

CATS WG
Internet-Draft
Intended status: Standards Track
Expires: 18 October 2026

CJ. Bernardos
UC3M
A. Mourad
InterDigital
16 April 2026

Site Mobility-Enabled Computing Aware Traffic Steering using IP address
anchoring
draft-bernardos-cats-anchoring-site-mobility-01

Abstract

The IETF CATS WG addresses the problem of how the network infrastructure can steer traffic between clients of a service and sites offering the service, considering both network metrics (such as bandwidth and latency), and compute metrics (such as processing, storage capabilities, and capacity).

This document defines new extensions and procedures for a terminal hosting a site, to benefit from transparent mobility management adapting to specific connectivity and computing requirements.

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 18 October 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document.

Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction and Problem Statement	2
1.1. Use case scenario	2
1.2. Problem statement	4
2. Terminology	5
3. Enabling site mobility with IP anchor mobility for CATS . . .	5
4. IANA Considerations	11
5. Security Considerations	11
6. Acknowledgments	11
7. Informative References	11
Authors' Addresses	12

1. Introduction and Problem Statement

1.1. Use case scenario

There are sensing scenarios and use cases that involve a distributed sensing task, in which one or multiple sensors participate, and that requires a supporting sensing service (e.g., fusing sensing measurements from different sensors and computing the sensing result). This sensing service needs to be executed by some kind of sensing processing/computing function, which might be hosted on a device that may be mobile (e.g., a terminal/User Equipment). The sensing service demands connectivity and computing requirements that are necessary to properly operate (and to timely deliver the sensing results).

The distributed task needs to be set up and the different participants (sensors and required sensing processing units) need to be selected. Both the sensors and processing functions might be on the move and could therefore change their respective points of attachment to the network. This requires mobility procedures that are aware of the associated sensing requirements in terms of latency and computing.

Note that this is just an example, other services would also benefit from compute and connectivity traffic steering. For the sake of having a simpler service, we can also consider an AR/VR/XR service where a terminal connected to the network needs to instantiate a service in the network to aid in the AR/VR/XR service by providing computing capabilities with latency constraints.

Note on terminology. In this document we use the old terminology in which by ICR we mean Ingress CATS-Forwarder [I-D.ietf-cats-framework], and by ECR we mean Egress CATS-Forwarder.

Figure 1 shows an exemplary scenario. There is a distributed sensing task (e.g., requested by an Application or Network Function). This involves four sensor functions (three hosted at UEs: UE #1, #2 and #3, and one at the RAN: Access Network #1) and a sensing processing function (hosted at a UE: UE #4). The selection of the composition of the sensing group is out of the scope of this document. The sensors might be of the same or different technology and might be connected to the same RAN or different ones (which might also be of different access technologies). In this example, we have the following configuration:

- * Three sensor (S) functions hosted at UEs/terminals: S#1@UE#1, S#2@UE#2 and S#3@UE#3.
- * One sensor function hosted at the access network (AN): S#4@AN#1.
- * A sensing processing (SP) function is hosted at a UE: SP@UE#4.
- * Access Networks are of different technology: AN#1 is 3GPPP and AN#2 is WiFi. There is also a third AN (AN#3) which is also a WiFi one.
- * UE#1 and UE#2 are connected to AN#1. UE#3 and UE#4 are connected to AN#2.

Since some of the sensors and the sensing processing function are mobile, a mobility event (e.g., a change of point of attachment of the UE hosting the sensing processing function, from AN#2 to AN#3) triggers updates of the connectivity and traffic steering used to forward the sensing data between the sensors and the sensing processing function.




Figure 1: Exemplary scenario

1.2. Problem statement

The main problem that this document tries to address is the following. Current distributed sensing solutions do not consider mobility of a device hosting a sensing processing function. Mobility solutions supporting sensing processing mobility and jointly considering computing and networking requirements are missing.

Based on the former, this document proposes solutions to enable devices (such as a UE) hosting a sensing processing function to be mobile and change its point of attachment to the network, while meeting the requirements in terms of computing and networking associated with the sensing service. In particular, this document addresses the following questions: what mechanisms are needed to support mobility of a sensing processing device participating in a distributed sensing task that has its own associated connectivity and computing requirement, leveraging the architecture defined in [I-D.bernardos-cats-ip-address-anchoring]?

2. Terminology

The following terms used in this document are defined by the IETF:

ECR: Egress CATS router. This refers to the Egress CATS-Forwarder as defined in [I-D.ietf-cats-framework].

ICR: Ingress CATS router. This refers to the Ingress CATS-Forwarder as defined in [I-D.ietf-cats-framework].

The following terms are used in this document:

(Distributed) Sensing Group: a group of devices participating on a sensing task.

Sensing Traffic: traffic used (after some processing) to generate a sensing result.

Sensing Processing Function: a function processing sensing traffic (potentially from different sources) to generate a sensing result (or something that can be further processed to generate a sensing result).

Sensing Signal: radio signal used in the processing.

Sensor Function: function running on a device participating on a sensing task that generates and/or processes a sensing signal.

ISF: integrated sensing function. This is a logical in charge of controlling the distributed sensing task.

CATS Agent: logical entity performing a function related to computing aware traffic steering.

Note that we use UE or terminal to refer to a mobile host.

3. Enabling site mobility with IP anchor mobility for CATS

We describe next an example of operation and signaling for a mobile terminal hosting a sensing processing function to support mobility (i.e., change of point of attachment) while meeting the connectivity and computing requirements associated to the sensing service. In addition to the functionality defined in [I-D.bernardos-cats-ip-address-anchoring] and [I-D.bernardos-cats-anchoring-service-mobility], this document defines a new functionality:

- * Site mobility: it deals with the procedures required to: (i) detect or predict a change of the current conditions due to a change in the point of attachment, jointly considering computing and networking, requiring of a service-hosting terminal (e.g., a terminal hosting a sensing processing function) mobility operation; (ii) decide whether a change of service instance location is needed, and, if so, signaling it to the appropriate network entity (e.g., the ISF) to trigger the service migration; and (iii) perform the required signaling to ensure traffic is not disrupted as a result of the change of point of attachment (e.g., IP tunnel(s) update(s) and traffic steering modifications).

Additionally, the procedures presented herein enable temporary change of PoA. For example, the terminal might decide to change PoA to a more "costly/expensive" network (e.g., with better capabilities), for a limited time to send data that require high bandwidth. Then, it may change PoA back to the initial PoA to continue its initial operation.

In some scenarios, where the UE has connectivity to more than one access network (AN), they may be assigned "Primary" and "Secondary" roles. For example, the secondary ANs may be used for establishing temporary PoAs.

Figure 2 shows the message sequence chart of the site mobility for CATS, which is explained next:

There is a distributed sensing task ongoing, which has been previously set up. This setup involved selecting the sensors (i.e., network locations hosting the sensing functions; in this example UE#1, UE#2, UE#3 and AN#1) that are part of the sensing task and selecting the sensing processing function (i.e., network location hosting the sensing processing function; in this example UE#4) in charge of processing the sensing data.

The sensing processing function impose its own requirements in terms of connectivity with the sensors, but also of computing capacity required for the sensing data processing (e.g., sensor fusion).

In terms of the distributed sensing task and the connection of the participating entities:

- * Each participating device has an IP (topological) address configured from the respective (access) networks it is attached to (probably benefitting from a local breakout approach to avoid too high network latencies).

- * In addition to the topological IP address, each function participating on the sensing task (i.e., sensors and processing functions) is provided with an IP prefix (anchored at the processing function) to configure an IP address to use for the sensing-related communications. This facilitates mobility management in case any of the nodes change its point of attachment.
- * Sensing data traffic is tunneled between the sensing functions and the sensing processing function, using the topological IP addresses of the devices hosting the functions as outer-header end-points. The use of IP-in-IP tunneling facilitates applying traffic handling policies to the packets as well as makes transparent the mobility to the actual sensing-related functions. This might be configured by an external controlling entity, such as the ISF.

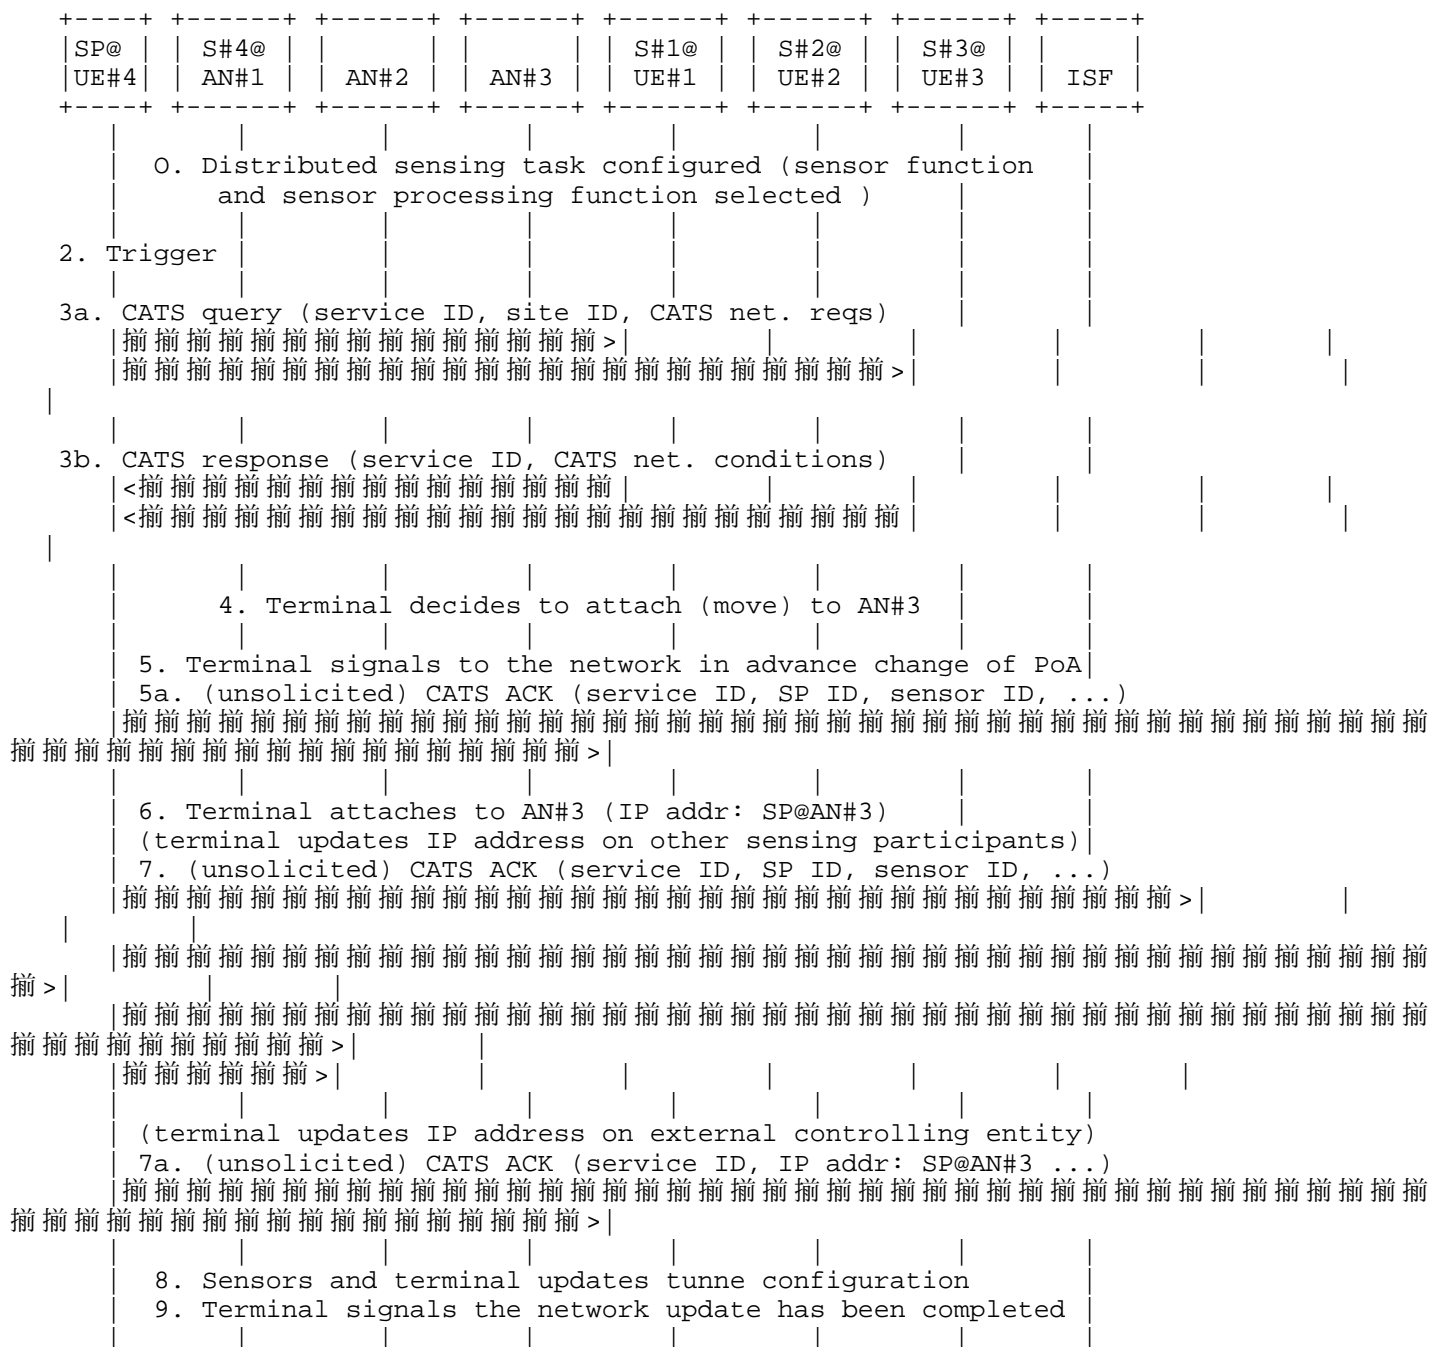


Figure 2: Exemplary signaling

Next we describe the different steps involved to support the mobility of the device hosting the sensing processing function, while ensuring that the requirements posed by the processing function are met for the particular distributed sensing task that is currently running.

0. The distributed sensing task is already configured (sensors selected, sensing processing function selected).
1. All the devices participating in the distributed sensing task may actively monitor the quality of the connection (e.g., using in-situ OAM or other types of telemetry/monitoring). This can help detect if the quality is degrading and approaching a level that is not good enough for the sensing service. Each device that is participating on the telemetry/monitoring may embed relevant information into the tunneling headers of the data packets transporting the sensing data (i.e. in-situ OAM) to support this task. Some examples of possible data items that can be included are:
 - * Sensing service tag/id.
 - * Timestamps, used to monitor latency.
 - * Sensing type labels (to prioritize traffic from specific types of sensors over others), etc.
 - * Sensing-quality related metadata, to aid recipients assessing if a sensing function is capable of perform as required.
2. The measured sensing service conditions or pure mobility reasons (e.g., detecting new neighbor points of an attachment) may trigger the terminal hosting the sensing processing function to look for candidate Points of Attachment (PoAs).
3. The terminal hosting the sensing processing function may start searching for information allowing to proactively prepare the distributed sensing task for an update due to the mobility of the sensing processing function:
 - * The terminal sends CATS query messages to potential PoAs in reach. These might include certain information regarding connectivity requirements of the sensing service (e.g., max latency to a given IP address) so the candidate PoA can measure it and provide estimations as part of the CATS response. This message may indicate whether the new PoA is intended to be used in temporary basis (and related parameters, e.g., time duration).
 - * The candidate PoAs respond with a CATS response message. This can be done through L2 (e.g., 802.11 Probe Requests and Probe Responses) or L3 (e.g., Router Solicitation and Router Advertisement) messages.

4. Based on the received information on the CATS response messages, the device hosting the sensing processing function decides to change its PoA from AN#2 to AN#3.
5. Before moving, the device hosting the sensing processing function prepares the network for the imminent handover. This might involve, as non-limiting examples: (i) setting up multicasting in the network so the sensing data is duplicated (for a limited amount of time) to the current location and the upcoming one (identified by the current and new topological IP addresses); (ii) configuring the transport network to perform PREOF to minimize the packet loss.
 - * The terminal may also notify an external sensing controlling entity, such as the ISF, about the upcoming change of PoA (and whether if the change is temporary). This might be used by the ISF to evaluate if a reconfiguration of the distributed sensing task is required.
6. The terminal changes its point of attachment and obtains a new topological IP address (SP@AN#3 in this example).
7. The terminal signals this new IP address to the other participants of the sensing task by sending an (unsolicited) CATS ACK message, which includes the following information:
 - * Service ID: an ID univocally identifying the sensing service.
 - * Sensing processing/site ID: an ID univocally identifying the terminal hosting the sensing processing function.
 - * Sensor ID: an ID univocally identifying the target sensor function.
 - * CATS conditions: networking and computing requirements associated with the sensing service.
 - * IP prefix: IP prefix shared by all the devices participating in the distributed sensing task.
 - * (Topological) IP addr: SP@AN#3.

Optionally, this terminal may also inform the external sensing controlling entity (e.g., an ISF). This can be used for example to evaluate (by the sensing controller/coordinating entity) if the sensing processing function should be migrated to a different location. In scenarios where the change to the new PoA is temporary, this message may indicate so and may include a time value, e.g., a timer, start time, end time, time duration, indicating when and/or how long the PoA may be used for.

8. All the sensors update the tunneling configuration with the new topological IP address of the device hosting the sensing processing function as end-point.
9. The terminal signals the network that the handover has been completed. This might involve stopping the temporal configuration of the transport network that was conducted in step 3 to minimize sensing data loss.
10. The sensing data traffic is now forwarded following the updated tunnels.

4. IANA Considerations

TBD.

5. Security Considerations

TBD.

6. Acknowledgments

The work of Carlos J. Bernardos in this document has been partially supported by the Horizon Europe MultiX (Grant 101192521), DISCO6G-CM (TEC-2024/COM-360) and UNICO I+D 6G-DATADRIVEN projects.

7. Informative References

[I-D.bernardos-cats-anchoring-service-mobility]

Bernardos, C. J. and A. Mourad, "Service Mobility-Enabled Computing Aware Traffic Steering using IP address anchoring", Work in Progress, Internet-Draft, draft-bernardos-cats-anchoring-service-mobility-04, 24 September 2025, <<https://datatracker.ietf.org/doc/html/draft-bernardos-cats-anchoring-service-mobility-04>>.

[I-D.bernardos-cats-ip-address-anchoring]

Bernardos, C. J. and A. Mourad, "Computing Aware Traffic Steering using IP address anchoring", Work in Progress,

Internet-Draft, draft-bernardos-cats-ip-address-anchoring-04, 24 September 2025,
<<https://datatracker.ietf.org/doc/html/draft-bernardos-cats-ip-address-anchoring-04>>.

[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-24, 2 April 2026,
<<https://datatracker.ietf.org/doc/html/draft-ietf-cats-framework-24>>.

Authors' Addresses

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
28911 Leganes, Madrid
Spain
Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: <http://www.it.uc3m.es/cjbc/>

Alain Mourad
InterDigital Europe
Email: Alain.Mourad@InterDigital.com
URI: <http://www.InterDigital.com/>