TECHNICAL REPORT

TR-378
Nodal Requirements for Hybrid Access Broadband Networks

Issue: 1
Issue Date: May 2019

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

Issue History

Issue Number	Approval Date	Publication Date	Issue Editor	Changes
1	10 May 2019	10 May 2019	Guiu Fabregas Nokia	Original

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

Editor	Guiu Fabregas	Nokia
Wireline-Wireless Convergence Work Area Director(s)	Dave Allan	Ericsson

Table of Contents

List of Figures

List of Tables

Executive Summary

This Technical Report specifies nodal requirements for Hybrid Access broadband networks, in support of TR-348 (Hybrid Access Broadband Architecture), to enable service providers to offer higher throughput and greater reliability to their subscribers by the use of both fixed broadband and 3GPP access networks.

Nodal requirements are provided for the three transport models defined in TR-348:

- L3 Overlay tunneling
- L3 Network-based tunneling
- L4 Multipath

1 Purpose and Scope

1.1 Purpose

The purpose of TR-378 is to define nodal requirements in support of TR-348 [4], which specified the Hybrid Access Broadband Network Architecture. This will enable service providers to offer higher throughput and greater reliability to their subscribers by means of simultaneous use of fixed broadband and 3GPP access networks.

1.2 Scope

This Technical Report defines nodal requirements for three Hybrid Access transport models connecting Hybrid CPE (HCPE) and Hybrid Access Gateway (HAG), as described in Section 5.4/TR-348:

- L3 Overlay tunneling
- L3 Network-based tunneling
- L4 Multipath

It defines the base Hybrid Access requirements, as well as those related to specific needs for each transport model.

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 [6].

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-101 Issue 2	*Migration to Ethernet-Based Broadband Aggregation*	BBF	2011
[2]	TR-134 Corrigendum 1	*Broadband Policy Control Framework (BPCF), Issue 1, Corrigendum 1*	BBF	2013

[3]	TR-178	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	2014
[4]	TR-348	*Hybrid Access Broadband Network Architecture*	BBF	2016
[5]	3GPP TS 29.274	*3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C); Stage 3*	3GPP	March 2015
[6]	RFC 2119	*Key words for use in RFCs to Indicate Requirement Levels*	IETF	1997
[7]	RFC 2698	*A Two Rate Three Color Marker*	IETF	1999
[8]	RFC 2890	*Key and Sequence Number Extensions to GRE*	IETF	2000
[9]	RFC 6824	*TCP Extensions for Multipath Operation with Multiple Addresses*	IETF	2013
[10]	RFC 8157	*Huawei's GRE Tunnel Bonding Protocol*	IETF	2017
[11]	draft-ietf-tcpm-converters	*0-RTT TCP Convert Protocol*	IETF	Ongoing

2.3 Definitions

The following terminology is used throughout this Technical Report.

Hybrid Access Per TR-348 [4], the coordinated and simultaneous use of two heterogeneous access paths (e.g., DSL and LTE).

Hybrid Access path Per TR-348, network connectivity instance between HCPE and HAG over a given access network; fixed broadband or 3GPP.

Hybrid Access path group Per TR-348, the set of paths in a Hybrid Access service instance.

Hybrid Access session Per TR-348, 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	Per TR-348, 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	Per TR-348, 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	Per TR-348, 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	Per TR-348, 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. TR-134 [2] in the Traffic Flow Identifier definition lists additional criteria to be considered for classification purposes.
Per-flow distribution	Per TR-348, a 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.
Per-packet distribution	Per TR-348, a traffic distribution scheme whereby packets in the same flow (see Flow definition) may be sent over different paths in the Hybrid access path group.
Transport Converter	Per draft-ietf-tcpm-converters [11], a Transport Converter is a function that is installed by a network operator to aid the deployment of TCP extensions and to provide the benefits of such extensions to clients. A Transport Converter may support conversion service for one or more TCP extensions. In the context of Hybrid Access, a HAG provides the Transport Converter function for the L4 Multipath transport model, when using MPTCP Explicit mode.

Converter Client	A TCP client that supports multipath transport capabilities, and uses the Converter Protocol to forward traffic to an end server through a Transport Converter. In the context of Hybrid Access, the Converter Client function is implemented by the HCPE, in support of the L4 Multipath transport model, when using MPTCP Explicit mode.
Transparent Proxy	A Transparent proxy is one where the clients are unaware of its existence. A Transparent proxy does not change the source or destination IP addresses of the forwarded packets.
Non-transparent proxy	A Non-Transparent proxy is one where the clients are aware of its existence. The clients either learn or are provisioned with the proxy IP address(es). The source IP address of a packet forwarded by a Non-Transparent proxy to a remote machine located in the Internet is an IP address of the proxy, not the address of the host.

2.4 Abbreviations

This Technical Report uses the following abbreviations:

3GPP	3rd Generation Partnership Project
AAA	Authentication, Authorization & Accounting
BBF	Broadband Forum
BNG	Broadband Network Gateway
BPCF	Broadband Policy Control Function
CGNAT	Carrier-Grade NAT
CIR	Committed Information Rate
CPE	Customer Premises Equipment
DHCP	Dynamic Host Configuration Protocol
DSL	Digital Subscriber Line
eNodeB	E-UTRAN Node B
GPRS	General Packet Radio Service
GTP	GPRS Tunnelling Protocol
GTP-U	GTP User Plane
GRE	Generic Routing Encapsulation
HA	Hybrid Access
HAG	Hybrid Access Gateway
HCPE	Hybrid CPE

IETF	Internet Engineering Task Force
IP	Internet Protocol
LAN	Local Area Network
LTE	Long Term Evolution
MPTCP	Multipath TCP
MS-BNG	Multi-Service BNG
MTU	Maximum Transmission Unit
NAPT	Network Address Port Translation
NAT	Network Address Translation
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PGW	Packet Data Network Gateway
PPPoE	Point-to-Point Protocol over Ethernet
QoS	Quality of Service
RADIUS	Remote Authentication Dial-In User Service
RTT	Round-Trip Time
SGW	Serving Gateway
SRA	Seamless Rate Adaptation
TCP	Transmission Control Protocol
TFO	TCP Fast Open
TR	Technical Report
UDP	User Datagram Protocol
WAN	Wide Area Network
WA	Work Area

3 Technical Report Impact

3.1 Energy Efficiency

TR-348/TR-378-based networks allow the use of both 3GPP and fixed broadband networks, and require a new Network element or function, the Hybrid Access Gateway (HAG). Further, the Hybrid Customer Premises Equipment (HCPE) will contain a receiver and transmitter for the mobile path. Therefore, some increase in energy consumption per subscriber will result.

3.2 Security

In general, TR-348/TR-378-based networks do not impact security as all mechanisms currently used by both wireless and wireline networks individually would continue to be used by a Hybrid Access (HA) system.

3.3 Privacy

TR-378 does not have a significant impact on the privacy offered by either fixed broadband or 3GPP networks. However, there may be a need to share customer information between the two networks to provide this service, which may introduce some privacy issues.

4 Architecture and Transport Models

Hybrid Access broadband networks enable converged network operators to offer their fixed access subscribers coordinated and simultaneous use of fixed broadband and 3GPP access networks, for higher throughput, increased access reliability, and faster service turn-up, as described in TR-348.

TR-348/TR-378 networks support flexible, policy-based traffic distribution over both networks using a variety of traffic distribution schemes (e.g., least cost first, load balancing, etc.), including both per-flow and per-packet traffic distribution. In addition, TR-348/TR-378 networks measure the performance of the Hybrid Access paths, allowing the dynamic modification of the traffic distribution based on the measured KPIs. Hybrid Access broadband networks also allow selected traffic to bypass the Hybrid Access traffic distribution function and be sent natively over one of the access networks.

Fixed network architectures, such as TR-101 [1] and TR-178 [3], use a Multi-Service Broadband Network Gateway (MS-BNG) as the fixed line access gateway. Hybrid Access broadband networks introduce an additional logical function located in the network, the HAG, which enables the simultaneous use of fixed broadband and 3GPP access for broadband services.

The existing network architectures are extended with additional reference points from the HAG to the different network elements of the fixed broadband and 3GPP networks.

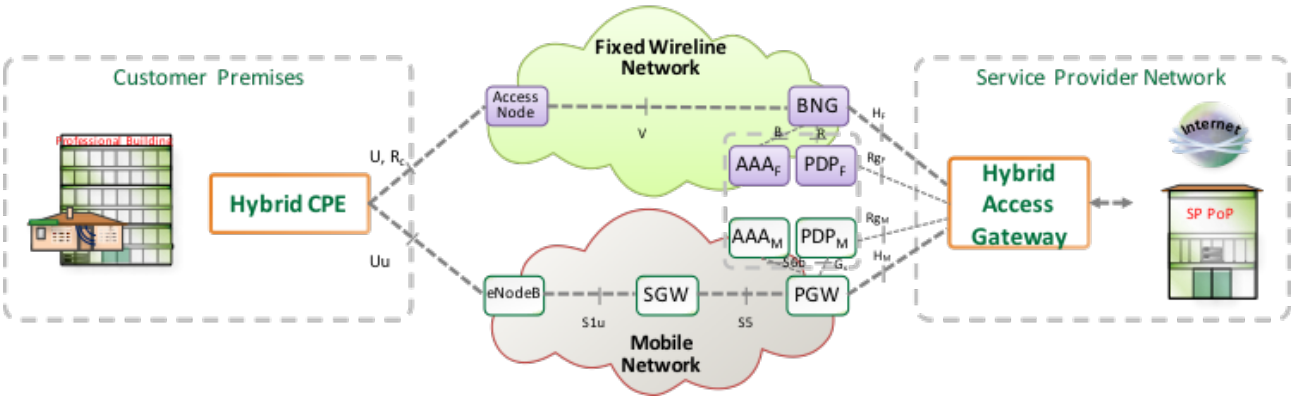

Figure 1 – Generic Hybrid Access network architecture and reference points

The HAG has two reference points to the fixed broadband network; the H_f from the MS-BNG to the HAG, enabling data transport between both functions; and the R_g interface from the fixed broadband Policy Decision Point (PDP) to the Policy Enforcement Point (PEP) function in the HAG, for policy control.

Similarly, the HAG has also two reference points to the 3GPP (Mobile) network; H_m, which identifies the interface between the Packet Data Network Gateway (PGW) and the HAG, and the R_g interface from the 3GPP PDP to the PEP function in the HAG, for policy control.

Table 1 lists the reference points for a Hybrid Access broadband network.

Reference Point	Description	Source
U	DSL Access	BBF TR-101
V	Access/Aggregation	BBF TR-101
B	MS-BNG to fixed broadband AAA	BBF TR-134
R	MS-BNG to fixed broadband PDP	BBF TR-134
R_g	HAG to fixed broadband PDP and/or 3GPP PDP	BBF TR-378
R_c	HCPE to fixed broadband PDP and/or 3GPP PDP	BBF TR-378
H_f	MS-BNG to HAG	BBF TR-378
Uu	Mobile Access	3GPP TS23.002
S1u	eNodeB to SGW	3GPP TS23.002
S5	SGW to PGW (non roaming)	3GPP TS23.002
S6b	PGW to 3GPP AAA	3GPP TS23.002
Gx	PGW to 3GPP PDP	3GPP TS23.002
R_m	HAG to 3GPP PDP	BBF TR-378
H_m	PGW to HAG	BBF TR-378

Table 1 Hybrid Access network architecture reference points summary

TR-348 describes three Hybrid Access transport models between HCPE and HAG:

- L3 Overlay tunneling
- L3 Network-based tunneling
- L4 Multipath

Section 5 of this Technical Report lists general solution and nodal requirements for Hybrid Access broadband networks. Sections 6, 7, and 8 define nodal requirements for each of these transport models. In addition, all architectural and high level nodal requirements in TR-348 are applicable to this Technical Report.

4.1 L3 Overlay Tunneling using GRE

As described in TR-348, the connectivity between the HCPE and the HAG for this transport model is established using Generic Routing Encapsulation (GRE) tunnels on top of the access infrastructure. The tunnels are established between the HCPE and the HAG over each of the access paths.

Figure 2 shows the logical architecture of this solution.

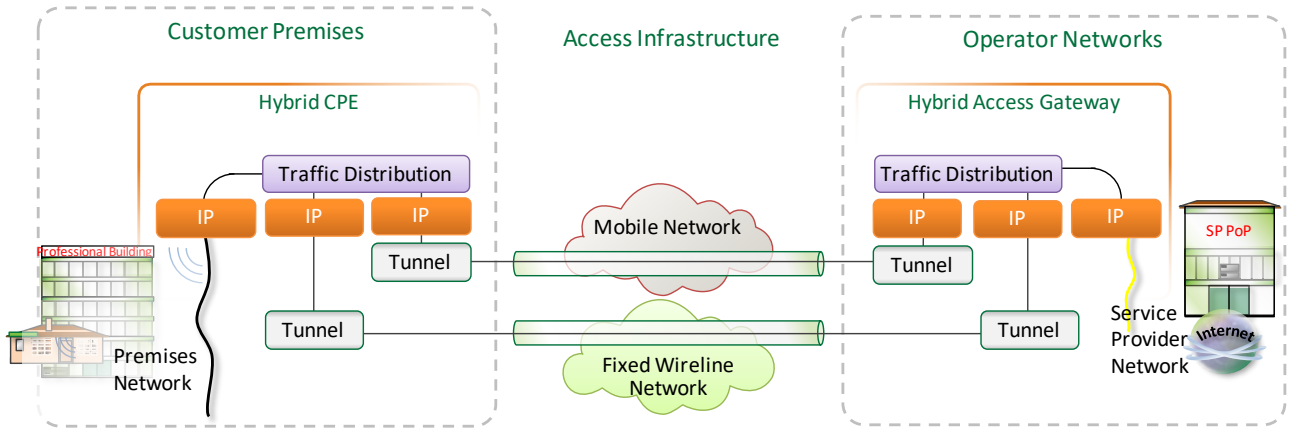

Figure 2 – L3 Overlay Tunneling

By default, 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, and can tear down the tunnels, in case of a network-based decision to terminate the service.

The L3 Overlay Tunneling transport model may operate in two different modes, depending on whether the Hybrid Access path over the 3GPP connection is kept up or torn down when the fixed broadband Hybrid Access path fails or is disabled. Whether the Hybrid Access path group operates in one mode or the other is a matter of policy, determined by the Service Provider.

The implementation itself is access network agnostic, therefore no changes to either the fixed broadband or the 3GPP access networks are necessary, with the possible exception of Maximum Transmission Unit (MTU) considerations, as described in section 4.1.5.

4.1.1 IPv4 Addressing

Figure 3 shows the end-to-end path between the HCPE and HAG. The HCPE fixed broadband interface is marked by “D” and the 3GPP interface is marked as “E”. Both interfaces are used as GRE tunnel endpoints in the HCPE. C is the service interface of the HCPE, and H is the GRE tunnel endpoint of the HAG. The customer Local Area Network (LAN) may use private IP addressing, in which case Network Address Translation (NAT)44 will be required at the HCPE, or public IP addressing, where NAT44 is not required. The “GRE Bonding” block comprises the policy enforcement and reordering.

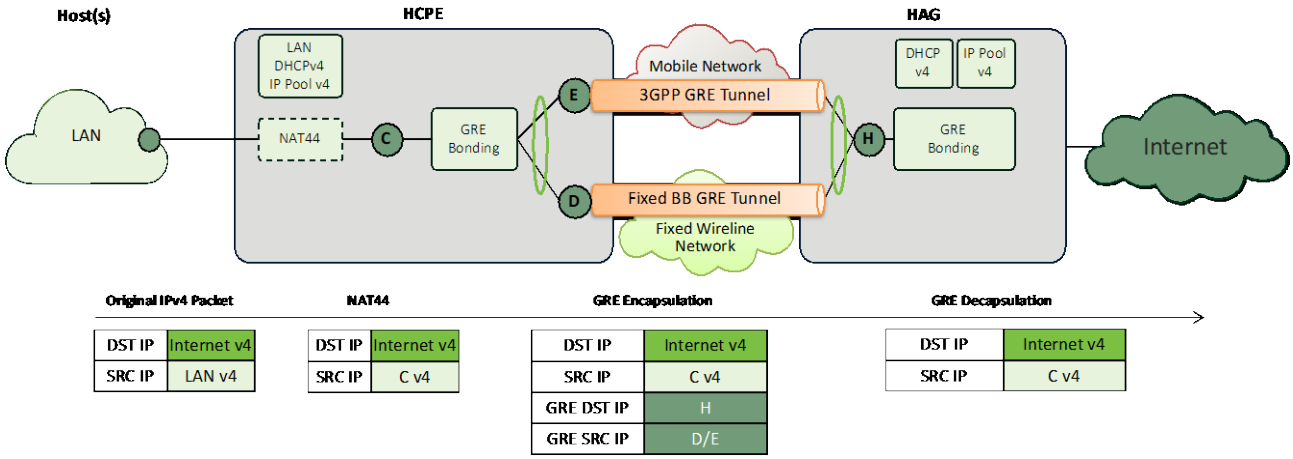

Figure 3 – IPv4 addressing for GRE-based L3 Overlay Tunneling

On the fixed broadband side, the MS-BNG assigns an IPv4 address to the HCPE fixed broadband Wide Area Network (WAN) interface during IPoE/PPPoE establishment. The subscriber's end device, connected to the LAN interface, gets an IPv4 address assigned via Dynamic Host Configuration Protocol (DHCP) by the HCPE. The subscriber may also assign addresses to the LAN devices manually. On the 3GPP access side, the PGW will assign an IPv4 address to the HCPE 3GPP WAN interface during PDP context establishment. Finally, the HAG acts as DHCPv4 server, assigning an IPv4 address to the HCPE to be used as the HCPE service IP (C), from one of the C address pools configured at the HAG.

Note that the use of public or private IPv4 addresses for each of the interfaces and subnets will be up to the Service Provider, based on the network architecture and requirements: e.g., private IPv4 LAN and NAT44 at the HCPE or NAT44 in the network, at HAG or separate platform.

Interface Name	Description	Assignment Mode	Device	Interface type
D	Fixed Broadband WAN Interface, GRE tunnel endpoint	Dynamic *IPoE/PPPoE*	HCPE	WAN
E	3GPP WAN Interface, GRE tunnel endpoint	Dynamic *Bearer Setup*	HCPE	WAN
C	HCPE service IP	Dynamic *DHCP from HAG (over tunnel)*	HCPE	WAN
H	HAG GRE tunnel endpoint	Static	HAG	WAN
LAN	LAN subnet(s)	Dynamic (DHCPv4) assigned by the HCPE or Static	Host	LAN

Table 2 IPv4 interfaces for GRE-based L3 Overlay Tunneling

Note that IP address C is not seen in the tunnel headers. In the case of using private IP addressing in the LAN and doing NAT44 at the HCPE, it is the source IP address of the packets going to Internet over the Hybrid Access path group.

For further details on IP address assignment for GRE-based L3 overlay tunneling, please see Section 6.1/RFC 8157 [10].

4.1.2 IPv4 NAT

After tunnel establishment, the HCPE uses DHCPv4 over the tunnel to get an IPv4 address from the HAG to be used for the Hybrid Access service (C). In those cases where the LAN subnet is using private IP addressing, the NAT function is expected to be present at the HCPE to translate the LAN private IPv4 addresses into the Hybrid Access service IPv4 address at HCPE (C).

However, alternatives are possible where NAT may be performed at the HAG or in a Carrier-Grade NAT (CGNAT) platform upstream of the HAG, which would allow flexible NAT models including 1:1, N:1 and N:M.

4.1.3 IPv6 Addressing

Figure 4 shows the IPv6 logical architecture of the GRE-based L3 Overlay Tunneling transport model. At the fixed broadband side, the MS-BNG assigns an IPv6 WAN prefix and an IPv6 Delegated Prefix to the HCPE. In addition, the HCPE uses DHCPv6 over the tunnel to get another IPv6 Delegated Prefix, assigned by the HAG.

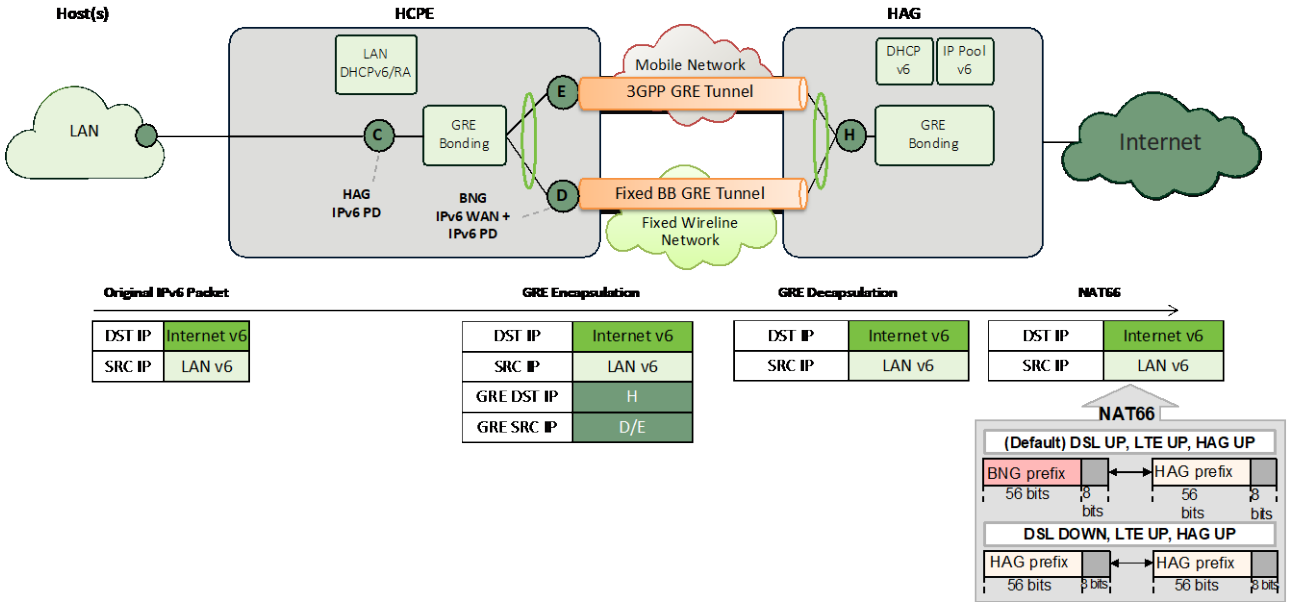

Figure 4 – IPv6 IP addressing for GRE-based L3 Overlay Tunneling

Table 3 lists the interfaces and their respective IPv6 prefixes.

Interface Name	Description	Assignment Mode	Device	Interface type
D	Fixed Broadband WAN Interface, GRE tunnel endpoint	Dynamic IPv6 WAN + IPv6 PD *IPoE/PPPoE*	HCPE	WAN
E	3GPP WAN Interface, GRE tunnel endpoint	Dynamic *Bearer Setup*	HCPE	WAN
C	HCPE service IP	Dynamic *DHCPv6 from HAG (over tunnel)*	HCPE	LAN
H	HAG GRE tunnel endpoint	Static	HAG	WAN
LAN	LAN subnet(s)	Dynamically (SLAAC/DHCPv6) assigned by the HCPE or Static	Host	LAN

Table 3 IPv6 interfaces for GRE-based L3 Overlay Tunneling

The IPv6 prefix on the HCPE fixed broadband (D) and 3GPP (E) WAN interfaces are to be used as the GRE tunnel endpoints.

4.1.4 IPv6 Delegated Prefixes and Prefix Translation

The IPv6 delegated prefixes are for IPv6 address allocation to LAN devices connected to the HCPE. Two prefixes are required, in order to allow attracting downstream traffic to two different nodes: the HAG for Hybrid Access traffic, and the MS-BNG for Hybrid Access bypass traffic. The use of two IPv6 prefixes for a given subscriber implies that IPv6 prefix translation will be required.

Note that Hybrid Access bypass traffic for the L3 Overlay Tunneling transport model never uses the 3GPP access interface.

Two alternatives are possible for the HCPE default IPv6 PD, depending on where the IPv6 translation will take place:

- IPv6 prefix translation at the HAG: the HCPE will use the IPv6 PD of the MS-BNG by default.
- IPv6 prefix translation at the HCPE: the HCPE will use the IPv6 PD of the HAG by default.

4.1.4.1 IPv6 Prefix Translation at the HAG

The HCPE uses the IPv6 PD received from the MS-BNG to allocate IPv6 addresses to LAN devices while the fixed broadband interface is active, which allows to directly forward Hybrid Access bypass traffic from the LAN to the fixed broadband WAN.

For traffic forwarded via the Hybrid Access path group, the HAG performs 1:1 IPv6 prefix translation from the MS-BNG IPv6 PD to the IPv6 PD allocated by the HAG.

Upon failure of the fixed broadband interface, the HCPE ages-out the IPv6 PD received from the MS-BNG and starts using the IPv6 PD received from the HAG for IPv6 address allocation to LAN devices. When the fixed broadband interface service is restored, the HCPE ages-out the IPv6 PD from the HAG and starts using the IPv6 PD received from the MS-BNG.

Note that in the case the Hybrid Access service is provided using unmanaged CPE, this option will give the Service Provider more control.

4.1.4.2 IPv6 Prefix Translation at the HCPE

The HCPE uses the IPv6 PD received from the HAG for IPv6 address allocation to LAN devices, which allows to directly forward Hybrid Access traffic over the tunnel and avoid IPv6 prefix translation at the HAG.

For Hybrid Access bypass traffic, the HCPE performs 1:1 IPv6 prefix translation from the LAN subnet, which is using the IPv6 PD allocated by the HAG, to the IPv6 PD allocated by the MS-BNG.

If the HCPE cannot establish connection to the HAG, it could use either the IPv6 PD allocated by the HAG or a Unique Local IPv6 Unicast prefix (RFC 4193) for IPv6 address allocation to LAN devices temporarily.

4.1.5 MTU Considerations

The maximum size of an IP packet that can be transported over the GRE tunnels depends on the MTU of the respective access networks and the overhead of the tunnel encapsulation.

Figure 5 shows an example of packet headers for the GRE-based L3 Overlay Tunneling transport model. As shown in the example (IPv4), the maximum size of an IPv4 packet that will be able to be transported through a GRE tunnel over the fixed broadband interface without fragmenting can be determined by subtracting the PPPoE (6 bytes), PPP (2 bytes), external IPv4 (20 bytes), and GRE (12 bytes) header sizes from the interface MTU. The Ethernet header is not shown in Figure 5.

For the 3GPP access, the external IP and GRE headers should be subtracted. Note that 3GPP traffic is encapsulated in General Packet Radio Service (GPRS) Tunnelling Protocol User Plane (GTP-U) which causes an overhead of 40 bytes (20 bytes IPv4 header, 8 bytes UDP header, and 12 bytes GTP-U header) not shown in the figure. If using IPsec, an additional 64 bytes of overhead must be considered.

Note that in the case of using IPv6 for the tunnel endpoints, the tunnel IPv6 header would amount to 40 bytes instead of 20 bytes, as is the case for an IPv4 header.

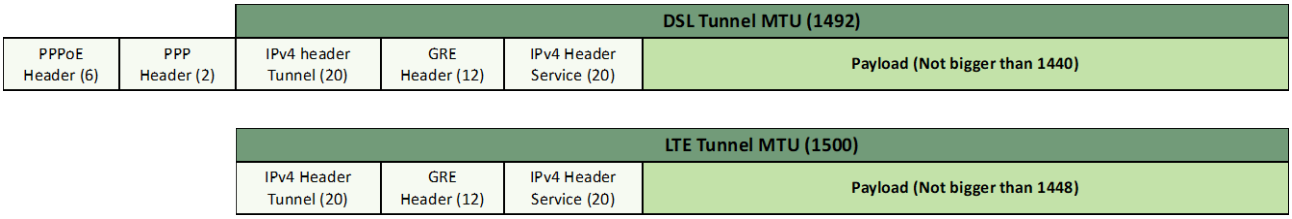

Figure 5 – Example frame encapsulations for GRE-based L3 Overlay Tunneling

Based on the above, the MTU of the GRE tunnel can be automatically calculated. However, it may also be set by means of configuration.

4.1.6 Authentication

An Authentication Server interface is used between the HAG and operator's backend servers for authentication and authorization. As shown in the Generic Hybrid Access network architecture and in Figure 6, the Authentication Server and the Policy server can be collapsed and run on the same platform. This allows policy information to be distributed to the HAG during the authentication phase, or at any time after that.

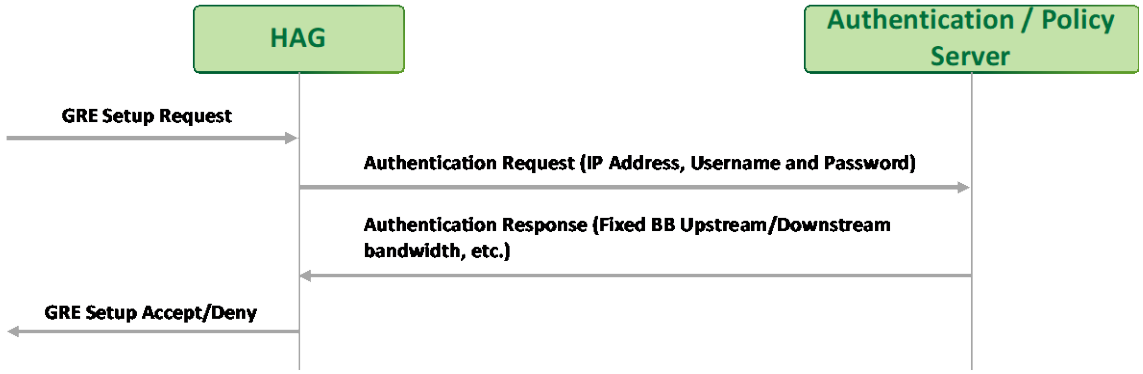

Figure 6 – GRE-based L3 Overlay Tunneling authentication

When the HCPE goes online, it initiates a tunnel establishment request to the HAG. The HAG sends the IP address, the user name and password carried in the request to the Authentication Server for user authentication. The Authentication Server sends an authentication response with the result of the operation, which may optionally include additional attributes, such as fixed broadband upstream/downstream bandwidth, etc. The HAG accepts or rejects the HCPE tunnel establishment request based on the result of the authentication phase.

Note that authentication is optional and may be enabled or disabled.

4.1.7 Bandwidth Reporting

The upstream and downstream fixed broadband WAN bandwidth values to be used by the HAG may be statically provisioned, or dynamically detected in the access network (using ANCP, PPP AVPs, or DHCP options) and reported to the Authentication/Policy Server, which in turn distributes them to the HAG.

For DSL access networks, the DSL sync rate can be used and is selected during the auto-negotiation process between the HCPE and the AN, and is reported to the network as described in Table 1/TR-147, Table 3/TR-101, and Table 4/TR-101. For GPON based access networks, similar parameters can be used, as described in Section 6.1.1.3/TR-207. The HAG reports the fixed broadband upstream and downstream rates to the HCPE in the Setup Accept message, as described in Section 5.2.10/RFC 8157.

The Hybrid Access bypass bandwidth is locally calculated at the HCPE based on statistics of non-GRE packets received over the fixed broadband interface and is reported to the HAG using a GRE Notify message, as described in Section 5.6.1/RFC 8157.

4.1.8 Traffic Classification

The coloring mechanism specified in RFC 2698 [7] is used to classify the customer's IP packets, both on upstream and downstream, to determine whether they should be forwarded over the fixed broadband or the 3GPP tunnels. Packets colored as green will be forwarded via the fixed broadband tunnel and packets colored as yellow will be forwarded via the 3GPP tunnel.

For upstream traffic, the Committed Information Rate (CIR) of the coloring mechanism of the HCPE is set to the bandwidth of the fixed broadband link minus the measured Hybrid Access bypass bandwidth consumed. For downstream traffic, the CIR of the coloring mechanism of the HAG is also set to the fixed broadband WAN bandwidth minus the bypassing fixed broadband bandwidth.

The mechanism above describes an implementation of a fixed broadband first traffic distribution scheme (e.g., for least-cost). Other traffic distribution schemes can be used, based on policy, as described in TR-348.

4.1.9 Traffic Recombination

The recombination function at the receiver is mandatory for per packet distribution based Hybrid Access. It provides for the in-order delivery of customers' traffic. As specified in RFC 2890 [8], the Sequence Number field in the GRE header is set for all packets distributed over the GRE tunnels of the Hybrid Access path group. In addition, the two GRE tunnels of the same Hybrid Access path group use the same Key value in the GRE header, which allows the receiver to correlate traffic from the Hybrid Access path group and reorder the packets according to the Sequence Number.

At the HCPE side, the Sequence Number field of the GRE header for upstream traffic is set based on incoming packet order to the outgoing queue buffer of the Hybrid Access path group interface.

The HAG classifies downstream traffic into hybrid access path groups. A sequence number per path group is applied prior to traffic distribution. The sequence number is encoded in the GRE header.

Reordering is not used for per-flow traffic distribution, as flow order is maintained by the nature of the traffic distribution mechanism.

4.1.10 Performance Measurement

The HCPE and HAG measure performance of the Hybrid Access paths by means of Hello and Notify messages, as described in RFC 8157.

The following are either exchanged or derived by means of the GRE protocol extensions defined in RFC 8157:

- Operating state of the hybrid access path groups
- The end-to-end delay and differential delay of the hybrid access path groups
- The maximum and available bandwidth of the hybrid access path groups
- Bypass bandwidth of a hybrid access path group

If the differential delay between the fixed broadband and 3GPP Hybrid Access paths becomes too large, throughput decreases and it could cause the overflow of the reordering buffers at the HCPE and/or HAG. To prevent this from happening, when the RTT difference exceeds a pre-defined threshold, traffic will be sent over only one of the Hybrid Access paths (over the GRE tunnel), based on policy. The other Hybrid Access path is not brought down but it is temporarily not used

until the RTT difference between both paths is below the threshold. The HAG controls this mechanism for downstream traffic and the HCPE is responsible in the upstream direction.

The HCPE or HAG first make a local decision to stop using one of the Hybrid Access paths and then notifies the other end to do so as well, using a GRE Notify message. When the RTT difference is within limits, a GRE message is sent once again to resume usage of the affected Hybrid Access path.

When either the fixed broadband or 3GPP access interfaces fail, the HCPE notifies the failure to the HAG using a GRE Notify message and the GRE tunnels are torn down by the HAG using a GRE Tear Down message, automatically disabling the Hybrid mode. In such event, all traffic is forwarded over the remaining access interface, if allowed by policy.

Further details of these mechanisms are described in Sections 5.6.3 and 5.6.4/RFC 8157.

4.2 L3 Network-based Tunneling using GTPv2

In this transport model, 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 MS-BNG and from HCPE to eNodeB respectively.

The following figure shows the high level logical architecture of this solution.

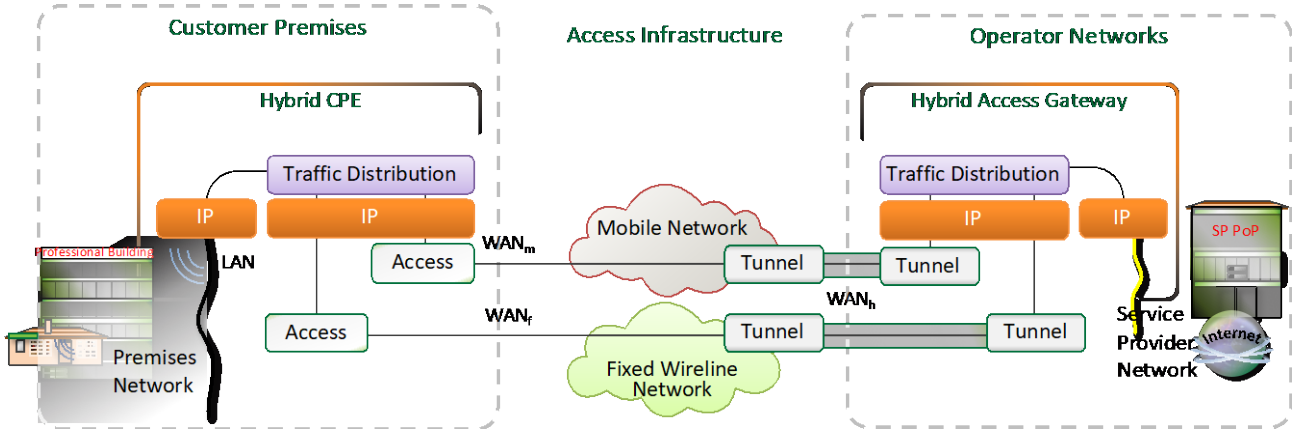

Figure 7 – L3 Network-based Tunneling

TR-378 supports two implementation options for L3 Network-based Tunneling, depending on where the HAG function resides:

- MS-BNG and HAG are integrated in a single node
- MS-BNG and HAG are in separate nodes

When the MS-BNG and HAG functions reside on separate nodes, the MS-BNG and eNB establish the tunnels to the HAG on behalf of the subscriber's HCPE and stitch 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 the context of TR-378, the tunnels established from the MS-BNG and eNodeB to the HAG use GTPv2, as per 3GPP TS 29.274 [5].

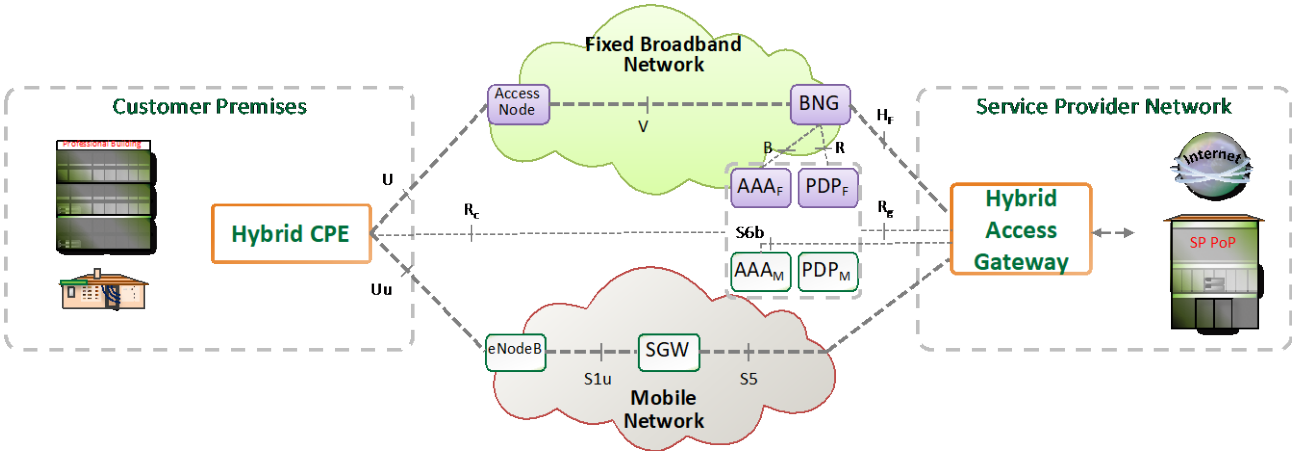

Figure 8 – L3 Network-based Tunneling with separate MS-BNG and HAG

Optionally, the Serving Gateway (SGW) function may also be integrated in the HAG (integrated S/PGW function), as shown in Figure 8:

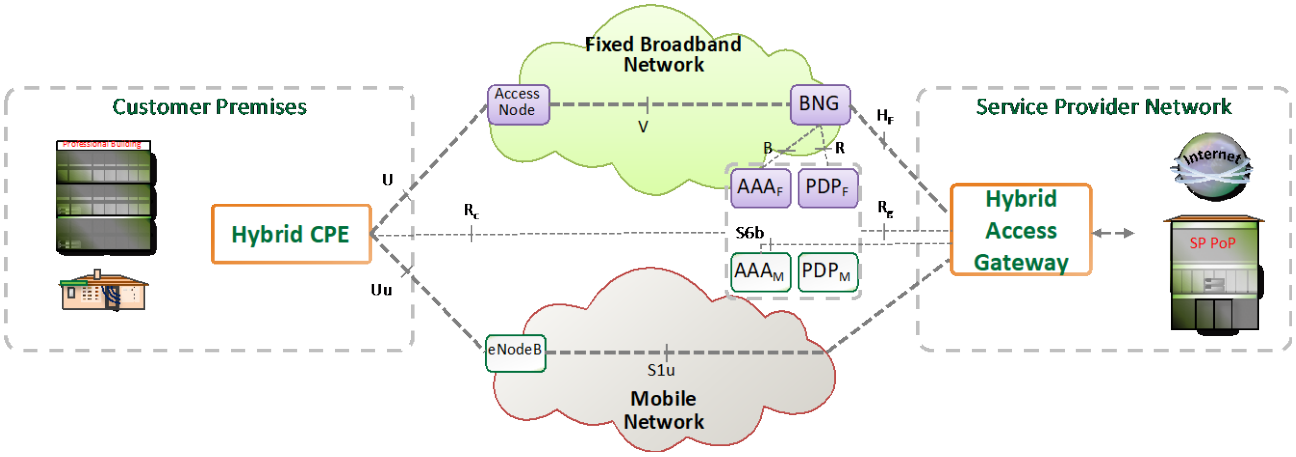

Figure 9 – L3 Network-based Tunneling with separate MS-BNG and HAG (integrated S/PGW)

In the case where the MS-BNG and HAG are in the same node, the HAG terminates the fixed broadband IP session (PPPoE or IPoE) and the 3GPP wireless session coming from the eNB (GTPv2) and treats them as a single Hybrid Access path group:

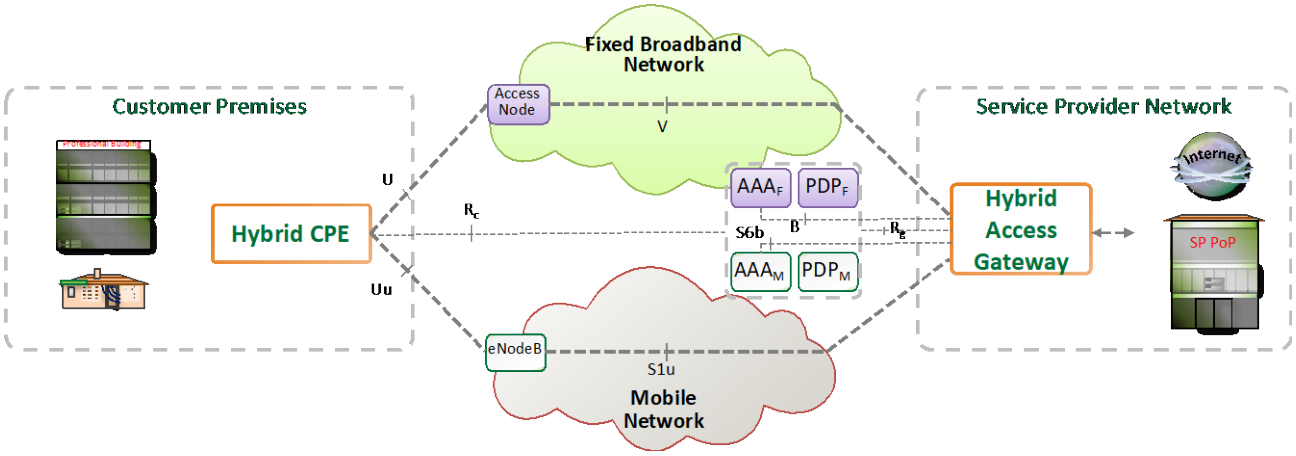

Figure 10 – L3 Network-based Tunneling with integrated MS-BNG and HAG (integrated S/PGW)

4.2.1 IP Addressing

In the network-based tunneling transport model, IP addresses are assigned to the HCPE fixed broadband and 3GPP access interfaces using the already existing IP allocation mechanisms in the respective networks, with support for dual-stack subscriber services.

Irrespective of whether the MS-BNG and HAG functions run in the same or separate nodes, IP address allocation is done by the HAG, which is responsible for the Hybrid Access path group.

For IPv4 addressing, the HCPE is assigned an IP address on the fixed broadband interface by starting an IPoE (DHCP) or PPPoE session, using existing address allocation procedures as described in existing Broadband Forum TRs (e.g., TR-101/TR-178). The 3GPP access interface of the HCPE is allocated an IPv4 address as part of the PDP procedures, as defined in existing 3GPP TSs.

The subnet behind the HCPE will normally use private IPv4 addressing. However, in the case a public IPv4 subnet is required behind the HCPE (e.g., for business services), the HCPE will be configured with an IPv4 prefix for local allocation and that same prefix will be activated via a Remote Authentication Dial-In User Service (RADIUS) Framed-Route in the HAG, pointing to the HCPE.

For IPv6 addressing, the HCPE is assigned either a /64 SLAAC address or a /128 DHCPv6 NA address for the fixed broadband interface, as described in TR-177 and TR-187, and a /64 prefix for the 3GPP access interface as part of the PDP procedures. For the subnet behind the HCPE, the HAG can assign an IPv6 Delegated Prefix to the HCPE using DHCPv6, which is used for address allocation within the premises. Alternatively, a Framed IPv6 Route for that prefix can be activated in the HAG, pointing to the HCPE. In the latter case the HCPE would be provisioned with the same prefix for local allocation.

In this transport model, the HCPE may use a single IP address for both Hybrid Access paths. The HAG can assign the same IP address to the HCPE via both access paths, reducing IP address consumption required for Hybrid Access to a minimum while also allowing seamless flow mobility between the two access paths.

Also note that in the case the MS-BNG and HAG functions are integrated, that Hybrid Access and bypass traffic share a common forwarding path, and as such NAT/prefix translation is not required at the HAG to attract downstream traffic, further reducing the IP address consumption required for Hybrid Access.

Interface Name	Description	IPv4 Addressing	IPv6 Addressing	Device	Interface type
WAN_f	Fixed Broadband WAN Interface	Dynamic *IPoE/PPPoE*	Dynamic *SLAAC/DHCPv6*	HCPE	WAN
WAN_m	3GPP WAN Interface	Dynamic *PDP*	Dynamic *PDP*	HCPE	WAN
WAN_h	HAG GTP tunnel endpoint	Static	Static	HAG	WAN
LAN	HCPE LAN Interface. Default gateway for LAN devices	Static	Dynamic *DHCPv6 PD* or Static	HCPE	LAN

Table 4 HCPE L3 Network-based Tunneling IP addressing

4.2.2 MTU Considerations

The network-based tunneling transport model makes use of the native technologies in both the fixed broadband and 3GPP access networks, without additional encapsulation. Therefore, this model does not introduce any additional MTU or fragmentation issues in the access network.

4.2.3 Traffic Distribution

The HAG spreads downstream traffic, by means of several different schemes:

- Weighted load-balancing
- Flow binding rules
- Active/standby
- Dynamic load-balancing, which attempts to saturate one link before using the other

Upstream traffic can enter through either connection, but it is recommended to keep flows identified by 5-tuple on the same link to avoid reordering.

By default, downstream traffic is hashed over two connections on a per-flow basis, allowing packets of the same flow to follow the same path and avoiding reordering issues. Flows are identified by the

5-tuple [src-ip, dst-ip, protocol, src-port, dst-port]. For non-L4 protocols, traffic is load-balanced with L3 criteria.

The default behavior can be altered based on policy to allow binding certain flows to a given Hybrid Access path.

Weights can be used to influence the hashing, allowing for percentage-based load balancing. In case of dynamic load-balancing, the weights can be dynamically adapted based on the load of the primary Hybrid Access path, enabling a least-cost first traffic distribution scheme.

4.2.4 Traffic Recombination

By default, the solution distributes traffic between the Hybrid Access paths on a per-flow basis. As such, there is no need for traffic recombination.

Per-packet traffic distribution is also possible, by leaving the re-ordering of traffic to the end hosts. If this effect is not desired, Multipath Transmission Control Protocol (MPTCP) can be run on top of L3 Network-based tunneling.

4.3 L4 Multipath using MPTCP

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.

The following figure shows the logical architecture of this solution.

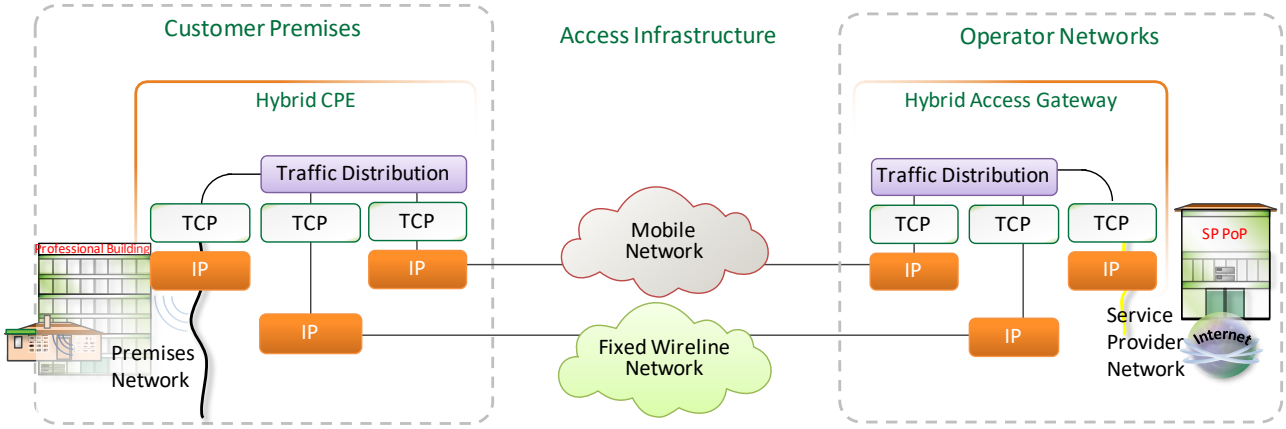

Figure 11 – L4 Multipath network using MPTCP

In the context of TR-378, the L4 multipath transport service uses MPTCP. It sets up multiple Transmission Control Protocol (TCP) subflows over the different access networks and utilizes real time HCPE to HAG flow control, using standard TCP flow and congestion control schemes. The HCPE and HAG are responsible for managing the MPTCP Hybrid Access paths, including establishment and tear down.

The HCPE and HAG may terminate the end user layer 4 sessions before transporting the data over the Hybrid Access paths, effectively executing a proxy function for these end user sessions. Policy for selecting which sessions will use MPTCP is vendor specific.

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

4.3.1 MPTCP Transport Models

An MPTCP-based Hybrid Access path group can be established between the HCPE and the HAG using either MPTCP Explicit Mode, as per draft-ietf-tcpm-converters [11] or MPTCP Implicit Mode as described in section 4.3.3 below.

The main advantages of MPTCP based Hybrid Access are:

- No additional tunneling. Minimal additional overhead
- No out-of-band signaling for each MPTCP subflow
- Uses standard TCP techniques for congestion detection and control
- Accommodates various deployment contexts (e.g., address sharing, preserve the external IP address, IPv6 addressing, etc.)

Note that only TCP traffic may benefit of distribution over an MPTCP Hybrid Access path group. Non-TCP traffic will be forwarded over one of the Hybrid Access paths.

HAG Deployment Scenarios:

- **On-path:** Requires the HAG to be on the forwarding path between the HCPE and the network. HCPE addressing for the HAG may be implicit (transparent) or explicit. In an on-path deployment model, all traffic for that path will pass through the HAG, irrespective of the destination IP address of the packets.

- **Off-path:** The HAG may be anywhere in the network, reachable by the HCPE using explicit proxy addresses. In an off-path deployment model, by default packets do not go through the HAG in either direction.

4.3.2 IP Addressing

IP addresses are assigned to the HCPE fixed broadband and 3GPP access interfaces using the already existing IP allocation mechanisms in the respective networks, with support for dual-stack subscriber services.

For IPv4 addressing, the HCPE is assigned an IP address on the fixed broadband interface by starting an IPoE (DHCP) or PPPoE session, using existing address allocation procedures as described in existing Broadband Forum TRs (e.g., TR-101/TR-178). The 3GPP access interface of the HCPE is allocated an IPv4 address as part of the PDP procedures, as defined in existing 3GPP TSs.

For IPv4, the subnet behind the HCPE will normally use private IPv4 addressing. However, alternative models are possible. For example, in the case a public IPv4 subnet is required behind the HCPE (e.g., for business services), the HCPE may be configured with an IPv4 prefix for local allocation and have that same prefix be activated (e.g., via a RADIUS Framed-Route) in the network, pointing to the HCPE.

For IPv6 addressing, the HCPE is assigned either a /64 SLAAC address or a /128 DHCPv6 NA address for the fixed broadband interface, as described in TR-177 and TR-187, and a /64 prefix for the 3GPP access interface as part of the PDP procedures. For the subnet behind the HCPE, the fixed network can assign an IPv6 Delegated Prefix to the HCPE using DHCPv6, used for address allocation within the premise. Other options, such as using Unique Local IPv6 Unicast addressing (RFC 4193) for the HCPE LAN, may also be used. Note that this would require prefix and port translation between the ULA prefix and the fixed broadband or 3GPP interfaces, for Hybrid Access bypass traffic.

Since IPv6 prefix delegation is not currently supported in 3GPP networks, only a /64 (SLAAC) is allocated by the 3GPP network to the HCPE during bearer activation. As such, for subflows over the 3GPP interface, the HCPE must do prefix address and port translation from the fixed network PD address to the /64 allocated to the HCPE 3GPP interface, while preventing multiple client addresses from overlapping on the 3GPP interface. This also applies to Hybrid Access bypass traffic over the 3GPP interface.

4.3.3 MPTCP Explicit Mode

In MPTCP explicit mode, the HAG may be deployed on-path or off-path. In this mode, the HCPE is aware of the presence of the HAG in the network and sends the upstream traffic received from the LAN that requires aggregation to the HAG. This is done by the HCPE forwarding the traffic subject to Hybrid Access to the proxy addresses of the HAG that are determined for both Hybrid Access paths.

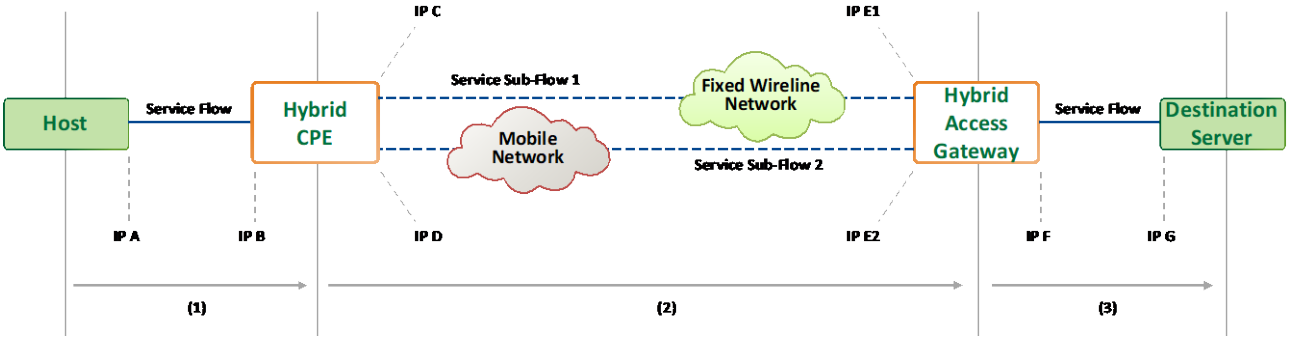

Figure 12 – HAG addressing using MPTCP explicit mode

The diagram presented at Figure 12 has the following upstream packet flow:

(1) Traffic corresponding to a certain service being provided to the End-User is sent to the public IP address (IP G) of a Destination Server through the customer premises default gateway, the HCPE.

(2) Based on the traffic distribution policy, the HCPE distributes traffic that requires aggregation at the HAG among Hybrid Access paths. Traffic that follows this policy is forwarded through Hybrid Access paths over the fixed broadband and 3GPP access networks to the HAG, which address E1 is previously provisioned at the HCPE. Address E2 may be provisioned at the HCPE or may be learned by using the MPTCP ADD_ADDR option provided by the HAG over a previously established Hybrid Access path.

(3) After the MPTCP subflow aggregation, the HAG forwards the traffic to the Destination Server, as originally intended.

4.3.4 MPTCP Implicit Mode

For implicit mode HAG deployments, the HAG must be deployed on-path for the primary path and may be on-path or off-path for the secondary path. Implicit HAG addressing is used for the primary path, and either implicit or explicit HAG addressing can be used for the secondary path. Explicit addressing is required in the case the HAG is off-path for the secondary path.

The HCPE may be configured with the proxy address of a HAG, or the HAG proxy address may be learned from the HAG using the MPTCP ADD_ADDR option. The HCPE should implement a mechanism that allows the HAG to differentiate between MPTCP sessions originated by the HCPE from MPTCP session originated from a client.

In MPTCP implicit mode, HCPE network addresses are configured as two unique IP addresses C and D on a dual IP HCPE, and the HAG is deployed on-path of the primary path, as shown in Figure 13. This avoids HCPE configuration with proxy IP addresses for the primary path, as the HAG proxy address (IP E) to be used by the HCPE the secondary path is provided by the HAG to the HCPE using a MPTCP option.

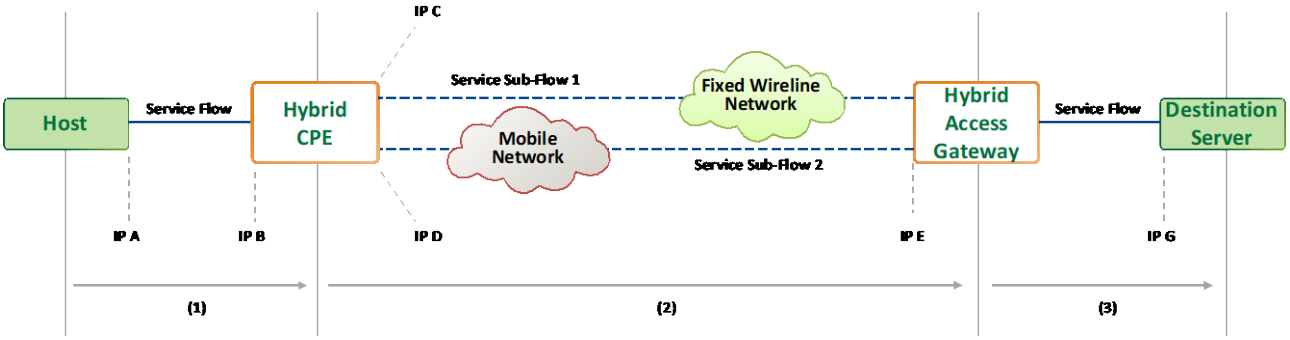

Figure 13 – HAG addressing using MPTCP implicit mode

As an example, flow establishment for the upstream direction would be as follows:

(1) Traffic corresponding to a certain service being provided to the End-User is sent to the public IP address (IP G) of a Destination Server through the customer premises default gateway, the HCPE.

(2) Initial TCP session setup establishes the first MPTCP subflow from address C, with MPTCP indicated in the TCP option field. This Subflow is transparent (no 5-tuple addressing change to destination G). The HCPE then initiates a TCP subflow on the secondary path from Address D to the HAG using destination address IP E, provided during the setup of the first subflow. The HAG MPTCP proxy uses MPTCP session ID to correlate the two subflows as one session.

Based on the traffic distribution policy, the HCPE and HAG distribute traffic that requires forwarding over both Hybrid Access paths.

(3) After the MPTCP subflow aggregation, the HAG forwards the traffic to the Destination Server, as originally intended

In the downstream (server to user) direction, transparent mode means that the HAG proxies are completely transparent to the network side; all traffic is destined to the HCPE IP address used on the transparent path, (IP C in this example). This avoids the need to implement any mechanism to attract traffic to the appropriate HAG. The HAG implements the downstream path distribution policy using subflow addresses C and D as needed.

4.3.4.1 Primary Path MPTCP Subflow Establishment

The HCPE must have addresses on its primary and secondary links (respectively named C and D in the example below), while the HAG requires an address (E) on its secondary link.

The HCPE maps a client-originated TCP connection onto a Hybrid Access MPTCP connection (and its associated subflows).

The client sends a TCP SYN segment addressed to the server. The TCP SYN segment is intercepted by the HCPE which in turn initiates an MPTCP connection on the primary path towards its on-path HAG. The HCPE creates a flow entry for the HAG connection and maps the Client side TCP connection onto the WAN side MPTCP connection. The destination address of the TCP SYN segment is the IP address of the Server, since presence of the HAG is transparent on the primary path. Over the primary access network, this TCP SYN appears as originating from the host IP address and being sent to IP address G.

The HAG acts as a transparent proxy for IP address G and intercepts the TCP SYN that contains the MP_CAPABLE option. It creates state for the MPTCP connection and initiates a TCP connection towards IP address G. The HAG does not perform source address translation. The remote server receives the TCP SYN as originating from the host address.

The HCPE considers the MPTCP connection to be active upon reception of the SYN+ACK segment from the HAG. The reception of this segment triggers the HCPE to confirm the establishment of the connection by sending a SYN+ACK segment towards the TCP Client. At this point, there are two established connections maintained by the HCPE:

1. LAN TCP connection: The endpoints are the Client and the HCPE.
2. Hybrid Access MPTCP connection: The endpoints are the HCPE and the HAG.

These two connections are bound by the HCPE. An example is shown in Figure 14.

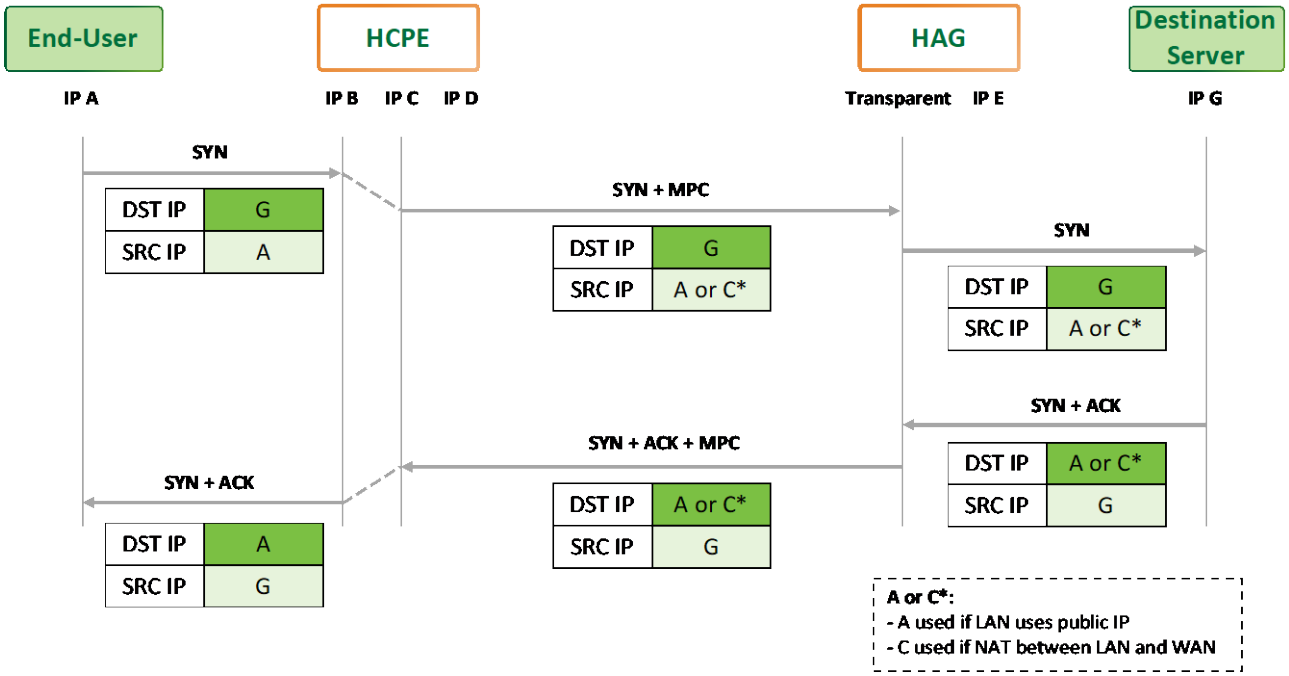

Figure 14 – Creation of the initial subflow with implicit mode

4.3.4.2 Secondary Path MPTCP Subflow Establishment

For implicit mode where the HAG is on-path for the primary link and off-path for the secondary link, the HCPE must be able to learn one address of the off-path interface of the HAG. This address (E) can be configured statically, dynamically distributed by means of a DHCP option, or provided to the HCPE from the HAG using MPTCP ADD_ADDR option, over the primary subflow.

In parallel, the HCPE should advertise to the HAG the IP address D it will use for the secondary subflow by sending an ADD_ADDR option on the primary subflow. Subsequently, the HCPE will establish an additional subflow from the HCPE over the second access network (arrows (2), (3), and (4) in Figure 14). The endpoints of this subflow are the IP address of the HCPE on the second access network, i.e., IP address D, and the IP address of the HAG, i.e., IP address E.

Note that the ADD_ADDR options shown in Figure 15 are optional. The HAG may or may not require knowledge of HCPE secondary path IP address D prior to accepting the subflow on the secondary path. If the HCPE already knows, e.g., by configuration or through other mechanisms, the IP address of the HAG, it can create the additional subflow without waiting for the ADD_ADDR option from the HAG containing its address on the secondary path.

If the HAG is also on-path for the secondary path, neither the HAG nor the HCPE need to advertise additional addresses to each other using the MPTCP ADD_ADDR option. The second sub-flow can also use transparent / implicit addressing.

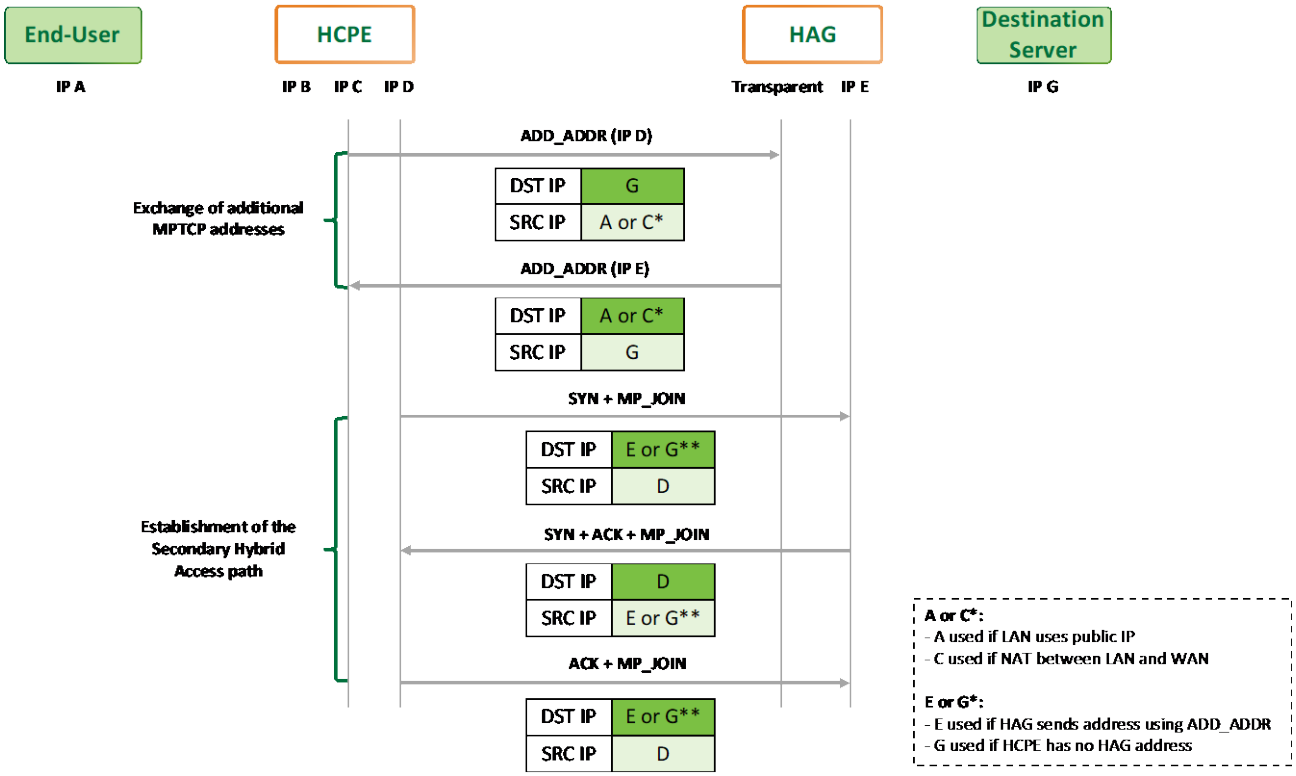

Figure 15 – Example creation of the second subflow by the HCPE with MPTCP implicit mode

4.3.5 Traffic Distribution

Per TR-348, the HCPE and HAG classify traffic into a set of HA classes, which are then forwarded as defined by policy.

For regular traffic, i.e., traffic subject to Hybrid Access traffic distribution, the HCPE and HAG terminate the customer side and network side L4 sessions respectively and forward the traffic over the MPTCP subflow corresponding to the selected Hybrid Access path for a given packet or flow, after performing a proxy function.

For Hybrid Access bypass traffic, the HCPE forwards the selected traffic to either the fixed broadband or 3GPP access interfaces, as dictated by policy.

When the performance of a Hybrid Access path is degraded (e.g., high packet loss, high RTT, low bandwidth), its capabilities are monitored and then considered by the Traffic Distribution function. Based on the new conditions, the Traffic Distribution function may then forward some or all of the existing flows to the other Hybrid Access path, depending on its current performance and the traffic distribution policy in place. The HCPE and the HAG are responsible for managing the dynamic changes to the traffic distribution for upstream and downstream traffic respectively. This process does not involve changes in IP routing.

4.3.6 Performance Measurement

The Hybrid Access performance measurement function estimates the service availability, bandwidth, and delay metrics as well as their variations by monitoring several TCP parameters of the TCP subflows over each Hybrid Access path, such as the RTT, packet loss ratio, etc.

The Hybrid Access performance measurement is agnostic of the underlying network, passive, and based on TCP connection data. There is no additional traffic generated, nor a need for additional specific equipment polling or using a specific path monitoring protocol.

4.4 Access Network Dynamic Rate Changes

Some fixed access network technologies support dynamic bandwidth control by means of Seamless Rate Adaptation (SRA), which may cause bandwidth changes from the CPE to the Access Node to occur dynamically. In the case of ADSL and VDSL, this mechanism can be activated or deactivated, depending on the Service Provider preference. In the case of G.fast, this mechanism is recommended in all cases.

The L3 Overlay Tunneling and the L3 Network-based Tunneling transport models supported in TR-378 require the HAG to be aware of the bandwidth available in the access network so that accurate traffic control and distribution policies can be enforced accurately. Therefore, a mechanism is required to convey information about access network bandwidth changes to the policy server. Note that the L4 Multipath MPTCP transport model dynamically adapts to the available bandwidth, and does not require this information.

The problem described here is not specific to Hybrid Access broadband networks and applies also to TR-101/TR-178 architectures, where the MS-BNG may need to be aware of the access network bandwidth to enforce accurate QoS policies to a given subscriber session.

As such, the definition of such mechanism is outside of the scope of TR-378, as it is not specific to Hybrid Access broadband networks.

5 General Nodal Requirements for Hybrid Access Broadband Networks

As mentioned in Section 4, all architectural and high level nodal requirements in TR-348 are applicable and imported by reference to this Technical Report.

In addition, the following general requirements apply to HCPE and HAG for all TR-378 transport models.

5.1 HCPE Requirements

5.1.1 Access Connectivity

[R-1] The HCPE MUST support PPPoE-based fixed access connectivity, as per TR-101 and TR-187.

[R-2] The HCPE MUST support IPoE-based fixed access connectivity, as per TR-101 and TR-177.

[R-3] The HCPE MUST support 3GPP wireless connectivity, as per 3GPP TS 23.401.

5.1.2 Performance Measurement

[R-4] The HCPE MUST be able to adjust the traffic classification and distribution according to the performance measurement results.

[R-5] When the Hybrid mode is disabled, the HCPE MUST only forward traffic over one of the access interfaces, based on policy.

5.2 HAG Requirements

5.2.1 Performance Measurement

[R-6] The HAG MUST be able to adjust the traffic classification and distribution according to the performance measurement results.

5.2.2 Charging and Billing Requirements

[R-7] The HAG MUST support at least one of the following:

- AAA-based charging
- 3GPP Gy/Gz-based charging
- 3GPP Gx usage-based charging

6 Nodal Requirements for L3 Overlay Tunneling

This section lists the set of requirements for network elements implementing the L3 overlay tunneling Hybrid Access transport model, shown in Figure 2.

6.1 HCPE Requirements

[R-8] The HCPE MUST support the control plane procedures for L3 GRE Overlay, as per Section 5/RFC8157.
[R-9] The HCPE MUST support the data plane procedures for L3 GRE Overlay, as per Section 6/RFC8157.
[R-10] If the GRE tunnel towards the HAG cannot be established, the HCPE MUST forward all traffic according to a predefined policy.

6.1.1 IPv4 Addressing

[R-11] The HCPE MUST support GRE over IPv4 tunnels.
[R-12] The HCPE MUST support transporting IPv4 traffic over GRE tunnels.
[R-13] The HCPE MUST support the IPv4 addressing scheme in Table 2.

6.1.2 IPv4 NAT

[R-14] The HCPE MUST support IPv4 Network Address Port Translation (NAPT).

If NAPT is enabled at the HCPE:
[R-15] For IPv4 Hybrid Access traffic, the HCPE MUST perform NAPT between the LAN subnet and the HCPE Service IP (C) before forwarding the traffic over the GRE tunnel.
[R-16] For IPv4 Hybrid Access bypass traffic, the HCPE MUST perform NAPT between the LAN subnet and the IP address of the selected outgoing WAN interface IP of the HCPE.

6.1.3 IPv6 Addressing

[R-17] The HCPE MUST support GRE over IPv6 tunnels.
[R-18] The HCPE MUST support transporting IPv6 traffic over GRE tunnels.
[R-19] The HCPE MUST support the IPv6 addressing scheme in Table 3.
[R-20] The HCPE MUST route traffic to the fixed broadband WAN interface for IPv6 Hybrid Access bypass traffic.

6.1.4 IPv6 Prefix Translation

[R-21] The HCPE MUST support IPv6 prefix translation.

If IPv6 prefix translation is enabled at the HCPE:
[R-22] For Hybrid Access bypass traffic, the HCPE MUST perform 1:1 IPv6 prefix translation between the prefix used in the HCPE LAN (HAG PD) to the MS-BNG PD, prior to forwarding traffic over the fixed broadband interface.

6.1.5 MTU Considerations

[R-23] The HCPE MUST support automatically deriving the MTU of the tunnel based on the interface MTU.
[R-24] The HCPE SHOULD support configuration of the tunnel MTU.
[R-25] The HCPE MUST support fragmenting IPv4 packets whose size is greater than the MTU of the outgoing tunnel interface.

Note that in IPv6 networks only the end nodes do the fragmentation, not the routers.

6.1.6 Performance Measurement

In addition to the general performance measurement requirements in Section 5.1.2, the following requirements apply to a HCPE implementing L3 Overlay Tunneling transport model:

[R-26] The HCPE MUST support performance measurement of the Hybrid Access path group using the procedures defined in Sections 4.6 and 5/RFC 8157.
[R-27] The HCPE SHOULD be able to deactivate the hybrid mode when difference in RTT between the Hybrid Access paths exceeds a threshold defined by the service provider.

6.1.7 Traffic Recombination

[R-28] The HCPE MUST use the same value in the GRE header Key field for both GRE tunnels in the Hybrid Access path group.
[R-29] The HCPE MUST set the GRE header Sequence Number field based on incoming packet order to the Hybrid Access path group egress queue buffer.
[R-30] The HCPE MUST support downstream packet re-ordering for all packets with the same GRE Key, based on the GRE header Sequence Number.

6.2 HAG Requirements

[R-31] The HAG MUST support the control plane procedures for L3 GRE Overlay, as per Section 5/RFC8157.
[R-32] The HAG MUST support the data plane procedures for L3 GRE Overlay, as per Section 6/RFC8157.

6.2.1 IPv4 Addressing

[R-33] The HAG MUST support GRE over IPv4 tunnels.
[R-34] The HAG MUST support transporting IPv4 traffic over GRE tunnels.
[R-35] The HAG MUST support the IPv4 addressing scheme in Table 2.

6.2.2 IPv4 NAT

[R-36] The HAG MUST support IPv4 1:1 NAT.
[R-37] The HAG MUST support IPv4 N:1 NAPT.

[R-38] The HAG MUST support IPv4 N:M NAPT.

If NAT is enabled at the HAG:
[R-39] The HAG MUST perform NAT for traffic received from the HCPE, using the selected NAT model by the Service Provider.

6.2.3 IPv6 Addressing

[R-40] The HAG MUST support GRE over IPv6 tunnels.
[R-41] The HAG MUST support transporting IPv6 traffic over GRE tunnels.
[R-42] The HAG MUST support the IPv6 addressing scheme in Table 3.

6.2.4 IPv6 Prefix Translation

[R-43] The HAG MUST support IPv6 prefix translation.

If IPv6 prefix translation is enabled at the HAG:
[R-44] The HAG MUST perform prefix translation between the IPv6 prefix used in the HCPE LAN (MS-BNG PD) to the IPv6 prefix reserved at the HAG for that customer.

6.2.5 MTU Considerations

[R-45] The HAG MUST support automatically deriving the MTU of the tunnel based on the interface MTU.
[R-46] The HAG SHOULD support configuration of the tunnel MTU.
[R-47] The HAG MUST support fragmenting IPv4 packets whose size is greater than the MTU of the outgoing tunnel interface.

Note that in IPv6 networks only the end nodes do the fragmentation, not the routers.

6.2.6 Traffic Classification

[R-48] The HAG MUST support static configuration of the downstream fixed broadband bandwidth.
[R-49] The HAG MUST support receiving the configuration of the downstream fixed broadband bandwidth from the policy server.
[R-50] The HAG MUST be able to adjust the downstream bandwidth of the fixed broadband Hybrid Access path using the values reported by the HCPE, as per Section 5.1.3/RFC 8157.

6.2.7 Performance Measurement

In addition to the general performance measurement requirements in Section 5.2.1, the following requirements apply to a HAG implementing L3 Overlay Tunneling transport model:

[R-51] The HAG MUST support performance measurement of the Hybrid Access path group using the procedures defined in Sections 4.6 and 5/RFC 8157.

[R-52] The HAG SHOULD be able to deactivate the hybrid mode when difference in RTT between the Hybrid Access paths exceeds a threshold defined by the service provider.

6.2.8 Traffic Recombination

[R-53] The HAG MUST use the same value in the GRE header Key field for both GRE tunnels in the Hybrid Access path group.
[R-54] The HAG MUST set the GRE header Sequence Number field based on incoming packet order to the Hybrid Access path group egress queue buffer.
[R-55] The HAG MUST support upstream packet re-ordering for all packets with the same GRE Key, based on the GRE header Sequence Number.

6.3 Policy Control Requirements

[R-56] The HAG MUST support the R_g interface to the fixed broadband PDP and 3GPP PDP.
[R-57] The HAG MUST support RADIUS for the R_g interface.
[R-58] The HAG SHOULD support SOAP for the R_g interface.

6.4 Charging and Billing Requirements

In addition to the general charging and billing requirements in Section 5.2.2, the following requirement applies to a HAG implementing L3 Overlay Tunneling transport model.

If the Hybrid Access path group is set to tear down the 3GPP Hybrid Access path upon failure or deactivation of the fixed broadband Hybrid Access path, then:
[R-59] The HAG MUST be able to stop accounting if the fixed line fails or is disabled.

7 Nodal Requirements for L3 Network-based Tunneling

7.1.1 HCPE Requirements

7.1.1.1 IPv4 Addressing

[R-60] The HCPE MUST support using a single IPv4 WAN address on both Hybrid Access paths.
[R-61] The HCPE MUST support using different IPv4 WAN addresses on the two Hybrid Access paths.
[R-62] The HCPE MUST support NAT (NAPT) between the private IPv4 prefix of the LAN and the IPv4 WAN address of the Hybrid Access path group.

7.1.1.2 IPv6 Addressing

[R-63] The HCPE MUST support using a single IPv6 WAN prefix on both Hybrid Access paths.
[R-64] The HCPE MUST support using different IPv6 WAN prefixes on the two Hybrid Access paths.
[R-65] The HCPE MUST support receiving an IPv6 PD from the HAG, for allocation to LAN devices.

7.1.1.3 Traffic Distribution

[R-66] The HCPE MUST support using hashing to determine the Hybrid Access path for a given traffic flow on upstream.
[R-67] The HCPE MUST support binding a 5-tuple to a Hybrid Access path, based on policy.
[R-68] The HCPE SHOULD support switching the upstream traffic to the Hybrid Access path chosen by the HAG for downstream traffic, when that path changes.
[R-69] In case of failure of a Hybrid Access path, the HCPE MUST support automatically switching traffic to the remaining path, if allowed by policy.

7.1.2 HAG Requirements

7.1.2.1 Connectivity

[R-70] The HAG MUST be able to associate the two Hybrid Access paths by virtue of a common subscriber identifier.

The following requirements apply when the HAG function is integrated with the MS-BNG:

[R-71] The HAG MUST support terminating a PPPoE session from the HCPE.
[R-72] The HAG MUST support terminating an IPoE session from the HCPE.

The following requirements apply when the HAG function resides in a separate node from the MS-BNG:

[R-73] The HAG MUST support terminating a GTPv2 session from the MS-BNG.
[R-74] The HAG SHOULD support terminating a stateless GRE tunnel from the HCPE.

The following requirements apply to the HAG to support establishment of Hybrid Access paths over the wireless network:

[R-75] The HAG MUST support the S1-U interface to the eNodeB, for direct connectivity.
[R-76] The HAG SHOULD support the S5 interface to the SGW, for indirect connectivity.
[R-77] The HAG MUST support the S11 interface to the MME.

7.1.2.2 IPv4 Addressing

[R-78] The HAG MUST support allocating the same IPv4 WAN address to the HCPE over both Hybrid Access paths.
[R-79] The HAG MUST support allocating different IPv4 WAN addresses to the HCPE over the two Hybrid Access paths.

If the customer has a public IP subnet allocated for the LAN of the HCPE:

[R-80] The HAG MUST support routing traffic to a public subnet in the HCPE LAN.
[R-81] The HAG MUST support activation of a route to the HCPE LAN upon reception of a RADIUS Framed-Route.

7.1.2.3 IPv6 Addressing

[R-82] The HAG MUST support allocating the same IPv6 WAN prefix to the HCPE over both Hybrid Access paths.
[R-83] The HAG MUST support allocating different IPv6 WAN prefixes to the HCPE over the two Hybrid Access paths.
[R-84] The HAG MUST support allocating an IPv6 PD to the HCPE, for address allocation to LAN devices.

7.1.2.4 Traffic Distribution

[R-85] The HAG MUST be able to distribute the downstream traffic received for a given HCPE over both Hybrid Access paths, using traffic distribution policies.
[R-86] The HAG MUST be able to do percentage-based distribution of traffic on a per flow basis for a given policy.
[R-87] The HAG MUST support binding a 5-tuple to a Hybrid Access path, based on policy.
[R-88] The HAG MUST support policy to fill a specific link first and then utilize another one.
[R-89] The HAG MUST be able to move flows from one link to another due to congestion or quality of the link.

7.1.3 Policy Control Requirements

[R-90] The HAG MUST support DIAMETER Gx for the Rg interface.
[R-91] The HAG MUST support RADIUS for the Rg interface.

[R-92] The HAG SHOULD support Network based IFOM (NBIFOM) for application of policies both at HAG and HCPE, as per 3GPP Rel. 13.

8 Nodal Requirements for L4 Multipath

8.1 MPTCP Plain Mode

8.1.1 HCPE Requirements

[R-93] The HCPE MUST support MPTCP, per RFC 6824 [9].
[R-94] The HCPE MUST be able to act as a Converter Client, turning regular TCP connections coming from the LAN into MPTCP connections to the HAG (Transport Converter), supporting draft-ietf-tcpm-converters [11].
[R-95] The HCPE MUST be able to disable/enable its Converter Client function, via device management, in support of Section 6.1.2/TR-348 [4].
[R-96] The HCPE MUST be able to carry TCP traffic over the MPTCP Hybrid Access path group.
[R-97] The HCPE MUST forward any non-TCP traffic outside of the MPTCP Hybrid Access path group, over one of the access interfaces, as directed by policy.
[R-98] The HCPE MUST support TCP Fast Open, per RFC 7413.
[R-99] The HCPE MUST support relaying the original destination IP address and port using TCP Fast Open (TFO), as per draft-ietf-tcpm-converters.
[R-100] The HCPE MUST simultaneously support both IPv4 and IPv6 MPTCP subflows.
[R-101] The HCPE SHOULD support sending the HCPE MPTCP termination IP address(es) to the HAG by means of the ADD_ADDR TCP option.
[R-102] The HCPE MUST support receiving the HAG MPTCP termination IP address(es) by means of the ADD_ADDR TCP option.
[R-103] The HCPE MUST support the configuration of the HAG MPTCP termination IP address(es) via device management.
[R-104] The HCPE SHOULD support receiving the HAG MPTCP termination IP address(es) by means of a DHCP option.
[R-105] The HCPE SHOULD support configurable MPTCP congestion control schemes and packet schedulers.
[R-106] The HCPE SHOULD be able to be configured to disable the MPTCP checksum procedures.
[R-107] If the HCPE cannot establish MPTCP subflows to the HAG, it MUST fallback to ordinary TCP operation.
[R-108] The HCPE SHOULD support selecting which TCP flows are subject to MPTCP-based Hybrid Access, as determined by policy.
[R-109] The HCPE SHOULD support gathering statistics and event logging for the MPTCP Proxy function.

8.1.2 HAG Requirements

[R-110] The HAG MUST support MPTCP, per RFC 6824.
[R-111] The HAG MUST be able to act as a Transport Converter, turning MPTCP connections coming from the HCPE (Converter Client) into regular TCP connections to the destination server, supporting draft-ietf-tcpm-converters [11].
[R-112] The HAG MUST be able to disable/enable the Transport Converter function on a per subscriber basis, as dictated by policy.

[R-113] The HAG MUST be able to carry TCP traffic over the MPTCP Hybrid Access path group.
[R-114] The HAG MUST support TCP Fast Open, per RFC 7413.
[R-115] The HAG MUST support receiving the original destination IP address and port of the TCP session in SYN messages with the Fast Open option, as per draft-ietf-tcpm-converters.
[R-116] The HAG MUST simultaneously support both IPv4 and IPv6 MPTCP subflows.
[R-117] The HAG MUST support sending the HAG MPTCP termination IP address(es) to the HCPE by means of the ADD_ADDR TCP option.
[R-118] The HAG MUST support receiving the HCPE MPTCP termination IP address(es) by means of the ADD_ADDR TCP option.
[R-119] The HAG SHOULD support configurable MPTCP congestion control schemes and packet schedulers.
[R-120] The HAG SHOULD be able to be configured to disable the MPTCP checksum procedures.
[R-121] The HAG SHOULD support gathering statistics and event logging for the MPTCP Proxy function.
[R-122] The HAG MUST be able to offer one public IPv4 address and/or one public IPv6 prefix per customer on the Internet-facing interface.

8.2 MPTCP Implicit Mode

8.2.1 HCPE Requirements

[R-123] The HCPE MUST support MPTCP, per RFC 6824.
[R-124] The HCPE MUST support Network-assisted MPTCP implicit mode, as described in in Section 4.3.4.
[R-125] The HCPE MUST be able to disable/enable its MPTCP Hybrid Access function, by means of device management, in support of Section 6.1.2/TR-348.
[R-126] The HCPE MUST be able to carry TCP traffic over the MPTCP Hybrid Access path group.
[R-127] The HCPE MUST be able to relay any non-TCP traffic over the primary path, without the use of MPTCP.
[R-128] The HCPE SHOULD implement a mechanism that allows the HAG to differentiate between MPTCP sessions originated by the HCPE from MPTCP sessions originated from a client.
[R-129] The HCPE MUST simultaneously support both IPv4 and IPv6 MPTCP subflows.
[R-130] The HCPE MUST support sending the HCPE secondary path MPTCP termination IP address(es) to the HAG by means of the ADD_ADDR TCP option.
[R-131] The HCPE MUST support receiving the HAG secondary path MPTCP termination IP address(es) by means of the ADD_ADDR TCP option.
[R-132] The HCPE MUST support the configuration of the HAG secondary path MPTCP termination IP address(es) via device management.
[R-133] The HCPE SHOULD support receiving the HAG secondary path MPTCP termination IP address(es) by means of a DHCP option.
[R-134] The HCPE SHOULD support configurable MPTCP congestion control schemes and packet schedulers.
[R-135] The HCPE SHOULD be able to be configured to disable the MPTCP checksum procedures.

[R-136] The HCPE MUST support at least one IPv4 address and/or one IPv6 address per Internet facing interface.
[R-137] If the HCPE cannot establish MPTCP subflows to the HAG, it MUST fallback to ordinary TCP operation over the primary path.
[R-138] The HCPE SHOULD support selecting which TCP flows are subject to MPTCP-based Hybrid Access, as determined by policy.
[R-139] The HCPE SHOULD support gathering statistics and event logging for the MPTCP Proxy function.
[R-140] HCPE SHOULD support native MPTCP connections.
[R-141] The HCPE MUST initiate the first subflow for new MPTCP sessions on the primary path, not on the secondary path.

8.2.2 HAG Requirements

[R-142] The HAG MUST support MPTCP, per RFC 6824.
[R-143] The HAG MUST support Network-assisted MPTCP on-path implicit mode, as described in Section 4.3.4
[R-144] The HAG MUST be able to disable/enable its MPTCP Hybrid Access function on a per subscriber basis, as dictated by policy.
[R-145] The HAG MUST be able to carry TCP traffic over the MPTCP Hybrid Access path group.
[R-146] The HAG SHOULD be able to relay any non-TCP traffic over the primary path, without the use of MPTCP.
[R-147] The HAG MUST simultaneously support both IPv4 and IPv6 MPTCP subflows.
[R-148] The HAG MUST be deployed on-path for the primary path used by the HCPE. The "primary" is the path used by all non-MPTCP traffic and the initial MPTCP subflow.
[R-149] The HAG MUST support sending the HAG secondary path MPTCP termination IP address(es) to the HCPE by means of the ADD_ADDR TCP option.
[R-150] The HAG MUST support receiving the HCPE secondary path MPTCP termination IP address(es) by means of the ADD_ADDR TCP option.
[R-151] The HAG SHOULD support configurable MPTCP congestion control schemes and packet schedulers.
[R-152] The HAG SHOULD be able to be configured to disable the MPTCP checksum procedures.
[R-153] The HAG SHOULD support gathering statistics and event logging for the MPTCP Proxy function.
[R-154] The HAG MUST retain the state of the TCP session source and destination address from the on-path subflow, and use these addresses on the network side TCP session.

8.2.3 Charging and Billing Requirements

For legal purposes, the network operator may have to provide data retention. If the HAG performs Source NAT or non-transparent proxy function, it hides the original customer IP addresses and so the mapping between the address visible in the Internet and the customer must be recorded at the HAG.

[R-155] The HAG MUST support lawful interception.
[R-156] The HAG MUST support logging the NAT'ed/proxied connections so that data retention obligations can be fulfilled.

End of Broadband Forum Technical Report TR-378