

IDR  
Internet-Draft  
Intended status: Standards Track  
Expires: 7 November 2025

Y.DY. Yuan  
Z.FL. Zhou  
D.H. Huang  
C.QD. Chen  
D.CN. Dai  
ZTE Corporation  
6 May 2025

Computing Aware Traffic Steering (CATS) with Generic Metric  
draft-yuan-idr-generic-metric-cats-01

Abstract

Steering traffic for computing-related services considering computing resources and circumstances is discussed in CATS WG. Correspondingly, publishing services and updating computing conditions turns out to be a significant issue. It SHOULD be realized that multiple same common metrics are required from both network and service instances in order to evaluate overall performance and further achieve and fulfill appropriate traffic steering and scheduling. Therefore, an implementation for distributed CATS with generic metrics delivery and distribution based on BGP is proposed and discussed in this draft.

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 7 November 2025.

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 . . . . .	2
2. Requirements Language . . . . .	4
3. Generic Metrics for CATS . . . . .	4
4. Senario 1: Minimum End-to-end Latency for Computing-related Service . . . . .	5
5. Senario 2: Minimum Cost for Computing-related Service with constrained latency . . . . .	8
6. Senario 3: Normalized Metrics in Distribution Process . . . . .	9
7. Conclusion . . . . .	12
8. Security Considerations . . . . .	13
9. Acknowledgements . . . . .	13
10. IANA Considerations . . . . .	13
11. Normative References . . . . .	13
Authors' Addresses . . . . .	14

## 1. Introduction

Since for computing related services, AR/VR, metaverse for instance, the performance experienced by clients and customers is determined not only by network metrics but also by computing circumstances. Relevant use cases and problem statements are discussed in [I-D.ietf-cats-usecases-requirements]. For CATS framework introduced in [I-D.ietf-cats-framework], it would be an essential and significant issue of computing metrics publishing and updating for CATS.

Generally, control plane for CATS could be organized and deployed in various patterns and forms depending on the specific schemes of computing metrics collection and notification, instance selection and path calculation and other workflows. Especially for distributed metrics collection and distributed control plane implementations, protocols including BGP, BGP-LS, IGP would be mentioned to extend their capabilities to support metrics distribution and collection.

Furthermore, for computing metrics, they could be classified into multiple types and categories. A typical instance for computing metric analysis and discussion is presented in

[I-D.ysl-cats-metric-definition]. Generally, there could be converted, abstract and generic metrics or explicit metadata. In another aspect, to achieve end-to-end service provisioning, metrics of same dimensions among network infrastructure and service instances SHOULD be considered together while unique types of computing metrics MAY be processed independently.

General considerations for metrics which MAY be distributed and utilized in CATS are discussed below.

- \* Generic and common metrics: Latency, bandwidth and converted abstract metrics or costs (TE metrics, Costs, etc) for instance. Service instances and computing resources share these same types of metrics with network infrastructure. The accumulation of latencies would reflect the end-to-end delay. Similarly, a minimum bandwidth of the forwarding paths would indicate the overall capacity. Thus, potential requirements for comprehensive considerations of overall generic metrics SHOULD be noted.
- \* Unique metrics originated from specific areas (computing-related services, clusters, etc.): Computing capabilities, available memories, existing connections for instance. Commonly, network devices and network links do not have these similar metrics. Thus, if these metrics are distributed to the network, they turn out to be unique types and are not natively recognized. To evaluate these metrics, they would be relatively considered independently.

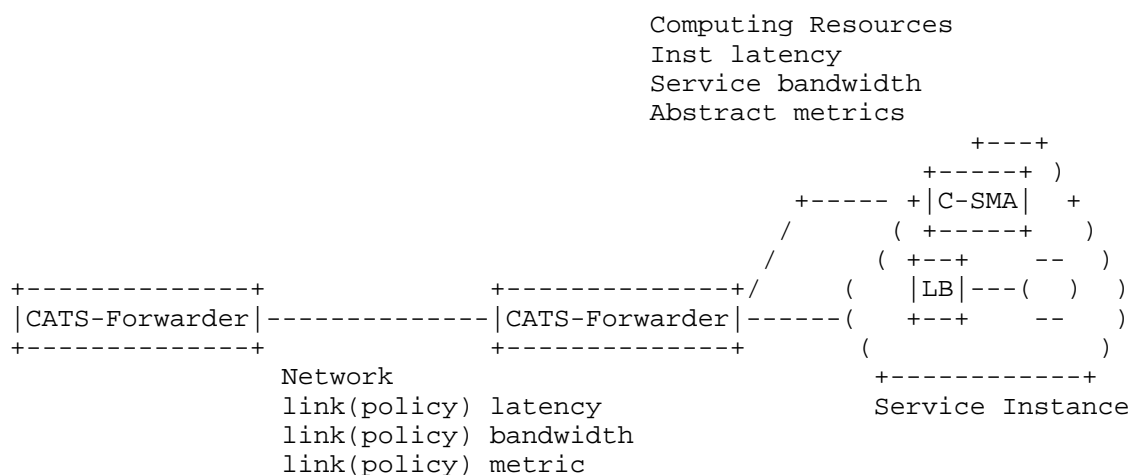


Figure 1: Network and Computing Metrics

In distributed control plane scenarios, especially when the service traffic needs to traverse multiple ASes, computing metrics SHOULD be distributed among CATS-Forwarders and be considered when performing ordered updates of routes. Thus, a distribution scheme based on generic metric introduced in [I-D.ietf-idr-bgp-generic-metric] is proposed in this draft. Generic metric is proposed to accumulate and propagate different types of metrics as it will aid in intent-based end-to-end path across BGP domains. Similarly, CATS SHOULD also be recognized as another intent-based end-to-end routing scenario. Computing-related services would be identified with multiple intents and thus these intents and relevant metrics SHOULD be able to be distributed. Furthermore, computing metrics, especially generic and common types of metrics, require to be accumulated and thus processed along the path of distribution. Detailed implementation will be introduced and discussed in the following sections.

## 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. Generic Metrics for CATS

In [I-D.ietf-idr-bgp-generic-metric], Accumulative Metric is defined.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Accumulated Metric Code | Accumulated Metric Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Accumulated Metric Data...
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 2: AMetric TLV

For the field of Metric Type in Accumulated Metric Data, values would be determined from IGP-Protocol registry for metric-types. Thus, parameters including latency, upstream/downstream bandwidth and configured TE metric of service instances could be encoded accordingly for a CATS scenario, in order to be processed in a general accumulative manner along the path.

```

      0              1              2              3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| Metric-type 1 | Metric-flags1 | Metric 1 value...
+-----+-----+-----+-----+
      0              1              2              3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| Metric-type 2 | Metric-flags2 | Metric 2 value...
+-----+-----+-----+-----+
      0              1              2              3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| Service Metric| Metric-flags  | Service Metric value...
+-----+-----+-----+-----+

```

Figure 3: Accumulative Metrics

Besides metric types defined with IGP registry, unique metric types would also be considered for a CATS scenario to extend and modify a current AMetric scheme. Suppose a general Service Metric or Cost would be proposed which specify the estimated or tested performance of a service instance with an abstract value. With normalized Service Metric and multiple dimensions of existing generic metrics, the implementations for CATS turn out to be various patterns. Regarding similar classifications for manifestations of discontinuity, typical scenarios will be displayed in the following sections.

#### 4. Scenario 1: Minimum End-to-end Latency for Computing-related Service

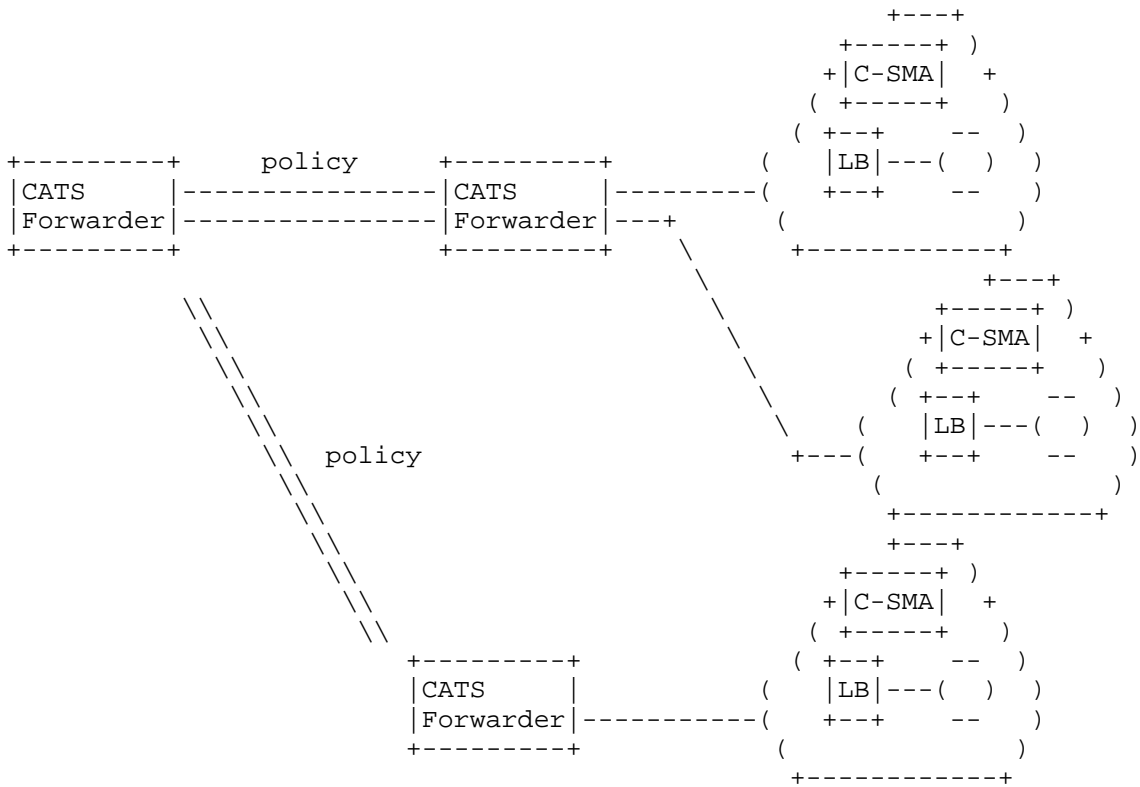


Figure 4: Minimum End-to-end Latency for Computing-related Service

1. C-SMAs collect computing-related metrics and pre-process relevant metadata. C-SMAs would be configured to establish BGP peers to CATS-Forwarders and thus distribute and update computing metrics with Generic Metric attribute. Suppose services deployed here require minimum end-to-end latency, delay would be filled in the update packets according to Generic Metric. Here, service routes MAY be distributed with next hop as a load balancer.
2. Services would be deployed in VRFs or a public VRF. CATS-Forwarders might be enabled to detect the latency to their correlated load balancers. Thus, service routes of same prefixes are updated with accumulated latency values. The value includes a processing delay of service instances and a detected delay between the CATS-Forwarder and the load balancer. Comparing among routes of same service prefixes, these routes would be re-ordered determined by the accumulated latency. When selecting a

best route, the service route will be distributed to the remote device and the next hop would be modified as the CATS-Forwarder itself.

3. Similarly, remote CATS-Forwarders would be able to detect the latency of policies or network links. Therefore, CATS-Forwarders could calculate the end-to-end latency values for each candidate service instance with resolved TE policies. Identically, ordered updates are performed and best routes are correspondingly determined. Since a delay parameter is accumulated along the path of service routes distribution, the accumulation would aid remote CATS-Forwarders to perform the specific latency-intent-based path selection.

The workflow also works for circumstances when service traffic needs to traverse multiple ASes. The end-to-end latency would be accumulated and calculated along the path of service routes distribution.




Figure 5: End-to-end Latency Accumulation among Multiple ASes

5. Senario 2: Minimum Cost for Computing-related Service with constrained latency

(For Inst 1 and 2)

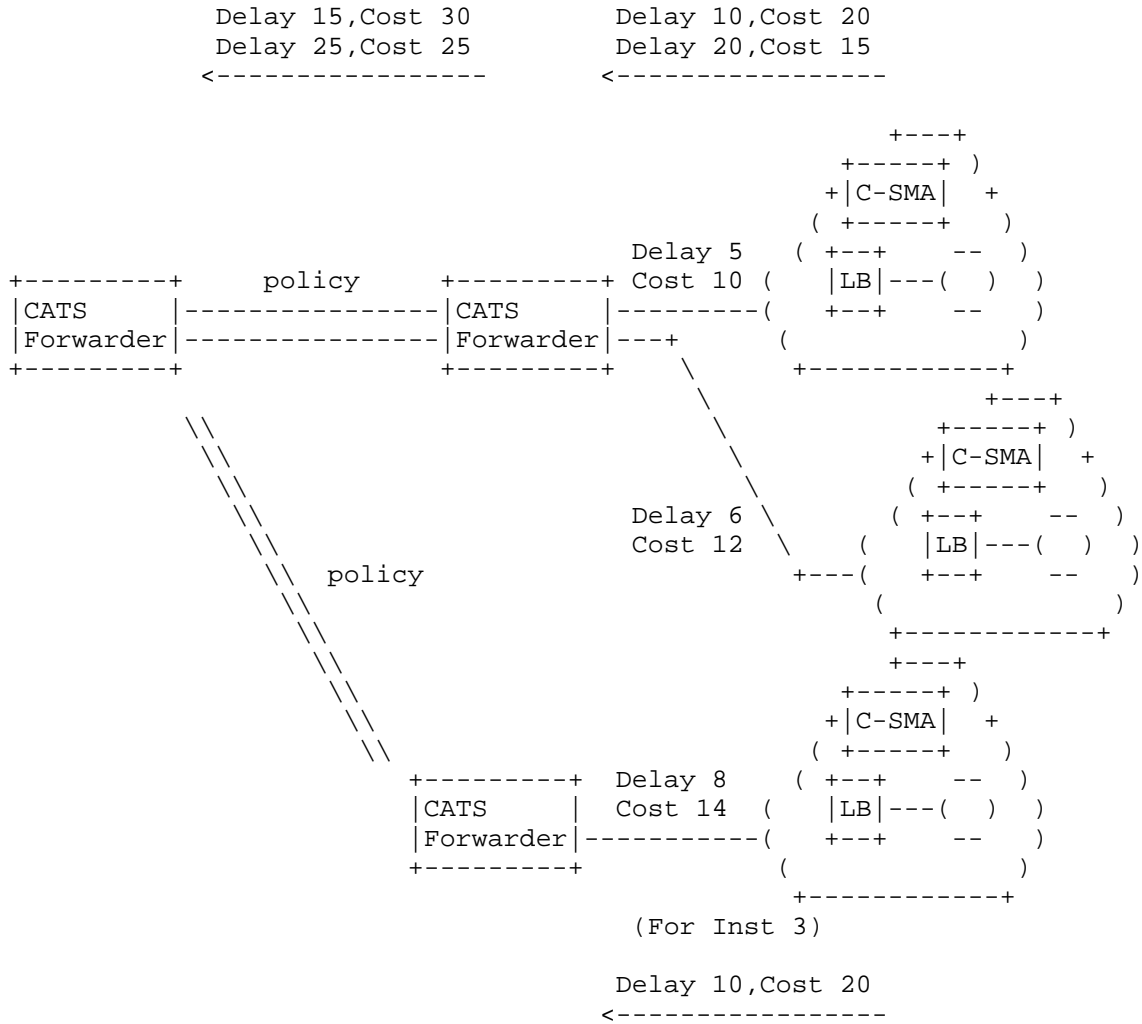




Figure 6: Minimum Cost for Computing-related Service with constrained latency

1. Similar to Scenario 1, C-SMAs collect computing-related metrics and distribute computing metrics with Generic Metric attribute. Suppose services deployed here require minimum end-to-end cost, TE metric for instance. Additionally, end-to-end latency is configured as constraints for ordered updates of routes. Converted costs and detected latency values would be filled in the update packets.
2. Service routes of same prefixes are updated with accumulated latency values and costs. The latency value includes a processing delay of service instances and a detected delay between the CATS-Forwarder and the load balancer. Similarly, The cost value includes a notified cost and a configured cost to the next hop. Additional path MAY be enabled at CATS-Forwarders, and thus service route will be distributed to the remote device and the next hop would be modified as the CATS-Forwarder itself.
3. Finally, remote CATS-Forwarders calculate the end-to-end latency values and overall costs for each candidate service instance with resolved policies or forwarding paths. Ordered updates with configured constraints are performed and best or appropriate routes are correspondingly determined.

Therefore, a generic metric scheme would work well for multi-factor scenarios.

#### 6. Scenario 3: Normalized Metrics in Distribution Process

It SHOULD be considered that generic metrics MAY be not always supported for each ASes and devices alongside the distribution process. Under certain circumstances, these metrics would be normalized or be transmitted unchanged.

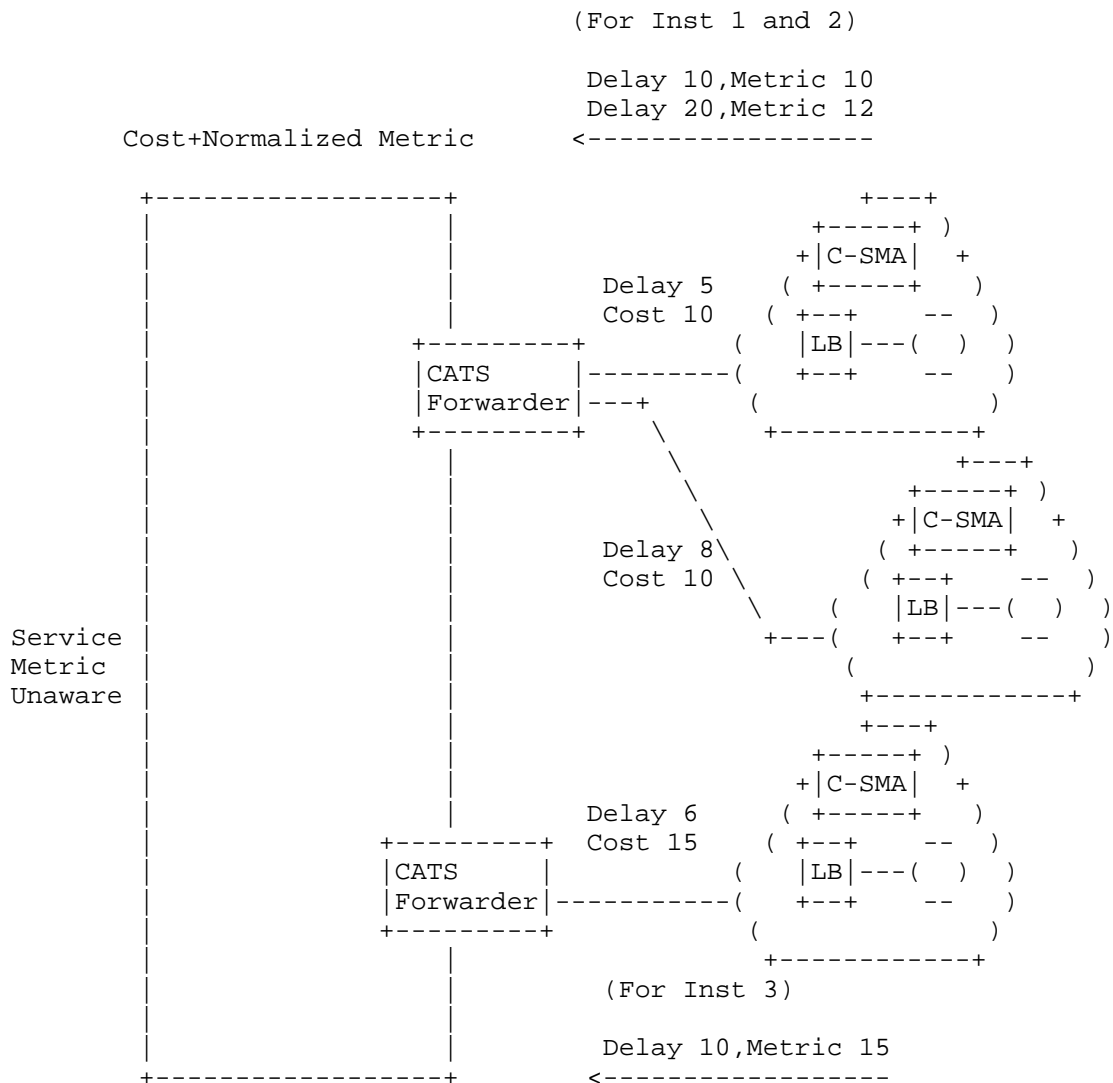


Figure 7: Minimum Cost for Computing-related Service with  
constrained latency

Normalization algorithms and strategies could be configured at CATS-Forwarders. When an AS or device is unaware of specific type of generic metric, a service metric displayed in the figure for instance, the metric value could be converted and normalized. For instance, service metric values could be magnified ten-fold to be common IGP Cost values. Afterwards, normalized values could be accumulated with IGP Costs to next hop. With the other implementation, unrecognized values would be transmitted unchanged if the remote devices are capable of analyzing such metrics. Ordered updates of service routes could be performed with a purpose of minimum service metric with constraints of end-to-end latency and cost.

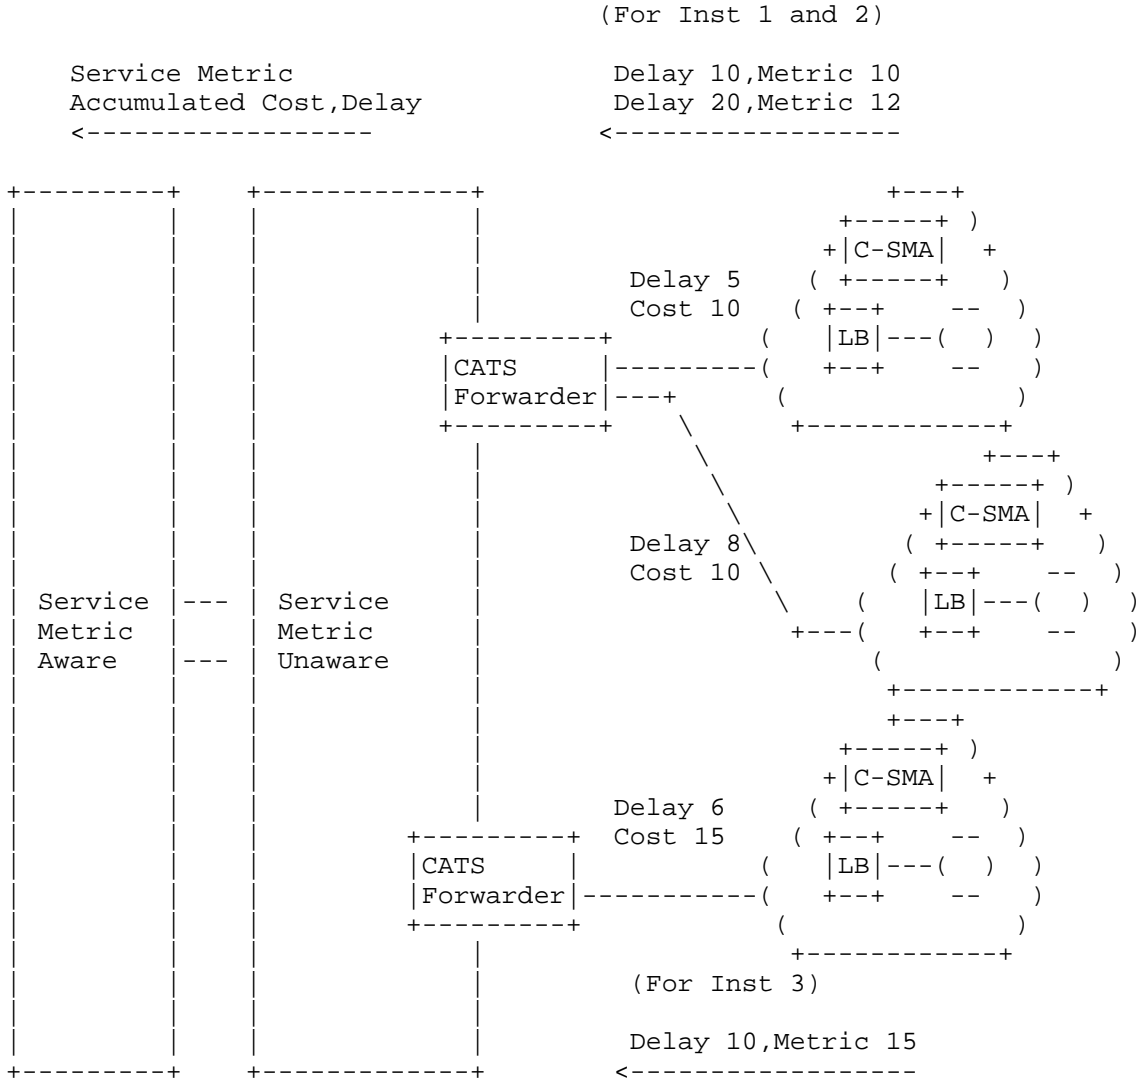


Figure 8: Minimum Service Metric for Computing-related Service with constrained latency and cost

## 7. Conclusion

About Computing Aware Traffic Steering (CATS) with Generic Metric, several considerations SHOULD be noted:

- \* It mainly applies for circumstances of distributed control plane for CATS. For a centralized control plane based on controllers or orchestrators, there might be existing interfaces for the collection of computing metrics.
- \* Generic common metrics between network and computing resources SHOULD be considered as significant factors which aid routes selection, especially for conditions of the provisioning of end-to-end services.
- \* Flexible and complex metadata or unique metrics are suggested to be normalized as simple and abstract factors which would restrain route oscillation and make route selection easier.

## 8. Security Considerations

TBA.

## 9. Acknowledgements

TBA.

## 10. IANA Considerations

TBA.

## 11. Normative References

### [I-D.ietf-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-ietf-cats-framework-07, 30 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-cats-framework-07>>.

### [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-06, 14 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-cats-usecases-requirements-06>>.

### [I-D.ietf-idr-bgp-generic-metric]

Sangli, S. R., Hegde, S., Das, R., Decraene, B., Wen, B., Kozak, M., Dong, J., Jalil, L., and K. Talaulikar, "Accumulated Metric in NHC attribute", Work in Progress,

Internet-Draft, draft-ietf-idr-bgp-generic-metric-00, 30 August 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-generic-metric-00>>.

[I-D.ysl-cats-metric-definition]

Yao, K., Shi, H., Li, C., Contreras, L. M., and J. Ros-Giralt, "CATS Metrics Definition", Work in Progress, Internet-Draft, draft-ysl-cats-metric-definition-03, 21 November 2024, <<https://datatracker.ietf.org/doc/html/draft-ysl-cats-metric-definition-03>>.

[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>>.

Authors' Addresses

Dongyu Yuan  
ZTE Corporation  
Nanjing  
China  
Phone: +86 13776623784  
Email: yuan.dongyu@zte.com.cn

Fenlin Zhou  
ZTE Corporation  
Nanjing  
China  
Phone: +86 15861819442  
Email: zhou.fenlin@zte.com.cn

Daniel Huang  
ZTE Corporation  
Nanjing  
China  
Phone: +86 13770311052  
Email: huang.guangping@zte.com.cn

Qiudong Chen  
ZTE Corporation  
Nanjing  
China  
Phone: +86 13813084885  
Email: chen.qiudong1@zte.com.cn

Chunning Dai  
ZTE Corporation  
Nanjing  
China  
Phone: +86 15150649854  
Email: dai.chunning@zte.com.cn