

Distributed Mobility Management (DMM)
Internet-Draft
Intended status: Informational
Expires: 3 September 2026

M. Liebsch
NEC
J. Mutikainen
NTT Docomo
Z. Zhang
Juniper Networks
T. Jiang
China Mobile
2 March 2026

Mobile Traffic Steering
draft-ietf-dmm-mts-01.txt

Abstract

The evolution of cellular mobile communication systems is aligned with an increasing demand for customized deployments, energy efficiency, dynamic re-configurability and the integration and use of other network technologies, such as non-cellular radio access technologies and non-terrestrial networks. In order to achieve and maintain the expected service quality and continuity, such systems should be designed and controllable end-to-end, taking all involved network domains and segments into account. This document discusses an end-to-end system from an advanced use cases perspective and substantiates the demand for solutions to share information and enable control interfaces between all connected network domains, including the mobile communication system and the transport network that stretches up to the data networks that host service instances. In the view of flexible implementations and deployment, two architectural principles, leveraging either a dedicated controller or a decentralized control plane, are described and discussed, accompanied by operational aspects and an associated information model that enable end-to-end mobile traffic treatment and steering in such complex and dynamically changing networks.

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 3 September 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. Terminology	3
2. Introduction	3
3. Reference Architecture in the view of advanced end-to-end operations	6
4. System Evolution and Use Cases	6
4.1. General directions and impact	7
4.2. MCS-proactive UPA relocation	8
4.3. MCS reactive UPA relocation	9
4.4. Inter-UPA communication	9
4.5. Node ephemerality	10
5. Framework and Deployment Options	13
5.1. Mobile User Plane and Data Plane aspects	14
5.2. Dedicated Control Plane	14
5.3. Decentralized Control Plane	15
5.4. Control Interfaces	16
6. Operational Aspects	17
6.1. User-/Data Plane - Topology Considerations	18
6.2. Mode I Operation - Dedicated Control Plane	19
6.3. Mode II Operation - Decentralized Control Plane	21
6.4. Mode III Operation - Hybrid	23
7. IANA Considerations	24
8. Security Considerations	24
9. Acknowledgments	24
10. References	24
10.1. Normative References	24

10.2. Informative References	24
Appendix A. Information Models	25
Appendix B. Exemplary Application of MTS to a 5G System	28
Authors' Addresses	30

1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Introduction

The evolution of cellular mobile communication systems resulted in a clear separation of control plane and user plane functions. The control plane comprises functions for security, mobile subscriber management, handover, and mobility management. The user plane functions represent anchor points for a user's traffic and enforce policies to user plane traffic, such as for traffic engineering or chargeable event monitoring and reporting. Compared to early standards, today's mobile communication systems support the demand of mobile subscribers as well as mobile network operators better in terms of deployment options and route optimization. This includes the decentralized deployment and selection of user plane anchors, which can, e.g., be topologically located on the path between a mobile user and a user's currently connected and used application service.

Such flexible deployment of user plane anchors is aligned with a rising interest in distributing compute resources. Edge Computing enables the provisioning of compute resources, which are topologically close to mobile users. This helps on the one hand to reduce end-to-end latency between a mobile user and a service or a compute resource, that performs a certain computation task for the user. On the other hand, it keeps data for certain use cases local, which can be leveraged for certain analytics tasks, where data is only of local interest, or for meeting some privacy and security objectives.

Figure 1 depicts the end-to-end view of a mobile user, which connects to an application service function (ASF) through a mobile communication system (MCS). The ASF serves a user's request at the data plane layer, while the service may have a dedicated application control function (ACF). The user, as an example, may connect to the ACF to configure the service and gets served by the associated ASF. The service components are deployed in a data network, which may be

represented by a central cloud, or a distributed data center with cloud computing resources. In alignment with mobile communication system standards, the MCS is separated into a radio access network (RAN) part and a mobile core network, which comprises the MCN control plane and the user plane anchors. The RAN offers cellular and non-cellular radio access technologies, such as WiFi, to a mobile user to connect to the MCS for mobile access to a target service.

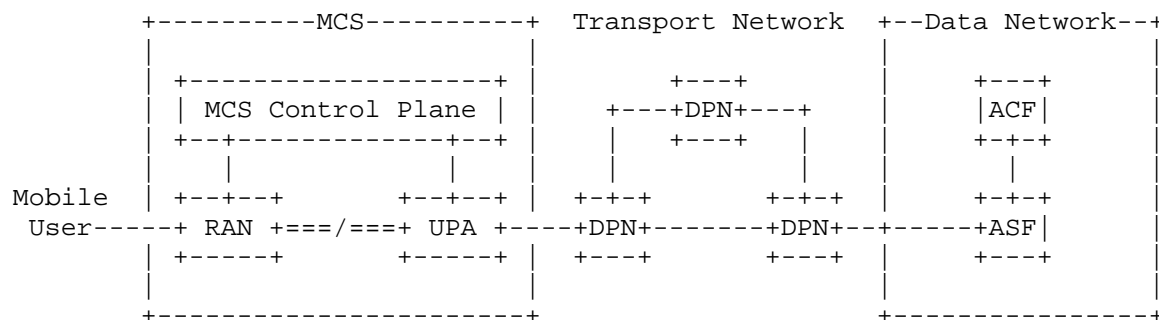


Figure 1: Reference End-to-End Architecture.

While mobile communication standards focus on the components of the MCS and their operation, details on the network in between the MCS and a data network are often treated out-of-scope of the mobile communication standard and deployment specific. Such network may indeed be deployment specific and differs in the type of network nodes and protocols that route or switch traffic in between the MCS and the data network. Figure 1 depicts this network as a set of data plane nodes (DPN), which can, for example, be routers with redundant paths and implementing MPLS or SRv6 in their transport plane.

This document addresses relevant system components in an end-to-end reference architecture to discuss advanced use cases and associated deployment options. In particular use cases are addressed, where components or network functions on the end-to-end path change, move or discontinue operation while the mobile user is still connected to the service. The objective of this document is to analyze these use cases and elaborate how existing or new IETF technology can serve as enabler to accomplish service continuity for a mobile user in such agile and dynamic system by means of controlled treatment of a mobile user's end-to-end traffic for steering, with an option to extend for advanced treatment policies, such as traffic engineering. The treatment for steering of traffic in between a mobile communication system and a data network should be according to computed and agreed policies taking mobile communication system-, transport network and

data network aspects into account. Advanced traffic treatment goes beyond path computation and can include quality-of-service or metering aspects, e.g. in the view of a mobile subscriber's service level agreements.

Two architectural principles are described to extend a mobile communication system towards the transport- and data network for mobile traffic steering, leveraging either a dedicated controller or a decentralized control plane. Each option is described and discussed regarding its strengths and shortcomings, accompanied by operational aspects and an associated information model that enable end-to-end mobile traffic treatment and steering in complex and dynamically changing networks.

Each of the described use cases is analyzed and discussed in the view of technical challenges and operational aspects in each of the two described architectural principles and deployment options. Advanced use cases are in-line with the view in industry, research community as well as pre-standards effort. The following deployment- and operational aspects of session and service continuity are included:

- o Mid-session relocation of a mobile user's UPA, e.g. due to user mobility, UPA failover or load balancing.

- o Deployments with moving or ephemeral system-relevant nodes on the end-to-end path. These nodes include system components, such as radio access network components and a MCS's UPA that are on-board of a moving resource, such as a low earth orbit (LEO) satellite, or an energy constrained node whose schedule enforces a power save- or inactive mode.

- o Mid-session relocation of a mobile user's serving data network, e.g. due to the service resource's mobility, service failover, energy/costs or quality of service reasons.

This document includes the description of two solution options, that complement a MCS without interference by means of well interconnected and collaborative control- and data planes. Operational aspects of the two solution options are described and semantics as well as information models, that apply to relevant reference points, are specified.

3. Reference Architecture in the view of advanced end-to-end operations

Figure 1 depicts a reference architecture for end-to-end operations, which includes communication between a mobile user and an Application Service Function (ASF) as well as user-to-user communications. The ACF can be a service control instance hosted in a central or a distributed cloud resources, or a workload placed by the mobile user to an assigned compute resource in either a central or a decentralized cloud. While the MCS ensures mobile connectivity and data services between a mobile user and its UPA, a Transport Network comprises a network of data plane nodes (DPN) and ensures routing of a mobile user's traffic between its UPA and one or multiple data networks.

The Transport Network may implement redundant paths and select the most suitable route based on the topological location and associated IP address of the ASF and mobile user, respectively. The mobile user's IP address may be topologically correct and fit the UPA's network, or it does not match the network where the mobile user's assigned UPA is located. The latter case applies, for example, to mobile user subscriptions, which have static IP addresses assigned even in a deployment of distributed UPAs. Routing of topologically incorrect mobile device IP addresses can be tackled by host routes, either on the PE routers of network-centric overlays, such as via encapsulation or label switching, or on all relevant on-path DPNs in the transport network in between the mobile user's UPA and its connected data networks.

With reference to Figure 1, the DPNs that are depicted close to the data network and to the MCS can represent the associated domains' PE router, or be independent routers in the domain of the transport network provider that connect to the PE routers of the MCS domain and the data network domain respectively.

The reference architecture's domains, the MCS, the transport network and the data network, may share the same AS or be located in different ASs. The relevance and scope of solutions for mobile traffic steering in advanced and anticipated use cases, that this document describes in Section 4, depend on the involved ASs.

4. System Evolution and Use Cases

4.1. General directions and impact

While recent standards for MCSs are still being evolved and extended, the research community and standards organizations started to elaborate on a new MCS generation and advanced use cases. Key objectives include tighter integration of a variety of radio access technologies, including non-terrestrial networks, e.g. satellites, support of energy-efficiency schemes, and runtime changes of service instances and network functions. Customized deployments place a set of UPAs topologically close to or even inside distributed data networks, which may host compute resources for service instances or temporary workload that has been offloaded from a mobile user.

In case of user mobility, the MCS may assign a new UPA to the mobile device, which can lead to sub-optimal routes between a data network and the mobile user. While the relocation of the UPA is handled by the MCS, awareness of such change is required in the transport network and the data network to update traffic treatment. For efficiency reasons, even the use of temporarily available and ephemeral resources is considered, which includes, for example, energy constrained or solar powered resources, mobile resources, such as vehicles, drones or satellites. Industry develops satellites with radio- and optical units to enable inter-satellite links and to connect to ground stations as well as mobile users on earth. Recent directions consider regenerative payload on satellites, which goes beyond the use of a satellite as a communication relay but offers on-board compute, storage and networking resources. This enables the deployment of functions from mobile communication system, such as a UPA, and Internet technology, such as a routing function, on satellites to enable mobile access to such satellite resources. Even the placement of suitable compute resources on-board of satellites is considered, which makes them pre-destined to host application service instances and other types of workload.

These directions enable highly efficient deployments in terms of customization, resources usage, energy saving, communication latency, connectivity and Internet coverage, as well as workload placement and distribution. On the other hand, the dynamics of system components' availability as well as their geographic and topological location requires sharing of information between the involved system components across domains, i.e. MCS, transport network and data network. Exposing such information and offering control interfaces between these domains allows timely configuration of alternative paths to continue a service and steer mobile data plane traffic between a mobile user and its service or between two mobile users. In the discussed end-to-end system, a solution should also enable traffic engineering in the traversed network segments and data networks.

UPA relocation and other dynamic changes in the network may result in reordering of the data plane packets. Solutions for mobile traffic treatment and steering should provide suitable enablers to mitigate the impact of such changes to packet re-ordering.

[I-D.zzhang-dmm-5gdn-end-marker] describes the use of End Markers to address this problem.

The following sub-sections describe selected use cases, their impact to the end-to-end system and the need for mobile traffic steering solutions. Section 5 introduces two architectural principles and options to tackle mobile traffic steering, which are subsequently discussed in the view of enabling technology, functional operation and required semantics on relevant control plane interfaces.

4.2. MCS-proactive UPA relocation

Figure 2 depicts a use case, where the MCS proactively relocates the mobile user's current UPA, UPA1, to a new UPA, UPA2. The resulting route between service instance ASF1 in the data network and UPA2 may differ from the previous route to UPA1. In case the IP address of the mobile user's device needs to survive in order to not break the current session with the used service at ASF1, the transport network needs to reactively enforce rules for traffic steering following the MCS's UPA relocation, e.g. by applying host routes on the path's DPNs, by applying source routing or an overlay route, i.e. IP tunnel. In case the user's IP address needs to continue only as long as the data session with ASF1 takes, the MCS may assign a new, topologically correct IP address from the network of UPA2 to the mobile user device. Details about whether a user device's IP address is persistent or can be updated after data sessions that use such IP address terminate, is out of scope of this document. The same applies to associated user profiles in the MCS, that could comprise details about how transient or persistent a user device's single or multiple IP addresses can be handled. After the mobile user terminates its data session with ASF1, new data sessions may use the new IP address and transient host routes in the transport network can be removed. Figure 2 depicts also the change of the serving ASF, which may happen also as a result of user mobility. As example, in an automotive use case, vehicles represent mobile users and connect to a most suitable, nearby ASF that is hosted by one of many available distributed cloud networks. Due to mobility, the MCS may relocate the mobile user's UPA from UPA1 to UPA2. In order to keep routing paths short and latency low, the ASF may be relocated from ASF1 to ASF2, as ASF2 is deployed in a more suitable data network.

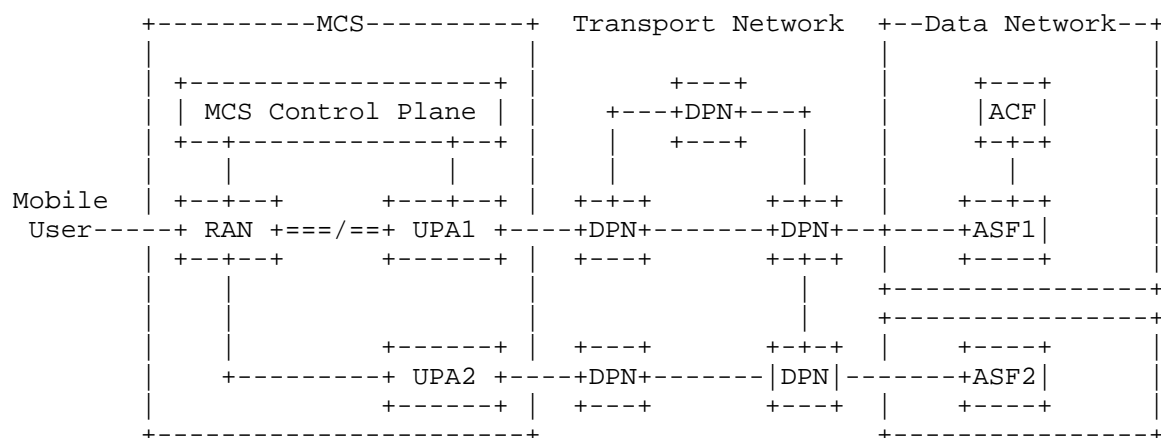


Figure 2: Mid-session UPA relocation, e.g. due to user mobility.

4.3. MCS reactive UPA relocation

In this case the source and rationale behind a UPA relocation in the MCS is different and results from a decision to change the serving ASF from ASF1 to ASF2 (Figure 2). This may be because of an anticipation that the targeted geographic region of the mobile user gets better service from a data network that hosts ASF2. As further example, the transport network applies changes in its routing strategy and as a result, a route between the MCS and a new ASF instance will result in better service quality and lower latency. The new situation in the service data network and the transport network can be exposed to the MCS, which may relocate the UPA from UPA1 to UPA2 to optimize the end-to-end path.

4.4. Inter-UPA communication

A MCS provide may deploy distributed UPAs in order to shorten the route and resulting latency for the communication between two mobile users. Figure 3 depicts such case where mobile user1 has UPA1 assigned, while mobile user2 has UPA3 assigned. Mobile traffic between the two mobile users traverses UPA1 and UPA3 as well as the transport network. In case the MCS relocates the mobile user1's UPA1 to UPA2, a different route may apply. In case of demand for IP address continuity or enforcement of particular traffic engineering rules in the DPNs, the change in the UPAs must be notified to the transport network to apply adjusted policies in the relevant DPNs.

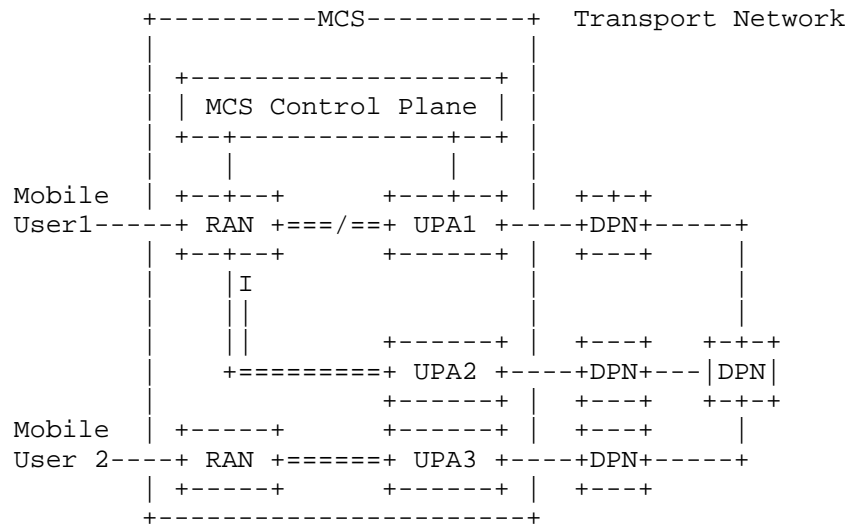


Figure 3: Mid-session UPA relocation, e.g. due to one user's mobility.

4.5. Node ephemerality

Components from the transport network or the MCS may be of ephemeral nature and discontinue service at a predictable or unpredictable point in time. Predictable discontinuity may be due to a scheduled power saving plan or mobility of the component. The latter case applies, as example, to LEO satellites in space, which are non-stationary. Advanced use cases consider the components of a radio base station and a UPA instance from the MCS as well as a DPN instance on-board of each satellite. In such case, the mobile user on earth or a High Altitude Platform Station (HAPS) is associated with the satellite's radio base station, the UPA and a DPN for further routing of the mobile traffic. At some point in time, the mobile user is not covered by the satellite's radio base station anymore and a different satellite, including its on-board radio base station, the UPA and DPN, takes over serving the mobile user. The end-to-end system needs to adapt to the changed triple of a radio base station, UPA and DPN in terms of traffic steering. Figure 4 depicts an example of such deployment, where the MCS's RAN and User Plane components are deployed on satellites, serving users through a mobile radio service link, while the MSC's control plane and further terrestrial transport network components as well as the serving data network are deployed on earth. Satellites connect to the terrestrial control- and data plane functions through a feeder link.

Taking an example from a 3GPP-based MCS, a joint deployment of a radio base station, UPA and virtual routing function can be realized per the description in [I-D.zhang-dmm-mup-evolution], while this draft addresses advanced options in the view of end-to-end service continuity and associated collaborative control between the MCS and the transport network.

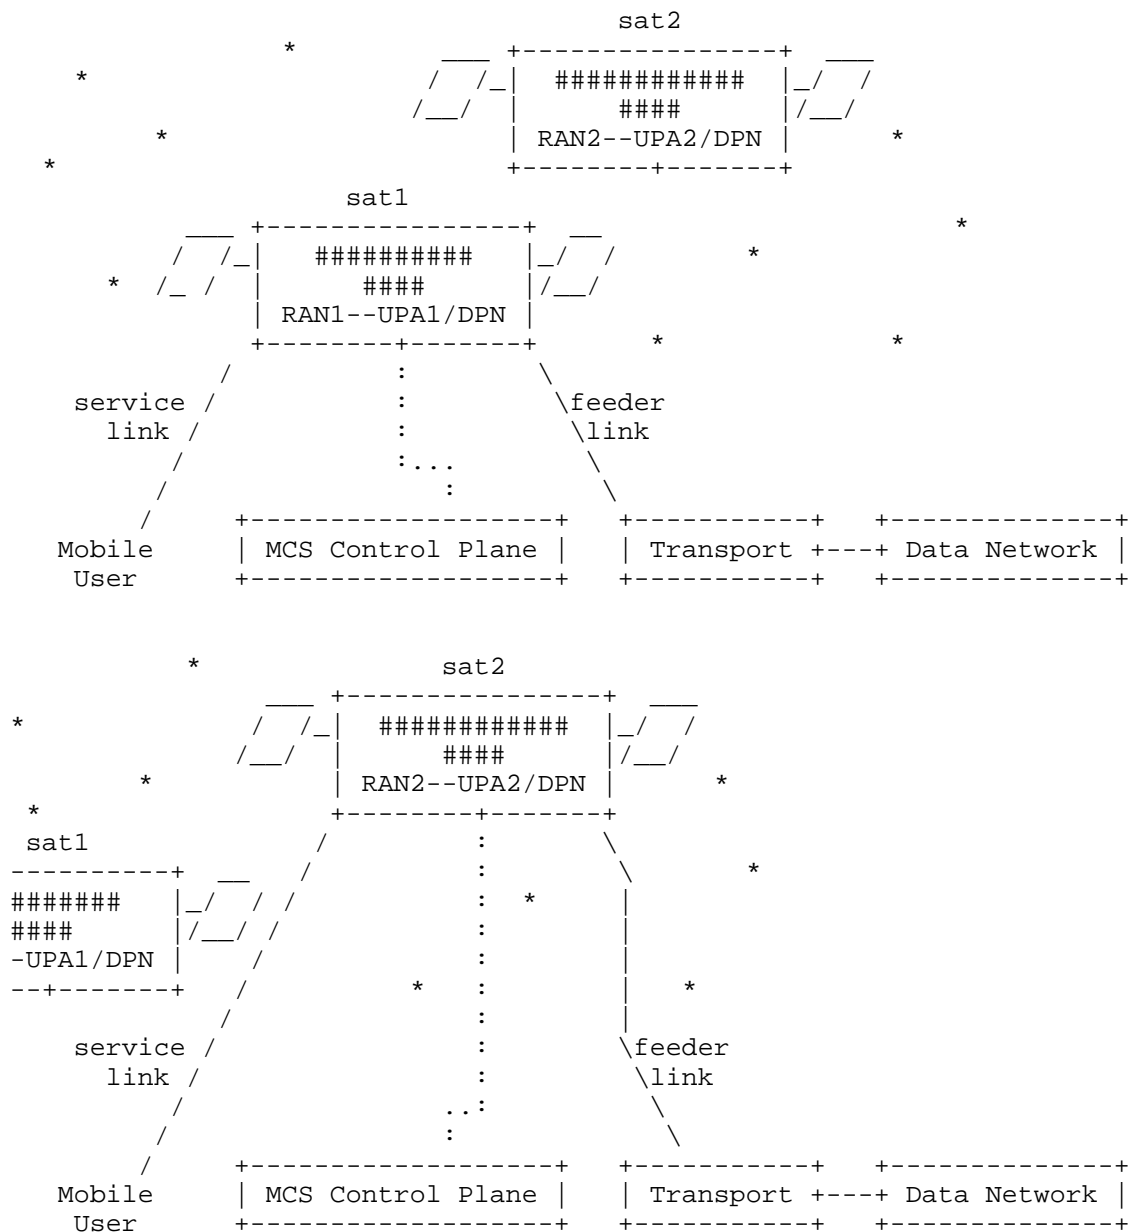


Figure 4: Example of node ephemerality: MCS- and Network components on board of moving satellites

5. Framework and Deployment Options

This section discusses two distinct architecture principles and deployment options that enable end-to-end awareness of changes in the network and its configuration that have impact on a mobile user's access to a service and the expected service quality. One option leverages a dedicated controller for mobile traffic steering, possibly built on top of a transport network controller (TN Controller), that shares relevant states and information with the MCS through a control plane interface. The same interface is used to receive notification from the MCS and to apply control commands in the MCS. The second option is based on decentralized control and requires tighter coupling of the transport network's routing functions with the MCS.

A dedicated controller for mobile traffic steering offers clearly defined access points towards a mobile communication system control plane for the alignment of traffic steering rules between the mobile communication system and the transport network. Further interfaces, which are complementary and out of scope of this document, from such dedicated controller can be exposed towards data networks for the end-to-end alignment of traffic steering and the enforcement of associated routing policies.

An implementation based on a decentralized control plane relies on DPNs that retrieve information about mobile traffic steering from the mobile communication system's UPAs. The need to update routing policies in the transport network is derived from such information and results in DPNs propagating these updates into the network of DPNs by means of routing protocols.

While such decentralized implementation of mobile traffic steering has no dependency on a dedicated controller, it enables updates of mobile traffic steering rules in the transport network mainly in a reactive way, following changes enforced by the mobile communication system's control plane into the UPAs. In contrast, a dedicated controller can leverage its interface to the mobile communication system control plane and also enforce updates in mobile traffic steering rules into the mobile communication system, e.g. due to changes in the transport network of serving data networks or DPNs. Such transport or data network-initiated update of end-to-end traffic steering rules may result, for example, in the mobile communication system relocating a single or a group of user plane sessions' UPA.

Note: The limitation of decentralized deployments in supporting mainly reactive mobile traffic steering, following the mobile communication system's rules about which UPA to use for which traffic, comes from the current limitations in mobile communication

systems that make a UPA to receive its configuration from the mobile communication system's control plane. Future systems may consider and support UPAs that can notify towards their control plane in the mobile communication system any updates in the transport network routing or serving data networks. The received information can be processed by the mobile communication system control plane for the computation and enforcement of updated UPA rules in alignment with the information received from the UPA. Such evolution in the mobile communication system mitigates the mentioned limitation of decentralized implementation and deployment of mobile traffic steering solutions and can enable bi-directional initiation of end-to-end mobile traffic steering rules, from the mobile communication system as well as from the transport network or the data network.

Based on the two distinct architecture principles, Section 6 describes operational aspects in the view of three modes, that (i) rely solely on the interface between the two dedicated control planes of the MCS and the TN, (ii) solely on the decentralized control plane, and (iii) a hybrid mode, that utilizes both, a dedicated control plane on the TN as well as the decentralized control plane.

5.1. Mobile User Plane and Data Plane aspects

A mobile user plane applies to the scope of the MCS and is managed by the MCS control plane to enable bi-directional data traffic between a mobile user's device and its assigned UPA. The network in between the UPA and a data network includes DPNs of the transport network and PE routers. This document focuses on IETF technology that applies to the control and data plane in the transport network and data network. The transport network's data plane transits into the MCS domain at a PE router or the UPA.

5.2. Dedicated Control Plane

Figure 5 illustrates a deployment with a dedicated control function in the transport network, which is denoted as Mobile Traffic Steering (MTS) Controller and may build on top of a TN Controller, leveraging its northbound API. The MTS Controller leverages an interface to the MCS control plane to receive notifications, such as during the change of a mobile user's UPA, and to apply changes in the configuration of a mobile user's states in the MCS, such as enforcement of a UPA change in alignment with a change in the transport network or data network. The data network may also use an interface either to the MTS controller or to the MCS control plane, which is used to enforce re-configurations in the transport network and/or the MCS in alignment with changes in the data network or a mobile user's service, e.g. ASF relocation or QoS settings.

The transport network's DPNs may make use of a dedicated ingress router, such as a PE router, to reach the MCS's UPAs. In this deployment option, the control of the UPAs is clearly separated from the DPN control, though the two control planes, the MCS control plane and the MTS Controller, are connected through a control plane interface. The MCS control plane may enforce rules in a UPA for uplink traffic treatment, e.g. to forward a mobile user's traffic to a particular DPN.

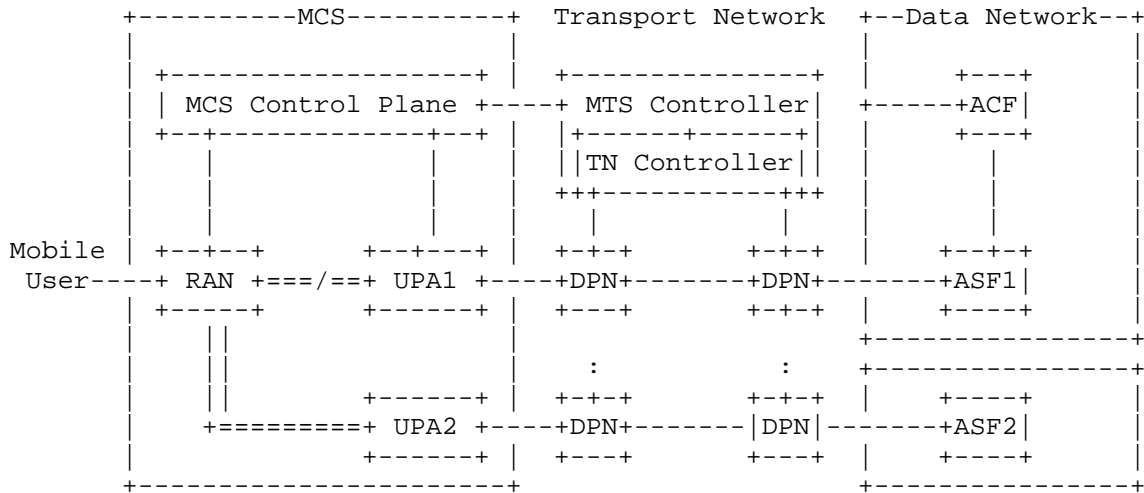


Figure 5: End-to-end architecture with a dedicated transport network controller and loose UPA-DPN coupling

5.3. Decentralized Control Plane

Figure 6 illustrates a decentralized deployment without a dedicated TN controller. DPNs expose routes through a routing protocol. To cover cases where the MCS relocates a mobile user's UPA, the DPN must either tightly couple with the UPAs and apply local APIs to learn about the mobile user and its changes configuration in the MCS, or the UPA applies a protocol to share relevant states and information with the DPN. This document does not consider a dedicated control interface between the MCS control plane and the DPN. Relevant semantic to enable the DPN to propagate updated routes towards the transport network and the data network must be available at the UPA. This may require an extension to the MCS control plane to configure the UPA with information that is relevant for mobile traffic treatment and steering in the transport network and the data network.




Figure 6: End-to-end architecture with a routing protocol-based propagation of traffic steering policies and tight UPA-DPN coupling

5.4. Control Interfaces

This document revolves around the specification of interface semantics that applies to control endpoints that operate in the MCS and the TN respectively. This control interface for mobile traffic steering is denoted `C_mts` and can apply in between (i) the MCS control plane and the MTS Controller in the TN, as well as between (ii) a UPA in the MCS and a DPN in the TN.

Figure 7 depicts the application of the `C_mts` interface and semantics per this document in between endpoints (c1) and (c2) of the MCS control plane and the MTS controller respectively. Figure 8 depicts the application of the `C_mts` interface and semantics per this document in between (c1) and (c2) of a UPA in the MCS and a DPN in the TN. MTS policies and resulting forwarding rules (fwd) apply in either case to the data plane endpoints (d1) and (d2) of a UPA in the MCS and a first hop DPN in the TN. For end-to-end mobile traffic steering, MTS policies may have impact on further DPNs in the TN between the MCS UPAs and data networks.

Note: (d2) per Figure 7 and Figure 8 applies to the first hop DPN from a UPA perspective towards the TN. A UPA may connect to multiple DPNs with a (d2) endpoint, and a DPN may connect to multiple UPAs with a (d1) endpoint. MTS policies are further propagated to the relevant DPNs in between a UPA and a data network by either mechanism, through a dedicated control- / data-plane protocol in between the MTS/TN Controller and DPNs, or a routing protocol that applies to the DPNs' control plane.

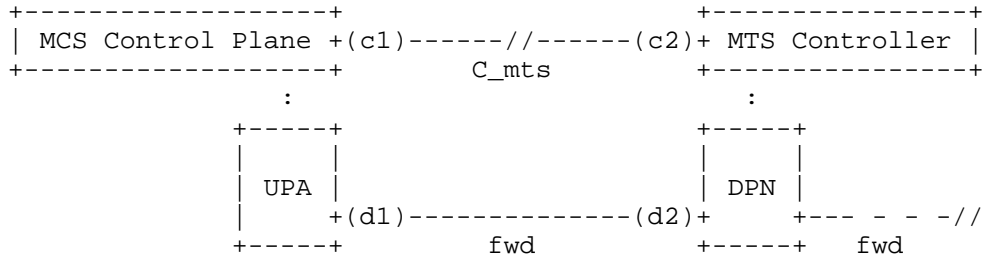


Figure 7: C_mts applies to reference point between MCS and MTS Control Plane




Figure 8: C_mts applies to reference point between an MCS's UPA and an associated DPN

6. Operational Aspects

This section describes operational aspects in the view of three modes, that (i) utilize solely a dedicated control plane for MTS, (ii) a decentralized control plane, or (iii) a hybrid mode, leveraging both, a dedicated MTS controller to make use of all control features enabled by the inter-action with the MCS control plane as well as a de-centralized routing protocol to propagate MTS route updates into the TN.

- Align traffic routing in the TN with updated TN conditions to match traffic's agreed service levels. Forwarding path may be updated in between a DN's first hop DPN and a UPA's first hop DPN. Such case may apply in case of changing TN conditions (load, delay, disruption, DPN mobility or energy saving mode) and be kept transparent to the MCS. Such case may also apply in case the TN needs to follow an event indication from the MCS, e.g. updated QoS policies that apply to identified traffic.

6.2. Mode I Operation - Dedicated Control Plane

This operational mode leverages the full feature set of the C_mts interface semantics to interact between an MTS Controller and the MCS control plane (Figure 10). MTS policies and their updates can be initiated by both planes, the MCS control plane or the MTS control plane. This mode assumes the existence of dedicated protocol planes between the MCS control plane and the UPAs (C/U_mcs), as well as between the MTS control plane and the TNs' DPNs (C/D_mts). The MCS control plane makes use of the MCS-specific C/U_mcs protocol to establish and maintain UPA states for the delivery of mobile users' traffic through the radio access network. The MTS Controller makes use of a suitable control-/data plane protocol C/D_mts to establish and maintain forwarding and traffic engineering states in the transport network's DPNs for the controlled steering of mobile users' data plane traffic between their current UPF(s) and service data network(s). A single mobile user may have concurrently assigned one or multiple UPAs and get served by a single or multiple data networks. Hence, routing for a mobile device's traffic that applies to the forwarding plane (fwd) between a UPA endpoint (d1) and a relevant first hop DPN endpoint (d2) is configured through the associated control interfaces of the MCS control plane towards the UPAs, and the MTS control plane towards the DPNs. The static assignment and configuration of one or multiple first hop DPNs to one or multiple UPAs is yet another option, while data plane forwarding between a pair of (d1)(d2) endpoints can follow static rules or any other, more dynamic algorithm. In such case the C_mts interface is then used to align end-to-end traffic steering and to configure mobile user traffic paths between a first hop DPN and the mobile data network. Details about a choice and implementation of the C/U_mcs and C/D_mts are out of scope of this specification.

While this document focuses on the C_mts semantics and an associated information model, REST and RPC are suitable enablers to implement operations on the C_mts interface in between the MCS control plane and a MTS controller. Other protocols or publish-subscribe frameworks may apply as an alternative.

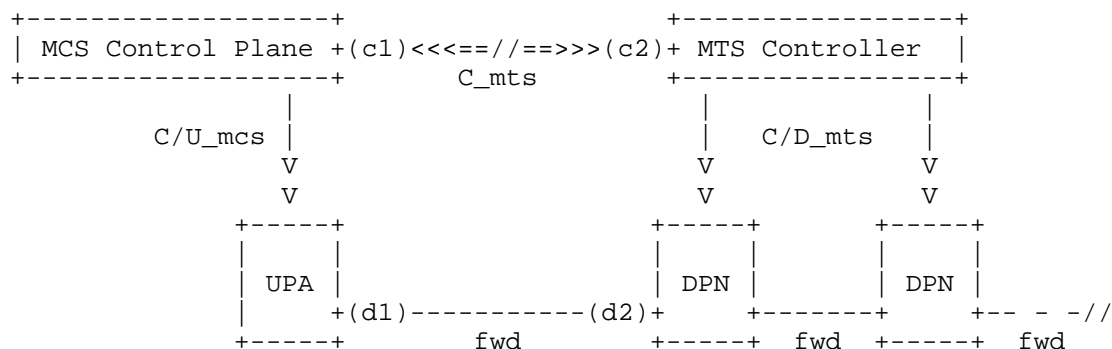


Figure 10: Use of dedicated MTS Controller that interacts with the MCS control plane for mobile traffic steering

Key operations in this mode include the following:

O The MTS Controller indicates to the MCS control plane that traffic associated with a single or a group of users will be assigned a new DPN and associated (d2) endpoint. In case further on-path DPNs will change on the traffic's end-to-end path up to the data network, the MTS Controller will enforce updated routing policies to the TN's DPN. These path changes are kept transparent to the MCS control plane and the UPA. The MTS Control will provide single DPN or a group of DPNs to the MCS control plane. In any case the MCS control plane may decide to keep the current UPA or change the UPA in alignment with the single or group of potential DPNs. In case the MTS control plane provides a group of potential DPNs with its indication, the MCS control plane can choose for the group of candidate DPNs. The MCS control plane provides details one the selected UPA and DPN back to the MTS Controller.

O The MCS control plane indicates to the MTS control that traffic associated with a single or a group of users will be assigned a new UPA, e.g. due to user mobility, load balancing or failover reasons. In case the MCS control plane identifies a particular UPA with the indication, the MTS Controller can either keep the current DPN or assign a new DPN to the associated traffic. Based on the identified UPA, the MTS Controller may interact with a backend system, resulting in a decision to change the data network associated with the single or group of users' traffic, e.g. due to route- and traffic latency, costs or load balancing reasons. The MTS Controller may compute a new path based on the newly assigned UPA and data network and provide information about the first hop DPN, to be used by the UPA, to the MCS control plane. Based on

the identified DPN, the MCS control plane may re-select a more suitable UPA and lets the MTS Controller know about this decision for final configuration of forwarding policies.

6.3. Mode II Operation - Decentralized Control Plane

This operational mode leverages the C_mts interfaces and semantics in between a UPA and DPN to make the DPN propagating MTS routing policies and updates into the TN (Figure 11). While the UPA receives MTS policies from the MCS control plane through the protocol applying to the C/U_mcs reference point, the DPN relies entirely on the C_mts interfaces with a UPA to receive or retrieve relevant information to propagate routes into the TN.

While this document focuses on the C_mts semantics and an associated information model, different protocols can apply in between the (c1) and (c2) endpoints of a UPA and a DPN. REST is a suitable enabler to implement operations in between a UPA and a DPN, while a particular routing protocol, such as RIP or BGP, can also be considered as suitable feed of C_mts semantics from a UPA to a DPN, leveraging a UPA's routing capabilities and possibly available support of basic routing protocols.

As mentioned before, in case of re-using existing protocols that apply in between a UPA and a DPN, additional information and semantics per this specification is desired for the routes announced by the UPAs to the DPNs. Any existing routing protocol with an extension to pass the additional information can be used, e.g., RIP/BGP. APIs provided by the DPN can also be used for the UPA to program the routes with the additional information into the DPN.

Note that the UPA running a routing protocol does not mean it needs to be a full-function router, and BGP is a reasonable choice to carry both the routes and the additional information as attributes. While BGP can support many advanced features with many of its extensions, only basic BGP functionality is needed here and that has become a commodity.

The distributed UPAs may want the TN to rate-limit or shape the DL traffic so that they are not overwhelmed by the traffic. This is much like that in the centralized UPA case the UPA applies traffic control. Therefore, traffic characteristics information may be distributed along with the routes as route attributes, as described in the "QoS Handling" section of [I-D.zhang-dmm-mup-evolution].

When a mobile user's UPA changes (which is more often with the distributed UPAs), packet re-ordering may happen and TN-triggered End-Marker can help mitigate that, as described in

6.4. Mode III Operation - Hybrid

This operational mode leverages the C_mts interfaces and semantics on both (c1) and (c2) endpoints, in between the MCS control plane and the MTS controller, as well as in between an applying UPA and DPN (Figure 12). While the dedicated control plane between MCS and a MTS controller is used to leverage the rich features enabled per Mode I operations, this mode does not assume the enforcement of MTS policies and their updates directly from the MTS controller, but leverages Mode II operations to receive or retrieve MTS policies and their updates through the C_mts interface that applies on the (c1) and (c2) endpoints between a UPA and a DPN and to propagate updated routes towards the TN by means of a routing protocol.

The high-level framework for a hybrid mode operation does not differ from Figure 5 but integrating both the dedicated TN controller (i.e., for centralized framework) and the exposure of the routes via routing protocol(s) between UPA & DPN (i.e., for distributed deployment). Similar to that in Section 5.2, the TN or MTS controller in transport network interfaces directly with the MCS control plane in Mobile network for configurations and provisioning as the dedicated control mode. While on the other aspect similar to Section 5.3, the (edge) DPNs in transport network couple closely with UPAs in mobile network to exchange routes via routing protocols (e.g., RIP, BGP, etc.) as the distributed framework. Here the traffic steering between a UPA in MN and a DPN in TN may be influenced by the combined policies of both the dedicated channel (i.e., MCS and MTS control) and the distributed logic (i.e., tight coupling of UPA & DPN).

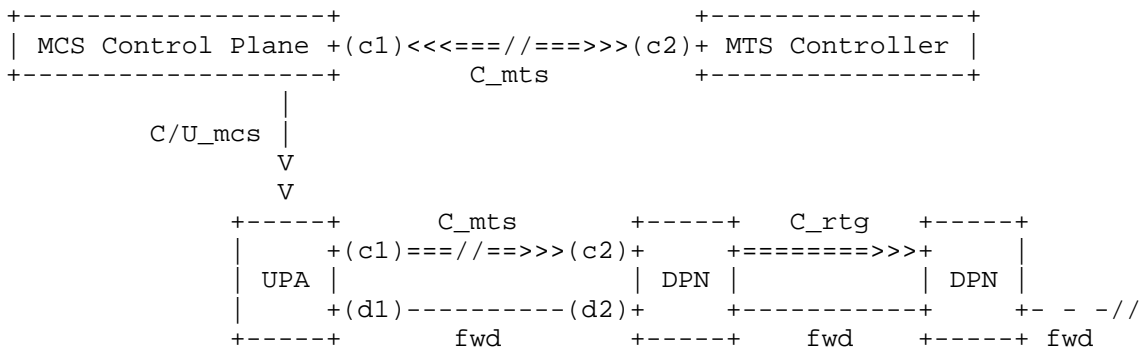


Figure 12: Use of MTS Controller to interact with the MCS control plane and of a routing protocols to propagate MTS routes in the transport network

7. IANA Considerations

This document has no request to IANA.

8. Security Considerations

TBD.

9. Acknowledgments

This document includes results from discussion with various IETF and non-IETF delegates. Many thanks to David Lake for his immediate interest in this topic and feedback at various opportunities and through different channels. Many thanks also to Joel Halpern for his clear and apt feedback during the public side meeting@IETF118 and to Sri Gundavelli and Satoru Matsushima for their support and advice on this matter. The authors want to thank John Kaippallimalil and Carlos Bernardos for the continuous discussion about this draft and in particular for their detailed review and constructive comments regarding advanced versions of this document.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

10.2. Informative References

- [I-D.zzhang-dmm-mup-evolution] Zhang, Z. J., Patel, K., Contreras, L. M., Islam, K., Mutikainen, J., Jiang, T., Jalil, L., Sejati, O. P., and S. Zadok, "Mobile User Plane Evolution", Work in Progress, Internet-Draft, draft-zzhang-dmm-mup-evolution-09, 2 July 2024, <<https://datatracker.ietf.org/doc/html/draft-zzhang-dmm-mup-evolution-09>>.

[I-D.zzhang-dmm-5gdn-end-marker]

Zhang, Z. J., Liebsch, M., and T. Jiang, "End Marker in 5G Data Network", Work in Progress, Internet-Draft, draft-zzhang-dmm-5gdn-end-marker-01, 6 February 2024, <<https://datatracker.ietf.org/doc/html/draft-zzhang-dmm-5gdn-end-marker-01>>.

[TS.29.561]

"3GPP TS 29.561: Interworking between 5G Network and external Data Networks; Stage 3", 3GPP TS 29.561, December 2024.

Appendix A. Information Models

In a deployment with a dedicated control plane, the information model applies to the reference points between the MCS Control Plane and the MTS Controller function, which builds on top of a Network Controller. This reference point is extracted from Figure 5 and depicted with suitable labels in Figure 7

In a decentralized deployment, the information model applies to the reference points between a MCS's UPA and an associated DPN. This reference points is extracted from Figure 6 and depicted in Figure 8.

Besides mobile device identifiers, MCS-specific session identifiers and traffic classifiers, the following Information Elements (IE) are considered relevant for mobile traffic steering:

IE	Description	Note about Use/Format
Node Information		
UPA_ID	Identifier of the UPA	String or numerical ID
UPA_SUPPL_URI	API to retrieve supplementary information about UPA	Source of suppl. info, e.g. map service, mobility pattern, etc.
UPA_GEO_LOC	UPA geographic locator	
UPA_TOPO_LOC	UPA topological locator	
UPA_TOPO_MAP	Reference to topology map	Can be map of physical or virtual topology
UPA_ROLE	Role={current, target, candidate}	

UPA_ND_STATUS	Node status info incl., load, expected changes	
UPA_CTRL_URI	Node controller URI	
Interface Information		
UPA_IF_ID	Identifier of a UPA Interface	
UPA_IF_TYPE	Interface type, i.e. control- / user plane, or controller	
UPA_IF_URI	URI of the identified Interface	
UPA_IF_IP	IP address of a UPA Interface	UPA Data Plane IP to be used by DPNs
UPA_IF_MAC	Layer2 address of a UPA Interface	UPA Data Plane MAC to be used by DPNs
UPA_IF_LABEL	Label for user plane traffic forwarding	Label or segment ID, port, type/format
UPA_IF_PORT	Port for control plane operations	Listen port for protocol or IPC operations

Figure 13: UPA Information Elements

IE	Description	Note about Use/Format
MTS Context	Identifies context of transaction	Associates context of multiple handshakes operation

Figure 14: MTS General Information Elements

IE	Description	Note about Use/Format
DPN_ID	Identifier of the DPN	String or numerical ID
DPN_CP_URI	Control Plane API for the identified DPN	MCS control plane access API or DPN control API
DPN_SUPPL_URI	API to retrieve supplementary information about DPN	Source of suppl. info, e.g. map service, mobility pattern, etc.
DPN_FWD_IP	IP address of a DPN	DPN next hop IP to be used by UPAs and DPNs
DPN_FWD_MAC	MAC address of a DPN	DPN next hop MAC to be used by UPAs and DPNs
DPN_FWD_ID	Lower layer ID for Forwarding	Label or segment ID, type/format
DPN_GEO_LOC	DPN geographic locator	
DPN_TOPO_LOC	DPN topological locator	
DPN_TOPO_MAP	Reference to topology map	Can be map of physical or virtual topology
DPN_ADJ_TYPE	Adjacency={UPA adj.,UPA attached, DN adj. }	Candidate router for UPA, DN/PE next hop
DPN_STAT_INF	Status info incl. load, expected changes	
DPN_IF_METRIC		Eases selection n case of multiple candidates

Figure 15: DPN Information Elements

IE	Description	Note about Use/Format
DN_ID	Identifier of the DN	String or numerical ID
DN_DPN_ID	Access to DN, e.g. PE	Reference to AP of DN, type DN_ID
DN_GEO_LOC	DN geographic locator	
DN_TOPO_LOC	DN topological locator	
DN_TOPO_MAP	Reference to topology map	Can be map of physical or virtual topology
DN_CAND_LIST	Candidate list of alternative DNs	List of DN_IDs
DN_TARGET	Target DN	Single DN_ID

Figure 16: DN Information Elements

Appendix B. Exemplary Application of MTS to a 5G System

The three architecture principles and deployment options as laid out in Section 5, i.e., the dedicated, the decentralized, and the hybrid framework, can perfectly fit into the 3GPP 5GS.

The Figure 17 shows the brief interaction between a 5GS (i.e., so-named MCS in the draft) and TN & DN. The 5GS control-plane or '5GS CP' interacts with the external control function, i.e., the TN controller in the TN or transport network. Application servers or ASFs are located in the (remote) DN or data network that is connected to the 5GS user-plane, e.g., UPF off the N6 interface, via the (intermediate) TN. Here, UPF acts as UPA. The N6 reference point is between the UPA (UPF) and the external IP network (or TN).

5GS CP: also the MCS controller
 C-/L- PSA: Central/Local PSA UPF (i.e., UPA)
 I-UPF: Intermediate UPF (also, UPA)
 TN: Transport Network DN: Data Network
 ASF: App Function in Data Domain (DN)

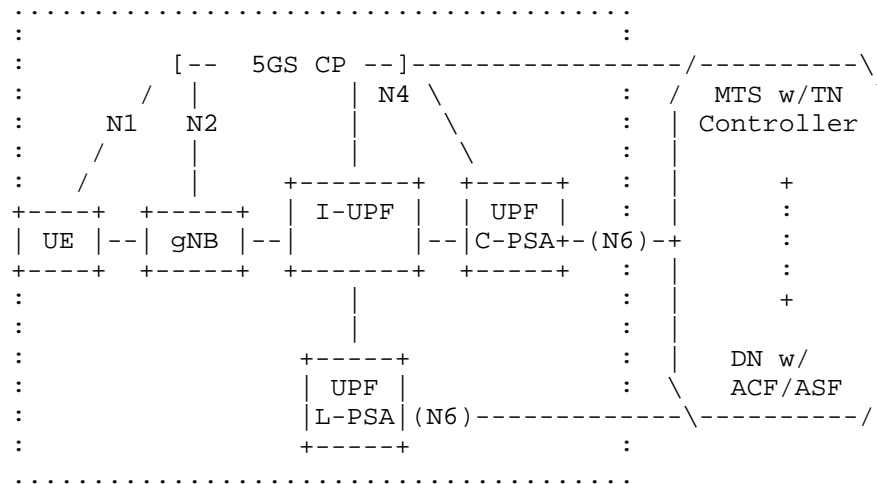


Figure 17: 5GS Architecture w/ DN & TN

The 3GPP document TS 29.561 [TS.29.561] specifies a UPF (i.e., UPA) can be seen as a normal IP router from the external IP network's point of view. Therefore, the control & steering of mobile traffic between UPA(s) in MCS and ASF(s) in DN via TN can be categorized based on the three options as in Section 5:

- * Dedicated CP: The 5GS CP (i.e., MCS controller), the TN/MTS controller and the ACF together determine intelligently the traffic steering from a UPA to an ASF. The 5G AF-provided session provisioning or control can be bucketed into the type. For example, the 5G Rel-19 eEdge has an enhancement to leverage the N6 path delay for (DN) ASF selection, through the interactions between the 5GS-CP (MCS) and the external controller (MTS).
- * Decentralized CP: Running native routing protocols. That is, as per [TS.29.561], a UPA acts as an IP router to engage & peer with the DPN PE(s) in the TN directly. The mobile traffic is steered natively over the transport network.

- * Hybrid CP: It is a combination of both the dedicated CP and the decentralized CP modes. One example is the mobile traffic steering based on the N6 point-to-point tunnel (e.g., PMIPv6/GRE, L2TP, etc.) between a UPA (i.e., UPF) and an ASF. The UPA acts as the forwarding node that might acquire its routing intelligence from multiple sources. One source could be from the interactions between the 5GS-CP (or the MCS controller) and the (external MTS) TN controller -- the dedicated mode. The second source could be based on the routing protocols operating natively between the UPA and the TN PEs -- the decentralized mode. Further, the third source could be the overlay transfer of compute information from the remote (DN) ACF (to the UPA). -- the dedicated mode.

Authors' Addresses

Marco Liebsch
NEC Laboratories Europe GmbH
Kurfuersten-Anlage 36
D-69115 Heidelberg
Germany
Email: marco.liebsch@neclab.eu

Jari Mutikainen
NTT Docomo
Email: mutikainen@docomolab-euro.com

Zhaohui Zhang
Juniper Networks
Email: zzhang@juniper.net

Tianji Jiang
China Mobile
Email: tianjijiang2012@gmail.com