

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

H. Fu  
D.H. Huang  
W. Duan  
B. Tan  
ZTE Corporation  
26 February 2026

Flow-Level Load Balancing of Computing-Aware Traffic Steering (CATS)  
draft-fu-cats-flow-lb-03

Abstract

This document specifies a flow-level load balancing mechanism for Computing-Aware Traffic Steering (CATS) that reduces control plane overhead and improves resource utilization through data plane autonomous flow distribution.

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 . . . . .	2
3. Terminology . . . . .	2
4. Problem Statement . . . . .	3
5. Architectural Model . . . . .	3
6. Control plane Operation . . . . .	5
7. Data Plane Operation . . . . .	6
8. Security Considerations . . . . .	7
9. Acknowledgements . . . . .	7
10. IANA Considerations . . . . .	7
11. References . . . . .	7
11.1. Normative References . . . . .	7
11.2. Informative References . . . . .	8
Authors' Addresses . . . . .	9

## 1. Introduction

Computing-Aware Traffic Steering (CATS) [I-D.ldb-cats-framework] directs traffic between service clients and providers based on real-time computing and network status. While CATS operates as an overlay system for selecting optimal service instances, the framework does not assume specific data plane or control plane solutions.

This document defines a flow-level load balancing mechanism addressing two operational challenges: control plane scalability limitations caused by reactive path computation, and resource utilization imbalances resulting from coarse-grained status reporting. The mechanism enables the control plane to pre-compute multiple viable forwarding alternatives while allowing the data plane to autonomously distribute traffic across these alternatives using flow-based affinity.

## 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 makes use of the terms defined in [I-D.ldb-cats-framework].

- \* CATS Routing Table: Control plane structure containing service identifiers, network paths, service instances, and load sharing weights.
- \* CATS Forwarding Table: Data plane structure derived from CATS Routing Table via UCMP-to-ECMP expansion.
- \* Flow Affinity Table: Data plane table maintaining flow-to-forwarding-entry bindings.

#### 4. Problem Statement

Current CATS implementations rely on periodic or threshold-triggered resource status reports to optimize service instance and path selection. This approach creates uneven computing resource utilization when status updates lag behind actual load changes, potentially directing multiple requests to already overloaded instances. Additionally, frequent metric fluctuations trigger repeated control plane path recalculation and policy updates, creating scalability constraints that incremental calculation alone cannot resolve.

#### 5. Architectural Model

The Flow-Level Load Balancing of Computing-Aware Traffic Steering is constructed based on the framework established in the CATS architecture [I-D.ldbc-cats-framework](Figure 1 for a visual representation).




Figure 1: CATS-Functional-Components

The mechanism operates through three functional entities: the Path Selector (C-PS) situated in the control plane collecting metrics via Metric Agents (C-SMA and C-NMA) and computing forwarding alternatives; the Forwarder operating in the data plane performing flow identification, affinity maintenance and packet forwarding; and the Metric Agents reporting service instance and network status to enable path computation.

The Path Selector maintains the CATS Routing Table containing for each CS-ID a set of Forwarding Alternatives comprising network path identifiers, service instance identifiers, and load sharing weights representing desired traffic distribution proportions. The Forwarder

generates the CATS Forwarding Table through UCMP-to-ECMP expansion, creating a uniform lookup structure where each Forwarding Alternative appears in proportion to its weight, and maintains the Flow Affinity Table binding flow identifiers to specific forwarding entries.

## 6. Control plane Operation

The Path Selector identifies for each CS-ID the Service Instance Set containing all Forwarding Alternatives satisfying SLA requirements, computing load sharing weights based on service instance metrics, network path characteristics and policy objectives. The resulting CATS Routing Table is translated into a CATS Forwarding Table and distributed to Forwarders..

Figure 2 shows an example of a representation of multi-next-hop CATS routing table designed for a specific CS-ID1.

VRF-ID	PREFIX	NEXT HOP		
		SR-Policy	Service SID	Load Sharing Ratio
100	CS-ID1	SR-Policy1(2ms)	END.DX-1	20%
		SR-Policy1(2ms)	END.DX-2	30%
		SR-Policy2(1.5ms)	END.DX-3	30%
		SR-Policy2(1.5ms)	END.DX-4	20%

Figure 2: An example of CATS routing table

Figure 3 shows an example of the CATS forwarding table following the changes.

VRF-ID	PREFIX	SR-Policy	Service SID	offset
100	CS-ID1	SR-Policy1(2ms)	END.DX-1	0
		SR-Policy1(2ms)	END.DX-1	1
		SR-Policy1(2ms)	END.DX-2	2
		SR-Policy1(2ms)	END.DX-2	3
		SR-Policy1(2ms)	END.DX-2	4
		SR-Policy2(1.5ms)	END.DX-3	5
		SR-Policy2(1.5ms)	END.DX-3	6
		SR-Policy2(1.5ms)	END.DX-3	7
		SR-Policy2(1.5ms)	END.DX-4	8
		SR-Policy2(1.5ms)	END.DX-4	9

Figure 3: An example of CATS forwarding table

To minimize control plane load, the Path Selector regenerates the CATS Routing Table only when metric changes cross predefined thresholds rather than reacting to every fluctuation, optionally employing hysteresis mechanisms to prevent rapid oscillation between configurations.

## 7. Data Plane Operation

Upon receiving a packet the Forwarder extracts the CS-ID and computes the flow identifier, typically derived from the five-tuple. The Forwarder queries the Flow Affinity Table and if an existing binding is found forwards the packet using the bound entry. For unbound flows the Forwarder computes a hash over the flow identifier, selects the CATS Forwarding Table entry at the resulting index, creates a Flow Affinity Entry binding the flow to this entry, and forwards the packet.

The CATS Forwarding Table remains stable during flow affinity establishment ensuring consistent forwarding for all packets of a flow. Table updates received from the Path Selector apply only to new flows while existing bindings remain valid until flow termination or timeout.

## 8. Security Considerations

The mechanism introduces potential denial of service vulnerabilities through Flow Affinity Table exhaustion if an attacker generates excessive new flows, mitigated by implementing binding creation rate limits. Predictable flow identifiers could enable binding hijacking requiring cryptographically robust hash functions. CATS Forwarding Tables distributed from Path Selector to Forwarder may expose topology information warranting protection during transmission.

## 9. Acknowledgements

To be added upon contributions, comments and suggestions.

## 10. IANA Considerations

TBA

## 11. References

### 11.1. Normative References

- [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>>.

## 11.2.   Informative References

### [I-D.fu-cats-hybrid-fwd]

Huakai.Fu, Yi, X., Pang, B., Yuan, D., Duan, W., and C. Miao, "Hybird Fordwarding of Computing-Aware Traffic Steering (CATS)", Work in Progress, Internet-Draft, draft-fu-cats-hybrid-fwd-00, 30 June 2025, <<https://datatracker.ietf.org/doc/html/draft-fu-cats-hybrid-fwd-00>>.

### [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>>.

### [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>>.

### [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.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>>.



[I-D.ldb-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-ldb-cats-framework-06, 8 February 2024, <<https://datatracker.ietf.org/doc/html/draft-ldb-cats-framework-06>>.

[I-D.li-dyncast-architecture]

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

[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](mailto:fu.huakai@zte.com.cn)

Daniel Huang  
ZTE Corporation  
Nanjing  
China  
Email: [huang.guangping@zte.com.cn](mailto:huang.guangping@zte.com.cn)

Wei Duan  
ZTE Corporation  
Nanjing  
China  
Email: [duan.weil@zte.com.cn](mailto:duan.weil@zte.com.cn)

Bin Tan  
ZTE Corporation  
ShangHai  
China  
Email: [tan.bin@zte.com.cn](mailto:tan.bin@zte.com.cn)