TECHNICAL REPORT

TR-348
Hybrid Access Broadband Network Architecture

Issue: 1
Issue Date: July 2016

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Technical Report has been approved by members of the Forum. This Technical Report is subject to change. This Technical Report is copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Technical Report may be copyrighted by Broadband Forum members.

Intellectual Property

Recipients of this Technical Report are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of this Technical Report, or use of any software code normatively referenced in this Technical Report, and to provide supporting documentation.

Terms of Use

1. License

Broadband Forum hereby grants you the right, without charge, on a perpetual, non-exclusive and worldwide basis, to utilize the Technical Report for the purpose of developing, making, having made, using, marketing, importing, offering to sell or license, and selling or licensing, and to otherwise distribute, products complying with the Technical Report, in all cases subject to the conditions set forth in this notice and any relevant patent and other intellectual property rights of third parties (which may include members of Broadband Forum). This license grant does not include the right to sublicense, modify or create derivative works based upon the Technical Report except to the extent this Technical Report includes text implementable in computer code, in which case your right under this License to create and modify derivative works is limited to modifying and creating derivative works of such code. For the avoidance of doubt, except as qualified by the preceding sentence, products implementing this Technical Report are not deemed to be derivative works of the Technical Report.

2. NO WARRANTIES

THIS TECHNICAL REPORT IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS TECHNICAL REPORT SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY IMPLEMENTER OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS TECHNICAL REPORT.

3. THIRD PARTY RIGHTS

Without limiting the generality of Section 2 above, BROADBAND FORUM ASSUMES NO RESPONSIBILITY TO COMPILE, CONFIRM, UPDATE OR MAKE PUBLIC ANY THIRD PARTY ASSERTIONS OF PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS THAT MIGHT NOW OR IN THE FUTURE BE INFRINGED BY AN IMPLEMENTATION OF THE TECHNICAL REPORT IN ITS CURRENT, OR IN ANY FUTURE FORM. IF ANY SUCH

RIGHTS ARE DESCRIBED ON THE TECHNICAL REPORT, BROADBAND FORUM TAKES NO POSITION AS TO THE VALIDITY OR INVALIDITY OF SUCH ASSERTIONS, OR THAT ALL SUCH ASSERTIONS THAT HAVE OR MAY BE MADE ARE SO LISTED.

The text of this notice must be included in all copies of this Technical Report.

Issue History

Issue Number	Approval Date	Publication Date	Issue Editor	Changes
1	18 July 2016	8 August 2016	Guiu Fabregas, Nokia	Original

Comments or questions about this Broadband Forum Technical Report should be directed to help@broadband-forum.org.

Editor	Guiu Fabregas	Nokia	guiu.fabregas@nokia.com
Wireline-Wireless Convergence Work Area Directors	Dave Allan	Ericsson	david.i.allan@ericsson.com
	Hongyu Li	Huawei Technologies	hongyu.li@huawei.com

TABLE OF CONTENTS

List of Figures

List of Tables

Executive Summary

TR-348 defines a framework that enables converged network operators to offer their fixed access subscribers coordinated and simultaneous use of fixed broadband and 3GPP access networks.

This Technical Report identifies faster service turn-up for new subscribers, increased access reliability and enabling higher throughput for subscribers as the main drivers for the work. The framework documents considerations and requirements for the hybrid access subscriber sessions, traffic classification, distribution and transport, performance measurement and policy aspects.

In the Hybrid Access architecture, the CPE needs to have a second WAN interface, offering 3GPP-based access connectivity in addition to the Multi-Service Broadband Network (MSBN) access, and is then termed a Hybrid CPE (HCPE). In addition, a new logical function, the Hybrid Access Gateway (HAG), is introduced in the operator network. The HAG implements the corresponding network side mechanisms for Hybrid Access services.

1 Purpose and Scope

1.1 Purpose

With the rise of network-hosted services, highly reliable network access is becoming critical to subscribers. In addition, service providers are looking to supply a higher throughput for their subscribers to provide a better user experience, especially in those cases where subscribers can only be offered a low bitrate DSL access. Simultaneous use of fixed broadband and 3GPP access networks can provide such solutions.

The purpose of TR-348 is to identify the business needs and use cases for hybrid access services and define an architectural framework and the different deployment options. The framework documents considerations and requirements for the hybrid access subscriber sessions, traffic classification, distribution and transport, performance measurement and policy aspects.

1.2 Scope

As outlined in TR-203 [9], Fixed Mobile Convergence (FMC) generally refers to permitting a subscriber to access services over fixed and wireless networks. TR-203 defines business requirements, use cases, high-level functional architecture, and deployment options for interworking. Building on that, TR-291 [10] provides nodal requirements in support of the scope of TR-203.

This Technical Report specifies the architectural requirements to allow coordinated and, when needed, simultaneous use of fixed broadband access and 3GPP access networks for converged operators, enabling further FMC use cases. As the CPE is a fixed device, mobility is not considered as part of the service model. In contrast to TR-203 and TR-291, this Technical Report focuses on enabling solutions for the broadband connected home or enterprise, and thus the access device will be a geographically pinned CPE, not a 3GPP UE.

The main drivers for this work are:

- Enable higher throughput for subscribers by leveraging existing networks
- Increased reliability of access to WAN services
- Faster service turn-up for new subscribers

To support the above use cases, the CPE needs to have a second WAN interface, offering 3GPP-based access, and is then termed a Hybrid CPE (HCPE). In addition, a new logical function, the Hybrid Access Gateway (HAG), is required in the operator network. The HAG implements the corresponding network side mechanisms for Hybrid Access services.

A number of new capabilities are required to support Hybrid Access:

- Simultaneous use of fixed broadband and 3GPP access paths
- Traffic distribution policies and schemes
- Communication of traffic distribution policies to the HCPE and HAG
- Instrumentation and retrieval of access path metrics (capacity, state, etc.)

These will be addressed in WT-378 [12].

2 References and Terminology

2.1 Conventions

In this Technical Report, several words are used to signify the requirements of the specification. These words are always capitalized. More information can be found be in RFC 2119 [16].

MUST	This word, or the term “REQUIRED”, means that the definition is an absolute requirement of the specification.
MUST NOT	This phrase means that the definition is an absolute prohibition of the specification.
SHOULD	This word, or the term “RECOMMENDED”, means that there could exist valid reasons in particular circumstances to ignore this item, but the full implications need to be understood and carefully weighed before choosing a different course.
SHOULD NOT	This phrase, or the phrase "NOT RECOMMENDED" means that there could exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications need to be understood and the case carefully weighed before implementing any behavior described with this label.
MAY	This word, or the term “OPTIONAL”, means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option.

2.2 References

The following references are of relevance to this Technical Report. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this Technical Report are therefore encouraged to investigate the possibility of applying the most recent edition of the references listed below.
A list of currently valid Broadband Forum Technical Reports is published at www.broadband-forum.org.

Document		Title	Source	Year
[1]	TR-069 Amendment 5	*CPE WAN Management Protocol*	BBF	2013
[2]	TR-101 Issue 2	*Migration to Ethernet-Based Broadband Aggregation*	BBF	2011
[3]	TR-134 Corrigendum 1	*Broadband Policy Control Framework (BPCF)*	BBF	2013

[4]	TR-144	*Broadband Multi-Service Architecture & Framework Requirements*	BBF	2007
[5]	TR-146	*Subscriber Sessions*	BBF	2013
[6]	TR-177	*IPv6 in the context of TR-101*	BBF	2010
[7]	TR-178	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	2014
[8]	TR-187 Issue 2	*IPv6 for PPP Broadband Access*	BBF	2013
[9]	TR-203	*Interworking between Next Generation Fixed and 3GPP Wireless Networks*	BBF	2012
[10]	TR-291	*Nodal Requirements for Interworking between Next Generation Fixed and 3GPP Wireless*	BBF	2014
[11]	TR-300	*Policy Convergence for Next Generation Fixed and 3GPP Wireless Networks*	BBF	2014
[12]	WT-378	*Nodal Requirements for Hybrid Access Broadband Networks*	BBF	2016
[13]	TS 23.203	*Policy and charging control architecture, Release 12*	3GPP	Sept. 2015
[14]	TS 29.212	*Policy and Charging Control (PCC); Reference points*	3GPP	Sept. 2015
[15]	TS 23.401	*General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access*	3GPP	Sept. 2015
[16]	RFC 2119	*Key words for use in RFCs to Indicate Requirement Levels*	IETF	1997
[17]	RFC 4006	*Diameter Credit-Control Application*	IETF	2005

2.3 Definitions

The following terminology is used throughout this Technical Report.

Hybrid Access The coordinated and simultaneous use of two heterogeneous access paths (e.g., DSL and LTE).

Hybrid Access path Network connectivity instance between HCPE and HAG over a given access network; fixed broadband or 3GPP.

Hybrid Access path group	The set of paths in a Hybrid Access service instance.
Hybrid Access session	A logical construct that represents the aggregate of network connectivity for a Hybrid Access subscriber at the HAG. It represents all traffic associated with a subscriber by a given service provider, with the exception of Hybrid Access bypass traffic, and provides a context for policy enforcement.
Hybrid Access bypass	Mechanism by which selected traffic bypasses the Hybrid Access traffic distribution function and is instead bound to either the fixed broadband or the 3GPP access. Hybrid Access bypass traffic is not forwarded through the HAG, and as such is not part of the Hybrid Access session.
HAG	Hybrid Access Gateway. A logical function in the operator network implementing the network side mechanisms for simultaneous use of both fixed broadband and 3GPP access networks.
HCPE	Hybrid Customer Premises Equipment (CPE). CPE enhanced to support the access side mechanisms for simultaneous use of both fixed broadband and 3GPP access.
HA Class	Hybrid Access Class. An abstract set of traffic that will be subject to the same traffic distribution policy and priority over a Hybrid Access path group.
Flow	Per TR-146 [5], a grouping of traffic identified by a set of header information and port information including, but not limited to: IP header, Layer 2 (L2) Header, Virtual and/or Physical interface Port, and/or Agent Circuit ID information for a remote port in the access network. In addition, TR-134 [3] lists additional criteria to be considered for classification purposes, in the Traffic Flow Identifier definition. **NOTE:** TR-348 focuses on a particular class of flow, which is a set of packets that share an ordering constraint, and therefore is a stricter definition than that implied by the classification options in TR-134 and TR-146.

Per-flow distribution	Traffic distribution scheme whereby packets in the same flow (see Flow definition) are always sent over the same path in the Hybrid access path group. For example a system doing per-flow traffic distribution of IP traffic would forward all packets with a common 5-tuple over the same path. A 5-tuple is made up of following header fields: source IP address, source port, destination IP address, destination port and protocol.
Per-packet distribution	Traffic distribution scheme whereby packets in the same flow (see Flow definition) may be sent over different paths in the Hybrid access path group.
Service Data Flow	Per TR-134, an aggregate set of packet flows that matches a specific criterion.
Rating Group	Per RFC 4006 [17], a Rating-Group gathers a set of services, identified by a Service-Identifier, and subject to the same cost and rating type.

2.4 Abbreviations

This Technical Report uses the following abbreviations:

3GPP	3rd Generation Partnership Project
AAA	Authentication, Authorization & Accounting
ACS	Auto Configuration Server
ARP	Allocation and Retention Priority
BBF	Broadband Forum
BNG	Broadband Network Gateway
BPCF	Broadband Policy Control Function
BSS	Business Support System
CIR	Committed Information Rate
CoA	RADIUS Change of Authorization
CoS	Class of Service
CPE	Customer Premises Equipment
CRM	Customer Relationship Management
CWMP	CPE WAN Management Protocol
DHCP	Dynamic Host Configuration Protocol
DPI	Deep Packet Inspection
DSL	Digital Subscriber Line

EAP	Extensible Authentication Protocol
eNodeB	E-UTRAN Node B
FMC	Fixed Mobile Convergence
GBR	Guaranteed Bit Rate
GGSN	Gateway GPRS Support Node
HA	Hybrid Access
HAG	Hybrid Access Gateway
HCPE	Hybrid CPE
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPTV	Internet Protocol Television
LAN	Local Area Network
LI	Lawful Interception
LTE	Long Term Evolution
MPTCP	Multipath TCP
MSBN	Multi-Service Broadband Network
MTU	Maximum Transmission Unit
NDR	xDSL Net Data Rate
OAM	Operations Administration and Maintenance
OSS	Operation Support System
PBR	Policy-Based Routing
PC	Personal Computer
PCC	Policy Charging and Control
PCEF	Policy and Charging Enforcement Function
PCRF	Policy and Charging Rules Function
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PGW	Packet Data Network Gateway
PIR	Peak Information Rate
PPPoE	Point-to-Point Protocol over Ethernet
QCI	QoS Class Identifier
QoS	Quality of Service
RADIUS	Remote Authentication Dial-In User Service
RTT	Round-Trip Time
SDO	Standards Developing Organization
SGW	Serving Gateway
TCP	Transmission Control Protocol
TR	Technical Report

TV	Television
UDP	User Datagram Protocol
UE	User Equipment
URL	Uniform Resource Locator
VPN	Virtual Private Network
WAN	Wide Area Network
WA	Work Area
WT	Working Text

3 Technical Report Impact

3.1 Energy Efficiency

TR-348 allows the use of both 3GPP and fixed broadband networks, and requires a new Network element, the HAG. Therefore, some increase in energy consumption per subscriber will result.

3.2 IPv6

As TR-348 allows for simultaneous use of two access paths, IP address consumption will be greater than in networks offering a single access to subscribers. In addition, some of the approaches described in TR-348 use tunneling techniques, which will further increase IP address consumption. TR-348 enabled networks would therefore benefit from the use of IPv6, given its greatly increased IP address space compared to IPv4.

TR-348 is defined so that existing IPv4, IPv6 and dual-stack approaches can be used at the service provider's discretion.

3.3 Security

In general TR-348 does not impact security as all mechanisms currently used by both wireless and wireline networks individually would continue to be used by an HA system. The only obvious impact is that the utility of a man in the middle attack occurring between the HAG and the HCPE is diminished for flows that use per-packet distribution across the set of HA paths. In this scenario the attacker will only see a portion of the flow and be unable to reconstruct the missing data.

3.4 Privacy

TR-348 does not diminish the quality of privacy offered by either fixed broadband or 3GPP networks.

4 Use Cases

4.1 Introduction

Due to the trend of centralization of IT resources and applications into private or public clouds, the WAN connectivity of the CPE has become critical to serve the home, and branch offices.

For home users, VoIP and entertainment services like on-demand music, gaming and video streaming are now widely used resulting in an ever-increasing demand for higher bandwidth residential services.

Insufficient throughput will have an effect on the home users' customer experience, causing discontent and a decline in the usage of network-based services, which will in turn diminish revenue for services operators.

For business services, unavailability of the WAN connection for Customer Relationship Management (CRM) applications, point of sale terminals, telephony services, email and video conferencing services prevents them from working and impacts the operations of a business subscriber. Inability to use these applications will have economic consequences for the business users as well.

Furthermore, deployment of professional applications, web apps, telephony, collaborative applications, etc. in centralized resources greatly increases the bandwidth requirements. Insufficient bandwidth for such applications can reduce work efficiency and productivity.

Even if programs for rollout of very high broadband services have been launched, these will take time to reach broad coverage. This is driving service providers to find ways to significantly accelerate the eligibility to higher bandwidth services for a larger subscriber base and deliver connectivity better suited for the shift to network-based services, while leveraging the existing network infrastructure.

Hybrid Access services, through simultaneous use of fixed broadband and 3GPP access networks, can provide such a solution for converged service providers.

This section describes the key use cases addressed by this Technical Report.

4.2 Increased access capacity

A subscriber is provided with two access connections to a single premise for the purpose of delivering a higher bandwidth, for either upstream, downstream or in both directions.

Several distribution schemes can be used to spread traffic across the access paths, the following being some examples:

- Least Cost First
- Least Loaded First

- Traffic Load balancing
- Application-aware
- Traffic binding

See section 6.4 for further details.

4.3 Improved WAN reliability

As described in TR-203 section 4.8, a subscriber is provided with two access connections to a single premise to improve the availability of access to WAN services. A specified access path is considered to be the primary one and the second path used as a backup. Under normal conditions, only the primary access is used to deliver service in both the upstream and downstream direction.

a) **Simple backup:** when the primary access path fails completely, all traffic is moved onto the backup path. When the primary access path returns to service, traffic should be returned to the primary whilst minimizing service impact.

b) **Performance based backup:** when the primary path performance is degraded such that it cannot support the subscriber's applications, if the backup access path will offer better performance, then all the traffic is moved onto the backup path. When the primary path can provide the required performance, traffic should be returned to the primary path, whilst minimizing service impact.

c) **Per application performance based backup:** when the primary access path performance is degraded such that it cannot support the needs of a defined set of the subscriber's applications, traffic is moved onto the backup path such that the performance demands of the applications at risk can once more be satisfied. This may be achieved by moving all or some of the at-risk application traffic, or otherwise redistributing traffic such that the performance needs can be met.

The capacity and underlying performance of the Hybrid Access paths is likely to be asymmetric. Furthermore, it is possible that the performance of a component access path could be degraded in only one direction. In these cases, it will be necessary to monitor the unidirectional status and performance of each Hybrid Access path and support an asymmetric Hybrid Access traffic distribution function in order to deliver the best possible performance for the end user.

4.4 Fast service turn-up

A new fixed access is to be provided to a subscriber premise, but has a longer lead-time than desired. An additional 3GPP access that has a shorter lead-time is delivered to provide the subscriber with service until the fixed access can be provided.

- It is expected that when the fixed access is activated, there will be minimal service interruption.
- The 3GPP access could either be disabled once the fixed access is activated, or used to improve reliability and/or add capacity, based on policy.
- An operator may choose to modify billing policies once the fixed access is activated.

5 Hybrid Access reference architectures

5.1 Fit with existing BBF architectural specifications

TR-348 builds on BBF's MSBN architecture. It reuses the BBF core documentation defining the MSBN architectural framework (TR-144 [4]; TR-146; TR-178 [7]), as referred to within the document. Briefly, these documents describe key concepts defining BBF MSBNs, as well as adding supporting aggregation layer technologies to the TR-101 [2] specification, extending its support into the access network. These TRs also describe the necessary requirements to support OAM and hierarchical QoS for a diverse set of transport services. Hence, these define the underlying transport infrastructure for the architectural approaches specified in TR-348.

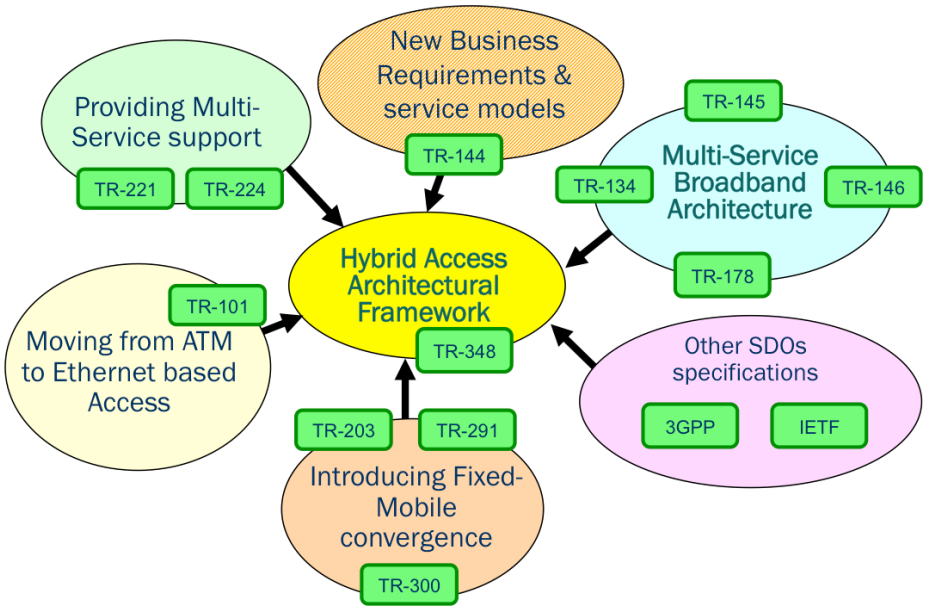

Figure 1 – Relation of TR-348 to BBF's and other SDOs' specifications

TR-348 addresses architectural requirements to allow coordinated and when needed, simultaneous use of fixed broadband access and 3GPP access networks for converged operators. As such, BBF specifications addressing control and user plane architectural options for MSBN integration with other networks, such as 3GPP networks (TR-134; TR-203; TR-291 and TR-300 [11]) are also re-used as necessary.

In particular, for those converged carriers aiming to re-use their existing 3GPP infrastructure for converged policy and charging control purposes, BBF TR-134 and TR-300 and the corresponding 3GPP and IETF specifications are referred to as options for realizing policy within the TR-348 architectural framework.

5.2 Hybrid Access deployment scenarios

Two different deployment scenarios are considered:

- Deployment Scenario #1: Hybrid CPE only
- Deployment Scenario #2: Hybrid CPE and Hybrid Access Gateway

The following subsections describe these scenarios.

5.2.1 Deployment scenario #1: Hybrid CPE (HCPE) only

In this scenario an enhanced CPE with support for both fixed broadband and 3GPP access interfaces is deployed, known as the Hybrid CPE (HCPE), which implements access side mechanisms for simultaneous use of both fixed broadband and 3GPP access networks. An example of an HCPE-only deployment is illustrated in Figure 2.

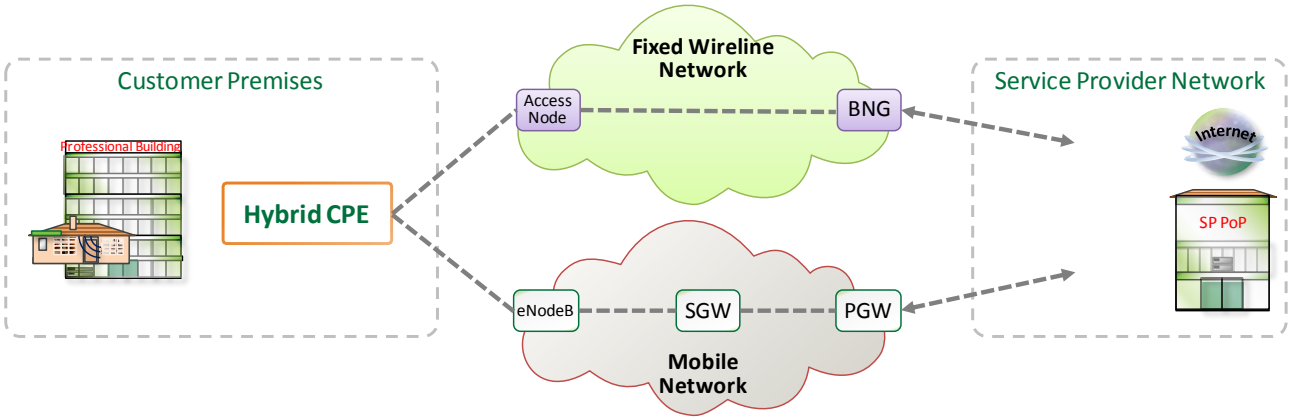

Figure 2 – Example of deployment Scenario #1: HCPE-Only

Several implementations are possible for this deployment model, depending on the IP addressing scheme selected. See Appendix I for further details.

5.2.2 Deployment scenario #2: HCPE and Hybrid Access Gateway (HAG)

As in the previous scenario, an HCPE is deployed that can make simultaneous use of two heterogeneous access connections. In addition, a new logical function is introduced in the operator network, the Hybrid Access Gateway (HAG). The HAG implements the network side mechanisms for simultaneous use of both fixed broadband and 3GPP access networks.

The HAG may be deployed stand-alone, integrated in the BNG or in the PGW. This will be discussed further in WT-378.

An example of deployment scenario # 2 is illustrated in Figure 3.

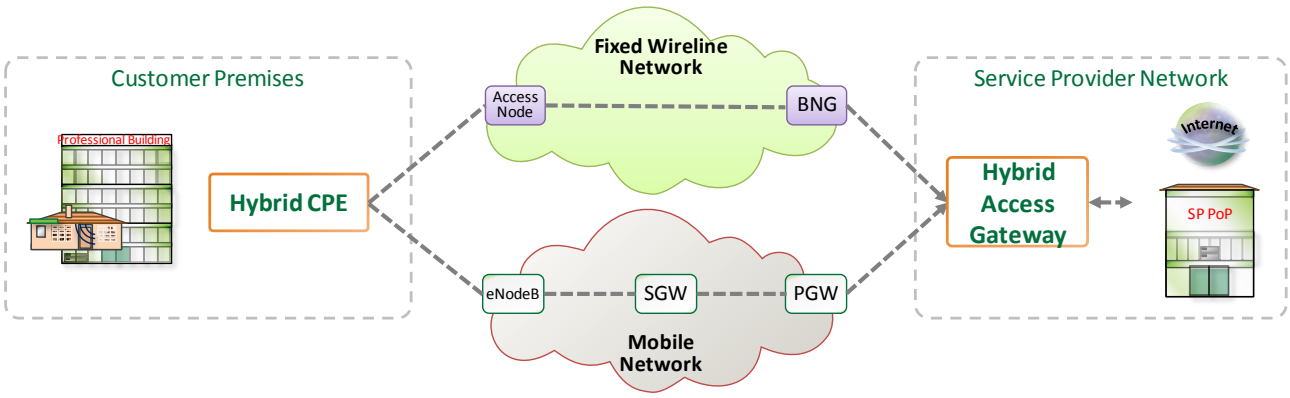

Figure 3 – Example of deployment Scenario #2: HCPE and HAG

In this scenario the HCPE and HAG provide enforcement points for QoS, traffic classification and traffic distribution between the available access paths. The HCPE acts as the enforcement point in the upstream direction while the HAG executes the equivalent functions on downstream traffic.

This allows converged service providers to provide advanced services, supporting elaborate and flexible traffic distribution schemes, including per-packet traffic distribution. For those cases where per-packet distribution is used, the HAG provides a re-ordering function for upstream traffic and the HCPE provides the corresponding function for downstream traffic.

5.2.3 Hybrid Access deployment scenario comparison

A Hybrid Access network requires extensive and coordinated control and enforcement of policies to ensure the most efficient usage of the network resources on both the fixed broadband and 3GPP access networks. This includes enforcement of QoS, filtering, traffic classification and distribution between the available access paths, etc. both in uplink and downlink directions.

In the HCPE-only deployment scenario, the HCPE is the sole entity responsible for providing the required functionality for a Hybrid Access service. While the HCPE has full visibility and control over the upstream traffic, it cannot directly control traffic in the downstream direction.

Downstream traffic forwarding happens through the BNG or PGW/GGSN, depending on the IP addressing model used and the traffic distribution done by the HCPE on upstream direction (see Appendix I). Downstream traffic cannot be easily and selectively forwarded over a given path, according to selected policy criteria.

QoS and filtering enforcement can be provided on downstream traffic, but require coordinated policy distribution to the PEPs in both edge gateways, the BNG and PGW/GGSN. In addition, an overall rate for the Hybrid Access subscriber cannot be enforced, unless static, pre-defined bandwidth allocation is done for each access.

Moreover, the HCPE-only solution cannot easily provide per-packet traffic distribution capabilities. Traffic on each path is sourced with a different IP address and would normally be interpreted as different flows from different hosts by the receiving endpoint. Support for MPTCP and similar solutions in the receiving host would allow using per-packet traffic distribution. However, it does

not allow service providers to have a self-contained solution that is independent of what may or may not be supported by other end devices. Even if support for MPTCP in software stacks is increasing, it is not yet widely available.

In contrast, the HCPE and HAG deployment scenario provides a single policy enforcement point for the downstream traffic of a Hybrid Access subscriber, the HAG.

Downstream traffic can be selectively forwarded through any of the available access paths, irrespective of the chosen path for upstream traffic, supporting more elaborated and flexible traffic distribution schemes. QoS and filtering can be provided at a single enforcement point, including the possibility of enforcing an overall rate for a given subscriber. Finally, the solution can provide flexible per-packet and/or per-flow traffic distribution based on policy and whether one or two paths are being used is irrelevant to the rest of the network.

Further definition and specification of requirements the HCPE-only deployment scenario are out of scope of this Technical Report. However example IP addressing options to support this deployment model are provided in Appendix I.

From this point forward, this Technical Report will only consider the HCPE and HAG model as the reference for Hybrid Access broadband networks.

5.3 Hybrid Access functional reference architecture

The following figure shows the Hybrid Access functional reference architecture.

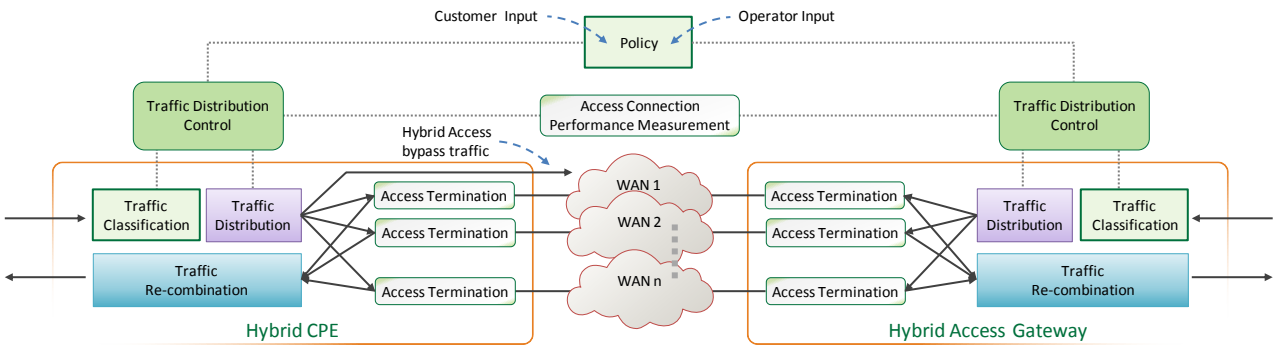

Figure 4 – Hybrid Access reference diagram

The various functional elements in Figure 4 perform the following:

- **Traffic Classification:** identifies and maps traffic flows into HA classes that may then be subjected to different distribution policies. Further details and requirements are given in section 6.3.
- **Traffic Control:** determines how traffic should be classified into HA classes and distributed across the set of paths in the Hybrid Access path group. Controls the rules enforced in the Traffic Classification and Traffic Distribution functions.
- **Traffic Distribution:** distributes traffic onto the appropriate Hybrid Access paths. Further details and requirements are given in section 6.4.

- **Access Termination:** terminates the Hybrid Access paths (which may be a physical link or logical tunnel).
- **Traffic Re-combination:** re-combines and re-sequences traffic flows received from the Hybrid Access path group (required when using per-packet traffic distribution)
- **Path Performance Measurement:** monitors the performance of the available Hybrid Access paths and feeds this information into the traffic distribution schemes. Further details and requirements are given in section 9.
- **Policy:** defines policies to manage and control the Hybrid Access functions based upon operator and subscriber input. Further details and requirements are given in section 7.

5.4 Hybrid Access Transport Models

This section describes a variety of options that may be used to transport traffic between the HCPE and HAG in Hybrid Access broadband networks:

1. L3 Overlay Tunneling, as described in section 5.4.1
2. L3 Network-based tunneling, as described in section 5.4.2
3. L4 Multipath, as described in section 5.4.3

The Hybrid Access path group is established between the HCPE, at the subscriber's site/home, and the HAG in the service provider's network. The Hybrid Access path group is transparent to the subscriber and the network beyond the HAG, as if a single access were present.

Some of the transport models described in this section use additional packet encapsulation (e.g. tunnels) in the uplink from the HCPE to the HAG through the access network. The access network is the section of the network that is usually most sensitive to MTU sizes (most constrained). As such, the MTU size of the access network will need to be considered to avoid fragmentation if possible. Alternatively, solutions such as fragmentation and reassembly, path MTU discovery or TCP MSS clamping could be used.

All the listed transport models above can support both packet-based and flow-based distribution.

5.4.1 L3 Overlay Tunneling

The following figure shows the logical architecture of this solution.

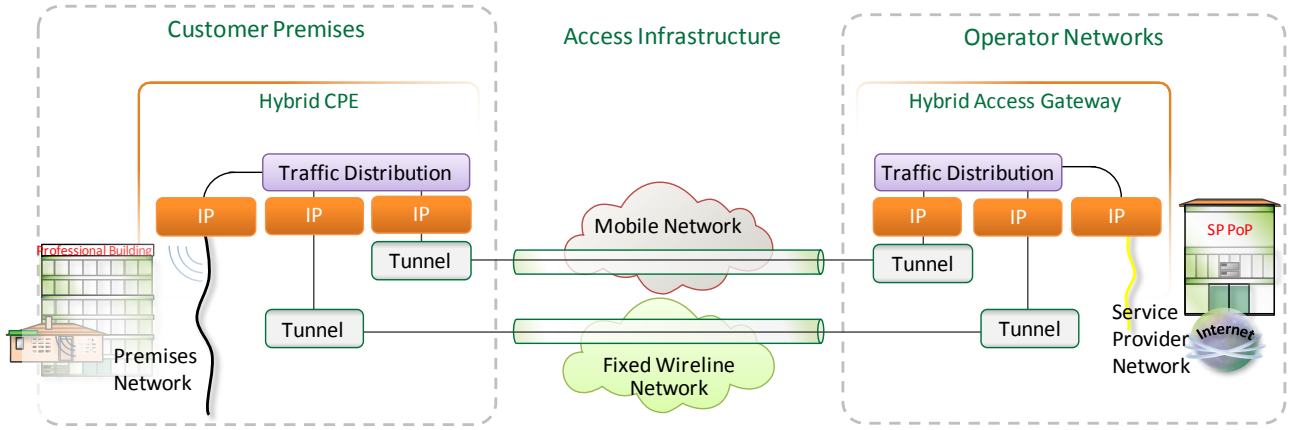

Figure 5 – L3 Overlay Tunneling

The connectivity between the HCPE and the HAG is established using tunnels on top of the access infrastructure. The tunnels are established between the HCPE and the HAG over each of the access paths. The HCPE is responsible for managing the tunnel (both establishment and tear down) as well as upstream forwarding decisions. The HAG is responsible for downstream forwarding decisions. The implementation itself is access network agnostic, therefore no changes to either the fixed broadband or the 3GPP access networks are necessary.

5.4.2 L3 Network-based Tunneling

The following figure shows the logical architecture of this solution.

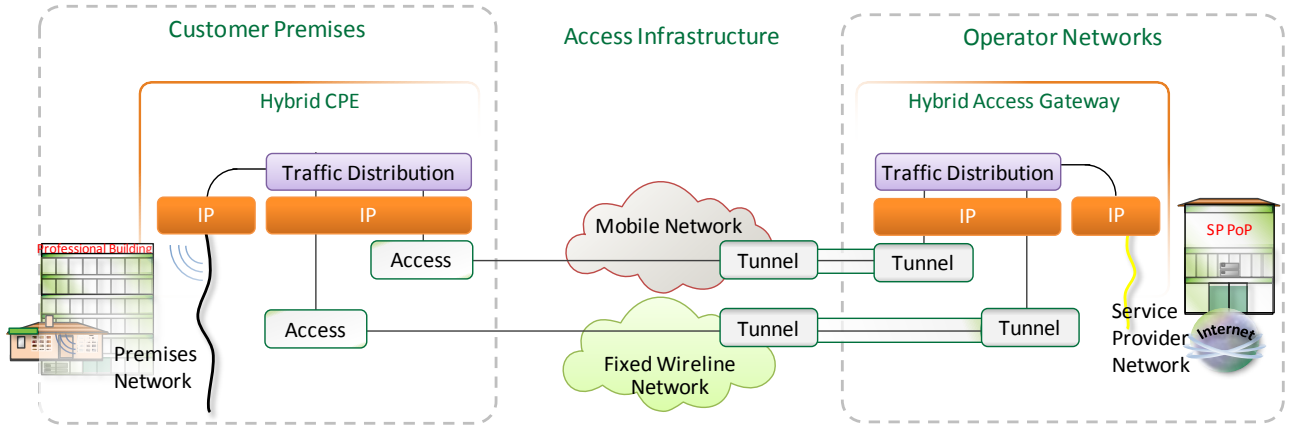

Figure 6 – L3 Network-based Tunneling

The connectivity between the HCPE and the HAG is realized by making use of the native technologies in both the fixed broadband (e.g. IPoE/PPPoE) and 3GPP access networks, from HCPE to BNG and from HCPE to eNodeB respectively. On setup, the network establishes the tunnels to the HAG on behalf of the subscriber's HCPE and stitches traffic from the access sessions

to those tunnels, in order to reach the HAG. Each Hybrid Access path is the end-to-end path resulting from stitching the access session in the respective access network with the corresponding tunnel from the access network to the HAG.

In this solution, the HCPE may use a single IP address for both Hybrid Access paths. The same address can be assigned via both access paths from the network.

5.4.3 L4 Multipath

The following figure shows the logical architecture of this solution.

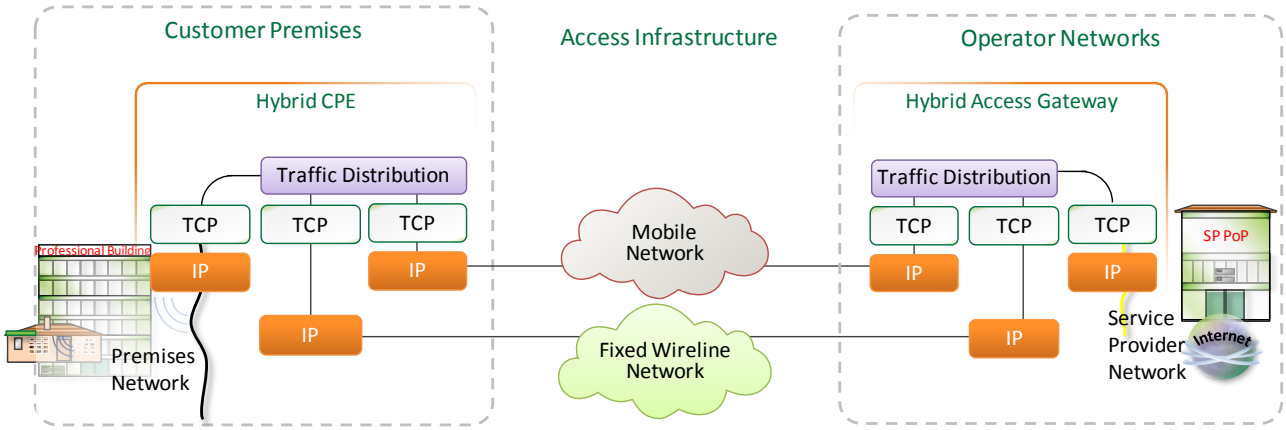

Figure 7 – L4 Multipath network

The connectivity between the HCPE and the HAG is established using a Layer 4 multipath transport service enabling IP flows to use multiple paths in the Hybrid Access path group simultaneously. As an example, a L4 multipath implementation using MPTCP sets up multiple TCP sub-flows over the different access networks and utilizes real time HCPE to HAG flow control. The HCPE and HAG are responsible for managing the MPTCP paths, including establishment and tear down.

The implementation itself is access network agnostic, therefore no changes at either the fixed broadband or the 3GPP access networks are necessary.

The HCPE and HAG terminate the end user layer 4 sessions before transporting the data over the Hybrid Access paths, effectively executing a proxy function for the end user sessions.

6 Hybrid Access architecture requirements

6.1 General requirements

This section lists the basic requirements that a Hybrid Access system must meet.

[R-1] The Hybrid Access system MUST support all the functional elements described in section 5.3.
[R-2] The Hybrid Access system MUST support both fixed broadband and 3GPP access.

6.1.1 IP Addressing

Hybrid Access broadband networks are designed so that existing address concepts for IPv4 and IPv6 can be used.

Fixed broadband and 3GPP networks support IPv4, IPv6 or dual-stack operation, both for infrastructure as well as subscriber addressing. It will be up to the service provider to use one, the other or both.

The IP addressing details and associated requirements for each of the Hybrid Access transport models will be addressed in WT-378.

[R-3] The Hybrid Access system MUST support both IPv4 and IPv6 for all Hybrid Access paths.
[R-4] The Hybrid Access system SHOULD allow the use of IPv4 or IPv6 for each Hybrid Access path independently.

6.1.2 Hybrid Mode de-activation

In certain circumstances, it might be necessary to disable Hybrid Access capabilities for a given subscriber. Examples of such situations include a change of contracted services by the subscriber, issues detected in the service, or even subscriber preference. This can be achieved by disabling the hybrid mode of the HCPE.

De-activation of the Hybrid Mode will cause the HCPE to fallback to operating as a single uplink CPE, using only the fixed broadband or the 3GPP access. Existing subscriber traffic may be impacted depending on the implementation.

[R-5] The HCPE MUST support deactivation of the hybrid mode and act as a traditional CPE using the fixed broadband access link.
[R-6] The HCPE MUST support deactivation of the hybrid mode and use only the 3GPP access link.
[R-7] The HCPE SHOULD support deactivation of the hybrid mode if it detects degradation of network performance, attributable to Hybrid Access.
[R-8] It MUST be possible for the Hybrid Access subscriber to de-activate the hybrid mode.

6.1.3 Access path monitoring and failover

Monitoring the availability of the Hybrid Access paths is an essential process in a Hybrid Access system. It allows reacting to failure conditions and redistributing traffic accordingly.

[R-9] The Hybrid Access system MUST be able to monitor the state of both Hybrid Access paths.
[R-10] In the case of IPoE connectivity, the HCPE MUST be able to monitor the HCPE to BNG connectivity using mechanisms described in BBF TR-146.
[R-11] In the case of PPPoE connectivity, the HCPE MUST be able to monitor the HCPE to BNG connectivity using LCP Echo-Request and Echo-Reply messages.
[R-12] The Hybrid Access system MUST support failover of all or part of the traffic through the 3GPP path according to policy, upon detection of failure of the fixed broadband access path.
[R-13] The Hybrid Access system MUST support failover of all or part of the traffic through the fixed broadband access path according to policy, upon detection of failure of the 3GPP path.

6.1.4 Authentication

The HCPE accesses both the fixed and mobile networks. Therefore, the HCPE needs to support the user credentials required for both networks.

[R-14] The HCPE MUST support authentication to 3GPP networks.
[R-15] The HCPE MUST support authentication to fixed broadband networks.
[R-16] The HCPE SHOULD support authentication to both fixed broadband and 3GPP networks using the same credentials.
[R-17] The Hybrid Access system SHOULD support authentication of the Hybrid Access session.

6.2 Hybrid Access State Machine

The state machine diagram for an individual HA Class is shown below for the deployment scenario of an end user with one fixed access path and one 3GPP access path:

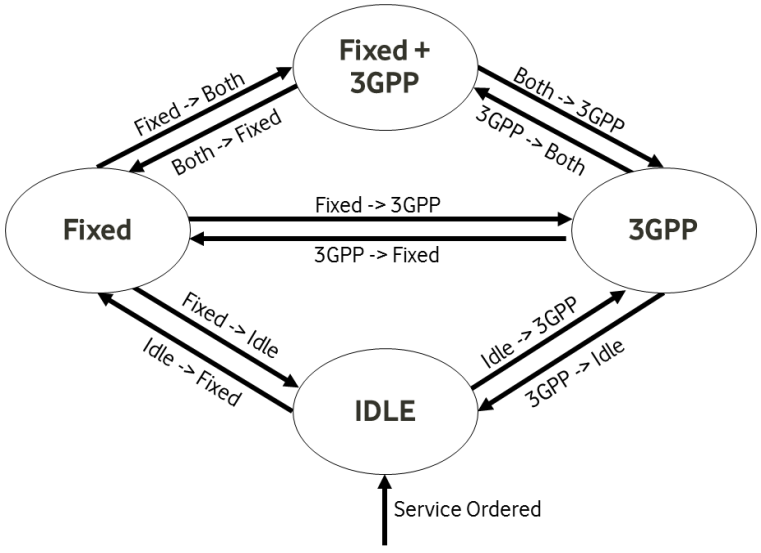

Figure 8 – Hybrid Access State Machine

The state diagram applies to a given HA class, for which there will be a defined classification policy. Thus, it could apply to all traffic, a subset of traffic or even a specific application or flow.

The available states are as follows:

1. IDLE: The HA system is considered to be down and traffic within the HA Class is not placed onto either access path. This will generally be when no access paths are available.
2. Fixed: Traffic within the HA class is placed only onto the Fixed access path.
3. 3GPP: Traffic within the HA class is placed only onto the 3GPP access path.
4. Fixed+3GPP: Traffic within the HA class is distributed onto both the Fixed and 3GPP access paths, according to the selected traffic distribution scheme for that HA class.

The triggers for transition between states will depend on the use case, and will be defined by operator and/or subscriber policy. It is expected that a degree of hysteresis and transition dampening will be incorporated into any implementation in order to avoid instability, or flapping between states.

[R-18] The Hybrid Access system MUST be able to transition between the Hybrid Access states.
[R-19] The Hybrid Access system MUST support enabling/disabling triggers for transition between states based on policy.
[R-20] The Hybrid Access system MUST support configurable stability criteria to avoid flapping between both paths.
[R-21] The Hybrid Access system must support live addition and removal of a Hybrid Access path in a Hybrid Access path group, whilst minimizing the service impact.

The following table lists the available triggers for transition, which may be enabled or disabled based on service provider and/or end-user policy, on a per HA class basis:

Transition	**Triggers for transition ***
Idle -> Fixed	• Fixed access becomes available • Fixed access is within performance threshold
Fixed -> Idle	• Fixed access fails and 3GPP access is not available • Fixed access fails and 3GPP access is outside of performance threshold • Fixed access fails and 3GPP access is not allowed for the HA class
Idle -> 3GPP	• 3GPP access becomes available • 3GPP access is within performance threshold
3GPP -> Idle	• 3GPP access fails and Fixed access is not available • 3GPP access fails and Fixed access is outside of performance threshold • 3GPP access fails and Fixed access is not allowed for the HA class
Fixed -> 3GPP	• 3GPP access becomes available and is preferred • 3GPP access becomes available and is considered lowest cost • Fixed access fails and 3GPP access is available and within performance threshold • Fixed access cannot meet the performance threshold and 3GPP access

	is available and within performance threshold • 3GPP access only will provide improved performance
3GPP -> Fixed	• Fixed access becomes available and is preferred • Fixed access becomes available and is considered lowest cost • 3GPP access fails and Fixed access is available and within performance threshold • 3GPP access cannot meet the performance threshold and Fixed access is available and within performance threshold • Fixed access only will provide improved performance
Fixed -> Both	• 3GPP access becomes available • 3GPP access is available and Fixed access only cannot meet the traffic demand • 3GPP + Fixed accesses will provide improved performance
Both -> Fixed	• 3GPP access fails • Fixed access only can meet traffic demand (Fixed access considered cheapest path) • Fixed access only will provide improved performance
3GPP -> Both	• Fixed access becomes available • Fixed access is available and 3GPP access cannot meet the traffic demand • 3GPP + Fixed accesses will provide improved performance
Both -> 3GPP	• Fixed access fails • 3GPP access only can meet traffic demand (3GPP access considered cheapest path) • 3GPP access only will provide improved performance
* **Note:** Triggers for transition are enabled/disabled by policy on a per HA class basis	

Table 1 Hybrid Access state machine transition triggers

6.3 Traffic Classification

Figure 9 shows the traffic classification function.

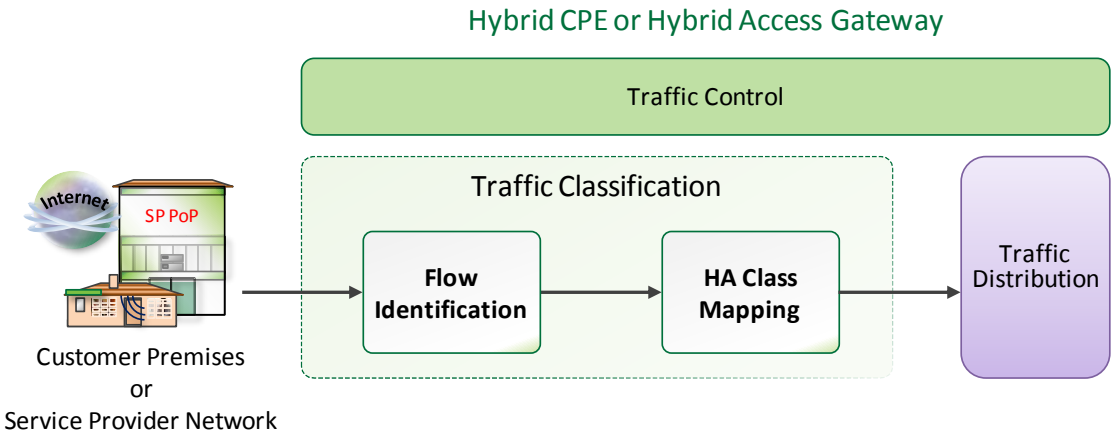

Figure 9 – Hybrid Access Traffic Classification

Upon receiving a packet from an interface that is not part of the Hybrid Access path group, the HCPE and HAG identify the traffic type and classify it into an aggregate HA class, which will later be given specific treatment by the traffic distribution function.

[R-22] The HCPE MUST support identifying a flow.
[R-23] The HCPE MUST support classifying a packet into an HA class after identifying the flow.
[R-24] The HCPE traffic classification function MUST be controlled by policy.
[R-25] The HAG MUST support identifying a flow.
[R-26] The HAG MUST support classifying a packet into an HA class after identifying the flow.
[R-27] The HAG traffic classification function MUST be controlled by policy.

6.4 Traffic Distribution

The traffic distribution function of both the HCPE and HAG spreads the traffic between the available Hybrid Access paths based on a combination of policy, the state of the network and the nature of the incoming traffic. More specifically, an incoming packet is sent over a given path according to:

- The policy rules set by the Policy function
- The performance of each of the access paths in the Hybrid Access path group
- The HA class of the packet

The system allows for specific types of traffic, e.g. video, voice, etc. to bypass the Hybrid Access Traffic Distribution function and forwards this traffic normally via fixed broadband or 3GPP accesses. Hybrid Access bypass and its implications are further described in section 6.4.1.

The Traffic Control function, present in the HCPE and HAG, is in charge of combining the policy input with the current state of the network to provide updated rules to the traffic distribution function. The Path Performance Measurement function provides input to the Traffic Control function to determine which paths may be used for a given traffic distribution scheme.

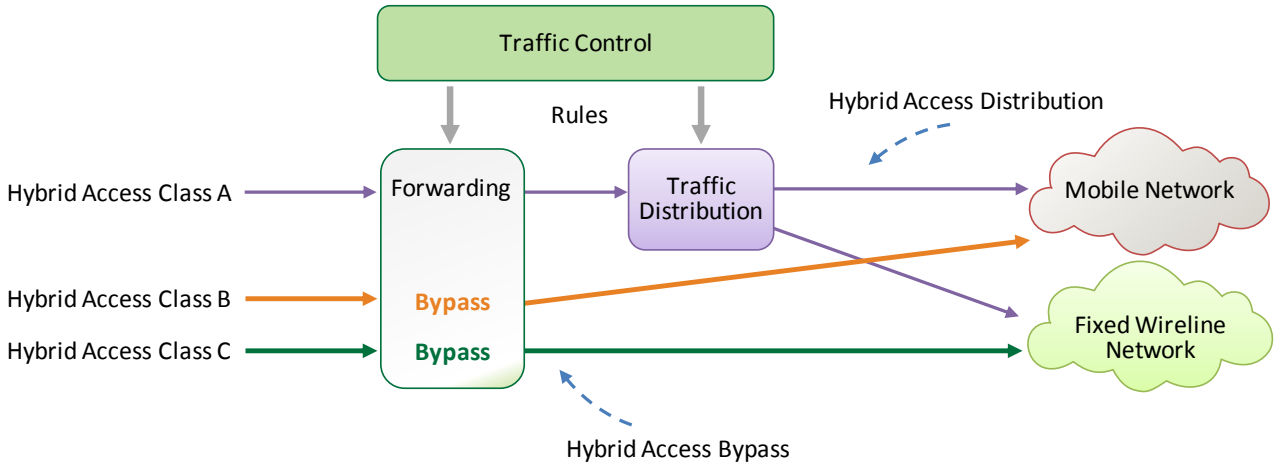

Figure 10 – Hybrid Access Traffic Distribution

The following are examples of traffic distribution schemes that may be used to distribute traffic between Hybrid Access paths:

- Least Cost First: The path with the lower associated cost is selected to forward traffic, allowing for overflow onto the higher cost path when congested. A possible usage for this distribution scheme would be to limit the use of the 3GPP access.
- Least Loaded First: The least loaded path is selected to forward traffic. In case of per-flow traffic distribution, the new flow stays on the selected path for the duration of its lifespan.
- Traffic Load balancing: Traffic is distributed between both access paths, allowing for equal or unequal traffic distribution, e.g. based on weights.
- Application-aware: The path selected for a given flow or application is that which is able to fulfill its performance demands, as defined by policy.
- Traffic binding: Selected traffic types or applications are bound to a given access path, as defined by policy. Policy may also specify if a given traffic type should be moved to the other access path, should the selected path fail.

When distributing traffic between Hybrid Access paths, a given HA Class may be subject to per-flow or per-packet based traffic distribution, based on policy.

A major challenge in Hybrid Access broadband networks is the distribution of traffic among paths with very different characteristics, e.g. delay, jitter, etc. It is essential that the traffic distribution schemes used do not degrade the user experience.

[R-28] The Hybrid Access system MUST support per-flow traffic distribution
[R-29] The Hybrid Access system MUST support per-packet traffic distribution
[R-30] The Hybrid Access system MUST support sufficient buffering to accommodate the different performance characteristic of the access paths, in the case of per-packet traffic distribution
[R-31] The Hybrid Access system MUST support traffic distribution according to policy.
[R-32] The Hybrid Access system MUST support using different traffic distribution schemes for each HA class in the same Hybrid Access path group.

6.4.1 Hybrid Access bypass

It must be possible for certain traffic flows to bypass the Hybrid Access Traffic Distribution function. This may be useful in the following cases:

- Access to services available only through fixed broadband access (e.g. IPTV) or 3GPP access (e.g. mobile services)
- End user experience preservation for given applications and traffic types

For instance, real time services like telephony are sensitive to latency and subscribers could perceive Hybrid Access as causing deterioration of end user experience, especially if using a per-packet distribution scheme, where reordering time could be substantial depending on the differential delay across the fixed broadband and 3GPP paths.

Hybrid Access bypass traffic is sent natively over one of the access paths, forwarded directly to the network after reaching the BNG or PGW, not forwarded through the HAG.

Policies related to Hybrid Access bypass can be applied at boot up time and modified at any time during the life of the session based on Subscriber and/or Operator policy input.

[R-33] The Hybrid Access system MUST permit traffic from selected HA Classes to bypass the Hybrid Access mechanism, controllable by policy.

6.4.1.1 Consequences of Hybrid Access bypass

Traffic that is not forwarded through the HAG traffic distribution block is not under its control. This impacts several aspects of traditional solutions. Lawful Interception Access Point and Charging functions for a given subscriber in fixed broadband networks are supported by a single network element, the BNG. The same is true for 3GPP networks, where these functions are supported by the PGW/GGSN. In contrast, providing these functions in a Hybrid Access network for a subscriber making use of Hybrid Access bypass will require executing them in up to three different network elements, i.e. the BNG, the PGW/GGSN and the HAG.

In addition, the amount of bandwidth consumed by Hybrid Access bypass traffic on each path in each direction is a priori unseen by the HAG. In order for the traffic distribution function of the HAG to be effective, the Hybrid Access system will need to continuously measure available bandwidth over each path.

[R-34] The Hybrid Access system's Traffic Distribution function MUST be adaptive to account for bypass traffic.

6.5 Hybrid Access configuration management

Once the HCPE establishes IP connectivity, it contacts the ACS using the procedures defined in TR-069 [1]. Upon successful establishment of the CWMP session between HCPE and the ACS, the ACS configures the HCPE as per existing BBF TRs, including Hybrid Access specific settings. Optionally, if TR-069 is used to relay Hybrid Access related policies to the HCPE, this step can be used to set the initial policies in the HCPE. In such a case, TR-069 may also be used for subsequent policy changes, as described in Section 7.3.2.1.

[R-35] Upon establishing IP connectivity, the HCPE MUST contact the ACS using TR-069.
[R-36] When the HCPE contacts the ACS, the ACS MUST update the HCPE with Hybrid Access related settings.
[R-37] The Hybrid Access system MUST support configuration of per-path attributes (e.g. MTU size).
[R-38] The Hybrid Access system MUST support enabling and disabling Hybrid Access paths.

6.6 Quality of Service

Many of the applications that run over broadband networks are sensitive to network conditions and have stringent requirements for service delivery. These include real-time applications such as VoIP, audio and video streaming services, etc. Supporting these applications over Hybrid Access subscriber services requires a QoS-enabled network.

A Hybrid Access broadband service makes use of both fixed and mobile access networks. As such, supporting QoS in the Hybrid Access architecture builds upon and makes use of the QoS mechanisms defined in existing BBF TRs and 3GPP specifications for the respective access networks.

6.6.1 QoS in BBF networks

TR-101 defines QoS mechanisms for Ethernet-based access and aggregation broadband networks, covering several deployment models including single and dual BNG deployments (joint or separate Video edge). Building on that, TR-144 defines the QoS needs in multi-service networks and TR-178 provides architectural and nodal requirements to support QoS in Multiservice Broadband Networks (MSBN) as well as support for a broader set of service architectures, introducing the concept of BNG hierarchies. TR-134 and TR-300 define mechanisms to relay QoS for subscriber sessions as part of the policy procedures.

6.6.2 QoS in 3GPP access networks

3GPP TS 23.401 [15] describes the QoS concepts for 3GPP mobile networks as well as the procedures and mechanisms required to support it. In addition to this, 3GPP TS 29.212 [14] describes the architecture and reference points used to convey QoS policies for 3GPP services.

Supporting multiple classes of service in a 3GPP network requires the use of dedicated resources in the 3GPP access and core elements. Specifically, each class of service with needs other than best effort will require signaling and maintaining a dedicated bearer. For those classes that require a higher scheduling priority but no bandwidth commitment, a non-guaranteed bit rate (non-GBR) bearer may be used. For those cases where bandwidth guarantees are required, a GBR bearer must be setup.

There are two key QoS attributes that characterize a traffic class, defined in 3GPP TS 23.203 [13]:

- QoS Class Identifier (QCI): An identifier representing QoS parameters, which can be used for packet forwarding treatment.
- Allocation and Retention Priority (ARP): Pre-emption priority in case there is not enough bandwidth.

Each bearer is associated with QCI and ARP attributes. The table below lists the standardized QCI values in 3GPP access networks and their characteristics.

QCI	Resource Type	Priority	Packet Delay Budget	Packet Error Loss Rate	Example Service
1	GBR	2	100 ms	10^{-2}	Conversational Voice
2		4	150 ms	10^{-3}	Conversational Video (Live Streaming)
3		3	50 ms	10^{-3}	Real Time Gaming
4		5	300 ms	10^{-6}	Non-Conversational Video (Buffered streaming)
65		0.7	75 ms	10^{-2}	Mission Critical user plane Push To Talk voice (e.g., MCPTT)
66		2	100 ms	10^{-2}	Non-Mission-Critical user plane Push To Talk voice
5	Non-GBR	1	100 ms	10^{-6}	IMS Signaling
6		6	300 ms	10^{-6}	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p file sharing, progressive video, etc.)
7		7	100 ms	10^{-3}	Voice Video (Live Streaming) Interactive Gaming
8		8	300 ms	10^{-6}	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p file sharing, progressive video, etc.)
9		9			
69		0.5	60 ms	10^{-6}	Mission Critical delay sensitive signaling (e.g., MC-PTT signaling)
70		5.5	200 ms	10^{-6}	Mission Critical Data (e.g. example services are the same as QCI 6/8/9)

Table 2 Standardized QCI characteristics in 3GPP TS 23.203

6.6.3 QoS in Hybrid Access broadband networks

Achieving end-to-end, consistent QoS requires the appropriate mechanisms to be in place in each of the networks involved. In a Hybrid Access network this includes the HCPE and HAG, as well as all the segments in both the fixed and 3GPP networks, i.e. the access, aggregation and core for the respective networks.

The more stringent QoS enforcement will be typically executed by the HCPE and HAG, as they are aware of the QoS requirements of each of the HA Classes in use for a given service and the traffic distribution scheme applied to the given class. Consistency between the QoS mechanisms in the Hybrid Access devices and existing intermediate network nodes is important. This can be achieved by aligning the required HA Classes with existing operationalized traffic classes in the intermediate nodes. However, new traffic classes could be operationalized for Hybrid Access, which might impact these nodes.

The framework enables service providers to offer equivalent functionality as in existing fixed broadband networks. From a QoS perspective, the framework allows subscriber-aware QoS enforcement. In addition, after traffic is classified into HA Classes according to a given set of classification rules, each HA Class may be subject to different QoS treatment, as per the nature of the traffic contained on it.

As described in the previous section, GBR, Non-GBR and Default bearers may be used in 3GPP access networks, depending on whether or not the service requires bandwidth guarantees and whether or not multiple classes of service are required. In Hybrid Access networks, if only best effort traffic is to be transported over the 3GPP access path then a Default bearer will be sufficient. If multiple classes of service are required over the 3GPP access path dedicated bearers will be required. For those classes with no bandwidth guarantees Non-GBR bearers should be used and for those with guaranteed bandwidth needs, GBR bearers will be required.

As in existing fixed broadband networks, a differentiated services QoS model will be used. In order to allow service providers to define flexible service offerings, it is desired to allow flexible bandwidth and priority allocation to traffic classes. It is also required that rate limits can be set at multiple levels:

- Per hybrid access session
- Per access path in a hybrid access session
- Per traffic class within a hybrid access path

Depending on the traffic distribution scheme being used and the network conditions at a given time dynamic QoS profile changes may be required, both for a given subscriber service or an HA Class within that service. It is desirable that any such changes are applied rapidly in the respective QoS enforcement points.

[R-39] The Hybrid Access system MUST support delay-sensitive services.
[R-40] The Hybrid Access system MUST support QoS.
[R-41] The Hybrid Access system MUST be able to differentiate between various Classes of Service (CoS) at the network level, for both upstream and downstream traffic.
[R-42] The Hybrid Access system MUST provide mechanisms for the classification of traffic into HA Classes.
[R-43] The Hybrid Access system MUST provide mechanisms for the distinct handling of traffic for each HA Class.
[R-44] The Hybrid Access system MUST support QoS policy enforcement per subscriber.
[R-45] The Hybrid Access system MUST support 2-level Hierarchical QoS (per subscriber / per path)
[R-46] The Hybrid Access system SHOULD support 3-level Hierarchical QoS (per subscriber / per path / per HA Class)
[R-47] The Hybrid Access system MUST support GBR bearers. See Table 2 for some examples.

6.7 Security

In order to access Hybrid Access services, an HCPE is configured with user credentials for both mobile and fixed access. For security and/or operational reasons, the operator might need to disable access to network services from a given HCPE, either through a given access or completely. In order to achieve this, existing mechanisms defined for fixed broadband and 3GPP networks can be used to revoke the credentials used by an HCPE.

[R-48] The Hybrid Access system MUST be able to revoke and reinstate the HCPE credentials used for fixed broadband access.

[R-49] The Hybrid Access system MUST be able to revoke and reinstate the HCPE credentials used for 3GPP access.
[R-50] When using Hybrid Access session authentication, the Hybrid Access system MUST be able to revoke and reinstate the HCPE credentials.

6.8 Lawful Interception

Lawful Interception allows law enforcement agencies to intercept upstream and/or downstream data associated with a subscriber, either all the data or a subset.

Interception of data of the Hybrid Access session associated with a given subscriber can be performed at the HAG, with the exception of Hybrid Access bypass traffic for that subscriber, as described in Section 6.4.1. In the event of a given subscriber making use of Hybrid Access bypass, interception will additionally need to be performed in the BNG and/or PGW/GGSN in the case of bypass to fixed broadband and/or 3GPP access networks respectively.

[R-51] The HAG MUST support Lawful Interception Access point capabilities for a Hybrid Access session.

7 Policy control in Hybrid Access networks

The Hybrid Access deployment scenarios presented in Section 5.2 introduce two new Policy Enforcement Points (PEPs): the HCPE; and the HAG (specifically for the deployment scenario presented in Section 5.2.2). According to the TR-134, the PEP is a logical entity that enforces policy decisions.

In the TR-348 context, some examples of the applicability of policy decisions at the HCPE and/or the HAG are: uplink/downlink traffic distribution between Hybrid Access paths; subscriber and charging differentiation; Hybrid Access bypass; and Hybrid Access path usage control.

Policy input may be sourced from both the service provider and the end-user sides, e.g. by making use of operator and subscriber service portals. Input from both sides is combined and subjected to mediation as per Section 7.2, resulting in a set of policy rules applying to the HCPE and a set of rules applying to the HAG for a given subscriber Hybrid Access path group.

According to the TR-134:

- **Static Policy Rules** are policy rules that are configured locally at the PEP and that may be activated or deactivated by default or remotely by the policy server.
- **Dynamic Policy Rules** are policy rules for which the definition is provided from the policy server to the PEP.

3GPP TS 23.203 has similar definitions for the existing policy rule types on 3GPP – the Policy and Charging Control (PCC) rule:

- **Dynamic PCC rules** are defined and provisioned to the PEP by the policy server.
- **Predefined PCC rules** are directly provisioned into the PEP and only activated by the policy server.

Please note that the 3GPP concept of predefined PCC rules is covered by the TR-134 definition of static policy rules.

There are several ways to convey a policy to the PEP: by activating a **Static Policy Rule**; or by sending a **Dynamic Policy Rule**. These policies might be conveyed to the HCPE and the HAG through out-of-band control interfaces, when managed by an external Policy Decision Point (PDP), or by in-band control interfaces between the HAG and the HCPE.

[R-52] The Hybrid Access system PEPs MUST support Dynamic Policy Rules taking precedence over the Static Policy Rules.

Figure 11 introduces the key elements involved in policy control for Hybrid Access, including the existing PEPs in the fixed and mobile networks, the new PEPs specific to Hybrid Access and the reference points between these elements and the policy management domain.

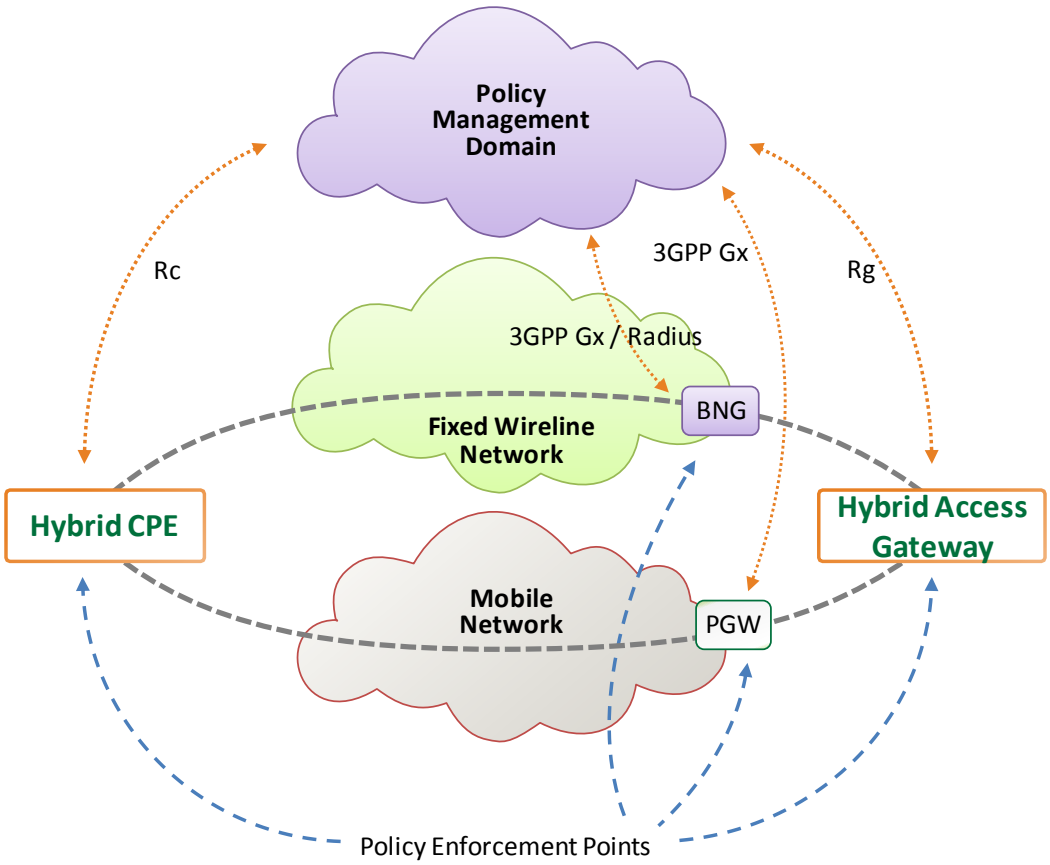

Figure 11 – Policy Control in Hybrid Access

In this context, two new logical reference points **Rc** and **Rg** are introduced between the policy management domain and the respective PEPs, the HCPE and the HAG.

Note: Rc and Rg are logical reference points and do not imply any particular implementation. Possible implementations are described in sections 7.1 and 7.3.

The following subsections describe the policy function, operator and subscriber policy input as well as the policy distribution.

7.1 Policy function

TR-134 defines the policy controller / server as the device where the PDP resides. Additionally, a PDP can serve a number of PEPs via the R interface. On a Hybrid Access system there might be more than one PEP in the network path. From TR-178 section 7.1.1, all MS-BNGs regardless of their deployment (centrally either stand alone or in a hierarchy; or at the edge in a hierarchy) need to support policy enforcement capabilities and requirements in TR-178 section 7.1.1, as received from the PDP through the R interface.

In MSBNs, two main types of policy control frameworks can be expected:

a) When the fixed and mobile domains each have their own distinct policy control functions.
b) When the fixed and mobile domains share the same policy control function.

For those use cases falling under type a), one solution is using TR-203 S9a interface to allow policy interworking between BBF MSBNs and 3GPP networks. In this scenario, it is necessary to enforce policy consistence. As such, one of the policy control functions will own the policy for Hybrid Access and relay policy to the other domain via its respective policy control function.

For type b) use cases where the common policy control function is the 3GPP Policy and Charging Rules Function (PCRF), TR-300 describes the requirements for MSBNs to interoperate with 3GPP control plane elements.

In TR-348, the policy function is the functional entity responsible for receiving operator (and/or subscriber) inputs, processing them as dictated by the logic functions and producing, managing and controlling the deployment of the required policies onto the hybrid access PEPs. Figure 4 shows the fit of the Policy Function in the Hybrid Access functional reference architecture.

The Hybrid Access Policy Function is a central function in the Hybrid Access architecture when it comes to ensuring robust and flexible operation of Hybrid Access systems. To that end, the Policy Function needs to be aware of the status of the fixed and mobile accesses according to a number of network parameters related to both HCPE and HAG functions.

The Hybrid Access policy function handles all the types of policies that the hybrid access system might require, including those related to enabling/disabling Hybrid Mode (e.g. based on minimum 3GPP RAT type), traffic distribution settings, applying policies for setting and/or changing those (and other relevant) parameters, and/or using those parameters as triggers to react to, and apply new policies.

The following requirements apply to the Hybrid Access framework.

[R-53] The Hybrid Access system MUST include a policy function
[R-54] The Hybrid Access policy function MUST be able to apply policies for both access networks in the Hybrid Access systems, both for policy interworking and policy convergence
[R-55] The Hybrid Access policy function MUST be able to apply policies to control the classification of traffic into HA Classes in the HA system
[R-56] The Hybrid Access policy function MUST be able to apply policies to control the distribution of traffic in the HA system
[R-57] The Hybrid Access system MUST support near real-time policy changes (including QoS) in response to changes in network conditions.

7.2 Policy input

The Hybrid Access Policy Domain provides a unified policy control point for the Hybrid Access system, which allows keeping and distributing a consistent policy across all involved PEPs.

The Hybrid Access Policy Domain mediates policy inputs from the following:

- Traditional OSS systems
- Operator portal(s)
- Subscriber portal

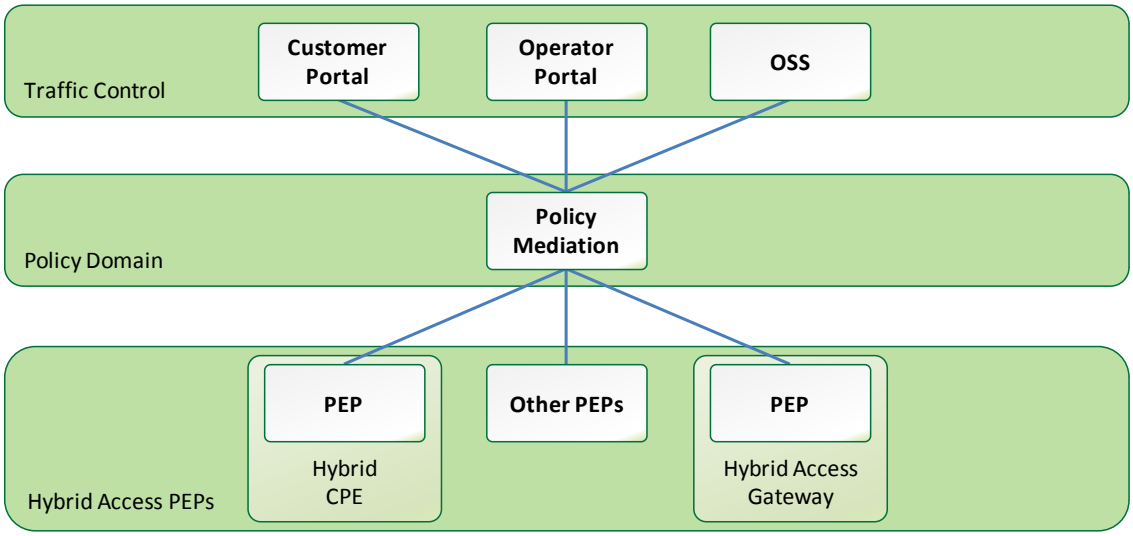

Figure 12 – Hybrid Access Policy Input

Mediation consists of the combination of the various policy inputs, subject to rules of precedence, etc. prior to distributing to the Hybrid Access system PEPs. Each PEP should only be sent the specific set of policy rules that apply to it, e.g. upstream traffic policy to the HCPE and downstream policy to the HAG.

The Subscriber portal should only expose those items that the Operator wishes to allow the subscriber to be able to modify.

[R-58] The Hybrid Access policy function MUST support open northbound interfaces for integration with OSS and portal platforms
[R-59] The Hybrid Access policy function MUST support updating policies during the life of a Hybrid Access subscriber session

7.3 Policy distribution

According to BBF TR-300 section 6.3.1 and TR-134, there are three different ways of providing policies to a Fixed Broadband PEP:

1. Policies configured locally in the PEP
2. Policies provided by the AAA via Remote Authentication Dial-In User Service (RADIUS) to the PEP
3. Policies provided by the PDP (i.e. PCRF in TR-300) via **R** reference point (i.e. Gx in TR-300) to the PEP

Option 1) corresponds to static policy rules configuration in the PEP and options 2) and 3) are mechanisms to distribute dynamic policy rules to the PEP.

As described in TR-134, policy distribution to PEPs can be achieved by either push or pull mechanisms. The Hybrid Access policy distribution to the HCPE and HAG needs to satisfy both of these models.

For any implementation protocol choice, and independent of whether it is an automated or a manual process, configuration and policy distribution are sensitive operations that require a secure environment.

[R-60] The Hybrid Access policy function MUST be able to distribute policies to the PEPs in the HCPE and HAG using both push or pull models.
[R-61] The Hybrid Access policy function MUST distribute policies to the PEPs in the HCPE and HAG via a secure and authenticated channel.

7.3.1 Policy distribution to HAG

The following subsections describe the potential mechanisms to be used in the context of Hybrid Access to distribute policy to the HAG.

7.3.1.1 Policy distribution using Gx

The use of a PCRF as PDP and the Gx reference point for providing policies to PEPs in the fixed and mobile networks is standardized in BBF TR-300 and in 3GPP technical specifications.

The Gx reference point resides between the PCEF and the PCRF and is specified in 3GPP TS 29.212. It supports both push and pull models. It may be used to:

- Provision, modify and remove rules for Policy and charging control (i.e. PCC Rules) from the PCRF to the PCEF
- Install event triggers at the PCEF to notify the PCRF when the PCEF detects such an event
- Communicate the AF Charging key in order to enable correlation of charging records by the charging systems

In the context of TR-134, the Diameter Gx application is an implementation of the R reference point. As per BBF TR-300 section 5.3.2, the Gx reference point enables the transfer of policies and charging control decisions from the PCRF to the PCEF for two types of devices connecting to the network:

- Fixed devices (e.g. RGs and business gateways)
- 3GPP UEs

7.3.2 Policy distribution to HCPE

At boot time, the HCPE is expected to perform policy initialization procedures to ensure that it has the necessary policies installed or if not yet installed, to get them. An example of such procedures would be to for the HCPE to contact an ACS after booting up to get policies using TR-069.

In the case where the policy initialization procedure fails, the HCPE needs to have a default policy to use; e.g. use a Least Cost First traffic distribution scheme.

[R-62] The HCPE MUST retrieve policies at boot time.
[R-63] The HCPE MUST support configuration of a default traffic classification and distribution policy.

7.3.2.1 Policy distribution using TR-069

In TR-134 the use case in 5.5.2 describes how a Policy Server interacts with the ACS in order to make a Policy change on the RG.

This can be expanded to allow applications and subscribers to interact directly or indirectly with the policy control function to allow dynamic changes to be made for both upstream policy and QoS in the HCPE.

The HCPE can be instructed by the ACS or by configuration to set a minimum interval time for requesting a new policy change triggered by an event at the subscriber side (e.g. performance issues in the form of packet loss) so that it can prevent situations such as the HCPE exhausting its available access bandwidth.

[R-64] The Hybrid Access policy function MUST be able to distribute policies to the PEP in the HCPE via TR-069.

8 Charging/Billing

The access paths that are part of a Hybrid Access path group use heterogeneous transmission media, and have different service delivery costs associated with them. As such, it is necessary to provide mechanisms to allow service providers to invoke differentiated charging for hybrid access service offerings.

TR-134 defines the accounting and charging needs for MSBNs and the associated requirements. TR-300 builds on TR-134 and 3GPP specifications and defines accounting and charging nodal requirements for wireline and wireless networks using a converged policy and charging system.

Further, to allow service providers to have the required flexibility when defining pricing tiers for hybrid access service offerings, TR-348 enabled networks should be capable of setting and enforcing differentiated charging criteria per access path and per HA class.

TR-348 supports the use of pre-/post-paid and per-use charging capabilities, including enforcing time or volume related quotas for those services making use of pre-paid charging models.

If charging per HA class is required on a given subscriber access, each HA class is mapped to a different Service Data Flow and assigned a specific rating group for charging purposes.

Since multiple PEPs are involved in a Hybrid Access system, e.g. HAG, BNG, PGW, etc., it is important to avoid duplicate charging to subscribers, by applying an appropriate charging strategy.

[R-65] The HAG MUST support at least one of the following:
- AAA based charging
- 3GPP Gy/Gz based charging
- 3GPP Gx usage-based charging

[R-66] The Hybrid Access system MUST support data volume counters on both the fixed broadband and 3GPP access paths for a given Hybrid Access session.

[R-67] The Hybrid Access system MUST support data volume counters in both the uplink and downlink direction, for each access path.

[R-68] The Hybrid Access system MUST be able to count data volumes per HA class.

[R-69] The Hybrid Access system MUST be able to assign specific rating groups to hybrid access classes for charging purposes.

9 Hybrid Access performance framework

9.1 Performance considerations for Hybrid Access

In contrast to a traditional single access service, the network performance of a Hybrid Access service depends upon multiple access paths and backhaul networks as well as the system, traffic distribution schemes, and policy selected to distribute traffic between them.

Objective end-to-end performance ultimately depends on the end-to-end packet delay characteristics and end-to-end packet loss characteristics. The source of these end-to-end characteristics in a hybrid access service can be decomposed into different components.

Figure 13 gives an overview of the subsystems that will largely dictate the performance characteristics of a Hybrid Access deployment. Note that in this case, the option to include a reactive mechanism between the HCPE and the HAG that may adapt the traffic distribution in response to variations in the underlying path performance makes it necessary to consider the characteristics of the combined path.

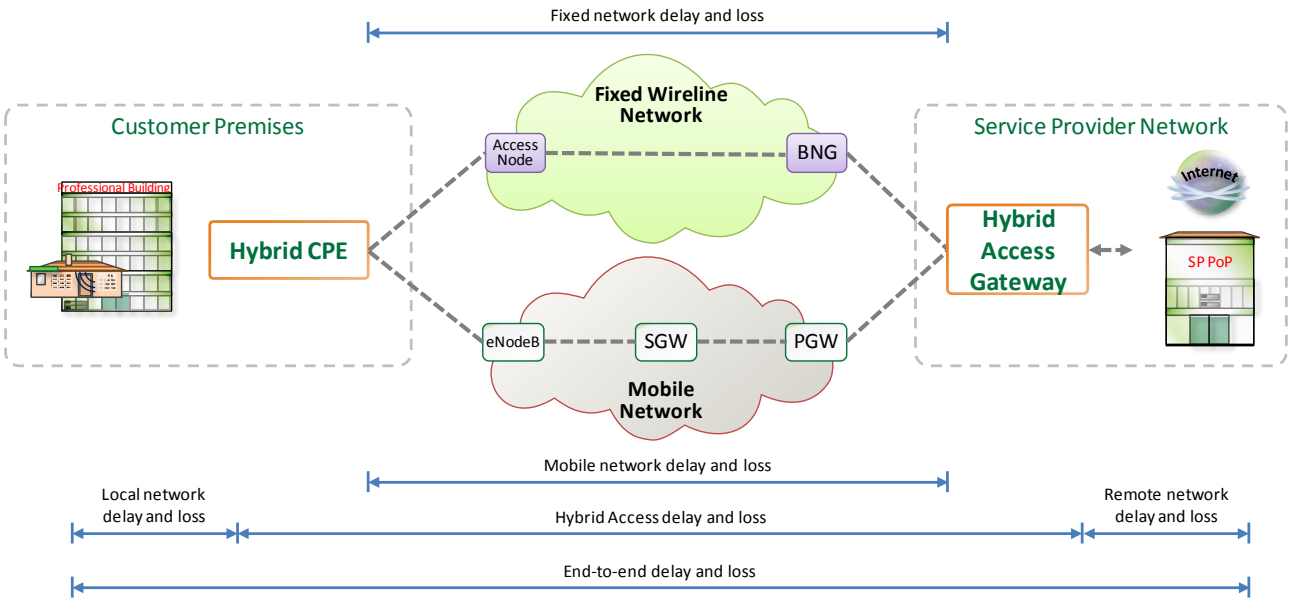

Figure 13 – Sources of end-to-end performance impairments in Hybrid Access deployments

Considering the different sources of packet delay and loss shown above, there are a number of reasons why special attention needs to be paid to the end-to-end performance on a Hybrid Access network:

- A Hybrid Access system may introduce additional packet delays into the system in its own right. Such delays can be caused by packet re-ordering, fragmentation and reassembly, etc.
- Packet performance may be different on each of the component access paths, therefore the user experience can be impacted by the distribution decisions of the HA system.

- Subscriber traffic moving between the different access paths during the lifetime of an application flow may result in a step change in performance that could trigger degradation in user experience below acceptable levels.
- Tools for measurement of end-to-end packet performance may not take the same access path as the applications it is intended to monitor.

In order to benchmark and support hybrid access solutions, a Service Provider should be able to understand and set bounds upon the performance of a Hybrid Access system as follows:

1. The relationship between the packet level performance of the available access paths (packet delay, delay variation and loss), and the delivered HA performance (and the corresponding impact upon user experience)
2. The short-term variation in end-to-end performance during changes in HA state such as enabling or disabling component access paths.
3. The short-term impact and response of the HA system to a sudden but short-term change in loss or latency on a component access path.
4. The impact that the performance characteristics of the HA system has upon the TCP congestion control algorithms.

9.2 KPIs

In order to assess the overall performance of a Hybrid Access system, certain information elements or KPIs (performance metrics, KPI, notification…), should be collected and analyzed periodically. These might in turn be further processed to obtain more complex metrics.

The following are examples of KPIs that could assist in characterizing the performance of each of the individual access paths:

- Available capacity (e.g. DSL NDR)
- Hybrid Access bypass bandwidth in use
- Latency (RTT or one-way)
- Jitter
- Loss (packet loss or BER)
- Queue utilization / congestion
- Link availability
- Hybrid Access traffic throughput (excluding bypass)

These may be used to obtain more complex metrics:

- Differential delay across both paths
- Hybrid Access latency: will include re-ordering time in those systems doing per-packet traffic distribution.
- Hybrid Access loss: loss percentage across both paths
- Hybrid Access throughput: measured as the sum of throughput processed on both access paths

- Efficiency: measured as Hybrid Access throughput divided by the sum of bandwidths of both paths
- Re-ordering latency: average time consumed in the Traffic Re-ordering function

All, or a subset of, these metrics may be used to characterize a Hybrid Access system. Performance characterization may need to be done in the upstream, downstream or both.

9.2.1 Measuring, reporting and reacting to KPIs

For metrics to be quantifiable, collection intervals and performance thresholds should be set for each of them. It must be possible to collect these performance metrics without taking the user out of service.

Collection intervals will define how often data records should be collected for a given metric and what sampling rates should be used.

Performance thresholds will allow detection of the performance of the system falling outside expected operational levels which might represent a potential problem. Several threshold levels may be defined, which will enable the system to provide different responses to different levels of criticality, from just generating an alarm, to taking preemptive action on the Hybrid Access system.

It is essential for the tools used for measurement of end-to-end packet performance to take the same access path as the applications it is intended to monitor for the measurements to be relevant.

[R-70] The Hybrid Access system MUST be able to monitor the performance of each individual path in the HA path group.
[R-71] The Hybrid Access system MUST be able to measure and collect data for a given set of performance metrics, as defined by policy or required by system operation.
[R-72] The Hybrid Access system MUST be able to export the collected data records for a given set of performance metrics.
[R-73] The Hybrid Access system MUST be able to compare the collected data for a given metric with the threshold(s) set for such metric.
[R-74] The Hybrid Access system MUST be able to generate an alarm, in the event of a violation of given performance criteria.

KPIs are needed to build up a long term view of network performance to assist setting semi-static policies, but they can also be used in near real-time to instigate more dynamic changes in path selection. For example, if the available bandwidth, delay or packet loss of one path go beyond a configured threshold, then a traffic stream could be switched to the other path. However, this needs to be done with care, particularly with regard to the following:

- The performance deterioration should be real and sustained, not a very short term glitch
- The alternative path performance needs to be better, and not beyond the same performance threshold itself
- There needs to be a defined reversion policy, with some hysteresis to avoid flapping between two paths

- Charging and volume caps may need to be taken into account when deciding whether or not to switch

Dynamic path switching can also be triggered immediately by a network event, for example losing a link, rather than waiting for this to impact a KPI.

[R-75] The Hybrid Access system MUST be able to take action, as set by policy, in the event of a violation of given performance criteria.

Appendix I. Informative HCPE-only implementation examples

I.1 Reference model

The following figure describes the reference architectures for deployment scenario #1: HCPE.

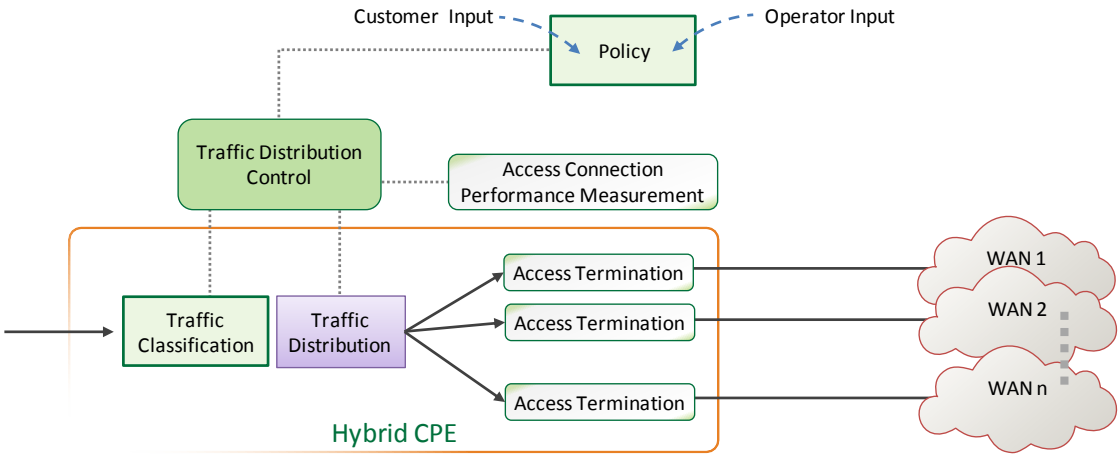

Figure 14 – Deployment scenario #1 reference diagram

Per-packet based traffic distribution is not possible for this scenario, as packet re-ordering is not possible on the network side.

I.2 IP Addressing implementation examples

This section provides example IP addressing options that can be used in an HCPE Only scenario, with the focus on IPv4. Similar models can be used when using IPv6 addressing.

The HCPE has two access interfaces on the WAN side (A and B in Figure 15). Each WAN interface will be assigned an IP address. In addition, an IP subnet will be assigned on the LAN side (C in Figure 15). These addresses may be allocated dynamically or statically by configuring the HCPE. The HCPE is the default gateway of the hosts on the LAN side.

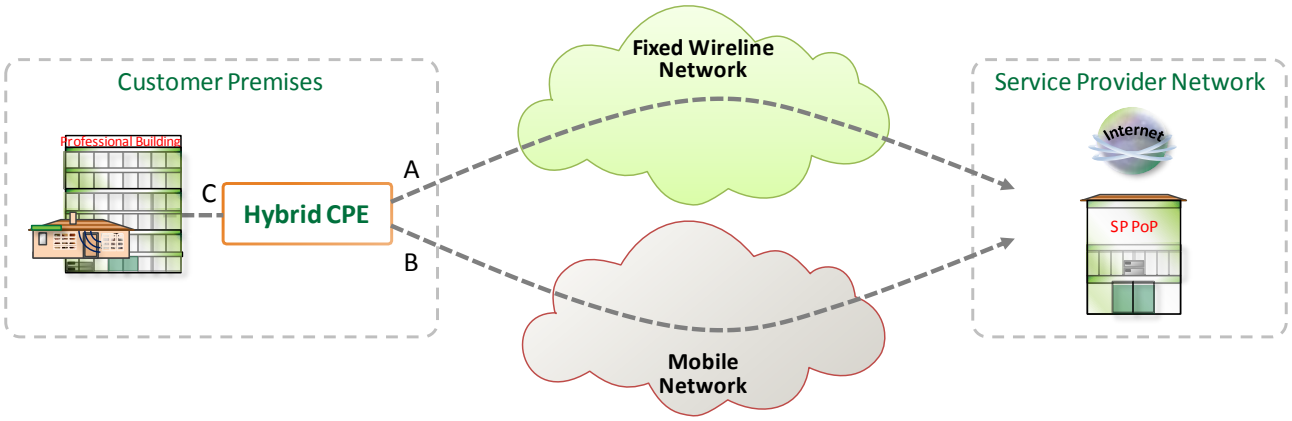

Figure 15 – HCPE only scenario

This appendix provides several implementation examples, covering several possible combinations of public and private IP addressing usage for each of these interfaces.

I.2.1 Public WAN IPs and private site/home subnet

In this scenario the HCPE has two public IP addresses on the WAN side and a private IP subnet on the LAN. In this case the HCPE translates the private addresses of the LAN to the public address of the either of the WAN interfaces, depending on the outgoing interface chosen for a given IP flow. In this model upstream and downstream packets of the same flow are guaranteed to use the same path, as downstream traffic is sent to the WAN interface that was chosen for the flow, due to NAT.

I.2.2 Private WAN IPs and public site subnet

In this scenario the HCPE has two private IP addresses on the WAN side and a public IP subnet on the LAN. A Framed-Route to the public subnet behind the HCPE should be dynamically enabled and advertised both on the BNG and the PGW/GGSN. This scenario may typically be used for business services.

I.2.3 Private WAN IPs and public site loopback

In this scenario the HCPE has two private IP addresses on the WAN side and a private IP subnet on the LAN. In addition, the HCPE is configured with a single public IP loopback address for the site. In this case the HCPE translates the private addresses of the LAN to the public address of the loopback. A Framed-Route to the public IP of the loopback in the HCPE should be dynamically enabled and advertised both on the BNG and the PGW/GGSN. This scenario may typically be used for business services. The main difference with the previous model is that it reduces IP address consumption to a single IP address per site.

End of Broadband Forum Technical Report TR-348