

cats
Internet-Draft
Intended status: Informational
Expires: 8 January 2026

C. Li, Ed.
Huawei Technologies
Z. Du
China Mobile
M. Boucadair, Ed.
Orange
L. M. Contreras
Telefonica
J. Drake
Independent
7 July 2025

A Framework for Computing-Aware Traffic Steering (CATS)
draft-ietf-cats-framework-11

Abstract

This document describes a framework for Computing-Aware Traffic Steering (CATS). Specifically, the document identifies a set of CATS components, describes their interactions, and provides illustrative workflows of the control and data planes.

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 8 January 2026.

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	3
2. Terminology	4
3. CATS Framework and Components	7
3.1. Assumptions	7
3.2. CATS Identifiers	7
3.3. Framework Overview	8
3.4. CATS Functional Components	9
3.4.1. Service Sites, Services Instances, and Service Contact Instances	10
3.4.2. CATS Service Metric Agent (C-SMA)	11
3.4.3. CATS Network Metric Agent (C-NMA)	11
3.4.4. CATS Path Selector (C-PS)	12
3.4.5. CATS Traffic Classifier (C-TC)	12
3.4.6. CATS-Forwarders	13
3.4.7. Underlay Infrastructure	13
4. CATS Framework Workflow	13
4.1. Service Announcement	14
4.2. Metrics Distribution	14
4.3. Service Access Processing	19
4.4. Service Contact Instance Affinity	20
5. Operational and Manageability Considerations	21
5.1. Provisioning of CATS Components	21
5.2. Deployment Considerations	22
5.3. Implementation Consideration on Using CATS Metrics	23
5.4. Verify Correct Operations	24
5.5. Impact on Network Operations	24
6. Security Considerations	24
7. Privacy Considerations	25
8. IANA Considerations	25
9. Informative References	25
Appendix A. Acknowledgements	28
Contributors	28
Authors' Addresses	29

1. Introduction

Computing service architectures have evolved from single service site to multiple, sometimes collaborative, service sites to address various issues such as long response times or suboptimal utilization of service and network resources (e.g., resource under-utilization or exhaustion).

The underlying networking infrastructures that include computing resources usually provide relatively static service dispatching, e.g., the selection of the service instances for a request. In such infrastructures, service-specific traffic is often directed to the closest service site from a routing perspective without considering the actual network state (e.g., traffic congestion conditions) or the service site state.

As described in [I-D.ietf-cats-usecases-requirements], traffic steering that takes into account computing resource metrics would benefit several services, including latency-sensitive services such as immersive services that rely upon the use of augmented reality or virtual reality (AR/VR) techniques. This document provides an architectural framework that aims at facilitating the making of compute- and network-aware traffic steering decisions in networking environments where computing service resources are deployed.

The Computing-Aware Traffic Steering (CATS) framework assumes that there might be multiple service instances that are providing one given service, which are running in one or more service sites. Each of these service instances can be accessed via a service contact instance, which is a client-facing service function instance. A single service site may host one or multiple service contact instances. A single service site may have limited computing resources available at a given time, whereas the various service sites may experience different resource availability issues over time. Therefore, steering traffic among different service sites can address the issues of lacking resources in a specific service site.

Steering in CATS is about selecting the appropriate service contact instance that will service a request according to a set of network and computing metrics. This selection may not necessarily reveal the actual service instance that will be invoked, e.g., in hierarchical or recursive contexts. Therefore, the metrics of the service contact instance may be the aggregated metrics from multiple service instances.

The CATS framework is an overlay framework for the selection of the suitable service contact instance(s) from a set of candidates. The exact characterization of 'suitable' is determined by a combination of networking and computing metrics.

Furthermore, this document describes a workflow of the main CATS procedures in Section 4, which are executed in both the control and data planes.

2. Terminology

This document makes use of the following terms:

Client: An endpoint that is connected to a service provider network.

Flow: A flow is identified by the 5-tuple transport coordinates (source address and destination address, source and destination port numbers, and protocol).

Computing-Aware Traffic Steering (CATS): A traffic engineering approach [RFC9522] that takes into account the dynamic nature of computing resources (e.g., compute and storage) and network state to optimize service-specific traffic forwarding towards a given service contact instance. Various relevant metrics may be used to enable such computing-aware traffic steering policies.

Metric: A piece of information that provides suitable input to a selection mechanism to determine a CATS egress node.

Computing metrics: Computing metrics are the metrics that come specifically from the computing side of the system as distinct from other metrics that may be used in a CATS system, such as the metrics from network side. For details, a set of metrics from computing side used for CATS has been defined in [I-D.ietf-cats-metric-definition].

Service: An offering that is made available by a service provider by orchestrating a set of resources (networking, compute, storage, etc.).

Which and how these resources are used is part of the service logic which is internal to the service provider. For example, these resources may be:

- * Exposed by one or multiple processes.
- * Provided by virtual instances, physical, or a combination thereof.

- * Hosted within the same or distinct nodes.
- * Hosted within the same or multiple service sites.
- * Chained to provide a service using a variety of means.

How a service is structured is out of the scope of CATS.

The same service can be provided in many locations; each of them constitutes a service instance.

Computing Service: An offering is made available by a service provider by orchestrating a set of computing resources.

CATS Service ID (CS-ID): An identifier representing a service, which the clients use to access it. See Section 3.2.

Service instance: An instance of running resources according to a given service logic.

Many such instances can be enabled by a service provider. Instances that adhere to the same service logic provide the same service.

An instance is running in a service site. Clients' requests are serviced by one or more of these instances.

Service site: A location that hosts the resources that are required to offer a service.

A service site may be a node or a set of nodes.

A CATS-serviced site is a service site that is connected to a CATS-Forwarder.

Service contact instance: A client-facing service function instance that is responsible for receiving requests in the context of a given service.

A service contact instance can handle one or more service instances.

Steering beyond a service contact instance is hidden to both clients and CATS components.

A service request is processed according to the service logic (e.g., handle locally or solicit backend resources).

A service contact instance is reachable via at least one Egress CATS-Forwarder.

A service can be accessed via multiple service contact instances running at the same or different locations (service sites).

A service contact instance may dispatch service requests to one or more service instances (e.g., a service contact instance that behaves as a service load-balancer).

CATS Service Contact Instance ID (CSCI-ID): An identifier of a specific service contact instance. See Section 3.2.

Service request: A request to access or invoke a specific service. Such a request is steered to a service contact instance via CATS-Forwarders.

A service request is placed using service-specific protocols.

Service requests are not explicitly sent by clients to CATS-Forwarders.

CATS-Forwarder: A network entity that steers traffic specific to a service request towards a corresponding yet selected service contact instance according to provisioned forwarding decisions. These decisions are supplied by a C-PS, which may or may not be on the CATS-Forwarder.

A CATS-Forwarder may behave as Ingress or Egress CATS-Forwarder.

Ingress CATS-Forwarder: An entity that steers service-specific traffic along a CATS-computed path that leads to an Egress CATS-Forwarder that connects to the most suitable service site that host the service contact instance selected to satisfy the initial service request.

Egress CATS-Forwarder: An entity that is located at the end of a CATS-computed path and which connects to a CATS-serviced site.

CATS Path Selector (C-PS): A functional entity that selects paths towards service locations and instances and which accommodates the requirements of service requests. Such a path selection engine takes into account the service and network status information. See Section 3.4.4.

CATS Service Metric Agent (C-SMA): A functional entity that is

responsible for collecting service capabilities and status, and for reporting them to a CATS Path Selector (C-PS). See Section 3.4.2.

CATS Network Metric Agent (C-NMA): A functional entity that is responsible for collecting network capabilities and status, and for reporting them to a C-PS. See Section 3.4.3.

CATS Traffic Classifier (C-TC): A functional entity that is responsible for determining which packets belong to a traffic flow for a specific service request. It is also responsible for forwarding such packets along a C-PS computed path that leads to the relevant service contact instance. See Section 3.4.5.

3. CATS Framework and Components

3.1. Assumptions

CATS assumes that there are multiple service instances running on different service sites, which provide a given service that is represented by the same service identifier (see Section 3.2). However, CATS does not make any assumption about these instances other than they are reachable via one or multiple service contact instances.

3.2. CATS Identifiers

CATS uses the following identifiers:

CATS Service ID (CS-ID): An identifier (ID) representing a service, which the clients use to access it. Such an ID identifies all the instances of a given service, regardless of their location.

The CS-ID is independent of which service contact instance serves the service request.

Service requests are spread over the service contact instances that can accommodate them, considering the location of the initiator of the service request and the availability (in terms of resource/traffic load, for example) of the service instances resource-wise among other considerations like traffic congestion conditions.

CATS Service Contact Instance ID (CSCI-ID): An identifier of a specific service contact instance.

3.3. Framework Overview

A high-level view of the CATS framework, without expanding the functional entities in the network, is illustrated in Figure 1.

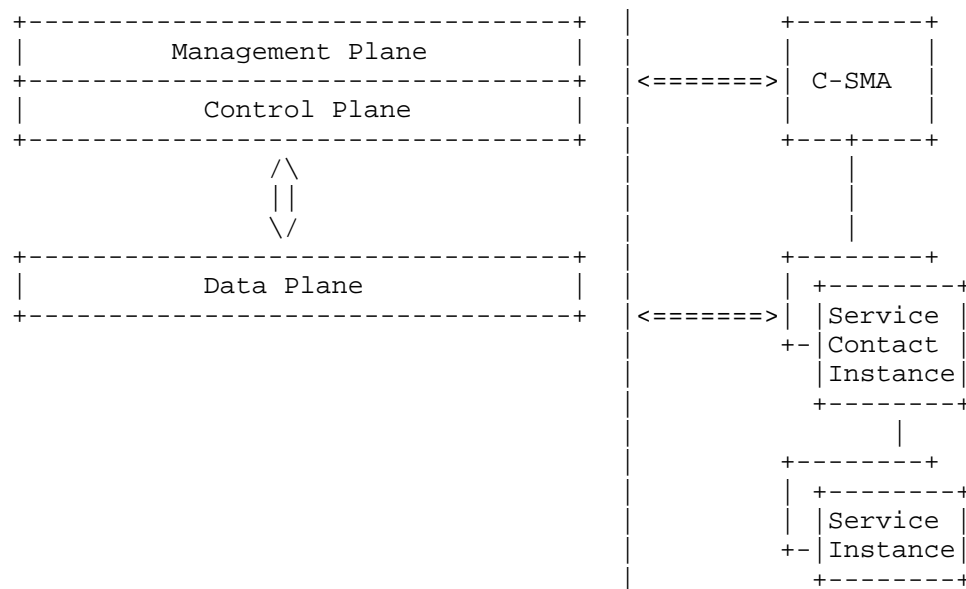


Figure 1: Main CATS Interactions

For the sake of illustration, "Service Instance" is shown as a single box in Figure 1. However, this does not imply that a service instance is hosted in a single node. Whether a service instance is realized by invoking resources within a same node or by chaining resources exposed by several nodes is deployment specific.

The following planes are defined:

- * CATS Management Plane: Responsible for monitoring, configuring, and maintaining CATS network devices.
- * CATS Control Plane: Responsible for scheduling services based on computing and network information. It is also responsible for making decisions about how packets should be forwarded by involved forwarding nodes and communicating such decisions to the CATS Data Plane for execution.
- * CATS Data Plane: Responsible for computing-aware routing, including handling packets in the data path, such as packet forwarding.

Depending on implementation and deployment, these planes may consist of several functional elements/components, and the details will be described in the following sections.

3.4. CATS Functional Components

CATS nodes make forwarding decisions for a given service request that has been received from a client according to the capabilities and status information of both service contact instances and network. The main CATS functional elements and their interactions are shown in Figure 2.

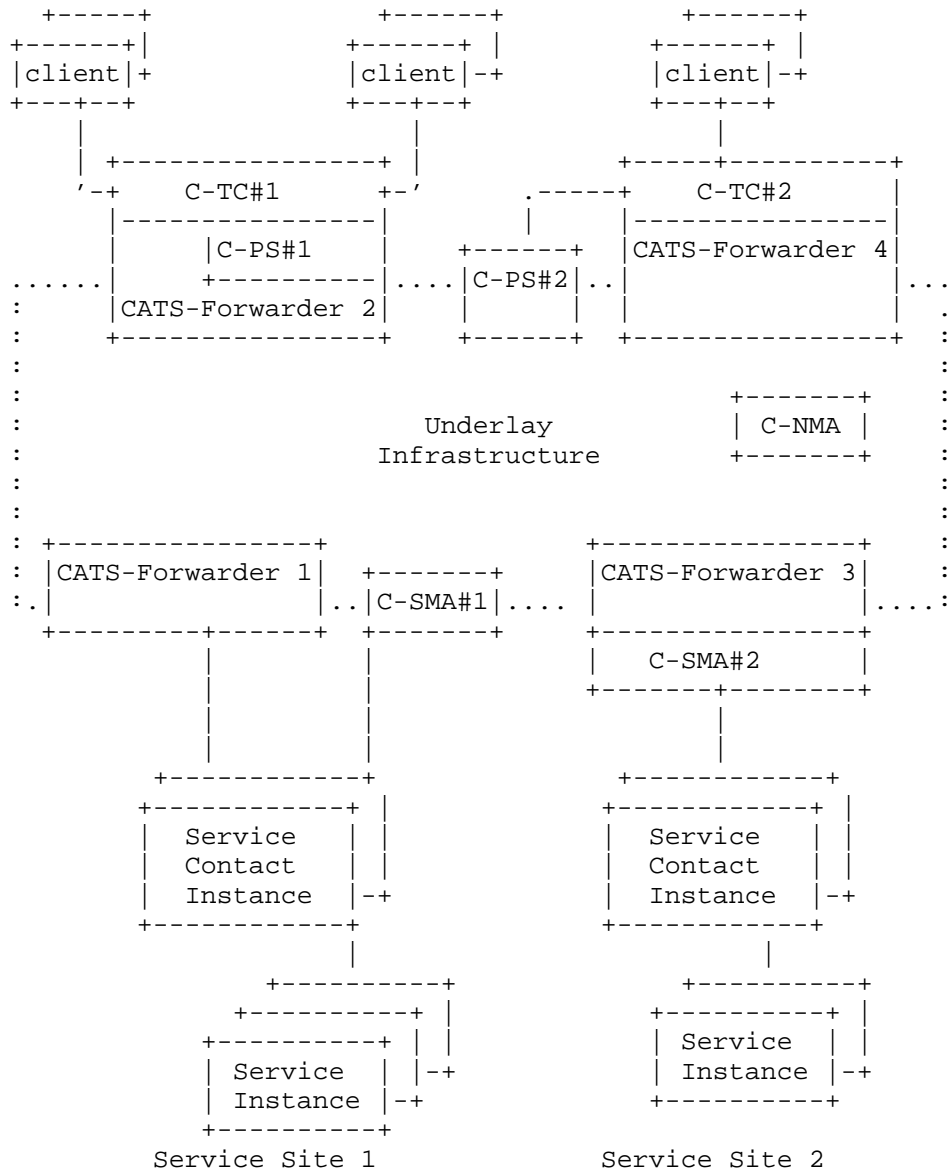


Figure 2: CATS Functional Components

3.4.1. Service Sites, Services Instances, and Service Contact Instances

Service sites are locations that host resources (including computing resources) that are required to offer a service.

A compute service (e.g., for face recognition purposes or a game server) is identified by a CATS Service Identifier (CS-ID). The CS-ID does not need to be globally unique, but must be sufficiently unique to unambiguously identify the service at all of the components of a CATS system.

A single service can be represented and accessed via several contact instances that run in same or different regions of a network.

As service instances are accessed via a service contact instance, a client will not see the service instances but only the service contact instance.

Figure 2 shows two CATS nodes ("CATS-Forwarder 1" and "CATS-Forwarder 3") that provide access to service contact instances. These nodes behave as Egress CATS-Forwarders (Section 3.4.6).

Note: "Egress" is used here in reference to the direction of the service request placement. The directionality is called to explicitly identify the exit node of the CATS infrastructure.

3.4.2. CATS Service Metric Agent (C-SMA)

The CATS Service Metric Agent (C-SMA) is a functional component that gathers information about service sites and server resources, as well as the status of the different service instances. A C-SMA may be co-located or located adjacent to a service contact instance, hosted by or adjacent to an Egress CATS-Forwarder (Section 3.4.6), etc.

Figure 2 shows one C-SMA embedded in "CATS-Forwarder 3", and another C-SMA that is adjacent to "CATS-Forwarder 1".

3.4.3. CATS Network Metric Agent (C-NMA)

The CATS Network Metric Agent (C-NMA) is a functional component that gathers information about the state of the underlay network. The C-NMAs may be implemented as standalone components or may be hosted by other components, such as CATS-Forwarders or CATS Path Selectors (C-PS) (Section 3.4.4).

C-NMA is likely to leverage existing techniques (e.g., [RFC7471], [RFC8570], and [RFC8571]).

Figure 2 shows a single, standalone C-NMA within the underlay network. There may be one or more C-NMAs for an underlay network.

3.4.4. CATS Path Selector (C-PS)

The C-SMAs and C-NMAs share the collected information with CATS Path Selectors (C-PSes) that use such information to select the Egress CATS-Forwarders (and potentially the service contact instances) where to forward traffic for a given service request. C-PSes also determine the best paths (possibly using tunnels) to forward traffic, according to various criteria that include network state and traffic congestion conditions. The collected information is encoded into one or more metrics that feed the C-PS path selection logic. Such information also includes CS-ID and possibly CSCI-IDs.

There might be one or more C-PSes used to select CATS paths in a CATS infrastructure.

A C-PS can be integrated into CATS-Forwarders (e.g., "C-PS#1" in Figure 2) or may be deployed as a standalone component (e.g., "C-PS#2" in Figure 2). Generally, a standalone C-PS can be a functional component of a centralized controller (e.g., a Path Computation Element (PCE) [RFC4655]).

Refer to Section 4.2 for a discussion on metric distribution (including, interaction with routing protocols).

3.4.5. CATS Traffic Classifier (C-TC)

The CATS Traffic Classifier (C-TC) is a functional component that is responsible for associating incoming packets from clients with service requests. CATS classifiers also ensure that packets that are bound to a specific service contact instance are all forwarded towards that same service contact instance, as instructed by a C-PS. To that aim, a C-TC uses CS-IDs (or their resolution of CS-ID to network locators) to classify service requests. Refer to Section 5.1 for more details about provisioning of classification rules.

Note that CS-IDs may be carried in packets if mechanisms such as TLS Server Name Indication extension (SNI) (Section 3 of [RFC6066]) are used.

CATS classifiers are typically hosted in CATS-Forwarders.

3.4.6. CATS-Forwarders

Ingress CATS-Forwarder are responsible for steering service-specific traffic along a CATS-computed path that leads to an Egress CATS-Forwarder. Egress CATS-Forwarders are the endpoints that behave as an egress for service requests that are forwarded over a CATS infrastructure. A service site that hosts service instances may be connected to one or more Egress CATS-Forwarders (e.g., multi-homing design). If a C-PS has selected a specific service contact instance and the C-TC has marked the traffic with the CSCI-ID related information, the Egress CATS-Forwarder then forwards traffic to the relevant service contact instance accordingly. In some cases, the choice of the service contact instance may be left open to the Egress CATS-Forwarder (i.e., traffic is marked only with the CS-ID). In such cases, the Egress CATS-Forwarder selects a service contact instance using its knowledge of service and network capabilities as well as the current load as observed by the CATS-Forwarder, among other considerations. In the absence of an explicit policy, an Egress CATS-Forwarder must make sure to forward all packets that pertain to a given service request towards the same service contact instance.

Note that, depending on the design considerations and service requirements, per-service contact instance computing-related metrics or aggregated per-site computing related metrics (and a combination thereof) can be used by a C-PS. Using aggregated per-site computing related metrics appears as a preferred option scalability-wise, but relies on Egress CATS-Forwarders that connect to various service contact instances to select the proper service contact instance. An Egress CATS-Forwarder may choose to aggregate the metrics from different sites as well. In this case, the Egress CATS-Forwarder will choose the best site by itself when the packets arrive at it.

3.4.7. Underlay Infrastructure

The "underlay infrastructure" in Figure 2 indicates an IP and/or MPLS network that is not necessarily CATS-aware. The CATS paths that are computed by a C-PS will be distributed among the CATS-Forwarders (Section 3.4.6), and will not affect the underlay nodes. Underlay nodes are typically P routers (Section 5.3.1 of [RFC4026]).

4. CATS Framework Workflow

The following subsections provide an overview of a typical CATS workflow. In order to enable CATS in a given domain, some provisioning is needed; see more details in Section 5.1. Section 5.2 describes several deployment options (distributed, centralized, and hybrid model) to accommodate a variety of contexts.

4.1. Service Announcement

A service is associated by the service provider with a unique identifier called a CS-ID. A CS-ID may be a network identifier, such as an IP address. The mapping of CS-IDs to network identifiers may be learned through a name resolution service (e.g., DNS [RFC1034]). Note that CATS framework does not assume nor preclude any specific name resolution service.

4.2. Metrics Distribution

As described in Section 3.4, a C-SMA collects both computing-related capabilities and metrics, and associates them with a CS-ID that identifies the service. The C-SMA may aggregate the metrics for multiple service contact instances, or maintain them separately or both.

The C-SMA then advertises CS-IDs along with metrics to related C-PSes in the network. Depending on deployment choice, CS-IDs with metrics may be distributed in different ways.

For example, in a distributed model, CS-IDs with metrics can be distributed from the C-SMA to an Egress CATS-Forwarder firstly and then be redistributed by the Egress CATS-Forwarder to related C-PSes that are integrated into Ingress CATS-Forwarders.

In the centralized model, CS-IDs with metrics can be distributed from the C-SMA to a centralized control plane, for instance, a standalone C-PS.

In the hybrid model, the metrics can be distributed to C-PSes in combination of distributed and centralized ways. The specific combination of metric distribution is an implementation choice, which is determined by the requirements of specific services.

The Computing metrics include computing-related metrics and potentially other service-specific metrics like the number of end-users who access the service contact instance at any given time, etc.

Computing metrics may change very frequently (see [I-D.ietf-cats-usecases-requirements] for a discussion). How frequently such information is distributed is to be determined as part of the specification of any communication protocol (including routing protocols) that may be used to distribute the information. Various options can be considered, such as (but not limited to) interval-based updates, threshold-triggered updates, policy-based updates, or using normalized metrics [I-D.ietf-cats-metric-definition]. Regarding using the normalized

metrics, Section 5.3 provides some suggestions to implement computing metrics normalization and aggregation in CATS framework, to reduce the impact of the routing system.

Additionally, the C-NMA collects network-related capabilities and metrics. These may be collected and distributed by existing measurement protocols and/or routing protocols, although extensions to such protocols may be required to carry additional information (e.g., link latency). The C-NMA distributes the network metrics to the C-PSes so that they can use the combination of service and network metrics to determine the best Egress CATS-Forwarder to provide access to a service contact instance and invoke the compute function required by a service request. Similar to computing-related metrics, the network-related metrics can be distributed using distributed, centralized, or hybrid schemes. This document does not describe such details since this is deployment-specific.

Network metrics may also change over time. Dynamic routing protocols may take advantage of some information or capabilities to prevent the network from being flooded with state change information (e.g., Partial Route Computation (PRC) of OSPFv3 [RFC5340]). C-NMAs should also be configured or instructed like C-SMAs to determine when and how often updates should be notified to the C-PSes.

Figure 3 shows an example of how CATS metrics can be disseminated in the distributed model. There is a client attached to the network via "CATS-Forwarder 1". There are three service contact instances of the service with CS-ID "1": two service contact instances with CSCI-IDs "1" and "2", respectively, are located at "Service Site 2" attached via "CATS-Forwarder 2"; the third service contact instance is located at "Service Site 3" attached via "CATS-Forwarder 3" and with CSCI-ID "3". There is also a second service with CS-ID "2" with only one service contact instance located at "Service Site 3".

In Figure 3, the C-SMA collocated with "CATS-Forwarder 2" distributes the computing metrics for both service contact instances (i.e., (CS-ID 1, CSCI-ID 1) and (CS-ID 1, CSCI-ID 2)). Note that this information may be aggregated into a single advertisement, but in this case, the metrics for each service contact instance are indicated separately. Similarly, the C-SMA agent located at "Service Site 3" advertises the computing metrics for the two services hosted by "Service Site 3".

The computing metric advertisements are processed by the C-PS hosted by "CATS-Forwarder 1". The C-PS also processes network metric advertisements sent by the C-NMA. All metrics are used by the C-PS to select the most relevant path that leads to the Egress CATS-Forwarder according to the initial client's service request, the

service that is requested ("CS-ID 1" or "CS-ID 2"), the state of the service contact instances as reported by the metrics, and the state of the network.

```
Service CS-ID 1, contact instance CSCI-ID 1 <computing metrics>
Service CS-ID 1, contact instance CSCI-ID 2 <computing metrics>
```

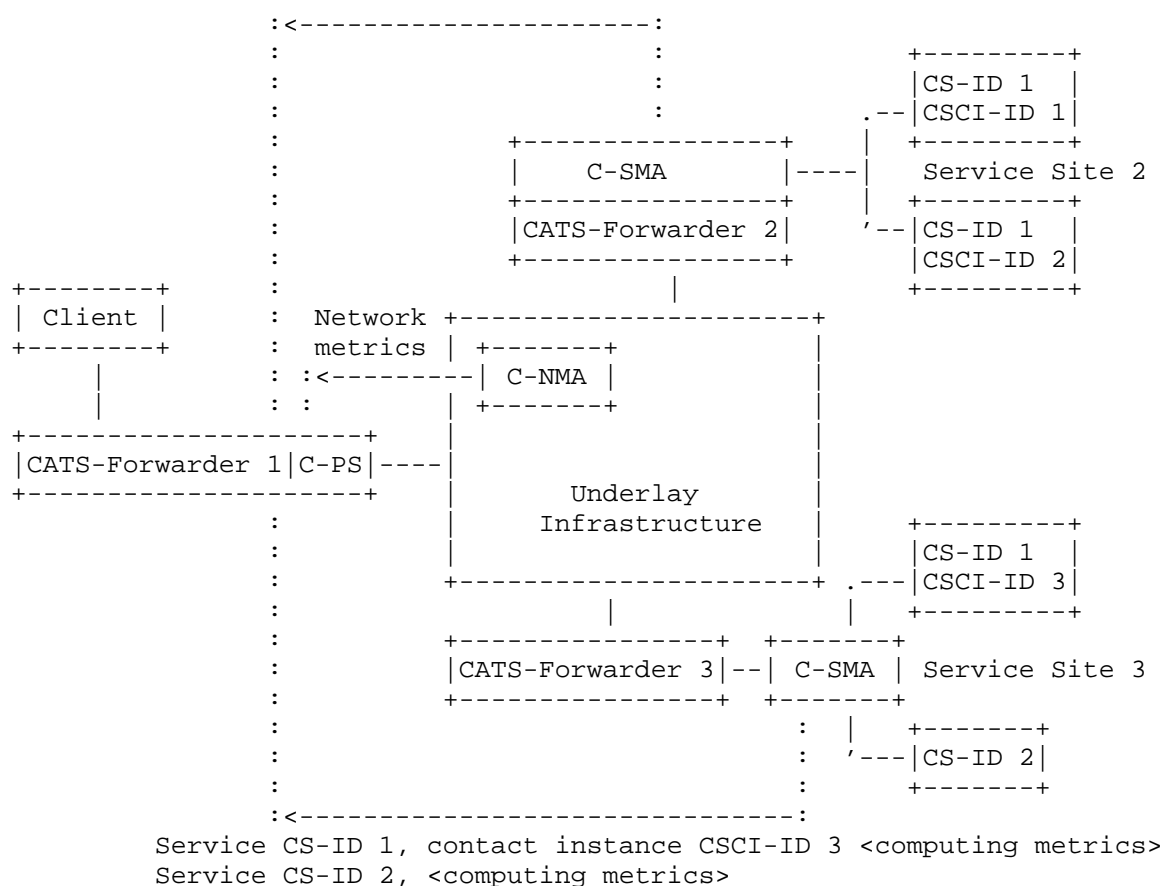


Figure 3: An Example of CATS Metric Dissemination in a Distributed Model

The example in Figure 3 mainly describes a per-instance computing-related metric distribution. In the case of distributing aggregated per-site computing-related metrics, the per-instance CSCI-ID information will not be included in the advertisement. Instead, a per-site CSCI-ID may be used in case multiple sites are connected to the Egress CATS-Forwarder to explicitly indicate the site from where the aggregated metrics come.

If the CATS framework is implemented using a centralized model, the metric can be, e.g., distributed as illustrated in Figure 4.

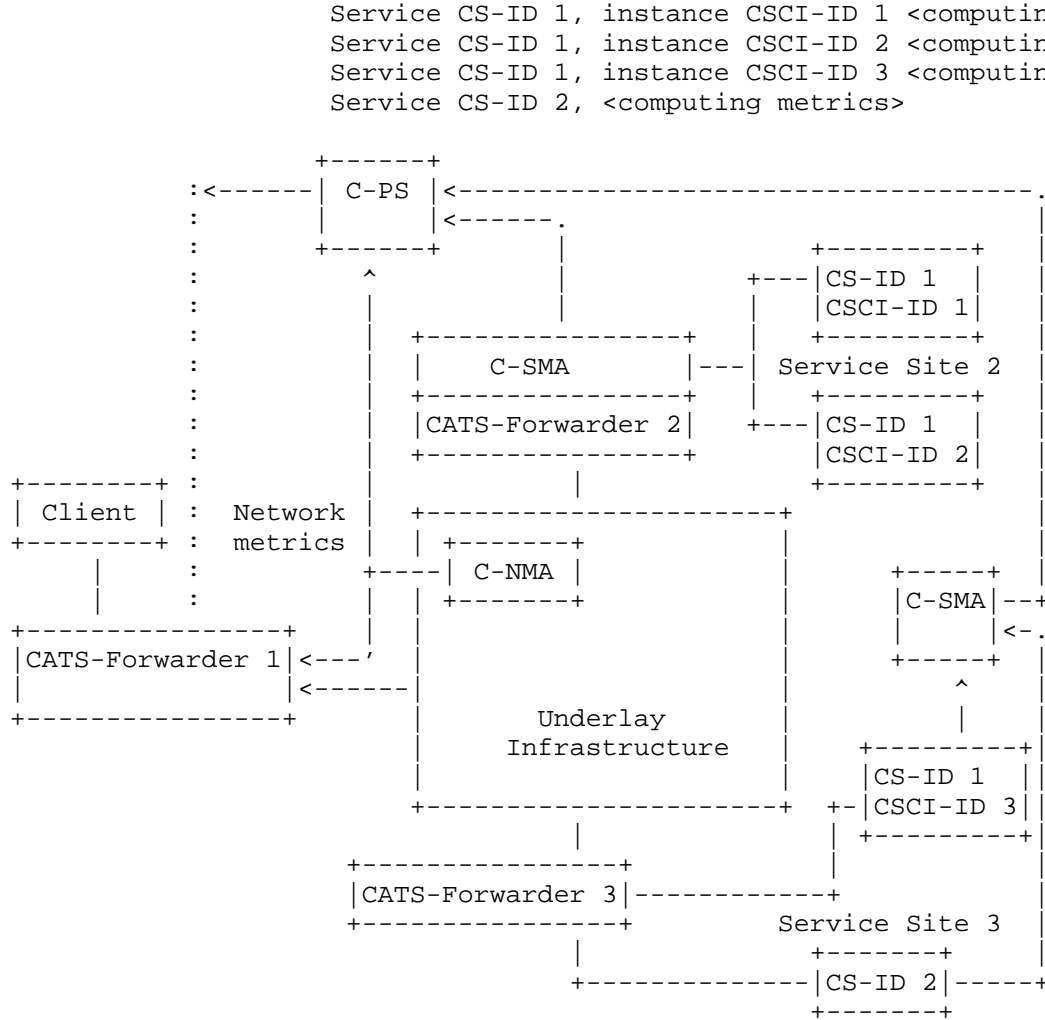


Figure 4: An Example of CATS Metric Distribution in a Centralized Model

In Figure 4, the C-SMA collocated with "CATS-Forwarder 2" distributes the computing metrics for both service contact instances (i.e., (CS-ID 1, CSCI-ID 1) and (CS-ID 1, CSCI-ID 2)) to the centralized C-PS. In this case, the C-PS is a logically centralized element deployed independently with the CATS-Forwarder 1. Similarly, the C-SMA agent located at "Service Site 3" advertises the computing metrics for the

two services hosted by "Service Site 3" to the centralized C-PS as well. Furthermore, the C-PS receives the network metrics sent from the C-NMA. All metrics are used by the C-PS to select the most relevant path that leads to the Egress CATS-Forwarder. The selected paths will be sent from the C-PS to CATS-Forwarder 1 to indicate traffic steering.

If the CATS framework is implemented using an hybrid model, the metric can be distributed, e.g., as illustrated in the Figure 5. For example, the metrics 1,2,3 associated with the CS-ID1 are collected by the centralized C-PS, and the metrics 4 and 5 are distributed via distributed protocols to the ingress CATS-Forwarder directly. For a service with CS-ID2, all the metrics are collected by the centralized C-PS. The CATS-computed path result will be distributed to the Ingress CATS-Forwarders from the C-PS by considering both the metrics from the C-SMA and C-NMA. Furthermore, the Ingress CATS-Forwarder may also have some ability to compute the path for the subsequent service accessing packets.

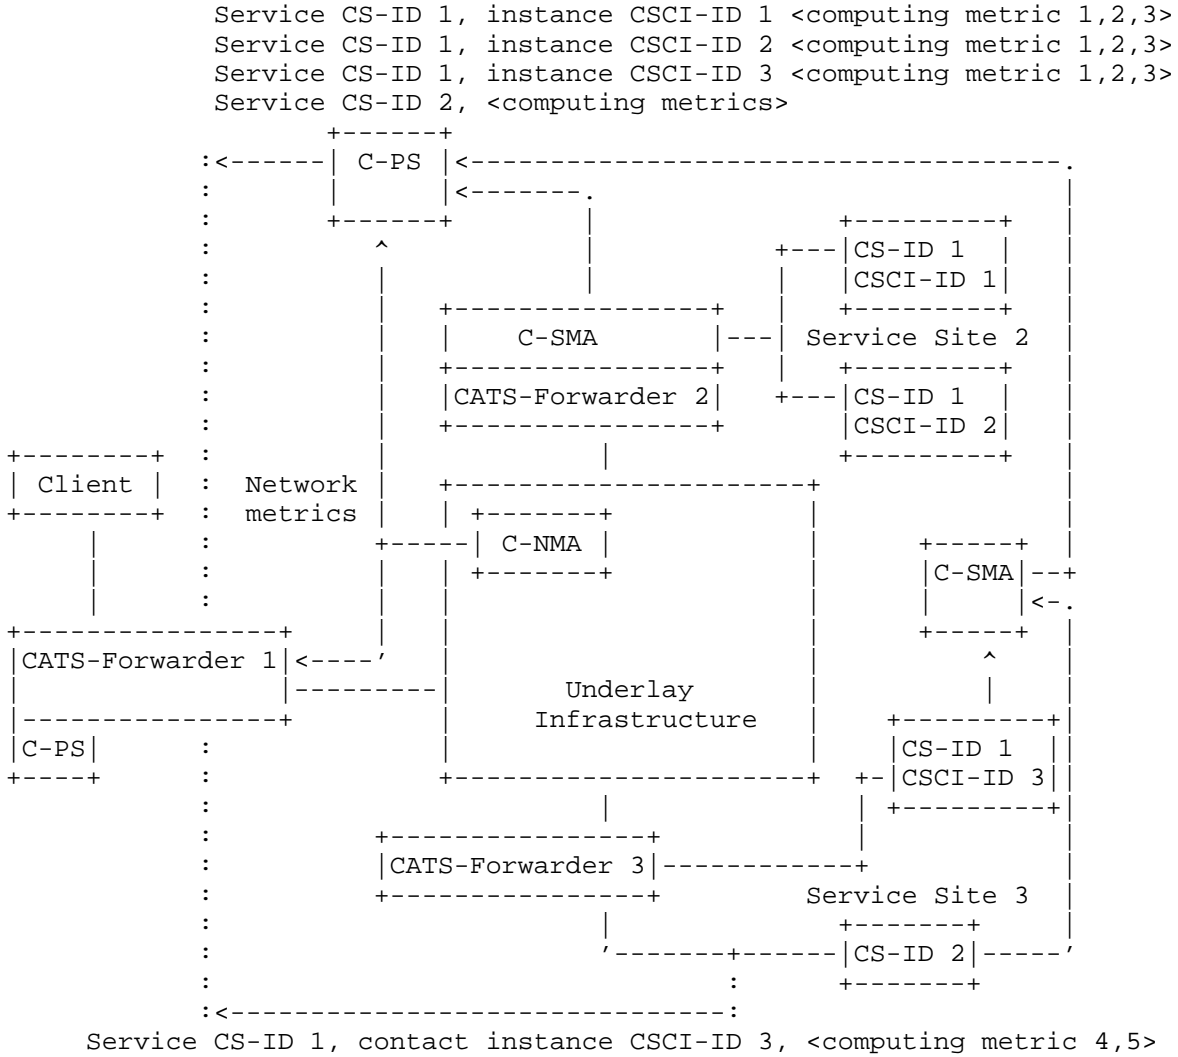


Figure 5: An Example of CATS Metric Distribution in Hybrid Model

4.3. Service Access Processing

A C-PS selects paths that lead to Egress CATS-Forwarders according to both service and network metrics that were advertised. A C-PS may be collocated with an Ingress CATS-Forwarder (as shown in Figure 3) or logically centralized (in a centralized model or hybrid model).

This document does not specify any specific algorithm for path selection purposes to be supported by C-PSes in order to not constrain the CATS framework to one possible selection only. Instead, it is expected that a service request or local policy may feed the C-PS with appropriate information on that selection logic that takes the suitable metric information as input and the selected service contact instance as output. Such "appropriate information" may be utilized to differentiate selection mechanisms to enable service-specific selections.

In the example shown in Figure 3, the client sends a service access via the network through the "CATS-Forwarder 1", which is an Ingress CATS-Forwarder. Note that, a service access may consist of one or more service packets (e.g., Session Initiation Protocol (SIP) [RFC3261], HTTP [RFC9112], IPv6 [RFC8200], SRv6 [RFC8754] or Real-Time Streaming Protocol (RTSP) [RFC7826]) that carry the CS-ID and potential parameters. The Ingress CATS-Forwarder classifies the packets using the information provided by the CATS classifier (C-TC). When a matching classification entry is found for the packets, the Ingress CATS-Forwarder encapsulates and forwards them to the C-PS selected Egress CATS-Forwarder. When these packets reach the Egress CATS-Forwarder, the outer header of the possible overlay encapsulation will be removed and the inner packets will be sent to the relevant service contact instance.

Note that multi-homed clients may be connected to multiple CATS infrastructures that may be operated by the same or distinct service providers. This version of the framework does not cover multihoming specifics.

4.4. Service Contact Instance Affinity

Service contact instance affinity means that packets that belong to a flow associated with a service should always be sent to the same service contact instance. Furthermore, packets of a given flow should be forwarded along the same path to avoid mis-ordering and to prevent the introduction of unpredictable latency variations. The CATS framework must ensure that service instance selection and path steering decisions remain consistent for a flow. Specifically, the same Egress CATS-Forwarder need to be solicited to forward the packets.

The affinity is configured on the C-PS when the service is deployed, or is determined at the time of newly formulated service requests.

Note that different services may have different notions of what constitutes a 'flow' and may, thus, identify a flow differently. Typically, a flow is identified by the 5-tuple transport coordinates

(source address and destination address, source and destination port numbers, and protocol). However, for instance, an RTP video stream may use different port numbers for video and audio channels: in that case, affinity may be identified as a combination of the two 5-tuple flow identifiers so that both flows are addressed to the same service contact instance.

Hence, when specifying a protocol to communicate information about service contact instance affinity, the protocol should support flexible mechanisms for identifying flows. Or, from a more general perspective, there should be a mechanism to specify and identify the set of packets that are subject to a service contact instance affinity.

More importantly, the means for identifying a flow for ensuring instance affinity should be application-independent to avoid the need for service-specific instance affinity methods. However, service contact instance affinity information may be configurable on a per-service basis. For each service, the information can include the flow/packets identification type and means, affinity timeout value, etc.

This document does not define any mechanism for defining or enforcing service contact instance affinity.

5. Operational and Manageability Considerations

5.1. Provisioning of CATS Components

Enabling CATS in a network can be done incrementally. That is, not all ingress routers need to be upgraded to support CATS.

In addition the CATS steering policies that are communicated by a C-PS to an Ingress CATS-Forwarder, some provisioning tasks are required. This includes, but not limited to:

- * Provide C-PS elements with the locators of available Ingress CATS-Forwarder. Such locators may also be discovered from the network.
- * Enable required setup to connect C-PS elements with C-NMA and C-SMA.
- * Allocate various identifiers CS-ID/CSCI-ID and bind them to specific service contact instances.
- * Provide C-PS element with the set of optimization metrics (per service) and an optimization policy.

- * Expose Encapsulation capabilities supported by CATS-Forwarders.
- * Configure specific encapsulation capabilities of CATS-Forwarders for use, including any credentials for mutual authentication between peer CATS-Forwarders.
- * Expose classification capabilities of C-TC elements.
- * Retrieve active classification table of C-TC elements.
- * Reset the classification table of C-TC elements.
- * Set the traffic counter at CATS-Forwarders to ease correlation between both Ingress and Egress CATS-Forwarders. Such correlation is needed to help identify issues induced by the underlying encapsulation.
- * Enable tools to check the correct behavior of various entities (e.g., classification rules, steering rules, and forwarding behavior)

The above task can be enabled using a variety of means (NETCONF [RFC6241], IPFIX [RFC7011], etc.). It is out of scope to discuss required CATS extension to these protocols.

5.2. Deployment Considerations

This document does not make any assumption about how the various CATS functional elements are implemented and deployed. Concretely, whether a CATS deployment follows a fully distributed design or relies upon a mix of centralized (e.g., a C-PS) and distributed CATS functions (e.g., CATS traffic classifiers) is deployment-specific and may reflect the preferences and policies of the (CATS) service provider.

For example, in a centralized design, both the computing related metrics from the C-SMAs and the network metrics are collected by a (logically) centralized path computation logic (e.g., a PCE). In this case, the CATS computation logic may process incoming service requests to compute paths to service contact instances. More generally, the paths might be computed before the service request comes. Based on the metrics and computed paths, the C-PS can select the most appropriate path and then synchronize with CATS traffic classifiers (C-TCs).

According to the method of distributing and collecting the computing related metrics, three deployment models can be considered for the deployment of the CATS framework:

- * ***Distributed model***: Computing metrics are distributed among network devices directly using distributed protocols without interactions with a centralized control plane. Service scheduling function is performed by the CATS-Forwarders in the distribution model, Therefore, the C-PS is integrated into an Ingress CATS-Forwarder.
- * ***Centralized model***: Computing metrics are collected by a centralized control plane, and then the centralized control plane computes the forwarding path for service requests and syncs up with the Ingress CATS-Forwarder. In this model, C-PS is implemented in the centralized control plane.
- * ***Hybrid model***: Is a combination of distribution and centralized models.

A part of computing metrics are distributed among involved network devices, and others may be collected by a centralized control plane. For example, some static information (e.g., capabilities information) can be distributed among network devices since they are quite stable (change infrequently). Frequent changing information (e.g., resource utilization) can be collected by a centralized control plane to avoid frequent flooding in the distributed control plane. Service scheduling function can be performed by a centralized control plane and/or the CATS-Forwarder. The entire or partial C-PS function may be implemented in the centralized control plane, depending on the specific implementation and deployment.

5.3. Implementation Consideration on Using CATS Metrics

According to the metric definition in [I-D.ietf-cats-metric-definition], computing metrics need to be normalized and/or aggregated in order to low down the scalability impact of the existing route system while providing sufficient detail for effective decision-making.

Depeding on the resources and processing capabilities of CATS components, the normalization and aggregation functions can be located in different CATS components. The suggested solution is to implement the normalization and aggregation functions located away from the decision maker, CATS Path Selector (C-PS), especially when C-PS is co-located with CATS-Forwarders. With this in mind, the normalization and aggregation functions of CATS metrics can be placed at Service contact instance or CATS Service Metric Agent (C-SMA).

When the C-SMA is co-located with CATS-Forwarders where there is limited resource for processing, the placement of normalization functions in the C-SMA may bring too much overhead and may influence the routing efficiency. Therefore, this document suggests to implement the normalization function at the service contact instance. Regarding the aggregation functions, it can be implemented in the C-SMA, or the service contact instance.

In the case of service contact instances and C-SMAs are provided by different vendors, it is needed to use the same common normalization function and aggregation functions, so that the service contact instance selection result can be fair among all the service contact instances.

5.4. Verify Correct Operations

CATS may be implemented by extending some existing control plane protocols, such as BGP or PCEP. A CATS implementation must log error events for better network management and operation. Means to assess the reachability and trace CATS paths should be supported.

5.5. Impact on Network Operations

Computing metrics are collected and distributed in CATS. A new function is needed to be deployed to manage the cooperation between network elements and computing elements. For example, this function may be provided by an orchestrator connecting with C-SMA and C-NMA. This might bring more complexity of the network management, especially if this function is not leveraged for other purposes beyond CATS.

6. Security Considerations

The computing resource information changes over time very frequently, especially with the creation and termination of service instances. When such information is carried in a routing protocol, too many updates may affect network stability. This issue could be exploited by an attacker (e.g., by spawning and deleting service contact instances very rapidly). CATS solutions must support guards against such misbehaviors. For example, these solutions should support aggregation techniques, dampening mechanisms, and threshold-triggered distribution updates.

The information distributed by the C-SMAs and C-NMAs may be sensitive. Such information could indeed disclose intelligence about the network and the location of compute resources hosted in service sites. This information may be used by an attacker to identify weak spots in an operator's network. Furthermore, such information may be

modified by an attacker resulting in disrupted service delivery for the clients, even including misdirection of traffic to an attacker's service implementation. CATS solutions must support authentication and integrity-protection mechanisms between C-SMAs/C-NMAs and C-PSes, and between C-PSes and Ingress CATS-Forwarders. Also, C-SMA agents need to support a mechanism to authenticate the services for which they provide information to C-PS computation logics, among other CATS functions.

7. Privacy Considerations

CATS solutions must support preventing on-path nodes in the underlay infrastructure to fingerprint and track clients (e.g., determine which client accesses which service). More generally, personal data must not be exposed to external parties by CATS beyond what is carried in the packet that was originally issued by the client.

In some cases, the service will need to know about applications, clients, and even user identity. This information is sensitive and should be encrypted. To prevent the information leaking between CATS components, the C-PS computed path information should be encrypted in distribution. The specific encryption method may be applied at the network layer, transport layer, or at the application/protocol level depending on the implementation, so this is out of the scope of this document.

For more discussion about privacy, refer to [RFC6462] and [RFC6973].

8. IANA Considerations

This document makes no requests for IANA action.

9. Informative References

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

Yao, K., Li, C., Contreras, L. M., Ros-Giralt, J., and H. Shi, "CATS Metrics Definition", Work in Progress, Internet-Draft, draft-ietf-cats-metric-definition-03, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-cats-metric-definition-03>>.

[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-07, 10 June 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-cats-usecases-requirements-07>>.

[I-D.yao-cats-awareness-architecture]

Yao, H., wang, X., Li, Z., Huang, D., and C. Lin,
"Computing and Network Information Awareness (CNIA) system
architecture for CATS", Work in Progress, Internet-Draft,
draft-yao-cats-awareness-architecture-02, 22 October 2023,
<<https://datatracker.ietf.org/doc/html/draft-yao-cats-awareness-architecture-02>>.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities",
STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987,
<<https://www.rfc-editor.org/rfc/rfc1034>>.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston,
A., Peterson, J., Sparks, R., Handley, M., and E.
Schooler, "SIP: Session Initiation Protocol", RFC 3261,
DOI 10.17487/RFC3261, June 2002,
<<https://www.rfc-editor.org/rfc/rfc3261>>.

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual
Private Network (VPN) Terminology", RFC 4026,
DOI 10.17487/RFC4026, March 2005,
<<https://www.rfc-editor.org/rfc/rfc4026>>.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path
Computation Element (PCE)-Based Architecture", RFC 4655,
DOI 10.17487/RFC4655, August 2006,
<<https://www.rfc-editor.org/rfc/rfc4655>>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF
for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008,
<<https://www.rfc-editor.org/rfc/rfc5340>>.

[RFC6066] Eastlake 3rd, D., "Transport Layer Security (TLS)
Extensions: Extension Definitions", RFC 6066,
DOI 10.17487/RFC6066, January 2011,
<<https://www.rfc-editor.org/rfc/rfc6066>>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed.,
and A. Bierman, Ed., "Network Configuration Protocol
(NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011,
<<https://www.rfc-editor.org/rfc/rfc6241>>.

[RFC6462] Cooper, A., "Report from the Internet Privacy Workshop",
RFC 6462, DOI 10.17487/RFC6462, January 2012,
<<https://www.rfc-editor.org/rfc/rfc6462>>.

- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/rfc/rfc6973>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/rfc/rfc7011>>.
- [RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", RFC 7471, DOI 10.17487/RFC7471, March 2015, <<https://www.rfc-editor.org/rfc/rfc7471>>.
- [RFC7826] Schulzrinne, H., Rao, A., Lanphier, R., Westerlund, M., and M. Stiemerling, Ed., "Real-Time Streaming Protocol Version 2.0", RFC 7826, DOI 10.17487/RFC7826, December 2016, <<https://www.rfc-editor.org/rfc/rfc7826>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/rfc/rfc8200>>.
- [RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March 2019, <<https://www.rfc-editor.org/rfc/rfc8570>>.
- [RFC8571] Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", RFC 8571, DOI 10.17487/RFC8571, March 2019, <<https://www.rfc-editor.org/rfc/rfc8571>>.
- [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/rfc/rfc8754>>.
- [RFC9112] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP/1.1", STD 99, RFC 9112, DOI 10.17487/RFC9112, June 2022, <<https://www.rfc-editor.org/rfc/rfc9112>>.

[RFC9522] Farrel, A., Ed., "Overview and Principles of Internet Traffic Engineering", RFC 9522, DOI 10.17487/RFC9522, January 2024, <<https://www.rfc-editor.org/rfc/rfc9522>>.

Appendix A. Acknowledgements

The authors would like to thank Joel Halpern, John Scudder, Dino Farinacci, Adrian Farrel, Cullen Jennings, Linda Dunbar, Jeffrey Zhang, Peng Liu, Fang Gao, Aijun Wang, Cong Li, Xinxin Yi, Jari Arkko, Mingyu Wu, Haibo Wang, Xia Chen, Jianwei Mao, Guofeng Qian, Zhenbin Li, Xinyue Zhang, Weier Li, Ines Robles and Nagendra Kumar for their comments and suggestions.

Some text about various deployment models was originally documented in [I-D.yao-cats-awareness-architecture].

Contributors

Guangping Huang
ZTE
Email: huang.guangping@zte.com.cn

Gyan Mishra
Verizon Inc.
Email: hayabusagsm@gmail.com

Huijuan Yao
China Mobile
Email: yaohuijuan@chinamobile.com

Yizhou Li
Huawei Technologies
Email: liyizhou@huawei.com

Dirk Trossen
Huawei Technologies
Email: dirk.trossen@huawei.com

Luigi Iannone
Huawei Technologies
Email: luigi.iannone@huawei.com

Hang Shi
Huawei Technologies
Email: shihang9@huawei.com

Changwang Lin
New H3C Technologies
Email: linchangwang.04414@h3c.com

Xueshun Wang
CICT
Email: xswang@fiberhome.com

Xuewei Wang
Ruijie Networks
Email: wangxuewei1@ruijie.com.cn

Christian Jacquenet
Orange
Email: christian.jacquenet@orange.com

Authors' Addresses

Cheng Li (editor)
Huawei Technologies
China
Email: c.l@huawei.com

Zongpeng Du
China Mobile
China
Email: duzongpeng@chinamobile.com

Mohamed Boucadair (editor)
Orange
France
Email: mohamed.boucadair@orange.com

Luis M. Contreras
Telefonica
Spain

Email: luismiguel.contrerasmurillo@telefonica.com

John E Drake
Independent
United States of America
Email: je_drake@yahoo.com