Operator Platform Evolution: Use Cases and Concepts

Version 0.1

February 2021

1. Introduction

OPG.01 [1] proposed an initial set of proposed requirements for the Operator Platform (OP), focusing on exposing resources to support Edge Computing workloads provided by 3rd party Application Providers. The Operator Platform Group (OPG, i.e. the group owning OPG.01), requested feedback on those requirements from the industry, Standards Development Organisations (SDO) and open source communities. The group is following up on that now, developing formal requirements that take into account this feedback.

To verify that no significant gaps exist in the proposed requirements, the OPG has defined a set of use cases, covered in section 2. Any such gaps will be closed in the formal requirements.

In parallel to incorporating feedback to the requirements and developing the use cases, the OPG has started activities to cover some areas that couldn't be covered in OPG.01 [1]. The OPG is looking into these areas in order of priority. The high-level concepts defining the direction for the first of these areas, i.e. those considered the most important, are stable by now. This document aims to inform on that direction.

1.2 Abbreviations

Term	Description
ePDG	Evolved Packet Data Gateway
EWBI	East-WestBound Interface
N3IWF	Non-3GPP InterWorking Function
NBI	NorthBound Interface
NEF	Network Exposure Function
OP	Operator Platform
OPG	Operator Platform Group
PGW	Packet GateWay
SBI-CR	SouthBound Interface – Cloud Resources
SBI-NR	SouthBound Interface – Network Resources
SDO	Standards Development Organisation
UNI	User-Network Interface
UPF	User Plane Function

1.3 References

Ref	Doc Number	Title
[1]	OPG 0.1	Whitepaper: Operator Platform Telco Edge Proposal – Version 1.0, 22 October 2020 https://www.gsma.com/operatorplatform/
[2]		Telco Edge Cloud: Edge Service Description & Commercial Principles Whitepaper, version 1.0, 27 October 2020 https://www.gsma.com/operatorplatform/

2 Use Cases

This section introduces a set of use cases that the Operator Platform Group developed to verify whether gaps exist in the requirements proposed in OPG.01 [1]. The OPG has selected these use cases for their breadth of functional coverage rather than embark on the impossible journey of defining an exhaustive set of use cases that benefit from federated edge computing. Collectively the use cases illustrate some of the critical capabilities that an OP will have to provide.

2.1 Automotive - Advanced Horizon

2.1.1 Description

A driver gets “look ahead” information about the local vicinity – for example, a patch of ice, a slow-moving tractor or red traffic lights. A driver’s ability to see “around the corner” could help safer and more economical driving.

The driver could be a human – as seen in today’s Advanced Horizon products from Bosch™ and Continental™ – or, in the future, it could be an automated driver.

2.1.2 OP Dependency

The service could be delivered through an application server on a cloudlet that gathers information from roadside sensors and nearby vehicles. The application server would aggregate this data and analyse it to send updates to vehicles in the vicinity. These updates can be more accurate and timely if the application server gets information from all nearby vehicles, potentially on several mobile operators. A federation of OPs would enable such information exchange between application servers on different operators or direct access from the devices.

Next to that, this service has essential security and trustworthiness requirements – both for the information reported by roadside sensors and other cars, and the analysis performed by the application server. An operator platform that authenticates the parties supplying the data, verifies applications and is involved in their discovery will provide the guarantees required for such a service.

2.2 Automotive – Remote Driving

2.2.1 Description

The second use case is remote driving or flying one or more vehicles or drones. This use case involves someone at a distance controlling the vehicle based on detailed information of its surroundings. Other vehicles might then follow the path set by the one driven or flown remotely without requiring control on an individual basis.

2.2.2 OP Dependency

This use case has similar requirements on trustworthiness and communication to other operators than the use case discussed in section 2.1.

The scenario requires strong guarantees on service assurance – about the network and compute's responsiveness, reliability, and security. Deploying the supporting application at the edge using the Operator Platform for discovery, potentially combined with Network Slicing that the Operator Platform intends to support in a future iteration, may provide those guarantees. Furthermore, a vehicle may have to pass borders and operate in a geographical region that requires other operators for coverage. The Operator Platform would help to ensure that the supporting edge application is available on those networks.

2.3 Multiplayer Augmented Reality Game

2.3.1 Description

The next use case is a multiplayer augmented reality game. Players participate in the real world, supplemented by online features, for example, a role-playing game. The players are thus all nearby but can be on different operators.

2.3.2 OP Dependency

For such a game, preference is that the players share the same application server, which is on a local cloudlet. A "shooter" game, for example, is moderately latency-sensitive, and fairness between players is crucial, requiring that the players all get similar server processing performance and similar network performance. An Operator Platform enabling the sharing of edge nodes between operators would be able to support this.

Some games need specialist compute (e.g. GPU). As indicated in the TEC whitepaper [2], a federated model to deliver an Operator Platform may require alignment between the federated operators to ensure that they offer similar resources. Thus, the party developing the game can expect the same specialist compute capabilities in all networks and consider them in their application design and dimensioning.

2.4 Privacy-preserving Health Assistant

2.4.1 Description

The following use case is a privacy-preserving health assistant. Already there are health-related personal monitors, such as smartwatches in use today. There will be many more personal IoT services, perhaps including actively controlled devices, for example, to automatically adapt an insulin dose based on its measurements.

These devices all provide their data to their dedicated backends without much user control over the handling of the provided data from that point onwards. An edge-based health assistant's appeal could be that it can act as a trusted third-party intermediate capable of aggregating the data from different devices and providing control over the access to that data. By design, the local cloudlet could store data only temporarily. For instance, an application in the cloud would be allowed only specific request types on the cloudlet (e.g. restrict exporting the complete data set).

2.4.2 OP Dependency

When the user roams onto another network, one solution approach is that the (trusted) home operator installs its application server on the local cloudlet.

2.5 Infrastructure sharing

2.5.1 Description

Infrastructure sharing is a technical use case where one operator uses infrastructure provided by the other. Possible examples could include:

- Two operators, each with a mobile network covering the whole country, agree to share edge compute infrastructure (say: one covering the North of the country and the other the South) – this similar to today's sharing of radio masts.
- A virtual OP, which buys access to compute infrastructure and networking capacity, from a ('real') OP.
- An OP has its own 'basic' edge infrastructure, but not the specialist compute or specialist hardware security that some application providers require.
- An OP whose edge compute is currently short of resource temporarily offloads new requests to another OP.

2.5.2 OP Dependency

The main requirement to enable this is for a commercial agreement between the involved OPs covering topics including security and trust, service level agreements and billing.
Note that the whitepaper defines home network control in the roaming case.

2.6 Use Case Overview

Capability	Interface	Whitepaper [1] section	UC 1 "Advance horizon" info for assisted driving	UC 2 Remote control of a vehicle (or drone etc.)	UC 3 Multiplayer AR location-dependent game	UC 4 Privacy-preserving Health assistant	UC 5 Infrastructure sharing
Application Provider **request** for Edge Cloud service	NBI	5.1.1.3 #1	Y	?	Y	Y	N
Provide info on UE's **location**	SBI-NR	5.1.3	Y	Y	Y (& verify?)		
Handover (UE moves in a mobile network) *(Implementation likely to require a move of the application server to a new cloudlet)*	SBI-NR	5.1.1.2.2 #9 5.1.3.2.2 #10	Y	Y	N		
Inter-network **Roaming** (UE roams to another operator) *(Preferably with local breakout, so application server on cloudlet in the visited operator)*	EWBI	5.2.2.4 5.1.2.3 #5	Y - Preferably	Y	Y	Y	
Application Provider requests **QoS** (typically latency)	NBI	5.1.1.3 #2	Y	Y - Critical	Y & 'Fair'	Y - Weak	
Establish a chain of **trust** between the elements	UNI & OP	3.5.3.2	Y	Y		Y - Critical	Extend over EWBI
Security Comms Compute Storage	UNI OP OP	2.1.4, 3.4.1 & missing	Y Y .	Y Y		Y Y Y	
Inter-OP Security		5.2.3.1.2					EWBI
Data sharing (Data is 'open' for use by multiple application providers)		missing	Y			Y - but highly filtered	
Specialist compute	SBI-CR	5.2.2.3			Y		
Shared Application Server	SBI-CR	missing			Y		

N.B. **Y** – indicates that the requirement is of particular importance in the use case
N – indicates that the requirement is not essential, or not needed, in the use case
Blank cell - indicates that the requirement is somewhat helpful for the use case, but not central to it

3 Operator Platform Evolution

The Operator Platform Group is defining high-level concepts to evolve the Operator Platform requirements to cover areas that couldn't be studied in depth when developing OPG.01 [1] (see also section 1.2 of that document). In order of decreasing priority, the group is looking into the following:

1. How the Operator Platform requirements can map to solutions available from or in the process of being developed by SDOs and Open Source organisations

2. A set of use cases for OP-enabled edge computing to assess better whether gaps remain in the requirements (see section 2)

3. Edge Node sharing defining how edge nodes provided by one network could serve a user using another network's radio access

4. Roaming investigating how to provide access to edge compute services for users that are outside of their home network's footprint or are using fixed access

5. Provide a continuous service to devices moving around within a network's footprint

6. Study of how the commercial principles defined in the Telco Edge Cloud whitepaper [2] impact the Operator Platform requirements

7. A more in-depth study of the security aspects of the Operator Platform

8. The call flows that the document should cover

9. The requirements to enable using container technology for applications deployed on edge resources using the Operator Platform

10. Ensuring consistency in the resource reservation for application providers and operators

11. Enabling low latency interactions between applications deployed on edge resources in different networks

12. Providing serverless deployment models

13. A better definition of the management plane for the Operator Platform and the resources exposed by it

14. Study the device-side architecture for enabling edge compute services with the least amount of

Not all of these concepts are stable yet. It cannot be excluded that some will not reach stability in time for inclusion in the first set of formal requirements provided by the Operator Platform Group. The following subsections will cover the most stable areas in more detail.

3.1 SDO Mapping

Following the publication of OPG.01 [1], the Operator Platform Group has engaged with several SDOs and open source communities using Liaison Statements and joint meetings. Through those engagements, the OPG has gathered input on the proposed requirements and how those parties' solutions aligned with them. Several organisations have indicated that their existing or under development solutions, could fit the Operator Platform requirements.
From studying that input, there is overlap between those solutions in some areas and that solutions might complement each other in other areas.

Next to that, the OPG identified gaps in different areas that are important for delivering on the Operator Platform's promise. The OPG is studying how to fill those gaps and whether the timeframe that can be achieved could still fit the Operator Platform market window.

Once that's clear, the next steps will involve bringing together the organisations whose solutions best fit the Operator Platform requirements to align on how to combine their solutions and how to close the gaps. The latter may require approaches to shorten the time to market, such as providing a reference implementation as open-source while SDOs develop formal specifications.

3.2 Edge Node sharing

The edge node sharing item studies how edge resources in one network can be made available to subscribers accessing the service using another network. Operators can use edge node sharing, for example, to improve the geographical coverage of their edge service offering.

The high-level concept chosen to enable edge node sharing is to keep the interaction between the device and the Operator Platform to discover and access those shared edge nodes between the device and the Operator Platform in the network providing access to the subscriber. That Operator Platform will then interact with the Operator Platform in the network sharing the resources to ensure that the requested application is available and obtain the data that the device requires to access those edge nodes.

FIGURE 1

EDGE NODE SHARING CONCEPT

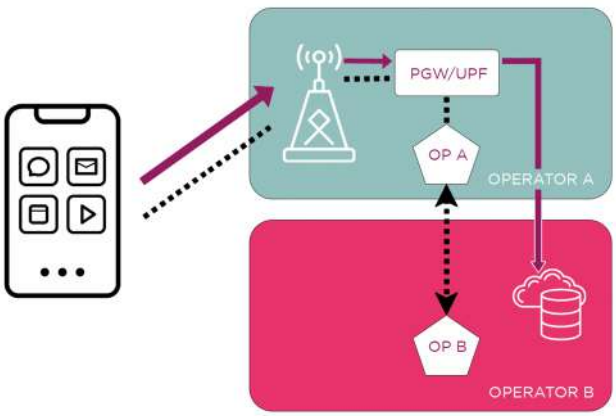

3.3 Roaming

For cellular roaming, the OPG studied two models that will likely exist alongside each other while the ecosystem develops:

1. Home routing, for scenarios where edge services provided by the visited network cannot be supported. The home network OP is the only OP involved in this case.

2. Local breakout, to access edge nodes available in the visited network. In this case, the OP in the visited network will handle the discovery of the actual edge nodes with the home network OP involved in the subscriber's authentication and authorisation.

FIGURE 2

ROAMING ACCESS TO OP AND EDGE RESOURCE - HOME ROUTING

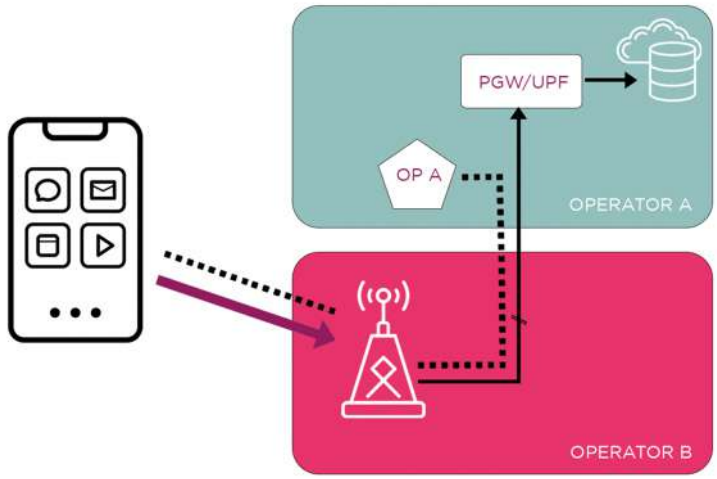

FIGURE 3

ROAMING ACCESS TO OP AND EDGE RESOURCE - LOCAL BREAKOUT

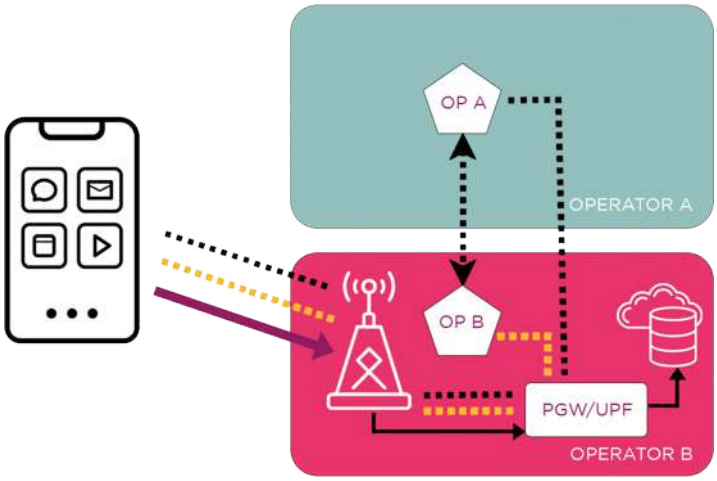

The OPG considers local breakout the preferred approach for the long term because only that approach allows delivering on the low latency promise of edge computing. Local breakout will require time to be universally available because it depends on device support and a technical and commercial model to enable it. The latter is relevant because most other services are using a home routed model. All mechanisms for settlement, monitoring, regulatory compliance, etc. have thus been set up based on that approach.

Both approaches also come with their challenges with regards to application availability. For home routing, the OP should limit application availability to those whose latency requirements can be met in this home routed context. For local breakout, the OP must ensure that the applications that a subscriber needs to have access to, are available in the visited network. The approach is followed that the visited operator is assumed to be in a federation with the subscriber's home operator, i.e. local breakout is linked to use within the federation.

This approach does simplify that problem because it allows using the federation to control an application's availability rather than defining roaming specific concepts. The more generic problem of roaming on a network that is not in the same federation has been deferred to a later phase. The OPG assumes that this more generic case's technical impact will be limited, but it will be much more complex from a commercial perspective

3.4 Fixed Access

When connected to a "fixed" network such as a Wi-Fi connection, access to OP services is to some extent similar to the roaming cases discussed in section. Services could be provided by edge resources in the mobile network (i.e. "home-routed"), enabling access over packet-core integrated Wi-Fi (i.e. via ePDG/N3iWF) or by resources in the fixed network (i.e. "local breakout"). A similar role split between the OP in the mobile network and the one on the fixed network side is possible for this latter case to the role split used for cellular roaming with local breakout (see section 3.3).

For this case of "local breakout on a fixed network", further alternatives exist because the device can also be considered as a fixed device with capabilities to use the mobile network. That allows being less dependent on the mobile subscription compared to the cellular roaming case. That looser dependency would enable the OP to in the mobile network to pass full control of the subscriber to the Operator Platform in the fixed network. The OP in the mobile network could provide just authentication services or leave even that to the fixed network's OP.

3.5 Mobility

The OPG considered mobility aspects as covered mostly in the requirements provided in OPG.01 [1] already. Therefore, further updates will be handled with lower priority and focus on delivering session continuity or enabling other ways to support stateful applications.

3.6 Alignment with Commercial Principles

Based on the study of the principles described in the Telco Edge Cloud whitepaper [2], the Operator Platform Group sees the need to enhance the requirements in the following ways:

1. Support for an Infrastructure as a Service model is needed

2. Enable Multi-access support including fixed, Wi-Fi and Mobile access networks (see sections 3.3 and 3.4)

3. Ensure that the Northbound and East-Westbound interfaces support the management and administration functionality defined in section 3.3 of the Telco Edge Cloud whitepaper [2].

4. Include requirements to ensure that the Northbound interface can be called directly through APIs from DevOps tools rather than being limited to a portal

5. Provide requirements ensuring that the East-Westbound and Northbound interfaces cover the charging principles defined in section 5 of [2]

The use of an Infrastructure as a service approach will mean that OP provided enablers like edge discovery and allocation, roaming support, etc. are not available. That limitation follows from an IaaS offering not including the OP's Southbound interface to the network resources and the interactions between the OP and the device on the UNI.

To support the proposed management functionality requirements need to be included on things like

- trouble ticketing in both directions between operator and platform provider,
- support of portals that enable the management and administration tasks of the operator and application provider
- Support for a security framework that allows verifying the onboarded applications
- Support for order management and
- Onboarding

The charging principles will result in most of the added complexity. Charging will require integrating various functions in the Operator Platform architecture with a charging function and billing functionality. The federation broker would be among the functions that need such integration. Next to that, the federation broker should support clearinghouse capabilities. The OP could either integrate with the charging function directly in the OP or access it through the NEF. Next to settling that approach, the OPG will look into the triggers for charging related events and what units to measure.

3.7 Security

The security topic is looking into what the OPG needs to do to ensure that the OP specification is fit for use from a security perspective. The OPG will verify all aspects of the OP from a security perspective, assessing the trade-offs between security, functionality and performance and will cooperate with the GSMA Fraud and Security Group.

The security topic will take into account existing threat models for cloud, edge and fog computing. The areas considered include, among others:

- Trust domains (between federated partners in particular)
- Key and certificate management
- Traffic routing
- Authentication of application providers, clients and applications
- Security of the SDO specifications considered as part of the SDO mapping (see section 3.1)
- Etc.

3.8 Containers

The OP intends to provide developers with a consistent deployment environment independent of the network in which they deploy their applications. The OPG will enhance the requirements to ensure that networks support a consistent environment for handling containers (e.g. image format, Host OS, CPU architecture, etc.). Requirements may not go as far as mandating specific container run-times or orchestration engines like Kubernetes because networks may need the flexibility to fit in the edge resources in their existing environments, but that is part of the assessment of this area.

3.9 Low-latency Interconnects

Various forms of low-latency interconnects between networks have been studied (e.g. to another operator's cloudlets covering a similar area, to foreign networks in support of applications requiring minimal latency, the interconnect required to support edge node sharing, etc.). For most, the OPG has concluded that they better fit into Network as a Service use cases that the group may cover after finalising the first version of the formal Operator Platform requirements.

The exception would be the interconnect required to support Edge Node sharing and possibly the home routed roaming scenario. Still, the details are considered out of scope because the nature of that interconnect will depend upon the logistics and/or business considerations of the involved operators' edge locations and mobile networks.