Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 8 January 2026

J. Dong
Huawei Technologies
L.M. Contreras
Telefonica

7 July 2025

IETF Network Slice Service Benchmarking
draft-dong-bmwg-network-slicing-00

Abstract

Network slicing aims to provide assurance of specific network performance objectives for network services which require both connectivity and specific performance commitment. Such network services are considered as network slice services. This document provides a benchmarking methodology for network slicing, focusing on evaluating the key functionalities of network slicing mechanisms and the performance of network slice services. The network slicing functionalities includes the data plane, control plane and management plane mechanisms for realizing network slice service, and the performance of network slice service includes the service level agreement (SLA) commitments (bandwidth, delay, and jitter), path constraints and resource guarantee.

The tests aim to demonstrate how network slicing can support competing services in a shared network, ensuring that critical network services in one network slice remain unaffected by congestion or unexpected behavior of other traffic in the same network.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

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

Copyright Notice

Table of Contents

1. Introduction

Network slicing aims to provide assurance of specific network performance objectives for network services which require both connectivity and specific performance commitment. Such network services are considered as network slice services. [RFC9543] describes the general framework for requesting and operating network slices built from IETF technologies. [RFC9543] also introduces the concept of Network Resource Partition (NRP), which is a subset of the buffer/queuing/scheduling resources and associated policies on each of a connected set of links in the underlay network. NRP can be created to support specific SLAs of one or a group of network slice services. In a default deployment, the network resources are not partitioned, which means all the network service are provisioned over the underlay network with shared network resources. Depending on the types of services carried in the network, this may or may not meet all the service requirements.

[RFC9732] describes a framework for Enhanced Virtual Private Networks (VPNs) based on Network Resource Partitions (NRPs) to support the needs of applications with specific traffic performance requirements (e.g., low latency, bounded jitter). NRP-based enhanced VPN can be used to deliver network slice services.

As outlined in [RFC9543], network slicing is crucial for 5G services and beyond, where diverse service requirements demand tailored performance guarantees. Benchmarking of network slicing is essential to understand its effect on service assurance, particularly in high-demand service environments. Evaluating performance aspects such as resource guarantee and SLA adherence helps service providers to understand different types of network slicing mechanisms and choose the suitable one for their network scenarios and demands.

This document provides benchmarking guidelines to evaluate the key functionalities and effectiveness of network slicing, referencing [RFC2544] for benchmarking principles. More specifically, it focuses on the benchmarking of 1) Network Resource partitioning for network slices; 2) Network Slice Topology Control; 3) SR Path with Resource Guarantee 4) Service SLA assurance, including bandwidth, latency etc..

2. Terminology

Network Slice: As defined in [RFC9543], An IETF Network Slice enables connectivity between a set of SDPs with specific Service Level Objectives (SLOs) and Service Level Expectations (SLEs) over a common underlay network.

Network Resource Partition (NRP): As defined in [RFC9543], an NRP is a subset of the buffer/queuing/scheduling resources and associated policies on each of a connected set of links in the underlay network.

End-to-End SLA: A guaranteed level of service performance across an entire network path.

3. Benchmarking Methodology

3.1. Test Setup

Benchmarking tests will be conducted in a controlled environment, consisting of:

Hardware: Network devices capable of supporting network slicing.

Software: Network slice control and management tools.

Traffic Generator: Configured to simulate background and high-value service traffic scenarios.

3.2. Traffic Profiles

Background Traffic: Continuous traffic aimed at saturating network links to simulate congestion.

High-Value Traffic: Critical service traffic which requires guaranteed bandwidth, low latency etc., such as voice or video streaming.

3.3. Metrics

Bandwidth of network slice service: The amount of bandwidth consumed by a network slice service.

Delay (Latency) of network slice service: Time taken for a data packet to traverse the network via a network slice.

Jitter of network slice service: Variation of packet arrival time in a network slice, which can affect time-sensitive applications.

4. Test Cases

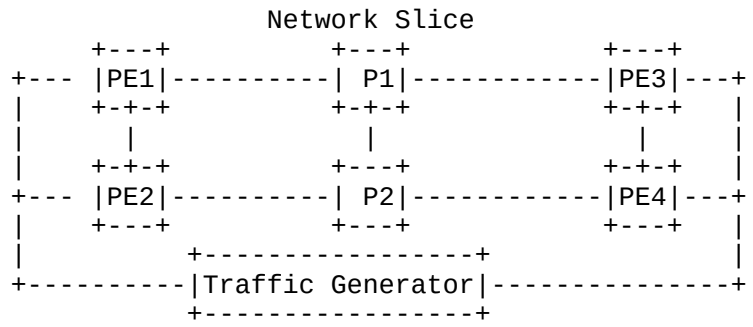

Test Topology

4.1. Network Slice Resource Partition

Objective: Verify that whether the device supports the partitioning of link resources to form an NRP.

Procedure:

1. Create an NRP instance, specify the set of links which belong to the NRP.

2. On each of the link of the NRP, configure the set of bandwidth to be reserved for the NRP.

3. Generate background traffic in the network.

4. Generate test traffic with NRP Selector ID [I-D.ietf-6man-enhanced-vpn-vtn-id] carried in data packet, test the NRP throughput on each link .

Expected Results:

1. The NRP instance is created successfully.

2. On each of the links in the NRP, the bandwidth resource is reserved successfully.

3. The throughput of test traffic is the same as the reserved bandwidth of the NRP.

4.2. Network Slice Topology Control

Objective: Verify the mechanism to control the topology of an NRP, so that network slice traffic will only be forwarded along the paths within the NRP topology.

Procedure:

1. Create an NRP instance, reserve link bandwidth on each link of the NRP.

2. Define a logical topology which aligns with the NRP topology using either Flexible-Algorithm [RFC9350] or Multi-Topology [RFC5120] with Segment Routing.

3. Generate background traffic in the network.

4. Generate test traffic with Algorithm/Topology -specific SR SID and NRP Selector ID carried in the packet, check the packet forwarding path and throughput.

Expected Results:

1. The NRP instance is created successfully, and the link resources for the NRP are reserved successfully.

2. The Flex-Algo or Multi-topology is created successfully, and aligns with the topology of NRP.

3. The test traffic is sent along the paths within the specified Flex-Algorithm or MT, no traffic goes to links out side the topology of NRP.

4. The throughput of test traffic is the same as the reserved bandwidth on the set of links of the NRP.

4.3. SR Policy with Resource Guarantee

Objective: Verify that an SR Policy can be associated with an NRP to provide guaranteed performance for network slice service

Procedure:

1. Provision an SR Policy [RFC9256] with one or multiple candidate paths, each consists of one or multiple segment lists.

2. Create an NRP instance, reserve link bandwidth on each link of the NRP. The topology of the NRP covers all the candidate paths and segment lists of the SR Policy.

3. Associate the SR Policy with the NRP for resource guarantee.

4. Generate background traffic in the network.

5. Generate test traffic with the SID list and NRP Selector ID carried in the packet, check the packet forwarding path and throughput.

Expected Results:

1. The NRP instance is created successfully, and the link resources for the NRP are reserved successfully.

2. The SR Policy is created successfully, and is associated with the NRP.

3. The test traffic is sent along the paths specified by the SID list of the SR Policy.

4. The throughput of test traffic is the same as the reserved bandwidth on the set of links of the NRP.

4.4. Network Slice Service with Bandwidth Guarantee

Objective: Verify that the bandwidth required by a network slice service can be guaranteed.

Procedure:

1. Provision an L3VPN or EVPN service with the requirement on high bandwidth.

2. Provision an SR Policy with one or multiple candidate paths, each consists of one or multiple SID lists according to the connectivity and bandwidth requirement of the L3VPN or EVPN service.

3. Create an NRP instance, reserve link bandwidth on each link of the NRP to meet the bandwidth requirement of the network slice service. The topology of the NRP covers all the candidate paths of the SR Policy.

4. Associate the SR Policy with the NRP for bandwidth guarantee.

5. Generate background traffic in the network, make the network congested.

6. Generate test traffic of the VPN service, steer the VPN service into the SR Policy. Check the packet forwarding path and throughput.

Expected Results:

1. The VPN service is provisioned successfully.

2. The NRP instance is created successfully, and the link resources for the NRP are reserved successfully.

3. The SR Policy is created successfully, and is associated with the NRP.

4. The VPN service test traffic is sent along the paths specified by the SID list of the SR Policy, and the packet is encapsulated with the NRP Selector ID.

5. The throughput of test VPN traffic is the same as the required bandwidth of the service.

4.5. Network Slice Service with Latency Guarantee

Objective: Verify that the latency required by a network slice service can be guaranteed.

Procedure:

1. Provision an L3VPN or EVPN service with specific requirement on bandwidth and latency.

2. Provision an SR Policy with one or multiple candidate paths, each consists of one or multiple SID lists which can meet the latency requirement of the L3VPN or EVPN service.

3. Create an NRP instance, reserve link bandwidth on each link of the NRP to meet the SLA requirement of the network slice service. The topology of the NRP covers all the candidate paths of the SR Policy.

4. Associate the SR Policy with the NRP for latency guarantee.

5. Generate background traffic in the network, make the network congested.

6. Generate test traffic of the VPN service, steer the VPN service into the SR Policy. Check the packet forwarding path, throughput and latency.

Expected Results:

1. The VPN service is provisioned successfully.

2. The NRP instance is created successfully, and the link resources for the NRP are reserved successfully.

3. The SR Policy is created successfully, and is associated with the NRP.

4. The VPN service test traffic is sent along the paths specified by the SID list of the SR Policy, and the packet is encapsulated with the NRP Selector ID.

5. The throughput of test VPN traffic is the same as the required bandwidth of the service, the latency of the traffic meet the required latency of the service.

5. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

6. Security Considerations

TBD

7. Acknowledgements

The authors would like to thank Xuesong Geng for the review and valuable suggestions.

8. References

8.1. Normative References

[I-D.ietf-6man-enhanced-vpn-vtn-id]
Dong, J., Li, Z., Xie, C., Ma, C., and G. S. Mishra, "Carrying Network Resource (NR) related Information in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-10, 2 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-6man-enhanced-vpn-vtn-id-10>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, DOI 10.17487/RFC2544, March 1999, <https://www.rfc-editor.org/info/rfc2544>.

[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>.

[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>.

[RFC9732] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for NRP-Based Enhanced Virtual Private Networks", RFC 9732, DOI 10.17487/RFC9732, March 2025, <https://www.rfc-editor.org/info/rfc9732>.

8.2. Informative References

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <https://www.rfc-editor.org/info/rfc5120>.

[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>.

Authors' Addresses

Jie Dong
Huawei Technologies
China
Email: jie.dong@huawei.com

Luis M. Contreras
Telefonica
Spain
Email: luismiguel.contrerasmurillo@telefonica.com