Enabling Service Oriented Principles for the 5G User Plane

ABSTRACT

This paper presents the architectural considerations of integrating the non-IP-based service routing solution, Name-based Routing, to the 5G user plane. While entirely preserving the control plane procedures on the terminal, the carefully crafted out considerations argue for a new Session Management Function functionality and the transitioning of N4 to Nupf. Furthermore, this paper presents User Plane Function provisioning procedures based on Software-defined Networking principles mitigating the need for any manual management procedures and enabling a cloud native orchestration of all 5G Core Network Functions. A planned multi-vendor 5G Core trial is described demonstrating the benefits of a service routing-enabled 5G user plane.

CCS CONCEPTS

• **Networks** → **Network architectures**; • **Computer systems organization** → **Architectures**.

KEYWORDS

5G, User Plane, Service Routing, Service-based Architecture, Software-defined Networking, Name-based Routing

ACM Reference Format:

. 2021. Enabling Service Oriented Principles for the 5G User Plane. In *Proceedings of 8th ACM Conference on Information-Centric Networking (ACMICN'21).* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

In today's 3GPP Systems, including the latest releases introducing 5G technology, there is a great deal of customization and pre-configuration. This pre-configuration and customization are intended to increase system efficiency, to e.g., execute node level Operations, Administration, Maintenance (OAM) procedures to define virtual links rather than allowing the system to discover system components as needed, and route messages using cloud-based principles. Thus, one can find examples where Access and Mobility Management Functions (AMFs) and NG-RAN nodes are setup through a manual configuration procedure, e.g., linking a specific Next Generation NodeB (gNB) with a specific AMF entity. The provisioning of 5G Core (5GC) Network Functions (NFs), in particular the Session Management Function (SMF) and User Plane Function (UPF), are manually conducted. Unfortunately, this is done at the expense of flexibility and extensibility.

Furthermore, 3GPP relies on the establishment of "data pipes" anchored to specific network entities, such as the UPF, which once set up, must be used to transport data from the User Equipment (UE) to the edge of the network or from UE to UE as is the case in 5GLAN type systems. These "data pipes" are referred to as Packet Data Unit (PDU) Sessions, are controlled by the 5GC which is composed of a set of dedicated NFs, responsible for establishing, maintaining, tearing down PDU Sessions, according to operator polices.

The integration of a service routing approach, such as Name-based Routing (NbR), to the user plane requires a carefully designed approach that a) allows the integration in the first instance and b) respects as many "Release 16 PDU session and node procedures" as possible to pave the way for a successful standardisation.

The remainder of the paper is structured as followed: Section 2 presents the technical background that enables the service oriented principles for the 5G user plane which are presented in Section 3. Before concluding the paper in Section 5, the planned trial demonstrating the advances presented in this paper is described.

2 TECHNICAL BACKGROUND

This subsection provides the technology background required for understanding the problem statement and innovative steps described.

2.1 5G System Architecture

Figure 1 provides the system architecture of a 5G system, according to Release 16, Version 16.7 [2]. The most relevant components and interfaces for the described inventive steps are as follows: A 5GC is composed of a set of NFs that offer Service-based Interfaces (SBIs) for a RESTful communication between consumers and producers. All NFs with a circle next to the component's interface line represent such SBI

ACMICN'21, July 2021, Paris, France

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.
https://doi.org/10.1145/nnnnnnn.nnnnnnn

with the ability to utilise the Service Communication Proxy (SCP) for service routing capabilities between consumers and producers. Most notably, even though the UPF is a 5GC NF, the interface to communicate with an UPF, i.e. N4, is not service-based and the SMF and UPF cannot utilise the SCP to reach each other.

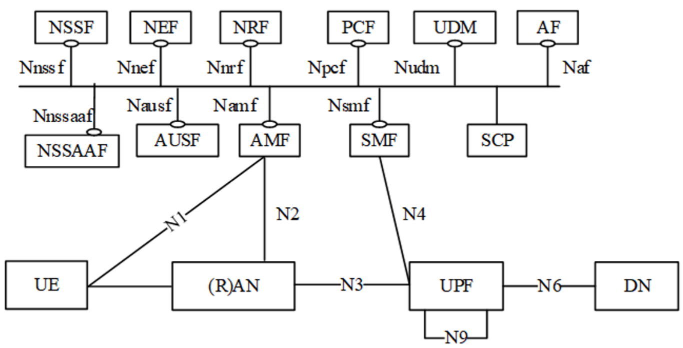

Figure 1: 5G system architecture

In January 2020 the 3GPP Technical Specification Group (TSG) Service and System Aspects (SA) Working Group 2 (WG2) removed the study item on Nupf (transitioning of N4 into an SBI) as part of Release 17 due to other higher prioritised items [1].

2.2 Name-based Routing

NbR is a routing technology that utilises Information-centric Networking (ICN) semantics to perform service routing for Hypertext Transfer Protocol (HTTP) endpoint while supporting all IP-based communication [6, 9]. The system architecture of NbR is provided in Figure 2 and depicts NbR components in orange, namely:

- Service Proxy (SP): Protocol translation from IP into ICN and vice versa, as described in [9]
- Path Computation Element (PCE): Performing Publish/Subscribe rendezvous and path calculation tasks for inter SP communication following the ICN semantics, as described in [9]
- Service Proxy Manager (SPM): a logically centralised component that manages the configurations for SPs and includes an interface to register/unregister Fully Qualified Domain Name (FQDN)-based service endpoints

Software-defined Networking (SDN) components are:

- SDN Controller: NbR, based on Bloom Filter-based forwarding identifiers, seamlessly integrates SDN controlling functions following OpenFlow 1.3 protocol [8] and supports the controllers OpenDaylight, Floodlight and ONOS. The PCE component acts as a northbound application to the SDN controller for topology management and rules injection purposes
- SDN Switch: NbR semantically overloads the IP header with its (Bloom Filter-based) Forwarding Identifiers (FIDs) for a path-based packet forwarding. The SDN controller uses OpenFlow 1.3 to configure the bit mask matching fields in the SDN switches as opposed to longest prefix matching (i.e. IP). For more information on the SDN integration of arbitrary bit mask matching, readers are referred to [8].

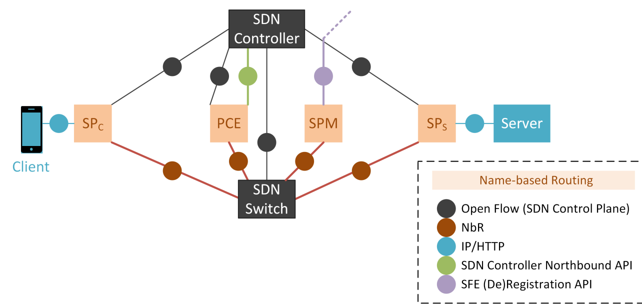

Figure 2: Name-based Routing architecture.

IP-only user plane components:

- Client: An IP-only endpoint that initiates an IP transaction.
- Server: An IP-only service endpoint that is listening on a transport layer port to handle incoming transactions by replying with another IP transaction.

The ability to perform service routing for HTTP services with the support for Transport Layer Security (TLS) as well as any other IP-based communication is available as an Internet Engineering Task Force (IETF) draft within the ICN Research Group (ICNRG) [4]. In a nutshell, NbR re-introduces multicast packet delivery behaviour for HTTP responses as well as switching of HTTP endpoints transparently and seamlessly at any point in the HTTP communication.

2.3 SMF and UPF Provisioning

3GPP allows the operation of a single UPF as well as distributed UPFs which are "located close to or at the Access Network (AN) site" [2, Clause 6.3.3.2]. In both scenarios, the UPFs are manually configured for N4 communication with an SMF (and vice versa), where either the UPF or the SMF initiates the association setup [3, Clause 4.4.3]. Alternatively, UPFs can be manually added to the Network Repository Function (NRF) which offers an Nnrf interface for an SMF to retrieve the list of available UPFs. If more than one UPF is available to an SMF, the UPF selection procedures for a specific PDU session are described in [2, Clause 6.3.3] and is done as part of a PDU session establishment procedure within the SMF.

3 ARCHITECTURE

This section introduces the architectural changes required to integrate an NbR-based user plane into a 5GC. Two integration modes are presented, infrastructure mode and UE mode, which differ from each other by excluding or including the UE in the NbR semantics.

3.1 Base Assumptions

Before the two integration modes are presented, a set of assumptions must be stated which form the basis for the architectural choices described hereafter:

- In order to leverage the true potential of NbR, it is assumed that the user plane is operated with highly distributed UPFs (see Section 2.3) forming some sort of mesh/star/tree network topology.
- The PDU session establishment, modification and release procedures on all nodes but the SMF remain untouched allowing the continuous usability of existing 5G handsets without any modification in their modem.
- Related to the first bullet point, no changes to Network Access Stratum (NAS) procedures by the UE and gNB are permitted to ease the standardisation and integration potential.
- Any additional NbR-related control plane interactions between UEs and 5GC to perform service routing is kept in-band over already established PDU sessions.
- UPF-SMF's N4 interface is upgraded to service-based semantics, i.e., HTTP/2 as the application layer protocol to enable the possibility to utilise the SCP to communicate with other NFs.

3.2 Infrastructure Mode

The proposed infrastructure mode allows UPFs to communicate via NbR procedures among each other, i.e., N9, while preserving the 3GPP standardised system and procedures for N2, N3 and N19.

3.2.1 System Architecture - Infrastructure Mode. In order to achieve the integration of NbR into the 3GPP user plane, the two NbR management components PCE and SPM are exposed to the 5GC, as illustrated in 3. For the PCE, which is responsible for rendezvous and FID calculation functionality, it is proposed to include this as part of the SMF via a dedicated SBI, Nsmf_NbR. This allows any consumer to utilise the functionality provided by the PCE, including other consumers that realise SMF functionality not covered by the PCE. Also, the architecture figure illustrates the N4 interface to become service based, Nupf, allowing any consumer to communicate with the UPF. For the integration of NbR the interface Nupf_NbR is proposed with dedicated Nupf primitives enabling the integration of Name-based Routing.

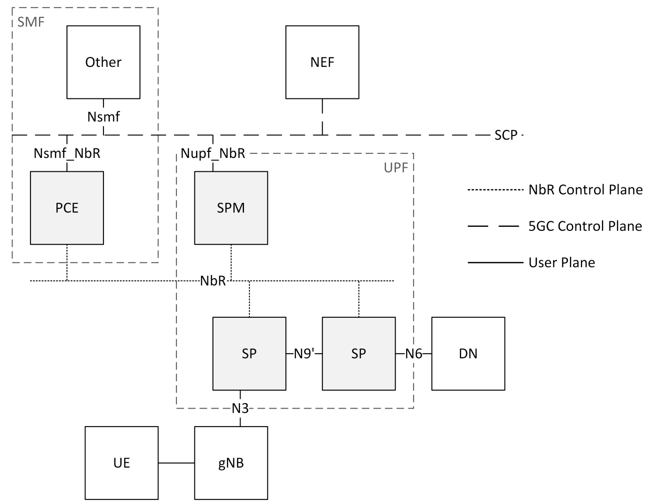

Figure 3: Name-based Routing architecture.

The reason for the transition from N4 to Nupf is twofold: With 5GCs being orchestrated (automated deployment) in a cloud native fashion, any post-orchestration configuration of point to point endpoints (i.e., SMF and UPF on N4) becomes obsolete and the provisioning and communication towards the UPF follows the same cloud native principles as the other 5GC NFs. Secondly, Nupf allows more than one NF to act as a consumer and communicate with the UPF, finally removing any hard binding between two NFs from the 3GPP 5GC architecture.

Inter-UPF communication over N9 is defined as an IP-based communication in Release 16, but with the introduction of NbR for this interface, the architecture figure annotates this interface as N9′, as NbR operates on top of 802.3 with a semantically overloaded IP header (IPv4 or IPv6).

For the NbR infrastructure mode, Figure 4 illustrates the protocol stack with NbR depicted as a layer on top of 802.3 for the N9′ interface. The payload originating from the UE or Data Network (DN) can be of type IP or 802.3 and are handled according to the NbR specifications for HTTP and non-HTTP IP-based communication [9].

3.2.2 Session Establishment - Infrastructure Mode. This section describes the procedures for the establishment of a PDU session over an NbR-based user plane in infrastructure mode as well as the actual uplink and downlink data exchange. Figure 5 provides the message sequence chart for the procedures. Note that the steps in this section cover both IP-based and 802.3-based PDU session types and starts with the communication of a session establishment request by an SMF

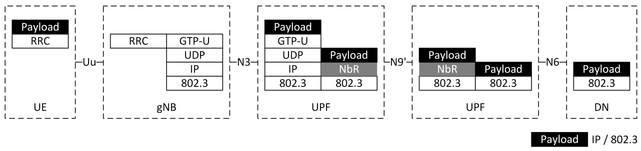

Figure 4: Name-based Routing user plane protocol stack for infrastructure mode for both Packet Data Unit session types IP and 802.3.

either as part of a proactive UPF configuration or as part of a PDU session establishment procedure.

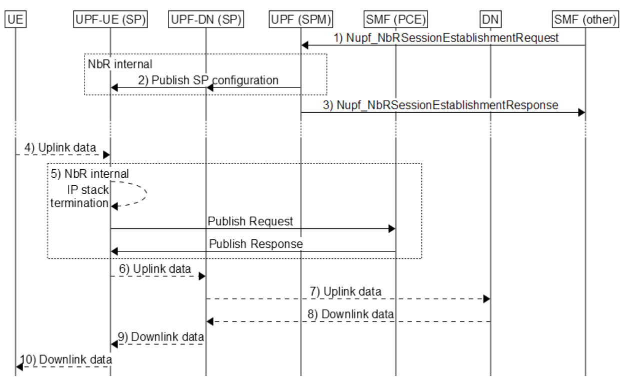

Figure 5: Session establishment for NbR-based UPFs operating in UE mode.

(1) The SMF communicates the Packet Detection Rules (PDRs) for the UPFs involved using the *Nupf_NbRSessionEstablishmentRequest()* primitive and covers the information provided by the *N4_SessionEstablishmentRequest()* primitive, as described in Release 16. The difference is that the consumer, e.g. SMF, must specify which distributed UPF is receiving which set of PDRs utilising the information obtained during the provisioning procedures. This information is known to the SMF from the provisioning procedures described in Section 3.4, where the UPF properties define the NbR mode an UPF implements (infrastructure or UE).
(2) Using the NbR-internal signalling pub/sub system, the Nupf information for each UPF is communicated to all SPs.
(3) The UPF (SPM) informs the consumer (SMF in the chart) about the status of the session establishment request from Step 1 using the *Nupf_NbRSessionEstablishmentResponse()* primitive.
(4) Following Rel.16 procedures, the UE eventually receives the confirmation that the PDU session has been established and is being able to send data over the user plane which arrives at the UPF-UE component.
(5) The UPF-UE then applies the Name-based Routing methods and procedures to translate IP into ICN, as disclosed in [9] depending on the traffic that arrives, i.e., TLS, HTTP or any other IP-based communication. As part of these procedures, the UPF-UE communicates over NbR with the SMF (PCE) in order to perform the ICN procedures, i.e., rendezvous and path calculation, informing the UPF-UE about the UPF-DN which will be able to serve the request received by the UE.
(6) The UPF-UE sends the uplink data to the UPF-DN where it is translated, transparently for both IP endpoints, into a standard IP-based communication.
(7) The uplink data is being sent to the DN where a vertical application will process the request.
(8) Upon the generation of the response to the received request by the UE, the vertical application in the DN issues the response using a standard IP-based communication stack.
(9) The response (aka downlink data) is translated into ICN by the UPN-DN and sent over an L2 switching fabric to the UPF-UE.
(10) The UPF-UE then translates the downlink data back into a standard IP-based communication and sends it to the UE via a standard N3 realisation.

3.3 User Equipment Mode

Figure 6 illustrates the system architecture for a 5GC integration based on the infrastructure mode with only two changes:

- The procedures for NbR over an Ethernet PDU session are disclosed in [4] and describes the SP residing on the UE where IP traffic is being translated into ICN and vice versa. The SP then communicates over a standard 802.3 frame header with the UPF. As part of the NbR header in the payload of the Ethernet header, [4] describes the additional fields it would add when another UE is the destination.
- While moving the SP functionality into the UE, the path-based forwarding approach NbR is based on is not extended to the UE directly using Bloomfilters. Instead, the UPF to UE communication follows an extended Ethernet header indicating which which SP is being addressed as the destination [4]. While the procedures for requesting an FID reside within the SP, the communication between UEs and UPFs are not part of this FID. Therefore, a Service Proxy Forwarder (SPF) is illustrated in Figure 6 as part of the UPF and performs UPF functionalities according to the N3 interface specification. The SPF component implements the counterpart of the UE procedures disclosed in [4].

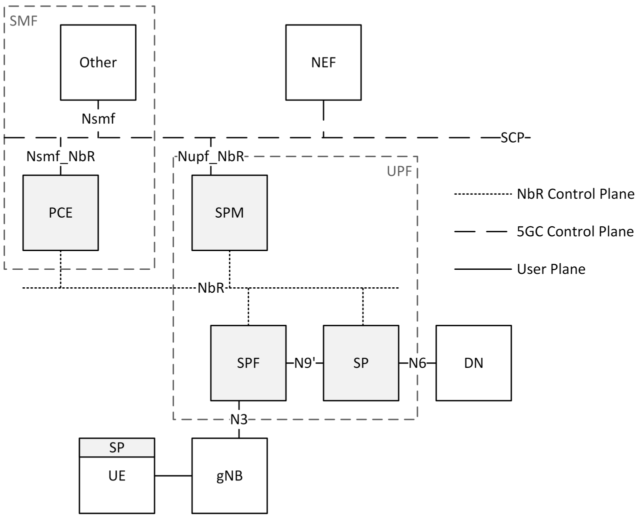

Figure 6: Name-based Routing architecture.

In summary, the architectural changes only affect the UE and UPF communicating with a gNB. All other components and interfaces remain identical to the infrastructure mode described above. However, while [7] describes the signalling for NbR internal pub/sub semantics as an extension to the PDU session establishment via NAS, the procedures and methods proposed in this work follow an "in-band" NbR control plane communication over an established PDU session (user plane) in order to enable the support for Release 16 compliant UEs and gNBs. Similar to the infrastructure mode, UPFs that implement either the SP or SPF communicate with each other over 802.3 and the interface is therefore annotated as N9′.

Figure 7 illustrates the user plane protocol stack for the UE mode. With the NbR layer extended to the UE, the UE mode only supports the PDU session type Ethernet. The payload can be any IP-based protocol with NbR offering special service routing capabilities for HTTP (including TLS-based HTTP communication).

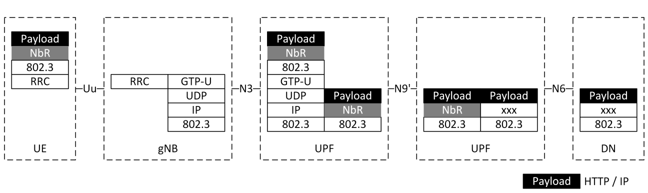

Figure 7: Protocol stack for 5G user plane in Name-based Routing UE mode.

3.3.1 Session Establishment - UE Mode. In principle, the session establishment for the UE mode follows the same assumptions as in the infrastructure mode, i.e., UPF configurations for N3 are communicated via the Nupf interface by SMF while the required pub/sub communication for finding the most appropriate service endpoint that will receive the packet is communicated in an in-band fashion.

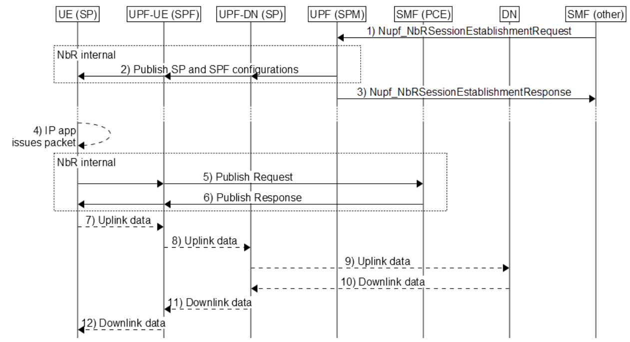

Figure 8: Session establishment for NbR UPFs operating in UE mode.

(1) The SMF communicates the PDRs for the UPFs involved using the *Nupf_NbRSessionEstablishmentRequest()* primitive and covers the information provided by the *N4_SessionEstablishmentRequest()* primitive, as described in Release 16. The difference is that the consumer (SMF) must specify which distributed UPF is receiving which set of PDRs utilising the information obtained during the provisioning procedures. In the UE mode, the UPF facing the gNB via N3 is the SPF that requires the MAC address of the UE for which a PDU session will be/is being established.
(2) Using the NbR-internal pub/sub signalling system, the N4 information for each SP and SPF is communicated by the SPM. For SPs located on UEs, the SPF acts as the forwarder for information related to NbR control plane procedures, e.g., where to find the PCE (in terms of the FID) for future publish requests. As the PDU session as not been configured in the gNB or UE at this moment of time, the SPF will not be able to reach the SP on the UE yet. It is foreseen that the SPF periodically communicates the SPM information to the UE until it explicitly acknowledges it. In another embodiment, the SP on the UE awaits the PDU session to be established and only then communicates with the SPF via a broadcast 802.3 frame that requires any SPM information which the SPF holds.
(3) The UPF (SPM) responds to the consumer that requested the session establishment with an *Nupf_NbRSessionEstablishmentResponse()* primitive.
(4) Once the application on the UE issues an IP-based communication, the SP - located on the UE - transparently intercepts any transaction according to the steps described in [9].

(5) The SP issues a publish request towards the PCE using the Ethernet PDU session towards the UPF-UE (SPF) with the FID provided by the SPM. The UPF-UE (SPF) then forwards the request to the SMF-PCE over the dedicated NbR signalling plane.
(6) The SMF-PCE response with the decision about the request. If a subscriber exists for the IP address or FQDN, the response comprises the information the SP on the UE requires to reach the destination.
(7) The SP on the UE now sends the uplink data via the Ethernet PDU session towards the UPF following the protocol stack illustrated in Figure 7.
(8) The UPF serving the UE (UPF-UE (SPF)) forwards the packet to the UPF that serves the DN where the IP service endpoint resides that can handle the transaction by the UE.
(9) The SP in the UPF-DN translates the payload received by the UE back to an IP-based communication and sends the transaction to the DN.
(10) Upon processing the request, the IP-based application inside the DN responds with an IP-based transaction which is sent to the UPF-DN.
(11) The UPF-DN then uses the methods and procedures described in [9] to transparently translate the IP-based traffic into an ICN communication. Following the protocol stack illustrated in Figure 11, the downlink data is sent to the UPF-UE.
(12) The SPF in the UPF-UE now applies the methods and procedures described in [4] to reach the UE over an Ethernet PDU session.
(13) Upon reception of an NbR communication, the SP inside the UE translates the ICN communication transparently back into an IP-based communication.

3.4 SMF and UPF Provisioning

The provisioning of SMF and UPF is foreseen as part of the automated deployment of a 5GC (aka orchestration) through an external technology such as OpenShift or Kubernetes integrated into ETSI Management and Orchestration (MANO) frameworks such as OpenSource Mano (OSM) or Open Network Automation Platform (ONAP). Current procedures to provision SMF and UPFs are of rather static nature even when involving the NRF for these purposes. With the introduction of Service-based Architecture (SBA) and the advances of cloud native software design and engineering[10], such static provisioning must be avoided. As the targeted distributed UPF scenario requires topology management procedures (i.e., bootstrapping, link discovery and link failure detection), SDN (e.g. OpenFlow) is leveraged for this purpose. Figure 9 illustrates the procedures for UPF topology management with the integrated NbR system architecture provided in this section.

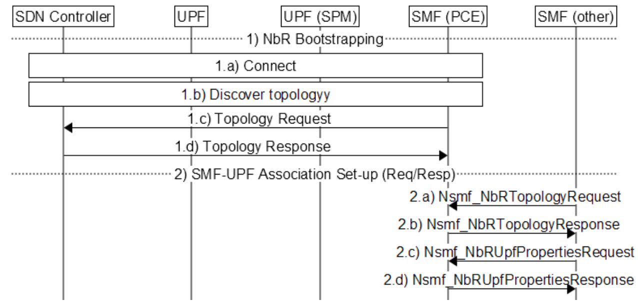

Figure 9: SDN-based Session Management Function and User Plane Function provisioning.

(1) The steps here allow all NbR components to join an SDN switching fabric and to form a topology:
1.a) All components that are attached to the NbR signalling plane (see Figure 3 and Figure 6) leverage SDN for their communication among each other and connect to the SDN controller in this step.
1.b) As part of standard SDN practices, the SDN controller requests all connected switches to discover their link-local neighbours via the Link Local Discovery Protocol (LLDP).
1.c) The PCE component periodically queries the SDN controller via its northbound Application Programming Interface (API).
1.d) The SDN controller responds with the known topology which includes switch identifiers and its neighbours among other information. As northbound APIs for SDN controllers are not standardised the information provided for each switch can vary but always includes switch identifier and its neighbours.
(2) The steps described here demonstrate the message exchange for 5GC consumers with the PCE-enabled SMF and the newly proposed SBI Nsmf_NbR. The steps hereafter describe a standard request response communication between a consumer and the SMF as the producer.
2.a) A consumer, e.g., another decomposed SMF service, requests an updated topology of UPFs using the *Nsmf_NbRTopologyRequest()* primitive.
2.b) The SMF producer serving the Nsmf_NbR primitives responds with an *Nsmf_NbRTopologyResponse()* which has the full topology. Note, only UPFs that implement the SP functionality and serve either N3, N6 or N9 interfaces will be returned in this primitive, as the NbR components PCE and SPM do not process user plane packets and therefore require no PDRs.

2.c) The consumer then requests the properties of a particular switch using the newly defined primitive *Nsmf_NbRUpfPropertiesRequest()*.
2.d) The producer responds to the request with an *Nsmf_NbRUpfPropertiesResponse()* primitive which provides various properties for a particular switch. Most importantly, the UPF properties provide the consumer with the information which routing technology a particular UPF implements.

Upon completing the steps outlined above, the SMF has all information required to determine PDR for future session requests by UEs.

3.5 Vertical Application Registration

NbR offers transparent service routing capabilities for stateless protocols, e.g. HTTP, allowing the switch of service endpoints (server) without affecting the requesting endpoint (client). In order to achieve that, the NbR layer - executing the routing decisions as an UPF - requires the information where service endpoints are located. Figure 10 presents the message sequence chart of how a vertical application located in an DN is registered against NbR-based UPFs.

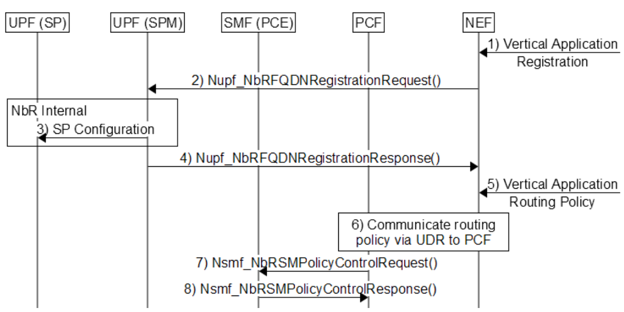

Figure 10: Vertical application registration.

(1) The Network Exposure Function (NEF) receives the information about the existence of a new application in a particular DN following standard Release 16 procedures.
(2) The NEF now uses the newly designed Nupf interface to communicate the existence of the new vertical application to the UPF using the *Nupf_NbRFQDNRegistrationRequest()* primitive.
(3) The UPF (SPM) distributes the information across all SPs that implement the user plane packet routing functionality. This is exchanged over the NbR internal pub/sub-based control plane.
(4) Upon distributing the registration information, the UPF (SPM) confirms the status of the registration to the consumer, i.e., NEF, using the *Nupf_NbRFQDNRegistrationResponse()* primitive.
(5) Following Release 16 specification[2], the NEF receives the request for specific routing policy instructions (e.g. from an Application Function (AF)).
(6) Together with the User Data Repository (UDR) and Unified Data Management (UDM), the Policy Control Function (PCF) establishes the desired routing policy.
(7) The outcome of the desired routing policy is communicated to the SMF (PCE) via the primitive *Nsmf_NbRSMPolicyControlRequest()*. The PCE allows to configure a range of service routing policies such as shortest path or weighted round robin.
(8) The SMF (PCE) consumer responds with an *Nsmf_NbRSMPolicyControlResponse()* indicating the processing state of the policy request.

4 SYSTEM AT TRIAL

In order to foster the adoption of the architectural changes proposed in this paper, the authors are going to realise the UPF and SMF changes described in the previous sections as part of one of the five customer-facing trials of a European multi-organisational project.

The narrative of the use-case is a media showroom with video content consumed by a range of different devices, such as handheld devices, Virtual Reality (VR) headsets and large monitors. The topology of the trial is illustrated in Figure 11 and depicts a single gNB and two DNs. The DN at the bottom is closer to the UEs then the DN at the right when comparing the number of hops. As mentioned, each DN has the same video application (and content) and serves all video via HTTP (either HTTP Live Streaming (HLS) or Dynamic Streaming over HTTP (DASH)). A vertical application orchestrator chartered with the life-cycle management of the two video application instances determines at which point in time the video application instances are spawn up and made available in the user plane. As explained in Section 3.5, the newly introduced *Nsmf_NbR* interface offers the programming of the service routing policy on the user plane, e.g., shortest path or weighted round robin.

As the UE implementation of NbR is available for Android-based 5G-enabled handsets only, the media showroom use case will see both architectural modes being realised allowing any 5G UE to join the video stream over NbR integrated over the user plane. Independently from which architectural mode is being utilised, Opportunistic Multicast (OMC) will be always observable between the UPF serving the DN and the UPF serving the gNB. Moreover, any lifecycle management change of the video application instances, e.g., shutting down the instance on the shorter hop count, will demonstrate the transparent service routing capabilities of NbR to seamlessly switch to the video application instances in the DN further away from the UEs.

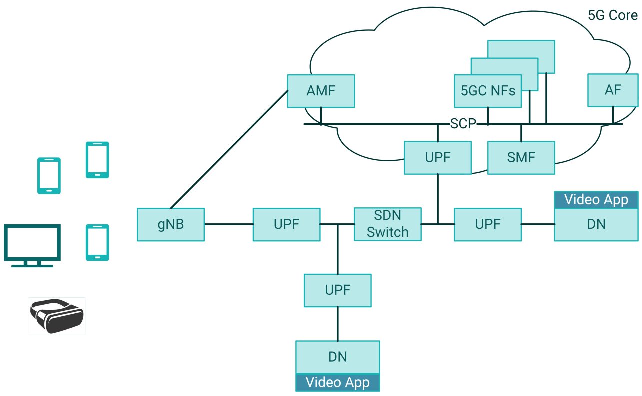

Figure 11: Deployment topology of media showroom validation trial.

The trial is going to showcase a multi-vendor 5GC with UPFs provided by Vendor A and all other NFs by Vendor B apart from the SMF. The SMF itself is going to be decomposed into Vendor B's SMF implementing Release 16 while Vendor A will provide the additionally required NbR functionality, i.e. PCE, including the newly proposed *Nsmf_NbR* interface. Moreover, in the trial all 5GC NFs instances, including the SBI-enabled UPF, utilise the SCP for any consumer-producer communication. The same service routing technology on the user plane which is presented in this paper powers the SCP, as it is already one of the three SCP deployment options in 3GPP [2].

5 CONCLUSION

This paper presented the integration proposal of NbR for the 5G user plane, a transparent service routing approach based on ICN principles. This has been achieved with minimal changes to the currently standardised 3GPP control plane procedures. A new SMF functionality has been proposed together with new Nsmf primitives disintegrating SMF functionality into individual microservices which are accessible via standardised SBIs. While trialling the presented approach at a showcasing event as part of a European 5G project, the standardisation ambition is to put up the N4 to Nupf transition as a study item in 3GPP and to standardise the UPF provisioning methods and newly introduced Nsmf_NbR and Nupf_NbR SBIs. Furthermore, the standardisation of NbR within the ICNRG is seen as a key driver for the adoption of this transaprent service routing approach.

REFERENCES

[1] 3rd Generation Partnership Project. [n.d.]. Draft Report of TSG SA meeting #86. https://www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_86/Report/Report_SA_86_Draft_v005.zip

[2] 3rd Generation Partnership Project. 2020. 23.501: Technical Specification Group Services and System Aspects; System architecture for the 5G System (5GS); Stage 2 (Release 16). Version 16.7.0.

[3] 3rd Generation Partnership Project. 2020. 23.502: Technical Specification Group Services and System Aspects; Procedures for the 5G System (5GS); Stage 2 (Release 16). Version 16.7.0.

[4] Dirk Trossen, Sebastian Robitzsch. 2020. Internet Services over ICN in 5G LAN Environments. https://datatracker.ietf.org/doc/draft-trossen-icnrg-internet-icn-5glan/

[5] FUDGE-5G Consortium. 2020. FUDGE 5G: FUlly DisinteGrated private nEtworks for 5G verticals. https://fudge-5g.eu

[6] K. Haensge, D. Trossen, S. Robitzsch, M. Boniface, and S. Phillips. 2019. Cloud-Native 5G Service Delivery Platform. In *2019 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*. 1–7. https://doi.org/10.1109/NFV-SDN47374.2019.9040042

[7] R. Ravindran, P. Suthar, D. Trossen, C. Wang, and G. White. 2019. Enabling ICN in 3GPP's 5G NextGen Core Architecture. https://tools.ietf.org/html/draft-ravi-icnrg-5gc-icn-04

[8] M. J. Reed, M. Al-Naday, N. Thomos, D. Trossen, G. Petropoulos, and S. Spirou. 2016. Stateless multicast switching in software defined networks. In *2016 IEEE International Conference on Communications (ICC)*. 1–7. https://doi.org/10.1109/ICC.2016.7511036

[9] Dirk Trossen, Sebastian Robitzsch, Scott Hergenhan, Janne Riihijärvi, Martin J. Reed, and Mays F. Al-Naday. 2019. Service-based Routing at the Edge. *CoRR* abs/1907.01293 (2019). arXiv:1907.01293 http://arxiv.org/abs/1907.01293

[10] Adam Wiggins. [n.d.]. The Twelve Factor App. https://12factor.net