TELECOM INFRA PROJECT

Open BNG
Technical Requirements

Authors:

- **Juan Rodriguez**
    - Technological Expert, CTIO, Telefónica
- **Rafael Canto Palancar**
    - Transport & IP Network Manager, CTIO, Telefónica
- **Paul Gunning**
    - Principal Researcher, BT
- **Peter Willis**
    - Senior Manager Software Based Networks, BT
- **Konstantinos Antoniou**
    - Senior Researcher, BT
- **Mario Kind**
    - Research & Development Engineer, Deutsche Telekom AG
- **Jonathan Newton**
    - Principal Engineer, Vodafone
- **Brendan Black**
    - Network Engineer, Vodafone

Contributors:

- **Rafael A. Lopez da Silva**
    - Technological Expert, CTIO, Telefónica
- **Victor Lopez Alvarez**
    - Technological Expert, CTIO, Telefónica
- **Jesús Luis Folgueira Chavarria**
    - Transport Senior Manager, CTIO, Telefónica

Table of Contents

Introduction **6**
Why Open BNG? 6
Scope of the document 7
Document structure 7
Network Architecture and Open BNG types **8**
Traditional transport networks for fixed services and new trends 8
Scenario BNG + router 8
Scenario BNG + router + enterprise PE 9
Regulatory requirements for wholesale and legal intercept 10
Fixed/mobile convergence 10
Open BNG types 10
Full functionality and maximum distribution 10
Full functionality and medium location 11
Service Only BNG Medium/Reduced Distribution 11
Platform Architecture & Hardware Specifications **1**2
Hardware Solution form factor and environmental conditions 12
Hardware Platform CPU and ASIC 12
Hardware Platform power supply and cooling requirements 13
Hardware Platform management 13
Hardware Platform Synchronization requirements 13
Hardware SKU network interfaces and forwarding capacity 13
BNG Software package (BSW): subscriber management capabilities **1**7
Encapsulation methods and protocols 17
PPPoE encapsulation 17
IPoE plus DHCP service 17
L2TP 17
VPN 17
PPPoE/IPoE over PW 18
IPv6 18
Authentication, Authorization and Accounting (AAA) 18
Quality of Service 19
Classification 19
Marking 19
Scheduling 20
Traffic rate limit 20
Per subscriber shaping values are part of the AAA policies. 20
Hierarchical QoS 20
Multicast 21
Lawful Interception 21
Router Software package (RSW): routing capabilities for Open BNG **22**
IP/MPLS routing 22
Link aggregation 23
BFD 23
Management 23
Monitoring 23
Security 23

Provider Edge Software package (PESW): Open BNG as a PE for enterprise services 26
L2VPN 26
L3VPN 26
E-VPN 26
IPSec 27
Open BNG programmability and SDN 28
Telemetry and SDN 28
Access security and Anti-Theft 29
Additional requirements 30
Port mirroring 30
Configuration and versions management 30
Zero-Touch Provisioning 30
Licensing 30
Local regulation compliance 31
Glossary 32

1. *Introduction*

This document represents the technical requirements for an open and disaggregated broadband network gateway (BNG) device that operators can deploy in current and future networks for the provision of fixed broadband services. It describes the required hardware and proposes non-mutually exclusive software packages for the support of additional services or functionalities.

It includes the role of software-defined networks (SDN) and the desired approach concerning fixed-mobile convergence. And it includes regulatory requirements to deploy Open BNG in networks of operators participating in this requirements document.

1.1. Why Open BNG?

The project goal is to develop a solution that overcomes some

of the most relevant issues operators presently face when deploying access services for fixed customers (e.g., residential, SOHO, SME). A short-term challenge for these services is continuous traffic growth, which not only affects required performance, but also determines the most adequate location for BNG functionality (i.e., more centralized vs. more distributed).

Other topics, such as the search for an appropriate fixed-mobile convergence strategy, or the inclusion of SDN methods for improved service provisioning, will also impact future BNG node specifications.

In this environment, operators find that current solutions:

- are based on monolithic platforms that make it extremely difficult to introduce innovation from other vendors in disparate parts of the device stack:
    - Lack of open hardware that can run various software types
    - Lack of open software that permits feature extensibility
    - Lack of fully open APIs that enable external components to interact with a device
- lack standardized (or at least globally agreed) features and tools for zero touch provisioning (ZTP) and initial auto-configuration
- demand per-vendor system integration that affects costs and time-to-market
- present a lock-in to same-vendor pluggable modules not based on technology
- represent a strong challenge in multi-vendor environments

This set of technical requirements aims to define an open and disaggregated platform that:

- is based on commercial, off-the-shelf components and open software that can perform traditional BNG functions
- reduces deployment and operational costs
- provides the scalability required for multiple scenarios

1.2. Scope of the document

The aim of this document is to describe:

- the desired Open BNG platform architecture—including requirements that will need to be met in relation to hardware and software features
- the end-to-end network architecture where the platform will need to operate, presenting disparate use cases affecting its interface layout and node capacity
- management of Open BNG combined with SDN controllers

Immediately following the publishing of this document, a detailed,

low-level TRS (technical requirement specification) will be undertaken as a basis for further technical discussions with platform hardware and software providers.

1.3. Document structure

This document is structured as follows:

- Chapter 1: Introduction
- Chapter 2: Network Architecture & Open BNG types
- Chapter 3: Platform Architecture & Hardware specifications
- Chapter 4: BNG Software package (BSW): subscriber management capabilities
- Chapter 5: Router Software package (RSW): routing capabilities for Open BNG
- Chapter 6: Provider Edge Software package (PESW): Open BNG as a PE for enterprise services
- Chapter 7: Open BNG and SDN
- Chapter 8: Additional requirements
- Chapter 9: Glossary

2. *Network Architecture and Open BNG types*

In this section, we introduce the Open BNG deployment models and the most typical network topologies where it will operate. In general, these aspects will have an impact in the hardware layout, the required capacity, and the demanded additional functionalities (i.e. software packages).

2.1. Traditional transport networks for fixed services and new trends

There are many different topologies deployed by different network operators. The intention of this section is not to cover all those in detail, not even a meaningful sample. Instead, we will start from a very typical network architecture (access, aggregation and core, Figure 1) to describe the reasoning behind the different deployment scenarios for the Open BNG (and the differentiated software packages of this specification).

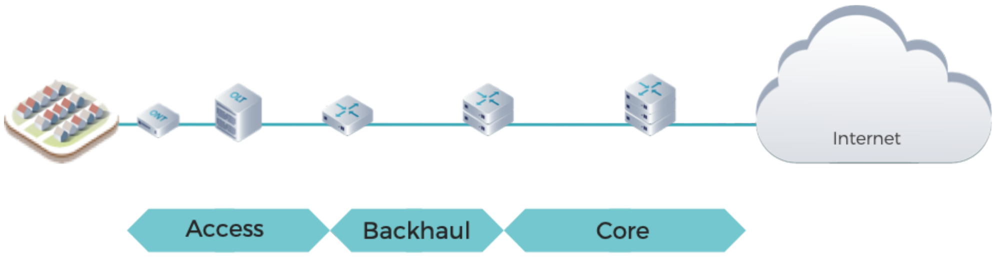

Figure 1. Traditional transport network for fixed services

In the figure above, access nodes (DSLAM[1] or OLT) in the access network are directly connected to a metropolitan-wide network segment, at which aggregation of regional traffic begins. The level of aggregation keeps growing upwards and, eventually, traffic is sent to a network backbone for national forwarding and exchange with external networks.

Technologies in each of these segments are varied. While we can easily assume that core networks are widely IP/MPLS based, aggregation networks present all sorts of possibilities: IP/MPLS, MPLS flavours like MPLS-TP, carrier Ethernet, the growing VxLAN, more recent trends like Segment Routing... The conclusion is that the location of the BNG functionality in the operators' networks has a great impact, not only in the required performance, but also in the type of features that will be demanded from the box.

2.1.1. Scenario BNG + router

A very typical scenario some years ago, when boxes were not so powerful, is the one presented in Figure 2. In the example figure, the BNG is located between the aggregation network and the core network, which was a very adequate location also to ensure isolation between the two. Indeed, the ingress interface of the BNG in these deployments was normally based on VLANs (typically QinQ), and that provided a beneficial segmentation of the transport protocols. However, the exact location of the BNG depended mainly on its scalability (i.e., total number of subscribers). This way, there were operators who deployed smaller BNG closer to the customers: the aggregation area before the BNG was normally smaller, and there probably was an IP-based regional network upwards, before the national core. However, in essence, these approaches were very similar.

In both cases, the key idea from the figure is that BNG boxes lacked, at that time, the capabilities to run the most modern transport protocols (i.e., MPLS). A very common solution was to deploy, jointly with the BNG, an Access Router towards the next segment, supporting the protocols the BNG was not capable of.

[1] For this specification, only IP DSLAMs will be considered. No ATM-related requirements will be demanded from the Open BNG.

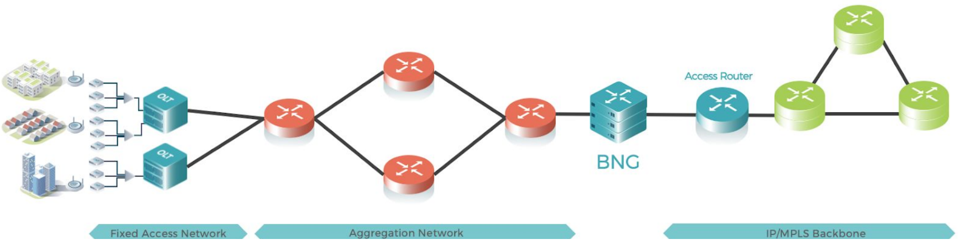

Figure 2. Scenario BNG + router

Today, this approach is very rare in the industry, as both functionalities (BNG and Router) have collapsed into a single box in most vendor implementations. However, when defining software requirements for the Open BNG, this traditional view explains the separation in this specification between the basic subscriber management features (BNG SW package – chapter 4) and the routing requirements (Router SW package – chapter 5). It is also fair to mention that, in past years, a lot of work has been done at TIP (and the telco ecosystem) to understand the maturity of routing protocols in disaggregated solutions (e.g. TIP's DCSG testing). On the other hand, evaluation of subscriber management features is less extensive. Thus, separating them simplifies the (expected) upcoming evaluation of proposals from the industry.

2.1.2. Scenario BNG + router + enterprise PE

In deployment scenarios like that of Figure 2, nation-wide enterprise services are typically provided by concatenating VPNs across the two (or more) regional networks and through the core. For this, new PE nodes exclusively for enterprise services may be deployed (Figure 3), in which case it is not desired to use the BNG plus router solution to terminate/start these VPN. Independently of whether there are technical or administrative reasons behind this separation, this possibility also exists when both features (BRAS and access router) are collapsed in a single box.

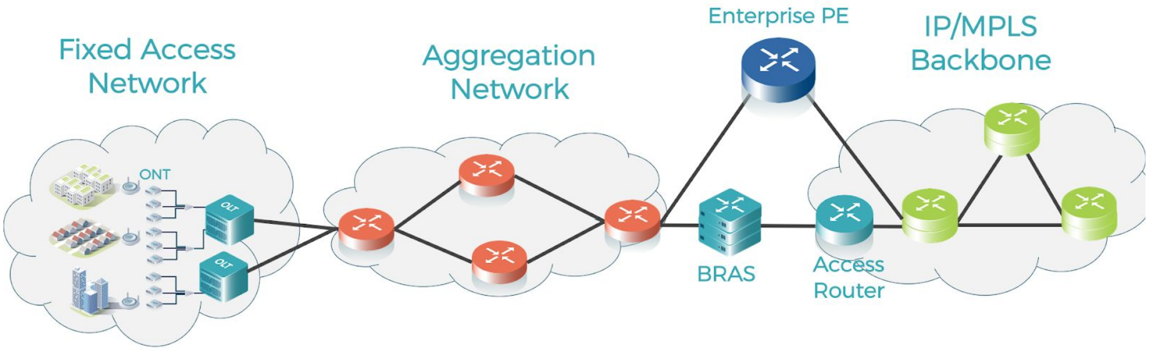

Figure 3. Scenario BNG and separated enterprise PE

However, there are alternative scenarios where operators demand PE functionalities integrated in the same box as the BNG (the Broadband Forum defines this integrated approach as a 'Multi-Service BNG' as described in TR-178).

In centralized deployments, implementing for example Seamless MPLS, the BNG can be considered a Service Node, supporting MPLS on both the uplink and downlink. In fact, the encapsulation of subscriber sessions over an MPLS pseudowire is widely supported. For scalability reasons, operators may easily still choose to provide national VPN services by dividing the VPN in several segments, using the BNG box to do the stitching.

Also, due to the big traffic growth experienced in the last few years, some operators have decided to distribute the BNG function, deploying it closer to the customers. In scenarios like this, it is common that the first aggregation node (i.e., where access nodes are connected) becomes the BNG, and also the enterprise PE (Figure 4). Furthermore, this node may participate also in the distribution of video services, where multicast capabilities are essential.

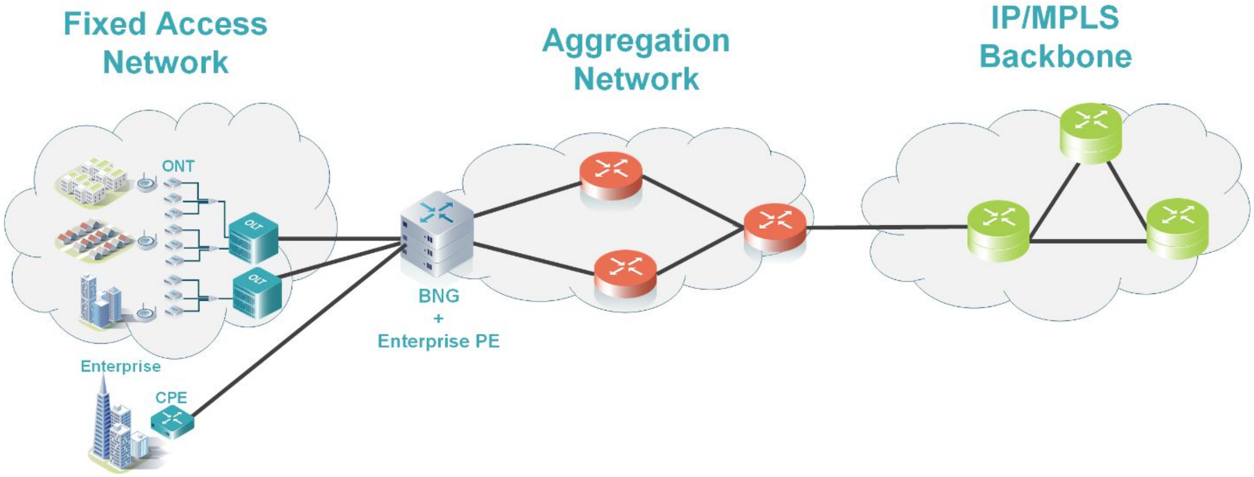

Figure 4. Scenario BNG + enterprise PE

Any of these reasons leads to the definition of the Provider Edge SW package for the Open BNG of this specification (chapter 6).

2.1.3. Regulatory requirements for wholesale and legal intercept

There are two important regulatory requirements for operators where a BNG is of importance.

First this is the requirement for granting wholesale access by operators holding significant market power. Typically, this is realised by Layer 2-bit stream access (L2BSA) in various implementation options with VLAN, L2TP or other tunneling protocol approach-based forwarding. Depending on deployment and regulatory models, the BNG is responsible for forwarding traffic to wholesale partner networks. Exact accounting plays an important role for wholesale connections. Potentially network connections from customers connected via wholesale partner networks are terminated at the BNG as well. Wholesale could be implemented as a function on the BNG or as a separate device connected to the BNG.

The other legal requirement is to provide access to customer traffic based on telecommunication law. This typically requires additional functionality like traffic duplication, encryption of traffic, time stamping, etc. It is typically implemented close to the BNG at the point where all traffic of a customer traverses (in contrast to other network locations because of the BNG routing function and redundant network paths between BNG and core network). Similar to wholesale this might be a function implemented in the BNG or as a separate device connected to the BNG.

2.1.4. Fixed/mobile convergence

There are multiple approaches that operators are following to try and take advantage of convergence techniques between fixed and mobile networks, elements or functions.

Convergent transport of fixed and mobile traffic is a common requirement. This means that it is possible that the BNG box participates in the mobile backhaul (MBH), or at least stands in the middle of the transmission path for the synchronization signals. The conclusion is that both Sync-E and PTP will be requirements for the Open BNG, the same as for any other node in the MBH. These will be covered in section 3, as they have a strong dependency with the available hardware.

In particular, based on Control and User Plane Separation (CUPS), there exist efforts in both planes, like:

- Implementing, for example, the required policy control of both services in a single network component in charge of a converged control plane.
- Implementing, for example, the mobile User Plane Function (UPF) and the data plane of a BNG in the same box.
- Implementing, for example, an access gateway function (AGF as per BBF TR-470) within the BNG to more fully integrate fixed subscribers with the 5G Core.

None of these requirements will be mandatory for solutions complying with this specification, although any advance in the convergence direction will be valued. Of great importance for fixed access solutions, however, is to make sure that no solution for convergence has an impact on the current complexity and cost of fixed services. Proposals that modify the architecture and processes of fixed broadband to assimilate mobile paradigms will be strongly discouraged, unless fully justified economically.
More discussion on SDN and network programmability like CUPS will be discussed in section 7.

2.2. Open BNG types

This section describes different flavours for Open BNG platforms, typically formed by the sum of different software (SW) packages and the deployment location, which will have an impact in the scalability numbers and the proposed standard configurations.

2.2.1. Full functionality and maximum distribution

In this flavour (Figure 5), the Open BNG is totally distributed, so access nodes are directly connected to this box, as are the business CE nodes and the MBH (or more generically mobile x-haul). The Open BNG will typically deal with VLANs towards the access network, but it must be possible to act as a full router towards the MBH, and of course, towards the aggregation network; this explains the demand for the associated SW packages. The Open BNG is also acting as PE for the enterprise VPN services, which explains the final PE SW package.
To account for those cases in which there is a Master Clock device upwards from the Open BNG, this node will have to support equivalent synchronization requirements as those in the MBH. Finally, compliance with the SDN principles provided later on in this specification will also be mandatory.
The required number of subscribers for this flavour is 32k (dual stack). A single box is preferred, but in case of scalability cannot be achieved, multiple parallel devices or a fabric-based solution are acceptable as well.

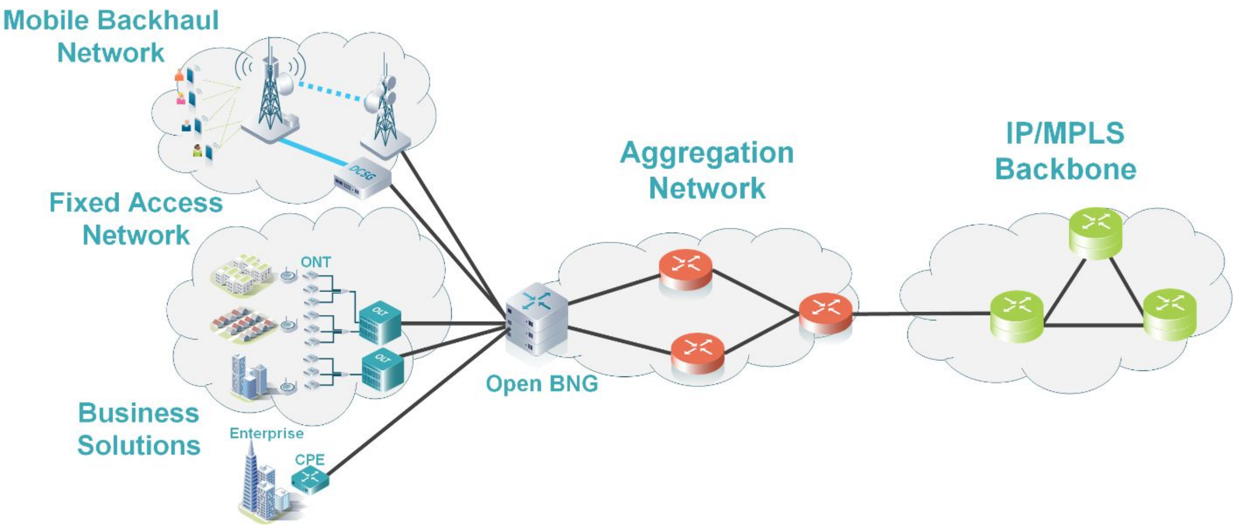

Figure 5. Open BNG with full functionality and maximum distribution

2.2.2. Full functionality and medium location

The main difference with the previous scenario is that, since the Open BNG is less distributed, it will have to aggregate more subscribers. The target is 128k (dual stack). TIP members are aware that these numbers may be quite a challenge considering current state-of-the-art, so multi-chassis or fabric solutions may be introduced if needed. If that is the case, though, it is highly recommended that actions are taken to simplify the complexity of working with multiple boxes (e.g. techniques that permit managing the multi-chassis solution as a single node). In addition, additional scalability might be needed for wholesale, legal intercept or existing and upcoming mobile network requirements.

Based on this scenario, also the encapsulation of subscriber sessions is subject to change. Since access nodes may no longer be directly connected to the Open BNG, other approaches may be possible (e.g. PW encapsulation of these sessions).

There are no changes with regards to SW packages or support of SDN in this flavour: all are required.

2.2.3. Service Only BNG Medium/Reduced Distribution

In this deployment case, the BNG is less distributed (e.g. central locations), but also does not form an integral part of the aggregation network, instead running as a service node that is connected to the aggregation network.

This 'Service Node' approach puts much more emphasis on terminating subscribers that are encapsulated within a tunnel of some form (e.g. an Ethernet pseudowire). Routing and PE functionality are required, but there is the option of deploying separate BNG for enterprise (with routing and Enterprise PE functions) and consumer (without Enterprise PE function).

This type of BNG will have to aggregate between 32k and 128k subscribers

As the 'Service BNG' is not in line with the aggregation network, it is not critical to support time/phase synchronization for mobile backhaul, as this can instead be passed through the aggregation network.

This case is also a candidate for a pure software BNG[2] that is deployed on a general purpose compute.

[2] TR-345 outlines how a virtualised (BNG) can be instantiated in software as a virtualized network function (VNF). https://www.broadband-forum.org/download/TR-345.pdf

3. *Platform Architecture & Hardware Specifications*

The Open BNG node consists of commercial off-the-shelf hardware (bare metal switch) and open software (Network Operating System – NOS) and interfaces (solutions that use software forwarding rather than being based upon bare metal switches may also be appropriate in the 'Service BNG' case. In that sense, disaggregation in this specification stands for separation of hardware and software.
Further vertical separation of the software stack for better programmability and SDN support is highly encouraged. There are different aspects and further detailed in section 7. Additionally, this specification does not define a strict position on how or where the software (or part of it) has to run (e.g. in the box, in an external VM). Again, this is considered an implementation decision to be taken by vendors and network operators..
The main modules/components of the platform are depicted in Figure 6.

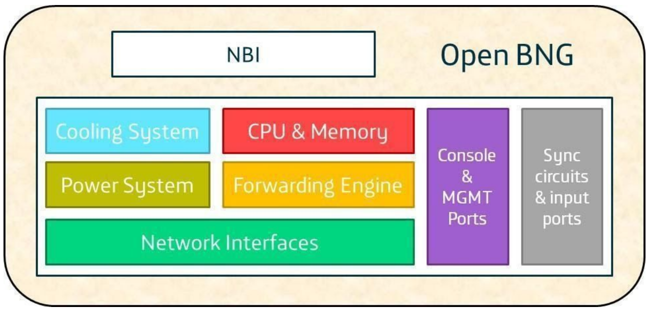

Figure 6. Open BNG high-level architectural components

The hardware system must not impose any restriction that limits the software that can run on it. In other words, the system must allow operators to install any operating system, even if its implementation comes from a third party. To ensure compatibility, it is highly recommended that Network Operating Systems for this platform are provided in the form of binary installers compatible with the Open Network Install Environment (ONIE) specification, as defined by the Open Compute Project (OCP). Equivalently, the Open BNG box will be equipped with ONIE.
If the platform provides the capability to verify the signature (via a particular certificate or a cryptographic key) of the software, it must be possible to disable such verification at any time, through software or firmware configuration, without the need for any specific or additional license. Moreover the hardware must be accompanied with a comprehensive toolkit of maintenance, configuration, diagnostic and repair information to facilitate modifications.

3.1. Hardware Solution form factor and environmental conditions

Open BNG hardware platform will be deployed exclusively in Central Offices. Installation will be done in standard 19" racks, typically 600mm or 800mm wide, and 800mm or 1000mm deep. The target form factor for the Open BNG hardware will be 2U. For higher scalabilities, the option of platform stacking could be considered.
The equipment shall conform to, or exceed, ETSI standard ETS 300 019-1-3 requirements (Class 3.1), in particular with regards to temperature (from -5°C to +45°C).
The equipment shall also conform to all applicable standards regarding mechanical, electrical and safety conditions, and must comply with all directives and certifications required in the countries where it will be deployed. In particular, solutions that improve the energy efficiency of the node will be extremely valued.

3.2. Hardware Platform CPU and ASIC

The CPU of the Open BNG hardware platform shall be based on an x86 architecture

(64bits).
The ASIC shall support line-rate forwarding across all ports without any limitations. The forwarding capacity of the platform shall be able to support all interfaces at full rate with no limitations. The ASIC shall support all the features and packet types defined in the software sections (or as a less preferred, not prevent the implementation of such features purely in software). Extensively, there shall be no limitations in the hardware to support any of the defined SW packages.
Programmability features of the ASIC will be valued. Redundancy of these components will also be valued.

3.3. Hardware Platform power supply and cooling requirements

The Open BNG hardware platform will include redundant power supplies and cooling components (fans); and replaceable filters to capture airborne dust and particulate matter from air intake and exhaust portals. It must be possible to substitute any of these with the device still in operation (hot swapping).
The equipment must support both AC and DC power supplies. In the event of a single power feed failure, the system shall be capable of maintaining full-service operation, without traffic loss. Associated alarms shall be generated for the duration of the fault.
The equipment must also include variable speed blowers, to adapt the cooling capacity to the exact demand. In general, the airflow in the Open BNG will be from front to back. Alternative mechanisms may be discussed. Again, associated alarms shall be generated in case of any fault in the cooling system.

3.4. Hardware Platform management

The Open BNG hardware must include, as a minimum, one console and one management port (both RJ45) and one USB port, for local configuration and debugging. Nonetheless, it must be possible to remotely disable the console and/or the management port and/or the USB port. It is therefore mandatory that the platform also supports in-band management.
The hardware platform must include status indicators, including per port LEDs.

3.5. Hardware Platform Synchronization requirements

As stated in section 2, the Open BNG hardware shall be able to propagate synchronization signals to other network elements directly or indirectly connected to it. Furthermore, it may also need to provide frequency, time and phase synchronization to 2G/3G/4G/5G base stations directly connected.
The following are the mandatory synchronization requirements:

- Support of IEEE 1588 profile as defined in ITU-T G.8275.1 (full-timing support; Telecom - Boundary Clock) and Sync-E for holdover purposes and Grandmaster redundant sources support. The requirement corresponds to the support of IEEE1588v2 – Precision Time Protocol – profile for telecoms (multicast mode preferred) and includes Sync-E in Ethernet interfaces as per ITU-T G.8261 (section 9.2.1), G.8262 and G.8264.
- Network quality model (microsecond precision) according to ITU-T G.8271.1.
- Node performance (noise generation, tolerance, transfer and holdover) according to ITU-T G.8273.2 (sections 7.1/7.2/7.3/7.4).
- Node performance (upon wander, failure and holdover) according to ITU-T G.8273.2 (sections 7.2/7.3/annex). Clock type class B (minimum) or C (desirable), as defined in ITU-T G.8273.2.
- Support of IEEE 1588 profile as defined in ITU-T G.8275.2 (partial-timing support).

The electronics used to build the synchronization regeneration capabilities are to be selected by the platform manufacturer (ensuring full performance compliance with standards), but detailed information should be shared within TIP. Hardware SKUs should include at least 2 time synchronization solutions to allow software (SW) providers to implement the above mentioned features using their preferred SW stack.

Finally, the hardware platform must include, at least, one GPS signal input interface (which may be based on SFP) and one 1PPS interface for external synchronization. These will be used in scenarios where the synchronization signals cannot be received from the network, and there are base stations directly connected. It may also be possible to select one signal among several (by means of a BMCA, desirably hardware implemented), in case there are multiple sources.

Note: These requirements are not applicable in scenarios where the BNG is a service only device (not being part of the aggregation network).

3.6. Hardware SKU network interfaces and forwarding capacity

The required capacity and the interfaces layout in the hardware platform heavily depend on the specific Open BNG flavour that is being considered. General requirements will be included first, with specifics per flavour defined later.

The solution must be able to support electrical and optical interfaces as per IEEE 802.3, and they shall be configurable to work either as UNI or as NNI. Pluggable optics will be preferred to fixed format connectors: they shall be able to operate at the same temperature ranges as the node, and it will be possible to configure them at different speeds (e.g. 1G/10G with SFP+) without reboot. There will be no limitations on the type of connectors that are used (SR, LR, etc.; LAN/WAN PHY), and the platform must be fully interoperable with third party optics. Additionally, the system shall be compatible with third-party coloured WDM pluggable optics (tuneable & fixed).

All physical network interfaces in the proposed platform must support multiple services simultaneously, independently on whether they are: multiple PWs each with its own set of VLANs; multiple VLANs associated to VRFs, VPLS, E-VPN, etc.; multiple native VLANs or VLANs associated with core interfaces (those configured with IS-IS, OSPF, LDP, etc.); multiple VLANs associated with BNG services (i.e., transporting PPPoE/IPoE sessions). Similarly, ethernet LAG interfaces must support the same type of services as single physical interfaces.

Accommodation of all the Open BNG flavours in a single SKU or chassis type will be extremely valued. If not possible, efforts will be done to limit the number of different SKUs. According to the flavours defined in section 2.2, the following standard configurations (SC) are envisioned:

- SC-1: 96x(1G/10G) + 16x25G +4x100G
- SC-2: 80x(1G/10G) + 16x25G +16x100G
- SC-3: Fabric with Leaf and Spine Switch
    - Leaf Switch: 32x(1G/10G/25G) + 8x100G
    - Spine Switch: 4x(10G/25G) + 20x100G + 4x400G

The 1G/10G Ethernet port cages should support both short-reach (SR) or long-reach (LR) pluggable optical transceivers, and direct attach copper (DAC) transceiver-cable assemblies.

It is insufficient to propose breakout cables, exclusively, for 1G/10G Ethernet ports. This is because for SC-1 and SC-2: 80% of the 1G/10G Ethernet ports will be short-reach (SR) optical connections; whilst the remaining 20% of the 1G/10G Ethernet ports will be longer distance (LR, ZR, ER) optical connections. In cases where breakout cables are proposed, they should include short-, medium-, and long-reach optical connection options.

The fabric configuration will use a mixture of interfaces for short-reach (SR) and longer distances (LR, ZR, ER as well as support for coherent interfaces) for the

connections between the leaf and spine switch (100G) as well as connections to the network core (400G) and access nodes (1G/10G/25G). It is expected that a number of these ports will be used (1G/10G/25G) for local connections to servers hosting control and management plane as well as other services like wholesale interfaces or legal intercept connections. It should be noted that 25G is an optional feature but expected to be relevant in the near term and valued to be available in solutions.
In the service node deployment model, there is not a requirement for high aggregation/fan-in and the BNG may be deployed as a software function on a general-purpose server or based upon a dedicated open hardware. In the case of general purpose server, it is anticipated that any number of high speed NIC Cards can be used up to the capacity of the server, and that Hardware Acceleration from the NIC cards is a valid option to improve performance (although acceleration must be made available to the BNG in a standard abstracted form). In the case of a dedicated open hardware BNG service node, as there is no longer the need for aggregation, we also define the following additional standard configuration:

- SC-4: 12x100G

With regards to the forwarding capacity, the main requirement is to present non-blocking forwarding architectures, as already commented in section 3.2. This leads to the following requirements:

- SC-1: capacity higher or equal than 1.84Tbps
- SC-2: capacity higher or equal than 2.8 Tbps
- SC3:
    - Leaf Switch: capacity higher or equal than 2.4 Tbps
    - Spine Switch: capacity higher or equal than 2.4 Tbps
- SC-4: capacity higher or equal than 1.2 Tbps

Note: Oversubscription of the spine is expected, as leafs will be run with less capacity for redundancy reasons

Finally, it is worth mentioning that the previous definition of standard configurations does not prevent the inclusion of additional variants with a higher number of ports, or higher capacities, to account for the future. As already said in section 3.1, the option of platform stacking could be considered to achieve this.

4. BNG Software package (BSW): subscriber management capabilities

The requirements for the BNG Software package are primarily based in the following Technical Report from the Broadband Forum: TR-101 Issue 2 (2011) – Migration to Ethernet-Based Broadband Aggregation. Any Open BNG platform proposal must provide a very high degree of compliance with such Technical Report and its list of requirements. Several of these will be explicitly mentioned in this specification for clarification; however, that does not exclude the needed compliance with the rest.
Together with TR-101, these three additional Technical Reports must be supported:

- TR-059 (2003): DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services
- TR-092 (2004): Broadband Remote Access Server (BRAS) Requirements Document
- TR-146 (2013): Subscriber Sessions

As derived from TR-101, the main access technology will be Ethernet. In particular, the Open BNG must support IEEE 802.1Q frames, and VLAN stacking (QinQ) mechanisms as per IEEE 802.1ad. It must be noted that QinQ must be supported even if the ethertype of the outer VLAN is 0x8100 instead of the standard 0x88A8.
It shall be possible to assign automatically one (or more) IP address to all subscribers in the Open BNG. These addresses may be: i) fixed and configured in the Open BNG, ii) dynamic from a pool configured in the Open BNG, iii) fixed defined by AAA server, iv) dynamic from a pool referenced by AAA server, or v) dynamic from a pool provided by AAA server.
Proposals for the support of legacy ATM access technologies will also be valued. These proposals can leverage on external equipment, like intermediate devices, specific pluggable optics, etc.

4.1. Encapsulation methods and protocols

4.1.1. PPPoE encapsulation

The platform must support PPP protocol as per RFC1661 (and later updates), permitting both PAP or CHAP authentication. Per-session keep-alive messages must be supported, with parameters like the interval time or the timeout being configurable.
Incoming PPP sessions will be encapsulated over Ethernet, either with one or with two VLAN headers.

4.1.2. IPoE plus DHCP service

The platform must support DHCP in combination with IPoE subscribers, as described in TR-101. It must also support the DHCP Relay Agent functionality, with the possibility to configure redundant DHCP Servers, and supporting Option 82 (DHCP Relay Agent Information Option).
Mechanisms to control the DHCP load (requests rate, filters, etc.) shall be implemented.

4.1.3. L2TP

Open BNG will also participate in wholesale scenarios, typically solved using L2TP. Therefore, the node must comply with the L2TP Access Concentrator (LAC) and L2TP Network Server (LNS) functionalities as per RFC2661.
According to this, it shall be possible to i) configure PPPoE subscribers and encapsulate their sessions over L2TP, and ii) terminate PPPoE subscribers that are encapsulated over L2TP. In the first scenario, regarding the decision on whether to terminate the session or encapsulate it, it must be possible either to configure it

statically, or receive it dynamically via RADIUS, on a per session basis. In any case, it shall be possible to activate RADIUS accounting for these L2TP LAC services.
L2TP tunnels in which Open BNG participates must be able to support multiple PPPoE sessions, and it must be possible to have redundant L2TP tunnels.

4.1.4. VPN

It must be possible to map dynamically PPPoE or IPoE sessions on top of a VPN, via RADIUS control.

4.1.5. PPPoE/IPoE over PW

The support of incoming PPPoE/IPoE sessions over MPLS PW will be valued. This requirement may be of special importance in environments like that in section 2.2.2, in which MPLS end-to-end strategies can be applied.

4.2. IPv6

The Open BNG must support both dual stack (IPv4 and IPv6) access and IPv6 only access over PPP. Some requirements for the BRAS functionality concerning IPv6 are:

- Wholesale model using L2TP must also be supported for IPv6 and IPv4/IPv6 sessions
- Prefix delegation (delegating router role) must be supported
- Support of DHCPv6 and DHCPv6 Relay is required
- The platform shall be able to interact with AAA servers including all IPv6-related attributes, in particular those dealing with prefix delegation
- In dual-stack environments, it shall be possible to steer IPv4 traffic from subscribers to external CG-NAT platforms, using for example VPN mechanisms.

4.3. Authentication, Authorization and Accounting (AAA)

Primary protocol for AAA in Open BNG will be RADIUS. Alternatively, support of DIAMETER will be valued. For those reasons, support of the most relevant RADIUS and DIAMETER RFCs will be requested. Some details for clarification are provided below.
First, it shall be possible to configure separate AAA servers for the different AAA processes, and it shall be possible to have redundancy for any of them. Similarly, to DHCP, it shall be also possible to configure control mechanisms in the AAA traffic, like specific rates, filters, etc.
In general, all subscriber profiles will reside in the AAA server. During the establishment of a subscriber session, the Open BNG will download[3] the specific profile(s) and will apply them to the session. The Open BNG will also receive the public IP address information from the AAA server. Two of the implications of this statement are:

- In general, there shall not be any specific manual configuration on a per-subscriber basis in the Open BNG.
- It shall be possible to apply more than one profile per subscriber. These will be stacked based on associated priorities so that it is clear how they apply at any moment. An example of this is provided in Figure 8. Profiles shall be executed consecutively, unless one profile determines otherwise (e.g. a captive portal for defaulting customers).

Apart from this, using Change of Authorization (CoA) messages, it shall be possible to modify the profiles applied to a subscriber without the need to restart the customer session.

[3] Cache mechanisms to maintain in the Open BNG a set of commonly used profiles will be valued.

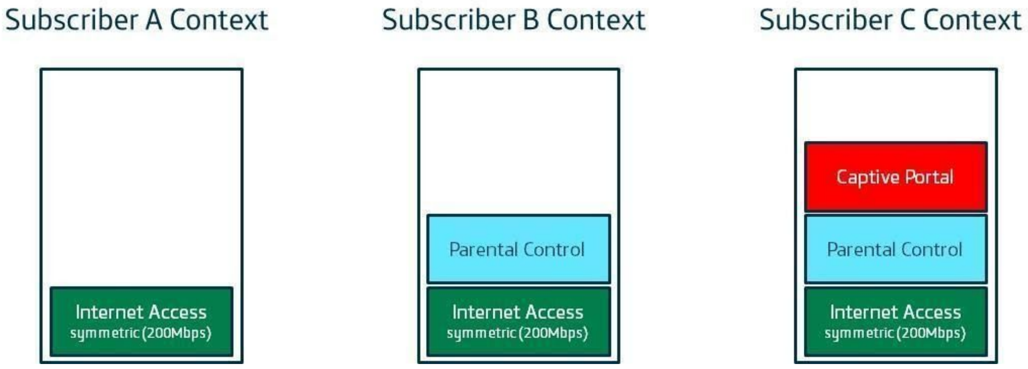

Figure 8. Simple example of stacked subscriber profiles

It shall be possible to base the authentication mechanisms in four methods, each univocally identifying a subscriber:

1. By the attributes NAS-Port (or NAS-Port-ID) plus NAS-IP-Address[4]
2. By username and password
3. By domain and/or agent_remote_id

In general, the first approach will be preferred, which implies that these attributes must appear in all authentication and accounting packets: Auth-req, Acc-req, Acc-Start, Acc-Update and Acc-Stop. NAS-Port will typically code the physical port ID, plus the VLAN tag values in the QinQ.
It is important to mention that the platform will generate as many accounting sessions as profiles/services are applied to a subscriber. In the case of Subscriber C of Figure 8, for example, there will be one accounting session for the subscriber session (as a whole), and then three additional accounting sessions, one for each of the applied profiles (service accounting). Each service accounting session will only take into account the traffic to which the associated profile has been applied.
Accounting ON/OFF messages shall also be implemented. Together with a clear record of the allocation/withdrawal of IPv4, IPv6 and IPv4/v6 addresses.
With regards to public IP address pools, Open BNG must be able to operate both with locally defined pools and with pools received from the AAA server.

4.4. Quality of Service

In multi-service nodes like the flavours of Open BNG that are being proposed, Quality of Service applies to all of the described services. However, requirements associated with the Internet Access are possibly the most complex, due to the strict hierarchical scheme that it is demanded. This is why the QoS section is included as part of the BRAS SW package.
It must be noted that Open BNG will not work with a single QoS header. Client interfaces are normally QinQ, so the platform will have to cope with Ethernet headers. On the contrary, network interfaces are normally IP/MPLS, so the node shall be able to interact at both levels.
The number of required traffic classes is 8 (eight). The platform shall be able to discriminate traffic and manipulate it based on the definition of those eight classes. The platform must also be able to match the internally generated traffic to any of these traffic classes, based on configuration.
In this section, we will start from the simpler features (classification, marking,

[4] NAS-IP-Address must be a configurable parameter.

scheduling, etc.), and later on we will describe the H-QoS scheme required for the Open BNG.

4.4.1. Classification

In the ingress direction, the platform will perform two main processes[5]: classification and assignment to ingress queues. The platform shall be able to classify and assign incoming traffic based on at least one, or the combination of, the following parameters: physical/logical port, VLAN IDs (including VLAN stacking), Ethernet CoS value, MAC source/destination address, IPv4/IPv6 source/destination address, TCP/UDP port values, IP DSCP/Precedence value and MPLS EXP value.
The above classification must be possible not only at the physical port level, but also at the sub-interface (or VLAN, or logical port, etc.) level.
In addition, the platform shall be able to re-classify packets in the egress direction, before assigning an egress-queue, based on the same parameters as above (except for physical and logical port).

4.4.2. Marking

For each of the traffic classes defined in the node, the platform shall be able to re-write the priority values at the following headers: Ethernet CoS (at any VLAN tag), IP DSCP/Precedence and MPLS EXP. This re-writing will be based on specific rules defined by configuration.
The platform shall also implement inheritance mechanisms, by which a certain header maps the QoS value of another header. The minimum requirement is that the QoS value of the IP Precedence header can be mapped both towards the Ethernet and MPLS QoS headers.

4.4.3. Scheduling

Similarly, to the classification process, the scheduling must also be possible at different levels, in particular, to individual logical connections (not only per port).
It is required that schedulers in the Open BNG can support up to 8 (eight) queues, with at least 2 (two)[6] different behaviours (priorities); it shall be possible for the highest priority to be configured as strict priority. The remaining queues shall support at least one weighted mechanism that ensures committed rates per queue (e.g. WFQ, WRR, etc.). Such committed rates may be defined either by numeric values, or by weights (percentages) of the total scheduler throughput.
An improved approach with regards to the number of queues and schedulers, which will be highly valued, is that each scheduler optimizes the usage of queues. In other words, each scheduler will only consume the number of queues that it requires (number that will be determined automatically and dynamically), and furthermore, if that number is not two or higher, the scheduler itself may not be needed[7].
With regards to exceeding traffic, it shall be possible to assign the remaining throughput to those queues that can still consume capacity beyond their committed rates, if available. It shall be possible to control the assignment of exceeding capacity for each traffic class. It shall also be possible to limit the maximum throughput that is assigned to the strict priority queue.
For dropping criteria, it shall be possible to define one of these policies on a per queue bases: Tail Drop (Weighted Tail Drop valued) or Weighted RED (with a minimum of two drop priorities).
As a result of the above description, each queue in the scheduler will be defined by the following parameters: priority, CIR, PIR, exceeding traffic policy (typically EIR) and drop policy.

[5] Potential ingress policing will be described in a later section.
[6] More than two priorities per scheduler will be valued.
[7] For example, if shaping is implemented at the scheduler itself, then the scheduler will be needed.

4.4.4. Traffic rate limit

In general, shaping will be preferred over policing for limiting subscribers' traffic rates, both in the uplink and in the downlink.
Shaping may be applied either as part of the schedulers or as a separate function. Independently of the implementation method, it must be possible to apply shaping to individual logical connections, not only on the physical port.
It shall be possible to define CIR and PIR (or EIR, where PIR = CIR + EIR). The possibility to define CBS and EBS will be valued.
Two items must be noted:

- It shall be possible to apply shaping to those subscribers that will be later on tunnelled via L2TP.

4.5. Per subscriber shaping values are part of the AAA policies.

It shall be possible to modify these values (e.g. speed upgrade operation) without the need to restart the session.
Independently of the support of shaping, policing must also be supported, both at the ingress and at the egress. Even if shaping is preferred for conditioning of subscribers traffic, many control mechanisms can still be implemented with policers. For that, it shall be possible to apply it to individual logical connections.
The following policer types will be required: i) single rate, two colours, and ii) two rate, three colours. Support of single rate, three colour policers will be valued. Actions applicable to policers will include, at least: forward, drop, re-classify, re-write and increment an associated counter.

4.5.1. Hierarchical QoS

The platform must permit the application (in cascade) of QoS policies at different hierarchical levels. Capabilities in each of these levels must be equivalent to those included in the sections above, and it shall be possible to configure the specifics of each level separately.
A minimum of three H-QoS levels is required, although four are recommended.

4.6. Multicast

The BRAS SW package must be able to handle source specific multicast (SSM) per subscriber and provide replication mechanisms according to the used encapsulation and protocols (PPPoE/IPoE and IPv4/v6). A special attention shall be paid to the use of IGMP controlled multicast with PPPoE encapsulation. In addition to SSM, the support for commonly used any-source multicast with IGMPv2 must be given in order to handle existing and widely deployed CPE implementations.

While the BRAS and the Router SW package (see section 5) must interwork, it must be based solely on source specific multicast and other variants must be adopted accordingly.

A special case is the deployment of a fabric and the distribution of the BRAS and Router SW package. Here an additional network segment must support multicast and the preferred option is based on multicast support in L3VPN (RFC 6514).

Where multicast packet replication is performed external to the OpenBNG – for example downstream within a separate head-end, or layer 2 switch - a dynamic policy control mechanism is required to shape (or police) unicast and multicast traffic flows that are distinguished by CVLAN. The requirement serves to maintain the contracted bandwidth ceiling that is applied to each end-user: the aggregated sum of the unicast and multicast traffic flows is dynamically clamped. In operation the policy control mechanism would decrement the unicast bandwidth allocated to the end-user to exactly match and rebalance each corresponding multicast bandwidth increment. The OpenBNG control plane on receipt of an IGMP join/leave datagram from an end-user (multicast receiver) would invoke the policy control mechanism – yet maintain a sufficient unicast bandwidth floor to maintain voice services.

Details about the multicast capabilities for routing are given in section 5.

4.7. Lawful Interception

The Lawful Interception mechanism supported by the Open BNG must be compliant with ETSI specifications like ETSI TR 101 943 (2004) and ETSI TR 102 528 (2006), in particular with regards to Reference Points INI1, INI2 and INI3 (X1, X2 and X3 in 3GPPP standards). The proposed platform must implement the Internal Intercept Function (IIF), while access and control of this function will only be possible from an Administration Function (ADMF). This access will be done ensuring that:

1. There is a single dedicated user profile, duly authenticated, with features not accessible to any other user profile (not even administrator). LI related configurations will only be visible for this profile.
2. Connectivity is limited to a single IP and port, with only one concurrent user, which is the ADMF.

Requirements for the IIF at the Open BNG include:

- Same availability as any other node feature.
- It shall be synchronized via NTP.
- Any audit data will be generated and sent to an external safe-keeping node. This communication will be encrypted. The Open BNG will not store any of the audit data.
- Audit data may be sent in-band or out-of-band. In particular, it shall be possible to send it over a VPN.
- There will be no actions by the IIF that permit the target to learn that it is being intercepted.
- It shall be possible to apply LI to PPPoE and IPoE subscribers, even those tunneled using L2TP.
- Identification of the target subscriber will be done based on any of the parameters that may govern the subscriber authentication, as defined in section 4.3
- Reboot of the Open BNG shall not affect to the LI operation and configuration.

5. Router Software package (RSW): routing capabilities for Open BNG

The following list of requirements is included to ensure that the Open BNG platform will be able to work in any type of network scenario potentially deployed in the production networks of the participating operators. Some of these requirements may be considered as pre-requisites for the support of VPN services (i.e., pre-requisites for the Open BNG operating as a PE). Examples of these may be MP-BGP or T-LDP. However, they will be included here to avoid too much dispersion in the location of requirements.

5.1. IP/MPLS routing

Firstly, it is important to note that the Open BNG must support dual stack (IPv4 and IPv6), or IPv6 only, in all its network interfaces[8]. In fact, for all the protocols mentioned below, it will be understood that they must support the applicable IPv6 extensions (when applicable).
The high-level list of routing protocols is the following:

- Static routing.
- IS-IS, including extensions for Traffic Engineering.
- OSPFv2/v3, including extensions for Traffic Engineering.
- BGP-4, including multiprotocol extensions, capabilities advertisement (RFC5492), communities (RFC1997), BGP-LU (RFC3107), deterministic-med (RFC 4721), graceful restart / non-stop Forwarding (RFC 4724), extensions for 4-byte AS number (RFC4893), confederations (RFC3065), route reflection (RFC4456), error-handling (RFC 7606), peer tracking (RFC 7854) and prefix-Independent Convergence (PIC)[9].

With regards to these protocols, the proposed platform must support distributing routes between any of them, based on defined policies, and also distributing local and static routes. It must also support modifying the priority (administrative distance) of the different routing protocols when populating the active FIB. Any mechanism for mutual authentication must support at least MD5 authentication.
MPLS must be supported as well, with at least four MPLS labels in the label stack. Signalling of MPLS labels must be possible both using:

- LDP, including T-LDP and IGP and LDP synchronization.
- RSVP-TE, with the capability of path computation based on CSPF and supporting FRR mechanisms.

Finally, other features that are required include:

- ECMP
- Entropy-label-based load balancing[10]
- IP LFA FRR mechanism
- VRRP
- Multicast capabilities, both for IPv4 and IPv6, for L3VPNs outlined in RFC6514 as the preferred option
- Policy Based Routing
- PWE3
- GRE as transport protocol, instead of MPLS

[8] Requirements for interfaces towards the CE, when the Open BNG acts as a PE node, are defined in section 6.
[9] Bashandy, A., Ed., Filsfils, C., and P. Mohapatra, "BGP Prefix Independent Convergence", Work in Progress, draft-ietf-rtgwg-bgp-pic-06, August 2020.
[10] The Use of Entropy Labels in MPLS Forwarding - https://tools.ietf.org/html/rfc6790

Segment Routing (and other mechanisms like TI-LFA) shall be supported. If segment routing is not supported, then proposals complying with this specification must demonstrate a clear roadmap to implement Segment Routing. Policy-based routing will be required to eliminate forwarding to non-adjacent next-hops i.e. a mechanism to prevent subscriber-to-subscriber traffic flow. On the other hand, the Open BNG must support PCEP and BGP-LS, to account for the deployment of PCE nodes together with RSVP-TE.
Finally, and despite the strong focus on layer-3 functionalities, support of layer-2 Ethernet switching is also required in the Open BNG.

5.2. Link aggregation

The proposed platform must support aggregating several physical ports into a single logical interface, based on LACP protocol as defined in IEEE 802.1ax. Schemes for balancing the load among the different ports that achieve a balanced share will be preferred. It shall also be possible to determine by configuration the number of interfaces within a LAG that determine the failure status of the whole logical interface (e.g. 1 out of 3, 2 out of 2…).
The support of multi-chassis LAG, meaning that a single device can connect using LACP, in a dual-home active/standby scheme to two ports in two different Open BNGs, will be valued.

5.3. BFD

The proposed platform must support Bi-directional Forwarding Detection (BFD) for failure detection as described in RFC5880. In particular, it is required to support BFD in Ethernet interfaces, and in LAG interfaces as described in RFC7130. It is also required to support BFD in the signalling protocols listed in section 5.1, and in MPLS LSPs.

5.4. Management

As stated in section 3.4, both out-of-band and in-band management must be supported. For this, mechanisms based on SSH, CLI and Netconf will be preferred to those based on web browser.
The CLI will permit the configuration of different access profiles, at different levels and with different permissions (e.g. operator, monitoring…). It will also permit the login of multiple concurrent users.
The following protocols shall be supported:

- SNMPv2c/v3, for management and configuration purposes
- Netconf for management and configuration purposes
- NTP, for time synchronization with the rest of the network
- SCP or SFTP, for file transfer (e.g. configuration files or software upgrade versions)

5.5. Monitoring

SNMP may also be used with monitoring purposes. In that sense, the platform shall support the required MIBs for monitoring of the main activated functionalities, and it must be possible to send monitoring traps to multiple destinations.
The support of SYSLOG is also a requirement for logging redirection: all relevant system events must generate SYSLOG messages, which may be sent to multiple SYSLOG servers.
Netflow/IPFIX is also a requirement for the Open BNG. It must be possible to define by CLI the sampling criteria (one out of “n” packets sent), and the vendor must ensure the values of “n” which will not affect the performance. It shall be possible to activate Netflow/IPFIX on a per interface basis, and to define several collectors for the

exporter.
Finally, for monitoring of quality of service, support of Y.1731, OWAMP and TWAMP mechanisms is requested.

5.6. Security

The following is a non-strict list of requirements having to do with security of the platform. Participating operators strongly encourage vendors to keep these capabilities up to date, incorporating new mechanisms as soon as they become available.
The proposed platform must be capable of implementing segment, packet and frame PDU filters in any physical or logical interface, and in the incoming and/or outgoing directions simultaneously with fine-grained – per subscriber - granularity. Filters will be able to take into account any combination of the following arguments: source/destination IP address, source/destination port, protocol, source/destination MAC address, VLAN, TCP flags, fragmentation flags, ICMP type and packet size. There must be counters available for each rule in the filter, increasing with each matching packet, and available for consultation via SNMP.
With regards to management operations, the platform must support the following:

- TACACS+ for authentication and authorization of the CLI features.
- Logging of all the commands executed by all the operators, for audit purposes.
- SSH sessions with 3DES encryption – and not preclude options to use more advanced encryption methods e.g. AES-128 or AES-256.
- Restriction of the management access only to a defined subset of IP addresses.

In the data plane, the platform must support Unicast Reverse Path Forwarding (URPF), as defined in RFC3704. Further, for PPP Termination and Aggregation (PTA) sessions, uRPF must be applied per subscriber. It must also be capable of limiting the maximum number of MAC addresses and limiting the BUM (broadcast, unknown unicast and multicast) traffic per physical or logical interface.
Protection mechanisms against Denial of Service (DoS) attacks is also required, at least for the following list: Tear Drop, Ping of Death, Smurf, Fraggel, UDP Flood and SYN-ACK.
Mechanisms to moderate control protocol (e.g. LCP) traffic load per subscriber (requests rate, filters, etc.) shall be implemented. Finally, the proposed architecture must ensure the maximum possible isolation between the control and management planes, and the data plane.

6. *Provider Edge Software package (PESW): Open BNG as a PE for enterprise services*

The following is a summary of the requirements for the Open BNG acting as a PE of L2 and L3 VPN services. As commented before, this summary relies on the support of certain transport protocols which are described in section 5.
A Multi-Service BNG is defined in the Broadband Forum Technical Report: TR-178 Issue 2 (2017) – Multi-service Broadband Network Architecture and Nodal Requirements. Any Open BNG platform proposal must provide a very high degree of compliance with such Technical Report and its list of requirements.

6.1. L2VPN

In the provision of L2 services for enterprise customers, the Open BNG shall be MEF 2.0/3.0 compliant, at least with regards to E-LAN and E-LINE services. E-TREE services will be valued. It will be very beneficial that the proposed platform is certified accordingly.
In particular, the platform must support the implementation of E-LINE services based on VLL technology, and E-LAN services based on VPLS. Support of H-VPLS topologies will also be valued.
Requirements below have already been tackled in previous sections, but they are included here as well for specific clarification of their usage in a L2VPN environment.

- The proposed platform must be able to forward traffic into the L2VPN at least based on the following encapsulations: no tagging, VLAN tagging (IEEE 802.1q), VLAN double tagging (IEEE 802.1ad) and MACinMAC (IEEE 802.1ah). For double tagging, the TPID value must also admit 0x8100 in the outer VLAN, together with the standard 0x88A8.
- The proposed platform will permit the establishment of CIR, EIR, CBS and EBS parameters over the L2VPN service and/or each specific L2VPN client connection.
- The proposed platform will permit configuring specific Layer-2 based filters per each L2VPN client connection. Other standard security measures shall be supported as well.
- For E-LAN (and E-TREE, in case they are supported) services, it must be possible to configure in the proposed platform a different rate limit per each type of BUM traffic (one for broadcast, one for unknown unicast and one for multicast). It must also be possible to configure limits in the number of learnt MAC addresses per service and client interface, and the lifetime for these learnt MAC addresses.

Finally, redundancy mechanisms, loop detection mechanisms and OAM functionalities shall be included in the proposed platform. IGMP snooping must also be supported.

6.2. L3VPN

Full-mesh and Hub & Spoke models of BGP/MPLS L3VPN must be supported for IPv4 and IPv6. In particular, the proposed platform must support the establishment of a L3VPN over multi-AS backbones.
The non-exclusive set of protocols outlined in Section 5.1 (RIP, IS-IS, OSPF and BGP-4),SSH including static routing, must be supported in the CE/PE link. BFD associated with these protocols and DHCP Relay must also be included.
With regards to the rest of the features, like QoS, multicast, filtering, redundancy, etc., the same requirements as listed in sections 4 and 5 for IP environments must be supported, both in the client and network interfaces. Additionally, control plane control mechanisms (e.g. BGP dampening) shall be available.

6.3. E-VPN

The proposed platform must support E-VPN services. PBB E-VPN will be valued. All the requirements specified above for MPLS based VPN services shall apply as well to E-VPN.

E-VPN must be supported with single and dual homing (with Hot-Standby and Load Balance in the latter case).

6.4. IPSec

As PE for enterprise services (and for other applications like the secure encapsulation of DIAMETER), it will be required that the Open BNG supports IPSec (applicable, for example, to the PE/CE interface). RFC4835 (updated by RFC7321 and later by RFC8221) defines the encryption and authentication algorithms that should be supported.

The implementation solution for IPSec must try to minimize the impact on the CPU usage and the introduced latency. HW-supported implementations will be permitted, although implementations that result in external specific components should be avoided (e.g. specific network cards).

IPSec must be compatible both with IPv4 and IPv6.

7. *Open BNG programmability and SDN*

As explained in previous sections, we encourage more modern and advanced management mechanisms and protocols with improved programmability and SDN support. The intention is to have a (centralised) control entity (an SDN Controller) managing OpenBNG in a smart and automatic way.

The SDN controller could be centralised and connected to the Open BNG in an out- or in-band management connection. In addition, the SDN controller could be collocated to the Open BNG, especially in fabric designs. We appreciate deployments where the OpenBNG domain may be orchestrated in common with adjacent domains e.g. access/regional aggregation domain, IP/MPLS backbone domain etc. to provision and maintain an end-to-end service. Also, there are situations where optical line termination (OLT) elements are incorporated within a combined OpenOLT / OpenBNG domain. But it is emphasised that regulatory oversight or market conditions can mandate that the OpenBNG is managed separately and distinctly from other network elements.

The SDN controller shall manage and optimize the OpenBNG domain in order to perform SLA fulfilment and service provisioning (Figure 9). All the configuration and management of the OpenBNG shall be done using industry wide well accepted protocols like Netconf (RFC7803) or gNMI. Netconf connections could be encrypted using Transport Layer Security (TLS) [RFC5246] or Secure Shell (SSH) [RFC4251], with TLS being the preferred option.

Additional programmability support is not an explicit requirement of this document, leaving it as an implementation decision, but capabilities are highly appreciated for further separation of the vertical software stack. Two prominent solutions exist with the Control-User Plane Separation (CUPS) approach and the P4/gNMI approach (e.g. an equivalent to ONF STRATUM/TASSEN). These target different levels of abstraction and solutions may reasonably support both, with P4/gNMI used to abstract between forwarding and a local control/NOS and CUPS as an onward interface to a more central subscriber control plane. For those vendors implementing CUPS, it is recommended that the TR-459 standard is followed.

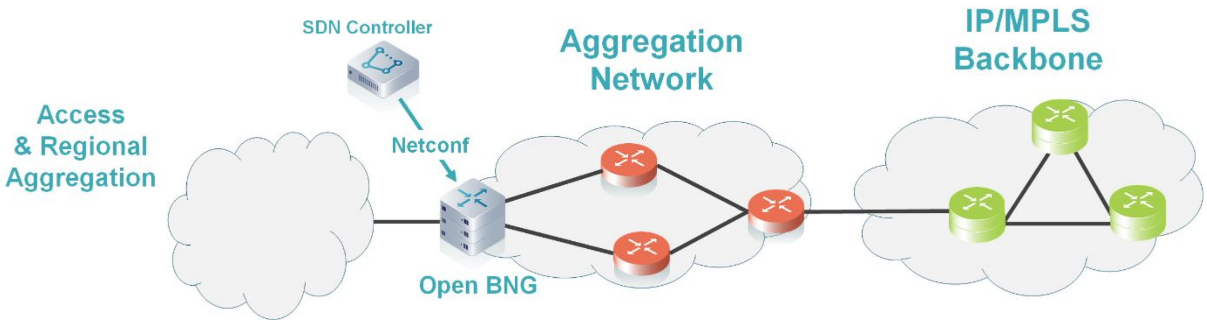

Figure 9. SDN controller and management

Standard data models shall be used as much as possible. The definition/selection of the target/needed models will be done in future releases of this specification but it can be anticipated that a generic requirement is that the node data structure be YANG compliant. There are three main organizations working on models that apply to OpenBNG: BBF, IETF and OpenConfig. Broadband Forum (BBF) is working on *YANG Modules for Broadband Network Gateways* (SD-460) that will cover the BSW aspects. OpenConfig and IETF are working mainly on RSW and PESW capabilities.

In the case of the YANG models not covered by OpenConfig today, operators will work to augment them to cope with requirements on this specification. Similarly, actions on pending OpenBNG requirements may be presented at BBF (for BSW) and IETF (for RSW and PESW).

7.1. Telemetry and SDN

OpenBNG shall support advanced monitoring and telemetry features, in particular:

- gRPC Network Management Interface (gNMI), gPB (Google Protocol Buffers) proto3 for encoding, and data exported (modelling) based on YANG models
- YANG push (RFC 8639, RFC 8640, RFC 8641, RFC 8650) with support for YANG QoS models[11]

Those will be used by the SDN controller (Figure 10) in order to monitor the status of the platform and the different services instantiated in the node.

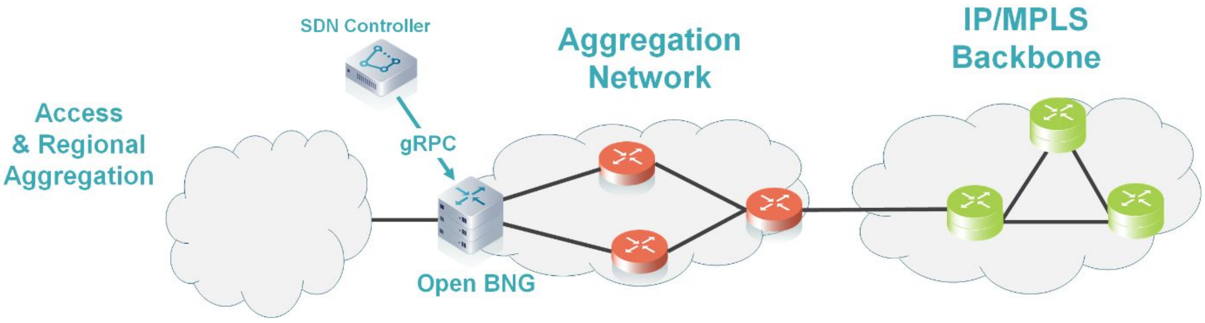

Figure 10. SDN controller and Management

7.2. Access security and Anti-Theft

In general, the solution must support the necessary security mechanisms to authenticate and encrypt communications between the network element and its management system or controller.
The Open BNG must offer the possibility of enabling this traffic only after it has been authenticated by the management platform/controller.
The system should also offer the possibility to enable anti-theft mechanisms that prevent the use of the equipment in any other environment than the one it was conceived for.

[11] Choudhary, A., Ed., Jethanandani, M., Strahle, N., Aries, E., and I. Chen, "YANG Model for QoS", Work in Progress, draft-ietf-rtgwg-qos-model-02", July 2020.

8. *Additional requirements*

To finish this specification, the following sections include additional requirements which do not fit any of the software packages defined above, but that are mandatory independently on the Open BNG flavour.

8.1. Port mirroring

The proposed platform must support mirroring capabilities, transparent (i.e., not impacting) to the rest of the Open BNG functionalities. In particular, specification of characteristics like the source granularity (LAG, port, VLAN, etc.), the source direction (RX, TX, both) and the destination (port, IP address, tunnel) is required. Mirrored traffic must be capable of redirection and appropriate encapsulation over physical or logical interfaces to nominated destinations.
It must be possible to activate multiple concurrent mirroring sessions, and to send the mirrored traffic to multiple destinations at the same time.
Port mirroring capabilities must coexist simultaneously, and will not interfere, with those for Lawful Interception.

8.2. Configuration and versions management

It shall be possible to modify the node configuration by means of configuration files accessible or directly copied in the node disk (via SFTP or SCP). Processes of file transfer, loading of a target configuration, and activation of the target configuration shall be three different processes. None of this shall require a reboot of the node, nor will affect its operational status.
It will be valued that more than two (active and target) configuration versions are accessible to the operator. This will be achieved by storing more than one target configuration in the node. It shall be possible to activate any target configuration from those stored.
Prior to the activation of a target configuration, the device will check its coherence, and will warn the operator of any incompatibility.
For configuration changes based on CLI:

- No configuration change will require a reboot of the system for its activation
- It will be valued that the introduction of the command, and its activation, are two different processes (e.g. using “save configuration” or “commit” commands).

On the other hand, notwithstanding automated procedures for Zero-Touch Provisioning as described in section 8.3, it shall also be possible to transfer a software package to the node for a version update or downgrade. Again, the process of transferring the file, and activating it, must be two different processes, and none must require a reboot or affect its operational status. Indeed, it is strongly encouraged that, even for those proposals that are built using a pizza-box approach, and do not implement redundancy of the control board, ISSU is supported.

8.3. Zero-Touch Provisioning

Zero-Touch Provisioning (ZTP) is the process to deploy a Network Operating System (NOS) and a base configuration in a network element, so it can enter in production environment without any human configuration. ZTP process is executed for the first time when the device is turned on in the network.
Periodically, system vendors release new versions of their NOS; a very similar process can be done for upgrade scenarios. Nowadays, the NOS upgrade process is a vendor-dependent process, which differs between each vendor solution. The ZTP and upgrade mechanisms should evolve to generate a common procedure for open

white-box scenarios.
This procedure will influence the requirements with regards to configuration management. The ZTP platform or the SDN controller will have to maintain the NOS images and a pointer for the current configuration files that the network element requires.

8.4. Licensing

It is not the goal to define a strict licensing agreement at this stage, but some important ideas, in line with the disaggregation philosophy, would include:

- With regards to SW packages defined in sections 4 to 6, it should be possible to differentiate between licenses associated with each of these packages. The figure of an “all included” license should also exist.
- Upgrades to features existing in one of the packages should not affect the cost of other SW packages.
- The platform shall be compatible with a pay-as-you-grow model not only from the hardware point of view, but also as defined by software licenses.

8.5. Local regulation compliance

Again, it is not a goal to include in this section the full list of local regulations that would apply to the commercialization of an Open BNG. It is worth mentioning, however, a few of them.
The solution shall be compliant with EU GDPR regulation: REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation).
For the Brazilian market, the solution must have a Certificate of Conformity issued by a Designated Certification Body and approved/homologated by ANATEL.

9. *Glossary*

AAA	Authentication, Authorization and Accounting	**LAC**	L2TP Access Concentrator
ADMF	Administration Function	**LACP**	Link Aggregation Control Protocol
API	Application Programming Interface	**LAG**	Link Aggregation Group
AS	Autonomous System	**LAN**	Local Area Network
ASIC	Application-Specific Integrated Circuit	**LDP**	Label Distribution Protocol
ATM	Asynchronous Transfer Mode	**LED**	Light Emitting Diode
BFD	Bidirectional Forwarding Detection	**LFA**	Loop Free Alternate
BGP-4	Border Gateway Protocol 4	**LI**	Lawful Intercept
BGP-LS	BGP Link State	**LNS**	L2TP Network Server
BGP-LU	BGP Labelled Unicast	**MAC**	Media Access Control
BNG	Broadband Network Gateway	**MBH**	Mobile Backhaul
BRAS	Broadband Remote Access Server	**MPLS**	Multiprotocol Label Switching
CBS	Committed Burst Size	**MPLS-TP**	MPLS Transport Profile
CE	Customer Edge	**NAS**	Network Access Server
CG-NAT	Carrier Grade Network Address Translation	**NNI**	Network-Network Interface
CHAP	Challenge Handshake Authentication Protocol	**NOS**	Network Operating System
CIR	Committed Information Rate	**NTP**	Network Time Protocol
CLI	Command Line Interface	**OAM**	Operations, Administration and Management
CoA	Change of Authorization	**OCP**	Open Compute Project
CoS	Class of Service	**OLT**	Optical Line Termination
CPU	Central Processing Unit	**ONIE**	Open Network Install Environment
CUPS	Control and User Plane Separation	**OSPF**	Open Shortest Path First
DHCP	Dynamic Host Configuration Protocol	**OWAMP**	One-Way Active Measurement Protocol

DSCP	DiffServ Code Point	**PAP**	Password Authentication Protocol
DSL	Digital Subscriber Line	**PBB**	Provider Backbone Bridge
DSLAM	Digital Subscriber Line Access Multiplexer	**PCEP**	Path Computation Element Protocol
E-VPN	Ethernet VPN	**PE**	Provider Edge
EBS	Excess Burst Size	**PIR**	Peak Information Rate
ECMP	Equal Cost Multi Path	**PPP**	Point to Point Protocol
EIR	Excess Information Rate	**PTP**	Precision Time Protocol
FIB	Forwarding Information Base	**PW**	Pseudowire
FRR	Fast Reroute	**PWE3**	PW Emulation Edge to Edge
gNMI	gRPC Network Management Interface	**QoS**	Quality of Service
GPS	Global Positioning System	**RED**	Random Early Detection
GRE	Generic Routing Encapsulation	**RFC**	Request For Comments
gRPC	gRPC Remote Procedure Call	**RIP**	Routing Information Protocol
ICMP	Internet Control Message Protocol	**RSVP-TE**	Reservation Protocol – Traffic Engineering
IGMP	Internet Group Management Protocol	**SC**	Standard Configuration
IIR	Internal Intercept Function	**SCP**	Secure Copy Protocol
IP	Internet Protocol	**SDN**	Software Defined Networks
IS-IS	Intermediate System to Intermediate System	**SFP**	Small Form-factor Pluggable
L2TP	Layer 2 Tunneling Protocol	**SKU**	Stock Keeping Unit
SLA	Service Level Agreement	**URFP**	Unicast Reverse Path Forwarding
SME	Small or Medium Enterprise	**USB**	Universal Serial Bus
SNMP	Simple Network Management Protocol	**VLAN**	Virtual Local Area Network
SOHO	Small Office, Home Office	**VLL**	Virtual Leased Line
SSH	Secure Shell	**VM**	Virtual Machine
T-BC	Telecom Boundary Clock	**VPLS**	Virtual Private LAN Service

T-LDP	Targeted LDP	**VPN**	Virtual Private Network
TACACS	Terminal Access Controller Access Control System	**VRF**	VPN Routing and Forwarding
TCP	Transport Control Protocol	**VRRP**	Virtual Router Redundancy Protocol
TI-LFA	Topology Independent LFA	**VxLAN**	Virtual Extensible Local Area Network
TLS	Transport Layer Security	**WAN**	Wide Area Network
TWAMP	Two-Way Active Measurement Protocol	**WFQ**	Weighted Fair Queueing
UDP	User Datagram Protocol	**WRR**	Weighted Round Robin
UNI	User-Network Interface	**ZTP**	Zero-Touch Provisioning
UPF	User Plane Function		