Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 2 August 2026

H. Zhu
Huazhong University of Science and Technology

29 January 2026

Network Service Type-Aware Traffic Labeling Protocol (NST-TLP)
draft-xsaopig-nsttlp-traffic-labeling-00

Abstract

This document specifies a protocol mechanism for embedding service type identifiers into network packets in order to enable intelligent traffic recognition, policy-based forwarding, and resource optimization by network devices. The protocol allows standardized service type labels to be carried in IPv4/IPv6 headers, MPLS labels, or Ethernet frame headers. It is applicable to a wide range of services, including immersive VR (e.g., 1080p, 4K), scientific computing, real-time communications, and Internet of Things (IoT) applications.

Status of This Memo

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

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

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

This Internet-Draft will expire on 2 August 2026.

Copyright Notice

Copyright (c) 2026 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. 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). They may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress".

2. Introduction

With the diversification of network service types, traffic flows exhibit significantly different requirements in terms of latency, bandwidth, jitter, and packet loss. Traditional network devices have limited visibility into application-layer semantics, resulting in coarse-grained and inefficient resource allocation.

This document proposes a lightweight and extensible traffic labeling protocol that enables network devices to apply differentiated forwarding and resource management policies based on service type, without relying on deep packet inspection (DPI) or maintaining complex per-flow state.

3. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] and [RFC8174].

4. Scope

This specification defines a standardized service type labeling mechanism to enhance network awareness of application requirements and to enable more precise and automated policy enforcement.

The protocol applies primarily to IP-based networks (IPv4 and IPv6), and can be extended to other data plane protocols that support label insertion, such as MPLS and selected Ethernet frame formats.

Labeling points may include end hosts, edge gateways, or SDN controllers. Label-aware nodes include routers, switches, firewalls, and virtualized network functions capable of interpreting NST-TLP labels and enforcing differentiated forwarding, QoS, traffic engineering, or resource scheduling policies.

The target services are those with significantly differentiated network requirements, such as immersive AR/VR, high-performance scientific computing, industrial control signaling, high-priority IoT telemetry, and interactive real-time communications.

This document does not fully specify the complete service type taxonomy. Service type registries are maintained through IANA as described in Section 8. Deployment of NST-TLP is incremental and does not require universal support within a network domain.

5. Terms and Definitions

Service Type: A standardized code identifying an application traffic category (e.g., "VR-1080p", "Scientific-Computing").

Traffic Label: A field embedded in a packet header that carries the service type identifier.

Labeling Point: A network entity responsible for inserting labels into packets.

Label-Aware Node: A network device capable of recognizing labels and applying policies accordingly.

6. Abbreviations

QoS: Quality of Service

7. Protocol Overview

NST-TLP provides a standardized mechanism to translate application semantics into compact, machine-readable labels that are carried in-band within packet headers. This enables network devices to perform service-aware forwarding without requiring DPI or complex flow classification.

7.1. Design Principles

In-band Signaling: Labels are carried within packets to ensure synchronization between traffic and policy.

Incremental Deployment: Networks may contain a mix of label-aware and unaware nodes. Unaware nodes forward packets according to existing protocol rules.

Separation of Semantics and Policy: Labels convey service semantics only. Forwarding actions are determined by local or controller-based policy tables.

8. Protocol Overview

The Network Service Type-Aware Traffic Labeling Protocol (NST-TLP) is a lightweight and extensible mechanism that provides standardized service type identification for network packets. Its fundamental objective is to translate application- or service-level semantics (e.g., service category and quality requirements) into a compact, machine-readable label carried in-band within packet headers.

By embedding service type information directly into packet headers, network devices along the forwarding path can recognize the service attributes of traffic flows and apply appropriate forwarding, scheduling, or resource management policies without relying on deep packet inspection (DPI) or maintaining complex per-flow state.

8.1. Design Principles

The design of NST-TLP follows these principles:

In-band Signaling: The service type label is carried as part of the packet itself, ensuring strict synchronization between traffic and policy enforcement. This avoids the latency and consistency issues associated with out-of-band signaling mechanisms.

Incremental Deployment: NST-TLP is designed to support coexistence of label-aware and label-unaware nodes within the same network. Nodes that do not support NST-TLP MUST ignore the label and forward packets according to the standard processing rules of the underlying protocol (e.g., IPv4 options or IPv6 extension header handling).

Separation of Semantics and Policy: The NST-TLP label conveys only standardized service type semantics and does not prescribe specific forwarding actions. Concrete actions (e.g., queue selection, path computation, or rate limiting) are determined by local policy or by controller-provided policy mapping tables. This separation provides operational flexibility for network administrators.

8.2. System Components and Workflow

A typical NST-TLP system involves the following logical components and their interactions:

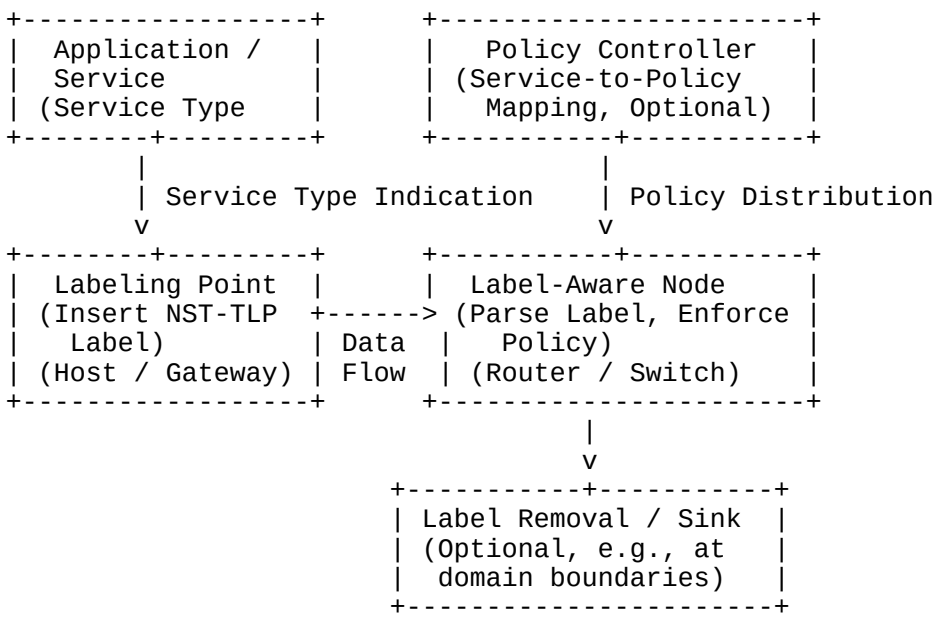

Figure 1

Service Type Identification and Label Generation: Service type information MAY be determined by the traffic source (e.g., an application or virtual machine network interface) or by an intelligent network ingress device (e.g., an SDN edge switch or service gateway) using traffic classification or control-plane instructions. Once identified, an NST-TLP label is generated according to the encoding rules defined in this specification.

Label Encapsulation: The labeling point inserts the generated NST-TLP label into outgoing packets. The specific encapsulation method depends on the underlying network layer protocol (e.g., IPv4, IPv6, MPLS, or Ethernet extensions).

Label-Based Forwarding and Policy Enforcement: NST-TLP label-aware nodes along the forwarding path parse the label carried in packets. These nodes maintain, or retrieve from a controller, a service-type-to-action mapping table. Based on the Service Class, Sub-Class, and Priority fields, the node applies corresponding actions, such as: assigning packets to specific priority queues, selecting low-latency or high-bandwidth paths, performing traffic metering or shaping, or triggering monitoring and logging events.

Label Lifetime: The lifetime of a label is typically bound to the lifetime of the associated traffic flow. A label MAY be removed at administrative domain boundaries (e.g., for security or policy reasons) or rewritten when service type mapping is required across domains. At the final destination, the label is normally ignored or stripped by the protocol stack.

9. Label Format

The NST-TLP label is encoded as a fixed-length or variable-length field carried within packet headers. A recommended basic format is shown below.

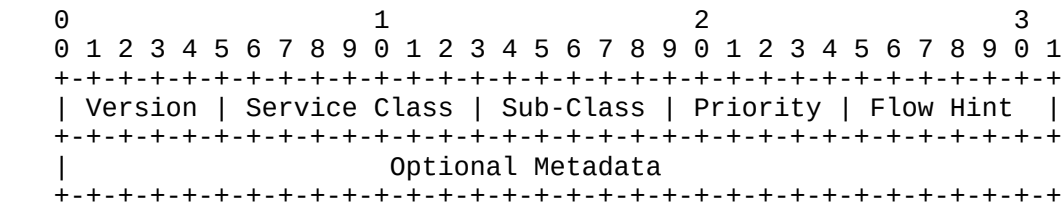

Figure 2

The fields are defined as follows:

Version (4 bits): Indicates the version of the NST-TLP format. This document defines version value 1.

Service Class (8 bits): Identifies the major service category, such as VR/AR, scientific computing, IoT, or real-time communication. Values are assigned from the IANA NST-TLP Service Class registry.

Sub-Class (8 bits): Identifies a sub-category within a given Service Class, such as 1080p, 4K, or 8K for VR video services. The interpretation of this field is specific to the associated Service Class and is defined in the corresponding sub-registry.

Priority (4 bits): Indicates the relative priority of the traffic within the same Service Class or Sub-Class. This field is intended to support QoS scheduling and differentiated forwarding behavior.

Flow Hint (8 bits): A sender-assigned hint value that assists network devices in flow identification or load balancing. For example, this field MAY be used to distinguish multiple parallel flows generated by the same application instance or to carry a simplified hash value for Equal-Cost Multi-Path (ECMP) forwarding. The processing of this field is local to a network domain and its semantics are not guaranteed to be consistent across domains.

Optional Metadata (variable length): Carries additional information related to the service type. The structure and semantics of this field are defined by the specification of the corresponding Service Class. For example, for scientific computing services, the metadata MAY include a compact Job Identifier or communication pattern indicator, while for financial trading services it MAY include a microsecond-level timestamp.

10. Label Encoding and Registry

To ensure global consistency and interoperability of NST-TLP labels, a public and authoritative encoding registry is required. This document specifies that IANA maintains the registries for NST-TLP Service Classes and Sub-Classes.

Once a value is assigned and published, its semantic meaning MUST remain stable across future versions of the protocol and MUST NOT be redefined. All assignments MUST be supported by stable and publicly available technical documentation describing the characteristics of the service type, its typical network requirements, and the format of optional metadata (if applicable).

10.1. NST-TLP Service Class Registry

The NST-TLP Service Class registry consists of 6-bit values (0-63). Each registry entry defines a major service category.

Each registry entry includes the following fields:

Value: A 6-bit numeric value (0-63).

Name: An English abbreviation and full name of the service category (e.g., "VRAR Virtual Reality/Augmented Reality").

Description: A brief description of the service category and its typical characteristics.

Sub-Class Specification: A reference to the RFC or stable document that defines the Sub-Class values for this Service Class.

Contact: The responsible IETF working group or designated expert.

Initial registry contents are as follows:

* 0: RESERVED

* 1: VRAR Virtual Reality / Augmented Reality

* 2: SCICOMP Scientific Computing

* 3: IOTCRIT Critical IoT Control

* 4: RTC Real-Time Communication (e.g., VoIP, Video Conferencing)

* 5-31: Unassigned (reserved for future IETF allocation)

* 32-63: Experimental/Private Use (not guaranteed to be globally interoperable)

10.2. NST-TLP Sub-Class Registries

For each Service Class that requires detailed sub-category definitions, IANA SHALL maintain an associated Sub-Class registry or a normative appendix referenced from the main registry entry.

For example, for Service Class 1 (VRAR), a "VRAR Sub-Class" registry is defined using 8-bit values.

Example entries for the VRAR Sub-Class registry include:

* 0: RESERVED

* 1: VIDEO_2K (1080p)

* 2: VIDEO_4K

* 3: VIDEO_8K

* 4: POSE_HIGH (high-rate pose and control signaling)

* 5: AUDIO_3D (spatial audio)

10.3. Example of a Complete Label Encoding

This section provides an illustrative example of constructing a complete NST-TLP label.

Consider a 4K VR video flow that is marked as high priority and uses the basic label format. The Flow Hint field is set to 0xAB.

Service Class (VRAR): IANA-assigned value 1 (binary 00000001).

Sub-Class (VIDEO_4K): Sub-registry value 2 (binary 00000010).

Priority: High priority = 0001.

Version: Basic format = 0001.

Flow Hint: 0xAB (binary 10101011).

The resulting 32-bit label is:

0001 00000001 00000010 0001 10101011

Which corresponds to the hexadecimal value 0x10121AB.

This 4-octet label can be inserted into an IPv6 Hop-by-Hop Options header or a dedicated extension header, enabling intermediate network devices to recognize the flow as a high-priority 4K VR service and apply appropriate low-latency and high-bandwidth forwarding policies.

11. Label Insertion and Handling

11.1. Labeling Points

NST-TLP labels MAY be inserted at different points in the network, depending on the network architecture, policy management model, and device capabilities.

End systems (e.g., hosts, servers, and user devices) are the preferred labeling points because they can provide the most accurate service type information. Applications that support the NST-TLP API MAY specify the service type when creating sockets or sending data, and the operating system network stack is responsible for generating and encapsulating the corresponding label. The operating system MAY also infer service type automatically based on packet characteristics such as destination port, transport protocol, or process context. In virtualized environments, virtual network interface drivers or hypervisors MAY label traffic based on the service type of the associated virtual machine or container.

When end systems do not support label insertion, network edge devices MAY perform this function. As the first hop in the network, an edge device MAY insert labels based on deep packet inspection (DPI), flow feature analysis, or policy instructions received from a control plane (e.g., an SDN controller). In mobile networks, a base station MAY map bearer-level QoS identifiers (e.g., QCI in LTE or 5QI in 5G) to NST-TLP labels. Security gateways MAY also insert service type labels based on application identification results while enforcing security policies.

In software-defined networking (SDN) and network function virtualization (NFV) environments, label insertion is more flexible. Using southbound interfaces (e.g., OpenFlow), controllers MAY

instruct switches to insert labels for specific flows. At the ingress of an NFV service chain, labels MAY be assigned according to service chain policies.

11.2. Label Processing Rules

NST-TLP label-aware nodes SHOULD process labels in a consistent manner to ensure predictable network behavior. A node MUST identify the label encapsulation location according to the packet type (e.g., IPv4, IPv6, or Ethernet) and validate the label format, including the Version field and length. The node MUST verify whether the Service Class and Sub-Class values are known. For unknown values, the node SHOULD apply a default policy.

A node MUST apply forwarding and resource management actions based on the label contents. For queue scheduling, packets SHOULD be assigned to appropriate priority queues according to the Service Class, Sub-Class, and Priority fields. For example, VR 4K traffic MAY be placed in a strict priority queue, while background data transfers MAY be assigned to a best-effort queue. For path selection, forwarding decisions MAY consider the service type, such as selecting high-bandwidth paths for scientific computing traffic and low-latency paths for interactive VR traffic. Nodes MAY reserve bandwidth or compute resources for specific service types.

By default, a node SHOULD preserve the NST-TLP label and forward it with the packet. In certain cases, a node MAY update the label, for example, by updating the Flow Hint field to reflect path changes, adjusting the Priority field based on congestion conditions, or performing service type mapping at administrative domain boundaries.

A node SHOULD consider removing the label under the following conditions: when traffic leaves an NST-TLP-capable domain; when required by security policy; or when the packet reaches its final destination.

If a label does not conform to the specified format, a node MAY discard the packet (in security-sensitive environments), remove the label and forward the packet according to default policy (in best-effort environments), and/or generate an error log or an ICMP error message (subject to local configuration).

If no explicit policy is configured for a given Service Class or Sub-Class, the node SHOULD apply the default policy associated with that Service Class. If no such default policy exists, the node SHOULD apply the lowest-priority best-effort policy.

12. Use Cases

By providing standardized service type identification, NST-TLP enables intelligent traffic management and resource scheduling in a wide range of network scenarios. This section describes representative use cases.

12.1. Immersive VR/AR Service Assurance

VR/AR traffic is highly sensitive to latency (e.g., less than 20 ms) and jitter, and different sub-flows within the same session (such as 4K video, pose tracking signaling, and haptic feedback) have heterogeneous bandwidth and reliability requirements. Compared with using only DSCP, NST-TLP allows the network to distinguish different sub-flows of the same VR session and provide more fine-grained resource guarantees.

A VR application or operating system marks video streams as Service Class: VRAR and Sub-Class: VIDEO_4K, and marks critical pose control signaling as Sub-Class: POSE_HIGH with Priority set to a high value. Access switches and core routers recognize these labels and apply strict priority forwarding and low-latency path selection to POSE_HIGH traffic, while ensuring high bandwidth and enhanced reliability (e.g., by using forward error correction) for VIDEO_4K traffic.

12.2. High-Performance and Scientific Computing Networks

High-performance computing (HPC) jobs (e.g., MPI communication) often generate bursty elephant flows and require high throughput and low job completion time, while management and control traffic is latency sensitive. NST-TLP improves overall cluster utilization by enabling the network to become computation-aware rather than a transparent pipe.

A job scheduler or compute node marks MPI bulk traffic as Service Class: SCICOMP and Sub-Class: MPI_BULK, and MAY include a compact Job Identifier hash in the optional metadata. Data center switches build job-aware scheduling policies, isolating SCICOMP traffic from web or background traffic. For collective operations (e.g., MPI_BARRIER), higher-priority sub-classes MAY be used to ensure fast synchronization.

12.3. Critical IoT and Industrial Control

Industrial IoT control signaling requires deterministic latency and high reliability, whereas sensor telemetry may tolerate higher delay and jitter. NST-TLP enables critical control traffic to receive deterministic service on a shared physical network, reducing deployment and maintenance costs.

Programmable logic controllers (PLCs) or gateways mark motion control commands as Service Class: IOTCRIT and Sub-Class: MOTION_CTRL, while periodic sensor data is marked as Sub-Class: TELEMETRY. Time-Sensitive Networking (TSN) switches or 5G user plane functions (UPFs) map IOTCRIT traffic to scheduled transmission gates or guaranteed bit-rate bearers to ensure latency and reliability requirements.

12.4. Real-Time Communication Quality Improvement

Real-time applications such as video conferencing and cloud gaming contain multiple traffic components, including audio, video, signaling, and file sharing, which require differentiated treatment. NST-TLP allows the network to prioritize core interactive media even under congestion, thereby improving user experience.

A real-time communication client marks audio streams as Service Class: RTC and Sub-Class: AUDIO_INTERACTIVE, video streams as VIDEO_CONF, and file transfers as DATA_FILE. Enterprise WAN optimization devices or provider edge routers use these labels to assign highest priority to audio traffic, guarantee bandwidth for video traffic, and rate-limit file transfers.

12.5. Cross-Domain Network Slicing and Service Chaining

In network slicing and service function chaining provided by operators and cloud providers, a generic flow-level identifier is required to steer traffic through specific virtual networks or function chains. NST-TLP provides a lighter-weight alternative to deep packet inspection and a more flexible mechanism than destination-based classification.

When user traffic enters the network, ingress devices or controllers mark packets with slice-related NST-TLP labels, such as Service Class values representing eMBB or URLLC. Sub-Class or optional metadata MAY indicate the required service function chain (e.g., firewall, intrusion detection, and video optimization). Network nodes forward traffic to the corresponding virtual network functions based on these labels.

13. Security Considerations

13.1. Label Forgery and Spoofing

Malicious endpoints or on-path attackers may forge or modify NST-TLP labels, for example, by marking bulk download traffic as high-priority VR traffic in order to obtain preferential treatment. This results in a form of QoS theft and may degrade the service quality of legitimate flows.

Mitigation strategies include rewriting or validating labels at trust domain boundaries (e.g., user-network interfaces), such that labels provided by untrusted user devices are not directly honored. In high-security environments, optional metadata MAY carry a lightweight message authentication code (MAC) generated and verified by trusted anchors or controllers using shared keys. This approach introduces processing overhead and MUST be carefully evaluated.

13.2. Information Disclosure and Privacy

NST-TLP labels may expose sensitive information about user behavior, application usage, or operational status. For example, frequent occurrence of SCICOMP labels may indicate that a user is performing scientific computing tasks.

To mitigate privacy risks, gateways MAY map fine-grained Sub-Class values to coarser Service Class values when traffic traverses public or untrusted networks. Labels MAY also be removed at administrative boundaries according to policy. Encryption mechanisms such as IPsec or TLS MAY be used to protect packets carrying labels against eavesdropping, noting that this may limit intermediate nodes' ability to process the labels.

14. IANA Considerations

This document requests IANA to create and maintain registries for NST-TLP Service Classes and Sub-Classes as described in Section 8.

15. Normative References

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

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", RFC 8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

Authors' Addresses

Huanxing Zhu, Huazhong University of Science and Technology, Wuhan, China, Email: huanxingzhu@hust.edu.cn

Author's Address

Huanxing Zhu
Huazhong University of Science and Technology
Wuhan
China
Email: huanxingzhu@hust.edu.cn