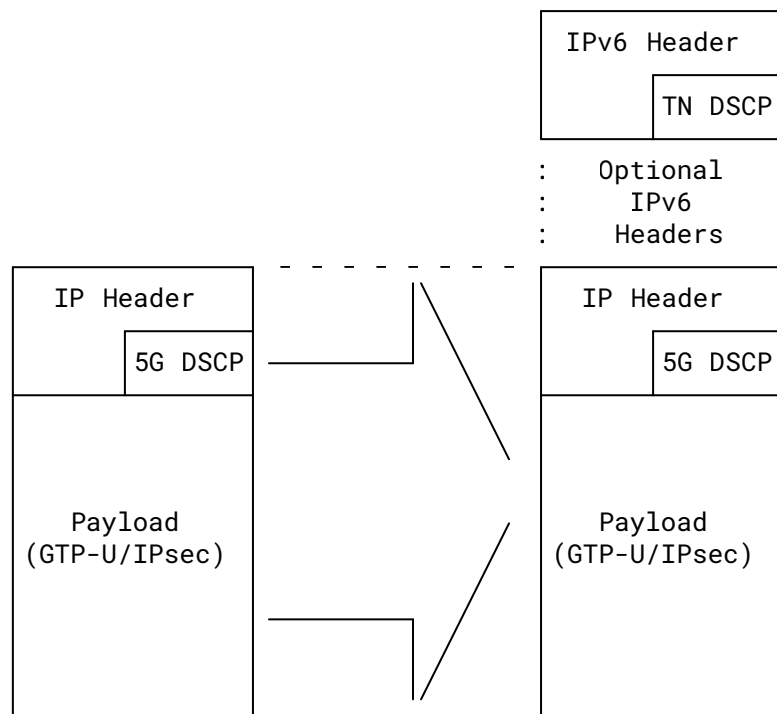


Figure 19: QoS with SRv6 Encapsulation

From a QoS perspective, both options are similar. However, there is one difference between the two options. The MPLS TC is only 3 bits (8 possible combinations), while DSCP is 6 bits (64 possible combinations). Hence, SRv6 provides more flexibility for TN CoS design, especially in combination with soft policing with in-profile and out-of-profile traffic, as discussed in [Section 5.2.1.1](#).

Provider network edge resources are controlled in a fine-grained manner, with dedicated resource allocation for each RFC 9543 Network Slice. Resource control and enforcement happens at each SDP in two directions: inbound and outbound.

5.2.1.1. Inbound Edge Resource Control

The main aspect of inbound provider network edge resource control is per-slice traffic volume enforcement. This kind of enforcement is often called "admission control" or "traffic conditioning". The goal of this inbound enforcement is to ensure that the traffic above the contracted rate is dropped or deprioritized, depending on the business rules, right at the edge of provider network. This, combined with appropriate network capacity planning/management ([Section 7](#)), is required to ensure proper isolation between slices in a scalable manner. As a result, traffic of one slice has no influence on the traffic of other slices, even if the slice is misbehaving (e.g., Distributed Denial-of-Service (DDoS) attacks or node/link failures) and generates traffic volumes above the contracted rates.

The slice rates can be characterized with the following parameters [NSSM]:

- CIR: Committed Information Rate (i.e., guaranteed bandwidth)
- PIR: Peak Information Rate (i.e., maximum bandwidth)

These parameters define the traffic characteristics of the slice and are part of the SLO parameter set provided by the 5G NSO to an NSC. Based on these parameters, the provider network's inbound policy can be implemented using one of following options:

- 1r2c (single-rate two-color) rate limiter

This is the most basic rate limiter, described in [Section 2.3](#) of [RFC2475] (though not termed "1r2c" in that document). At the SDP, it meters a traffic stream of a given slice and marks its packets as in-profile (below CIR being enforced) or out-of-profile (above CIR being enforced). In-profile packets are accepted and forwarded. Out-of-profile packets are either dropped right at the SDP (hard rate limiting) or re-marked (with different MPLS TC or DSCP TN markings) to signify "this packet should be dropped in the first place, if there is congestion" (soft rate limiting), depending on the business policy of the provider network. In the latter case, while packets above CIR are forwarded at the SDP, they are subject to being dropped during any congestion event at any place in the provider network.

- 2r3c (two-rate three-color) rate limiter

This was initially defined in [RFC2698], and an improved version is defined in [RFC4115]. In essence, the traffic is assigned to one of the these three categories:

- Green, for traffic under CIR
- Yellow, for traffic between CIR and PIR
- Red, for traffic above PIR

An inbound 2r3c meter implemented with [RFC4115], compared to [RFC2698], is more "customer friendly" as it doesn't impose outbound peak-rate shaping requirements on CE devices. In general, 2r3c meters give greater flexibility for provider network edge enforcement regarding accepting the traffic (green), deprioritizing and potentially dropping the traffic on transit during congestion (yellow), or hard-dropping the traffic (red).

Inbound provider network edge enforcement for the 5QI-unaware model, where all packets belonging to the slice are treated the same way in the provider network (no 5G QoS Class differentiation in the provider), is outlined in [Figure 20](#).

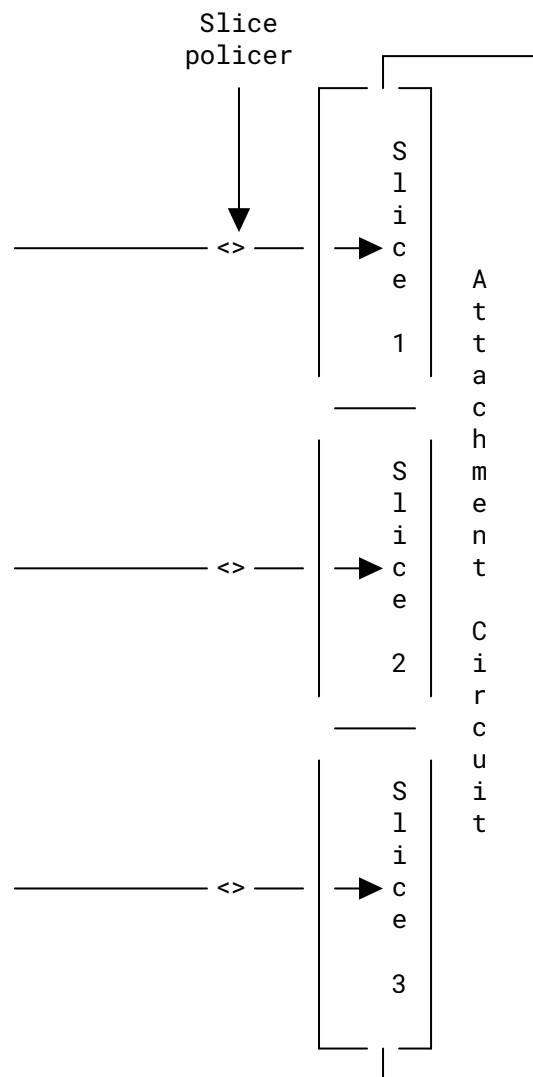


Figure 20: Ingress Slice Admission Control (5QI-Unaware Model)

5.2.1.2. Outbound Edge Resource Control

While inbound slice admission control at the provider network edge is mandatory in the architecture described in this document, outbound provider network edge resource control might not be required in all use cases. Use cases that specifically call for outbound provider network edge resource control are:

- Slices use both CIR and PIR parameters, and provider network edge links (ACs) are dimensioned to fulfill the aggregate of slice CIRs. If, at any given time, some slices send the traffic above CIR, congestion in the outbound direction on the provider network edge link (AC) might happen. Therefore, fine-grained resource control to guarantee at least CIR for each slice is required.

- Any-to-Any (A2A) connectivity constructs are deployed, again resulting in potential congestion in the outbound direction on the provider network edge links, even if only slice CIR parameters are used. This again requires fine-grained resource control per slice in the outbound direction at the provider network edge links.

As opposed to inbound provider network edge resource control, typically implemented with rate-limiters/policers, outbound resource control is typically implemented with a weighted/priority queuing, potentially combined with optional shapers (per slice). A detailed analysis of different queuing mechanisms is out of scope for this document but is provided in [\[RFC7806\]](#).

[Figure 21](#) outlines the outbound provider network edge resource control model for 5QI-unaware slices. Each slice is assigned a single egress queue. The sum of slice CIRs, used as the weight in weighted queueing model, should not exceed the physical capacity of the AC. Slice requests above this limit should be rejected by the NSC, unless an already-established slice with lower priority, if such exists, is preempted.

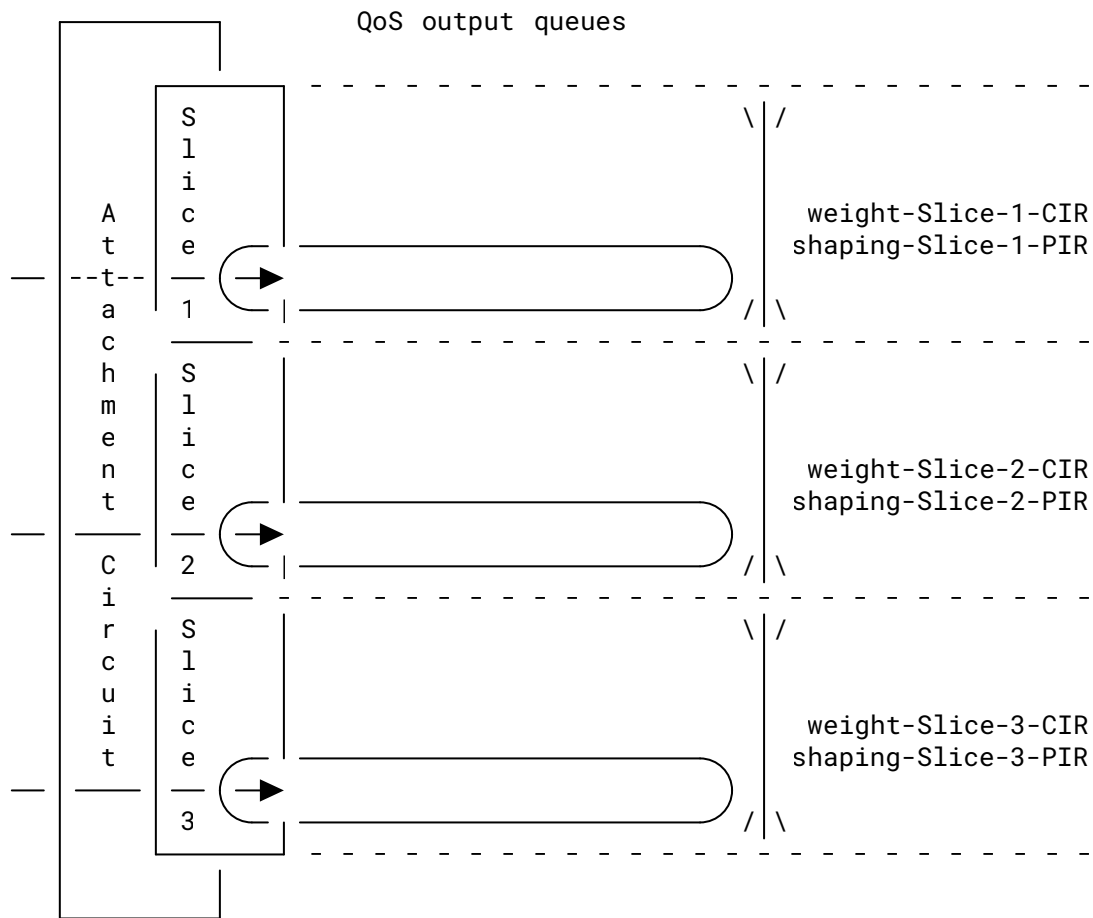


Figure 21: Ingress Slice Admission Control (5QI-Unaware Model) - Output

5.2.2. 5QI-Aware Model

In the 5QI-aware model, a potentially large number of 5G QoS Classes, represented via the DSCP set by NFs (the architecture scales to thousands of 5G Network Slices), is mapped (multiplexed) to up to 8 TN QoS Classes used in a provider network transit equipment, as outlined in [Figure 22](#).

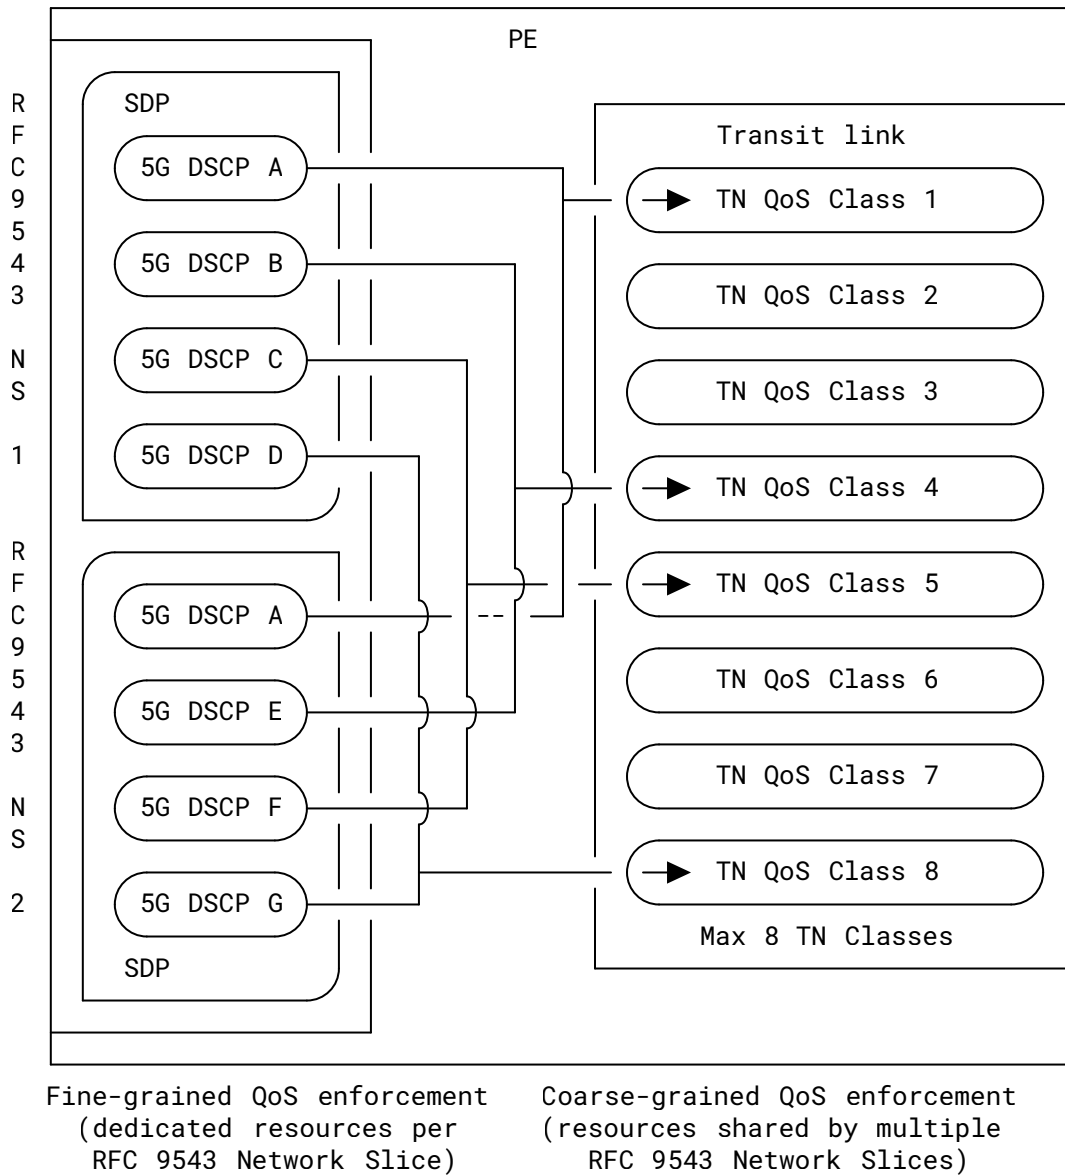


Figure 22: Mapping of Slice 5G QoS to TN QoS (5QI-Aware Model)

Given that in deployments with a large number of 5G Network Slices, the number of potential 5G QoS Classes is much higher than the number of TN QoS Classes, multiple 5G QoS Classes with similar characteristics -- potentially from different slices -- would be grouped with common operator-defined TN logic and mapped to the same TN QoS Class when transported in the provider network. That is, common Per-Hop Behavior (PHB) [RFC2474] is executed on transit provider network routers for all packets grouped together. An example of this approach is outlined in Figure 23. A provider may decide to implement Diffserv-Intercon PHBs at the boundaries of its network domain [RFC8100].

Note: The numbers indicated in Figure 23 (S-NSSAI, 5QI, DSCP, queue, etc.) are provided for illustration purposes only and should not be considered as deployment guidance.

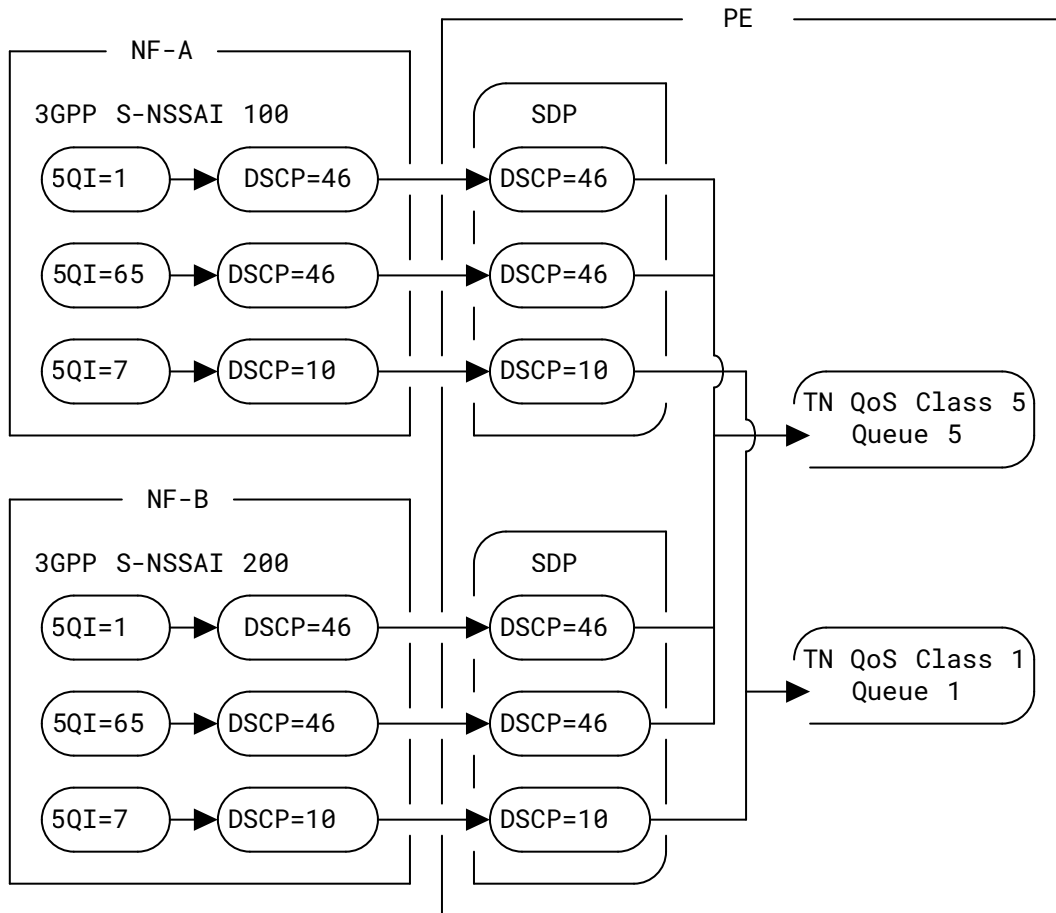


Figure 23: Example of 3GPP QoS Mapped to TN QoS

In current SDO progress of 3GPP (Release 19) and O-RAN, the mapping of 5QI to DSCP is not expected to be in a per-slice fashion, where 5QI-to-DSCP mapping may vary from 3GPP slice to 3GPP slice; hence, the mapping of 5G QoS DSCP values to TN QoS Classes may be rather common.

Like in the 5QI-unaware model, the original IP header retains the DSCP marking corresponding to 5QI (5G QoS Class), while the new header (MPLS or IPv6) carries the QoS marking related to TN QoS Class. Based on the TN QoS Class marking, per-hop behavior for all aggregated 5G QoS Classes from all RFC 9543 Network Slices is executed on the provider network transit links.

Provider network transit routers do not evaluate the original IP header for QoS-related decisions. The original DSCP marking retained in the original IP header is used at the PE for fine-grained inbound/outbound enforcement per slice and per 5G QoS Class on the AC.

In the 5QI-aware model, compared to the 5QI-unaware model, provider network edge resources are controlled in an even more fine-grained manner, with dedicated resource allocation for each RFC 9543 Network Slice and for a number of traffic classes (most commonly, up to 4 or 8 traffic classes, depending on the hardware capability of the equipment) within each RFC 9543 Network Slice.

5.2.2.1. Inbound Edge Resource Control

Compared to the 5QI-unaware model, admission control (traffic conditioning) in the 5QI-aware model is more granular, as it not only enforces per-slice capacity constraints, but may also enforce the constraints per 5G QoS Class within each slice.

A 5G Network Slice using multiple 5QIs can potentially specify rates in one of the following ways:

- Rates per traffic class (CIR or CIR+PIR), no rate per slice (sum of rates per class gives the rate per slice).
- Rate per slice (CIR or CIR+PIR), and rates per prioritized (premium) traffic classes (CIR only). A best-effort traffic class uses the bandwidth (within slice CIR/PIR) not consumed by prioritized classes.

In the first option, the slice admission control is executed with traffic class granularity, as outlined in [Figure 24](#). In this model, if a premium class doesn't consume all available class capacity, it cannot be reused by a non-premium class (i.e., best-effort).

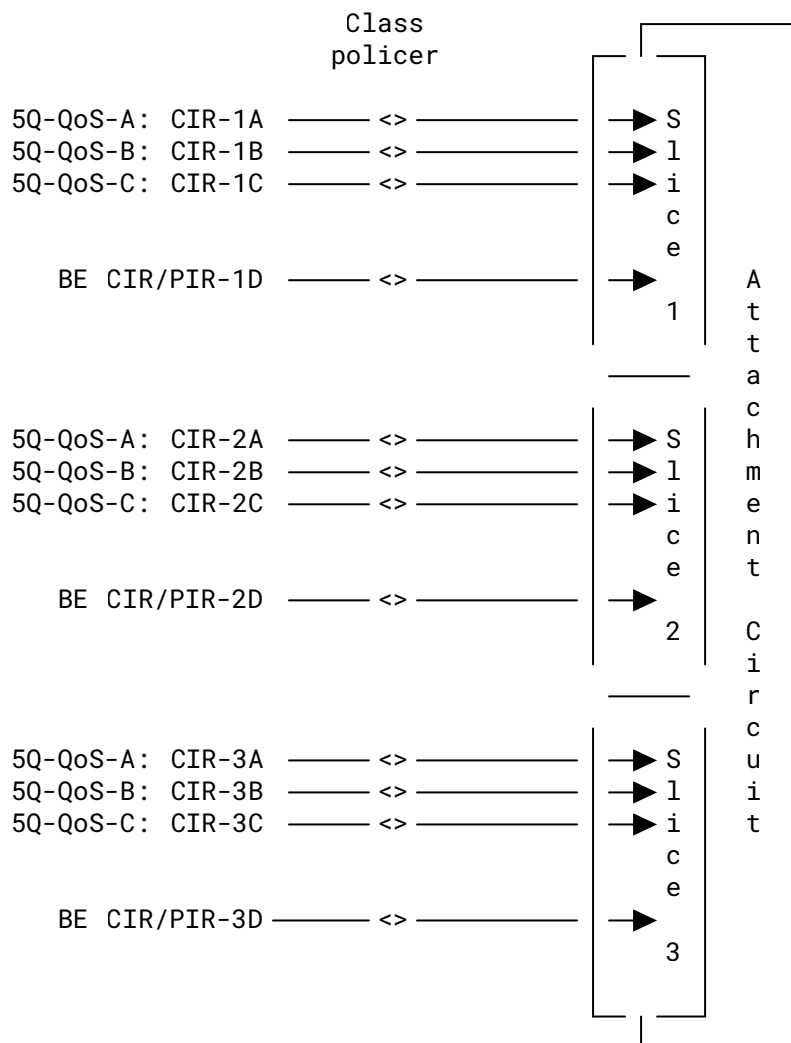


Figure 24: Ingress Slice Admission Control (5QI-Aware Model)

The second option combines the advantages of the 5QI-unaware model (per-slice admission control) with per-traffic-class admission control, as outlined in Figure 24. Ingress admission control is at class granularity for premium classes (CIR only). A non-premium class (i.e., best-effort class) has no separate class admission control policy, but it is allowed to use the entire slice capacity, which is available at any given moment (i.e., slice capacity, which is not consumed by premium classes). It is a hierarchical model, as depicted in Figure 25.

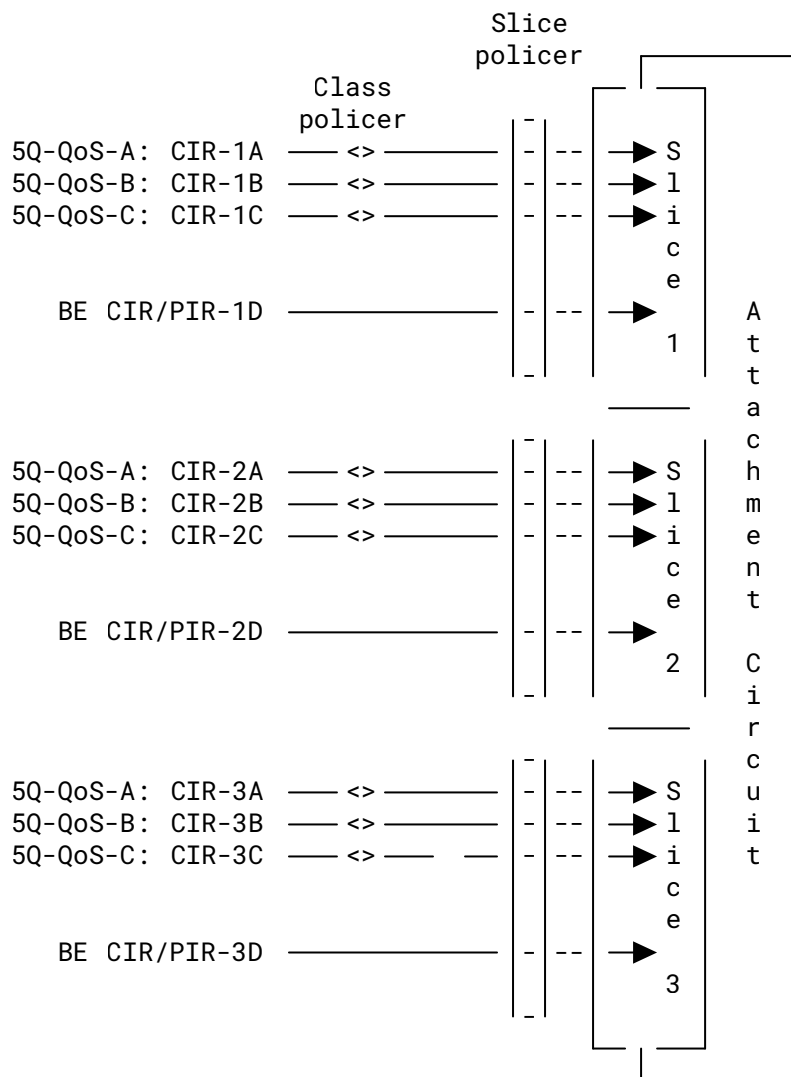


Figure 25: Ingress Slice Admission Control (5QI-Aware Model) - Hierarchical

5.2.2.2. Outbound Edge Resource Control

Figure 26 outlines the outbound edge resource control model at the Transport Network layer for 5QI-aware slices. Each slice is assigned multiple egress queues. The sum of queue weights, which are 5G QoS queue CIRs within the slice, should not exceed the CIR of the slice itself. And, similar to the 5QI-aware model, the sum of slice CIRs should not exceed the physical capacity of the AC.

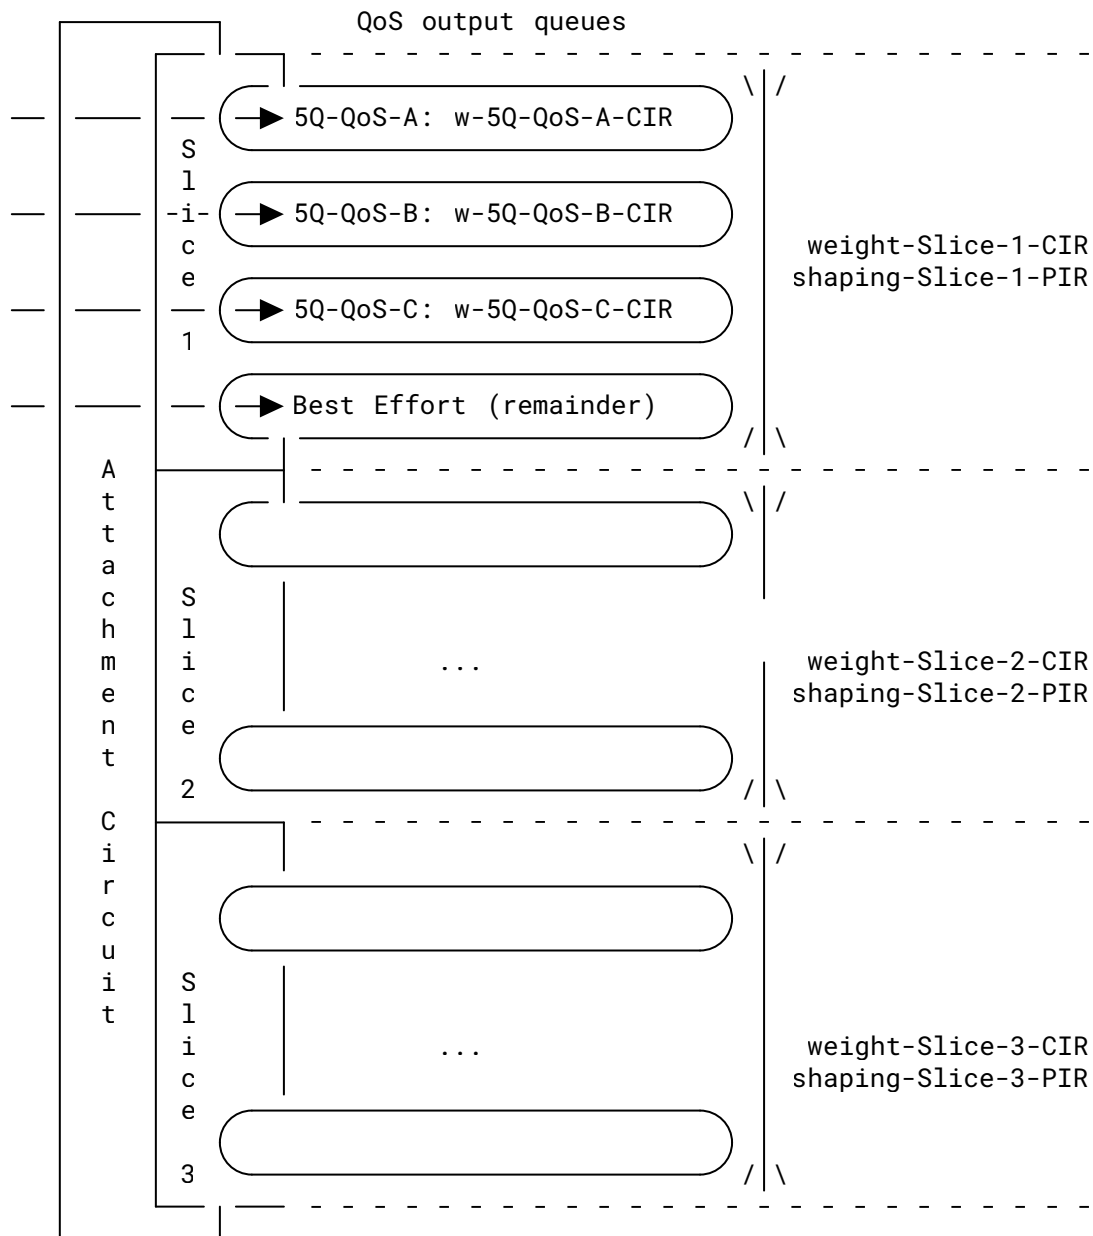


Figure 26: Egress Slice Admission Control (5QI-Aware Model)

5.3. Transit Resource Control

Transit resource control is much simpler than edge resource control in the provider network. As outlined in Figure 22, at the provider network edge, 5G QoS Class marking (represented by DSCP related to 5QI set by Mobile Network functions in the packets handed off to the TN) is mapped to the TN QoS Class. Based on TN QoS Class, when the packet is encapsulated with an outer header (MPLS or IPv6), the TN QoS Class marking (MPLS TC or IPv6 DSCP in outer header, as depicted in

Figures 18 and 19) is set in the outer header. PHB in provider network transit routers is based exclusively on that TN QoS Class marking, i.e., original 5G QoS Class DSCP is not taken into consideration on transit.

Provider network transit resource control does not use any inbound interface policy but only uses an outbound interface policy, which is based on the priority queue combined with a weighted or deficit queuing model, without any shaper. The main purpose of transit resource control is to ensure that during network congestion events (for example, events caused by network failures or temporary rerouting), premium classes are prioritized, and any drops only occur in traffic that was deprioritized by ingress admission control (see Section 5.2.1.1) or in non-premium (best-effort) classes. Capacity planning and management, as described in Section 7, ensures that enough capacity is available to fulfill all approved slice requests.

6. PE Underlay Transport Mapping Models

The PE underlay transport (underlay transport, for short) refers to a specific path forwarding behavior between PEs in order to provide packet delivery that is consistent with the corresponding SLOs. This realization step focuses on controlling the paths that will be used for packet delivery between PEs, independent of the underlying network resource partitioning.

It is worth noting that TN QoS Classes and underlay transport are each related to different engineering objectives. For example, the TN domain can be operated with 8 TN QoS Classes (representing 8 hardware queues in the routers) and two underlay transports (e.g., a latency-optimized underlay transport using link-latency metrics for path calculation and an underlay transport following IGP metrics). The TN QoS Class determines the per-hop behavior when the packets are transiting through the provider network, while underlay transport determines the paths for packets through the provider network based on the operator's requirements. This path can be optimized or constrained.

A network operator can define multiple underlay transports within a single NRP. An underlay transport may be realized in multiple ways such as (but not limited to):

- A mesh of RSVP-TE [RFC3209] or SR-TE [RFC9256] tunnels created with specific optimization criteria and constraints. For example, mesh "A" might represent tunnels optimized for latency, and mesh "B" might represent tunnels optimized for high capacity.
- A Flex-Algorithm [RFC9350] with a particular metric-type (e.g., latency), or one that only uses links with particular properties (e.g., a Media Access Control Security (MACsec) link [IEEE802.1AE]) or excludes links that are within a particular geography.

These protocols can be controlled, e.g., by tuning the protocol list under the "underlay-transport" data node defined in the L3VPN Network Model (L3NM) [RFC9182] and the L2VPN Network Model (L2NM) [RFC9291].

Also, underlay transports may be realized using separate NRPs. However, such an approach is left out of the scope given the current state of the technology at the time of writing this document.

Similar to the QoS mapping models discussed in [Section 5](#), for mapping to underlay transports at the ingress PE, both the 5QI-unaware and 5QI-aware models are defined. Essentially, entire slices can be mapped to underlay transports without 5G QoS consideration (5QI-unaware model). For example, flows with different 5G QoS Classes, even from same slice, can be mapped to different underlay transports (5QI-aware model).

[Figure 27](#) depicts an example of a simple network with two underlay transports, each using a mesh of TE tunnels with or without Path Computation Element (PCE) [[RFC5440](#)] and with or without per-path bandwidth reservations. [Section 7](#) discusses in detail different bandwidth models that can be deployed in the provider network. However, discussion about how to realize or orchestrate underlay transports is out of scope for this document.

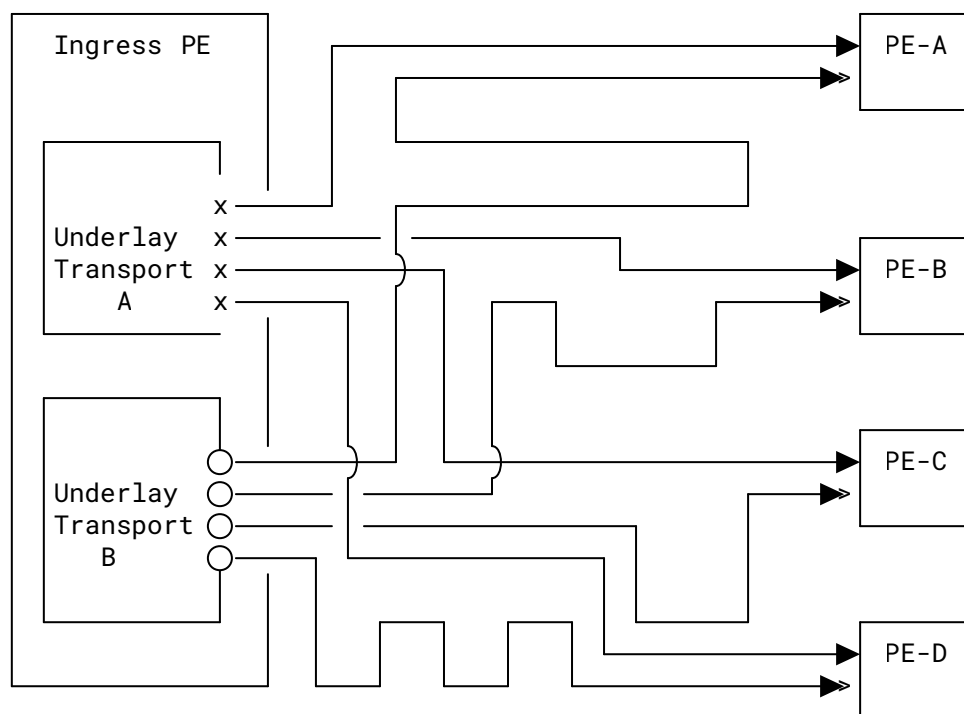


Figure 27: Example of Underlay Transport Relying on TE Tunnels

For illustration purposes, [Figure 27](#) shows only single tunnels per underlay transport for an (ingress PE, egress PE) pair. However, there might be multiple tunnels within a single underlay transport between any pair of PEs.

6.1. 5QI-Unaware Model

As discussed in [Section 5.2.1](#), in the 5QI-unaware model, the provider network doesn't take into account 5G QoS during execution of per-hop behavior. The entire slice is mapped to a single TN QoS Class; therefore, the entire slice is subject to the same per-hop behavior. Similarly, in the 5QI-unaware PE underlay transport mapping model, the entire slice is mapped to a single underlay transport, as depicted in [Figure 28](#).

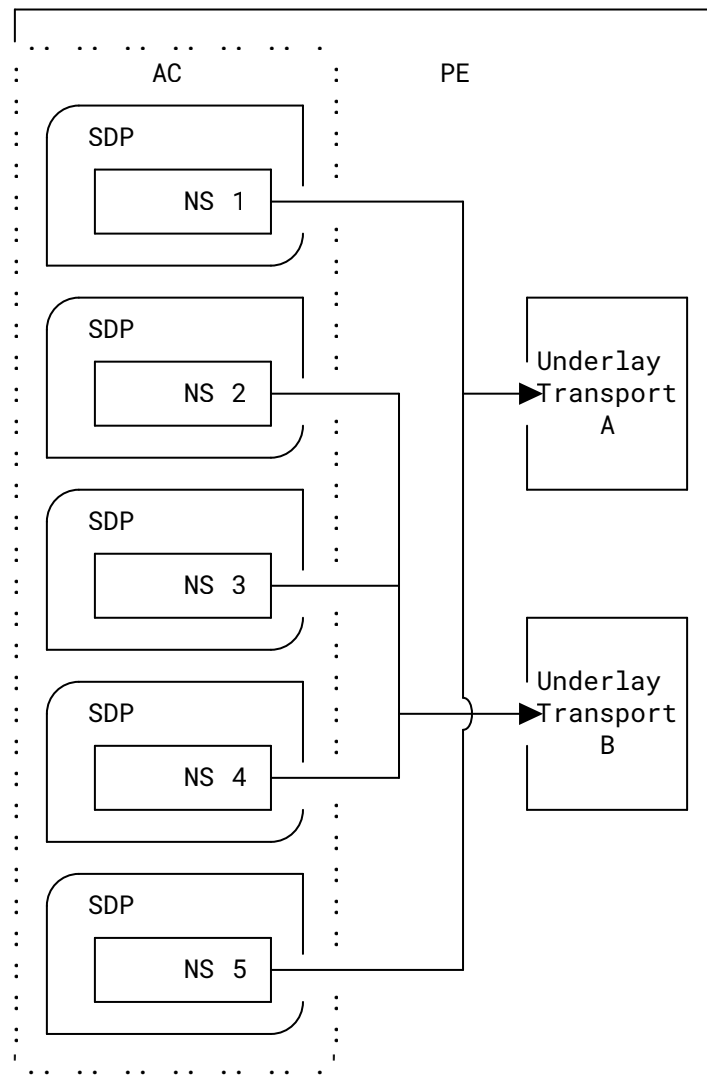


Figure 28: Mapping of Network Slice to Underlay Transport (5QI-Unaware Model)

6.2. 5QI-Aware Model

In the 5QI-aware model, the traffic can be mapped to underlay transports at the granularity of 5G QoS Class. Given that the potential number of underlay transports is limited, packets from multiple 5G QoS Classes with similar characteristics are mapped to a common underlay transport, as depicted in [Figure 29](#).

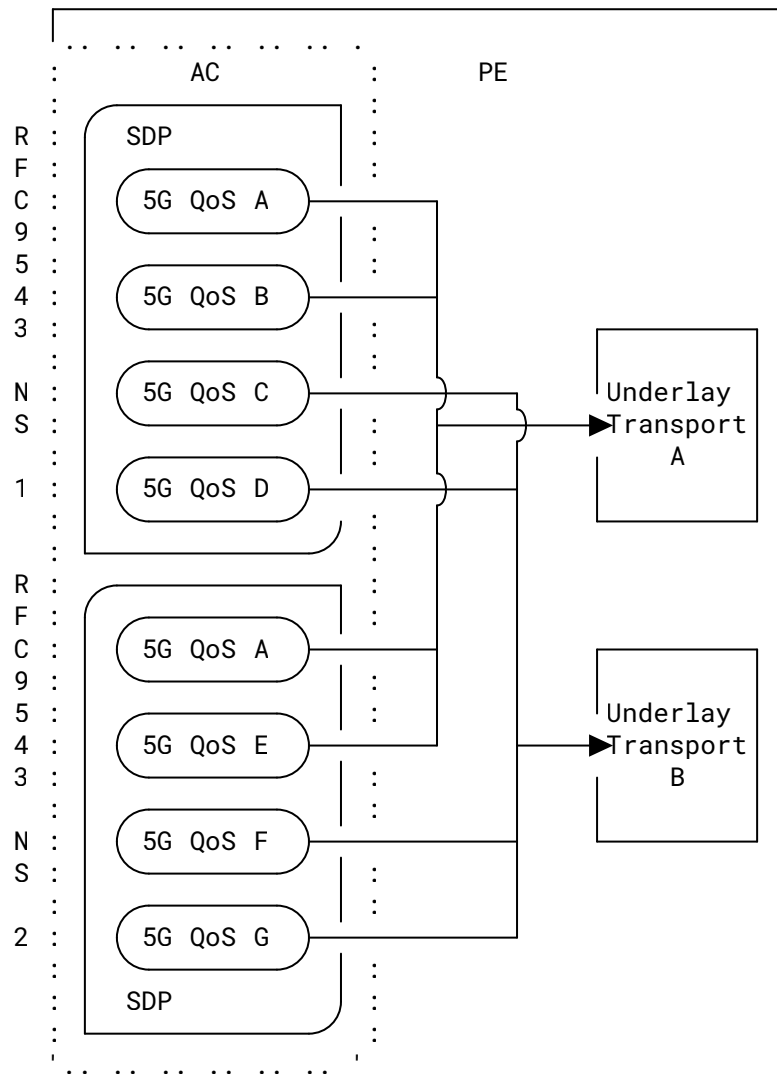


Figure 29: Mapping of Network Slice to Underlay Transport (5QI-Aware Model)

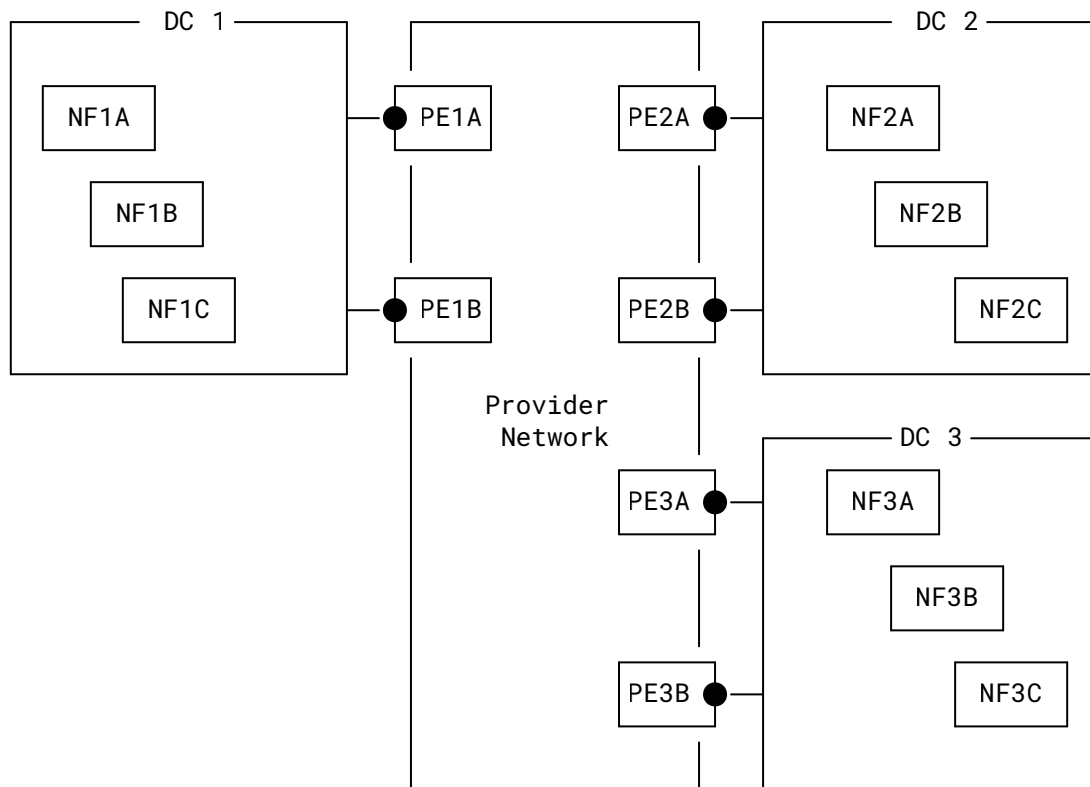
7. Capacity Planning/Management

7.1. Bandwidth Requirements

This section describes the information conveyed by the 5G NSO to the NSC with respect to slice bandwidth requirements.

[Figure 30](#) shows three DCs that contain instances of network functions. Also shown are PEs that have links to the DCs. The PEs belong to the provider network. Other details of the provider network, such as P-routers and transit links, are not shown. In addition, details of the DC infrastructure in customer sites, such as switches and routers, are not shown.

The 5G NSO is aware of the existence of the network functions and their locations. However, it is not aware of the details of the provider network. The NSC has the opposite view -- it is aware of the provider network infrastructure and the links between the PEs and the DCs, but it is not aware of the individual network functions at customer sites.



- SDP, with fine-grained QoS (dedicated resources per RFC 9543 Network Slices)

Figure 30: Example of Multi-DC Architecture

Let us consider 5G Network Slice "X" that uses some of the network functions in the three DCs. If this slice has latency requirements, the 5G NSO will have taken those into account when deciding which NF instances in which DC are to be invoked for this slice. As a result of such a placement decision, the three DCs shown are involved in 5G Network Slice "X", rather than other DCs. For its decision-making, the 5G NSO needs information from the NSC about the observed latency between DCs. Preferably, the NSC would present the topology in an abstracted form, consisting of point-to-point abstracted links between pairs of DCs and associated latency and, optionally, delay variation and link-loss values. It would be valuable to have a mechanism for the 5G NSO to inform the NSC which DC-pairs are of interest for these metrics; there may be thousands of DCs, but the 5G NSO will only be interested in these metrics for a small fraction of all the possible DC-pairs, i.e., those in the same region of the provider network. The mechanism for conveying the information is out of scope for this document.

Table 1 shows the matrix of bandwidth demands for 5G Network Slice "X". Within the slice, multiple NF instances might be sending traffic from DC_i to DC_j. However, the 5G NSO sums the associated demands into one value. For example, "NF1A" and "NF1B" in "DC 1" might be sending traffic to multiple NFs in "DC 2", but this is expressed as one value in the traffic matrix: the total

bandwidth required for 5G Network Slice "X" from "DC 1" to "DC 2" (8 units). Each row in the right-most column in the traffic matrix shows the total amount of traffic going from a given DC into the Transport Network, regardless of the destination DC. Note that this number can be less than the sum of DC-to-DC demands in the same row, on the basis that not all the NFs are likely to be sending at their maximum rate simultaneously. For example, the total traffic from "DC 1" for slice "X" is 11 units, which is less than the sum of the DC-to-DC demands in the same row (13 units). Note, as described in [Section 5](#), a slice may have per-QoS class bandwidth requirements and may have CIR and PIR limits. This is not included in the example, but the same principles apply in such cases.

From/To	DC 1	DC 2	DC 3	Total from DC
DC 1	n/a	8	5	11.0
DC 2	1	n/a	2	2.5
DC 3	4	7	n/a	10.0

Table 1: Inter-DC Traffic Demand Matrix (Slice X)

The YANG data model defined in [\[NSSM\]](#) can be used to convey all of the information in the traffic matrix to an NSC. The NSC applies policers corresponding to the last column in the traffic matrix to the appropriate PE routers, in order to enforce the bandwidth contract. For example, it applies a policer of 11 units to PE1A and PE1B that face DC 1, as this is the total bandwidth that DC 1 sends into the provider network corresponding to slice "X". Also, the controller may apply shapers in the direction from the TN to the DC if there is the possibility of a link in the DC being oversubscribed. Note that a peer NF endpoint of an AC can be identified using "peer-sap-id" as defined in [\[RFC9408\]](#).

Depending on the bandwidth model used in the provider network ([Section 7.2](#)), the other values in the matrix, i.e., the DC-to-DC demands, may not be directly applied to the provider network. Even so, the information may be useful to the NSC for capacity planning and failure simulation purposes. On the other hand, if the DC-to-DC demand information is not used by the NSC, the IETF YANG data models for L3VPN service delivery [\[RFC8299\]](#) or L2VPN service delivery [\[RFC8466\]](#) could be used instead of the YANG data model defined in [\[NSSM\]](#), as they support conveying the bandwidth information in the right-most column of the traffic matrix.

The provider network may be implemented in such a way that it has various types of paths, for example, low-latency traffic might be mapped onto a different transport path from other traffic (for example, a particular Flex-Algorithm, a particular set of TE paths, or a specific queue [\[RFC9330\]](#)), as discussed in [Section 5](#). The 5G NSO can use the YANG data model defined in [\[NSSM\]](#) to request low-latency transport for a given slice if required. However, the YANG data models in [\[RFC8299\]](#) or [\[RFC8466\]](#) do not support requesting a particular transport-type, e.g., low-latency. One option is to augment these models to convey this information. This can be achieved by reusing the "underlay-transport" construct defined in [\[RFC9182\]](#) and [\[RFC9291\]](#).

7.2. Bandwidth Models

This section describes three bandwidth management schemes that could be employed in the provider network. Many variations are possible, but each example describes the salient points of the corresponding scheme. Schemes 2 and 3 use TE; other variations on TE are possible as described in [RFC9522].

7.2.1. Scheme 1: Shortest Path Forwarding (SPF)

Shortest path forwarding is used according to the IGP metric. Given that some slices are likely to have latency SLOs, the IGP metric on each link can be set to be in proportion to the latency of the link. In this way, all traffic follows the minimum latency path between endpoints.

In Scheme 1, although the operator provides bandwidth guarantees to the slice customers, there is no explicit end-to-end underpinning of the bandwidth SLO, in the form of bandwidth reservations across the provider network. Rather, the expected performance is achieved via capacity planning, based on traffic growth trends and anticipated future demands, in order to ensure that network links are not over-subscribed. This scheme is analogous to that used in many existing business VPN deployments, in that bandwidth guarantees are provided to the customers but are not explicitly underpinned end to end across the provider network.

A variation on the scheme is that Flex-Algorithm [RFC9350] is used. For example, one Flex-Algorithm could use latency-based metrics, and another Flex-Algorithm could use the IGP metric. There would be a many-to-one mapping of network slices to Flex-Algorithms.

While Scheme 1 is technically feasible, it is vulnerable to unexpected changes in traffic patterns and/or network element failures resulting in congestion. This is because, unlike Schemes 2 and 3, which employ TE, traffic cannot be diverted from the shortest path.

7.2.2. Scheme 2: TE Paths with Fixed Bandwidth Reservations

Scheme 2 uses RSVP-TE [RFC3209] or SR-TE [RFC9256] paths with fixed bandwidth reservations. By "fixed", we mean a value that stays constant over time, unless the 5G NSO communicates a change in slice bandwidth requirements, due to the creation or modification of a slice. Note that the "reservations" may be maintained by the transport controller; it is not necessary (or indeed possible for current SR-TE technology at the time of writing this document) to reserve bandwidth at the network layer. The bandwidth requirement acts as a constraint whenever the controller (re)computes a path. There could be a single mesh of paths between endpoints that carry all of the traffic types, or there could be a small handful of meshes, for example, one mesh for low-latency traffic that follows the minimum latency path and another mesh for the other traffic that follows the minimum IGP metric path, as described in Section 5. There would be a many-to-one mapping of slices to paths.

The bandwidth requirement from DC_i to DC_j is the sum of the DC_i-DC_j demands of the individual slices. For example, if only slices "X" and "Y" are present, then the bandwidth requirement from "DC 1" to "DC 2" is 12 units (8 units for slice "X" (Table 1) and 4 units for slice "Y" (Table 2)). When the 5G NSO requests a new slice, the NSC, increments the bandwidth requirement according to

the requirements of the new slice. For example, in [Figure 30](#), suppose a new slice is instantiated that needs 0.8 Gbps from "DC 1" to "DC 2". The transport controller would increase its notion of the bandwidth requirement from "DC 1" to "DC 2" from 12 Gbps to 12.8 Gbps to accommodate the additional expected traffic.

From/To	DC 1	DC 2	DC 3	Total from DC
DC 1	n/a	4	2.5	6.0
DC 2	0.5	n/a	0.8	1.0
DC 3	2.6	3	n/a	5.1

Table 2: Inter-DC Traffic Demand Matrix (Slice Y)

In the example, each DC has two PEs facing it for reasons of resilience. The NSC needs to determine how to map the "DC 1" to "DC 2" bandwidth requirement to bandwidth reservations of TE LSPs from "DC 1" to "DC 2". For example, if the routing configuration is arranged such that in the absence of any network failure, traffic from "DC 1" to "DC 2" always enters "PE1A" and goes to "PE2A", the controller reserves 12.8 Gbps of bandwidth on the path from "PE1A" to "PE2A". On the other hand, if the routing configuration is arranged such that in the absence of any network failure, traffic from "DC 1" to "DC 2" always enters "PE1A" and is load-balanced across "PE2A" and "PE2B", the controller reserves 6.4 Gbps of bandwidth on the path from "PE1A" to "PE2A" and 6.4 Gbps of bandwidth on the path from "PE1A" to "PE2B". It might be tricky for the NSC to be aware of all conditions that change the way traffic lands on the various PEs and therefore know that it needs to change bandwidth reservations of paths accordingly. For example, there might be an internal failure within "DC 1" that causes traffic from "DC 1" to land on "PE1B" rather than "PE1A". The NSC may not be aware of the failure and therefore may not know that it now needs to apply bandwidth reservations to paths from "PE1B" to "PE2A" and "PE2B".

7.2.3. Scheme 3: TE Paths without Bandwidth Reservation

Like Scheme 2, Scheme 3 uses RSVP-TE or SR-TE paths. There could be a single mesh of paths between endpoints that carry all of the traffic types, or there could be a small handful of meshes, for example, one mesh for low-latency traffic that follows the minimum latency path and another mesh for the other traffic that follows the minimum IGP metric path, as described in [Section 5](#). There would be a many-to-one mapping of slices to paths.

The difference between Scheme 2 and Scheme 3 is that Scheme 3 does not have fixed bandwidth reservations for the paths. Instead, actual measured data plane traffic volumes are used to influence the placement of TE paths. One way of achieving this is to use distributed RSVP-TE with auto-bandwidth. Alternatively, the NSC can use telemetry-driven automatic congestion avoidance. In this approach, when the actual traffic volume in the data plane on a given link exceeds a threshold, the controller, knowing how much actual data plane traffic is currently traveling along each RSVP or SR-TE path, can tune one or more paths using the link such that they avoid that link. This approach is similar to that described in [Section 4.3.1](#) of [\[RFC9522\]](#).

It would be undesirable to move a path that has latency as its cost function, rather than another type of path, in order to ease the congestion, as the altered path will typically have a higher latency. This can be avoided by designing the algorithms described in the previous paragraph such that they avoid moving minimum-latency paths unless there is no alternative.

8. Network Slicing OAM

The deployment and maintenance of slices within a network imply that a set of OAM functions [RFC6291] need to be deployed by the providers, for example:

- Providers should be able to execute OAM tasks on a per-network-slice basis. These tasks can cover the "full" slice within a domain or a portion of that slice (for troubleshooting purposes, for example).

For example, per-slice OAM tasks can consist of (but not limited to):

- tracing resources that are bound to a given network slice,
- tracing resources that are invoked when forwarding a given flow bound to a given network slice,
- assessing whether flow isolation characteristics are in conformance with the Network Slice Service requirements, or
- assessing the compliance of the allocated network slice resources against flow and customer service requirements.

[RFC7276] provides an overview of available OAM tools. These technology-specific tools can be reused in the context of network slicing. Providers that deploy network slicing capabilities should be able to select whatever OAM technology or specific feature that would address their needs.

- Providers may want to enable differentiated failure detection and repair features for a subset of network slices. For example, a given network slice may require fast detection and repair mechanisms, while others may not be engineered with such means. The provider can use techniques such as those described in [RFC5286], [RFC5714], and [RFC8355].
- Providers may deploy means to dynamically discover the set of network slices that are enabled within its network. Such dynamic discovery capability facilitates the detection of any mismatch between the view maintained by the control/management plane and the actual network configuration. When mismatches are detected, corrective actions should be undertaken accordingly. For example, a provider may rely upon the L3NM [RFC9182] or the L2NM [RFC9291] to maintain the full set of L2VPN/L3VPNs that are used to deliver Network Slice Services. The correlation between an L2VPN/L3VPN instance and a Network Slice Service is maintained using the "parent-service-id" attribute (Section 7.3 of [RFC9182]).
- Providers should provide the means to report a set of network performance metrics to assess whether the agreed Slice Service objectives are honored. These means are used for SLO monitoring and violation detection purposes. For example, the YANG data model in [RFC9375] can be used to report the one-way delay and one-way delay variation of links. Both conventional active/passive measurement methods [RFC7799] and more recent telemetry methods (e.g., YANG Push [RFC8641]) can be used.

- Providers should have the means to report and expose observed performance metrics and other OAM state to customer. For example, the YANG data model defined in [NSSM] exposes a set of statistics per SDP, connectivity construct, and connection group.

9. Scalability Implications

The mapping of 5G Network Slices to Transport Network Slices (see [Section 3.5](#)) is a design choice of service operators that may be a function of, e.g., the number of instantiated slices, requested services, or local engineering capabilities and guidelines. However, operators should carefully consider means to ease slice migration strategies. For example, a provider may initially adopt a 1-to-1 mapping if it has to instantiate just a few network slices and accommodate the need of only a few customers. That provider may decide to move to an M-to-1 mapping for aggregation/scalability purposes if sustained increased slice demand is observed.

Putting in place adequate automation means to realize network slices (including the adjustment of the mapping of Slice Services to network slices) would ease slice migration operations.

The realization model described in this document inherits the scalability properties of the underlying L2VPN and L3VPN technologies ([Section 3.7](#)). Readers may refer, for example, to [Section 13](#) of [RFC4365] or [Section 1.2.5](#) of [RFC6624] for a scalability assessment of some of these technologies. Providers may adjust the mapping model to better handle local scalability constraints.

10. IANA Considerations

This document has no IANA actions.

11. Security Considerations

[Section 10](#) of [RFC9543] discusses generic security considerations that are applicable to network slicing, with a focus on the following considerations:

Conformance to security constraints:

Specific security requests, such as not routing traffic through a particular geographical region can be met by mapping the traffic to an underlay transport ([Section 6](#)) that avoids that region.

NSC authentication:

Per [RFC9543], underlay networks need to be protected against attacks from an adversary NSC as this could destabilize overall network operations. The interaction between an NSC and the underlay network is used to pass service provisioning requests following a set of YANG modules that are designed to be accessed via YANG-based management protocols, such as NETCONF [RFC6241] and RESTCONF [RFC8040]. These YANG-based management protocols have to use (1) a secure transport layer (e.g., SSH [RFC4252], TLS [RFC8446], and QUIC [RFC9000]) and (2) mutual authentication.

The NETCONF access control model [RFC8341] provides the means to restrict access for particular NETCONF or RESTCONF users to a preconfigured subset of all available NETCONF or RESTCONF protocol operations and content.

Readers may refer to documents that describe NSC realization, such as [NSC-MODEL].

Specific isolation criteria:

Adequate admission control policies, for example, policers as described in Section 5.2.1.1, should be configured in the edge of the provider network to control access to specific slice resources. This prevents the possibility of one slice consuming resources at the expense of other slices. Likewise, access to classification and mapping tables have to be controlled to prevent misbehaviors (an unauthorized entity may modify the table to bind traffic to a random slice, redirect the traffic, etc.). Network devices have to check that a required access privilege is provided before granting access to specific data or performing specific actions.

Data Confidentiality and Integrity of an RFC 9543 Network Slice:

As described in Section 5.1.2.1 of [RFC9543], the customer might request a Service Level Expectation (SLE) that mandates encryption.

This can be achieved, e.g., by mapping the traffic to an underlay transport (Section 6) that uses only MACsec-encrypted links.

In order to avoid the need for a mapping table to associate source/destination IP addresses and the specific S-NSSAIs of slices, Section 4.2 describes an approach where some or all S-NSSAI bits are embedded in an IPv6 address using an algorithm approach. An attacker from within the Transport Network who has access to the mapping configuration may infer the slices to which a packet belongs. It may also alter these bits, which may lead to steering the packet via a distinct network slice and thus to service disruption. Note that such an attacker from within the Transport Network may inflict more damage (e.g., randomly drop packets).

Security considerations specific to each of the technologies and protocols listed in the document are discussed in the specification documents of each of these protocols. In particular, readers should refer to the "Security Framework for Provider-Provisioned Virtual Private Networks (PPVPNs)" [RFC4111], the "Applicability Statement for BGP/MPLS IP Virtual Private Networks (VPNs)" (Section 6 of [RFC4365]), and the "Analysis of the Security of BGP/MPLS IP Virtual Private Networks (VPNs)" [RFC4381] for a comprehensive discussion about security considerations related to VPN technologies (including authentication and encryption between PEs, use of IPsec tunnels that terminate within the customer sites to protect user data, prevention of illegitimate traffic from entering a VPN instance, etc.). Also, readers may refer to Section 9 of [RFC9522] for a discussion about security considerations related to TE mechanisms.

12. References

12.1. Normative References

[RFC4364]

Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.

- [RFC7608] Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", BCP 198, RFC 7608, DOI 10.17487/RFC7608, July 2015, <<https://www.rfc-editor.org/info/rfc7608>>.
- [RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration Access Control Model", STD 91, RFC 8341, DOI 10.17487/RFC8341, March 2018, <<https://www.rfc-editor.org/info/rfc8341>>.
- [RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <<https://www.rfc-editor.org/info/rfc9543>>.

12.2. Informative References

- [Book-5G] Peterson, L., Sunay, O., and B. Davie, "Private 5G: A Systems Approach", 2023, <<https://5g.systemsapproach.org/>>.
- [ECPRI] Common Public Radio Interface, "Common Public Radio Interface: eCPRI Interface Specification", <https://www.cpri.info/downloads/eCPRI_v_2.0_2019_05_10c.pdf>.
- [IEEE802.1AE] IEEE, "802.1AE: MAC Security (MACsec)", <<https://1.ieee802.org/security/802-1ae/>>.
- [MAPPING] Contreras, L. M., Ed., Bykov, I., Ed., and K. G. Szarkowicz, Ed., "5QI to DiffServ DSCP Mapping Example for Enforcement of 5G End-to-End Network Slice QoS", Work in Progress, Internet-Draft, draft-cbs-teas-5qi-to-dscp-mapping-05, 20 October 2025, <<https://datatracker.ietf.org/doc/html/draft-cbs-teas-5qi-to-dscp-mapping-05>>.
- [NG.113] GSMA, "NG.113: 5GS Roaming Guidelines", Version 4.0, May 2021, <<https://www.gsma.com/newsroom/wp-content/uploads//NG.113-v4.0.pdf>>.
- [NS-APP] Geng, X., Contreras, L. M., Ed., Rokui, R., Dong, J., and I. Bykov, "IETF Network Slice Application in 3GPP 5G End-to-End Network Slice", Work in Progress, Internet-Draft, draft-ietf-teas-5g-network-slice-application-05, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-5g-network-slice-application-05>>.
- [NS-IP-MPLS] Saad, T., Beeram, V., Dong, J., Halpern, J., and S. Peng, "Realizing Network Slices in IP/MPLS Networks", Work in Progress, Internet-Draft, draft-ietf-teas-ns-ip-mpls-06, 20 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ns-ip-mpls-06>>.

- [NSC-MODEL]** Contreras, L. M., Rokui, R., Tantsura, J., Wu, B., and X. Liu, "IETF Network Slice Controller and its Associated Data Models", Work in Progress, Internet-Draft, draft-ietf-teas-ns-controller-models-06, 20 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ns-controller-models-06>>.
- [NSSM]** Wu, B., Dhody, D., Rokui, R., Saad, T., and J. Mullooly, "A YANG Data Model for the RFC 9543 Network Slice Service", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slice-nbi-yang-25, 9 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slice-nbi-yang-25>>.
- [O-RAN.WG9.XPSAAS]** O-RAN Alliance, "Xhaul Packet Switched Architectures and Solutions", O-RAN.WG9.XPSAAS, Version 09.00, October 2025, <<https://specifications.o-ran.org/specifications>>.
- [RFC1997]** Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <<https://www.rfc-editor.org/info/rfc1997>>.
- [RFC2474]** Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.
- [RFC2475]** Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, DOI 10.17487/RFC2475, December 1998, <<https://www.rfc-editor.org/info/rfc2475>>.
- [RFC2698]** Heinanen, J. and R. Guerin, "A Two Rate Three Color Marker", RFC 2698, DOI 10.17487/RFC2698, September 1999, <<https://www.rfc-editor.org/info/rfc2698>>.
- [RFC3209]** Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC4026]** Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, DOI 10.17487/RFC4026, March 2005, <<https://www.rfc-editor.org/info/rfc4026>>.
- [RFC4111]** Fang, L., Ed., "Security Framework for Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4111, DOI 10.17487/RFC4111, July 2005, <<https://www.rfc-editor.org/info/rfc4111>>.
- [RFC4115]** Aboul-Magd, O. and S. Rabie, "A Differentiated Service Two-Rate, Three-Color Marker with Efficient Handling of in-Profile Traffic", RFC 4115, DOI 10.17487/RFC4115, July 2005, <<https://www.rfc-editor.org/info/rfc4115>>.
- [RFC4176]** El Mghazli, Y., Ed., Nadeau, T., Boucadair, M., Chan, K., and A. Gonguet, "Framework for Layer 3 Virtual Private Networks (L3VPN) Operations and Management", RFC 4176, DOI 10.17487/RFC4176, October 2005, <<https://www.rfc-editor.org/info/rfc4176>>.

-
- [RFC4252] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Authentication Protocol", RFC 4252, DOI 10.17487/RFC4252, January 2006, <<https://www.rfc-editor.org/info/rfc4252>>.
 - [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
 - [RFC4365] Rosen, E., "Applicability Statement for BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4365, DOI 10.17487/RFC4365, February 2006, <<https://www.rfc-editor.org/info/rfc4365>>.
 - [RFC4381] Behringer, M., "Analysis of the Security of BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4381, DOI 10.17487/RFC4381, February 2006, <<https://www.rfc-editor.org/info/rfc4381>>.
 - [RFC4664] Andersson, L., Ed. and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, DOI 10.17487/RFC4664, September 2006, <<https://www.rfc-editor.org/info/rfc4664>>.
 - [RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, DOI 10.17487/RFC4761, January 2007, <<https://www.rfc-editor.org/info/rfc4761>>.
 - [RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, DOI 10.17487/RFC4762, January 2007, <<https://www.rfc-editor.org/info/rfc4762>>.
 - [RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, DOI 10.17487/RFC5286, September 2008, <<https://www.rfc-editor.org/info/rfc5286>>.
 - [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.
 - [RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, DOI 10.17487/RFC5714, January 2010, <<https://www.rfc-editor.org/info/rfc5714>>.
 - [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, DOI 10.17487/RFC5952, August 2010, <<https://www.rfc-editor.org/info/rfc5952>>.
 - [RFC6136] Sajassi, A., Ed. and D. Mohan, Ed., "Layer 2 Virtual Private Network (L2VPN) Operations, Administration, and Maintenance (OAM) Requirements and Framework", RFC 6136, DOI 10.17487/RFC6136, March 2011, <<https://www.rfc-editor.org/info/rfc6136>>.

-
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <<https://www.rfc-editor.org/info/rfc6291>>.
- [RFC6624] Kompella, K., Kothari, B., and R. Cherukuri, "Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling", RFC 6624, DOI 10.17487/RFC6624, May 2012, <<https://www.rfc-editor.org/info/rfc6624>>.
- [RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, DOI 10.17487/RFC7276, June 2014, <<https://www.rfc-editor.org/info/rfc7276>>.
- [RFC7422] Donley, C., Grundemann, C., Sarawat, V., Sundaresan, K., and O. Vautrin, "Deterministic Address Mapping to Reduce Logging in Carrier-Grade NAT Deployments", RFC 7422, DOI 10.17487/RFC7422, December 2014, <<https://www.rfc-editor.org/info/rfc7422>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC7510] Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black, "Encapsulating MPLS in UDP", RFC 7510, DOI 10.17487/RFC7510, April 2015, <<https://www.rfc-editor.org/info/rfc7510>>.
- [RFC7623] Sajassi, A., Ed., Salam, S., Bitar, N., Isaac, A., and W. Henderickx, "Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)", RFC 7623, DOI 10.17487/RFC7623, September 2015, <<https://www.rfc-editor.org/info/rfc7623>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [RFC7806] Baker, F. and R. Pan, "On Queuing, Marking, and Dropping", RFC 7806, DOI 10.17487/RFC7806, April 2016, <<https://www.rfc-editor.org/info/rfc7806>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [RFC8100] Geib, R., Ed. and D. Black, "Diffserv-Interconnection Classes and Practice", RFC 8100, DOI 10.17487/RFC8100, March 2017, <<https://www.rfc-editor.org/info/rfc8100>>.
- [RFC8214] Boutros, S., Sajassi, A., Salam, S., Drake, J., and J. Rabadan, "Virtual Private Wire Service Support in Ethernet VPN", RFC 8214, DOI 10.17487/RFC8214, August 2017, <<https://www.rfc-editor.org/info/rfc8214>>.

-
- [RFC8299] Wu, Q., Ed., Litkowski, S., Tomotaki, L., and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", RFC 8299, DOI 10.17487/RFC8299, January 2018, <<https://www.rfc-editor.org/info/rfc8299>>.
- [RFC8355] Filsfils, C., Ed., Previdi, S., Ed., Decraene, B., and R. Shakir, "Resiliency Use Cases in Source Packet Routing in Networking (SPRING) Networks", RFC 8355, DOI 10.17487/RFC8355, March 2018, <<https://www.rfc-editor.org/info/rfc8355>>.
- [RFC8365] Sajassi, A., Ed., Drake, J., Ed., Bitar, N., Shekhar, R., Uttaro, J., and W. Henderickx, "A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)", RFC 8365, DOI 10.17487/RFC8365, March 2018, <<https://www.rfc-editor.org/info/rfc8365>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8466] Wen, B., Fioccola, G., Ed., Xie, C., and L. Jalil, "A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery", RFC 8466, DOI 10.17487/RFC8466, October 2018, <<https://www.rfc-editor.org/info/rfc8466>>.
- [RFC8641] Clemm, A. and E. Voit, "Subscription to YANG Notifications for Datastore Updates", RFC 8641, DOI 10.17487/RFC8641, September 2019, <<https://www.rfc-editor.org/info/rfc8641>>.
- [RFC8969] Wu, Q., Ed., Boucadair, M., Ed., Lopez, D., Xie, C., and L. Geng, "A Framework for Automating Service and Network Management with YANG", RFC 8969, DOI 10.17487/RFC8969, January 2021, <<https://www.rfc-editor.org/info/rfc8969>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/info/rfc9000>>.
- [RFC9099] Vyncke, É., Chittimaneni, K., Kaeo, M., and E. Rey, "Operational Security Considerations for IPv6 Networks", RFC 9099, DOI 10.17487/RFC9099, August 2021, <<https://www.rfc-editor.org/info/rfc9099>>.
- [RFC9182] Barguil, S., Gonzalez de Dios, O., Ed., Boucadair, M., Ed., Munoz, L., and A. Aguado, "A YANG Network Data Model for Layer 3 VPNs", RFC 9182, DOI 10.17487/RFC9182, February 2022, <<https://www.rfc-editor.org/info/rfc9182>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

- [RFC9291] Boucadair, M., Ed., Gonzalez de Dios, O., Ed., Barguil, S., and L. Munoz, "A YANG Network Data Model for Layer 2 VPNs", RFC 9291, DOI 10.17487/RFC9291, September 2022, <<https://www.rfc-editor.org/info/rfc9291>>.
- [RFC9330] Briscoe, B., Ed., De Schepper, K., Bagnulo, M., and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture", RFC 9330, DOI 10.17487/RFC9330, January 2023, <<https://www.rfc-editor.org/info/rfc9330>>.
- [RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <<https://www.rfc-editor.org/info/rfc9350>>.
- [RFC9375] Wu, B., Ed., Wu, Q., Ed., Boucadair, M., Ed., Gonzalez de Dios, O., and B. Wen, "A YANG Data Model for Network and VPN Service Performance Monitoring", RFC 9375, DOI 10.17487/RFC9375, April 2023, <<https://www.rfc-editor.org/info/rfc9375>>.
- [RFC9408] Boucadair, M., Ed., Gonzalez de Dios, O., Barguil, S., Wu, Q., and V. Lopez, "A YANG Network Data Model for Service Attachment Points (SAPs)", RFC 9408, DOI 10.17487/RFC9408, June 2023, <<https://www.rfc-editor.org/info/rfc9408>>.
- [RFC9522] Farrel, A., Ed., "Overview and Principles of Internet Traffic Engineering", RFC 9522, DOI 10.17487/RFC9522, January 2024, <<https://www.rfc-editor.org/info/rfc9522>>.
- [RFC9834] Boucadair, M., Ed., Roberts, R., Ed., Gonzalez de Dios, O., Barguil, S., and B. Wu, "YANG Data Models for Bearers and Attachment Circuits as a Service (ACaaS)", RFC 9834, DOI 10.17487/RFC9834, September 2025, <<https://www.rfc-editor.org/info/rfc9834>>.
- [RFC9835] Boucadair, M., Ed., Roberts, R., Gonzalez de Dios, O., Barguil, S., and B. Wu, "A Network YANG Data Model for Attachment Circuits", RFC 9835, DOI 10.17487/RFC9835, September 2025, <<https://www.rfc-editor.org/info/rfc9835>>.
- [TS-23.501] 3GPP, "System architecture for the 5G System (5GS)", Version 19.5.0, Release 19, 3GPP TS 23.501, September 2025, <https://www.3gpp.org/ftp/Specs/archive/23_series/23.501/23501-j50.zip>.
- [TS-28.530] 3GPP, "Management and orchestration; Concepts, use cases and requirements", Version 19.0.0, Release 19, 3GPP TS 28.530, March 2025, <https://www.3gpp.org/ftp/Specs/archive/28_series/28.530/28530-j00.zip>.

Appendix A. Example of Local IPv6 Addressing Plan for Network Functions

Different IPv6 address allocation schemes following the approach in [Section 4.2](#) may be used, with one example allocation shown in [Figure 31](#).




Figure 31: Example of S-NSSAI Embedded into an IPv6 Address

In reference to [Figure 31](#), the most significant 96 bits of the IPv6 address are unique to the NF but do not carry any slice-specific information. The S-NSSAI information is embedded in the least significant 32 bits. The 96-bit part of the address may be structured by the provider, for example, on the geographical location or the DC identification. Refer to [Section 2.1](#) of [[RFC9099](#)] for a discussion on the benefits of structuring an address plan around both services and geographic locations for more structured security policies in a network.

[Figure 32](#) uses the example from [Figure 31](#) to demonstrate a slicing deployment, where the entire S-NSSAI is embedded into IPv6 addresses used by NFs. Let us consider that "NF-A" has a set of tunnel termination points with unique per-slice IP addresses allocated from 2001:db8:a::/96, while "NF-B" uses a set of tunnel termination points with per-slice IP addresses allocated from 2001:db8:b::/96. This example shows two slices: "customer A eMBB" (SST=1, SD=00001) and "customer B MIoT" (SST=3, SD=00003). For "customer A eMBB" slice, the tunnel IP addresses are auto-derived as the IP addresses {2001:db8:a::100:1, 2001:db8:b::100:1}, where {0100:0001} is used as the last two octets. "customer B MIoT" slice (SST=3, SD=00003) tunnel uses the IP addresses {2001:db8:a::300:3, 2001:db8:b::300:3} and simply adds {0300:0003} as the last two octets. Leading zeros are not represented in the resulting IPv6 addresses as per [[RFC5952](#)].

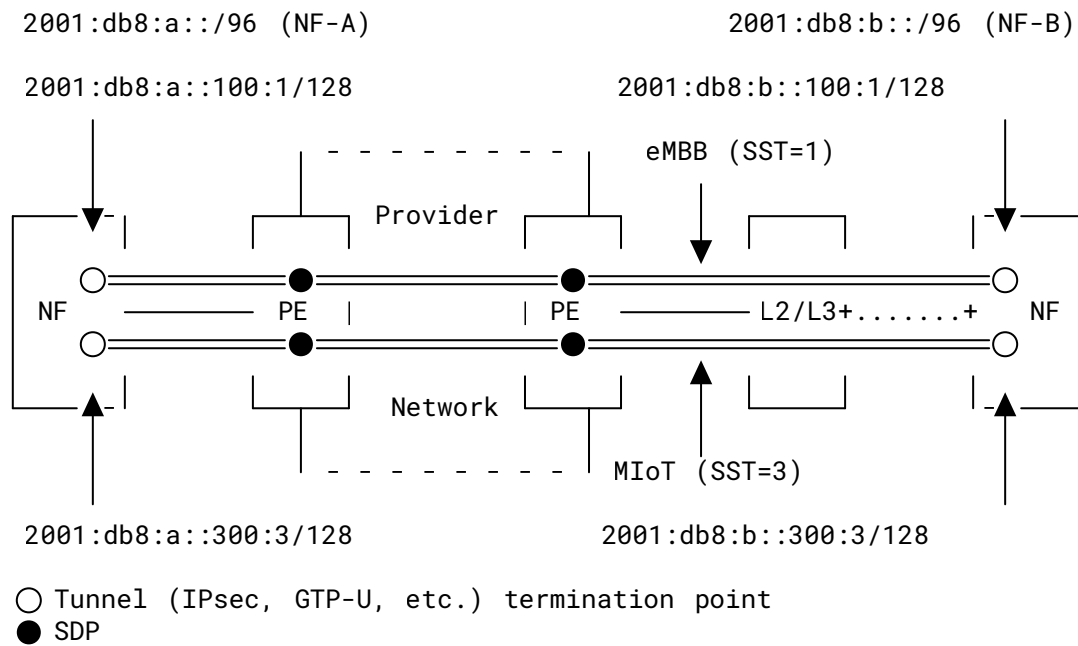


Figure 32: Deployment Example with S-NSSAI Embedded into IPv6 Addresses

Acknowledgments

The authors would like to thank Adrian Farrel, Joel Halpern, Tarek Saad, Greg Mirsky, Rüdiger Geib, Nicklous D. Morris, Daniele Ceccarelli, Bo Wu, Xuesong Geng, and Deborah Brungard for their review of this document and for providing valuable comments.

Special thanks to Jie Dong and Adrian Farrel for the detailed and careful reviews.

Thanks to Alvaro Retana and Mike McBride for the rtg-dir reviews, Yoshifumi Nishida for the tsv-art review, Timothy Winters for the int-dir review, Lars Eggert for the genart review, Joseph Salowey for the secdir review, and Tim Wicinski for the opsdire review.

Thanks to Jim Guichard for the AD review.

Thanks to Erik Kline, Ketan Talaulikar, and Deb Cooley for the IESG review.

Contributors

John Drake
 Sunnyvale, CA
 United States of America
 Email: je_drake@yahoo.com

Ivan Bykov

Ribbon Communications
Tel Aviv
Israel
Email: ivan.bykov@rbbn.com

Reza Rokui

Ciena
Ottawa
Canada
Email: rrokui@ciena.com

Luay Jalil

Verizon
Dallas, TX
United States of America
Email: luay.jalil@verizon.com

Beny Dwi Setyawan

XL Axiata
Jakarta
Indonesia
Email: benyds@xl.co.id

Amit Dhamija

Rakuten
Bangalore
India
Email: amitd@arccus.com

Mojdeh Amani

British Telecom
London
United Kingdom
Email: mojdeh.amani@bt.com

Authors' Addresses**Krzysztof G. Szarkowicz (EDITOR)**

HPE
Wien
Austria
Email: kszarkowicz@juniper.net

Richard Roberts (EDITOR)

Nokia
Rennes
France
Email: richard.roberts@nokia.com

Julian Lucek

HPE
London
United Kingdom
Email: jlucek@juniper.net

Mohamed Boucadair (EDITOR)

Orange
France
Email: mohamed.boucadair@orange.com

Luis M. Contreras

Telefonica
Ronda de la Comunicacion, s/n
Madrid
Spain
Email: luismiguel.contrerasmurillo@telefonica.com
URI: <https://lmcontreras.com/>