

Distributed Mobility Management
Internet-Draft
Intended status: Informational
Expires: 2 January 2026

N. Tran
Y. Kim
Soongsil University
1 July 2025

Computing Aware Traffic Steering Consideration for Mobile User Plane
Architecture
draft-dcn-dmm-cats-mup-06

Abstract

The document [I-D.draft-ietf-dmm-mup-architecture] describes the Mobile User Plane (MUP) architecture for Distributed Mobility Management. The proposed architecture converts the user mobility session information from the control plane entity to an IPv6 dataplane routing information. When there are multiple candidate instances located at different location to serve an user request, the MUP Provider Edge (PE) might prioritize the closest service location. However, the closest routing path might not be the optimal route.

This document discusses how the mentioned MUP architecture can be leveraged to set up dataplane routing paths to the optimal service instance location with the assistance of computing-aware traffic steering capabilities. For each session request, based on the up-to-date collected computing and network information, the MUP controller can convert the session information to the optimal route.

Status of This Memo

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

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

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

This Internet-Draft will expire on 2 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	2
2. Terminology used in this draft	4
3. MUP enhancement requirements for supporting CATS	4
4. MUP enhancement considerations for supporting CATS	5
4.1. CATS-MUP Centralized Deployment case	5
4.1.1. MUP Route enhancements	6
4.1.2. Deployment architecture	7
4.1.3. New UE request underlay routing setup procedure	8
4.1.4. UE mobility handling procedure	10
4.2. CATS-MUP Distributed Deployment case	12
4.2.1. MUP Route enhancements	12
4.2.2. Deployment architecture	13
4.2.3. New UE request underlay routing setup procedure	14
4.2.4. UE mobility handling procedure	16
5. References	18
5.1. Informative References	18
Authors' Addresses	19

1. Introduction

The document [I-D.draft-ietf-dmm-mup-architecture] describes the Mobile User Plane architecture for Distributed Mobility Management. This architecture is composed of a MUP controller and multiple MUP PEs. When applying the MUP architecture in 5G network, the MUP PEs accommodate the N3 RAN network as Interwork Segment or the N6 DN network as Direct Segment. The MUP PEs advertise the Interwork and Discovery Segment dataplane network reachability (e.g. Segment Routing IPv6 segment identifier (SRv6 SID)) to the MUP network via the interwork and direct segment Discovery routes. Meanwhile, the MUP controller transformed the received user mobility session information to the corresponding interwork and direct segment information. Then, it advertises the transformed information to MUP

PEs via Session Transformed routes. The MUP PEs use the matching Discovery routes to resolve the Session Transformed routes and forward the packet through the MUP SRv6 network.

This document discusses the optimal route configuration problem when applying the mentioned MUP architecture in a network scenario where an user request can be served by multiple computing instances of the same service located at different locations. The closest geographical service location to users might not be the optimal service instance's location as pointed out in the problem statement document of IETF Computing-Aware Traffic Steering (CATS) working group [I-D.draft-ietf-cats-usecases-requirements]. In this scenario, an optimal service instance location can be decided at the mobile control plane or data plane.

In the control plane case, it is possible to use an Application Function (AF) to determine the optimal service instance and influence the 5G control plane to select the DN corresponding to the chosen instance. The MUP-C only needs to transform the optimal DN information in the session information into the corresponding route. Meanwhile, in the data plane approach, the MUP-C should decide the optimal service instance location by itself and transform the unoptimal session information into the optimal route based on its decision. The data plane approach can avoid additional signalling procedure at the control plane of the other approach. It also supports IP Routing paradigm benefit of SRv6 mobile user plane as mentioned in the edge computing use case of the document [I-D.draft-ietf-dmm-srv6mob-arch].

Therefore, a solution to integrate CATS capabilities into the mentioned MUP architecture is presented in this document. By considering service computing and network information of all candidate service instances, the MUP controller can convert the session information into the optimal dataplane route.

This document is proposed to discuss a possible extension consideration of the original MUP architecture document[I-D.draft-ietf-dmm-mup-architecture]. Regarding the Distributed Mobility Management requirements described in [RFC7333], the MUP architecture can partly address the "Non-optimal routes" problem and the "Multicast considerations" requirement by integrating CATS capabilities. As described in [RFC4786], anycast is the practice of making a particular service address available in multiple locations. Anycast support could be in the scope of multicast support for distributed mobility management.

2. Terminology used in this draft

CATS-MUP-C: Computing-aware traffic steering MUP-C which integrates CATS path selection and MUP-C features.

Besides, this document uses the following terminologies which has been defined in [I-D.draft-ietf-cats-framework]

CATS: Computing-Aware Traffic Steering takes into account the dynamic nature of computing resource metrics and network state metrics to steer service traffic to a service instance.

Service: An offering that is made available by a provider by orchestrating a set of resources (networking, compute, storage, etc.). The same service can be provided in many locations; each of them constitutes a service instance.

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

Service contact instance: A client-facing service function instance that is responsible for receiving requests in the context of a given service. A single service can be represented and accessed via several contact instances that run in different regions of a network.

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

CATS Service Metric Agent (C-SMA): A functional entity that is responsible for collecting service capabilities and status, and for reporting them to a C-PS.

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

3. MUP enhancement requirements for supporting CATS

This section presents 3 enhancement points that need to be added in MUP for selecting an optimal service instance for serving an user request.

First, the MUP architecture should be capable of identifying the service and its candidate service instances. These service identifiers are well defined in CATS framework document [I-D.draft-ietf-cats-framework], CATS Service ID (CS-ID) is used to differentiate between different services. CATS Instance Selector ID (CIS-ID) is used to differentiate between different service instances of the same service.

Second, the MUP architecture should be capable of advertising service deployment information among its components. The egress MUP PE attaching to the MUP direct segment should gather the corresponding service and service instance information (CS-ID and CIS-ID) and advertise to the MUP environment. Different methods can be considered for this requirement.

Third, the MUP architecture should be capable of advertising the computing and network metrics (CATS metrics) related to the each service instance. The egress MUP PE attaching to the MUP direct segment should gather the corresponding service CATS metrics and advertise to the MUP environment. Different methods can be considered for this requirement.

This document only discusses the requirements and architecture. Different methods for service information and CATS metrics distribution to the network can be applied. BGP extension is an example approach that can be referred from related IETF documents such as [I-D.draft-lin-idr-distribute-service-metric] or [I-D.draft-ietf-idr-5g-edge-service-metadata]. Solutions for mitigating routing instability issue that might be caused by varied and frequent changed CATS metrics can also be referred in related IETF documents. For example, updating frequency can be set to at least 30 seconds according to [I-D.draft-ietf-idr-5g-edge-service-metadata] or converting low-level fine-grained metrics into a normalized one as mentioned in [I-D.draft-ietf-cats-metric-definition]. Collecting metrics by a centralized entity as described below in this document is another candidate method for reducing the metric advertisement burden in the network.

4. MUP enhancement considerations for supporting CATS

4.1. CATS-MUP Centralized Deployment case

4.1.1. MUP Route enhancements

Compared with the original route definition introduced in [I-D.draft-ietf-dmm-mup-architecture], the Direct Segment Discovery Route (DSD) and the Type 2 Session Transformed Route (T2ST) need modifications to support the centralized CATS-MUP deployment case. Another CATS Metrics Update Route (CMU) is also introduced.

The Direct Segment Discovery route advertises the reachability information of the direct segment. This route is advertised from the PEs attaching to the direct segments to the PEs attaching to the mobile network access side. In CATS scenario, the direct segment is a specific instance of a service. The service identifier CS-ID and service instance identifier CIS-ID information are required in this route. The CS-ID can be used as the direct segment BGP extended community attribute. The list below shows the information that should be included in the BGP NLRI of the DSD route in CATS-MUP centralized deployment case:

- * CS-ID
- * CIS-ID
- * Attached PE SID

The Type 2 Session Transformed Route convert the session information into dataplane routing information. This route is advertised from the CATS-MUP-C to the PEs attaching to the mobile network access side. In CATS scenario, the direct segment is a specific instance of a service. This route type includes the target service identifier CS-ID and the tunnel endpoint identifier on the mobile network core side information. The optimal service instance identifier CIS-ID determined by the CATS-MUP-C is also required in this route information. The list below shows the information that should be included in the BGP NLRI of the T2ST route in CATS-MUP centralized deployment case:

- * CS-ID
- * Optimal CIS-ID
- * Tunnel Endpoint Identifier on the core side

The CATS Metric Update route convert the session information into dataplane routing information. This route is advertised from the PEs attaching to the direct segments to the CATS-MUP-C. This route type update the CATS metrics related to the attaching service instance of each PE to the CATS-MUP-C. The list below shows the information that should be included in the BGP NLRI of the CMU route in CATS-MUP centralized deployment case:

- * CS-ID
- * CIS-ID
- * CATS metrics

4.1.2. Deployment architecture

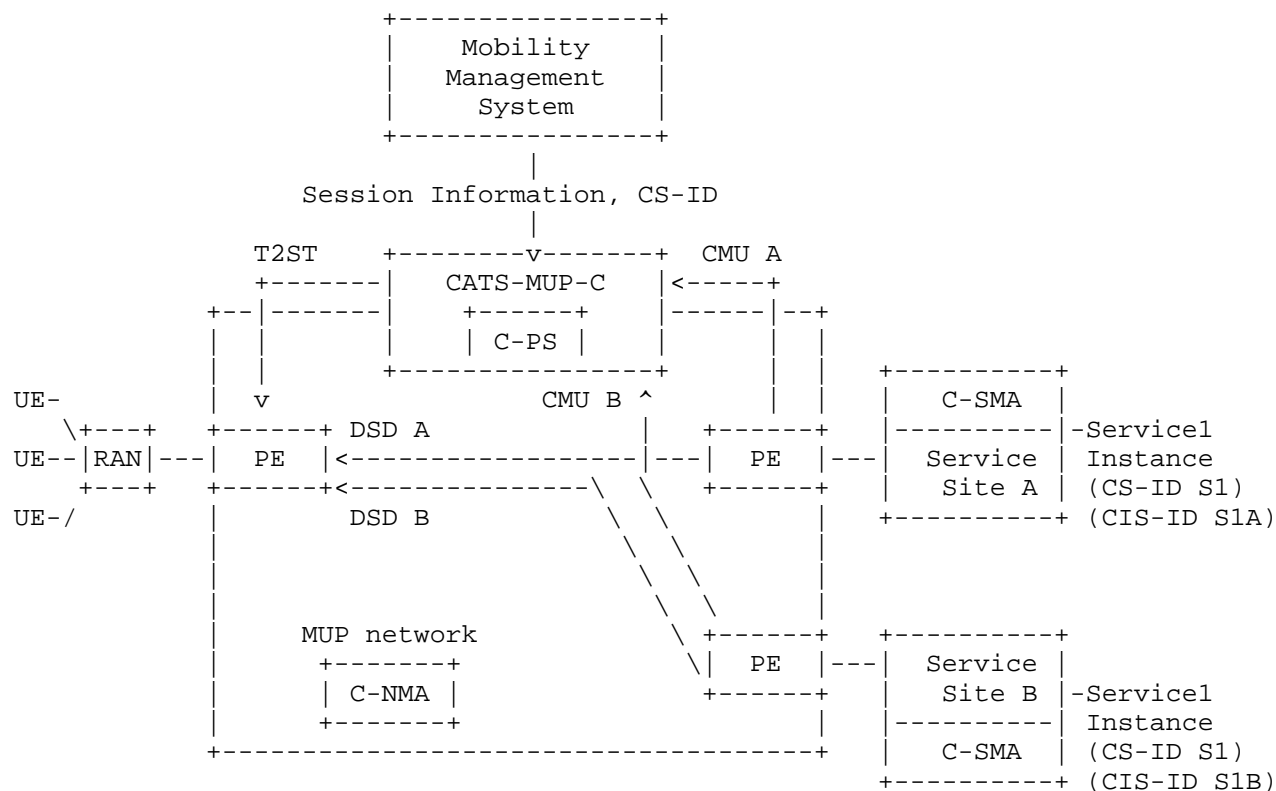


Figure 1: CATS-MUP Centralized deployment option

Figure 1 describes the CATS-MUP Centralized deployment architecture. The controller MUP-C in previous mentioned document is enhanced with CATS path selection capability and renamed to CATS-MUP-C. The Centralized deployment option has the following key features:

- * The DSD routes advertise service and service instance identity information (CS-ID and CIS-ID) from the Service Site PEs to the RAN PEs. This happens only one time when a service instance is deployed or removed at a service site
- * The CMU routes periodically advertise CATS metrics corresponding to each service instance from the Service Site PEs to the CATS-MUP-C
- * The CATS-MUP-C decides the optimal service instance for each UE session based on the metrics from the CMU routes
- * The T2ST routes convert the UE session information into the optimal underlay routing information and provide them to the RAN PEs.

4.1.3. New UE request underlay routing setup procedure

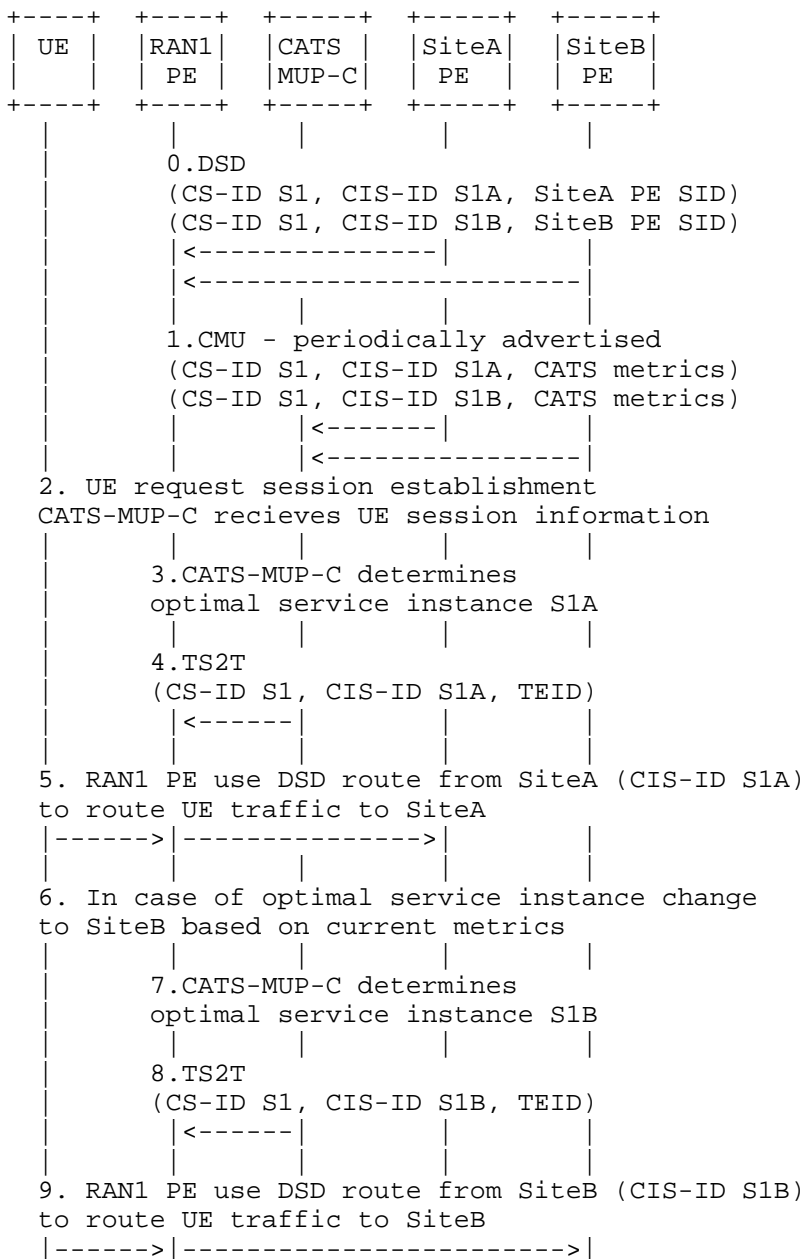
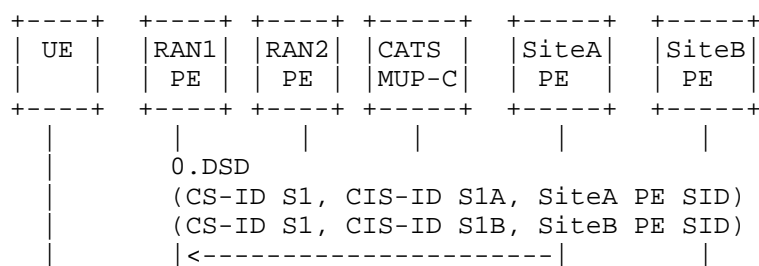


Figure 2: New UE request underlay routing setup procedure

Figure 3 describes the sequence of how the CATS-MUP Centralized deployment approach setup the underlay routing path for an UE request.

- * First, when a service instance is available at a service site, a DSD route is advertised from the corresponding PE of that site. In the example in Figure 3, DSD routes are advertised from SiteA and SiteB PEs to all of the RAN PEs. These routes announce the service and service instance identities (CS-ID, CIS-ID) running at each service site and the corresponding PE SRv6 SID
- * Then, the CATS metrics corresponding to each service instance is periodically updated from the Site PEs to the CATS-MUP-C via the CMU routes
- * When a UE request session establishment to the mobile management system, the CATS-MUP-C receives the session information which includes the requested service CS-ID and the TEID of the session. The CATS-MUP-C determines the optimal service instance for the session based on current CATS metrics.
- * Then, the CATS-MUP-C maps the session information TEID with the optimal service instance CIS-ID and advertises this information to the RAN PE via the T2ST route.
- * Upon receiving the T2ST route, the RAN PE selects the DSD route that has the same CIS-ID information with the received T2ST route (CIS-ID S1A in this example). The RAN PE uses this D2D route for steering the UE traffic to the optimal service site.
- * In case of the optimal service instance change to another site (SiteB in this example) due to current CATS metrics update, the CATS-MUP-C advertises a new T2ST route to the RAN PE, which maps the session information TEID with the new optimal service instance CIS-ID (S1B in this example). The RAN PE resolves the new T2ST route with the corresponding DSD route of the new CIS-ID.

4.1.4. UE mobility handling procedure



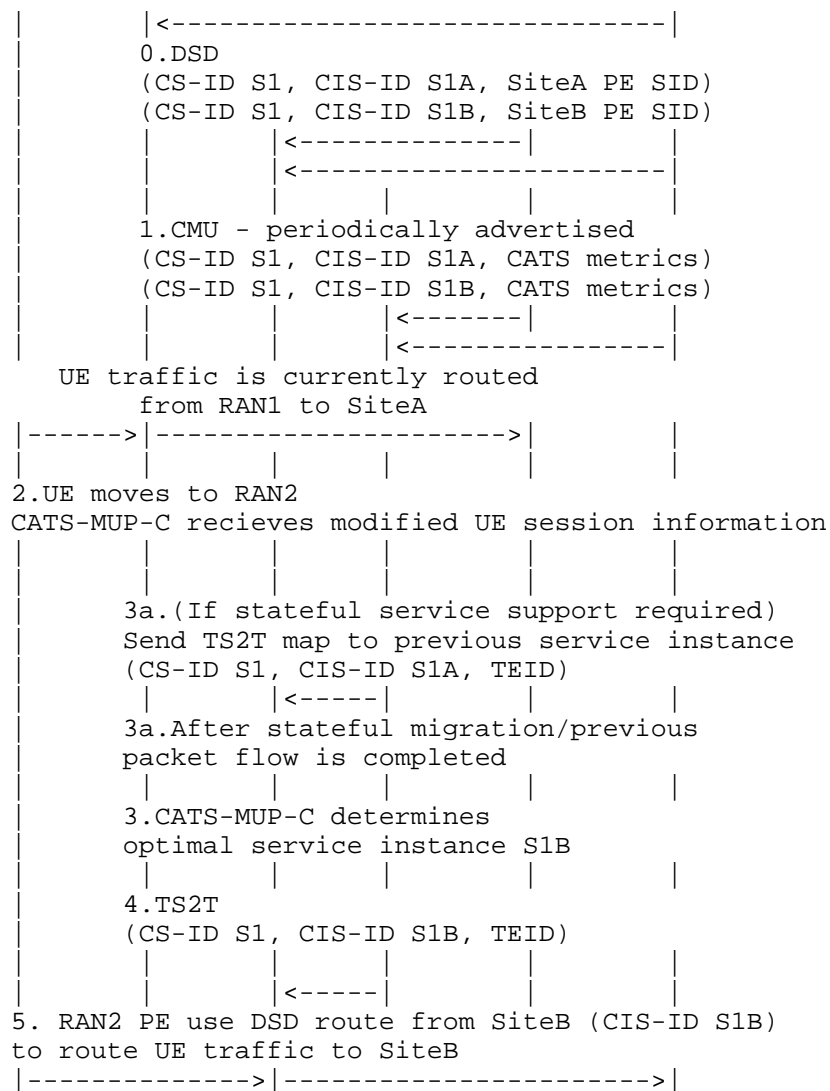


Figure 3: UE mobility handling procedure

Figure 4 describes the sequence of how the CATS-MUP Centralized deployment approach setup the underlay routing path when UE moves to a new RAN.

- * DSD routes are advertised from service site PEs to all RAN PEs. CMU routes are periodically advertised from the site PEs to the CATS-MUP-C
- * When UE moves to a different RAN, the CATS-MUP-C receives modified UE session information from the mobility management system
- * The CATS-MUP-C selects the new optimal service instance based on the current CATS metrics and sends a corresponding T2ST route to the new RAN PE that maps the session information with the new optimal service instance CIS-ID (CIS-ID S1B in this example).
- * In case of stateful/sticky service support are required, where connection to the previous service instance needs to be maintained, the CATS-MUP-C should firstly send a T2ST route that maps the session information with the previous service instance CIS-ID. Then, depends on the session continuity policy, the CATS-MUP-C can update the PE with a new T2ST route later. For example, new T2ST can be sent to PE after service packet flow to the previous instance is completed or after service state migration is completed.
- * Upon receiving the T2ST route, the RAN PE selects the DSD route that has the same CIS-ID information with the received T2ST route (CIS-ID S1B in this example). The RAN PE uses this D2D route for steering the UE traffic to the optimal service site.

4.2. CATS-MUP Distributed Deployment case

4.2.1. MUP Route enhancements

Compared with the original route definition introduced in [I-D.draft-ietf-dmm-mup-architecture], the Direct Segment Discovery Route (DSD) and the Type 2 Session Transformed Route (T2ST) need modifications to support the distributed CATS-MUP deployment case.

The Direct Segment Discovery route advertises the reachability information of the direct segment. This route is advertised from the PEs attaching to the direct segments to the PEs attaching to the mobile network access side. For the distributed CATS-MUP deployment case, in addition to the CS-ID and the CIS-ID, the CATS metrics of the corresponding service instance of the PE is also included. The CS-ID can be used as the direct segment extended community ID. The list below shows the information that should be included in the BGP NLRI of the DSD route in CATS-MUP centralized deployment case:

- * CS-ID

- * CIS-ID
- * CATS metrics
- * Attached PE SID

The Type 2 Session Transformed Route convert the session information into dataplane routing information. This route is advertised from the CATS-MUP-C to the PEs attaching to the mobile network access side. For the distributed CATS-MUP deployment case, this route type only includes the target service identifier CS-ID and the tunnel endpoint identifier on the mobile network core side information. The list below shows the information that should be included in the BGP NLRI of the T2ST route in CATS-MUP centralized deployment case:

- * CS-ID
- * Tunnel Endpoint Identifier on the core side

4.2.2. Deployment architecture

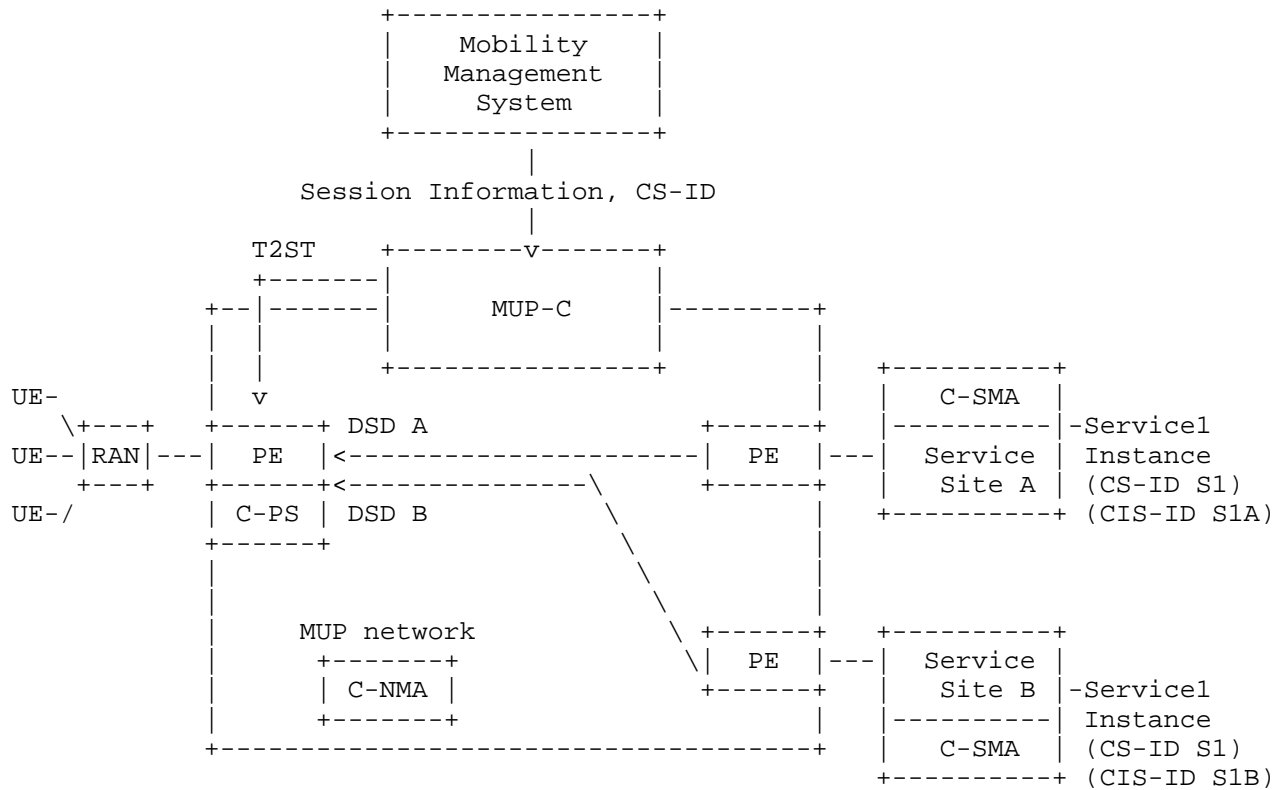


Figure 4: CATS-MUP Distributed deployment option

Figure 5 describes the CATS-MUP Distributed deployment architecture. This option has the following key features:

- * The DSD routes periodically advertise service, service instance identity information (CS-ID and CIS-ID) and CATS metrics from the Service Site PEs to the RAN PEs.
- * The T2ST routes provides the UE session information and requested service CS-ID to the RAN-PE.
- * The C-PS at the RAN-PE selects the optimal service instance based on current CATS metrics provided from the DSD routes.

4.2.3. New UE request underlay routing setup procedure

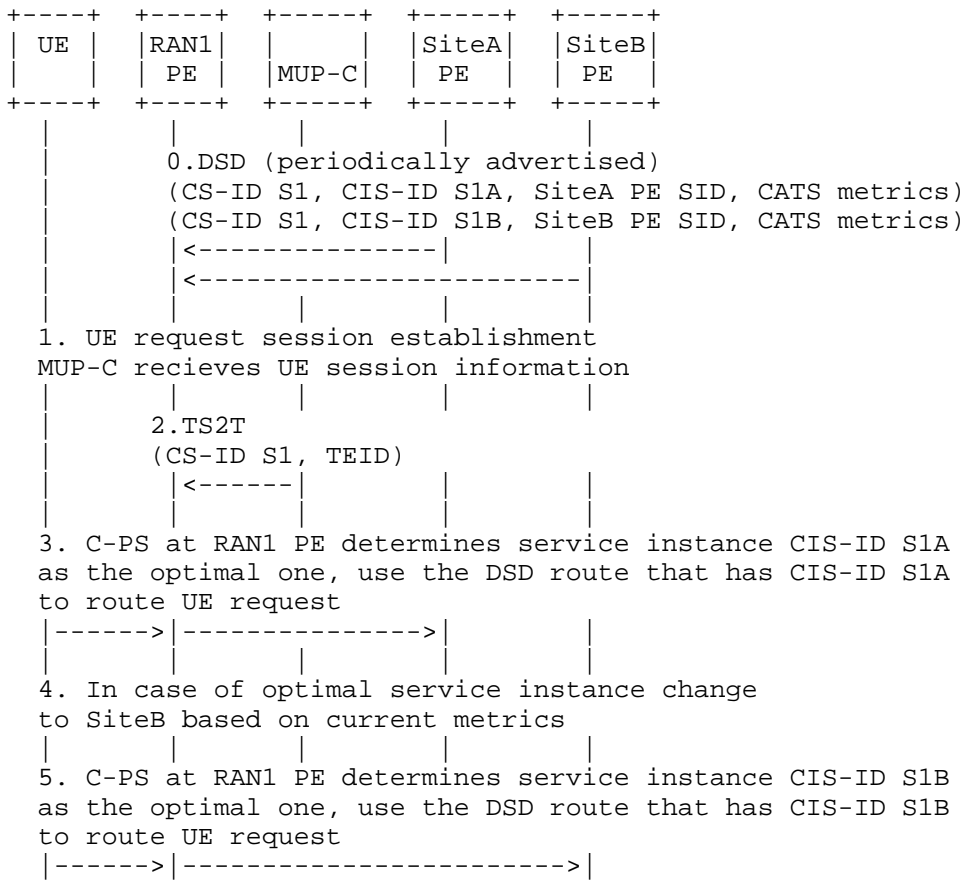


Figure 5: New UE request underlay routing setup procedure

Figure 6 describes the sequence of how the CATS-MUP Distributed deployment approach setup the underlay routing path for an UE request.

- * DSD routes are advertised from the PEs of the service sites that hosts the service instances to all of the RAN PEs. Because CATS metrics are periodically updated, these routes are also periodically advertises to the RAN PEs with the latest CATS metrics alongside the service, service instance identities (CS-ID, CIS-ID) and the PE SRv6 SID.

- * When a UE request session establishment to the mobile management system, the MUP-C receives the session information which includes the requested service CS-ID and the TEID of the session. The MUP-C provides this information to the RAN PEs via the T2ST routes.
- * Upon receiving the T2ST route, the C-PS at the RAN PE selects optimal service instance for the requested service (identified by the CS-ID) based on the current CATS metrics of the service advertised from the DSD routes. The RAN PE uses the D2D route that has the optimal service instance CIS-ID for steering the UE traffic to the optimal service site.
- * In case of the optimal service instance change to another site (SiteB in this example) due to current CATS metrics update, the C-PS at the RAN-PE selects a new optimal service instance based on current CATS metrics. Then, the RAN PE uses the corresponding DSD route of the new optimal service instance to route the UE request.

4.2.4. UE mobility handling procedure

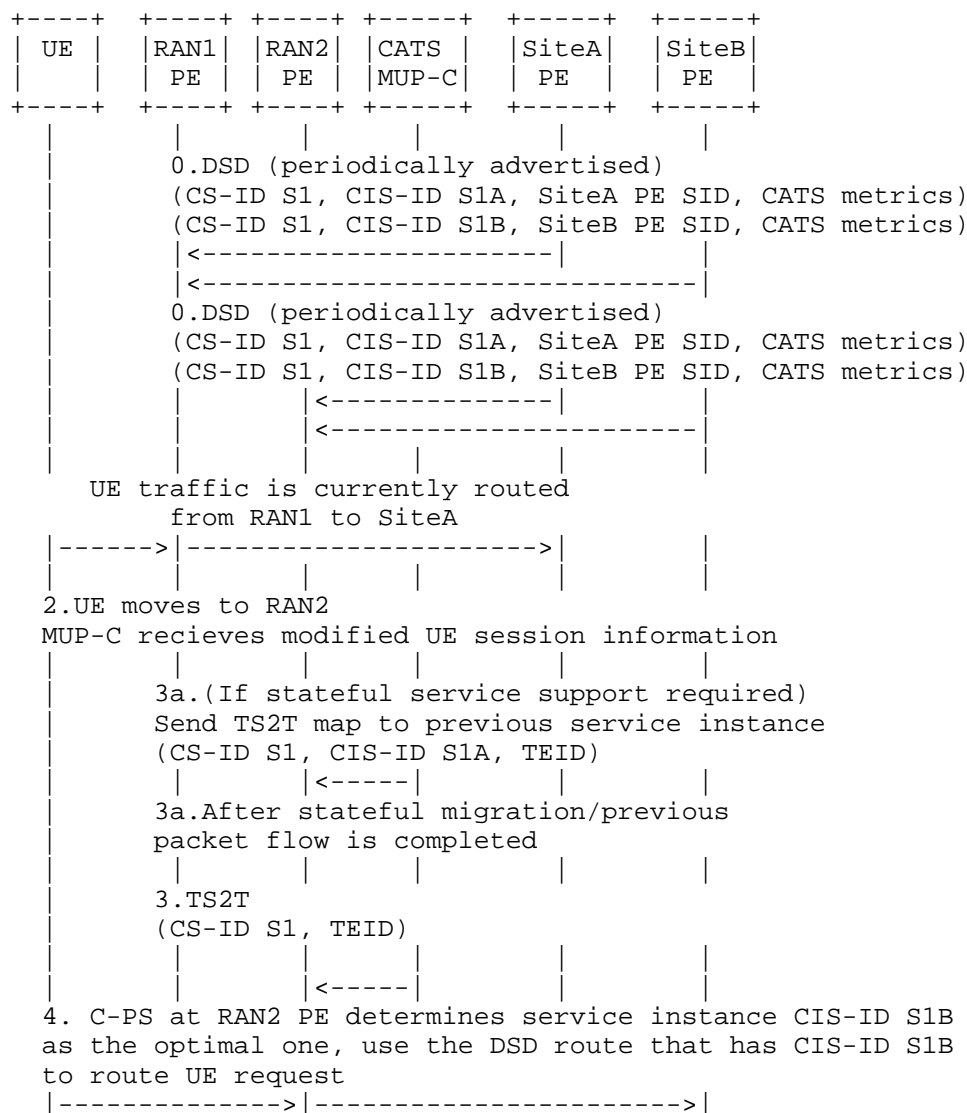


Figure 6: UE mobility handling procedure

Figure 7 describes the sequence of how the CATS-MUP Distributed deployment approach setup the underlay routing path when UE moves to a new RAN.

- * DSD routes are periodically advertised from service site PEs to all RAN PEs with the corresponding CATS metrics
- * When UE moves to a different RAN, the MUP-C receives modified UE session information from the mobility management system, and provided the session information TEID and the requested service CS-ID to the new RAN PE via the T2ST route.
- * The C-PS at the new RAN PE selects new optimal service instance based on current CATS metrics provided by the DSD routes. Then it uses the DSD route that maps to the optimal service instance CIS-ID to route the UE request.
- * In case of stateful/sticky service support are required, where connection to the previous service instance need to be maintained, the MUP-C should firstly send a special T2ST route that additionally contains the previous service instance CIS-ID information. The RAN-PE then uses the DSD route that maps to the previous instance CIS-ID. Then, depends on the session continuity policy, the CATS-MUP-C can update the PE with a normal T2ST route later. For example, new TS2T can be sent to PE after service packet flow to the previous instance is completed or after service state migration is completed. Once the RAN-PE receives the normal T2ST route, it uses the DSD route that maps to the optimal service instance CIS-ID to route the UE request.

5. References

5.1. Informative References

[I-D.draft-ietf-cats-framework]

Li, C., Du, Z., Boucadair, M., Contreras, L. M., Drake, J., Huang, D., and G. S. Mishra, "A Framework for Computing-Aware Traffic Steering (CATS)", Work in Progress, Internet-Draft, draft-ietf-cats-framework, 10 February 2025, <<https://datatracker.ietf.org/doc/draft-ietf-cats-framework/>>.

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

Yao, K., Shi, H., Li, C., Contreras, LM., and J. Ros-Giralt, "Mobile User Plane Architecture using Segment Routing for Distributed Mobility Management", 5 February 2025, <<https://datatracker.ietf.org/doc/draft-ietf-cats-metric-definition/>>.

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

Yao, K., Trossen, D., Boucadair, M., Contreras, LM., Shi, H., Li, Y., Zhang, S., and Q. An, "Mobile User Plane

Architecture using Segment Routing for Distributed Mobility Management", 14 February 2025, <<https://datatracker.ietf.org/doc/draft-ietf-cats-usecases-requirements/>>.

[I-D.draft-ietf-dmm-mup-architecture]

Matsushima, S., Horiba, K., Kawakami, Y., Murakami, T., Patel, K., and J. Horn, "Mobile User Plane Architecture using Segment Routing for Distributed Mobility Management", Work in Progress, Internet-Draft, ietf-dmm-mup-architecture, 30 May 2025, <<https://datatracker.ietf.org/doc/draft-ietf-dmm-mup-architecture/>>.

[I-D.draft-ietf-dmm-srv6mob-arch]

Kohno, M., Clad, F., Camarillo, P., Ali, Z., and L. Jalil, "Architecture Discussion on SRv6 Mobile User plane", 15 February 2024, <<https://datatracker.ietf.org/doc/draft-ietf-dmm-srv6mob-arch/>>.

[I-D.draft-ietf-idr-5g-edge-service-metadata]

Dunbar, L., Majumdar, K., Li, C., Mishra, G., and Z. Du, "Distribute Service Metric By BGP", 22 July 2024, <<https://datatracker.ietf.org/doc/draft-ietf-idr-5g-edge-service-metadata/>>.

[I-D.draft-lin-idr-distribute-service-metric]

Lin, C. and H. Yao, "Distribute Service Metric By BGP", 6 June 2024, <<https://datatracker.ietf.org/doc/draft-lin-idr-distribute-service-metric/>>.

[ieee-access-cats-mup]

Tran, M-N., Duong, V-B., and Y. Kim, "Design of Computing-Aware Traffic Steering Architecture for 5G Mobile User Plane", 24 June 2024, <<https://doi.org/10.1109/ACCESS.2024.3418960>>.

[RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", December 2006, <<https://datatracker.ietf.org/doc/rfc4786/>>.

[RFC7333] Chan, H., Liu, D., Seite, P., Yokota, H., and J. Korhonen, "Requirements for Distributed Mobility Management", August 2014, <<https://datatracker.ietf.org/doc/rfc7333/>>.

Authors' Addresses

Minh-Ngoc Tran
Soongsil University
369, Sangdo-ro, Dongjak-gu
Seoul
06978
Republic of Korea
Email: mipearlskal307@dcn.ssu.ac.kr

Younghan Kim
Soongsil University
369, Sangdo-ro, Dongjak-gu
Seoul
06978
Republic of Korea
Phone: +82 10 2691 0904
Email: younghak@ssu.ac.kr