

CATS
Internet-Draft
Intended status: Standards Track
Expires: 30 August 2026

H. Fu
ZTE Corporation
X. Yi
China Unicom
B. Pang
Beijing Jiaotong University
D. Yuan
W. Duan
C. Miao
ZTE Corporation
26 February 2026

Hybrid Forwarding Mechanism for Computing-Aware Traffic Steering (CATS)
draft-fu-cats-hybrid-fwd-01

Abstract

This document specifies a hybrid forwarding mechanism for Computing-Aware Traffic Steering (CATS). The mechanism integrates the CATS forwarding table with the IP forwarding table to optimize forwarding table capacity utilization. By customizing the forwarding model based on service identifiers, it accommodates both experience-sensitive and non-experience-sensitive services. Additionally, it supports flow-granularity load balancing, enhancing the utilization of computing and networking resources while ensuring differentiated service requirements are satisfied.

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 30 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. Introduction	2
2. Requirements Language	3
3. Terminology	3
4. Problem Statement	4
5. Hybrid Forwarding Model	5
5.1. Table Management and Working Mechanism	6
5.2. Considerations for Hybrid Forwarding	8
6. Security Considerations	11
7. Acknowledgements	11
8. IANA Considerations	11
9. References	11
9.1. Normative References	11
9.2. Informative References	12
Authors' Addresses	13

1. Introduction

Computing-Aware Traffic Steering (CATS) [I-D.ldbcb-cats-framework] targets efficient routing at the network edge, directing traffic between service clients and providers. It relies on real-time computing and network status data for informed decisions. CATS operates as an overlay system, choosing optimal service instances for requests. The CATS framework does not assume any specific data plane or control plane solutions.

As the integration of networking and computing deepens, users are placing increasingly higher demands on both the experience of network services and the efficiency of resource utilization. This is particularly evident for services with low latency requirements, such as AR/VR, and services with explicit requirements for computing resources. Traditional traffic steering mechanisms are not capable of meeting these evolving demands. Relevant studies have

demonstrated that traffic steering that takes into account computing resource conditions can significantly improve service performance and user experience [I-D.ietf-cats-usecases-requirements].

To address these challenges, the Computing-Aware Traffic Steering (CATS) framework has been proposed in [I-D.ldbc-cats-framework]. The primary goal of the CATS framework is to efficiently steer traffic at the edge device of the network instructed by traffic steering decisions referring to both computing and network conditions. The CATS framework is designed as an overlay mechanism for selecting the optimal service instances, without relying on specific data-plane or control-plane solutions.

This document specifies a hybrid forwarding mechanism for CATS with the following key features:

- * **Integration of Forwarding Tables:** The CATS forwarding table is integrated with the IP forwarding table to achieve shared and efficient utilization of the forwarding table capacity.
- * **Customized Forwarding Model:** The forwarding model is customized based on service identifiers to meet the demands of both experience-sensitive and non-experience-sensitive services.
- * **Flow-Level Load Balancing:** The mechanism supports flow-granularity load balancing, enhancing the utilization of computing and networking resources while meeting diverse requirements for differentiated services.

This hybrid forwarding mechanism provides a flexible and efficient solution for network vendors and service providers, facilitating the deployment of computing-aware traffic steering.

2. Requirements Language

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] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Terminology

This document uses terms defined in [I-D.ldbc-cats-framework], [I-D.lbdd-cats-dp-sr], and [I-D.fu-cats-flow-lb].

This document defines the following additional terms:

CATS-IND: An indicator in the forwarding table to identify CATS service packets.

MODE-IND: An indicator to differentiate service requirements (e.g., experience-sensitive vs. non-sensitive services).

Flow Affinity Table: A data plane table maintaining flow-to-next-hop bindings to ensure consistent forwarding for packets belonging to the same flow.

CATS Routing Table: A control plane data structure containing network paths, service instances, and associated load sharing ratios for a given service identifier.

CATS Forwarding Table: A data plane construct derived from the CATS Routing Table, optimized for packet forwarding decisions.

4. Problem Statement

To effectively steer service request packets to appropriate forwarding paths and service instances, CATS faces several issues and challenges with existing technical solutions:

Flow Affinity Mechanism Complexity: To ensure continuous access to the same computing resource during a service session, a flow affinity mechanism based on five-tuples or three-tuples is required. However, incorporating such lookups into the general routing process introduces additional processing overhead, which significantly increases packet forwarding latency and consumes chip resources. This additional processing burden can degrade the performance of latency-sensitive services.

Control Plane Overhead: Computing services are characterized by ubiquity and state changes over time. If the control plane indiscriminately performs periodic or event-driven route computation and table updates for all managed computing service types, it will excessively consume CPU resources and increase the burden of system management and maintenance.

Initial Packet Handling and Latency: Some solutions attempt to alleviate the control plane load by throttling and sending the first computing request packet to the control plane to generate a flow affinity table before the packet hits the flow affinity table. However, there is a risk of packet loss before the table entry takes effect, and the control plane processing of the packet can also lead to high latency for the first packet, which affects the quality of experience and service continuity.

5. Hybrid Forwarding Model

To address the aforementioned issues, this document specifies a hybrid forwarding mechanism. This solution features high flexibility, enabling customized service strategies tailored to different service characteristics:

Experience-Sensitive Services: Utilize direct table-based forwarding to guarantee stringent QoS requirements and optimal user experience.

Non-Sensitive Services: Employ packet-driven table updates to minimize resource overhead while maintaining forwarding efficiency.

Key advantages of this mechanism include:

- (i) **Reduced Control/Forwarding Plane Overhead:** Mitigates update pressure through optimized table management.
- (ii) **Decoupled Legacy Router Dependency:** Enhances the flexibility of deployment in modern network environments.
- (iii) **High Implementability and Scalability:** Supports evolving service demands while maintaining backward compatibility.

This approach provides an efficient and scalable solution for heterogeneous service delivery in next-generation networks.

The Hybrid Forwarding mechanism operates within the framework defined in the CATS architecture [I-D.ldb-cats-framework] (see Figure 1).

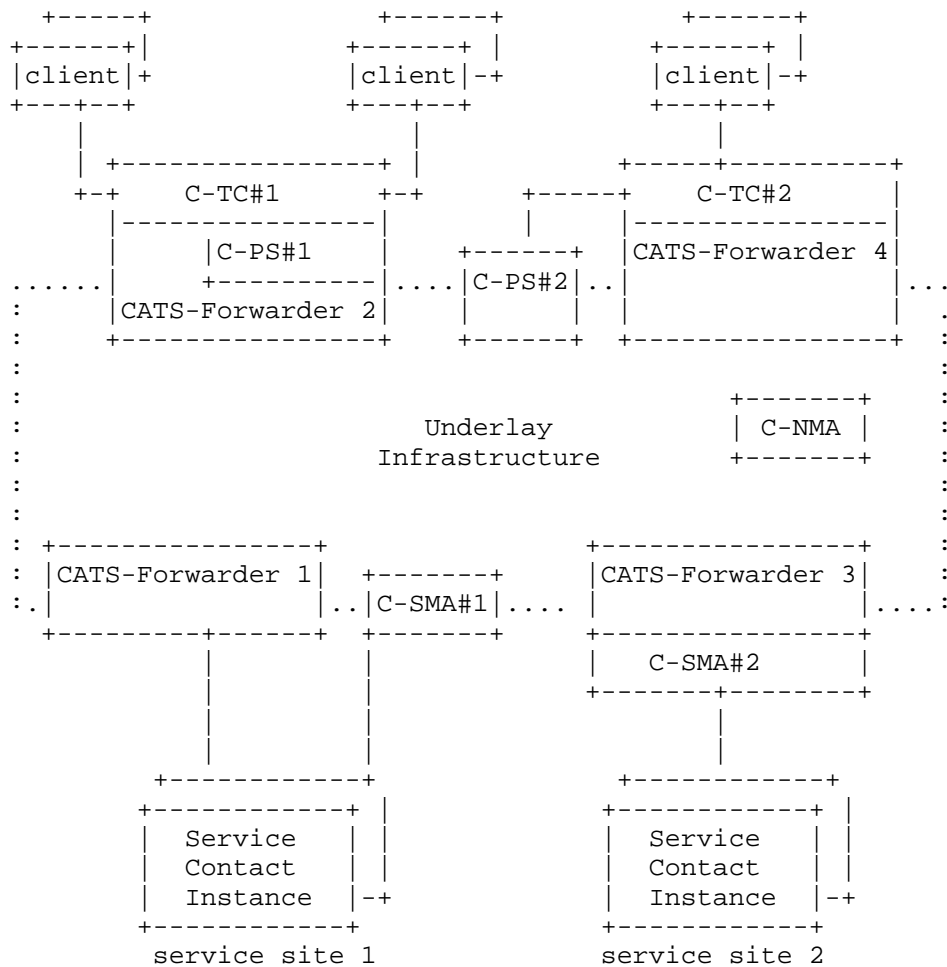


Figure 1: CATS Functional Components

5.1. Table Management and Working Mechanism

CATS ROUTING TABLE

NextHop	VPN ID	Prefix	IP	CATS-IND	MODE-IND	SR-Policy	Service	SID
1	CS-ID-1			1	1	SR-P-1	END-DX6-1	
1	CS-ID-1			1	0	SR-P-2	END-DX6-2	

|
v

CATS FORWARDING TABLE

NextHop	VPN ID	Prefix	IP	CATS-IND	MODE-IND	SR-Policy	Service	SID
1	CS-ID-1			1	1	SR-P-1	END-DX6-1	
1	CS-ID-1			1	0	NULL	NULL	

Figure 2: CATS Table Management by the Control Plane

Figure 2 presents a specific form of managing the CATS table entries on the control plane. The explanation for the entries in the table is as follows:

CS-ID Usage: The CS-ID typically employs anycast IP addresses to identify services. As specified in [I-D.lbdd-cats-dp-sr] and [I-D.fu-cats-muti-dp-solution], the forwarding mechanism varies based on service instance connectivity: (1) when an egress gateway is connected to multiple service instances, traffic is forwarded to a specific instance via either a tunnel or a general interface (on the same CATS-Forwarder) using an END.DX4/6 Service SID; (2) when a single service instance is involved, an END.DT4/6 Service SID enables direct forwarding via anycast IP. Consequently, each Service SID maintains a strict one-to-one correspondence with its associated service instance.

C-PS Deployment and Function: The C-PS component is deployed on the head node or a centralized computation-network controller. It collects and reports CS-ID, CIS-ID, and metric information related to service instances via C-SMA, as well as network-related capabilities and status information via C-NMA. Based on the collected information, the C-PS calculates the optimal network forwarding path and service instance, which is then referred to as the CATS routing table. The process of inserting or downloading these table entries to the forwarding plane is called the CATS forwarding table update.

Design and Function of the CATS Forwarding Table: The CATS forwarding table shares space with the traditional IP forwarding table. To distinguish CATS service packets from other types of packets, CS-ID is used for marking and identification, with CATS-IND serving as an indicator. Additionally, to differentiate service requirements (e.g., delay-sensitive or other types), MODE-IND can be used for marking and labeling services based on CS-ID. As shown in Figure 2, relevant table entries are provided for reference. For example, MODE-IND = 1 indicates experience-sensitive services, while MODE-IND = 0 indicates non-sensitive services. It should be noted that although Figure 2 shows only one NextHop for the same CS-ID in the CATS routing and forwarding tables, multiple NextHops can still be supported to achieve flow-level load balancing according to [I-D.fu-cats-flow-lb].

CATS Flow Affinity Table Mechanism: To ensure that service sessions consistently access the same service instance, CATS introduces the flow affinity table. This table uses a five-tuple or three-tuple as the key, with network path and service instance as related attributes. By looking up the table with the five-tuple or three-tuple learned from the service packet, the relevant attributes are retrieved, and the forwarding action is executed accordingly.

5.2. Considerations for Hybrid Forwarding

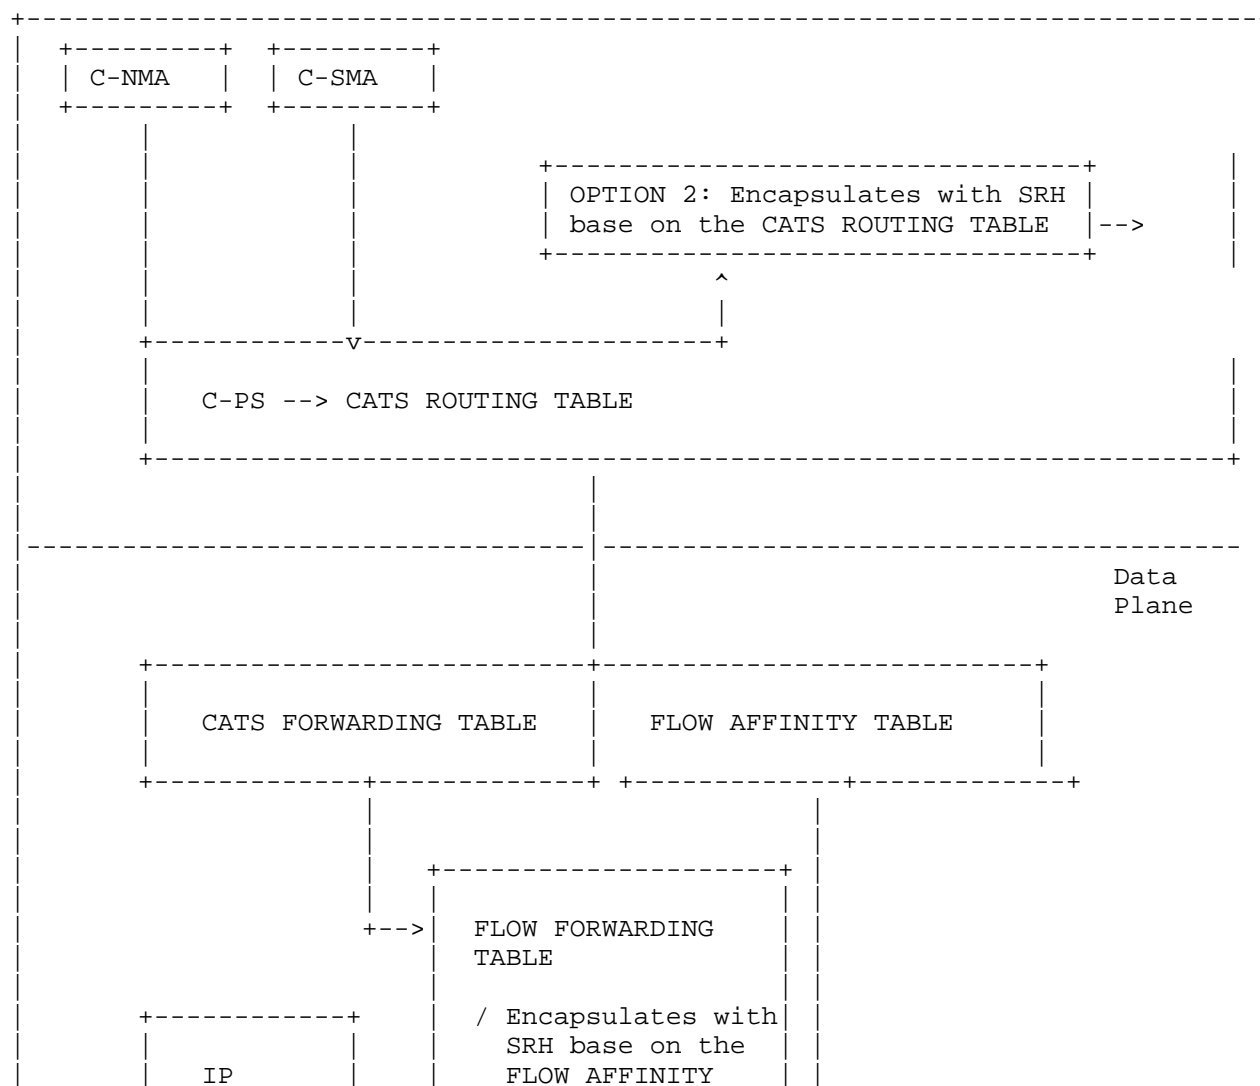
To meet the demands of both experience-sensitive and non-sensitive services, this document specifies the following forwarding mechanism in conjunction with the key table entries discussed in Figure 2.

Since the CS-ID is encoded in the form of anycast addresses, the CATS forwarding table and the IP forwarding table can share the same space. Typically, traffic is distinguished by looking up the IP forwarding table. Once CATS packets are identified, further flow forwarding processing is carried out. This approach reduces invasive modifications to the forwarding mechanism while maintaining high-performance forwarding for both traditional services and CATS services.

Specifically:

- (i) For the case in which MODE-IND = 1, the issued CATS forwarding table includes SR-policy and Service SID information. The control plane only sends updates to the data plane when table entries change, ensuring that both the first packet and subsequent packets are processed directly at the forwarding plane, while the control plane is solely responsible for generating the flow affinity table.

- (ii) For the case in which `MODE-IND = 0`, the issued CATS forwarding table does not include SR-policy and Service SID information. In this case, regardless of how the attributes of the CATS routing table change dynamically, the C-PS will not update the forwarding plane to reduce the overhead of table updates. The first packet is uploaded to the control plane, which queries the CATS routing table, encapsulates and forwards the packet, and generates the flow affinity table to be sent to the forwarding plane. Subsequent packets are forwarded based on the flow affinity table in the data plane.



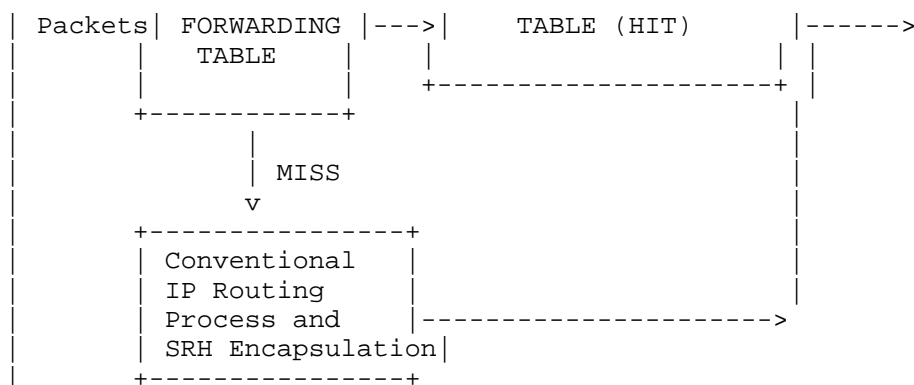


Figure 3: Hybrid Forwarding Mechanism

As shown in Figure 3, the relevant forwarding process is as follows:

- (i) Upon arrival, a packet first looks up the IP forwarding table. If the CATS_IND in the lookup result is 0, the packet is identified as a conventional non-CATS service packet; otherwise, it is identified as a CATS service packet, and the flow forwarding process is initiated.
- (ii) After entering the flow forwarding processing, the MODE-IND identifier in the lookup result is further recognized, and different processing is performed based on its value.
- (iii) If MODE-IND = 1, the first packet extracts five-tuple/three-tuple information, sends it to the control plane to generate the flow affinity table and send it to the forwarding plane. At the same time, the first packet is encapsulated and forwarded directly at the forwarding plane based on the SR-policy and Service SID information carried in the lookup result (Option 1).
- (iv) If MODE-IND = 0, the forwarding table entry does not carry SR-policy and Service SID information. The first packet is sent directly to the control plane, which queries the CATS routing table, generates the flow affinity table, and sends it to the forwarding plane. At the same time, the first packet is encapsulated and forwarded at the control plane based on the SR-policy and Service SID information carried in the lookup result (Option 2).

Based on this, for experience-sensitive services, the first packet can be forwarded directly at the forwarding plane to reduce first-packet latency; for non-experience-sensitive services, the first packet can be processed at the control plane to reduce the processing overhead of control-plane table updates.

By adopting the above measures, a flexible CATS hybrid forwarding solution can be selected according to the characteristics of different services associated with the CS-ID, ensuring optimal performance for different service types.

6. Security Considerations

This document does not introduce new security considerations beyond those discussed in [I-D.ldb-cats-framework]. The hybrid forwarding mechanism relies on the security properties of the underlying CATS framework and the data plane technologies used (e.g., SRv6).

Implementations SHOULD ensure that:

- * Flow affinity table entries are protected against exhaustion attacks through rate limiting of new flow creation.
- * Control plane updates to the CATS forwarding table are authenticated and authorized.
- * First packets processed by the control plane (MODE-IND = 0) are subject to appropriate validation to prevent control plane overload.

7. Acknowledgements

TBD.

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. Normative References

[I-D.ldbcs-cats-framework]

Li, C., Du, Z., Boucadair, M., Contreras, L. M., and J. Drake, "A Framework for Computing-Aware Traffic Steering (CATS)", Work in Progress, Internet-Draft, draft-ldbcs-cats-framework-06, 8 February 2024, <<https://datatracker.ietf.org/doc/html/draft-ldbcs-cats-framework-06>>.

[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>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

[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>>.

9.2. Informative References

[I-D.ietf-cats-usecases-requirements]

Yao, K., Contreras, L. M., Shi, H., Zhang, S., and Q. An, "Computing-Aware Traffic Steering (CATS) Problem Statement, Use Cases, and Requirements", Work in Progress, Internet-Draft, draft-ietf-cats-usecases-requirements-14, 2 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-cats-usecases-requirements-14>>.

[I-D.li-dynicast-architecture]

Li, Y., Iannone, L., Trossen, D., Liu, P., and C. Li,
"Dynamic-Anycast Architecture", Work in Progress,
Internet-Draft, draft-li-dynicast-architecture-08, 16
January 2023, <[https://datatracker.ietf.org/doc/html/
draft-li-dynicast-architecture-08](https://datatracker.ietf.org/doc/html/draft-li-dynicast-architecture-08)>.

[I-D.huang-service-aware-network-framework]

Huang, D., Tan, B., and D. Yang, "Service Aware Network
Framework", Work in Progress, Internet-Draft, draft-huang-
service-aware-network-framework-01, 22 November 2022,
<[https://datatracker.ietf.org/doc/html/draft-huang-
service-aware-network-framework-01](https://datatracker.ietf.org/doc/html/draft-huang-service-aware-network-framework-01)>.

[I-D.lbdd-cats-dp-sr]

Li, C., Du, Z., Drake, J., and shangyuxiang, "Computing-
Aware Traffic Steering (CATS) Using Segment Routing", Work
in Progress, Internet-Draft, draft-lbdd-cats-dp-sr-06, 13
October 2025, <[https://datatracker.ietf.org/doc/html/
draft-lbdd-cats-dp-sr-06](https://datatracker.ietf.org/doc/html/draft-lbdd-cats-dp-sr-06)>.

[I-D.fu-cats-muti-dp-solution]

Huakai.Fu, Liu, B., Li, Z., Huang, D., Yuan, D., Ma, L.,
and W. Duan, "Analysis for Multiple Data Plane Solutions
of Computing-Aware Traffic Steering", Work in Progress,
Internet-Draft, draft-fu-cats-muti-dp-solution-03, 19
August 2025, <[https://datatracker.ietf.org/doc/html/draft-
fu-cats-muti-dp-solution-03](https://datatracker.ietf.org/doc/html/draft-fu-cats-muti-dp-solution-03)>.

[I-D.fu-cats-flow-lb]

Fu, H., Huang, D., Duan, W., and B. Tan, "Flow-Level Load
Balancing of Computing-Aware Traffic Steering (CATS)",
Work in Progress, Internet-Draft, draft-fu-cats-flow-lb-
03, 26 February 2026,
<[https://datatracker.ietf.org/doc/html/draft-fu-cats-flow-
lb-03](https://datatracker.ietf.org/doc/html/draft-fu-cats-flow-lb-03)>.

[RFC7094] McPherson, D., Oran, D., Thaler, D., and E. Osterweil,
"Architectural Considerations of IP Anycast", RFC 7094,
DOI 10.17487/RFC7094, January 2014,
<<https://www.rfc-editor.org/info/rfc7094>>.

Authors' Addresses

Huakai Fu
ZTE Corporation
Wuhan
China

Email: fu.huakai@zte.com.cn

Xinxin Yi
China Unicom
Beijing
China
Email: yixx3@chinaunicom.cn

Bo Pang
Beijing Jiaotong University
Beijing
China
Email: bopang@bjtu.edu.cn

Dongyu Yuan
ZTE Corporation
Nanjing
China
Email: yuan.dongyu@zte.com.cn

Wei Duan
ZTE Corporation
Nanjing
China
Email: duan.weil@zte.com.cn

Chuanyang Miao
ZTE Corporation
Nanjing
China
Email: miao.chuanyang@zte.com.cn