Operator Platform: Requirements and Architecture

Version 11.0

19 February 2026

Security Classification: Non-Confidential

Copyright Notice

Disclaimer

Compliance Notice

The information contain herein is in full compliance with the GSMA Antitrust Compliance Policy.

This Permanent Reference Document is classified by GSMA as an Industry Specification, as such it has been developed and is maintained by GSMA in accordance with the provisions set out GSMA AA.35 - Procedures for Industry Specifications.

Table of Contents

1 Introduction

1.1 Overview

Operators in the 5G era have a significant opportunity to monetise the capabilities of their networks. Moreover, with the existing relationships that Operators have with enterprises, their vast local footprint, their ability to support digital sovereignty principles and their competence to provide high-reliability services, the missing piece is the ability to package and expose their networks in a scalable fashion across multiple Operators. The Operator Platform (OP) concept, as introduced in this document, described the architecture of a generic platform to fill this gap, identifying technical requirements, functional blocks and interface characteristics.

1.1.1 Relationship to existing standards

1.1.1.1 3GPP

Unless otherwise stated, the requirements listed in this document are based on the open and published 3GPP specifications listed in Section 1.5. 3GPP Release 17 is taken as the basis.

1.2 Scope and objectives

This document covers requirements and architecture specifications that guide the entire industry ecosystem to define a common solution for exposing network capabilities and edge compute resources, referred to as Operator Platform. The ecosystem includes Operators, vendors, OEMs, and service providers.

This document covers the following areas:

- Requirements to enable Operator Platform-based service exposure
- Architecture, functional levels and components
    - **Reference architecture for enabling Service Exposure**: Definition of modular architecture suitable for the exposure of capabilities.
    - **Federation:** Enable federation between Operator Platform instances allowing an Application Provider using an Operator Platform to access also the capabilities and resources exposed by the Operator Platform instances that have federated with that platform.
    - **Reference interfaces**: Definition of interconnection for the end-to-end service, between service providers to End-Users, network elements and federated platforms. This document focuses on Northbound, Southbound, East/West (i.e. Operator Platform Federation), and User to Network interfaces as a first approach.
    - **Mobility**: Network and terminal integration should allow Service Continuity of the exposed capabilities against End-User mobility in the home and visited networks.

This document provides a target architecture and requirements to enable an end-to-end delivery chain for different services. The interaction of the entire ecosystem involved in the application delivery should be covered, including Application Providers, Aggregators, Operators, and the Subscribers.

1.3 Definitions

Term	Description
Aggregation Platform	A platform through which the Aggregator offers the services.
Aggregator	An actor who provides (or combines) services exposed by different Operators and exposes them for use to the Application Providers. Note: Exposed services by the Aggregator may differ from the services provided by the Operators. Synonyms: Channel Partner, Hyperscaler(one possible role)
Alternative QoS Performance Profiles	A prioritised list of alternative Quality of Service (QoS) profiles which refers to set of QoS parameters e.g., bit rate, packet delay budget etc. which OP should use in case specific QoS Performance Profiles requested by the Application Provider cannot be met.
API Catalogue	A repository that organises and provides detailed information about Operator offered APIs. That are documented, categorized, and version-controlled, making it easy for developers to find, understand, and integrate with API efficiently
Application	A software implementation that supports services to End-Users. An Application can consist of multiple architectural parts (e.g. an Application Client running on the served End-User's UE). Some of those parts might be deployed and managed through an OP.
Application Client	A specifically developed client component of an application.
Application Edge Part	An architectural part of an Edge Application that is to be deployed on edge compute Cloudlets. An End-to-End Application may include multiple Application Edge Parts (e.g. microservices).
Application Identifier (Application ID)	Identification of an Application (owned by an Application Provider) towards the End-User as part of the Consent capture
Application Identity Profile	Application information that describes the Application that an Application Provider (via the Aggregator) wishes to register for Operator-offered Service API consumption.
Application Instance	An instantiation of an Application Edge Part on a Cloudlet.
Application Privacy Profile	Information held within the CSP domain about subscribed legal usage and required scopes and purposes as well as additional application specific privacy considerations. This information must consider all the local regulations in place. It may also contain information related to the person entitled to consent access to privacy-sensitive data.
Application Provider	The provider of the Application that accesses an OP. An Application Provider may be part of a larger organisation, like an enterprise, enterprise customer of an OP, or be an independent entity. Synonym: Developer
Application Provider Vetting Information	Subset of the Vetting Information provided by the Application Provider during Application Provider onboarding which needs to be vetted, as required by local Operator needs and local regulations.
Application Provider Vetting Proof	A proof issued by a verification service indicating that the Application Provider Vetting Information has been verified.

Term	Description
Application Vetting Information	Subset of the Vetting Information provided by the Application Provider during Application onboarding which needs to be vetted, as required by local Operator needs and local regulations.
Application Vetting Proof	A proof issued by a verification service indicating that the Application Vetting Information has been verified.
Availability Zone	An OP Availability Zone is the equivalent of an Availability Zone on Public Cloud. An Availability Zone is the lowest level of abstraction exposed to an Application Provider who wants to deploy an Application on Edge Cloud. Availability Zones exist within a Region. Availability Zones in the same Region have anti-affinity between them in terms of their underlying resources - this ensures that in general terms, when an Application Provider is given a choice of Availability Zones in a Region, they are not coupled which ensures separation and resilience.
Callback	An endpoint defined by the Application Provider in a service request to receive event notifications from the OP.
Certificate Authority	An entity that issues digital certificates.
Cloudlet	A point of presence for the Edge Cloud. It is the point where Edge Applications are deployed. A Cloudlet offers a set of resources at a particular location (either geographically or within a network) that would provide a similar set of network performance.
Communication Service	a service that enables transmission and receipt of information between two or more points/entities
Communication Service Provider	an entity that provides Communication Services. Designs, builds and operates its Communication Services. The provided Communication Service can be built with or without a Network Slice.
Consent	The agreement of a Subscriber to allow the usage of their personal data. This agreement can be revoked at any time.
Consent Capture	Process through which a Subscriber grants permission to the OP to share certain personal data with an Application for processing under a defined Purpose of Data Processing.
Converged Charging System	The element within the Operator’s Network that allows to do the real time rating and charging for the services that are provided by that Operator [10]
Data Collection Interval	A common interval for data reporting that should be negotiated to facilitate federation.
Data Protection	Legal control over access to and use of data stored in computers.
East/Westbound Interface	The interface between instances of an OP that extends an Operator's reach beyond their footprint and Subscriber base.
Ecosystem Party	In the context of GSMA OP, [OP] Ecosystem Party represents either an Application Provider, Aggregator, Partner OP or corresponding synonyms.
Edge Application	An Application whose architecture includes one or more Application Edge Parts (e.g. microservices) Note: an Application doesn’t necessarily need to be an Edge Application to use capabilities exposed by an OP.

Term	Description
Edge Cloud (EC)	Cloud-like capabilities located at the network edge including, from the Application Provider's perspective, access to elastically allocated compute, data storage and Network Resources. In the context of this document, the Edge Cloud management function /domain is accessed through the Operator Platform. The phrase "located at the infrastructure edge" is not intended to define where an Operator deploys its Edge Cloud. The Edge Cloud is expected to be closer (for example, latency, geolocation, etc.) to the Application Clients than today's centralised data centres, but not on the User Equipment, and could be in the last mile network. Note: This definition is based on that in "Open glossary of edge computing", v2.0 [1].
Edge Cloud Resources	In the context of this document, resources of the Edge Cloud Service whose management function / domain is accessed through the Operator Platform Southbound Interface – Cloud Resources (SBI-CR).
Edge Interconnection Network (EIN)	A direct and dynamically managed (optionally pre-existing) logical interface between two EC instances. The interface shall use existing network infrastructure for connection establishment, and may have security and rules applied based on the EC and OP requirements.
Edge Node	A resource in a physical data centre. The term Edge Node used in context with the Edge Node Sharing refers to the compute resources offered by the Partner OP to the Leading OP. The Leading OP may use such resources to serve its own End-Users in scenarios such as not having the Edge Clouds footprint in locations where the End-Users requesting access to edge services but a Partner OP is offering Edge Cloud Resources in those locations. [9]
Edge Resource	Sum of compute, network, and storage capabilities made available for workload deployment and processing in Edge Nodes.
Edge Site	A physical location where an Edge Node is deployed.
EIN Establishment	A procedure to create an EIN connection between two ECs. The process also establishes application and traffic security and filtering rules on the EIN as part of the establishment.
End-User	A human participant who uses the Application. A customer of the Application Provider. Note: The End-User is not always the Subscriber.
Home OP	The Operator Platform instance belonging to the Subscriber's Operator; that is, whose PLMN identity (Mobile Country Code, MCC, and Mobile Network Code, MNC) matches with the MCC and MNC of the Subscriber's International Mobile Subscriber Identity (IMSI), as defined in 3GPP TS 23.122 [6]. Note: the use of this term relates to Subscriber roaming and is used to differentiate from the Visited OP of the roaming Subscriber. Note: non-SIM devices are for further study
Leading OP	The Operator Platform instance connected to the Application or Aggregation/Marketplace Platform and receiving the Application's requests, sharing them to the selected federated Partner OPs.

Term	Description
	Note: the use of this term relates to federation and is used in relation to the Partner OPs with which the Leading OP is federated. Note: this can include use in the context of Subscriber roaming where a Leading OP acts as Home OP for a given Subscriber and needs to interact with a Partner OP serving the visited network to which the Subscriber is connected to deliver the service requested by the Application.
Local Breakout	Edge Cloud Services are provided to a roamed UE by the Visited OP, rather than by the Home OP
Marketplace Platform	A platform where services (and APIs) are published and offered to 3rd parties (e.g. Application Providers).
Network Communication Service	The external representation of a Communication Service orderable by an Application Provider
Network Resource Location	The Network Resource Location is how near to the edge or the centre of the network an Application Edge Part is instantiated and Cloud resources are consumed. Whilst typically, an OP deploys an Application Edge Part on a Cloudlet at the edge of the network, it may choose to deploy it, for example, at a Region level or centrally (but within the OP). An OP decides on the Network Resource Locations.
Network Resources	In the context of this document, the network services and capabilities provided by the Operator whose management function /domain is accessed through the Operator Platform Southbound Interface – Network Resources (SBI-NR).
Network Slice	A logical network that provides specific network capabilities and network characteristics [4]
Network Slice Instance	A set of Network Function instances and the required resources (e.g. compute, storage and networking resources) which form a deployed Network Slice [4]
Network Subscription	An entry in the network user database that regulates authentication and authorization (including policy aspects) of a user trying to access network services. It includes a mapping from its General Public Subscription Identifiers (GPSI) to SUbscription Permanent Identifier (SUPI), only a single SUPI or IMSI/Mobile Subscriber Integrated Services Digital Network Number (MSISDN, in case 3G/4G access) can be active at a time for the Network Subscription.
Northbound Interface	The interface that exposes an Operator Platform to Application Providers
Operate API	APIs to be used for the business management of APIs exposed by the GSMA Operator Platform on its NBI. Example TM Forum OGW Operate API
Operator	In the context of GSMA OP, an Operator is an entity that exposes capabilities and/or resources of their network (IT, mobile) to Application Providers, provides connectivity to User Equipment and has an Operator Platform. Synonyms: CSP (Communication Service Provider), MNO (Mobile Network Operator)

Term	Description
Operator Capability	Any capability that can support services that an Operator may offer to Applications allowing them to enhance the service that they provide to their End-Users.
Operator Platform	An Operator Platform (OP) facilitates access to the Edge Cloud and other capabilities of an Operator or federation of Operators and Partners. It follows the architectural and technical principles defined in this document.
Partner	An entity or other party that offers and provides a service or resource, in the context of the Operator Platform's federation, to other Partners. Each Partner hosts an OP and offers the resources through its East/Westbound Interface (E/WBI) federation. For example, a Partner can be an Operator that provides network, Subscribers and cloud services or a hyperscaler / cloud provider that offers cloud services only.
Partner OP	An Operator Platform that federates with another Operator Platform. It uses the E/WBI to offer its OP Capabilities-based services to the other Operator Platforms that could offer them to the Applications for which they act as Leading OP.
Privacy Information	Information held within the CSP domain used for keeping evidence/records of the lawfulness of privacy-sensitive data processing and sharing (including user consent data). The capture of this information is performed according to the Application Privacy Profile.
Privacy Management	Service within the CSP domain supporting management of the Application-related Privacy Information. The service supports also notifying (to the interest parties) when Privacy Information has changed.
Purpose of Data Processing	Reason for which processing personal information is required by the Application (owned by an Application Provider). It declares what the Application intends to do with a set of personal information resources. Each Purpose of Data Processing maps to a set of Scopes (usually defined in an API specification).
QoS Performance Profile	Information about QoS network capabilities that can be delivered by the operator network for APIs where QoS is relevant.
Region	An OP Region is equivalent to a Region on a public cloud. The higher construct in the hierarchy exposed to an Application Provider who wishes to deploy an Application on the Edge Cloud and broadly represents a geography. A Region typically contains one or multiple Availability Zones. A Region exists within an Edge Cloud.
Representational Consistency	Representational Consistency means that the information elements that the Application Provider exchanges with an Edge Cloud do not change as a function of the Partner OP with which it is ultimately interacting. This implies that a function of the Exposure functional level is to provide a consistent information model.
Service Continuity	The uninterrupted user experience of a service, including in those cases where the Internet Protocol (IP) address or anchoring point change
Service APIs	APIs abstracting Operator services exposed for use by Applications or Aggregation/Marketplace/Enterprise Platforms. Service APIs are defined by CAMARA.

Term	Description
Session Continuity	The continuity of a Protocol Data Unit (PDU) Session. For PDU Session of the IPv4 or IPv6 or IPv4v6 types, "Session Continuity" implies that the IP address is preserved for the lifetime of the PDU Session
Southbound Interface	Connects an OP with the specific Operator infrastructure that delivers the network, cloud and charging services and capabilities.
Subscriber	A client/customer of the Operator, identified by a unique identifier. The Subscriber is usually also the End-User, but this is not always the case. For example, a parent may be the Subscriber of a mobile subscription for their child (End-User)
User Client	Functionality that manages on the user's side the interaction with an OP. The User Client represents an endpoint of the User to Network Interface (UNI) and is a component on the User Equipment. Note: Different implementations are possible, for example, Operating System component, separate application software component, software library, Software Development Kit (SDK) toolkit and so on.
User Equipment (UE)	Any device with a SIM used directly by an End-User to communicate. User Clients and Application Clients are deployed on the User Equipment. By default, the term “UE” means UE with the explicit SIM-based Telecom wireless network connectivity throughout the document.
User Identity Token	A security token that serves as proof of authentication, confirming that a user is successfully authenticated with the parameters provided at token generation.
User Identity Token Manager	An OP (external) function that generates User Identity Tokens for UE Application Clients and verifies the User Identity Token when received from the AP through the Northbound Interface (NBI).
User Mobility	The process of a Subscriber moving from one place to another that may trigger changes in their network connectivity
User to Network Interface	Enables the User Client hosted in the User Equipment to communicate with an OP.
Vetting Information	Information provided by the Application Provider during Application Provider and Application onboarding as required by local Operator needs and local regulations.
Vetting Process	Business process (e.g. part of TM Forum Business Partner Agreement Management) whereby an Aggregator and / or Operator verify Vetting Information.
Visited OP	The Operator Platform instance that belongs to the Operator providing access to a roaming Subscriber; that is, whose PLMN identity (MCC and MNC) matches with the MCC and MNC of a roaming Subscriber's current Visited Public Land Mobile Network (VPLMN). Note: the use of this term relates to Subscriber roaming and is used to differentiate from the Home OP of the roaming Subscriber. Note: non-SIM devices and non-3GPP access are for further study

1.4 Abbreviations

Term	Description
3GPP	Third Generation Partnership Project
5G	5th Generation Mobile Network
5GC	5G Core
AAPrM	Authentication, Authorization and Privacy Management
AF	Application Function
API	Application Programming Interface
CAPIF	Common API Framework
CCS	Converged Charging System
CHF	Charging Function
CN	Core Network
CPU	Central Processing Unit
CRUD	Create, Read, Update and Delete
CSP	Communication Service Provider
DDoS	Distributed Denial of Service
DNAI	Data Network Access Identifier
DNN	Data Network Name
DNS	Domain Name System
DoS	Denial of Service
EC	Edge Cloud
EIN	Edge Interconnection Network
E/WBI	East/Westbound Interface
gNB	gNodeB
GPSI	Generic Public Subscription Identifier
GPU	Graphic Processing Unit
GST	Generic network Slice Template
HR	Home Routing
HTTP	HyperText Transfer Protocol
ID	IDentifier
IEC	Immediate Event Charging
IETF	Internet Engineering Task Force
IMSI	International Mobile Subscriber Identity
IP	Internet Protocol
LBO	Local BreakOut
MCC	Mobile Country Code
MNC	Mobile Network Code
MSISDN	Mobile Subscriber Integrated Services Digital Network Number
NAS	Non-Access Stratum

Term	Description
NAT	Network Address Translation
NBI	Northbound Interface
NDS	Network Domain Security
NEF	Network Exposure Function
NSI	Network Slice Instance
NSMF	Network Slicing Management Function
NSSAI	Network Slice Selection Assistance Information
NWDAF	Network Data Analytics Function
OP	Operator Platform
OSS	Operation Support System
PDN	Packet Data Network
PDU	Protocol Data Unit
PEC	Post Event Charging
PGW	PDN (Packet Data Network) GateWay
PII	Personally-Identifiable Information
PRD	(GSMA) Permanent Reference Document
QoE	Quality of Experience
QoS	Quality of Service
SBI	Southbound Interface
SBI-CR	Southbound Interface – Cloud Resources
SBI-NR	Southbound Interface – Network Resources
SBO	Session BreakOut
SCEF	Service Capability Exposure Function
SDK	Software Development Kit
SLA	Service Level Agreement
SLI	Service Level Indicators
S-NSSAI	Single Network Slice Selection Assistance Information
SSC	Session and Service Continuity
SUPI	SUbscription Permanent Identifier
TAC	Tracking Area Code
UE	User Equipment
UNI	User to Network Interface
UPF	User Plane Function
URL	Uniform Resource Locator
URSP	UE Route Selection Policy
VPLMN	Visited Public Land Mobile Network
Wi-Fi	Wireless network protocols, based on the 802.11 standards family published by the IEEE.

1.5 References

Ref	Doc Number	Title
[1]		Open Glossary of Edge Computing, Linux Foundation Edge, https://github.com/State-of-the-Edge/glossary/blob/master/edge-glossary.md
[2]	3GPP TS 29.522	5G System; Network Exposure Function Northbound APIs, Stage 3 https://www.3gpp.org/DynaReport/29522.htm
[3]	3GPP TS 29.122	T8 reference point for Northbound APIs https://www.3gpp.org/DynaReport/29122.htm
[4]	3GPP TS 23.501	System architecture for the 5G System (5GS) https://www.3gpp.org/DynaReport/23501.htm
[5]	GSMA PRD FS.31	Baseline Security Controls, GSM Association Official Document FS.31
[6]	3GPP TS 23.122	Non-Access-Stratum (NAS) functions related to Mobile Station (MS) in idle mode
[7]	IETF RFC 4122	A Universally Unique IDentifier (UUID) URN Namespace https://datatracker.ietf.org/doc/html/rfc4122
[8]	3GPP TS 23.503	Policy and charging control framework for the 5G System https://www.3gpp.org/DynaReport/23503.htm
[9]	GSMA PRD OPG.04	East-Westbound Interface APIs
[10]	3GPP TS 32.240	Charging management; Charging architecture and principles https://www.3gpp.org/DynaReport/32240.htm
[11]	3GPP TS 28.202	Charging management;Network slice management charging in the 5G System (5GS);Stage 2 https://www.3gpp.org/DynaReport/28202.htm
[12]	3GPP TS 32.257	Telecommunication management; Charging management; Edge computing domain charging https://www.3gpp.org/DynaReport/32257.htm
[13]	IETF RFC 2119	“Key words for use in RFCs to Indicate Requirement Levels”, S. Bradner, March 1997. Available at http://www.ietf.org/rfc/rfc2119.txt
[14]	GSMA PRD WA.101	Open Gateway Channel Partner Onboarding Guide
[15]	GDPR	“General Data Protection Regulation”. Available at https://gdpr-info.eu/
[16]	3GPP TR 33.867	Study on user consent for 3GPP services https://www.3gpp.org/DynaReport/33867.htm
[17]	RFC 6749	The OAuth 2.0 Authorization Framework https://datatracker.ietf.org/doc/html/rfc6749
[18]	CIBA Flow	OpenID Connect Client-Initiated Backchannel Authentication Flow - Core 1.0 https://openid.net/specs/openid-client-initiated-backchannel-authentication-core-1_0.html

Ref	Doc Number	Title
[19]	GSMA PRD OPG.09	NBI APIs Realisation in the SBI
[20]	GSMA PRD OPG.10	Open Gateway Technical Realisation Guidelines
[21]	GSMA PRD OPG.11	OP Requirements For Edge Services
[22]	TMF 931	TM Forum Open Gateway Operate API: Onboarding and Ordering User Guide
[23]		CAMARA Consent Info API https://github.com/camaraproject/ConsentInfo
[24]		CAMARA Security and Interoperability Profile https://github.com/camaraproject/IdentityAndConsentManagement/blob/main/documentation/CAMARA-Security-Interoperability.md

Note: Some documents in this list (e.g., [5]) may not be released as public documents.

1.6 Conventions

The key words “must”, “must not”, “required”, “shall”, “shall not”, “should”, “should not”, “recommended”, “may”, and “optional” in this document are to be interpreted as described in RFC2119 [13].

2 High-Level Architectural Requirements

This section defines the requirements that an OP's architecture should fulfil. Section 2.1 defines the high-level requirements for the platform in general that are independent of the services/capabilities exposed. Section 2.2 defines the functionality that the platform should offer to the different parties interacting with it. Section 2.4 provides the security requirements of an OP at a similar service-agnostic level. Finally, section 2.3 provides requirements on how capabilities should be exposed.

Note: The section does not define the architecture itself nor what the platform's interfaces should support. Those aspects are introduced in sections 3 and 4 of this document respectively.

2.1 General

An OP and its architecture shall comply with the following requirements:

1. An OP shall expose supported network-related capabilities to the Application Providers.
2. For each Operator supporting the OP architecture, there shall be an OP instance that has the sole responsibility for managing the resources and services that the OP exposes in that Operator's network.

Note: This instance may be operated by the Operator or be outsourced.

3. The interfaces that an OP instance offers to other parties shall be provided using common definitions based on the requirements in this document.

2.2 Functionality offered to OP Ecosystem Party

2.2.1 Functionality offered to Application Providers

An OP and its architecture shall fulfil the following requirements related to the functionality offered to Application Providers:

1. The OP architecture shall allow an Application Provider to use a common interface to access capabilities exposed by multiple Operators subject to an agreement with the Operators involved.

Note: such an agreement could result in the federation of OPs between involved Operators.

2. An OP shall hide the complexity of the OP architecture, the involved Operator networks and client access to those networks from the Application Providers.
3. There shall be a "separation of concerns" of the OP and the Application Providers, meaning that the Application Providers and OP do not require knowledge of each other's internal workings and implementation details, for instance:

   a) An OP does not expose its internal topology and configuration, physical locations of related resources (see note), internal IP addressing, and real-time knowledge about detailed resource availability.
   b) An OP does not know how the application works (for instance, it does not know about the application's identifiers and credentials).
   c) It is the responsibility of the OP to provide to the Application Providers the required tools for them to operate accordingly to the local regulation.
   d) The OP shall provide to the Application Providers all the required information, e.g. a clear identification of the geographical location of the resources, with a level of granularity that allows the Application Providers to manage their services in a way that is compliant with local regulations.
   e) It is responsibility of the Application Providers to use the tools and the information provided by the OP to operate their applications in compliance with the local regulations.

Note: An OP provides information on the geographical Region(s) where a service is available. The Application Provider provides information sufficient for an OP to process the request and (if accepted) fulfil it.

4. The OP architecture shall allow an Application Provider to monitor the application's usage across the networks where it uses capabilities and/or resources exposed by the OP.
5. The OP architecture shall allow an Application Provider to discover and set the available Performance Profiles for the End-User to select the Performance Profile(s) that suit the application needs.

2.2.2 Functionality offered to Aggregators

An OP needs to expose information to enable an Aggregator to find the Home OP where to route an Application Programming Interface (API) call to on behalf of an Application Provider’s application. The Application Provider’s application typically has the MSISDN, IP

address or a User Identity Token that will be provided to the Aggregator in the API call. Such a User Identity Token should allow to identify both the Network Subscription and the Application using the subscribed network services. The OP architecture shall fulfil the following requirements related to the functionality offered to Aggregators:

1. An Aggregator shall be able to identify the Home OP based on the identifiers in an API request.
2. The OP architecture shall allow an Aggregator to use a common interface to find the target OP with a single API call to identify the target Operator ID.
3. The OP architecture shall enable the Aggregator to discover the End-User's Home OP when roaming.
4. The OP architecture shall enable the Aggregator to identify the target OP for non-subscription-related APIs
5. The OP architecture shall support automated onboarding and lifecycle management of Application Providers and Applications by Aggregators, for provisioning and ongoing operational control.
6. The OP architecture shall support API Catalogue synchronisation / management, so that Aggregator can keep an up-to-date list of services made available.
7. The OP architecture shall support reporting to Aggregators on service usage by Applications.
8. The OP architecture shall support the ordering of Service API product offerings by the Aggregator.

2.2.3 Functionality offered to End-Users/Devices

An OP and its architecture shall fulfil the following requirements related to the functionality offered to End-Users and their devices:

1. The OP shall allow the End-User to benefit from the services exposed by the OP securely, both in their home and in visited networks.

Note: Availability in visited networks is dependent on the Operator’s roaming agreements.

2. The OP, in coordination with the Privacy Management Function, located in the CSP domain, shall provide the means for the End-User to express their Consent for the usage of their sensitive data with respect to the services provided by the OP, in accordance with the local regulation.
3. When presenting information on the Application or Application Provider to the End-User (e.g. while capturing Consent), the OP shall do so based on approved Vetting Information (if applicable).

    a) If the channel through which this information is presented allows it, the OP shall provide an indication to the End-User to confirm this vetting status.

2.2.4 Functionality offered to Operators

An OP and its architecture shall fulfil the following requirements related to the functionality offered to Operators:

1. The OP architecture shall allow an Operator to monitor and track the usage by an OP of their resources and network-related capabilities by Application Providers.
2. The OP architecture shall enable Operators to monitor their Subscribers' usage of Resources and network-related capabilities allocated to an Application Provider in a visited network.
3. The OP architecture shall enable Operators to establish connections between resources allocated to Application Providers, and monitor their usage by the applications for charging purposes.
4. The OP architecture shall allow an Operator to charge for the services and capabilities provided by OP to application providers, Subscribers, and other Operators.
5. If an Operator Platform is part of the Operator's security domain (see Note 2), it shall access the network and other resources through the Southbound Interface (SBI) and any other operating interface.

   Note 1: An Operator may choose to outsource some of its functionality to another party. That external party may need to know some details about the Operator's internal workings, such as its Cloudlets' physical locations. This approach would require an agreement covering commercial, Data protection, security, legal issues, etc.

   Note 2: Security Domains administer and determine the classification of an enclave of network equipment/servers/computers. Networks using different security domains are isolated from each other. Security Domains are managed by a single administrative authority. Within a security domain, the same level of security and usage of security services is typical. For example, a network operated by a single Operator or a single transit provider typically constitutes one security domain, although an Operator may subsection their network into separate sub-networks. See 3GPP TS 33.210 Network Domain Security (NDS); IP network layer security.

6. Where an Operator Platform is not part of an Operator's security domain, there is also a "separation of concerns" of the Operator from the OP. "Separation of concerns" again means that they do not require knowledge of each other's internal workings and implementation details.

   a) It is responsibility of Operators to provide to the OP the required tools for the OP to operate accordingly to the local regulation.
   b) The Operator shall provide all the required information to the OP, e.g. a clear identification of the geographical location of the Edge Resources, with a level of granularity that allows the OP to be compliant with local regulations.
   c) It is the responsibility of the OP to use the tools and the information, provided by the Operator, to offer its services, to its customers, in compliance with the local regulations.

7. The OP shall provide support for O&M capabilities allowing Operator(s) to monitor, measure and analyse the performance and health of the OP.

2.2.5 Functionality offered to other OPs

An OP and its architecture shall fulfil the following requirements related to the functionality offered to other OPs:

1. The OP architecture shall allow an OP to access services offered by another OP (when there is a federation agreement between the OPs) on behalf of the Application Providers that it serves.
2. A federation of independently operated Operator Platforms shall enable additional capabilities, such as:

    a) the continuation of a User Equipment's (UE) use of the services enabled by the OP when moving into a "visited network".

3. The OP architecture shall allow a Leading OP to monitor and track capability and resource usage of an application in Partner OPs .
4. Similarly, there shall be a "separation of concerns" between OPs. In terms of responsibilities:

    a) It is the responsibility of "Partner OP" to provide to the "Leading OP" the required tools for the "Leading OP" to operate accordingly to the local regulation.
    b) The "Partner OP" shall provide to the "Leading OP" all the required information, e.g. a clear identification of the geographical location of the Edge Resources, with a level of granularity that allows the "Leading OP" to be compliant with local regulations
    c) It is responsibility of the "Leading OP" to use the tools and the information, provided by the "Partner OP", to offer its services, to its customers, in compliance with the local regulations.

2.3 Exposure requirements

This section provides the requirements that an OP's architecture should fulfil for the exposure to Application Providers. Section 2.3.1 provides high-level requirements for that exposure and section 2.3.2 defines requirements for the different types of capabilities that could be exposed and their management as well as their usage by an OP. Section 2.3.3 defines the requirements around Privacy Management. Finally, the aspects around mobility of Subscribers for which an OP is managing capabilities are covered in the requirements in section 2.3.4.

2.3.1 High-Level requirements

The following requirements apply for an OP related to the exposure of capabilities:

1. An OP shall allow an Operator to expose capabilities within the Operator or Partner network to applications consuming such capabilities.
2. The OP architecture shall allow an Application Provider to consume capabilities for their end-to-end application.

Note: The capabilities are those of different network domains, such as

- Radio Access Network
- Core Network
- Transport Network (including Intra-Cloudlet network and Inter-Cloudlet network)
- Non-3GPP Access networks
- Operation and Maintenance systems
- Edge Compute resources

3. An OP needs to handle the situation that specific capabilities are not or differently available for all networks or network domains at any time and in any federated network.
4. An OP shall allow an Operator to expose capabilities based on network functions exposure and resources sharing.
5. An OP shall allow an Operator to manage capabilities exposure depending on different factors (type, interface, application provider and Operator).
6. The OP shall support the exposure of capabilities to the federated Partner OPs.

Note: For clarity, availability of a network capability in a visited network is only possible by federation between networks. Roaming on a non-federated Operator's network does not allow exposing the visited network's capabilities.

7. Access to the capabilities in the visited network shall be subject to authorisation by the Home OP and the Visited OP.

2.3.2 Capability Management requirements

2.3.2.1 General principles

An OP shall be able to access capabilities exposed by an Operator network and use those capabilities either to optimise its OP specific services or expose those to the Application Provider.

An OP shall be able to access capabilities exposed by Partner OPs that offer such capabilities (indicated through the catalogue), and use those capabilities either for optimising its own OP-specific services or expose those to the Application Provider.

An OP shall allow the consumption of the capabilities by the Application as applicable.

An OP shall support the exposure of four categories of capabilities:

- Control capabilities: these are capabilities that allow controlling certain behaviours or characteristics of the network, e.g. applying a QoS Performance Profile to a session or Traffic Influence
- Event-based capabilities: these are typically notification-based network services allowing to inform the application upon specific network-related events, e.g. UE reachability
- Transactional information capabilities: these are capabilities allowing to obtain information from the network in a request-response pattern

- Analytics Capabilities: these are capabilities allowing to obtain analytical information from the network through notifications or in response to a request.

2.3.2.2 Control Capabilities

An OP shall be able to expose control capabilities to the Application Provider. Those capabilities may be exposed by matching a declared intent expressed by the application provider for using specific capabilities.

Note: Any control capability includes a specific set of Service Level Indicators (SLIs) used to exchange information on particular levels of a network capability between the Operator network, Partner networks, the respective OPs and the application.

2.3.2.3 Event-based capabilities

An OP shall be able to request to receive notifications about network events. The information elements obtained may be used for other use within the OP, e.g. for the orchestration of Application Instances, or be exposed to the Application Provider and to a Partner OP.

An OP shall provide a generic event notification mechanism for those events exposed to the Application Provider as well as those exposed to a Partner OP.

An OP shall be able to provide as part of the Network Event-Based Capabilities, “status type” information to the Application Provider regarding, e.g. UE mobility, UE communication, UE Statistics, Status Reports or QoS sustainability.

2.3.2.4 Transactional capabilities

An OP shall allow using transactional capabilities for purposes within the OP itself or exposure through appropriate APIs to the Application Provider.

2.3.2.5 Analytics capabilities

An OP shall be able to request network analytics information and to be notified about events.

Note: The information elements obtained may be used within the OP, e.g., for the orchestration of Application Instances, or be exposed to the Application Provider or a Partner OP.

An OP shall provide a generic event notification mechanism for those events exposed to the Application Provider as well as those exposed to a Partner OP.

Note: An OP can have access to the 3GPP Nnwadaf_AnalyticsInfo Network Data Analytics Function (NWDAF) service to obtain through the Network Exposure Function (NEF) a specific analytics request.

2.3.2.6 Operator Capabilities Service lifecycle

An OP shall be able to support the following requirements:

1. An OP shall allow an Application Provider to access Operator Capabilities, directly or through an Aggregator. The Operator Capabilities are accessed via APIs, known as Service APIs, which are consumed by the Application Provider’s applications.

2. An OP shall expose functionalities for Privacy Management to access the Operator Capabilities, when those capabilities require the treatment of personal information of the Operator's Subscribers.
3. An OP shall enable OAM functionalities for an Application Provider or Aggregator to configure, monitor, measure and control the Operator Capabilities. The OAM functionalities are accessed via APIs, known as Operate APIs. OAM capabilities requirements include:

    a) An OP shall allow an Operator to expose the catalogue of products, including the available Operator Capabilities, to an Application Provider or Aggregator.
    b) An OP shall allow an Aggregator to register a new Application provider so that they can access Operator Capabilities.
    c) An OP shall allow an Application Provider or Aggregator to register a new application so that they can access Operator Capabilities.
    d) An OP shall allow an Application Provider or Aggregator to request a specific Operator Capabilities product to be accessed by a specific registered application.
    e) An OP shall allow an Application Provider or Aggregator to access the usage reporting of Operator Capabilities from their registered applications.
    f) An OP shall allow an Application Provider or Aggregator to supervise the status of the Operator Capabilities accessed by their registered applications.
    g) An OP shall allow an Application Provider or Aggregator to report platform fault or issues of the Operator Capabilities accessed by their registered applications.

4. An OP shall provide routing information to the Application Provider or Aggregator for identifying which Operator is responsible for handling the required Operator Capability, as related to the Operator's Subscriber.

2.3.3 Privacy Requirements

An OP shall:

1. Ensure that there is a legal basis under the local privacy regulation (if any) for sharing personal data on Subscribers with an Application (potentially through an Aggregator) for the Purpose of Data Processing indicated by the Application Provider.
2. Ensure that no restriction on data sharing is in place when the Purpose of Data Processing is covered by a legal basis (different from Consent) that is compliant with local regulations.
3. Interface with an external Privacy Management Function to retrieve Application-related Privacy Information (if already available) or if the Subscriber's Consent to share data with a particular Application Provider must be captured.
4. Ensure that authorization (see Note 1) and Consent (if applicable) are obtained prior to an API invocation that would trigger the sharing of the personal data with an Application Provider for a specific Purpose of Data Processing.
5. Support Consent (to share personal data with an Application) being revoked by the Subscriber.
6. Support Legitimate Interest legal basis (when applicable) being objected by the Subscriber (e.g., the Subscriber may object marketing campaigns at any time).

7. Keep records of access to personal data through Logging, Tracing and Auditing functions.
8. Have the technical mechanisms in place to guarantee mutual authentication with an Application Provider so that the Consent Capture can be based on reliable information elements, e.g., Application Provider ID, Application ID, Purpose of Data Processing, etc.
9. Present available APIs, Scopes, Purpose of Data Processing to the Application Provider at any time (and to other OPs in federated environments).
10. If the integration of an external party (e.g., Application Provider) includes more parties, e.g., an Aggregator in-between, the identity of the party triggering the capture of the Consent (if applicable legal basis) shall be the identity of the Application Provider.
11. Have in place the mechanisms for verifying the identity of the Application Provider before capturing the Consent.

Note: In non-zero trust scenarios, the Operator may delegate the Application Provider identity verification to an Aggregator.

12. Leverage the information elements to enable Vetting Process before capturing the Consent. These information elements may include as part of the Vetting Information: Application Provider name, Application Provider identifier, Application Provider logo, address, organisation type, Application name, Application identifier, Application logo, Application type, redirect URL, etc.
    a) Whenever the Vetting Information contains a URL (e.g., when the Application logo is given as a URL), the OP should download the information before triggering the Vetting Process.
13. For End-User facing operations (e.g., Consent capture), the OP's handling shall consider only the latest approved Vetting Information for an Application Provider (see Table 8) and for an Application (see Table 9). Therefore, an update of the Vetting Information (triggered by an Application Provider), should not replace the already approved Vetting Information in the OP until the Vetting Process (triggered by the update) has completed.

Note 1: Whether that authorization is obtained in batch or as part of individually triggered requests depends on the Application logic and similarly what identifiers are used to identify the Subscriber on whom they would obtain the data.

2.3.4 Mobility Requirements

2.3.4.1 High-Level Roaming Requirements

An OP and the OP architecture shall support UE's accessing the service from outside their home Operator's footprint (i.e. roaming Subscribers). For those scenarios, the following applies:

1. UEs, while roaming, shall be able to access resources within the visited network with the specified characteristics (e.g. for access to applications deployed on Edge Resources).

Note 1: This requires Local Breakout (LBO) or Session Breakout (SBO) of the Subscriber's Protocol Data Unit (PDU)/Packet Data Network (PDN) connection to a User Plane Function (UPF)/PDN Gateway (PGW) in the visited network.

Note 2: To allow LBO/SBO for the Subscribers, the agreements between Operators need to be in place. Part of the information shared between the Operators is the APN/ Data Network Name (DNN)/ Network Slice Selection Assistance Information (NSSAI) used for LBO/SBO.

2. Access to resources in the visited network shall be subject to authorisation by the Home OP and the Visited OP.
3. An Application Provider shall be able to indicate whether their application is available to inbound/outbound roaming UEs and, if so, in which networks.

Note 3: Availability of the applications a UE wishes to access is currently assumed to be covered by the federation between networks. Roaming on a non-federated Operator's network is not in scope.

4. If an OP is not available in the visited network or the OP managing the resources in that network is unavailable to the UE (e.g. the required federation or LBO roaming agreements are missing), the UE shall still be routed to the most favourable resources. This would be the resources in the network closest to the UE that provide the required capabilities. Because the visited network cannot provide those, the Subscriber shall be routed to resources in the UE's home network, i.e. the next most favourable location.
5. An Application Provider shall be able to indicate whether access by UEs connected to visited networks to those resources should be possible, given that such access may result in significant increases in latency.

Note 4: Seamless handover from home or visited network to another visited network is not in the scope of the current version of this document.

2.3.4.2 Bearer Change Requirements

The OP shall track the network bearers within the different network domains and ensure that requested capabilities are maintained. E.g. if a session between an Application Client and a UPF over a 5G is using network capabilities exposed to the OP via the NEF and this session is transferred to 4G, the OP shall manage the same capability over the Service Capability Exposure Function (SCEF), if available.

2.4 High-Level Security Requirements

The OP architecture shall comply with the following security requirements:

1. An OP shall expose network capabilities and resources data (e.g., compute and storage) to Application Providers and federated Partner OPs following the 'need-to-know' principle and only for the legitimate scenarios expected in the PRD.
2. An OP shall not expose its configuration data and internal topology (referred to as topology hiding) to Application Providers and federated Partner OPs.

3. An OP shall apply data protection mechanisms to assure data availability, confidentiality, authenticity, and integrity. Data shall be protected both during storage and processing and be transported in a secure way. This means:

   a) protecting the data in transit, via encrypted and integrity protected channels, to prevent data interception and manipulation, as well as to prevent intervening attacks, while also assuring user privacy protection;
   b) in storage and execution, via technological means, e.g., log file or database access controls, trusted enclaves.

4. The permitted data (i.e., data that may be shared on the need-to-know principle as in requirement 1) shall be exposed only to authorised and authenticated entities.
5. An OP shall implement role-based access control for configuring users, with policies defined and enforced, ensuring a secure binding between services and authorised entities.
6. An OP shall adopt an integrity protection mechanism for the various identifiers in use (such as resource IDs, User/Subscriber IDs, session IDs, application IDs).
7. An OP should support TCP proxies to avoid server IP address guessing and TCP connection hijacking.
8. An OP should support flow-control on invoking application services control plane APIs to protect federated services from abuse of these APIs.
9. An OP should support different role-based privileges for such roles as Application Providers and network/infrastructure Operators to control unauthorised access to Network Slice management of shared/virtualised resources.
10. Security mechanisms (e.g., Certificate Authorities) used to protect tracking and logging of information on an OP's different interfaces should be independent of each other.
11. Validate the Application Provider’s ownership and correct format of the callback URL(s) before sending notifications to the Application.

3 OP Architecture and interfaces

This section defines the architectural model for the OP. This model is introduced in section 3.1 with the interfaces covered in section 3.2 and the OP functional levels and components in section 3.3 .

3.1 Architecture

The primary goal of the OP’s architecture is providing a global and common way of exposing an Operator’s services to external Application Providers or Aggregators, whether through a direct connection from the service producer towards the final service consumer or by employing intermediate integration platforms (e.g. Marketplace).

The OP environment hosts multiple business actors who may need to interwork to complete end-to-end service delivery, resource sharing and footprint expansion. This interworking implies defining a common way of enabling actors to interact with each other. The business actors may be:

- Operator: The owner of an OP. The Operator may act as an Aggregator or even an Application Provider.
- Application Provider: The owner of an Application.

- Aggregator: The owner of a Marketplace. The Aggregator may act as an Application Provider.
- End-User: The user of an Application.

To satisfy its goals, an OP shall enforce a multi-layer architecture with functional separation of the requirements presented in Chapter 2. For a system as complex as an OP, a target architecture is needed to localise and inter-relate the requirements. Such a target architecture is presented in this section.

The target architecture is described at a relatively high level. Where OP-specific concepts are presented, they are defined as OP functional components and interfaces. This is done to capture the essential behaviour needed by OPs without constraining the ability of the architecture to conform to prevailing standards or the ability of vendors to innovate.

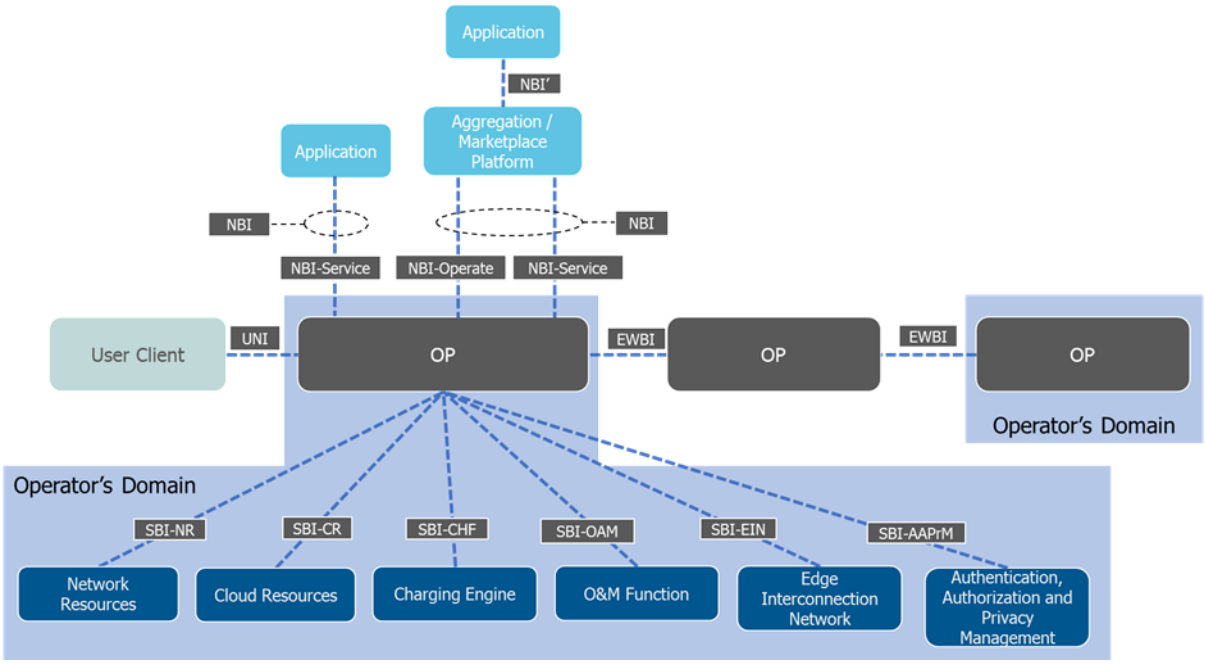

Figure 1: OP Interface Reference Architecture

Note: NBI' is not in the scope of this document.

The Operator's Domain contains all the endpoints for the Southbound Interfaces (SBI) needed for a given service.

GSMA PRD WA.101 [14] describes different models how the service can be exposed and offered to the Application Provider, either directly or via a Marketplace. Figure 1 depicts how an Application can consume capabilities exposed by the OP. The services and APIs exposed by the Aggregator, who owns the Marketplace, may differ from the services provided by the Operators.

The following sections cover the OP functional levels, functional components and interfaces.

3.2 Interfaces

An OP provides the following interfaces:

- Northbound Interface (NBI): to Applications and Aggregation/Market Place Platforms
- East/Westbound Interface (E/WBI): interface between OP instances that extends an Operator's reach beyond their footprint and Subscriber base.
- Southbound Interface (SBI): A set of interfaces that connects an OP with the specific Operator infrastructure that delivers the network, cloud and charging services and Operator Capabilities.
- User to Network Interface (UNI): a service-specific interface that enables the User Client hosted in the UE to communicate with the OP that is providing the Operator Capabilities supporting Application Clients active on the same UE. The interface is optional, depending on whether it is required for the Operator Capabilities that are offered (e.g. Edge Services).

3.3 Functional Levels and Components

3.3.1 General

The OP is realised via multiple functional components (or functions). These functions enable an OP instance to interact with other endpoints in the OP ecosystem, namely other OP instances, the Cloud Resources and the Network Resources and execute scenarios from/towards Application Providers or Aggregators. Figure 2 shows the high-level architecture and functions in an OP.

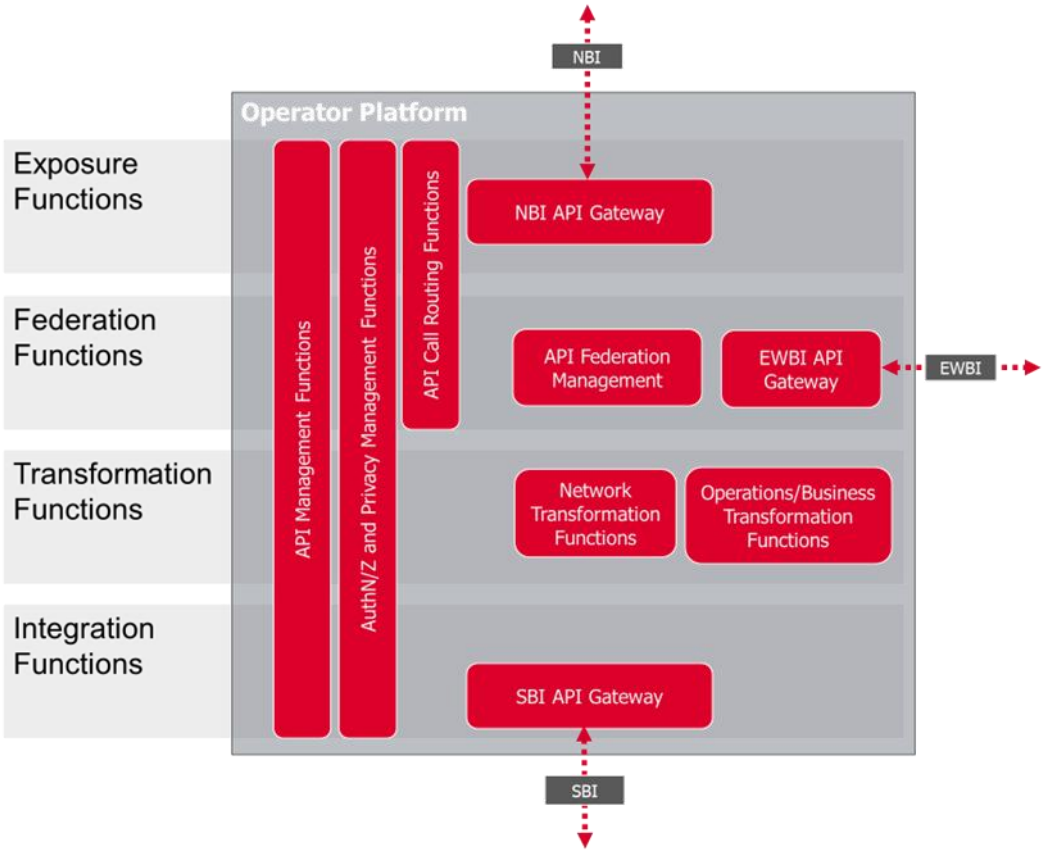

Figure 2: OP Functional levels and components

As shown in Figure 2, the functions can be grouped into four functional levels:

1. Exposure Functions,
2. Federation Functions,
3. Transformation Functions, and
4. Integration Functions.

It is worth mentioning that some “common” functions may span multiple functional levels (see e.g., API Management or API call Routing Functions in Figure 2).

The functional components in Figure 2 may be deployed in a distributed manner (as an architectural pattern that goes beyond monolithic realisations) enabling also flexible functional composition (for instance, if federation is not a scenario to be considered, the Federation-related functions do not need to be deployed)

This section elaborates on the functional levels and components depicted in Figure 2.

3.3.2 Common Functions

This category groups functions that may apply to more than one functional level.

3.3.2.1 API Management Functions

Providing (among others) the following functions:

- API Catalogue
- Application Provider management
- Application Onboarding
- API Subscription management
- API Usage management
- API Monitoring
- API Service Level Agreement (SLA) management
- API Provider management
- API Lifecycle management
- API Access Policy management

3.3.2.2 API Gateway Functions

API Gateway Functions are available in all of the interfaces in Figure 2, namely NBI, EWBI and SBI. They include (among others) the following functions:

- API Registry
- API Access Control / Security enforcement
    - Authentication (see below section 3.3.2.3)
    - Authorisation (see below section 3.3.2.3)
    - Plan control
- API Usage Data Generation
- API Logging and Tracing
- API Metrics Generation
- API Audit Logging

- API Traffic Management
    - Spike arrest
    - Usage throttling / Rate limiting
    - Traffic prioritization
- Interface translation
    - Format translation (e.g., from XML to JSON)
    - Protocol translation (e.g., from SOAP to REST)
- Caching
- Handling NB/SB/EWB interface connectivity aspects
    - For instance, providing Heartbeat/Keep-Alive mechanisms over the interfaces

3.3.2.3 Authentication, Authorization and Privacy Management

Providing (among others) the following:

- Authentication and Authorization.
    - Identity Management (if applicable)
    - User Identity Token management (if applicable)
- Privacy Management (if applicable)
    - key and certificate management
    - whenever Consent is the applicable legal basis:
        - Consent enforcement point (for NBI or EWBI)
        - Caching relevant Application Privacy Profile(s)retrieved from the Privacy Management Function in the CSP domain (if allowed by local regulations)
        - Triggering Consent Capture by the Privacy Management in the CSP domain
        - In federated scenarios, triggering Consent Capture by the Privacy Management function in the CSP domain of the federated Partner

Note: OP may rely on procedures regarding Authentication (including Identity) / Authorization / Privacy Management through interfaces already in place in the CSP domain which are mapped to the SBI-AAPrM interface.

3.3.2.4 API Call Routing Functions

The API call routing functions provides (among others) the following:

- Load balancing
- Telco Finder service which is responsible for resolving the Operator associated with a target user identifier (e.g. based on a specific phone number) and returning information about the associated Operator. More details regarding Telco Finder can be found in GSMA PRD OPG.10 [20].

3.3.2.5 Routing of Requests in federation environments

An OP may have E/WBIs with multiple other OPs, as shown in Figure 1. Therefore, when an OP needs the support of another OP, it will need to make a routing decision determining what Partner OP's support would be required and what E/WBI would be best suited to reach that OP. Depending on the nature of the request, this decision could be based on combining information related to the request with information that the OP has on its Partner OPs. For example:

- For UE or Subscriber-related requests,
    - On the UE's IP address combined with information on the IP Address ranges that a Partner OP is covering,
    - On the MSISDN associated to a subscription combined with either information on the MSISDN ranges covered by the Partner OP or MSISDN-specific information from a Number Portability Database,
    - Domain information from a token or GPSI related to the Subscriber combined with the domain(s) covered by the Partner OP,
    - As the Home OP of an UE or Subscriber roaming on another network, information on what network they are roaming on combined with information on what Visited OP would serve that network.
- For resource-related requests,
    - On the Availability Zone that the request refers to combined with information on the Availability Zones supported by the Partner OP.

Note: An OP may have multiple routes towards a Partner OP, direct and indirect where the request might pass through one or more other Partner OPs. The criteria to select the most suitable route in that case are out of scope of this specification.

3.3.3 Exposure Functions

The Exposure Functions enable exposing Service APIs (to Applications or Aggregation/Marketplace/Enterprise Platforms) via the NBI-Service interface and Operate APIs (to Aggregation/Marketplace/Enterprise Platforms) via the NBI-Operate interface. The termination points for the NBI-Service and NBI-Operate API calls are provided by the corresponding API Gateway function as described below.

The Exposure Functions in an OP enable an Application Provider or an Aggregator to operate applications. Operating an application includes discovering the capabilities of the OP, both in functionality (e.g., how an application may be onboarded or instantiated) and in range of functionality (e.g., where may an application be run, and what QoS Performance Profiles are possible). The Exposure Functions also enable the exposure of other Operator Capabilities that can be consumed, via APIs, by the Application Providers or the Aggregators.

The Application accesses the Exposure Functions via the Northbound Interface (NBI). The Application Provider or the Aggregator expects to use APIs implemented at the NBI to carry out required functions.

3.3.3.1 NBI API Gateway

In addition to the common API Gateway functions provided in above section 3.3.2.2, the NBI API Gateway supports the following functions:

- Providing termination points for Service API calls from Applications (owned by Application Providers) or Aggregation/Marketplace/Enterprise Platforms (owned by an Aggregator or a 3rd party)
- Mapping to Transformation functions / SBI Gateway
- Routing to API Federation Management function / EWBI Gateway in case of API call Federation

Additionally, the NBI API Gateway supports:

- Providing termination points for Operate API calls from Aggregation/Marketplace/Enterprise Platforms
- Mapping to Operations and Business Transformation Functions / SBI Gateway

3.3.3.2 Exemplary scenarios

Based on the aforementioned capabilities, the NBI is expected to enable the following non-exhaustive list of scenarios:

- Service Endpoint Exposure: the Application Provider uses an authenticated and authorised endpoint to carry out service-specific scenarios (e.g. involving Application Instances on Edge Clouds);
- Operator Capabilities exposure: the Application Provider or the Aggregator inventories the Operator's capabilities, like Network Analytics, nominally available to applications as telco services via Service APIs (aka Open Gateway Services) (see section 3.3.2.1).
- Privacy Management for accessing Operator Capabilities: certain Operator services consider personal information of the Subscribers, the management of which shall remain under the Operator's control, based on functionalities exposed by the OP on the NBI (see section 3.3.2.3).
- Operator Capabilities OAM: the Aggregator or the Application Provider, when integrated through an Aggregator, will require certain functionalities to configure and control the Operator Capabilities, nominally available to an Aggregator via Operate APIs.
- Telco routing information exposure: the Aggregator or the Application Provider use the NBI of different Operators to identify which Operator is responsible for handling an Operator Capability, as related to a specific Operator's Subscriber (see section 3.3.2.4).
- Roaming capabilities: the Application Provider directly or via the Aggregator uses the NBI to consume information about service availability in the visited networks.

Note: The Aggregator consumes the OP services on behalf of the Application Provider. This interaction only refers to a technical service consumption, but the final service consumer remains in the Application Provider, e.g. for the Privacy Management.

Note: An Application Provider accessing OAM capabilities is considered as implementing an Aggregator's functionalities, therefore the Application Provider is indeed taking an Aggregator role in that case.

3.3.4 Federation Functions

The Federation Functions (see Figure 2) in an OP deal with establishing and sustaining the Federation Interconnect (E/WBI) between the OP instances.

3.3.4.1 NBI API Federation Management

Providing (among others) the following services supporting the E/WBI:

whenever a Leading OP is originating an E/WBI request:

- For incoming NBI-service requests, routing to the Partner OP selected to fulfil the request
- For incoming API responses from other OPs, routing to the NBI API Gateway.

whenever Partner OP is originating the E/WBI request

- For incoming API requests from Partner OPs, routing to the relevant Transformation function to map to SBI API calls through SBI API Gateway. This includes:
    - Forwarded Partner OP NBI-service requests
    - Partner OP Service and Resource Management requests, for instance, for reservation / monitoring of Edge Cloud Resources

3.3.4.1.1 Capability Exposure in a visited network

The exposure of capabilities in a federated/visited network, such as applying QoS Performance Profiles or obtaining certain network information, is crucial to provide the desired quality of experience to the Application Client in the roaming scenario. Therefore, the goal is to provide the same capabilities and SLIs in the visited network as in the home network. To achieve that, the Visited OP has to inform the Home OP about the capabilities available, including the SLIs. This may be subject also to the specific federation agreement.

If the visited network cannot fulfil a requested capability, the Home OP shall provide this information to the Application Provider.

- Capabilities like the NEF’s Network Capabilities will not be exposed directly on the E/WBI. The catalogue of available capabilities that can be exposed to the federated applications of the Leading OP shall be shared over the E/WBI.

3.3.4.2 Edge API Federation Management

The Federation Management functionality within an OP enables it to interact with other OP instances, often in different geographies, thereby providing access for the Application Providers to a more extensive set of Subscribers and multiple Operator capabilities

The following are prerequisites to enable the federation model:

- Operators need to have an agreement to share Edge Cloud and/or Network resources;

- Operators need to agree on an Edge Cloud and/or Network Resource sharing policy;
- Operators need to enable connectivity between the OP instances over which East/West Bound Interface signalling flows.

Federation Management provides the Management plane. The Management Plane covers the set of functionalities offered to Application Providers and OPs to control and monitor the resources and applications within the federation under their responsibility.

The Management Plane functionality is realised via the multiple functional blocks within an OP instance listed in the subsections below. The management actions are relayed between these different functional blocks using the NBI, SBI and E/WBI interfaces that have been defined for communication between them in section 3.1.

The Management plane works at two domain levels: application and infrastructure (resources). Each of these domains supports management at two distinct stages in the managed entities' life-cycle: the configuration and the run time management. Table 1 lists the functionality provided by the Management Plane in each domain and stage.

Domain	Stage	Management Functionality
Resources	Configuration	Resource Catalogue Synchronisation and Discovery
		Edge Node Sharing
		Resource sharing policies
		Automation of Orchestration
	Run Time	Edge Cloud Resource monitoring
		Lifecycle Automation
Application	Configuration	Application Management
		Service Availability on Visited Networks
		Automation of Orchestration
	Run time	Lifecycle Automation

Table 1: Management Functionalities

Note: Edge Cloud is used as an example. More details on the service are provided in GSMA PRD OPG.11 [21].

Note: There may be legal constraints restricting the distribution of specific applications to certain Regions that would need to be considered in the agreement when the federation is planned among multiple Operators. The technical impact of such legal constraints on OP is for further study.

3.3.4.2.1 Resource Catalogue Synchronisation and Discovery

Operators can include the Edge and Network Resources in the OP's set of available resources using the SBI.

OPs shall exchange and maintain the types of resources offered to each other (E/WBI).

This information may change and can be updated via the E/WBI whenever the types of resources offered to an OP by a Partner change due to Operational or Administrative events (e.g. due to scheduled maintenance).

A notification mechanism is supported over the E/WBI to achieve the above.

3.3.4.2.2 Resource sharing policies

An OP shall provide controls to the Operator to specify Availability Zones to be made available to a Partner OP. These controls shall allow all or part of the resources of an Availability Zone to be shared. Availability Zone sharing is dependent on the Federation agreement that exists between the OPs.

The information elements and data model used to represent Availability Zones in the Partner OP shall provide Representationally Consistency with the NBI data model.

3.3.4.3 Network Communication Service API Federation Management

3.3.4.3.1 UE Provisioning of URSP rules

To correctly provision the UE with UE Route Selection Policy (URSP) rules, the serving network deployment and configuration need to be considered. The home network defines the URSP rules, while the Visited OP needs to direct the visited network to connect the application to the correct DNN/NSSAI.

Operators need to exchange information to populate the URSP rules that map an Application to the right DNN/NSSAI; please refer to 3GPP TS 23.503 [8] for more details. The OP may facilitate the exchange of relevant deployment and configuration information between the serving and home networks to influence the construction of the URSP rules to be provided to the UE. This can apply to all connectivity models.

3.3.4.4 EWBI API Gateway

In addition to the common API Gateway functions provided in above section 3.3.2.2, the EWBI API Gateway supports (among others) the following functions:

- Federation interconnect management
- Providing initiation / termination of E/WBI API calls to / from other OPs
- Routing to API Federation Management Function (when applicable)

The Federation Interconnect uses secure transport, plus capabilities such as integrity protection for the E/WBI messaging between OP instances.

During the Federation Interconnect establishment, the Federation Functions of the participating OPs need to verify each other's identities through mutual authentication (being the security enforcement point of the EWBI API Gateways).

Federation Functions (see Figure 2) also ensure that the Partner OP is authorised to establish and maintain the interconnect according to the federation agreement between the Partner OPs/Operators.

3.3.4.4.1 Partner OP Provisioning

An OP shall allow mechanisms to provision Partner OP information used for Federation Interconnect establishment and management. This information would include, at a minimum, the following:

- The Partner Name;
- The Partner's geographical area (e.g. Country of operation, Regions, and Availability Zones);
- The Partner's description of shared resources;
- The Partner identifiers;
- The Partner's federation interconnect E/WBI endpoint;
- The federation agreement validity duration.

Between any two Partner OPs, the provisioning information shall be mutually consistent.

The Exposure Functions of a Leading OP (of an Application Provider) are responsible for providing Representational Consistency of the view of Regions, Availability Zones, and Resources as might be required by the Application Provider to perform application lifecycle actions.

3.3.4.4.2 Configurations

An OP shall provide various configuration capabilities to establish and manage the Federation Interconnect (e.g., via the API Federation Management function in Figure 2).

API Federation Management Functions interact with the Exposure Functions and the NBI. Cases of this interaction are detailed as appropriate in the following subsections.

3.3.4.4.3 Service Availability on Visited Networks Management

Note: Service Availability highly depends on the service. Generic Platform Requirements are for further study (e.g. Privacy Management).

3.3.4.4.4 Operational visibility.

The OPs shall have an operational view of each other, allowing Fault Management and Performance management within the limits of their agreements in the federation contracts.

This fault and performance management is based on the information obtained through the OP's monitoring of the exposed resources.

Due to the amount of exchanged information, a notification mechanism should be available to allow the above filtering for the information relevant for Fault and Performance management.

3.3.4.5 EWBI Authentication and Authorisation

When an OP connects to a Partner OP via the federation interconnect, it needs to authenticate itself to that Partner OP. This authentication requires that authentication information (e.g. digital certificate or passphrase) is provisioned in the OP. This mechanism can be mutually agreed between the involved Operators as a first step. A more generic solution based on a Certificate Authority could be considered going forward within the GSMA.

Authentication and authorisation between Partner OPs do not reach an Application Provider via the NBI. An Application Provider is expected to authenticate and authorise with its Leading OP, but not with any Partner OPs on which their applications run. The “chain of trust” required for an Application Provider to deploy an application on a Partner OP is composed of the authentication of the Application Provider on the Leading OP and the authentication of the Leading OP to the Partner OP.

3.3.4.6 Exemplary scenarios

Typical scenarios enabled by the Federation Functions across the E/WBI role are:

- Capabilities Exposure towards Partner OPs;
- Service Availability in visited networks.

Most of the above scenarios are supported by an Operator’s Operation Support System (OSS) capabilities (Service and Resource domains) which would need exposure of corresponding APIs (e.g., TMF Open APIs for Service / Resource management) through the OP. Exposure would mean publishing these APIs in the OP Catalogue, or creating an operational transformation function if one wants to expose simpler APIs and publish those APIs instead.

Depending on the deployment choice, an OP may perform a brokering function to simplify the interaction across multiple OPs, for which a subset of functions introduced in this section may be considered.

3.3.5 Transformation Functions

The OP Transformation Functions enable NBI and E/WBI API calls to be mapped to the required SBI API calls. One incoming API call may be mapped to one or more API calls on the SBI to provide a suitable response.

Transformation may concern mapping of the following request:

- NBI-Service API requests towards the underlying network capabilities
- Service and Resource Management requests towards the corresponding Service and Resource Management domain(s). A Transformation Function may need information about available Services and Resources from those domains before performing any transformation logic.
    - Service and Resource Management requests may originate from the NBI-Service, NBI-Operate or EWBI.

Depending on the destination of the SBI calls, the Transformation Functions can be split into:

3.3.5.1 Network Transformation Functions

Providing the following services:

- Transformation Functions for the realisation of the Service APIs in the lower levels of the architecture (e.g., as in GSMA PRD OPG.09 [19]).

3.3.5.2 Operations and Business Transformation Functions

Providing the following services:

- Transformation Functions for the realisation of the TM Forum Operate APIs in the lower levels of the architecture (e.g., on the SBI-OAM interface).

3.3.5.3 Exemplary scenarios

Typical scenarios enabled by the Transformation Functions are:

- SBI:
    - Exposure of usage and monitoring information to the Operator's charging engine via the SouthBound Interface – Charging Function (SBI-CHF) to enable Operators to charge for the OP's services.
    - Interaction with the mobile network via the Southbound Interface – Network Resources (SBI-NR), for example to:
        - Request and receive notifications on UE Mobility events from the network to assist applications.
        - Configure traffic steering in the mobile network;
        - Receive statistics/analytics.
        - Receive information related to the network capabilities, such as QoS, policy, network information, etc.
        - Receive the End-User’s profile data (e.g. S-NSSAI, DNN, etc.).

3.3.6 Integration Functions

3.3.6.1 SBI API Gateway

In addition to the common API Gateway functions provided in above section 3.2.2.2, the SBI API Gateway provides (among others) the following functions:

- Termination of the SBI towards:
    - Network Resources (SBI-NR)
    - Operations and Management systems (SBI-OAM)
    - Authentication, Authorization and Privacy Management functions in the CSP domain (SBI-AAPrM)
    - Cloud Resources (SBI-CR)
    - Edge Interconnection Network (SBI-EIN)
    - Charging (SBI-CHF)

4 Detailed Interface requirements

4.1 Northbound Interface (NBI)

4.1.1 High-level requirements

NBI covers NBI-Service functionality provided to Application Providers after registration and activation. Aggregators have additional needs to automatically onboard and manage Application Providers and Applications that connect to OP via the Aggregator, this process is handled via the NBI-Operate.

4.1.1.1 Generic NBI requirements

1. All Operators and Operator Platforms shall offer the service, resource and End-User provisioning capabilities through the same NBI.
2. The NBI shall offer Operator Capabilities to Application Providers and Aggregators.
3. In deployment, the NBI shall use profile-based access control to provide appropriate restrictions on the amount of functionality that the NBI offers to a particular system or person, according to the operational profile. For example, profile-based access control such as Role-Based Access Control restricts the degree of access depending on the person's (or system's) defined privilege and role.

   Note: Not all profiles have access to all the functions listed below. For example, monitoring information would not necessarily be accessible during onboarding. In addition, the detail of monitoring information may depend on the operational profile (for example, first-line vs second-line support).

   Note: Access control to NBI Service and Operate APIs can be based on an Access Token mechanism.

   Editor's note: Access control mechanisms are FFS.

4.1.1.2 NBI-Service requirements

1. The NBI shall allow to easily identify the target resource (e.g. UE, Subscriber, Availability Zone) to which the request relates.
2. For End-User identity, an OP shall accept MSISDN, IP-address or User Identity Token as identifier and for routing purposes

   a) End-User identification based on an MSISDN requires the Operator to register the MSISDNs using number ranges that can be shared or Number Portability Databases.
   b) End-User identification using IP-address ranges requires the Operator to share its service IP ranges amongst participating Operators and Aggregators for routing purposes.
   c) End-User identification using User Identity Token requires the Operator to include an User Identity Token generating and check function as described in section 4.3.6.2

4. An OP shall expose routing information to Aggregators to find the Home OP of the End-User for all subscriptions served by the OP
5. The Home OP shall be found regardless of whether its OP serves a MVNO or MNO, including sub brands owned by these organizations.
6. The Home OP shall be discoverable when a device is connected to Wi-Fi. Note that in this case the IP address will not point to the home Operator.
7. The Operator shall ensure that MSISDNs including non dialable IOT MSISDNs can be used for identification of the Home OP if these are needed for routing purposes.
8. For an API routing lookup method with User Identity Token check, the end-to-end latency including device, network and OP shall aim for sub 1 second and not exceed 4 seconds

9. The lookup mechanisms shall support devices connected without an MSISDN such as game consoles, laptops and IOT devices (e.g. based on IP address or User Identity Token)
10. The service should return the routing information in case of secondary devices owned by the Subscriber.

Note: Device ownership being different from subscription owner (e.g. connecting someone else's laptop to your UE) is FFS.

11. The OP shall only accept NBI-Service API requests from Applications that are registered.
12. In the event of any error conditions that occurred following the NBI-Service API request by the Application or the Aggregation/Marketplace Platform, the OP shall respond with the appropriate API error codes.

Note: Such error conditions can occur in the NBI API flow request including within the authorization flow, the Service API and the Operate API request.

13. In the event of any error conditions that occurred in the authorization flow request, the OP shall respond with the appropriate error codes.

Note: These error conditions can occur within the specific steps in the authorization flow such as when the OP handles the routing, the End-User validation and authentication, the End-User consent management, the client authentication or the network-based tokens.

14. If the NBI API supports Callbacks such as API responses for subscriptions and notifications, the OP should factor in the required security considerations and protection against abuse.
15. The OP should ensure that the NBI API error response should not disclose an End-User's privacy sensitive information.
16. In the event of a mismatch in network technology such as when the NBI API request relates to a user being on 5G but the End-User is connected or downgraded to 3G or 2G, the OP should respond with the appropriate error codes.

4.1.1.3 NBI-Operate requirements

NBI-Operate serves Operate APIs to registered Aggregators. Using the NBI-Operate the Aggregator can provide and manage the needed information about its customer Application Providers and its Applications that want to use NBI-Service APIs provided by the Operator.

Note: The registration of Aggregators is not in scope of this document.

1. The OP shall allow the Operator to perform end-to-end lifecycle management including validation of the legal entity, modification of existing Applications or management of Terms and Conditions acceptance.
2. The OP shall expose APIs to Aggregators to automate the onboarding of Application Providers and Applications and requesting access to their wanted APIs.
    a) For the Application Provider identity, an OP shall accept an information structure with details about the Application Provider and confirm its creation to the Aggregator once approved.

3. For the Application Identity, an OP shall accept an Application Identity Profile connected to an approved Application Provider and confirm its creation to the Aggregator once approved.
4. The OP shall process subscription access requests to Service APIs after requirement steps 2 and 3 are approved for the Application and confirm the request once approved.
5. The OP shall provide the capability for Aggregators to retrieve a list of Service API product offerings from the API Catalogue in the API Management function.
6. The OP shall allow the Aggregator to subscribe to updates of the API Catalogue to ensure updates on Service API product offerings (e.g. description, prices, etc) are received.
7. The OP shall provide Aggregators with information related to consumption of Service APIs exposed by the OP (and accessed through the Aggregator) by the Applications of the Application Providers.

Note: This information is exposed towards Aggregators for consolidation and billing management.

8. The OP shall provide the capability for Aggregators to supervise the status of running services, using appropriate performance management (e.g., monitoring) and fault management (i.e., issue/ticket management) mechanisms.

4.1.2 Onboarding Profile

4.1.2.1 General

When an Application Provider is onboarded via an OP portal or via an NBI-Operate, the OP shall be in charge of:

- gathering the necessary information to onboard the Application and fulfil the request,
- mapping the Application Provider's request for exposed capabilities to the available capabilities in the target network(s),
- accept detailed information structures about the Application Provider and confirming its creation once approved
- accept an Application Identity Profile linked to an approved Application Provider and confirming its creation once approved
- process and confirm subscription access requests for Service APIs for the Application
- Service Catalogue exposure covering the OP’s available services

Thus, the onboarding management shall allow onboarding the application while meeting different criteria provided by the Application Providers and the Operators that own the OP instance and the underlying resources.

4.1.3 Management Profile

An OP shall offer a uniform view of management profile(s) to Application Providers:

1. An OP shall describe the exposed capabilities of the Leading OP's network(s) and those of the federated target networks
2. An OP shall allow an Application Provider to specify whether or not it requires service availability on visited networks (that is, when a UE roams away from its home network Operator).

3. An OP shall provide an Application Provider with telemetry information concerning the performance of the requested service, including fault reporting.
4. An OP shall allow an Application Provider to request a particular granularity for the telemetry information they receive.

Note: Possibly using a publish and notification approach.

Note: Different operational profiles require different granularity about the telemetry information (how fine-grained and how often).

4.1.4 Void

4.1.5 Security Requirements

The following security requirements shall be considered:

1. The NBI shall provide an authentication mechanism to enable access only to authenticated and authorised entities.
2. All interactions over the NBI interface shall use an application layer security protocol that runs over a reliable transport and guarantees mutual authentication between the OP (Exposure Functions) and the Application Provider. The security policies are enforced by the NBI API Gateway.
3. This authentication shall rely on commonly used API authentication mechanisms (e.g. OpenID Connect, etc.).
4. The NBI shall provide security mechanisms to guarantee the confidentiality, integrity and authenticity of the exchanged data. The security protocol used over the NBI shall also guarantee security properties such as perfect forward secrecy and mechanisms to prevent intervening attacks, such as replay, relay, and man-in-the-middle attacks.

4.2 East/Westbound Interface

The E/WBI connects Partner OP instances with the primary goal of allowing Application Providers of an OP to utilise the Operator Capabilities exposed by another OP.

The E/WBI is not exposed to the Application Providers and is primarily managed by the Federation Functions within the OP (see Figure 2).

4.2.1 High-level requirements

1. The E/WBI is universal, meaning that all Operators and Operator Platforms provide Edge Cloud and other capabilities to each other through the same E/WBI.
2. The E/WBI shall allow establishing the identity of the Partner OPs.
3. An OP shall be able to identify the UEs among OP instances.
4. An OP shall be able to identify the Application Providers among OP instances.
5. An OP shall be able to identify the applications among OP instances.
6. The E/WBI shall allow to easily identify the target resource (e.g. UE, Subscriber, Availability Zone) to which the request relates.
7. The E/WBI shall allow to update Partner OPs on changes related to the resource identifiers for which the OP can offer services (e.g. IP address ranges for UEs, Availability Zones offering Edge Resources).

4.2.2 E/WBI Services

The following sections provide a list of services that would be executed on the East/West Bound Interface.

4.2.2.1 East/West Bound Interface Management Service

The East/West Bound Interface Management Service shall be used for setting up and maintaining the East/West Bound interface between OPs.

The service would include APIs for the following:

- Setup of the East/West Bound Interface between OPs;
- Update parameters of the East/West Bound Interface;
- Heartbeat/Keep-Alive of the East/West Bound Interface;
- Termination of the East/West Bound Interface.

4.2.2.2 Application Management

The E/WBI needs to replicate the behaviour and functions available on the NBI to transmit the workload, requirements, mobility decisions, privacy considerations (if already in place on the originating OP side), and policies across all the operators' instances required to deploy the application.

1. The E/WBI shall allow forwarding the requests to any federated OP whose footprint has to be covered.

    a) If privacy considerations are already in place in the Leading OP side, e.g., Privacy Information (when Consent is the applicable legal basis), the E/WBI shall support forwarding e.g. the relevant Privacy Information to any partner OP (provided local regulation frameworks allow it).
    b) If the Leading OP avails Privacy Information, the E/WBI shall allow notifications (to any partner OP) on changes of that configuration on the Leading OP-side (provided local regulation frameworks allow it).

2. An OP receiving a request through its E/WBI shall get in charge of the management of the application:

    a) An OP receiving an instantiation request through its E/WBI shall apply its own policies and criteria for processing the request.
    b) An OP receiving an instantiation request through its E/WBI shall be responsible of the Network Communication Service management for the service and the (to be) connected subscribers.
    c) An OP receiving an instantiation request through its E/WBI along with Privacy Information (when Consent is the applicable legal basis), should cache this information for Authorizing future invocations and be able to process any future notification on changes of the Privacy Information.

3. The E/WBI shall forward the management procedures, information and statistics to be shared with the Leading OP of the Application Provider.
4. The E/WBI shall be employed for managing the service continuity on visited networks.

5. The E/WBI shall forward the network and analytics information to be shared with the Leading OP of the Application Provider.
6. The E/WBI shall forward the Capture Consent message (whenever Consent is the applicable legal basis) to the Leading OP, which in turn should forward it to the associated Privacy Management Function to capture the consent from the End-User.

4.2.2.3 Events and Notifications Service

The Events and Notifications Service shall be used to set up, send and receive Events and Notifications from one OP to another OP over the E/WBI.

As indicated under the Availability Zone Information Synchronisation Service, each OP publishes the information about the resource levels provided to each Partner. An OP shall send Notifications to Partner OPs related to these published resources. For example, in the following scenarios:

- The availability state of these resource changes;
- The consumption of resources reaches a pre-defined threshold (e.g. warning notifications when consumption reaches 80% of the agreed threshold value);
- Imminent Federation Agreement expiry.

To enable this, the Events and Notifications Service provides the following APIs over E/WBI:

- Setup Event reporting (e.g. resource threshold levels);
- Update Event reporting parameters;
- Notifications for Events.

4.2.2.4 Service Availability in Visited Network Management Service

This service shall be used to support information exchange between the OPs to enable service availability for UEs in the visited network.

Information elements that need to be shared over E/WBI to support this scenario include:

- Discovery Service URL for a Partner OP.
- Authorisation information for User Clients.

Note: In this version of the document, it is assumed that the applications available to roaming Subscribers have been provided to the Visited OP through a federation including both OPs. Future versions of this document may extend to roaming outside of a federation.

This service shall include APIs over the E/WBI for the following:

- Setup Service Availability in Visited Network related parameters towards Partner OPs;
- Update Service Availability in Visited Network related parameters towards Partner OPs;
- Enable User Client authentication information and provide authorisation for a visiting User Client from the Home OP.

4.2.3 Security Requirements

The following security requirements shall be considered.

OPs can belong to different Operators/players, so special requirements shall be considered for managing the relations and the resources/information sharing.

1. The E/WBI shall maintain the topology hiding policy between Operators/players.
2. An OP shall only expose the resources to its Partner OPs previously agreed with each specific Partner.
3. The E/WBI shall provide an authentication mechanism to enable access only to authenticated and authorised entities. Therefore, mutual authentication is recommended between the instances of the OP.
4. This authentication shall rely on commonly used API authentication mechanisms (e.g. Public Key Certificates, etc.).
5. The E/WBI shall provide a security mechanism to safeguard the confidentiality, integrity and authenticity of the exchanged data
6. The E/WBI shall support the adoption of strong security algorithms that guarantee forward secrecy and prevent intervening attacks such as replay, relay, and man-in-the-middle attacks.

4.3 Southbound Interface

The Southbound Interface of an OP includes all interfaces the OP is consuming from other parts of the service provider's infrastructure to create the capabilities of the different roles described in section 3.3. Therefore, the SBI includes interfaces for:

- Infrastructure manager functions of a cloud or edge cloud infrastructure (e.g. resource management for compute and Network Resources);
- Orchestrator functions facilitating the application lifecycle management and scheduling;
- Service management functions (e.g. platform services, network services, mobility support, etc.);
- Other external functions that are providing services to the OP.

In many cases, close interworking between resource management, application lifecycle management, platform services and traffic management services is needed.

The SBI is not defined by the OP but by the systems consumed.

4.3.1 SBI-CR

If the OP is used to expose cloud resources (e.g. Edge) or capabilities based on such resources, the integration with cloud resources APIs on SBI allows OP to support the needed functionalities for application and resources management.

An Operator Platform shall be able to access the cloud resources of its Operator/cloud provider. This access shall allow the OP to fulfil request/response transactions regarding an application's lifecycle, catalogue the resources/capabilities and get feedback about the status of the different Cloudlets or Edge Nodes.

4.3.2 SBI-NR

4.3.2.1 Network

The Network Exposure APIs on the SBI-NR, optionally, can help an OP to obtain various mobile core network information of a UE and may enable the OP to perform some of the tasks. Some task examples are as given below:

- UE location information retrieval;
- Request specific Quality of Service (QoS);
- Apply local routing and traffic steering rules for LBO/SBO for the Application traffic;
- Application relocation on most adequate Edge Nodes;
- Influence Data plane attachment point (re)selection for Service Continuity;
- Collect radio network information, e.g. cell change notification, measurement reports etc. for mobility decisions;
- Support the profile data for the End-User.

Some of the functions, namely location info retrieval or requesting specific QoS, can be performed in a 4G network, while others are introduced in 3GPP Release 15. They will be guided by further developments in the specifications in future revisions.

The functionalities mentioned above are optional, and an OP implementation can choose to use the available interfaces to optimise the platform functionalities.

The above list is not exhaustive but indicates some of the main informational elements and functions an OP is expected to perform. The SBI-NR interface enables an OP to meet the required SLA agreed with the external actors like Application Providers and may help optimise the utilisation of available Network Resources in a mobile Operator network.

The mobile core network may provide all, or a subset of, the above information via the SBI-NR APIs to the OP. In a 5G mobile core network, an OP, in the role of an Application Function (AF), may communicate with the 5G Core (5GC) network over the standardised interfaces as defined by 3GPP, for example, using the services of the NEF network function.

Additionally, an OP, apart from using the SBI-NR APIs for self-decision, may also provide (indirect and abstracted) access to some of the API's capabilities to authorised applications through the NBI and E/WBI. For enabling that, the Transformation functional level (and the respective Transformation Functions) depicted in Figure 2 are involved.

4.3.2.2 General

The SBI-NR connects an OP with the specific Operator infrastructure that delivers the network services and capabilities to the user.

An OP shall be able to access network capabilities that the Operator has chosen to expose through the SBI-NR interfaces of the Operator. However, an Operator need not implement the NEF/SCEF interfaces, in which case these capabilities have to be provided in some other way or else may not be available.

OP integration to Network Resources shall allow:

- The OP to authenticate and authorise the End-Users to access the services in the home and visited network scenarios.
- The OP to access network capabilities that the Operator has chosen to expose, e.g. QoS Performance Profiles, Network Events/Statistics.
- The OP to access the location information of the End-Users in the network.
- The OP to access policy control capability exposed by the network, e.g. for charging or quality of service handling.
- The OP shall be made aware of the data connection status (e.g. if a user has a data session or not).
- The home network OP shall be the only entity able to control home Network Resources.
- The OP shall be able to retrieve network analytics information (when available) in a standardised way: load level information, network performance, service experience, etc.
- The OP shall be able to retrieve resource analytics information (when available) in a standardised way.
- The OP shall be able to access an End-User's data profile.
- The OP shall be able to retrieve End-User's roaming access details (e.g. status, network connected).

4.3.2.3 OP integration to 5G Core/4G Core via Exposure Functions

4.3.2.3.1 Introduction

The NEF/SCEF APIs [4] [5] are a set of APIs defining the related procedures and resources for the interaction between NEF/SCEF and AF/Services Capability Server (SCS). The APIs allow the AF/SCS to access the services and capabilities provided by 3GPP network entities and securely exposed by the NEF/SCEF. Some APIs are applicable for both 5G Core and 4G Core.

Figure 3 shows a functional mapping that describes how an OP accesses features and services exposed by the NEF/SCEF.

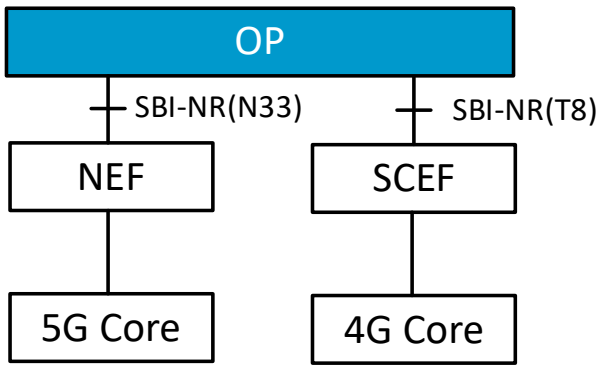

Figure 3: Functional mapping between OP and NEF/SCEF

A combined SCEF+NEF may be deployed by the MNO to hide the specific network technology from applications for user devices having capabilities for 4G and 5G access. In those scenarios it is expected that an OP should be able to support the communication with the combined SCEF+NEF on the SBI-NR interface.

Different sets of APIs can be supported by the two network types i.e., EPC and 5GC. From this perspective an OP should be able to discover the API capability differences while interacting over the SBI-NR with mobile core network.

4.3.2.3.2 General Requirements

1. An OP's SBI-NR shall be able to interact with 5G Core/4G Core via the NEF or SCEF to access network capabilities.
2. An OP's SBI-NR shall support the exposure interface [4] [5] for interacting with the 5G Core/4G Core.
3. If the NEF/SCEF returns an error response to an OP's SBI-NR, the OP shall perform error-handling actions.
4. An OP's SBI-NR shall be able to report the functionality available from the network.
5. An OP shall be able to deal with the situation where the network is not providing the expected functionality.
6. An OP's SBI-NR may be able to configure the user traffic to be routed to the applications in the local data network.
7. An OP's SBI-NR may be able to interact with the NEF for configuring and influencing the traffic routing policies.

    a) An OP may be able to specify the request for routing, influencing network mobility and routing, including but not limited to:

        i. UE and application identities
        ii. Traffic filtering and routing criteria,
        iii. Possible locations of the Application Instances
        iv. Whether the UE network data plane can be relocated.
        v. Whether validation on UE network data plane relocation is required.
        vi. Whether the UE IP address shall be preserved in data plane relocation
        vii. The type of Session and Service Continuity (SSC) mode
        viii. Whether inter-operator handover is required.

    b) An OP may be able to request to be informed on UE data plane mobility events.
    c) An OP may be able to receive UE data plane mobility events, receiving the target node identifier where the UE should re-attach because of the network mobility process.
    d) An OP may be able to receive UE data plane mobility events, receiving and processing the target IP of the UE that will be assigned.
    e) An OP may be able to negotiate the UE data plane mobility process based on the Application Instance relocation process.

8. An OP's SBI-NR may be able to collect information on network congestion or access concentration in a specific area.
9. An OP's SBI-NR may be able to retrieve a UE mobility analytics report.
10. An OP's SBI-NR may be able to retrieve a UE communication pattern report (e.g. Uplink/Downlink volume per application).
11. An OP's SBI-NR may be able to retrieve a network performance report (e.g. gNB active ratio, gNB computing resource usage).
12. An OP's SBI-NR may be able to report QoS change statistics in a specific area.

13. An OP's SBI-NR may be able to retrieve UE status reports (e.g. location information, reachability, roaming status).
14. An OP's SBI-NR may be able to control the transfer of data in the background for UEs.
15. An OP's SBI-NR may be able to configure QoS session parameters to allow an Application Provider to communicate with a UE with an improved QoS level (e.g. more suitable QoS Performance Profile).
16. An OP's SBI-NR may be able to configure the Alternative QoS Performance Profiles applicable to the user data session; e.g. when using access technologies where the specific QoS Performance Profilerequested by the Application Provider cannot be met.
17. An OP's SBI-NR may be able to receive QoS relevant notifications based on UE connection statistics.
18. An OP's SBI-NR may be able to configure the charging party of the UE data sessions.
19. An OP's SBI-NR may be able to configure service-specific parameters for UEs (e.g. Network Slice).
20. An OP's SBI-NR may be able to initiate a device trigger to a UE for performing application-specific actions (e.g. starting communication with the OP's SBI-NR).
21. An OP's SBI-NR shall be able to influence the URSP rules sent to the UE to allow the UE to select the DNN/NSSAI of the serving network based on the Application's needs.
22. An OP's SBI-NR shall be able to trigger URSP rules to be sent to the UE connected to subscribed QoS Performance Profiles.
23. An OP's SBI-NR may be able to influence the 5G mobile core network to establish a user plane for PDU sessions requiring access to edge services based on OP-provided criteria.
24. An OP's SBI-NR may be able to report access type change notifications for UEs due to user mobility.
25. For the APIs that are common to EPC and 5GC, an OP's SBI-NR shall be able to support operations to be informed on their availability or the expected level of support.
26. An OP's SBI-NR shall be able to work with the Common API Framework (CAPIF) when available.

Note: An OP's SBI-NR can work without CAPIF. If CAPIF is not supported, the SBI-NR API will provide an alternate means of providing these functions.

4.3.2.4 Security Requirements

The following security requirements shall be considered:

1. The SBI-NR shall provide an authentication mechanism to enable access only to authenticated and authorised entities. Therefore, mutual authentication shall be provided between the OP and the NR.
2. The SBI-NR shall support the use of authorisation mechanisms by its endpoints that grant access to only the necessary authorised services and data.
3. The SBI-NR shall provide security mechanisms to safeguard the confidentiality, integrity and authenticity of the exchanged data.
4. The SBI-NR shall support the adoption of strong security algorithms that guarantee forward secrecy and prevent intervening attacks such as replay, relay and man-in-the-middle attacks.
5. The SBI-NR shall support security mechanisms to protect the network functions discovery procedure of the NEF/SCEF by an OP.

4.3.3 SBI-CHF

An OP shall provide a set of capabilities that will enable the charging and billing for the usage of the Operator's services exposed to third party providers. Although these services and capabilities are quite heterogeneous and in constant evolution, they can be classified into a set of categories that share common characteristics from a charging perspective and that are described in Annex E of this document.

The Operator that runs an OP decides on its commercial model and how it charges for OP services. There are many potential choices. Two simple examples are subscription-based and pay per use, whilst a more complex example is demand-based pricing. The OP architecture, therefore, defines various information to support a variety of commercial models. However, a particular commercial model may only require a subset of the information, while another may require additional details. When a service uses federated resources, the two Operators need to agree in advance on what charging information to report. Note that this is independent of the commercial model between the application provider and its OP.

Finally, OP shall expose all of that information to an external charging engine through an SBI for charging (SBI-CHF) under Operator or resource owner control so that each stakeholder can define its commercial strategy, models and offers. As shown in Figure 2, the OP termination of the SBI-CHF interface is within the Integration functional level (e.g., in a SBI API gateway instance).

4.3.3.1 General Charging Integration Requirements

1. An OP shall be able to integrate with the CCS (Converged Charging System) deployed in the Operator's network through the SBI-CHF interface. This integration will allow doing the rating and charging for the usage of the services and capabilities exposed by the OP.
2. Considering that there could be different CCS instances deployed in the Operator's network (e.g. dedicated instances for a particular service/customer segment, geo-redundant deployment, etc.) the OP will be able to select the CCS instance that will be used to do the rating and charging of the service.

Note: The criteria used for this CCS instance selection (e.g. CCS discovery mechanism, OP local configuration, etc.) are for further study but as a general approach an OP will provide mechanisms to configure the target CCS instance depending on a combination of different parameters (e.g. type of service used, application provider identifier, etc.)

3. An OP shall support different charging integration models with the CCS. The charging integration models to be supported will be the ones standardised by 3GPP and defined in 3GPP TS 32.240 [10]. As a reference, the following charging models shall be supported:

    - Event Based Charging:

      This charging model is based on a request/response pattern, where an OP would trigger a charging request when an event occurs (e.g. an API invocation) including all the information relevant for rating and charging for the CCS.

The CCS would use the information provided in the charging request to do the rating and charging for that event and will send the response to the OP with the result.
The following charging model, defined by 3GPP, will be supported by the OP:

- PEC (Post Event Charging): a charging request is sent after the service is delivered. (e.g.. an OP receives an API call, makes several API calls through the SBI to deliver the service and a Charging request is sent after the OP makes these API calls through the SBI)

Note: Although 3GPP also defines IEC (Immediate Event Charging) charging model, where a charging request Is sent before the service that is associated to the event is delivered, the support for this charging model in the OP is not mandatory and is left for further analysis as it has dependencies on the evolution of 3GPP standards in the context of 5G SA charging.

- Session Based Charging:

Note: the usage of this charging model in the context of the OP requirements is for further study. The description of this charging model will be expanded in next releases of this document).

The charging model to be used by an OP in the integration with the CCS will depend on the particular service to be charged.

4. An OP shall provide mechanisms that will allow doing the charging for the services in the case of unavailability of the connection with a CCS through the SBI-CHF interface. These mechanisms are for further study but as a reference the following approaches could be used:

   - Usage of a primary/secondary/pool of CCS instances as the result of the CCS instance selection procedure, so that in case the primary instance is not available a secondary one could be used.
   - Ability to log/store charging requests when no CCS instances are available so that this information could be used to do the rating and charging when communication is re-established.

4.3.3.2 Services and capabilities exposure charging requirements

1. An OP shall support rating and charging for the following service categories described in Annex I of this document:

   a) Category 1: Network capabilities exposure services with no impact on the device's data usage.
   b) Category 2: Network capabilities exposure services with impact on the device's data usage.
   c) Category 3: Network provisioning services.
   d) Category 4: Edge application management services.

Note: Category 4 is relevant only if the OP exposes such capabilities.

2. An OP shall support the following charging factors/events for triggering charging for the services included in Category 1:

    a) Service activation charging.
    b) Charging per Service API invocation (and related notifications):

        An OP shall be able to send a charging request to the Operator's CCS through the SBI-CHF following the PEC charging model defined by 3GPP. This charging request will be sent once the service is delivered upon confirmation from the Network.

Note: As a reference, the integration between an OP and the Operator's CCS for this charging factor is shown in section G.2.1 of this document.

3. An OP shall support the following charging factors/events for triggering charging for the services included in Category 2:

    a) Service activation
    b) Charging per Service API invocation (and related notifications):

        The same requirements as in requirement 2.b) (charging per Service API invocation for category 1 services) of this section are applicable for this case.

    c) Charging based on data traffic consumption in the Operator's Network as a result of a previous Service API invocation.

        An OP shall be responsible for providing the Operator's Network (through the SBI-NR) with the information that allows the correlation between a Service API invocation and a data traffic flow from a device in the Operator's Network.

Note: A charging dialogue will take place between the Operator' Network and the Operator's Charging engine following the regular procedure used in the Operator to do the data sessions charging (out of the scope of this document). The Operator's Network will include the correlation information provided by the OP in the charging requests sent to the CCS to indicate the API Invocation's impact on charging.

Note: As a reference, the integration between an OP, the Operator's Network and the Operator's CCS for this charging factor is shown in section G.2.2 of this document.

4. An OP shall support the following charging factors/events for triggering charging for the services included in Category 3:

    a) Service activation
    b) Charging per Service API invocation (and related notifications):

        The same requirements as in 3.b) above (charging per Service API invocation for category 2 services) are applicable for this case.

    c) Charging per Service API invocation (service lifecycle modification charging):

An OP shall be able to send a charging request to the Operator's CCS through the SBI-CHF following PEC charging model defined by 3GPP after a Service lifecycle modification API request is received from the NBI and the service is delivered.
For the specific case of Network Slice Instance management services, the charging models defined in 3GPP TS 28.202 [11] will be supported. In other scenarios, e.g. realisation through QoS services or dedicated APNs where no specific definition is available, the Operator can fall back to traditional ways of charging per event (e.g. QoS) or lifecycle management of a needed APN.

d) Charging based on data traffic consumption in the Operator's Network as a result of a previous Service API invocation:

The same requirements as in 3.c) above (charging based on data traffic consumption for category 2 services) are applicable for this case.

5. An OP shall support the following charging factors/events for triggering charging for the services included in Category 4:

a) Service activation
b) Charging per Service API invocation (application lifecycle management operations):

An OP shall be able to send a charging request to the Operator's CCS through the SBI-CHF following PEC charging model defined by 3GPP after an Application lifecycle management API request is received from the NBI or from the E/WBI (for the case of federated scenarios) and the service is delivered.

Note: As a reference, the integration between an OP and the Operator's CCS for this charging factor is shown in sections G.2.1 and 4.10.1 (for federated scenarios) of this document.

c) Charging per Service API invocation (charging for edge enabling infrastructure resources usage based on subscribed capacity in API request):

An OP shall be able to send a charging request to the Operator's CCS through the SBI-CHF following PEC charging model defined by 3GPP after an API request is received from the NBI or from the E/WBI (for the case of federated scenarios) requesting for capacity reservation in the Operator's Cloud Resources and the API request is processed.

Note: As a reference, the integration between an OP and the Operator's CCS for this charging factor is shown in sections G.2.1 and G.3.1 (for federated scenarios) of this document.

d) Charging based on edge enabling infrastructure resources usage:

An OP shall be able to send a charging request to the Operator's CCS through the SBI-CHF including the information about the effective resources' usage in the Operator's CR over a period of time.

Charging models defined by 3GPP in 3GPP TS 32.257 [12] will be supported for this purpose.

Note: An OP will periodically retrieve the actual resource usage information from the Operator's CR based on the agreed Data Collection Interval from the SBI-CR or from the E/WBI in the case of federated scenarios. This information will be the one included in the charging request sent to the Operator's CCS through the SBI-CHF. The definition of this mechanism is out of the scope of this section.

Reference to diagram flow in section G.2.3 and in section G.3.2 (for federated scenarios) of this document is provided for clarifications.

e) Charging based on data traffic consumption in the Operator's Network as a result of a previous Service API invocation (in non-federated scenarios):

The OP shall be responsible for providing the Operator's Network (through the SBI-NR) with the information that allows to do the correlation between a Service API invocation and a data traffic flow from a device in the Operator's Network that is accessing an application.

Note: A charging dialogue will take place between the Operator' Network and the Operator's Charging engine following the regular procedure used in the Operator to do the data sessions charging (out of the scope of this document). The Operator's Network will include the correlation information provided by the OP in the charging requests sent to the CCS.

Reference to diagram flow in section G.2.2 of this document is provided for clarifications.

6. An OP shall allow the configuration of charging factor/factors to be used for the services applicable to each category on a per service basis. It will also be possible to configure different charging factor/factors per service depending on the scenario: federated / non-federated.

   In the case of federated scenarios, this configurability is applicable both to the Leading and Partner Ops.

Note: It is an Operator's decision to decide which charging factors - from the ones that are applicable to a category – will be configured/used to do the charging and billing for the usage of a service/capability.

This decision will be dependent on the commercial model chosen by the Operator for the commercialisation of that service (and out of the scope of this document).

The configuration in the OP will need to be aligned with this decision.

7. An OP shall maintain security and data/topology privacy requirements when reporting consumptions to the Operator's Charging Engine (both in federated and non-federated scenarios.

In case that - for technical reasons - an OP needs to report information disclosing privacy-sensitive data/topology to the Operator's Charging Engine (e.g., to allow correlation in scenarios where data traffic consumption in the Operator's Network needs to be correlated with a Service API invocation) the responsibility to guarantee these security/privacy requirements will be on the Operator's Charging engine side.

Note: the specific information to be provided for correlation and the mechanisms used in the Operator's Charging Engine are for further study.

4.3.3.3 Charging information

1. The charging requests sent by an OP to the Operator's Charging Engine through the SBI-CHF shall include any information usable by the Operator's CCS to address the rating and charging of the services and enable the final billing process in the Operator. An OP creating or sharing charging data shall guarantee the security, integrity, availability, and non-repudiation of charging data.
2. Charging information to be provided by an OP in the charging requests shall include the identification of the different parties that are involved in the transaction, from the Application Provider to the UE. These identifiers could be used for different purposes by the Operator's CCS (e.g., to determine the chargeable parties, to have end-to-end traceability of the transaction, etc.).
3. An OP shall at least include the following identifiers in the charging requests sent to the Operator's CCS through the SBI-CHF interface:

    a) Party identifiers involved in the transaction: Application ID, Application Provider ID, Customer Device ID.
    b) Operator Platform ID
    c) Partner Operator Platform ID (only applicable in the case of federated scenarios)

4. An OP shall include a correlation identifier of the NBI Service API invocation in the charging requests sent to the Operator's CCS. This correlation identifier is a unique identifier for that particular transaction. This ID will allow end-to-end traceability and will assist in the correlation required to enable charging factors where data traffic charging in the Operator's Network needs to be correlated with the Service API invocation by the Operator's CCS.

Note: The mechanism used to generate this correlation identifier is out of the scope of this section and for further study.

5. An OP shall include specific information that will depend on the service category and the charging factor in use in the charging requests sent to the Operator's CCS.

Note: The information to be collected is described in the next requirements.

A summary table showing the list of potential charging factors per service category is shown in Annex E of this document.

6. For the services included in categories 1, 2, 3 and in case the charging factor chosen by the Operator is the one based on API invocations or on service lifecycle modification

operations received, the OP shall be able to include the following information in the charging requests:

a) Mandatory information:

i. API type (identification of the Service API that was invoked through the NBI e.g., device location)

b) Optional information:

i. A subset of the parameters included in the Service API invocation. The list of parameters to be included (if any) will be configurable per service.
ii. A subset of parameters retrieved from the Network (e.g., device ID in the Operator's Network) after the service is delivered. The list of parameters to be included (if any) will be configurable per service.
iii. API result code

7. An OP shall expose the QoS Performance Profiles associated to the Network Communication Service in the charging request as part of the optional information.

   In the specific case where the Network Communication Service is realised with slicing in a 5G network, the charging for the Network Slice management operations will use the charging models and charging information defined in 3GPP TS 28.202 [11].
   The concrete list of mandatory/optional parameters is for further specification but as a general approach any of the parameters included in the GST (Generic Network Slice Template) could be included by an OP based on configuration.
   The Operator needs to fallback to traditional ways of charging when the Network Communication Service is realised in a 4G network with QoS or a dedicated APN.

8. In the case of service categories 2 and 3, and if the charging factor chosen by the Operator is based on API invocations to enable simple time-based charging models (charging per unit of time the service is delivered where this time is not measured in the Operator's Network), the OP shall be able to include the time parameter in the charging requests to be used for charging purposes.

   The procedure used in OP to measure this service delivery time is out of the scope of this section.

9. In the case of service categories 2 and 3, and if the charging factor chosen by the Operator is based on data traffic consumption in the Operator's Network, the OP shall include the following information in the charging requests sent through the SBI-CHF:

   a) Mandatory information:

   i. API type (identification of the Service API that was invoked through the NBI e.g., QoS influence)
   ii. Correlation information: this information will allow the CCS to correlate the charging requests associated to the devices data traffic consumption received from the Operator's Network with the Service API invocation (to distinguish this traffic from the regular data traffic navigation of a customer).

The OP will also be responsible for providing this correlation information to the Operator's Network through the SBI-NR. The Operator should have in place the mechanisms to guarantee that this correlation information is provided to the CCS in the charging requests sent from the Operator's Network. This mechanism is out of the scope of this document.
The list of parameters to be included (if any) will be configurable per service and is left for further specification.

b) Optional information:

i. A subset of the parameters included in the Service API invocation. The list of parameters to be included (if any) will be configurable per service.
ii. A subset of parameters retrieved from the Network (e.g., service flow id in the Operator's Network) after the service is delivered. The list of parameters to be included (if any) will be configurable per service.

10. In the case of services in categories 3 and 4 and if the charging per Service API invocation is chosen by the Operator to enable lifecycle management API requests charging, the OP shall include the following information in the charging requests sent through the SBI-CHF:

a) Mandatory information:

i. API type and operation (identification of the Service API that was invoked through the NBI: e.g., application instantiation)

b) Optional information:

i. A subset of the parameters included in the Service API invocation. The list of parameters to be included (if any) will be configurable per service.
ii. API result code

11. In the case of services in category 4 and if the charging factor chosen by the Operator is for edge enabling infrastructure resources usage based on subscribed capacity in API request, the OP shall include the following information in the charging requests sent through the SBI-CHF:

a) Mandatory information:

i. API type and operation

b) Optional information:

i. A subset of the parameters included in the Service API invocation that include the detailed information about the resources to be reserved (independent from the effective usage):
    1) Subscribed compute capacity:
        a. vCPU
        b. Memory
        c. Network Resource Location

d. Availability Zone

2) Subscribed storage capacity:
   a. Storage
   b. Type
   c. Network Resource Location
   d. Availability Zone
3) Subscribed Network capacity:
   a. Input
   b. Output
   c. Label (internet traffic, intra-cluster, inter Edge Cloud traffic …)
4) Subscribed accelerators capacity:
   a. Accelerator name (Example: GPU)
   b. Type
   c. Network Resource Location
   d. Availability Zone

ii. A reservation time period

12. In the case of services in category 4 and if the charging factor chosen by the Operator is for edge enabling infrastructure resources usage (information about effective consumption), the OP shall include the following information in the charging requests sent through the SBI-CHF:

a) Mandatory information:

i. API type and operation

b) Optional information:

i. A subset of the parameters included in the Service API invocation that include the detailed information about the resources to be reserved (independent from the effective usage):
   1) Effective compute usage:
      a. vCPU
      b. Memory
      c. Network Resource Location
      d. Availability Zone
   2) Effective storage usage:
      a. Storage
      b. Type
      c. Network Resource Location
      d. Availability Zone
   3) Effective Network usage:
      a. Input
      b. Output
      c. Label (internet traffic, intra-cluster, inter Edge Cloud traffic)
   4) Effective accelerators usage:
      a. Accelerator name (Example: GPU)
      b. Type

c. Network Resource Location
d. Availability Zone

ii. Covered usage time period.

4.3.3.4 Security Requirements

The following security requirements shall be considered:

1. The SBI-CHF shall provide an authentication mechanism to enable access only by authenticated and authorised entities. Therefore, mutual authentication shall be provided between the OP and the Charging Engine element.
2. The SBI-CHF shall provide an authorisation mechanism to grant access to only the necessary services to which previous authorisation has been granted.
3. The SBI-CHF shall support the use of security mechanisms by its endpoints that safeguard the exchanged data's confidentiality, integrity, and authenticity.
4. The SBI-CHF shall support the adoption of strong security algorithms that guarantee forward secrecy and prevent intervening attacks such as replay, relay, and man-in-the-middle attacks.
5. An OP shall maintain security and data/topology privacy requirements when reporting federated consumption.

4.3.4 SBI-EIN

To execute operations where applications hosted on resources exposed by the OP can communicate directly with each other, an OP shall enable EIN Establishment between exposed resources, even if those are not in the same physical location. Example of such operations are:

1. Application relocation.
2. Application context relocation.
3. Application load sharing or failover handling.

The above example operations can be executed over the EIN. The OP will enable the Applications to communicate over the EIN by providing the right information and applying appropriate rules over the SBI-EIN interface.

4.3.5 SBI-OAM

The APIs exposed on the SBI-OAM interface can help an OP to determine the status of a Network Slice in its life cycle. In some cases, the OP needs to inform the Application Provider if a Network Slice status has changed or can request such change. An Operator can make a choice on which APIs to use for Network Slice lifecycle management, this is depending on where the Network Slice Management Function (NSMF) sits in the Operator’s architecture, i.e. whether this function is located outside OP.

4.3.6 SBI-AAPrM

The SBI-AAPrM is a reference point enabling interactions among the OP and the Authentication (including Identity) Management, Authorization Management and Privacy Management Functions in the CSP domain. The interactions and information flows on the SBI-AAPrM interface shall allow an OP:

1. To securely and confidentially create, and manage Application Privacy Profiles,
2. To securely and confidentially retrieve, and update Privacy Information
3. Trigger the capture of the Consent from the End-User via Privacy Management Function within an authentication and authorization context,
4. Subscribe to and request to get notifications from the Privacy Management Function related to any changes in the Privacy Information,
5. Subscribe to and get notified by the Privacy Management Function about any changes in the Application Privacy Profile(s).
6. To interact with the User Identity Token Manager function , if applicable

4.3.6.1 Southbound Interface for Privacy Management

The integration with the Privacy Management Function in the CSP domain enables an OP to verify whether a suitable legal basis allows sharing personal data with an Application (owned by an Application Provider) within an Authentication and Authorization context. Depending on the legal basis associated with the Purpose of Data Processing, an explicit interaction with the End-User (whom personal data belongs) to grant access to the protected resources (for instance for Consent legal basis) might be needed.

4.3.6.1.1 High-Level Requirements

The OP integration to the Privacy Management Function shall allow:

1. The OP to retrieve whether Privacy Informationfor a specific API call is already in place in the Privacy Management Function,
2. Cache retrieved Privacy Information and request to receive notifications from the Privacy Management Function related to the change of Privacy Information,
3. To trigger the capture of the Consent if there is no Privacy Informationin the local cache or in the Privacy Management Function,
4. When applicable, trigger an update of the Privacy Informationin the Privacy Management Function to enable the exercise of End-User privacy rights management (see Annex F),
5. The OP to validate User Identity Tokens used for identifying End-Users, if applicable.

4.3.6.1.2 Security requirements

The following security requirements shall be considered:

1. The SBI-AAPrM shall be confidentiality and integrity protected.
2. The SBI-AAPrM shall support mutual authentication between the OP and the Privacy Management Function within the CSP domain.
3. The SBI-AAPrM shall support the use of authorisation mechanisms by its endpoints that grant access to only the necessary authorised services and data.
4. The SBI-AAPrM shall support the adoption of strong security mechanisms that guarantee forward secrecy and prevent intervening attacks such as replay, relay and man-in-the-middle attacks.

4.3.6.1.3 Managing Consent

The OP shall be able to:

1. trigger the capture of the Consent when dictated by the legal basis, and when no valid Privacy Information is found in either the local cache or the Privacy Management Function.
2. receive notifications when a Subscriber revokes Consent via NBI (and forward that notification over the SBI-AAPrM and possibly EWBI in federation scenarios.)
3. receive notifications via SBI-AAPrM when a Subscriber revokes Consent (and forward that notification over the NBI and possibly EWBI in federation scenarios).
4. keep records of attempts to capture or revoke Consent from the Subscribers through Logging, Tracing and Auditing functions

4.3.6.2 Southbound Interface for User Identity Token check

Network API invocation must be secured in a way to prevent intruders from taking the identity of applications and / or devices to get hold of network and device specific information. Most use cases deal with applications that are used from a device and that leverage Network APIs that target the very same device. The mechanism for API invocation must therefore support a means of authenticating the End-User and the Application via some User Identity Token that is supplied by OP or retrieved this elsewhere in the Operator domain (e.g. entitlement server…) and that is used throughout the API invocation chain. The integration with the User Identity Token Manager in the Operator domain enables an OP to verify whether the application and user inside the token is correct. Integration with the User Identity Token Manager using the SBI-AAPrM (Privacy Management) reference point shall be supported.

4.3.6.2.1 High level requirements

The OP must support using an User Identity Token for End-User and Application authentication, which has been provided by the User Identity Token Manager function to the Application running on the device and passed through the API invocation chain to the OP platform:

1. The OP shall support application registration with globally used ids (e.g. ID of the device Application Client as known in the device vendor app store)
2. The OP shall provide the Operator ID to the User Identity Token Manager function for inclusion in the User Identity Token for call routing purposes.
3. The OP may support the User Identity Token as login hint on the different variants of authorization (e.g. Oauth 2.0, OIDC, Mobile Connect) protocol
4. The OP shall decompose the User Identity Token , extract the Network Subscription ID (e.g. MSISDN), Operator ID, and application ID and authenticate these against the registered data
5. The OP shall be able to map the device Application Client ID against the backend application invoking the API to ensure that these belong to each other.
6. Secondary devices like smartwatch share the same external MSISDN but have an internal technical MSISDN which is not revealed to the Subscriber. An OP shall be able to use this technical MSISDN for API resolution.

Note: Dual SIM situations are for FFS.

7. The User Identity Token shall contain identifiers with Operator ID in the token to allow an Application Provider or an Aggregator to route the request towards the Home OP

8. The identity of the issuing User Identity Token Manager may be obfuscated in the token for privacy reasons, in which case trusted Aggregators and APs should be able to deobfuscate the routing information for routing purposes.

4.3.6.2.2 Security requirements

The following security requirements shall be considered:

1. The SBI-AAPrM shall be confidentiality and integrity protected.
2. The SBI-AAPrM shall support the use of authorisation mechanisms by its endpoints that grant access to only the necessary authorised services and data.
3. The SBI-AAPrM shall support the adoption of strong security mechanisms that guarantee forward secrecy and prevent intervening attacks such as replay, relay and man-in-the-middle attacks.
4. The User Identity Token shall be confidentiality and integrity protected.
5. The User Identity Token shall be End-User and application specific.
6. The User Identity Token shall contain an encrypted section for End-User identifiers as well as application identity and a potentially an unencrypted section indicating the domain of the home Operator.

    a) There shall be the option to encrypt the User Identity Token section indicating the domain of the home Operator.

7. The User Identity Token may be used through the whole API invocation chain, even if there are multiple Partners between Application Backend and the OP.

4.3.6.2.3 Checking the User Identity Token

The OP shall be able to:

1. Interact with the User Identity Token Manager function for token validation over the SBI-AAPrM.
2. Optionally trigger User Identity Token invalidation in case validation has failed
3. Keep records of User Identity Tokens used on API invocations through Logging, Tracing and Auditing functions.
4. OP shall be able to interact with the User Identity Token Manager function over the SBI-AAPrM to retrieve the appropriate device/End-User identifier (such as MSISDN) and Application Identifier from the supplied User Identity Token.

4.4 User to Network Interface

The primary function of the User to Network interface (UNI) is to enable a User Client to interact with an OP, to enable the matching of an Application Client with an Application Instance on a Cloudlet exposed through an Edge Cloud service.

5 Detailed Requirements on functional elements

This section defines the requirements of the functional elements that make up the OP architecture.

5.1 Exposure Functions

5.1.1 High-level requirements

The Exposure Functions serve as intermediary layer between the Application Provider and the Leading OP and transitively to those OPs federated with the Leading OP. To carry out this function, it shall satisfy the requirements listed below.

Note: In some cases, a requirement associated with the Exposure Functions specifically applies to its endpoint to the Application Provider, i.e. the NBI. In those cases, the requirement will be specified for the NBI.

1. The Exposure Functions shall present an information model to the Application Provider that is consistent among the Leading OP and the Partner OPs federated with it.
2. The Exposure Functions shall support a secure means of authentication and authorisation, operating over the NBI.
3. The Exposure Functions shall support a common model for telemetry data (i.e., data arising from resource monitoring) and a means of configuring telemetry data collection.
4. The telemetry system should be consistent with the SBI-CHF interface of section 4.3.3.

5.1.2 Security Requirements

The following security requirements shall be considered:

1. The Exposure Functions shall provide an authorisation mechanism to grant access to only the necessary authorised services and data. The security enforcement point is the NBI API Gateway.
2. The Exposure Functions shall provide a fine-grained authorisation mechanism to grant authenticated entities selective access to the NBI exposed services and functionalities.
3. The Exposure Functions shall use profile-based access control to provide appropriate restrictions on the amount of functionality that the OP offers to a particular Application Provider, according to their operational profile and the type of access requested.
4. When defining and assigning the authorisation profiles, the Exposure Functions shall apply the principle of least privilege, ensuring that any entity should have only the minimum profile roles necessary to perform its function.
5. Given the external exposure of the NBI, the Exposure Functions shall provide security mechanisms to counteract/prevent attacks aimed to undermine the availability of the NBI, such as Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks, reconnaissance attacks (attempts to identify service or API vulnerabilities) and brute force attacks.
6. The Exposure Functions should provide isolation between resources of different Application Providers (e.g. when providing telemetry data or when accessing and managing Edge Applications configuration data).
7. The Exposure Functions should provide security mechanisms to protect accounting and guarantee safe logging (e.g. integrity, non-repudiation, etc.) of the activity over the NBI.

5.2 Federation Functions

5.2.1 Federation and Platform Interconnection

5.2.1.1 General

One of the Operator Platform's primary purposes is offering Application Providers an extended Operator footprint and capabilities through interconnecting with other Operators' resources and Subscribers. This capability is achieved by the federation E/WBI interface; to interconnect OPs belonging to different Operators, enterprises or others.

The communication between federated entities shall support a distributed tracking mechanism that allows end-to-end tracking across these federated entities. For example, requests may contain identifiers that are propagated and used in every communication.

5.2.1.2 Authentication/authorisation

Federating OPs are likely to belong to different entities in different security domains. Therefore, the capability to exchange authentication and authorisation between federated OPs is required:

1. There shall be a mechanism to register and authenticate different OP instances.
2. An OP shall be able to identify unequivocally any federated OP instance.
3. An OP shall be able to authorise a registration request from another OP instance.
4. An OP shall exchange a token or "federation key" on the association handshake, identifying each federation integration.
5. User authentication/authorisation shall remain independent from the OP to OP authentication/authorisation.

5.2.2 Settlement

Federation interfaces shall expose management and settlement data. This data allows the charging systems of each Operator to account for the services consumed.

1. An OP shall share usage statistics through the E/WBI for the services requested by the federated connection.
2. An OP shall provide any needed information that is useful for billing/settlement among Operators, e.g.:

    a) Type of resources used;
    b) Quantity of resources employed on the service.
    c) The number of Application Instances used.
    d) The number of user sessions served.
    e) Usage time of the resources.
    f) Additional services employed, e.g. network location query.

These services will be provided over the SBI-CHF where the CDRs generated by the Leading and Partner Ops Charging Engine are input to settlement and reconciliation processes outside of charging and hence not in scope. Reference to diagram flows in section G.3 of this document are provided for clarifications.

5.2.3 Resources management via interconnection

One of the essential points to be solved through the federation interfaces is sharing the Resource Catalogue between instances.

1. An OP shall allow the Operators/resource owners to select the capabilities and resources to be shared via federation.
2. An OP shall be able to share the exposed network-related capabilities.

5.2.4 Security Requirements

The following security requirements shall be considered:

1. The Federation Functions shall provide an authorisation mechanism to grant access only to the necessary authorised services and data for a Partner OP. The security enforcement point is the EWBI API Gateway.
2. The Federation Functions shall provide security mechanisms to counteract attacks aimed to prevent the availability of the E/WBI, such as DoS attacks
3. The Federation Functions should provide security mechanisms to protect accounting and guarantee safe logging (e.g., integrity, non-repudiation, etc.) of the activity over the E/WBI.

5.2.5 Routing of Requests

When having relationships with one or more Partner OPs the following requirements apply

1. The OP shall be able to determine whether it needs the support of a Partner OP to provide a service.
2. When needing the support of a Partner OP, the API Federation Management Function (depicted in Figure 2) shall be able to determine what E/WBI the OP needs to use to reach that Partner OP.
3. For cases where routing depends on the Subscriber or UE to which the service request relates, the API Federation Management Function shall be able to identify the appropriate Partner OP and E/WBI based on

    - a public IP address through which the UE or Subscriber is identified,
    - the MSISDN associated to the subscription, and
    - A network-specific token that identifies the Subscriber (e.g. an external GPSI including a domain).

4. For cases where routing depends on network or cloud resources to which the service request relates, the API Federation Management Function shall be able to identify the appropriate Partner OP and E/WBI based on

    - Identifiers for the Availability Zone where the resources would be located.

5. For this identification, the API Federation Management Function shall take into account the information provided by the Partner OPs over the E/WBI that the OP has with them (see section 4.2.1).
6. An OP's API Federation Management Function shall update Partner OPs on changes in the resource identifiers for which the OP can offer services (e.g. IP address ranges for Subscribers or UEs, Availability Zones offering Edge Resources).

5.3 Transformation Functions

No general requirement have been identified so far for the Transformation Functions in Figure 2.

Note: Future general requirements for the Transformation Functions (e.g., identified in other groups), may be transposed into this document.

5.4 Integration Functions

5.4.1 Service Availability on Visited Networks

5.4.1.1 General

Service availability on visited networks shall be considered to allow the users to use a service provided through the OP when outside of their Operator network. This condition includes international situations and the inter-operator handovers that occur, for example, when connecting to the End-User's home Wi-Fi network, which a different Operator may provide.

5.4.1.2 Requirements

1. When a device first attaches to a visited network, there shall be messaging between the User Client, Home OP and Visited OP. The messaging's purpose is for the Home OP to authenticate the User Client and authorise it to use the Edge Cloud and Network Capabilities on the Visited OP.

    a) The messaging shall not be repeated for each application session or each application.
    b) The authorisation shall be valid for a finite period.
    c) The Home OP and Visited OP shall have a separate process to agree about charging /settlement for the use of Cloudlets by UEs of the Home OP. It is not the intention to define a granular charging /settlement mechanism ("granular" meaning, for example, per UE or per Application Instance).

2. User plane LBO/SBO shall be available for the UE in the visited network.

    a) If no LBO/SBO is available or there is no service availability agreement among Operators, the UE receives service from home resources and Home OP without Visited OP interaction.

3. The Visited OP shall match the Application Provider's requirements on Network Capabilities to the exposed capabilities in the visited Operator network.
4. The Visited OP shall be able to provide the abstract application Service and Session Continuity capabilities over the E/WBI for roaming users to their Home OP

    Note: UE mobility management is handled with existing mobility management mechanisms.

5.4.2 User Mobility support

5.4.2.1 General principles

A mobile Subscriber actively engaged with an Application may, during their movement from one place to another (i.e. User Mobility), trigger changes in network connectivity (i.e. a different network attachment point for their session). This may affect the services based on Operator Capabilities provided by the OP that the Application relies on to support its End-User. This is due to various network access factors like poor radio connectivity, network congestion, etc.

The impact of these connectivity changes on an Application is dependent on the nature of the Application and the Operator Capabilities that it relies on. Some categories of Applications (e.g. video streaming)may be able to maintain a seamless user experience despite interruptions in connectivity through application domain-specific algorithms while for other categories, e.g. gaming applications relying on low latency, such interruptions may affect the user experience significantly. Such an Application might need uninterrupted transport-level Session Continuity for a TCP session with improved QoS for example.

An OP shall provide support to enable an Application to provide a consistent user experience during User Mobility.

As general principles, the following are essential requirements to provide Application-level continuity in the OP architectural model:

- An OP shall rely upon the 5G core network capabilities for supporting continuity of the services that it provides to Applications.
- An OP, based on the network capabilities, shall expose abstract continuity models for services requiring that towards the Application Provider or an Aggregator over the NBI interface
- An OP shall interact with the mobile network and the 3GPP-defined standard services over the SBI-NR interface to synchronise with the 5G core network procedures to support service continuity.
- When required, an OP shall inform Applications about upcoming or past User Mobility events.

5.4.2.2 Access technologies support for service continuity

To support the continuity of the services offered, an OP may rely on the core network’s Service and Session Continuity (SSC) capabilities. The SSC capabilities in a mobile network depend considerably on the type of the radio network, i.e. 4G, 5G, Wi-Fi etc. and on the support for Session Continuity defined for these networks in standards like the 3GPP's. It also depends on whether the Operator has deployed such services for their Subscribers.

Note: The abstract Service and Session Continuity modes corresponding to 3GPP defined SSC modes 1, 2 and 3 are typically described as “IP Preservation”, “Break-Before-Make”, and “Make-Before-Break" respectively.

Depending on their access hardware and software capabilities, UEs may attach to mobile networks following the access policies configured for the subscription and network capabilities deployed and operated by the mobile service providers.

The UE may perform its network attachment to the radio networks available in the UE's location. Those networks could be broadly segregated into 3GPP or non-3GPP (trusted or untrusted) access technologies. As part of the SIM configuration, an Operator can configure their preference for the selection of access technologies to the UE. The network to which a UE is currently attached would also determine the level of support available for Session Continuity in that network what an application can expect.

Handovers and associated SSC procedures may be triggered by the mobility of UEs within the mobile network coverage area. These procedures or capabilities are defined for devices attached to a mobile network using 3GPP's 5G radio technologies. Table 2 describes the SSC that an OP shall support in the current version of this document when 5G capable UEs attached to a 5G radio network are served by the 5G core network (i.e. 5G Standalone (SA)).

	Support in Home NW	Support in Visited NW
5G to/from 5G	Supported	Supported
5G to/from 4G	Supported	Supported
5G to/from non-3GPP trusted access	Not Supported	Not Supported
5G to/from non-3GPP untrusted access	Not Supported	Not Supported

Table 2: Access Technologies Supported In OP Architecture For Service Continuity

Note: For the above scenarios where an OP supports SSC, the cases involving mobility from one Operator network to another Operator's network are for future study.

Note: For non-3GPP access technologies, the SSC capabilities continue to evolve and, therefore, are not supported.

5.4.2.3 Network and OP responsibilities for service continuity

Assuming a Subscriber actively engaged with an Application starts moving in a network operated by their home Operator, this may result in network procedures to reselect a network attachment point for the UE to maintain the requested services (e.g. QoS Performance Profiles).

The mobile core network may activate SSC mode (starting with 3GPP Release 15 for 5G's Standalone Architecture (SA)) specific procedures based on the user's subscription and the network policies defined by the Operator.

Due to the SSC mode procedures execution in the core network, the following events may occur that require external entities to take application-specific actions:

1. For SSC mode 1, which could be named as "IP preservation mode", in which the network may assign a different attachment point while keeping the IP address for the UE unchanged:

- The mobile network may assign SSC mode 1 to a PDU session considering factors such as user subscription information, Operator configured local policy, an indication from authorised Application Functions (AF), e.g. an OP, if a PDU session cannot be relocated (application relocation indication) and the address should be preserved
- In this situation, the mobile core network may be unable to provide the desired service needed by the application (e.g. QoS Performance Profile).
- In such cases, an OP should have access to information related to the network attachment point change (user plane reconfiguration) event and the service (e.g. QoS Performance Profiles) that the network can provide for the UE PDU session
- An Application may need to adapt its behaviour according to the service (e.g. QoS Performance Profile) that the network can deliver end-to-end. If the mobile network cannot maintain the requested service during the mobility period, then based on notifications from the NEF, an OP can timely notify the events to the Application. Applications can gracefully adapt their behaviour based on notifications, e.g. switching to a lower frame rate for video streaming. An OP may also allow Applications to request to be notified about this kind of event, allowing them to take appropriate actions to provide consistent quality of experience to their users.
- An OP shall also publish over NBI the monitoring information regarding the change in service behaviour (e.g. new QoS Performance Profile for the application sessions).

2. For SSC mode 2, which could be named “Break-Before-Make” mode, the network may change the existing user plane and assign an optimum user plane in the new location of the UE, which would cause the IP address of the UE to change. It may be possible for the mobile network to provide the desired service (e.g. QoS Performance Profile) as needed by the application without preserving the Session Continuity

    - An OP should have access to information related to the user plane change preparation event for the UE PDU session in the mobile network via notifications related to user plane change events requested over the SBI-NR interface
    - An OP could use these events to notify the Application to be prepared for a possible connectivity break (e.g. a possible application session context relocation).
    - Based on interacting with the application and ensuring that any application/service specific actions have completed, an OP should relay the completion of the Application-level relocation procedures to the core network via the SBI-NR interface and the NEF (as per 3GPP NEF specified procedures).
    - On receiving the UE user plane change progress indication over the SBI-NR interface, the OP, in response to the network, shall provide the SBI-NR API parameters, e.g. description of the traffic steering rules for the application traffic, QoS reference, a period of time or a traffic volume, etc. to the mobile core network over the SBI-NR interface to ensure continuity of the services provided related to the UE-s Application traffic

Note: It is important to note that 3GPP specifications do not put any time constraints for external AFs to respond to the core network notifications and acknowledge the application's readiness for the session/context relocation.

Therefore, any OP implementation shall follow the behaviour described in 3GPP specifications and treat the acknowledgements towards the core network independent of any specific OP procedures, e.g. session/context relocation, application instantiation etc.

Note: It is important to note that due to user mobility in mobile networks, events like a user plane change may result in a UE IP address change managed by the core network. Similarly, circumstances outside the mobile network could change Application endpoints, i.e., an IP address change triggered by the OP. Any implementation of an OP that supports application SSC will need to consider such aspects from both the application's and UE's perspective.

3. For SSC mode 3, which could be named “Make-Before-Break” mode, the network may, similarly to SSC mode 2, assign a different user plane to UE due to its mobility. This user plane change would cause a modification of the UE's IP address later. However, in this mode, UE application traffic can still use the existing connection in the meantime.
    - It may be possible for the mobile network to provide the desired service (e.g. a QoS Performance Profile) needed by the Application and more time for the OP to facilitate the migration of the service to the new anchor point and synchronise this transfer for stateful applications. An OP shall have the mechanisms to minimise the time simultaneous sessions remain active to optimise the network and compute resources.
    - An OP shall relay completion of all the relocations tasks to the mobile network over the SBI-NR interface, allowing the network to reclaim the Network Resources of the previous session and start steering the UE traffic towards the new anchor point.

SSC Modes	Key Characteristics	Capability Name	Network Capability Description	Key Mobility Events Handling in OP
1	UE IP Preserved	IP Preservation	Preserve UE IP agnostic to user location change for active sessions	• Request notifications on UE Mobility events over SBI-NR • Monitor application session QoS • Enforce Application Provider policies and Operator defined policies

SSC Modes	Key Characteristics	Capability Name	Network Capability Description	Key Mobility Events Handling in OP
2	UE IP Not Preserved	Break-Before-Make	PDU session modification with a new PDU Session Anchor(PSA) and IP connectivity disruption	• Request notifications on UE mobility events over SBI-NR • Coordinate OP activities, e.g., application session relocation in synchronism with mobile network • Enforce Application Provider policies and Operator-defined policies • Notify Session Continuity events to the Application
3	UE IP Not Preserved, Concurrent Sessions	Make-Before-Break	PDU session modification with a new PDU Session Anchor(PSA) and with simultaneous connectivity with the previous session anchor	• Request notifications on UE mobility events over SBI-NR • Indication of simultaneous connectivity temporarily maintained for the source and target PSA based on app criteria • Coordinate OP activities, e.g., application session relocation in synchronism with mobile network • Enforce Application Provider policies and Operator defined policies • Notify Session Continuity events to the Application

Table 3: Summary of OP responsibilities for supporting 3GPP-defined SSC modes

Note: It is expected that to support Session Continuity in 5G mobile networks, the Operator would need to support features like UL CL (Uplink Classifier) or IPv6 multi-homing as defined by 3GPP for the UPF

Note: Based on some of the events on the SBI-NR interface, e.g. location monitoring events, QoS status notification events etc., an OP may determine the QoS Performance Profile provided by the mobile network to application sessions against the QoS Performance Profile requested by the application. In such cases, the OP may initiate the user plane relocation (e.g., by using Traffic Influence APIs) services on the SBI-NR interface. Possibly this may result in the triggering of session mobility procedures in the mobile network.

5.4.2.4 Void

5.4.2.5 Void

5.4.2.6 Session Continuity Support for Roaming Users

An OP shall support services related to roaming users when they roam into locations served by a Partner OP.

To provide these services for roaming users, the Partner OP shall provide the following information to the Leading OP over E/WBI interface (Not an exhaustive list),

- Supported Abstract Session Continuity Modes (as described in section 5.4.2.1)
- LBO Capability
- Supported Service APIs
- Relocation Failure Events
    - Network attachment relocations denied by OP
    - Network attachment relocation execution failures and causes

Note: Some network capabilities and applications relocation event monitoring information shared by a Partner OP over E/WBI can be published over the NBI to inform application providers on the Partner OP capabilities before deploying the applications. This information can be helpful if the applications are sensitive to Session Continuity capabilities supported by the Partner OP.

5.4.2.7 Application Session Continuity Support for handovers between 4G and 5G

An OP shall support continuity of the offered services for Applications when the user devices support both 4G and 5G capabilities. The mobile network may provide the interoperability between 4G and 5G for UEs that support both 5GC Non-Access-Stratum (NAS) and EPC NAS and may also offer the network capability exposure APIs based on combined SCEF+NEF via CAPIF (see section 4.3.2.3.1).

An OP shall request notifications on the SBI-NR to be informed about the expected level of support for network services or network capability exposure APIs. Based on the UE's serving network, the OP shall use these APIs as per the level of support available.

For devices attached to the 5GC with SSC mode 1 or in the EPC with an IP preservation session, an OP shall request notifications on the SBI-NR for the Core Network (CN) type (EPC, 5GC) change events for the PDU sessions used by Applications that are mobility sensitive. An OP shall interact with the mobile network to monitor the QoS level provided by the mobile network for a given PDU session.

Depending on the monitored QoS level notified over the SBI-NR, an OP may provide the Application Provider requested Alternative QoS Performance Profiles to the mobile network over the SBI-NR. These Alternative QoS Performance Profiles are defined in relation to a CN type. An OP shall determine the set of QoS Performance Profiles according to the CN type that the UE is attached to. Also, based on the OP receiving notifications related to QoS level change events for a PDU session, QoS level information as received over the SBI-NR shall be made available over the NBI to the Application Providers.

In the scenario when a user with an application session in a 5G network with SSC mode 2 or 3, is handed over from the 5G to a 4G network it may not be possible for an OP to ensure seamless Session Continuity. To support these scenarios, the SBI-NR should provide early notifications during the 5GC to EPC handover initiation process. The OP shall use these notifications to inform Applications that requested those notifications of the upcoming handover allowing those to take appropriate application-level actions to ensure the most optimal user experience.

5.4.3 Security Requirements

The following security requirements shall be considered:

1. The OP shall provide security mechanisms to counteract attacks on the OP's Southbound Interfaces (i.e. the SBI-CR, the SBI-NR, the SBI-CHF, the SBI-EIN and the SBI-OAM) aiming to prevent data availability, such as DoS attacks.
2. The OP shall protect Personally identifiable information (PII) of Subscribers while in storage.
3. Privacy and tracking protection: Information originating in the UE should be protected for integrity, privacy, confidentiality, nonrepudiation.
4. The OP shall provide an authorisation mechanism for the UNI requests to grant access to only the previously authorised services. The authorisation mechanism shall ensure that the EC is authorised to access the provisioned services and that the UE can access the edge data network.
5. Given the external exposure of the UNI, the OP shall provide security mechanisms to counteract attacks on the OPs UNI aimed to prevent the availability of the interface, such as DoS or DDoS attacks.

5.5 User Client

Detailed requirements on the User Client will be provided in a future version of this document.

5.6 Common Functions

5.6.1 Authentication, Authorization and Privacy Management

The following requirements for the Authentication, Authorization and Privacy Management Functions (introduced in clause) shall be considered

5.6.1.1 High-level requirements

The following high-level requirements for the Authentication, Authorization and Privacy Management Functions has been identified:

1. The Authentication, Authorisation and Privacy Management function may trigger an update on the Privacy Information whenever there is a change on the status of the Vetting Information. Depending on the specific case, that change may also imply authorisation information (e.g., access token) invalidation and triggering Application Provider / Application offboarding procedures.

   For example, Application Provider Vetting Information may change in the below cases.

a) If Application Provider 1 is no longer a valid legal entity (e.g., due to bankruptcy), then the Privacy Information on the CSP domain is no longer valid (for all the Application Provider 1's Applications), previously issued authorisation information should be removed, and the Application Provider 1 and all its Applications shall be offboarded, or
b) If Application Provider 1 is absorbed by Application Provider 2 (e.g., due to mergers), then the following steps are dependent on local choices:
    i. If the Application Provider 2 Vetting Information was not yet verified, the the Vetting Process needs to be done.
    ii. If the Application Provider 2 Vetting Information was already verified, then all / some of the Application Provider 1's Applications need to be associated to Application Provider 2.
    iii. If the Privacy Information on the CSP domain is considered no longer valid, Consent (if applicable) shall be captured to have a legal basis in place to enable the OP sharing information with Application(s) of Application Provider 2.

2. The Authentication, Authorisation and Privacy Management function shall validate the Application Provider's ownership and correct format of the redirection URL(s) before capturing Consent with authorization flows that imply redirection. The redirection URL(s) are provided to the OP during Application onboarding [22].

3. The Authentication, Authorisation and Privacy Management function shall follow the guidelines and recommendations for authorization flows in CAMARA [24].

Annex A Deployment Scenario

This section provides an overview of deployment options of an Operator Platform.

A.1 Relationship with OP and Operator

An OP's deployment scenario can have two options depending on whether each Operator has its OP.

In Figure 4, the OP manages at least the resources of a single Operator. OP A run by Operator 1 can federate with OP B run by Operator 2.

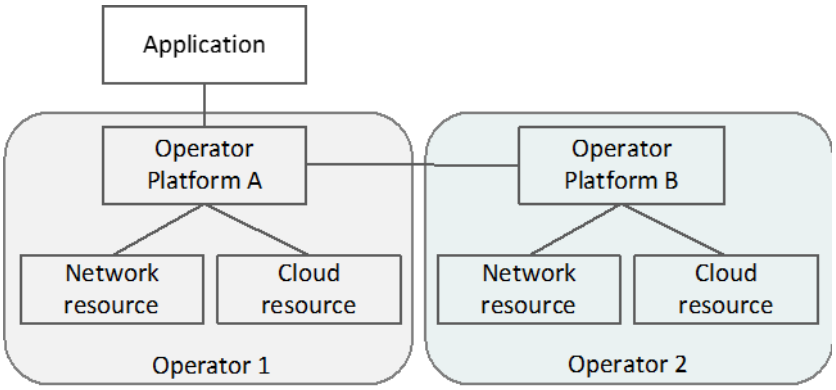

Figure 4: Each Operator has an own Operator Platform

In Figure 5, an OP manages multiple Operators' resources. Because one OP manages the resources of multiple Operators, when receiving a federation request from OP B or a deployment request from an Application Provider, Operator 1 or Operator 2 is selected based on OP A's policy.

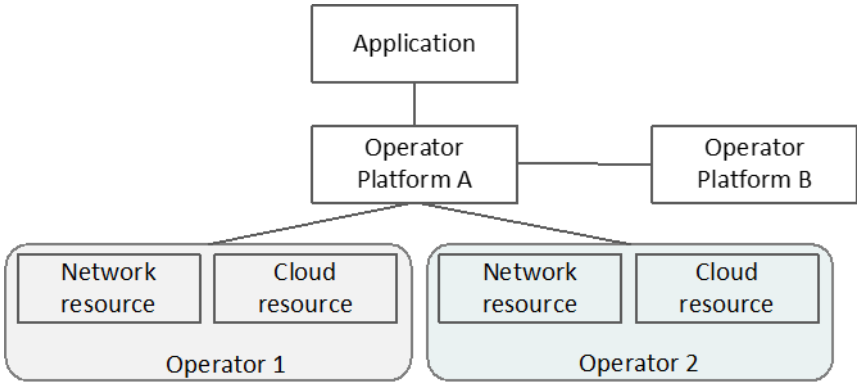

Figure 5: Multiple Operators share the same OP

A.2 Relationship with hyperscalers from a single Operator perspective

An Operator can have their own cloud resource and collaborate with a hyperscaler simultaneously. An OP can integrate hyperscalers with the same features as it does with its own cloud resources and support APIs of hyperscalers.

There are two ways for Hyperscaler integration via an OP. First, hyperscalers can be considered enterprise customers to the OP and can interact via the NBI. The second is that hyperscalers can implement an OP and become a Partner connecting via the E/WBI.

The SBI-CR is likely to match the interface that hyperscaler is exposing to its direct enterprise customers (i.e. Application Provider 2). In addition, Hyperscaler resources can be available for OP A to offer its enterprise customers (i.e. Application Provider 1).

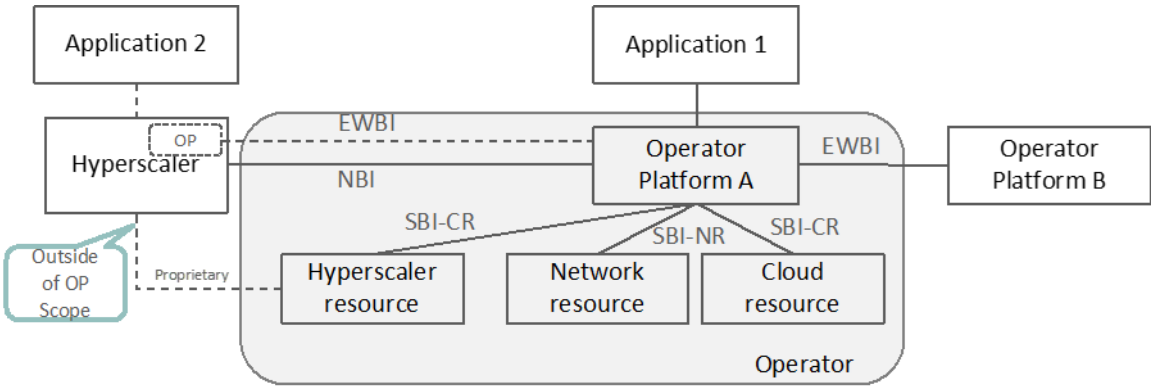

Figure 6: Relationship with hyperscalers

Annex B Aggregation / Marketplace Platform

An Aggregation/Marketplace Platform is a store the services exposed by the Operators via an OP. In addition, there may be OPs that offer additional services beyond those specified in this document, for example, specialist AI or media encoding. The purpose of the Aggregation/Marketplace Platform would be to enable Application Providers to discover such additional services and possibly buy them.

The following are potential functionalities supported by Aggregation/Marketplace Platform:

- authenticates and authorises Application Providers
- aggregates the additional APIs offered by OPs and exposes them to Application Providers
- receives requests from Application Providers for the additional services and requests the appropriate OP to fulfil them
- provides a public repository for storage and validation of the application package for edge that Application Providers upload for deployment.

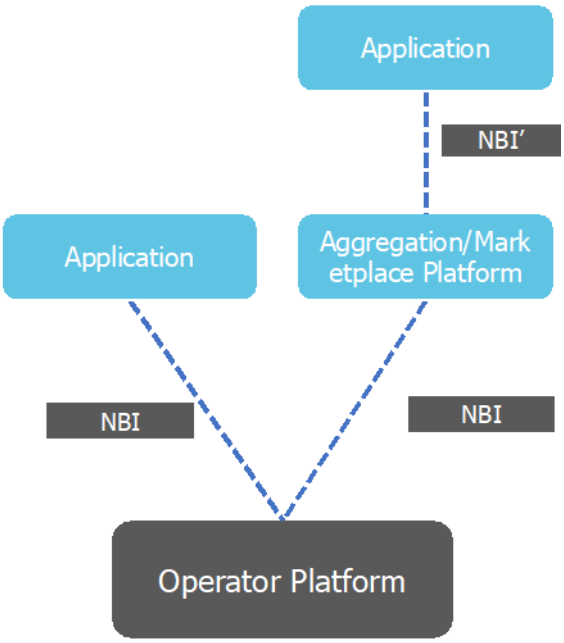

Figure 7: Operator platform with Marketplace

Annex C Operator Platform Security

C.1 Guidance for the implementation, deployment and operation

Some threats identified in this Annex cannot be mitigated through the OP's architecture and interface definitions. Therefore, this section provides guidance for the implementation, deployment and operation of an OP and the Operator Capabilities that it exposes. The following guidance is to be taken into account at a high-level:

1. The implementation and deployment of an OP needs to use operational procedures to carry out security hardening. This hardening includes, e.g., auditing to ensure that software patches are up to date, publishing regular security audits.
2. An OP implementation needs to apply protection mechanisms to ensure service availability to prevent attacks targeting the availability of exposed applications/services, e.g., denial of service attacks and brute force attacks.
3. An OP implementation is recommended to support telemetry for intrusion detection.
4. An OP deployment and its operation are recommended to follow best practices for DevSecOps (i.e., the practice of introducing security practices into DevOps), as described in GSMA FS.31 [5].
5. An OP implementation needs to employ telemetry and analytics to detect and report application security policy violations at runtime to localise and isolate malicious application behaviour.
6. An OP implementation needs to employ telemetry and analytics to detect DDoS attacks against the network and enable rate-limiting and traffic isolation in network segments and endpoints.
7. An OP implementation is recommended to support hardware-root-of-trust (e.g. Trusted Platform Module) based security keys for platform integrity checks, mutual authentication, and the establishment of secure tunnels with Application Providers.

Note: A future phase of this work will investigate defining security levels between Operators.

8. An OP implementation is recommended to support a secure Domain Name System (DNS) service to avoid attacks that exploit DNS, such as impersonation attacks.
9. An OP implementation is recommended to enable resource isolation, sharing authorisation, and residual data clean-up to protect shared Network Resources/slices from tampering and data theft.
10. An OP implementation is recommended to employ message filtering of HTTP control plane signalling and firewall configurations to protect Network Resources from spoofing attacks from roaming interconnections.
11. An OP deployment is recommended to enable security audits on the access privilege management to avoid identity theft or fraud.
12. An OP implementation is recommended to employ secure storage of account credentials to avoid identity theft or fraud.
13. An OP implementation needs to employ secure initialisation and secure configuration data storage to avoid the exploitation of network configuration data weaknesses.
14. An OP deployment should provide hardware root-of-trust based tools to guard network configuration status.
15. An OP deployment is recommended to support centralised and unified log management to protect from any tampering, whether malicious or inadvertent,
16. An OP implementation is recommended to support the automation of security operations.
17. An OP implementation needs to provide secure tracing and logging of charging and billing data requests.

Annex D Void

Annex E Service and capability exposure charging concepts

As described in section 2.2.4, the Operator Platform architecture needs to allow Operators to charge for the services and capabilities that are exposed by that Operator to Application Providers, Subscribers, and other Operator Partners.

Any decision relating to charging and/or billing for the usage of the services as described in this Annex is for an individual Operator to decide.

A set of technical requirements are necessary to enable these charging and billing capabilities. These technical requirements will support potential commercial models defined by Operators –for federation and towards end customers/developers.

Note: The definition of commercial models is out of the scope of this document.

An Operator Platform exposes different Operator’s services and capabilities to third parties. Although this set of services and capabilities is quite heterogeneous and is in constant evolution, it is possible to establish a classification of these services/capabilities from a charging perspective. The following service categories can be considered:

- Network capabilities exposure services with no impact on the device’s data usage.

- Network capabilities exposure services with impact on the device's data usage.
- Network provisioning services.
- Edge Application management services.

A detailed description of these categories together with examples of potential charging factors used for services/capabilities will be provided in the next sections of this Annex.

In addition to the categories listed above, there is one more that can be considered that groups "General purpose services" into its own category. This category would include the set of services/capabilities that are exposed by the Operator as "enabler" services (e.g., to manage the connection from the Application Provider to the OP, to manage permissions/consents, etc.). This category may require generation of file records (e.g., XDRs) that could be used by the Operators for charging and/or reporting purposes.

E.1 Network capabilities exposure services: with no impact on device's data usage

This category includes the group of services that are consumed by the Application Providers to access the capabilities exposed by the Operator's Network and that have no impacts on the device's data traffic usage as a result of the service invocation. These services are normally used to get information from the Operator's Network and some potential examples are:

- Network information retrieval related services: for example, to get or verify the location of a device that is registered in the Operator's Network, to get or check the device's registration status, to be informed about a device's location changes etc.
- Services to receive notifications related to analytics information provided by the Operator's Network.

From a technical perspective, the following potential charging factors, supporting charging and billing, could be used:

- Service activation charging:

  A one-time fee or a periodical fee (e.g., a monthly fee) could be charged by an Operator to third parties. This fee would enable the access to a particular service (different fee per service/group of services). This fee will not be dependent on the service usage.

- Charging per API invocation received:

  In this case charging would be based on the Service API that is invoked by the Application Provider. Depending on the Operator's decision, this charging factor would allow the Operator to charge based on:

  - The particular API (operation) that is invoked by the Application Provider, without considering the parameters in the payload included in the Service API invocation.
  - The particular API (operation) that is invoked by the Application Provider and considering some parameters included in the Service API invocation (selected API payload).

Note that in this case only a subset of parameters, that will be dependent on the service, would be considered (e.g., in a device location service request, the precision included in the API payload could be used to use that level of precision as a potential parameter to consider in the rating and charging).
The reason for considering only a subset of the parameters is to avoid unnecessary complexity and potential latency/dimensioning issues.

This charging factor would allow the Operators to have the possibility to do the charging and billing based on:

- The number of API invocation requests for Network information retrieval received (e.g., Charging per device location query request received)
- The number of API invocation requests for a notification service received (e.g., Charging for requests to receive notifications from an analytics information service during a period of time)

- Charging per notification sent to the Application Provider (as a result of a request for such notifications):

  In this case charging would be based on the type of notification that is sent (e.g., Charging per analytics information notification delivered to the Application Provider)

The list of charging factors are the potential ones that the Operator can choose to support the commercial models for the services included in this category.

The related technical requirements that need to be supported by the Operator Platform for these charging factors are described in section 4.3.3 of this document.

E.2 Network capabilities exposure services: with impact on device's data usage

This category includes the group of services that are consumed by the Application Providers to access the capabilities exposed by the Operator's Network and that have an impact on the device's data traffic usage. Some potential examples of these services are:

- Services that influence the device's QoS (e.g., to request a specific QoS Performance Profile – 'High' QoS – to be delivered to a specific PDU data traffic session of a device)
- Services that allow sponsorship of data traffic usage (e.g., A particular PDU data traffic session of a device is sponsored by an Application Provided)
- Services that influence how the data traffic of a device is steered in the Operator's Network.

From a technical perspective, the following potential charging factors, supporting charging and billing, could be used:

- Service activation charging:

  A one-time fee or a periodical fee (e.g., a monthly fee) could be charged by an Operator to third parties for enabling the access to a particular service (different fee

per service/group of services). This fee would not be dependent on the service usage.

- Charging per API invocation received:

  In this case charging could be based on the Service API that is invoked by the Application Provider.
  As in the previous category, depending on the Operator's decision, charging can be based on the operation that is invoked (API type) or on a combination of the operation invoked and a subset of parameters included in the API invocation payload.
  Through this charging factor, the Operators would have the possibility to use time-based charging models to do the charging and billing of a service (e.g., charging per unit of time that a particular QoS Performance Profile is provided to a device/PDU session)

- Charging based on data traffic usage as a result of a previous service invocation:

  In this case charging could be based on the data traffic consumption of a device in the Operator's Network as a result of a previous Service API invocation (e.g., charging per each unit of traffic volume that is carried over an active QoS Performance Profile session)
  Using this charging factor, it would be possible to enable volume-based charging models to do the charging and billing of a service.
  The feasibility of using this charging factor depends on the ability to correlate the data traffic in the Operator's Network that is impacted by a particular API call with the API invocation.

The list of charging factors described above are the potential options that an Operator could use to support the commercial models that an Operator chooses to carry out the charging and billing for the services included in this category.

The technical requirements that need to be implemented by an Operator Platform to support each of these charging factors are described in section 4.3.3 of this document.

E.3 Network provisioning services

This category includes the group of services that are consumed by the Application Providers to manage different aspects of Network Services Provisioning in the Operator's Network.

In this category, the Application Providers are also accessing services and capabilities provided by the Operator with impact on the devices data traffic. The main difference compared to the previous category (*Network capabilities exposure services: with impact on device's data usage*), is that the exposition of these services requires previous provisioning activities in the Operator's Network (e.g., to provision a particular APN or Network Slice Instance in the Operator's Network).

Note: The Operator's BSS/OSS should be involved during these services provisioning flows. How this is done is Operator-dependent and out of the scope of this document.

From a technical perspective, the following potential charging factors, supporting charging and billing, could be used:

- Service activation charging:

  A one-time fee or a periodical fee (e.g., a monthly fee) could be charged by the Operator's to third parties for enabling the access to a particular service (different fee per service/group of services). This fee would not be dependent on the service usage.

- Charging per API invocation received (service lifecycle modification):

  In this case charging would be based on the Service API that has been invoked by the Application Provider and, depending on the Operator's decision, charging could be based just on the operation that is invoked or a combination of the operation invoked, and a subset of parameters included in the API invocation payload.
  Some potential examples of this charging factor include:

  - Charging per service allocation/deallocation operation received (e.g., Charging per operation received to assign an Application Provider to an existing Network Slice Instance, NSI)
  - Charging based on the number of devices that are provisioned with a particular Network service (e.g., charging per number of devices that are using a particular APN based on the number of subscription/unsubscription operations received)

- Charging based on data traffic usage:

  In this case charging would be based on the device's data traffic usage using a particular Network Service (e.g., an APN or an NSI).
  Depending on the Operator's decision, the following charging models could be used:

  - Time based charging: charging per each unit of time that the devices are using a particular Network Service (e.g., charging per each unit of time devices are using an APN)
  - Volume based charging: charging per each unit of traffic volume that devices are consuming using a particular Network Service (e.g., charging per each unit of traffic volume that is carried over an APN)

  The feasibility of enabling this charging factor depends on the ability to correlate the data traffic in the Operator's Network that is impacted by a particular API call with the API invocation.

The technical requirements that need to be implemented by an Operator Platform to support each of these charging factors are described in section 4.3.3 of this document.

E.4 Edge Application management services

This category includes the group of services that are consumed by the Application Providers to manage their applications and is relevant only if the OP exposes such services. Some potential examples are:

- Resources reservation services.

- Application onboarding and lifecycle management services.
- Application resources monitoring services.

From a technical perspective, the following potential charging factors, supporting charging and billing, could be used:

- Service activation charging:

  A one-time fee or a periodical fee (e.g., a monthly fee) could be charged by the Operator's to third parties for enabling the access. This fee would not be dependent on the service usage.

- Charging based on Edge enabling infrastructure resources usage:

  Depending on the Operator's decision the following possibilities could be considered:

  - Charging based on the subscribed capacity requested by the Application Provider in the API request (a subset of the parameters included by the Application Provider in the API request payload).

    As an example, the following resources could be considered to do the rating and charging: number of vCPUs, memory, storage, incoming/outgoing data volume, etc.

  - Charging based on real resources usage (based on periodical information retrieved from the Operator's Network)

- Charging based on Application lifecycle management API requests:

  In this case charging could be based on the type of operation that is invoked (instantiation/upgrade/termination) and potentially, depending on Operator's decision, a subset of the parameters received in the API payload (e.g., Availability Zone).

- Charging based on data traffic usage in the Operator's Network:

  In this case charging would be based on the data traffic consumption of the devices in the Operator's Network accessing an Edge Application.
  The feasibility of enabling this charging factor depends on the ability to correlate the data traffic in the Operator's Network that is impacted by a particular API call with the API invocation.

The technical requirements that need to be implemented by an Operator Platform to support each of these charging factors are described in section 4.3.3 of this document.

E.5 Charging factors summary

The table below summarises the list of potential events/charging factors that could be used by the Operator to carry out the charging and billing, depending on the service category exposed. The factors marked with "YES" are the ones potentially applicable for the service category.

Potential Events/ Triggers for Charging	Service Categories				- Technical Complexity +
	Network capabilities exposure: no impact on data traffic	Network capabilities exposure: with impact on data traffic	Network services provisioning	Edge capabilities management	
Service activation	YES	YES	YES	YES	
Service API invocation (and related notifications)	YES (API + payload)	YES (API+payload)	YES (API+payload) Service lifecycle management	YES (API+payload) Reserved Infra resources App lifecycle management	
Data traffic usage in the Operator's Network	NO	YES Only if volume-based charging (info provided by the Network) (*)	YES Only if volume-based charging (info provided by the Network) (*)	YES Only if volume-based charging (info provided by the Network) (*)	
Edge enabling infrastructure resources usage	NO	NO	NO	YES Only in case charging based on effective use infra resources (**)	

Table 4: Charging factors summary

(*): although information for charging is provided by the Operator's Network it will have implications for the Operator Platform that will have to provide the Network with information – via the SBI-NR – that can be used by the Operator's Charging engine to correlate API invocations with Data Traffic usage.

(**): Event-based charging model to be used for this purpose.

As already mentioned in this Annex, the Operator will be responsible for selecting the charging factors to use for a particular service depending on the selected commercial model for that service. Therefore, any decision relating to charging and/or billing for the usage of the services as described in this table is for an individual Operator to decide.

The table also shows the level of technical complexity that would be required for the implementation of the different charging triggers/factors, where the first rows have lower level of complexity than the rows at the end. Each charging trigger/factor is independent from one another.

Annex F Privacy Management considerations

F.1 General

From an OP perspective, data processing is limited to sharing data to an Application Provider (potentially through an Aggregator), so an Application (owned by the Application Provider) can perform any further processing on the shared personal data. To declare what an Application Provider wants to do with a set of personal information resources, a Purpose of Data Processing must be declared. Each Purpose of Data Processing is associated with a legal basis which must be compliant with local regulations. Only pre-defined Purposes of Data Processing can be used by the Application Provider and the use of personal data cannot go beyond that Purpose of Data Processing. Whenever the legal basis dictates direct interaction with End-Users (e.g., Consent legal basis), the Application signals the Purpose of Data Processing to the End-User through the OP (an Aggregator could be involved) and Privacy Management Function in the CSP domain. In turn, the End-User must opt-in and the Privacy Management Function must capture the result of that operation. It is expected that the data processing in the Application takes place exclusively under the indicated Purpose of Data Processing.

Several legal bases are well-established as indicated e.g., in [15]:

- Consent

  Note: Throughout this document, the terms “Consent” and “Application-related Consent” are interchangeable

- Context of a contract (to which the End-User is party)
- Compliance with a legal obligation (to which the controller is subject)
- Protect vital interest (of the data subject or of another natural person)
- Performance of a task carried out in the public interest
- Legitimate interests (pursued by the controller or by a third party)

Although most of the technical interest on legal bases revolves around the Consent for processing personal data (e.g., there is dedicated 3GPP study to deal with Consent for accessing 3GPP services [16]), other legal bases for processing personal data could be used. Nevertheless, it is noteworthy that some definitions related to other legal bases are rather relative to local regulations, e.g., what public interest or vital means could vary around the globe, so having universal mechanisms to fulfil local regulations is challenging.

If Consent is the applicable legal basis for processing, users must actively agree through an affirmative action (opt in). How Consent can be captured depends on the concrete use case and on the laws of the jurisdictions which govern the use case. Even though Consent can be obtained through a variety of methods and techniques (e.g., ticking a box on a website or writing/accepting a letter confirming the grant for processing personal data), having the Consent captured during runtime is also a well-established approach for some scenarios as elaborated in PRD OPG.10 [20].

There is not a single universal solution for Consent Management. It depends for instance in the controllership of the device (e.g., only one application should have control on the device), or the type of service provided by an application running on a generic device. For the former

case, a device identifier could be considered (along with other information) for granting access to personal data, whereas for the latter case, depending on the type of service, a server-side IP address could be considered (along with other information) for granting the application access to personal information.

Additionally, the Consent can be granted for one device or several at the same time. For the latter case, many mechanisms could be in place, e.g., providing a list of devices upfront while signing a contract, or via API calls through a portal.

F.2 Requirements for supporting relevant End-User rights

General requirements to OP can be derived from analysis of a subset of rights that an End-User could be entitled to. Direct interactions among End-Users and the Application Provider, as well as direct interactions among Subscribers and the CSP domain are considered out-of-the-scope of this document. Table 5 presents technical requirements for the OP, SBI-AAPrM, NBI and EWBI to support a subset of privacy-related rights in which OP plays a relevant role.

		Requirements on:		
		Privacy Management Function	**OP**	**Application Provider**
End-User privacy-related rights	Information (to understand what will be done with their data)	• Upon indication from the OP, capture the Consent from the End-User signaling the Purpose of Data Processing for an Application ID • Store the Privacy Information in the Privacy Management Function • Create a subscription to get the OP notified about any possible change on the Privacy Information	• Ensure that a suitable legal basis (compliant with local regulations) supports sharing personal data with an AP. • If the applicable legal basis is Consent: 1. Trigger the Consent Capture carrying the intended Purpose of Data Processing, AP ID and Application ID 2. Cache Privacy Information (if allowed by local regulations) 3. Get notified about changes on Privacy Information • If a Vetting Process took place already, the OP should present an indication to the End-User about the Vetting Process status while capturing the Consent (if applicable).	• -Signal the Purpose of Data Processing (associated to a legal basis) while onboarding and obtaining authorization process
	Access (to get a confirmation whether their personal data is being processed)		• Logging, Tracing and Auditing functions	• Logging, Tracing and Auditing functions

		Requirements on:		
		Privacy Management Function	**OP**	**Application Provider**
	• Restricting of processing (to request ceasing of all processing of their data) • Object	• Upon notification, update relevant Privacy Management Function entries • Log, trace and audit any further attempt to capture Consent • Keep the Application Privacy Profilestored during the time indicated by the local regulation	• Upon notification, stop sharing associated personal information and (potentially) notify other entities: 1. If applicable legal basis is Consent, notify Privacy Management Function 2. Notify a Partner OP • Log, trace and audit any further invocation involving personal data	• Upon End-User request, stop any associated processing of personal data and notify OP or Leading OP

Table 5: Requirements for supporting End-User privacy-related rights

F.3 Considerations from the architecture perspective

The Privacy Management Function in the CSP domain holds both:

- Privacy Information which is filled in when a person entitled to consent access to protected data opts-in / opts-out, and
- Application Privacy Profile which is populated during API subscription time.

Figure 8 presents a subset of functional components of the OP and illustrates how the Application Privacy Profiles may be populated to enable subsequent runtime operations (e.g., getting an access token).

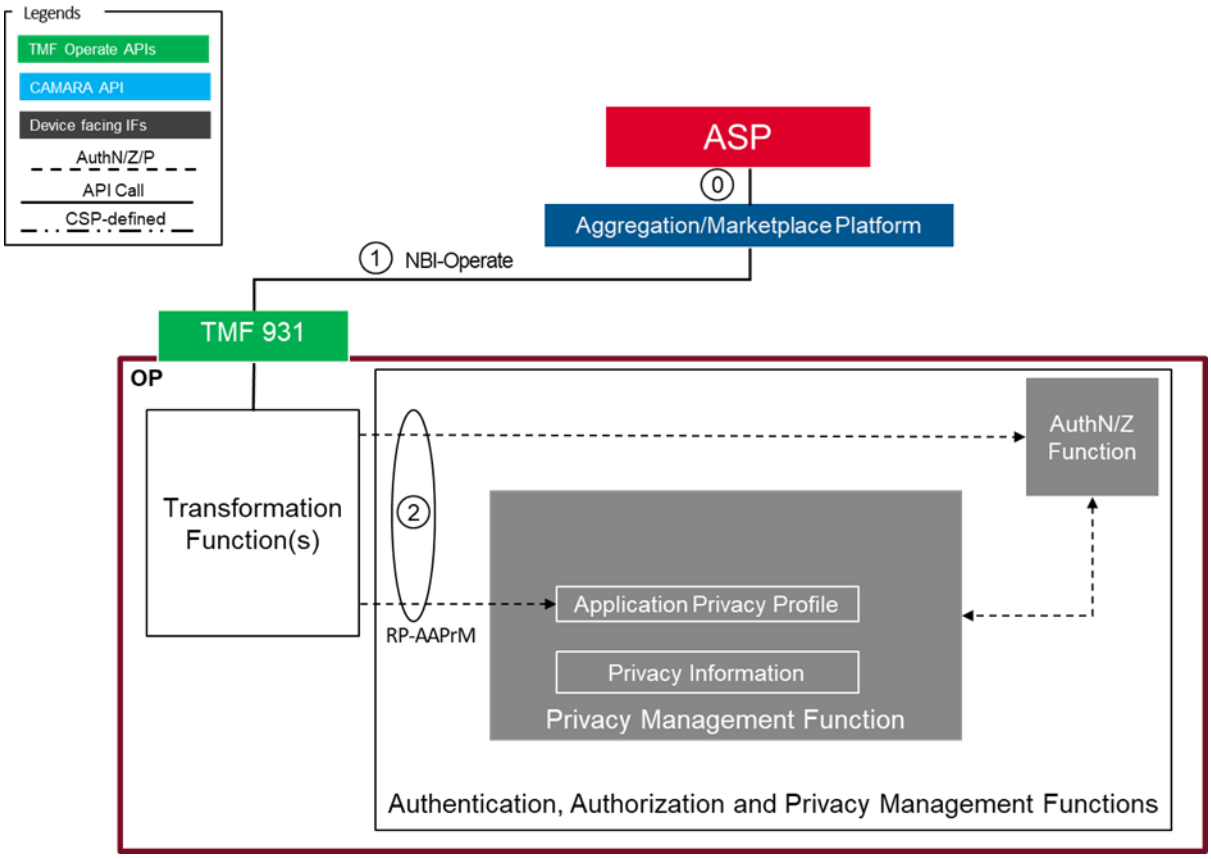

Figure 8: Filling in Application Privacy Profile

Note: Having both the Privacy Management Function or Authentication/Authorization Function (grey boxes in Figure 8) outside of OP is also a valid deployment option that implies interactions over the SBI-AAPrM (rather than the internal reference point RP-AAPrM).

- (step 0) in Figure 8, during the Application Onboarding and API subscription process [22], the Application Provider (referred to as ASP in the figure) is presented a list of products (aggregated APIs), including information about supported scopes, Purpose of Data Processing, and applicable grant type(s) and legal bases (which depends on local regulations)
- (step 1) once the Application Provider has agreed the terms and conditions for the API subscription, and as part of the API ordering to the CSP,
- (step 2) the Transformation Functions (see section 3.3.5.2) may perform two actions:
    - Capture the relevant Application Privacy Profile (as in section G.4.2.1), and
    - Provide the relevant information to the Authentication/Authorization Function for handling authorization requests.

Before an Application can consume a Service API, it needs to obtain a valid access token via authorization/authentication request. Considering Figure 9,

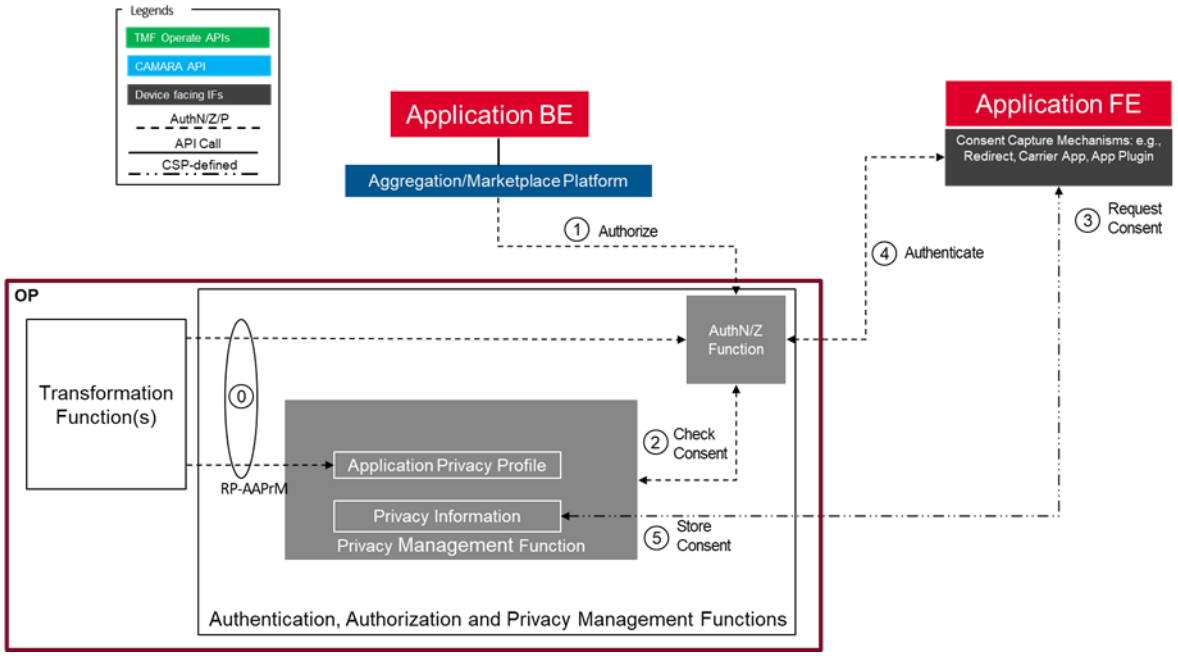

Figure 9: (High-level) Consent capture trigger

Note: Having both the Privacy Management Function or Authentication/Authorization Function (grey boxes in Figure 9) outside of OP is also a valid deployment option that implies interactions over the SBI-AAPrM as already considered in section 3.3.2.3.

- (step 0) the process depicted in Figure 8 was carried out
- (step 1) the Authentication/Authorization Function receives an authorization/authentication request carrying information about the scope and Purpose of Data Processing,
- (step 2) the Authentication/Authorization Function checks the Application Privacy Profile to determine if Consent needs to be captured based on the local regulatory considerations and if so, whether the Consent is already in place querying the Privacy Information. If Consent is in place, the Authentication/Authorization Function may provide the access token needed by the Application.
- (step 3) if Consent needs to be captured, the person entitled to consent access to protected data will be notified (e.g., via redirects or out-of-band mechanisms) about the need for an explicit opt-in.

Note: the actual mechanisms for interacting with the End-User for notifying and capturing the Consent are left to the CSP.

- (step 4) an authentication step may be needed depending on the use case.
    - depending on the scenario, an interaction with a User Identity Token Manager may be needed.
- (step 5) the results of that operation will be stored on the Privacy Information.

Unlike the interactions depicted in Figure 9 in which the Consent capture is triggered while processing an authorization/authentication request, the Consent may be obtained upfront

using the Application execution context to improve the user experience. For that, CAMARA has defined a Service API called Consent Info API [23].

Note: To use the Consent Info API, the Application must obtain an access token. It is assumed that obtaining an access token for Consent Info API does not imply Consent capture itself.

The high-level process for capturing Consent using the CAMARA Consent Info API is as shown in Figure 10:

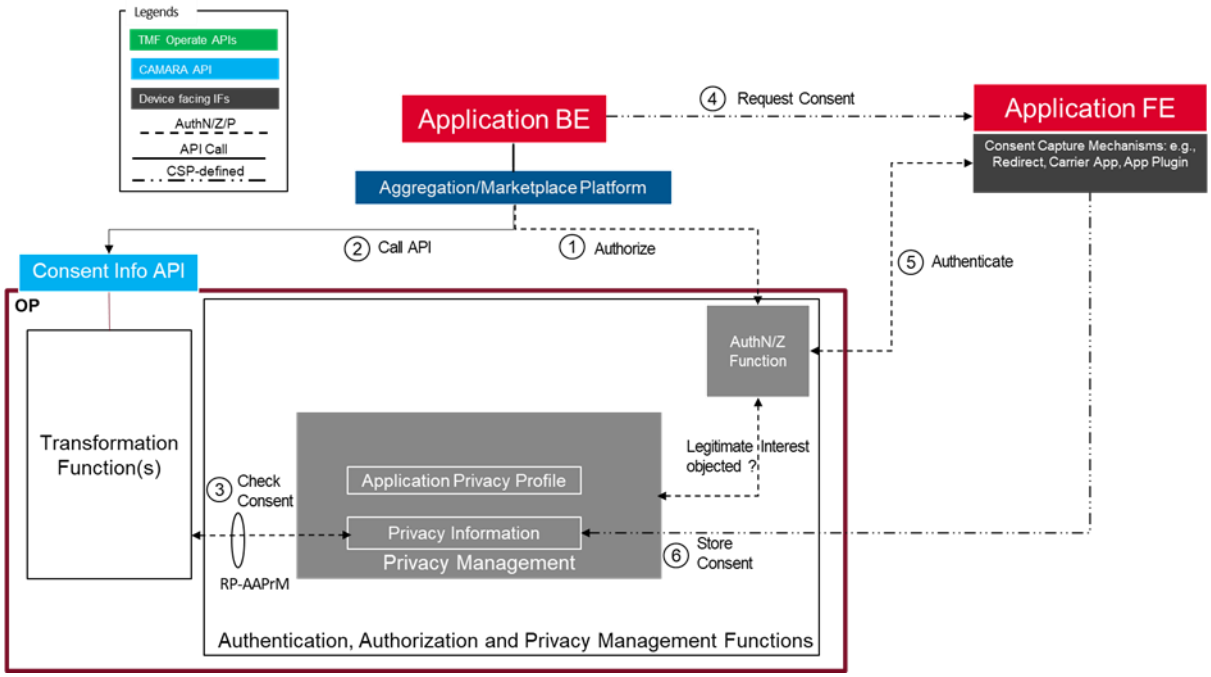

Figure 10: (High-level) Upfront Consent capture

Note: Having both the Privacy Management Function or Authentication/Authorization Function (grey boxes in Figure 10) outside of OP is also a valid deployment option that implies interactions over the SBI-AAPrM as already considered in section 3.3.2.3.

- (step 1) a decision for checking the validity of data processing (which is held in the CSP domain within the Privacy Information) for a specific user, scope and purpose is made within the Application. The Application Backend will get an access token (following the high-level process in Figure 9) providing a Purpose of Data Processing that implies a legal basis different from Consent (e.g., Legitimate Interest, Compliance with a legal obligations). The Authentication/Authorization Function may check whether the End-User has objected Legitimate Interest previously for the Application. If that is not the case (and depending on local regulatory conditions), the Authentication/Authorization Function issues an access token.
- (step 2) the Application Backend sends a request to the Aggregation/Marketplace Platform for checking the validity of data processing is in place within the CSP domain. The request may indicate that a Consent capture URL needs to be provided in the response. The Aggregation/Marketplace Platform discovers the CSP that

provides the connection to the user identifier received from the Application. The Aggregation/Marketplace Platform forwards the request to the corresponding OP.

- (step 3) upon receiving the request, in step 2 the OP uses the Transformation Functions (see section 3.3.5.1) to query via the Privacy Management Function to check if there is Privacy Information associated with the provided filtering criteria (user identifier, scope and purpose) in place. The result of this check is sent back to the Application Backend conditionally including the Consent capture URL. Based on the Consent check result (e.g., Consent expired or revoked),
- (step 4) the Application will present a Consent capture dialog within the context of the Application execution (e.g., presenting an in-app webview or redirecting the End-User to the Consent capture portal).

Note: the actual mechanisms for interacting with the End-User for notifying and capturing the Consent are left to the CSP.

- (step 5) an authentication step may be needed depending on the use case.
- (step 6) based on the presented dialog, the person entitled to consent access to protected data will decide and select the Consent preferences which will be persisted in the Privacy Management Function within the CSP domain.

More details about the procedure in Figure 10 can be found in GSMA PRD OPG.10 [20].

Once an access token is retrieved, it is possible to consume a Service API. Figure 11 presents a high-level flow for consuming a Service API.

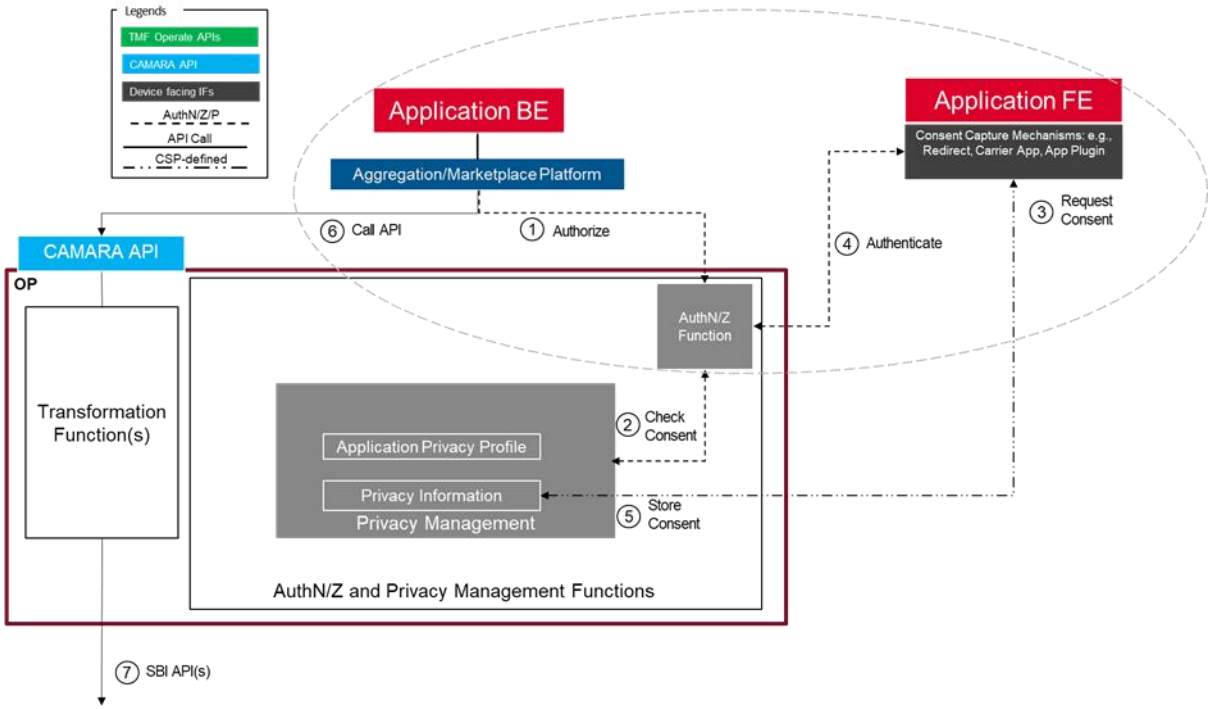

Figure 11: (High-level) Service API invocation

Note: Having both the Privacy Management Function or Authentication/Authorization Function (grey boxes in Figure 11) outside of

OP is also a valid deployment option that implies interactions over the SBI-AAPrM as already considered in section 3.3.2.3.

- (steps 1-5) The Application Backend obtains a valid access token,
- (step 6) The Application Backend sends a Service API request to the Aggregation/Marketplace Platform. The Aggregation/Marketplace Platform discovers the CSP the request should be forward to. Once the corresponding CSP is discovered, the Aggregation/Marketplace Platform forwards the request to the corresponding OP.
- (step 7) upon receiving the request, the OP uses the Transformation Functions (see section 3.3.5.1) for mapping the request on the NBI to the corresponding SBI services.

Annex G Service Flows

This section describes how an Operator Platform could interact with network elements, UEs and other parties to realise various service use cases that it enables and supports.

G.1 Service delivery by the OP (without UNI)

G.1.1 Service delivery to UE attached to the Home Network

In most cases there will be no UE registration directly with the OP. The UE will register with its Application Backend and that Application Backend will be authenticated and authorised by the OP to use the services.

Note: GSMA PRD OPG.11 [21] covers further flows where the UE registers with the OP over an Edge service-specific UNI.

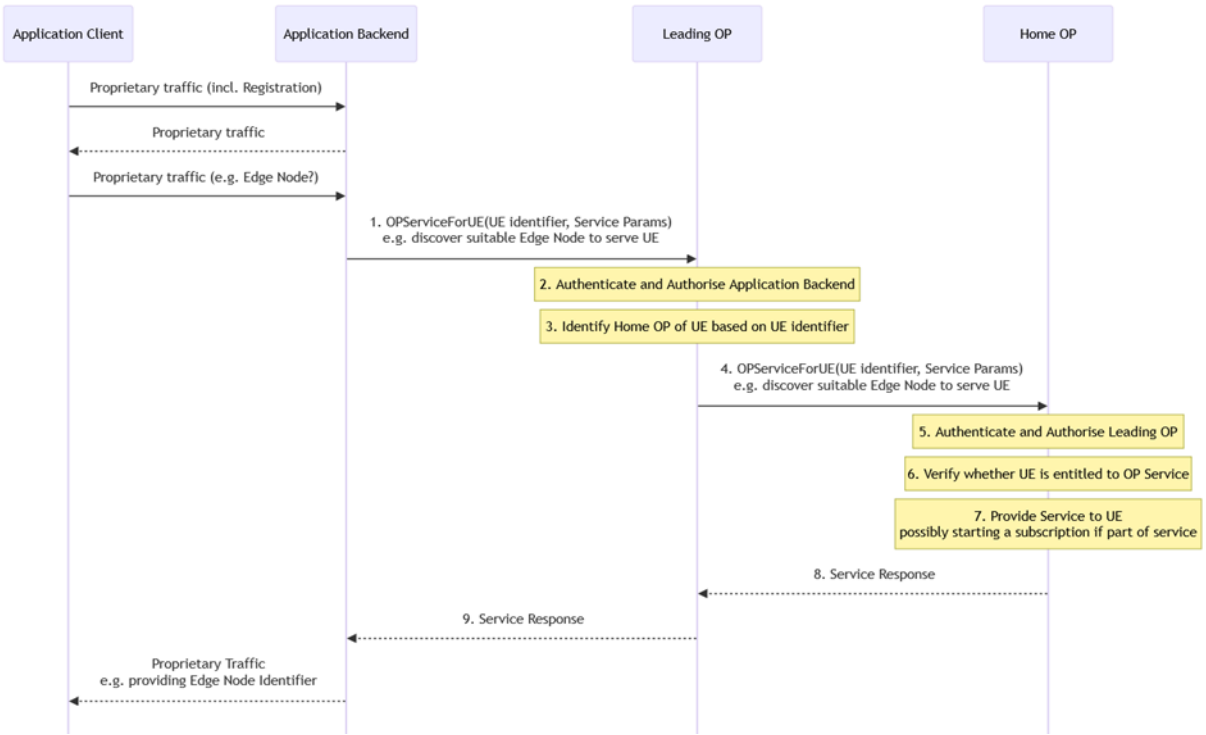

Figure 12: UE Registration to the OP without UNI

In this case, after some proprietary application-specific exchanges, the Application Backend may contact the Leading OP for the Application to obtain an OP-provided service related to the UE that would support the service that the Application is providing. The steps would then be as follows:

1. The Application Backend contacts the Leading OP, identifying itself and the UE to which the service should be delivered (e.g. providing its IP address or MSISDN).
2. The Leading OP authenticates and authorises the Application Backend
3. The Leading OP determines the target network to serve the request. In this non-Roaming case that would be the Home OP of the UE.

Note: The Leading OP is not necessarily the same as the Home or Visited OP for the UE. Therefore, the flow separates these.

4. The Leading OP provides the Application Backend's request to the UE's Home OP over the E/WBI between those OPs identifying itself and the target UE.
5. The Home OP authenticates the Leading OP.
6. The Home OP verifies whether the subscription of the UE that is referred to is entitled to use the requested service and whether the Application is authorised to request it for that Subscriber.

Editor's Note: Privacy Management is for further study.

7. The Home OP delivers the requested service. Depending on the request, this may involve starting an event subscription to inform the Application of any events related to the service delivery to the UE.
8. The Home OP provides a response on the service delivery to the Leading OP (e.g. indicating a suitable Edge Node.that was discovered)
9. The Leading OP provides a response on the service delivery to the Application (e.g. indicating a suitable Edge Node.that was discovered).

The Application Backend can then use the information in the response received from the Leading OP to enhance its service based on the agreements that the Application Provider has with the Leading OP for using the OP services (e.g. indicating to the UE in an application-specific exchange how to reach the most suitable Edge Node).

G.1.2 Service delivery to UE attached to a Visited Network

This procedure describes the service delivery to a UE (UE-A) with an Operator Platform (Operator A) without UNI while attached to a visited network (Operator B). In this case the flow for requesting the service will be similar to the scenario described in section G.1.1.

Depending on the service requested and capabilities available, the following roaming approaches are considered:

- **Home routing (HR)** – for services that depend entirely on the home network (e.g. subscription-related requests) and scenarios where service delivery provided by the visited network cannot be supported (e.g. Edge Resources in the visited network cannot be accessed).

The Home OP is the only OP involved in this case, with the service request handled as defined in section G.1.1. Figure 13 below shows the relations between the networks for the access to Edge Resources in this case. This scenario comes with limitations on application availability due to increased latency (see GSMA PRD OPG.11 [21]).

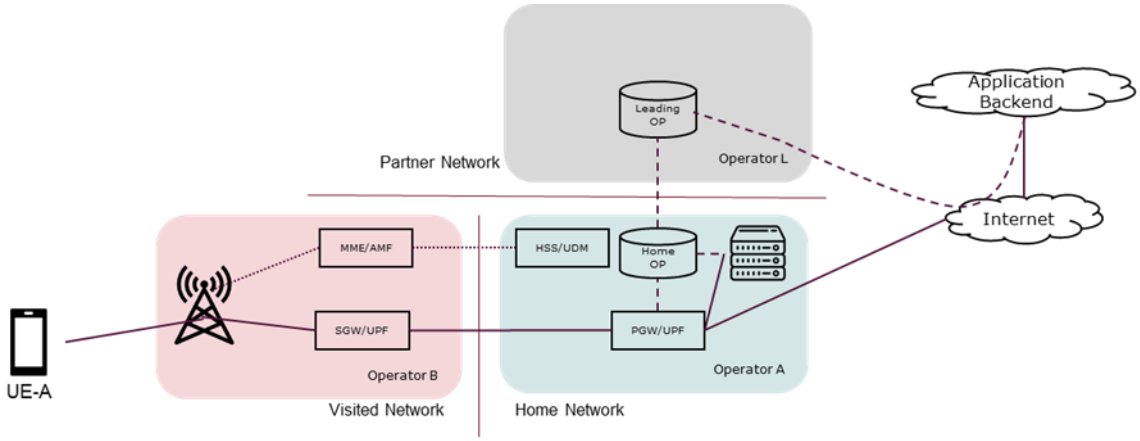

Figure 13: Roaming access to OP and Edge Resource - HR

- **Local Breakout (LBO)** – for services that depend (partly) on service delivery by the visited network. After receiving the service request as defined in section G.1.1, the Home OP will identify the visited network on which their Subscriber is roaming and determine what Partner OP to contact to support the service delivery. The Home OP will then request the service(s) to be delivered by the Partner OP through its EWBI with that Partner OP identifying its Subscriber as a roaming user.

Note: What service(s) should be delivered by the Partner OP depends on the service that the Home OP was requested to provide and would thus be service-specific.

Note: How to identify the Subscriber to the Partner OP is for further study.

In some cases, the service to be delivered might involve the UE accessing resources (e.g. Edge Nodes) available in the visited network directly. This model is preferred because the service depending on these resources is provided closer to the UE then. It is expected that the UE will have one data session routed to the home network (Operator A). In this case, the Home OP will receive the identifiers for UE access to those resources over the E/WBI as part of a discovery (e.g. Edge Discovery) or other service invoked on the Visited OP through the E/WBI and pass them on to the Leading OP and the Application Backend over the E/WBI and/or the NBI. Figure 14 below shows the relations between the networks in this case with the following clarifications:

- The solid lines: UE data session (PDU)/PDN) that carries the OP and Application Identifiers. It is also used for any data traffic and UE access to the resources on Operator B (through the UPF/PGW in Operator B network).

- The dotted lines: Signalling traffic between the visited network (Operator B) and home network (Operator A) as defined by 3GPP TS 23.501 [4]. It is used for the UE registration into the network
- The dashed line: NBI, SBI and EWBI connection for enabling the access to Operator B's resources.

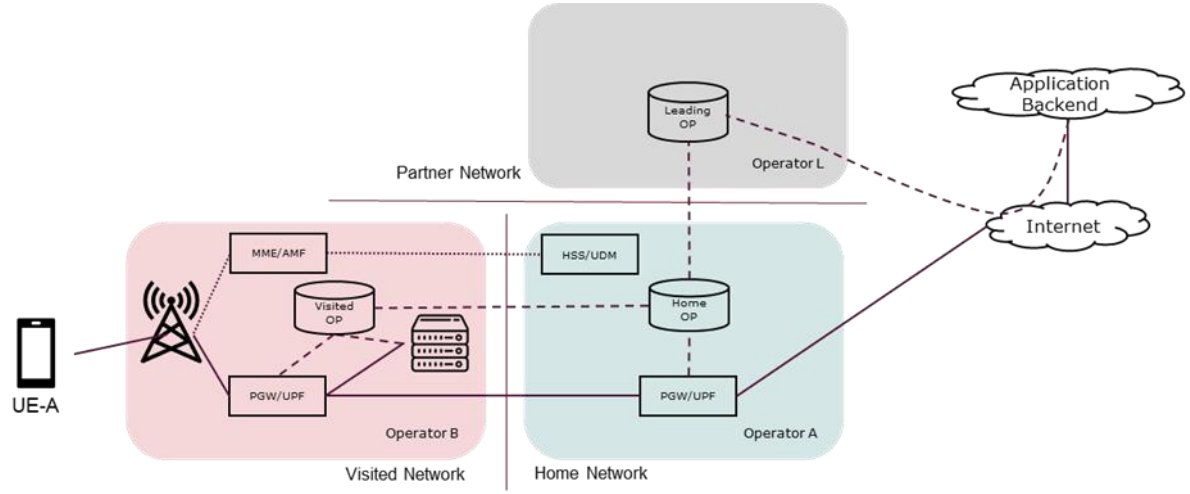

Figure 14: Roaming access to OP and Edge Resource – HR+LBO

G.2 Charging Concepts

The following flows describe how charging factors can be used for different services in non-federated scenarios. Please see section 4.3.3 for the requirements related to the SBI-CHF and Annex E.5 for the charging factors and the service categories underpinning the scenarios.

G.2.1 Charging for Service API Invocation

This flow describes concepts of charging for different type of Service API invocations based on API type and payload. Depending on the API type, a subset of the parameters included in the payload will be of interest to enable rating and charging. This model also supports rating and charging based on the API plus the operation used without payload analysis.

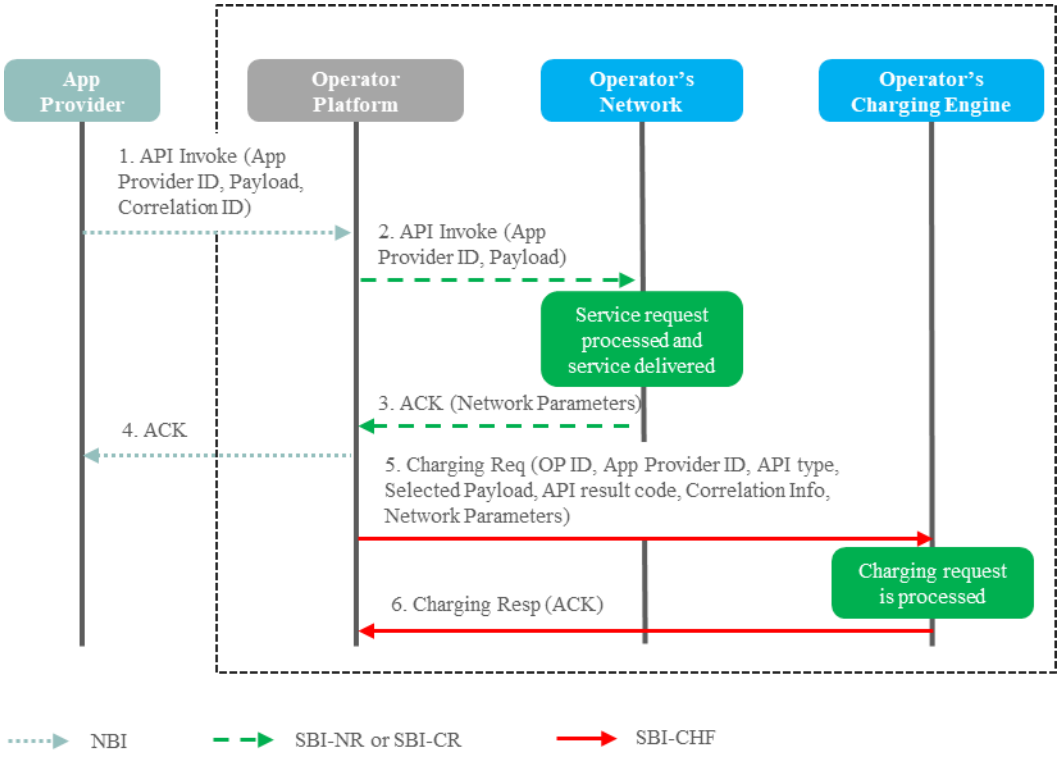

Figure 15: Charging for Service API Invocation

1. The Application Provider invokes an API. Information may include:
    - Party identifiers: App ID, Application Provider ID, Customer device ID
    - API Payload
    - Correlation ID
2. The Operator Platform invokes the corresponding API(s) using SBI-NR.
3. The service request is processed in the Operator's Network and the service is delivered. An ACK is sent back to the Operator Platform. The network response can include relevant parameters for rating and charging.
4. The Operator Platform sends back the response (ACK) to the Application Provider using the NBI
5. The Operator Platform sends a charging request to the Operator's Charging Engine using the SBI-CHF. A charging request includes at least:
    - Party Identifiers: OP ID, App ID, App Provider ID, Customer device ID
    - API type, selected API payload (not mandatory to include and dependent on the service) and API result code
    - Correlation Information
    - Selected Network parameters coming from SBI NR response (not mandatory to include and dependent on the service)
6. The Operator's Charging Engine processes the charging request (rating and charging is done based on provided information in the charging request) and sends a response back to the Operator Platform using the SBI-CHF.

G.2.2 Charging for Data Traffic Usage in Operator Network

This flow describes how the charging concept for data traffic usage will be performed as a result of Service API invocation.

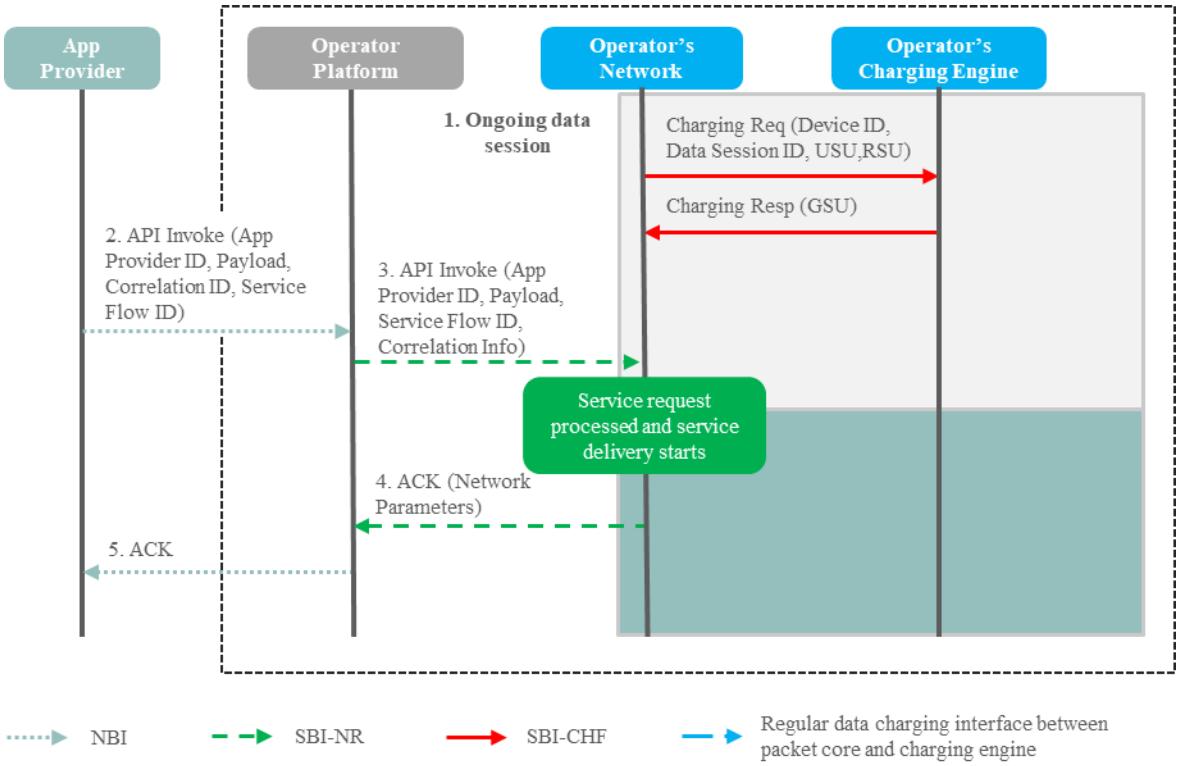

Figure 16: Charging for Data Traffic Usage in Operator Network – Part 1 of 2

1. There is an ongoing data session for a customer. A regular dialogue (session-based charging) is performed between the Operator's data packet core and the Operator's Charging Engine. Online or offline charging could be used depending on the Operator's decision (the online mode shown here as an example) and is out of the scope of this document.
2. An Application Provider invokes an API that influences data traffic usage of a device. Information may include:
    - Party identifiers: App ID, Application Provider ID, Customer device ID.
    - API Payload
    - Correlation ID
    - Data Service flow ID
3. The Operator Platform invokes the corresponding API(s) using the SBI-NR. The Operator Platform provides the required parameters to enable correlation, Correlation Information, between API invocation and the data session in the Operator's Charging Engine.
4. The service request is processed in the Operator's Network and the service is starting to be delivered. An ACK is sent back to the Operator Platform optionally including some additional information which may be relevant for charging.
5. The Operator Platform sends back the response (ACK) to the Application Provider using the NBI

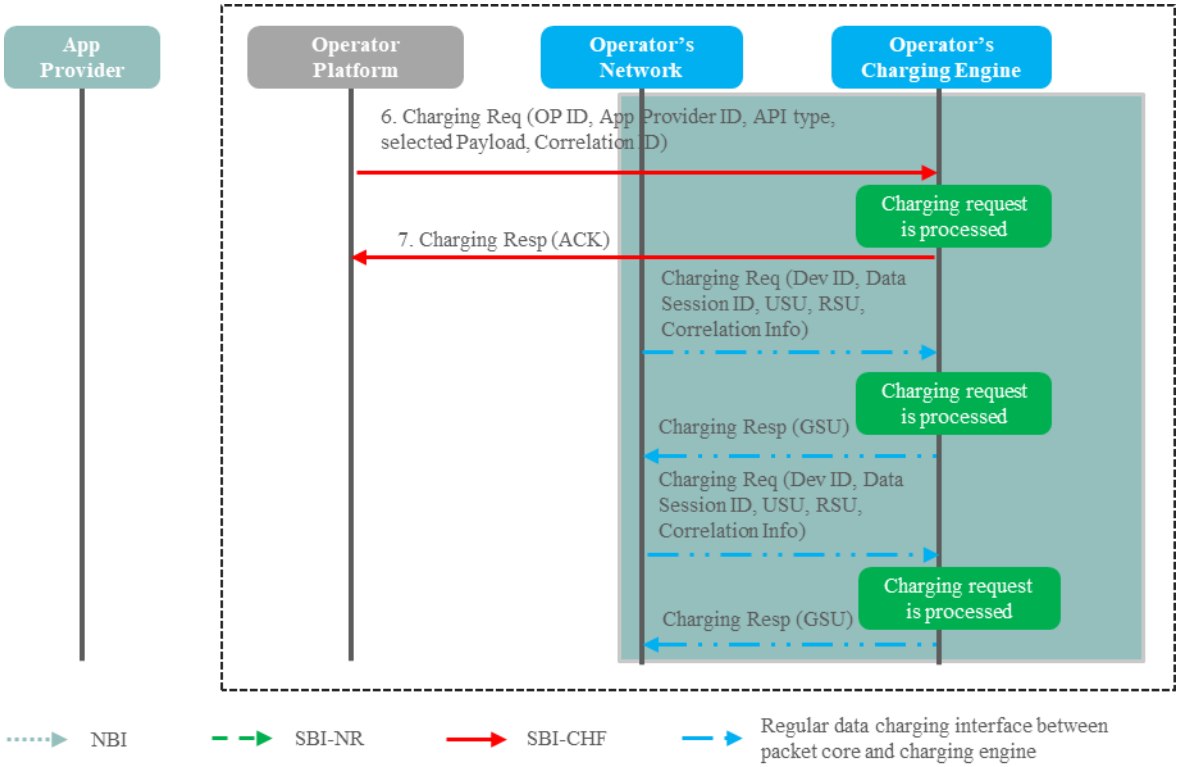

Figure 17: Charging for Data Traffic Usage in Operator Network – Part 2 of 2

6. The Operator Platform sends a charging request to the Operator's Charging Engine using the SBI-CHF to ask for API invocation charging. A charging request may include:
    - Party Identifiers: Leading OP ID, App ID, App Provider ID, Customer device ID
    - API type + selected API payload (optional)
    - Correlation Information
    - DataSessionID
    - Time/duration (optional)
    - Network parameters (optional)
7. The Operator's Charging Engine processes the charging request (rating and charging is done based on provided information in the charging request) and sends a response back to the Operator Platform.

During the service delivery (time/volume where the data traffic is impacted by the API call) the charging dialogue between the Operator's data packet core and the Operator's Charging Engine will continue.

Note: Besides the regular information exchanged for the ongoing charging of the data session, The Packet core will include the Correlation Information required by the Charging Engine to be able to identify the data traffic volume impacted by the Service API Invocation. The contents of this correlation information are for further study.

G.2.3 Charging for Edge Enabling Infrastructure Resource Usage

This flow describes charging for Edge enabling infrastructure resource usage.

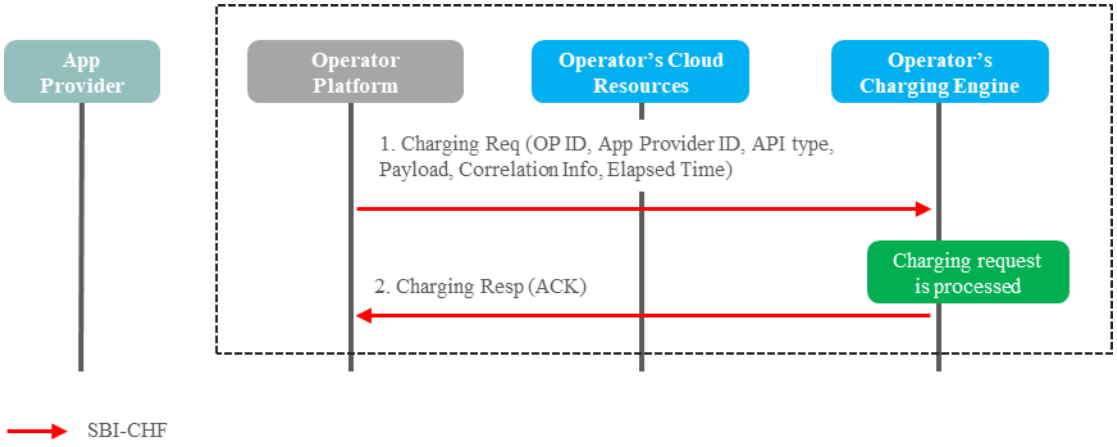

Figure 18: Charging for Edge Enabling Infrastructure Resource Usage

1. The Operator Platform monitors the usage of Edge infrastructure resources and sends a charging request to the Operator's Charging Engine using the SBI-CHF. This can be done periodically based on a configurable timer or one request for the whole period. A charging request may include:
    - Party Identifiers: OP ID, App ID, App Provider ID, Pool ID
    - API type + selected API payload (vCPUs, memory, storage, incoming/outgoing data volume, time period)
    - Correlation Information
2. The Operator's Charging Engine processes the charging request (rating and charging is done based on provided information in the charging request) and sends a response back to the Operator Platform.

G.3 Charging Concepts in Federated Scenarios

The following flows describe charging concepts for Edge computing services in federated scenarios.

Note: Further analysis of other federation services will be included in future versions of this document. This includes any scenarios with intermediate parties.

The E/WBI will be used in communication between the Operators and a pre-requisite is that there is an existing federation agreement in place. In federation scenarios the focus of charging is to ensure that both Leading and Partner Operators have records of the service use as part of their record keeping. This will be the input to settlement and reconciliation between the two Operators later as well as for any wholesale charging between the Leading Operator and the Application Provider.

Please see section 4.3.3 for the requirements related to the SBI-CHF and Annex E.5 for the charging factors and the service categories underpinning the scenarios.

G.3.1 Federated Service API Invocation

This flow describes charging for different type of Service API invocations based on payload in a federated scenario for Edge Computing.

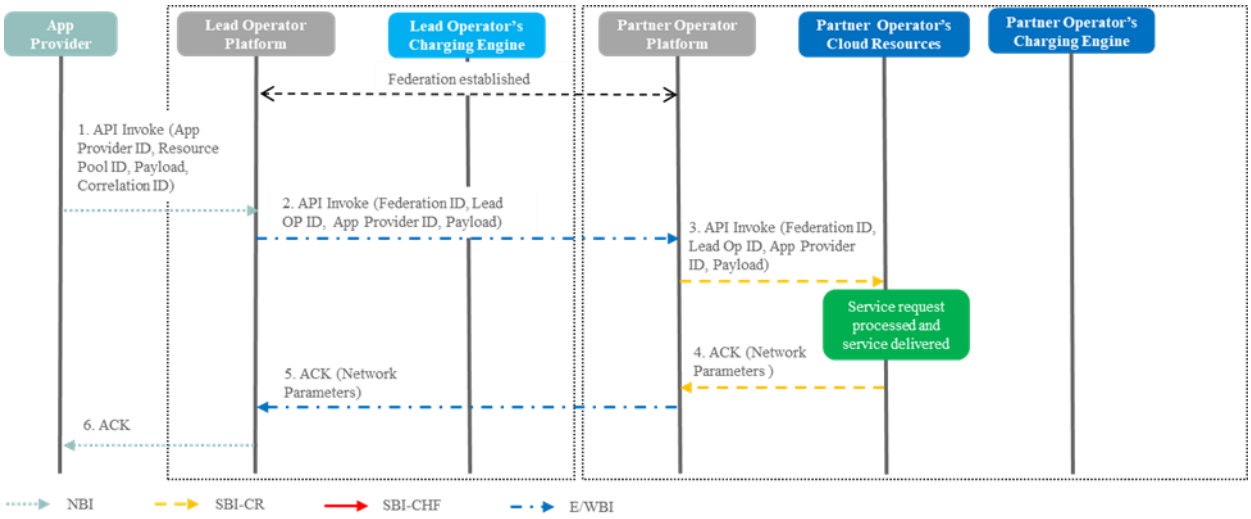

Figure 19: Federated Service API Invocation – Part 1 of 2

1. The Application Provider invokes a Service API
2. The Leading Operator Platform forwards the API request using the E/WBI towards the Partner Operator Platform
3. The Partner Operator Platform invokes the corresponding API(s) using the SBI-CR
4. The service request is processed in the Partner Operator's Cloud Resources and the service is delivered. An ACK is sent back to the Partner Operator Platform
5. The Partner Operator Platform sends an ACK back to the Leading Operator Platform
6. The Leading Operator Platform sends back a response (ACK) to the Application Provider using the NBI

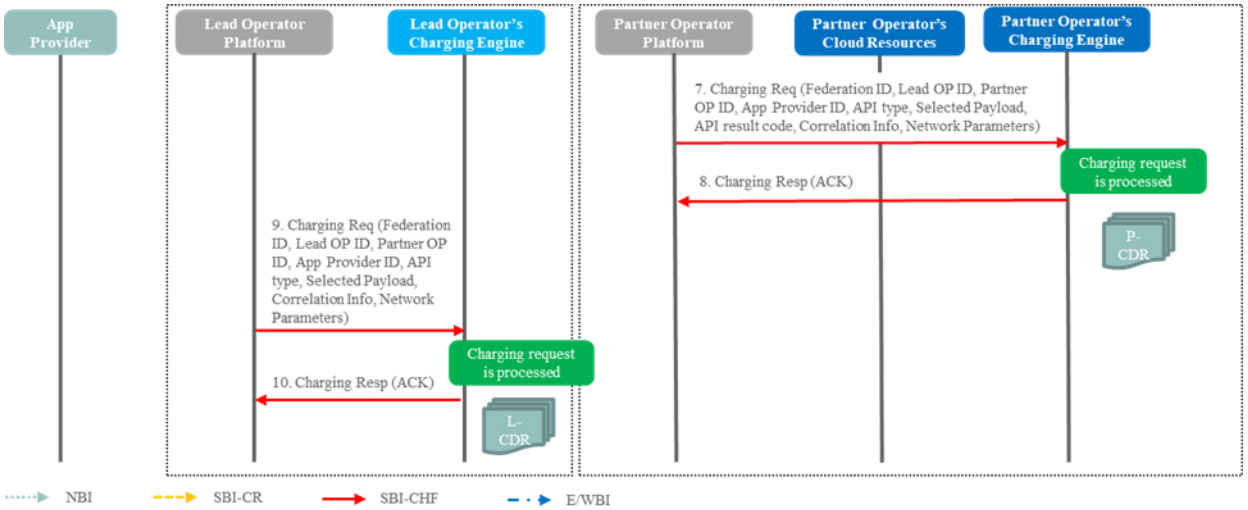

Figure 20: Federated Service API Invocation – Part 2 of 2

7. The Partner Operator Platform sends a charging request to the Partner Operator's Charging Engine using its SBI-CHF. This can be done periodically based on a

configurable timer or be one request for the whole period. A charging request may include:

- Party Identifiers: Federation ID, Leading OP ID, Partner OP ID, App ID, App Provider ID, Pool ID
- API type + selected API payload (not mandatory to include and dependent on the service) + API result code
- Correlation Information

8. The Partner Operator's Charging Engine processes the charging request (rating and charging is done based on provided information in the charging request) and sends a response back to the Operator Platform. The result of charging is stored by the Partner OP as events used for settlement and reconciliation. How and where the events are stored is up to the Operator to decide.
9. The Leading Operator Platform sends a charging request to the Leading Operator's Charging Engine using the SBI-CHF. This is based on the results received from the Partner OP. A charging request includes:

    - Party Identifiers: Federation ID, Leading OP ID, Partner OP ID, App ID, App Provider ID, Pool ID
    - API type + selected API payload (not mandatory to include and dependent on the service) + API result code
    - Correlation Information

10. The Leading Operator's Charging Engine processes charging request (rating and charging is done based on provided information in the charging request) and sends a response back to the Operator Platform. The result of the charging is stored by the Leading OP as events used for settlement and reconciliation. How and where the events are stored is up to the Operator to decide.

Note: The CDRs generated by the Leading and Partner OPs Charging Engines are input to settlement and reconciliation processes outside of charging and hence not in scope.

G.3.2 Federated Edge Enabling Infrastructure Resource Usage

For federated scenarios, it will be possible to periodically exchange information around the effective Edge Resource usage over the E/WBI. Consideration needs to be taken to ensure that the resource consumption used for charging on the Partner and Leading Operators are synchronised to reduce risk of reconciliation issues.

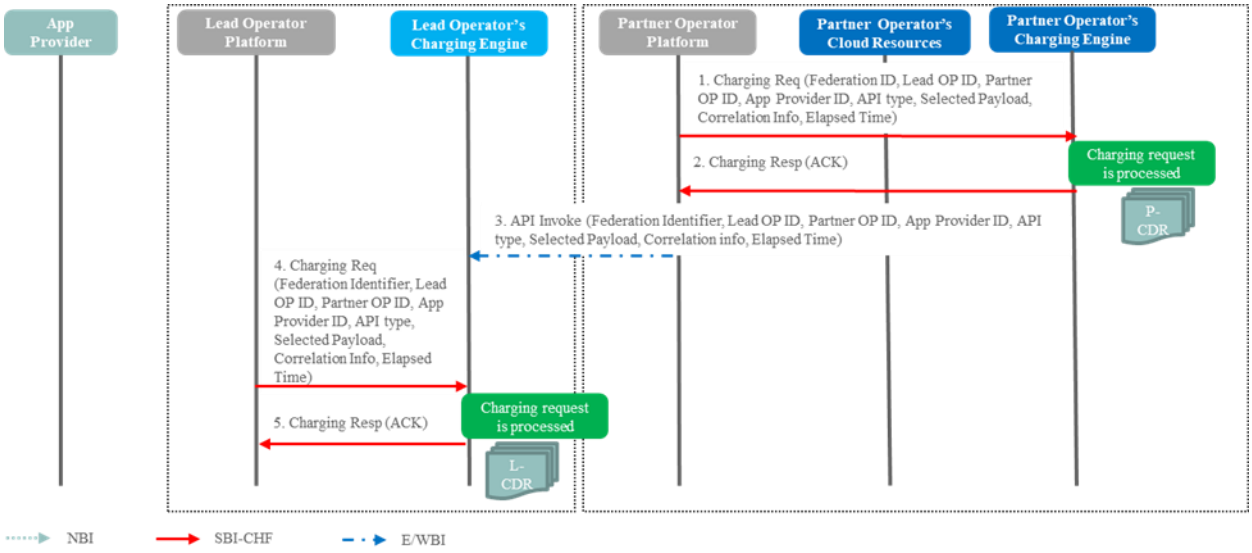

Figure 21: Federated Edge Enabling Infrastructure Resource Usage

1. The Partner Operator Platform monitors the usage of Edge infrastructure resources and sends a charging request to the Partner Operator's Charging Engine using the SBI-CHF. This can be done periodically based on a configurable timer or be one request for the whole period. A charging request may include:
    - Party Identifiers: Federation ID, Leading OP ID, Partner OP ID, App ID, App Provider ID, Pool ID
    - API type + selected API payload (vCPUs, memory, storage, incoming/outgoing data volume, time period)
    - Correlation Information
2. The Partner Operator's Charging Engine processes the charging request (rating and charging is done based on provided information in the charging request) and sends a response back to the Operator Platform. The result of the charging is stored by the Partner OP as events used for settlement and reconciliation. How and where the events are stored is up to the Operator to decide.
3. The Partner Operator Platform invokes the API to exchange the actual resource usage using the E/WBI towards the Leading Operator Platform
4. The Leading Operator Platform sends a charging request to the Leading Operator's Charging Engine using its SBI-CHF. This is based on the actual resource usage received from the Partner OP. A charging request includes:
    - Party Identifiers: Federation ID, Leading OP ID, Partner OP ID, App ID, App Provider ID, Pool ID
    - API type + selected API payload (vCPUs, memory, storage, incoming/outgoing data volume, time period)
    - Correlation Information
5. The Leading Operator's Charging Engine processes the charging request (rating and charging is done based on provided information in the charging request) and sends a response back to the Operator Platform. The result of the charging is stored by the

Leading OP as events used for settlement and reconciliation. How and where the events are stored is up to the Operator to decide.

The CDRs generated by the Leading and Partner Ops Charging Engine are input to settlement and reconciliation processes outside of charging and hence not in scope.

G.4 Privacy Management

Depending on the legal basis associated with the Purpose of Data Processing, it might be needed to interact with a person entitled to consent access to protected data to obtain authorization for sharing personal data with an Application (e.g., for Consent legal basis – see Annex F). For an OP to trigger the capture of Consent by the Privacy Management Function in the CSP domain, several procedures could be in place depending on the use case and scenario.

G.4.1 Consent capture: use cases and flows

Detailed use cases and flows for Consent capture are presented in GSMA PRD OPG.10 [20].

G.4.2 Relevant information

G.4.2.1 Application Privacy Profile

Information used for controlling the applicable legal basis (if any) for all APIs used by an Application. The Application Privacy Profile is highly dependent on the local regulations and it is recommended to support relevant ApiProducts' sub-resource fields [22], as reproduced in Table 6. Due to the local validity of the information, it is recommended that this information is held within the CSP domain.

Data type	Description
Country Code	Country Code
Application Provider ID	Unique identifier of the Application Provider
Application ID	Unique identifier for the Application requesting access to personal information
API Name	Standardized name of the API
API version	API Version
Base Path	Base path of the API
Scope	Represent a resource the Client requests access to (e.g., “sim-swap:retrieve-date”)
Purpose	Purpose of data processing (e.g., “Fraud Prevention and Detection”, “Direct Marketing”)
Grant Type	Grant type for obtaining authorization (e.g., “authorization_code” [17], “client_credentials” [17], “urn:openid:params:grant-type:ciba” [18])
Legal Basis	Applicable legal basis (e.g., “Consent”, “Legitimate Interest”) [15]
Additional Privacy Considerations	(Based on local regulations) placeholder to keep potential provisions for specific data processing scenarios e.g., transfer of personal data to third countries [15].

Table 6: Application Privacy Profile

Note: Table 6 suggests information placeholders rather than implementation.

G.4.2.2 Privacy Information

Information held within the CSP domain used for keeping evidence/records of the lawfulness of privacy-sensitive data processing and sharing. This information is populated via explicit End-User opt-in / opt-out actions.

Data type	Description
Consent ID	Identifier (on the Privacy Management Function) of the Consent entry
Authorizing Party ID	Identity of the party granting the Consent for processing personal data
Matching Criteria	Individual or list of Device ID(s), or PDU filter(s), or Subscription ID(s) for which the personal information processing is allowed
Application Provider ID	Unique identifier of the Application Provider
Application ID	Unique identifier for the Application requesting access to personal information
API Name	Standardized name of the API
API version	API Version
Scope(s)	Reference to a set of resources being protected defined in an API specification
Purpose of Data Processing	Predefined/standardized Purpose of Data Processing
Capture Method	Mechanism by which consent was obtained (Batch, Frontend based, SMS, API calls, e-mail, etc).
Status	Granted, Denied, Revoked, Pending, Expired, Objected (applicable for “legitimate interest” legal basis)
Consent Grant Timestamp	Timestamp at which the Consent was granted
Consent validity	Timestamp until which the consent record is valid
Revocation Method	Mechanism by which revocation was requested (Batch, Frontend based, SMS, API calls, e-mail, etc.)
Revocation Timestamp	Timestamp at which the Consent was revoked
Retention Period	Duration of time for which the personal data needs to be retained following receipt of revocation request

Table 7: Privacy Information

Note: On federation scenarios it might not be needed or allowed to share information about the party who granted the Consent.

Note: Table 7 suggests information placeholders rather than implementation.

G.5 Vetting Process

G.5.1 High-level flow for Vetting process

Figure 22 presents a high-level flow of the Vetting Process triggered by a CSP after initial Application Provider and Application onboarding. The vetting service(s) in Figure 22 enable the Vetting Process and may be offered by a CSP itself, an Aggregator or a third party (vetting service(s)) provider.

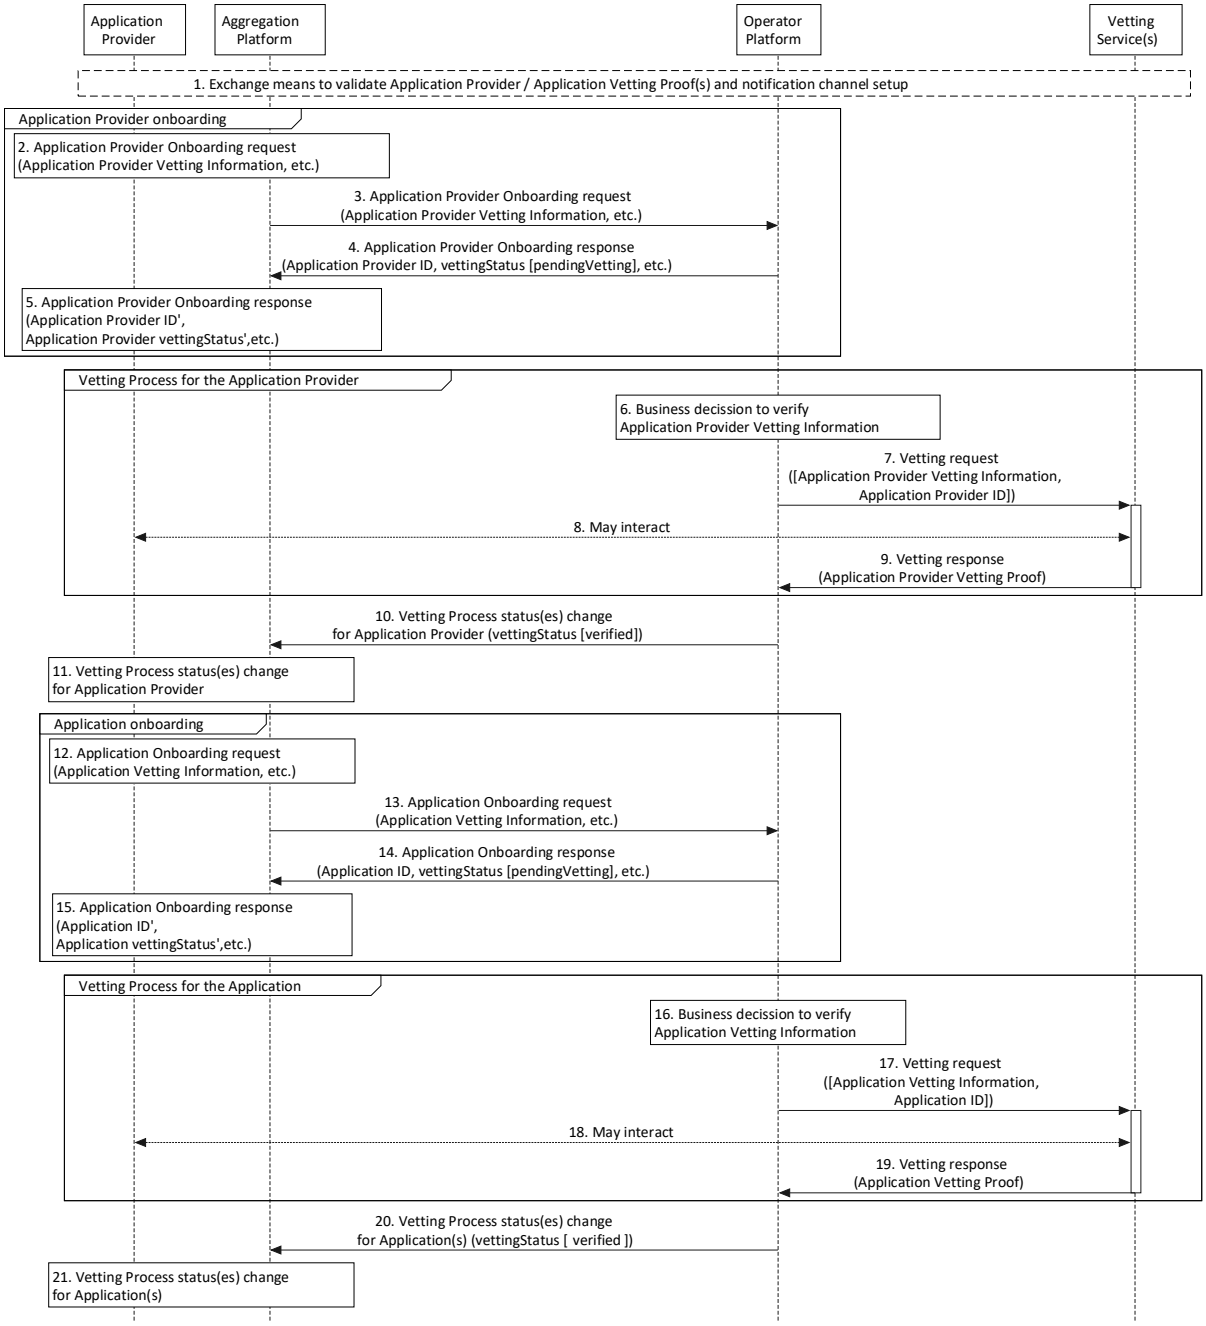

Figure 22: High-level flow for Vetting Process during initial onboarding

1. A trust relationship and notification channels are established among Application Provider(s), Aggregator(s), CSP(s) and vetting service providers (if applicable) to enable a) creating, providing and validating the Vetting Proofs for both Application Provider(s) and Application(s) and b) prompt notification of relevant events associated to the Vetting Process.
2. The Application Provider triggers the Application Provider onboarding process with the Aggregator (via Aggregator's channels) sending, among others, the Application Provider Vetting Information.
3. The Aggregator continues the Application Provider onboarding process towards the CSP forwarding the Application Provider Vetting Information.

4. The CSP will process the Application Provider onboarding request, which includes generating an Application Provider identifier and will inform the Aggregator that the Vetting Process of the Application Provider Vetting Information is pending.

Note: How some of the mechanisms in [22] can be reused for signalling the Vetting Process status is for further study / CSP's prerogative

5. The Aggregator will continue the process of the Application Provider onboarding request, which includes generating an Aggregator-specific Application Provider identifier and will let the Application Provider know that the Vetting Process is pending (via Aggregator's channels).
6. The CSP makes a business decision to trigger the Vetting Process for the Application Provider Vetting Information.
7. The CSP will submit the Vetting Process (via offline interactions) to the vetting service (which may be provided by the CSP itself).
8. The Vetting Process will be triggered and may require offline interactions with the Application Provider.
9. Once the Vetting Process is done, the vetting service will issue the Application Provider Vetting Proof.
10. The CSP will let the Aggregator know that the Vetting Process is finished and provide the updated Vetting Process status(es).

Note: How some of the mechanisms in [22] can be reused for signalling the Vetting Process status is for further study / CSP's prerogative

11. The Aggregator will inform the Application Provider that the Vetting Process is finished and provide the updated Vetting Process status(es) (via Aggregator's channels).
12. The Application Provider performs the Application onboarding process with the Aggregator (via Aggregator's channels) sending, among others, the Application Vetting Information.
13. The Aggregator continues the Application onboarding towards the CSP forwarding the Application Vetting Information.
14. The CSP will process the Application onboarding request, which includes generating an Application identifier and will inform the Aggregator that the Vetting Process of the Application Vetting Information is pending.

Note: How some of the mechanisms in [22] can be reused for signalling the Vetting Process status is for further study / CSP's prerogative

15. The Aggregator will continue the process of the Application onboarding request, which includes generating an Aggregator-specific identifier and will let the Application Provider know that the Vetting Process is pending.
16. The CSP makes a business decision to trigger the Vetting Process for the Application Vetting Information.
17. The CSP will submit the Vetting Process (via offline interactions) to the vetting service (which may be provided by the CSP itself).

Note: The vetting service verifying the Application Vetting Information may be different from the one verifying the Application Provider Vetting Information.

18. The Vetting Process will be triggered and may require offline interactions with the Application Provider.
19. Once the Vetting Process is done, the vetting service will issue the Application Vetting Proof.
20. The CSP will let the Aggregator know that the Vetting Process is finished and provide the updated Vetting Process status(es).

Note: How some of the mechanisms in [22] can be reused for signalling the Vetting Process status is for further study / CSP's prerogative

21. The Aggregator will inform the Application Provider that the Vetting Process is finished and provide the updated Vetting Process status(es) (via Aggregator's channels).

An already onboarded Application Provider may also inform about changes on the Application Provider / Application Vetting Information to the Aggregator and CSP which in turn may trigger a Vetting Process for validating the updated information, as shown in Figure 23.

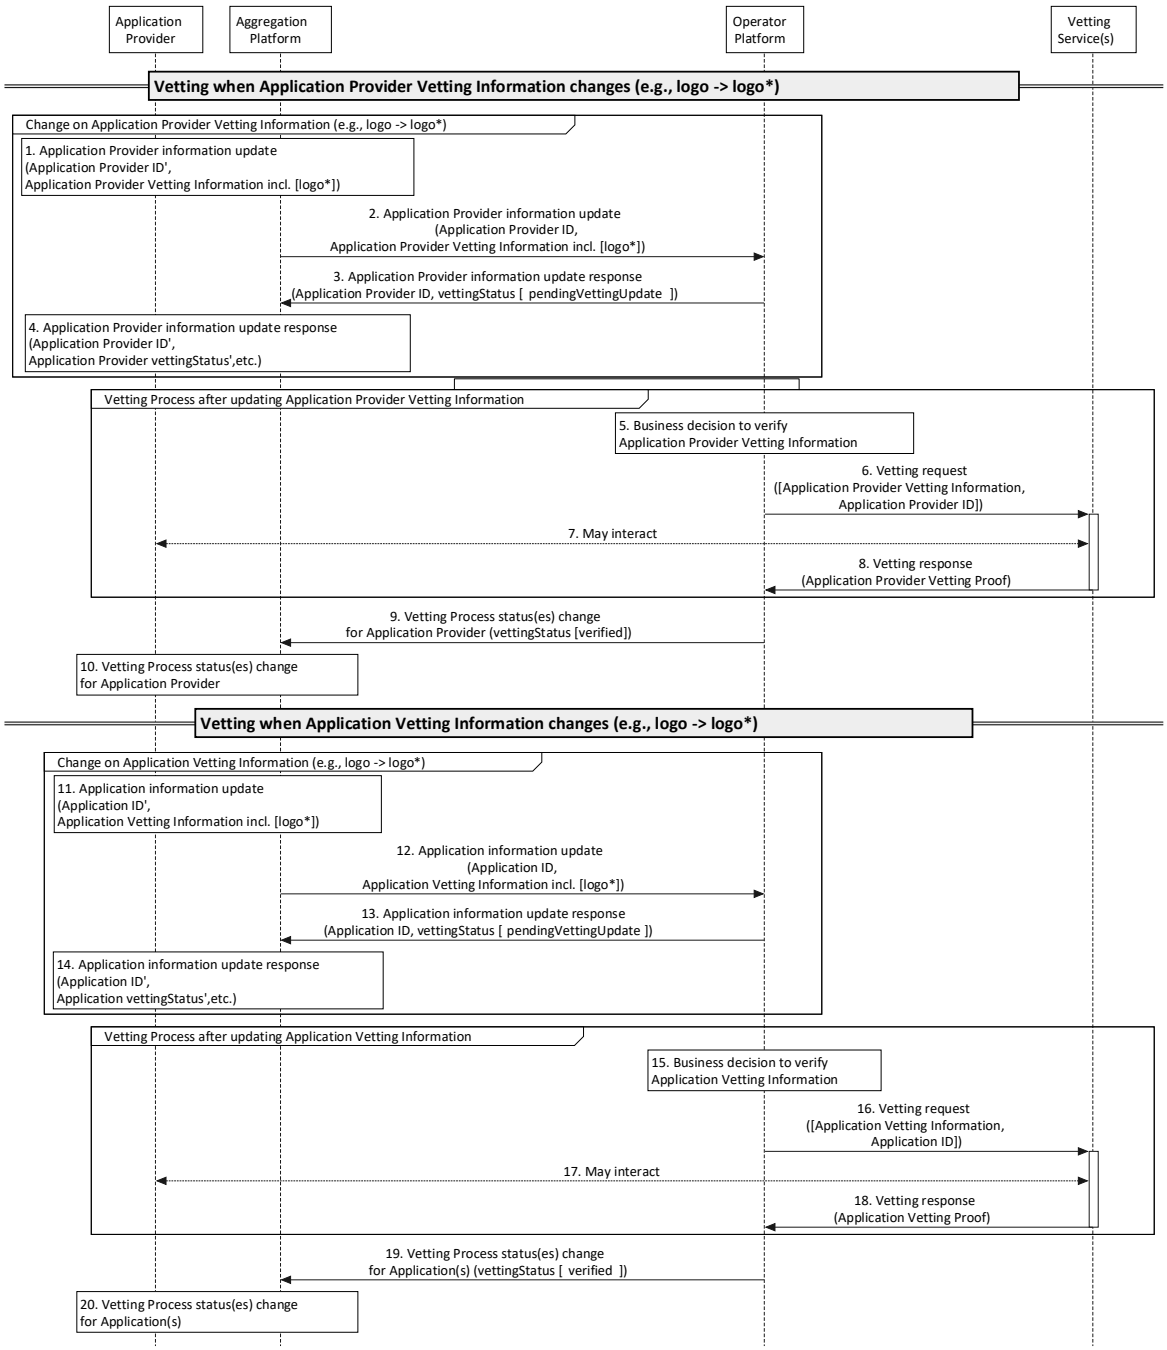

Figure 23: High-level flow for Vetting Process after initial onboarding

If an Application Provider has already been onboarded, and there is a change of the Application Provider Vetting Information afterwards, e.g., change in the logo of the Application Provider,

1. The Application Provider will inform the Aggregator about that change (via Aggregator's channels).
2. The Aggregator will inform the CSP about the changes of the Application Provider Vetting Information.
3. The CSP will process the change of the Application Provider Vetting Information and inform the Aggregator that a new Vetting Process due to changes of the Application Provider Vetting Information is pending.

Note: How some of the mechanisms in [22] can be reused for signalling the Vetting Process status is for further study / CSP's prerogative

Note: The handling of End-User-facing operations by the OP should consider only the latest approved Vetting Information.

4. The Aggregator will inform the Application Provider that a new Vetting Process due to changes of the Application Provider Vetting Information is pending (via Aggregator's channels).
5. The CSP makes a business decision to trigger the Vetting Process due to the change of the Application Provider Vetting Information.
6. The CSP will submit the Vetting Process (via offline interactions) to the vetting service (which may be provided by the CSP itself).
7. The Vetting Process will be triggered and may require offline interactions with the Application Provider.
8. Once the Vetting Process is done, the vetting service will issue the Application Provider Vetting Proof.
9. The CSP will let the Aggregator know that the Vetting Process is finished and provide the updated Vetting Process status(es).

Note: How some of the mechanisms in [22] can be reused for signalling the Vetting Process status is for further study / CSP's prerogative

10. The Aggregator will inform the Application Provider that the Vetting Process is finished and provide the updated Vetting Process status(es) (via Aggregator's channels).

A similar procedure will take place when there is a change in the Application Vetting Information:

11. If an Application has already been onboarded, and there is a change of Application Vetting Information afterwards, e.g., change in the logo of the Application, the Application Provider will inform the Aggregator about that change (via Aggregator's channels).
12. The Aggregator will inform the CSP about the changes of the Application Vetting Information.
13. The CSP will process the change of the Application Vetting Information and inform the Aggregator that a new Vetting Process due to changes of the Application Vetting Information is pending.

Note: How some of the mechanisms in [22] can be reused for signaling the Vetting Process status is for further study / CSP's prerogative

Note: The handling of End-User-facing operations by the OP should consider only the latest approved Vetting Information.

14. The Aggregator will inform the Application Provider that a new Vetting Process due to changes of the Application Vetting Information is pending (via Aggregator's channels).
15. The CSP and the Aggregator make a business decision to trigger the Vetting Process due to the change of the Application Vetting Information.

16. The CSP will submit the Vetting Process (via offline interactions) to the vetting service (which may be provided by the CSP itself).
17. The Vetting Process will be triggered and may require offline interactions with the Application Provider.
18. Once the Vetting Process is done, the vetting service will issue the Application Vetting Proof.
19. The CSP will let the Aggregator know that the Vetting Process is finished and provide the updated Vetting Process status(es).

Note: How some of the mechanisms in [22] can be reused for signaling the Vetting Process status is for further study / CSP's prerogative

20. The Aggregator will inform the Application Provider that the Vetting Process is finished and provide the updated Vetting Process status(es) (via Aggregator's channels).

G.5.2 Relevant information

G.5.2.1 Vetting Process for Application Providers

Information used for keeping records of the latest Application Provider Vetting Information (Application Provider ID, Application Provider name, Application Provider logo), the corresponding Vetting Process status and some metadata associated to the Vetting Process.

Information	Description
Application Provider ID	Unique identifier (generated by the CSP) of the Application Provider
Application Provider name	Part of the Application Provider Vetting Information indicating the commercial name of the Application Provider
Application Provider logo	Part of the Application Provider Vetting Information indicating the logo of the Application Provider
Vetting Process status(es)	Part of the Application Provider Vetting Proof(s) indicating the result of the Vetting Process (e.g., "Verified", "Unverified", "Pending – triggered by onboarding", "Pending – triggered by information update", "Replaced", "Deprecated", etc.)
Vetting Service(s) provider identity	(If applicable) Part of the Application Provider Vetting Proof(s) indicating the identity of the party providing Vetting services
Vetting date / time(s)	Part of the Application Provider Vetting Proof(s) indicating the dates / times when the Vetting process was triggered / finished
Vetting expiration date(s)	(If applicable) Part of the Application Provider Vetting Proof(s) indicating the expiration date of the Vetting Process status

Table 8: Relevant information for the Vetting Process of Application Providers

Note: Table 8 suggests information placeholders rather than implementation.

G.5.2.2 Vetting Process for Applications

Information used for keeping records of the latest Application Vetting Information (Application ID, Application name, Application logo, Information for technical Vetting), the

corresponding Vetting Process status and some metadata associated to the Vetting Process.

Information	Description
Application ID	Unique identifier (generated by the CSP) of the Application
Application name	Part of the Application Vetting Information indicating the commercial name of the Application
Application logo	Part of the Application Vetting Information indicating the logo of the Application
Information for technical Vetting	Part of the Application Vetting Information indicating technical information associated to the Application, for instance redirect URLs format / ownership (applicable for authorization flows that imply HTTP redirects), etc.
Vetting Process status(es)	Part of the Application Provider Vetting Proof(s) indicating the result of the Vetting process (e.g., “Verified”, “Unverified”, “Pending – triggered by onboarding”, “Pending – triggered by information update”, “Replaced”, “Deprecated”, etc.)
Vetting Service(s) provider identity	(If applicable) Part of the Application Vetting Proof(s) indicating the identity of the party providing Vetting services
Vetting date / time(s)	Part of the Application Vetting Proof(s) indicating the dates / times when the Vetting process was triggered / finished
Vetting expiration date(s)	(If applicable) Part of the Application Vetting Proof(s) indicating the expiration date of the Vetting Process status

Table 9: Relevant information for the Vetting Process of Applications

Note: Table 9 suggests information placeholders rather than implementation.

Note: While capturing Consent, the CSP shall present to the End-User only a subset of the non-technical information elements in Table 9. Thus, the End-User is informed in a simple manner on which entity is trying to obtain access to privacy-sensitive information.

Annex H Document Management

H.1 Document History

Version	Date	Brief Description of Change	Approval Authority	Editor / Company
1.0	30 Jun 2021	New PRD OPG.02, based on requirements proposed in OPG.01.	ISAG	Tom Van Pelt / GSMA
2.0	14 Apr 2022	Update implementing OPG.02 CR1002	ISAG	Tom Van Pelt / GSMA
3.0	03 Oct 2022	Update implementing OPG.02 CR1003	ISAG	Tom Van Pelt / GSMA
4.0	29 Mar 2023	Update implementing OPG.02 CR1004	ISAG	Tom Van Pelt / GSMA
5.0	26 Jul 2023	Update implementing OPG.02 CR1005	ISAG	Tom Van Pelt / GSMA
6.0	16 Feb 2024	Update implementing OPG.02 CR1006	ISAG	Tom Van Pelt / GSMA
7.0	20 Sep 2024	Update implementing OPG.02 CR1007	ISAG	Tom Van Pelt / GSMA
8.0	28 Feb 2025	Update implementing OPG.02 CR1008	ISAG	Tom Van Pelt / GSMA
9.0	09 May 2025	Update implementing OPG.02 CR1009: moving requirements into new GSMA PRDs OPG.11 and OPG.12	ISAG	Tom Van Pelt / GSMA
10.0	18 July 2025	Update implementing OPG.02 CR1010	ISAG	Tom Van Pelt / GSMA
11.0	19 Feb 2026	Update implementing OPG.02 CR1011	ISAG	Tom Van Pelt / GSMA

H.2 Other Information

Type	Description
Document Owner	Operator Platform Group
Editor / Company	Tom Van Pelt / GSMA

It is our intention to provide a quality product for your use. If you find any errors or omissions, please contact us with your comments. You may notify us at prd@gsma.com

Your comments or suggestions & questions are always welcome.