TECHNICAL REPORT

TR-146
Subscriber Sessions

Issue: 1
Date: May 2013

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Broadband Forum Technical Report has been approved by members of the Forum. This Broadband Forum Technical Report is not binding on the Broadband Forum, any of its members, or any developer or service provider. This Broadband Forum Technical Report is subject to change, but only with approval of members of the Forum. 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.

This Broadband Forum Technical Report is provided AS IS, WITH ALL FAULTS. ANY PERSON HOLDING A COPYRIGHT IN THIS BROADBAND FORUM TECHNICAL REPORT, OR ANY PORTION THEREOF, DISCLAIMS TO THE FULLEST EXTENT PERMITTED BY LAW ANY REPRESENTATION OR WARRANTY, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTY:

(A) OF ACCURACY, COMPLETENESS, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR TITLE;
(B) THAT THE CONTENTS OF THIS BROADBAND FORUM TECHNICAL REPORT ARE SUITABLE FOR ANY PURPOSE, EVEN IF THAT PURPOSE IS KNOWN TO THE COPYRIGHT HOLDER;
(C) THAT THE IMPLEMENTATION OF THE CONTENTS OF THE TECHNICAL REPORT WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

By using this Broadband Forum Technical Report, users acknowledge that implementation may require licenses to patents. The Broadband Forum encourages but does not require its members to identify such patents. For a list of declarations made by Broadband Forum member companies, please see http://www.broadband-forum.org. No assurance is given that licenses to patents necessary to implement this Technical Report will be available for license at all or on reasonable and non-discriminatory terms.

ANY PERSON HOLDING A COPYRIGHT IN THIS BROADBAND FORUM TECHNICAL REPORT, OR ANY PORTION THEREOF, DISCLAIMS TO THE FULLEST EXTENT PERMITTED BY LAW (A) ANY LIABILITY (INCLUDING DIRECT, INDIRECT, SPECIAL, OR CONSEQUENTIAL DAMAGES UNDER ANY LEGAL THEORY) ARISING FROM OR RELATED TO THE USE OF OR RELIANCE UPON THIS TECHNICAL REPORT; AND (B) ANY OBLIGATION TO UPDATE OR CORRECT THIS TECHNICAL REPORT.

Broadband Forum Technical Reports may be copied, downloaded, stored on a server or otherwise re-distributed in their entirety only, and may not be modified without the advance written permission of the Broadband Forum.

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

History

Issue Number	Date	Publication Date	Issue Editor	Changes
1	May 28, 2013	June 4, 2013	Frederic Klamm, France Telecom-Orange	Original

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

Editors	Alan Kavanagh	Ericsson
	Frederic Klamm	France Telecom
	Wojciech Dec	Cisco Systems
E2E Architecture	David Allan	Ericsson
WG Chairs	David Thorne	BT
Vice Chair	Sven Ooghe	Alcatel-Lucent

TABLE OF CONTENTS

List of Figures

List of Tables

No table of figures entries found.

Executive Summary

TR-146 describes Subscriber Sessions as an evolution of broadband access networks from PPP-based to Ethernet and IPoE-based networks. TR-146 describes Subscriber Session Management for Ethernet, IPv4 and IPv6 Subscriber Sessions based on TR-101[4], TR-145 [7], TR-156 [8], TR-167 [9], TR-177 [10] and TR-187 [11] architectures. TR-146 leverages AAA and IP address/prefix allocation of IPoE connected endpoints in order to create Subscriber Session and provide specific profile management for them.

Purpose and Scope

1.1 Purpose

The purpose of TR-146 is to define Subscriber Sessions and Flow classifiers, along with Subscriber Session authentication and management applicable to a broadband access environment. This should allow service providers to provide a diversified set of Ethernet and IP services, whilst still having the network tools to control and account for them.

TR-146 presents basic network element requirements to ensure overall functionality and vendor interoperability.

1.2 Scope

TR-146 describes Subscriber Sessions and Subscriber Session Grouping in the context of TR-145 [7], TR-177 [10], WT-178 and TR-187[11]. A Subscriber Session description covers the following stages:

- Subscriber Session creation and triggering (see section 5.6)
- The application and change of Subscriber Session profiles, including authentication, authorization, accounting, monitoring and Grouping
- Subscriber Session termination (see section 5.7)

For each element of this cycle, TR-146 describes the usage scenarios, mechanisms, protocols and interactions necessary to Subscriber Sessions. Furthermore, the document also describes the treatment of a subset of a Subscriber Session traffic known as IP Flows. These are limited to those spanning the IP network and transport protocols. Traffic classification methods at higher layers of the protocol stack are not in the scope of TR-146.

Application layer sessions and flows are outside the scope of this document.

TR-146 re-uses architectural notions introduced in TR-59 [1], TR-101 [4] and TR-177 [10], and places requirements on RG and various Service Edges for the purpose of supporting Ethernet, and single IPv6 or IPv4 stack, or dual stack IP Sessions on the Service Edge.

This TR supports both bridged and routed RG modes.

TR-146 acknowledges the existence of a BPCF (Broadband Policy Control Framework) as defined in TR-134 [6]. It provides a set of recommendations that a Service Edge interface should support concerning Subscriber Sessions and flows encompassed by these Subscriber Sessions, along with their delegated policies. However, TR-146 does not define the interface protocol, just the information that is expected to be conveyed across such an interface.

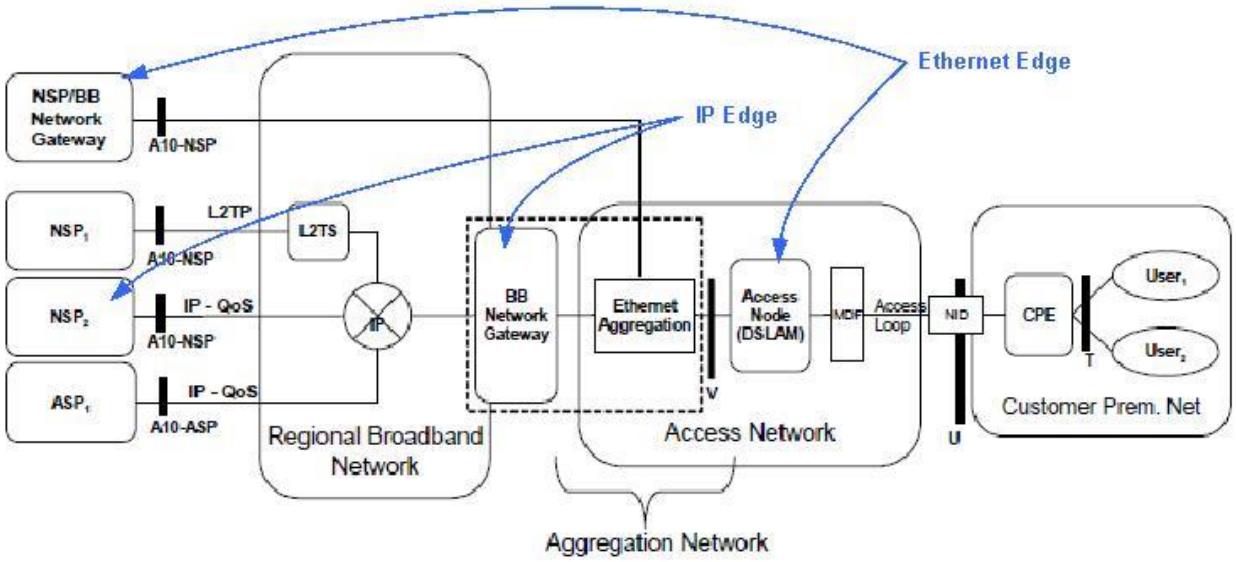

Figure 1: Service Edges in TR-101architecture

Note: As shown in Figure 1 and in Figure 2, there can be multiple Service Edges, each managing a Subscriber Session at L2 or L3, and multiplexed on lower layer Ethernet Sessions. So, an Ethernet Edge may manage Ethernet usage, while at the same time an IP Service Edge may manage IP usage. Further, IPv4 and IPv6 usage may be managed by different IP Service Edges

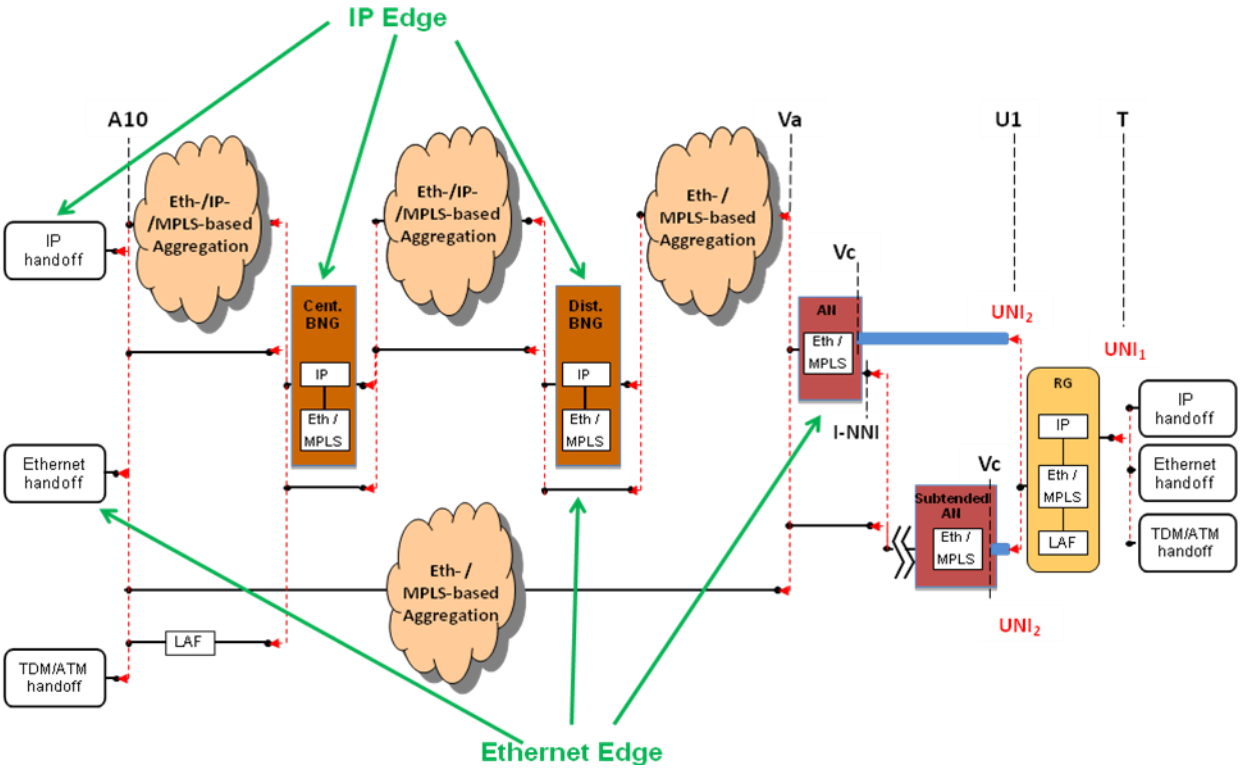

Figure 2: Service Edges in WT-178 architecture

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 [15].

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 adjective “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 adjective “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-059	*DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services*	Broadband Forum	2003
[2] TR-069	*CPE WAN Management Protocol*	Broadband Forum	2004
[3] TR-092	*Broadband Remote Access Server (BRAS) Requirements Document*	Broadband Forum	2004
[4] TR-101i2	*Migration to Ethernet-Based Broadband Aggregation*	Broadband Forum	2011
[5] TR-124i3	*Functional Requirements for Broadband Residential Gateway Devices*	Broadband Forum	2012
[6] TR-134	*Broadband Policy Control Framework*	Broadband Forum	2012
[7] TR-145	*Multi-service Broadband Network Functional Modules and Architecture*	Broadband Forum	2012
[8] TR- 156i3	*Using GPON Access in the context of TR-101*	Broadband Forum	2012
[9] TR-167i2	*GPON-fed TR-101 Ethernet Access Node*	Broadband Forum	2010
[10] TR-177	*IPv6 in the context of TR-101*	Broadband Forum	2010
[11] TR-187	*IPv6 for PPP Broadband Access*	Broadband Forum	2010
[12] 802.1X	*Port Based Network Access Control*	IEEE	2010
[13] 802.1ad	*Provider bridging*	IEEE	2005
[14] 802.1Q	*Virtual LANs*	IEEE	2005
[15] RFC 2119	*Key words for use in RFCs to Indicate Requirement Levels*	IETF	1997
[16] RFC 2131	*Dynamic Host Configuration Protocol*	IETF	1997
[17] RFC 2684	*Multiprotocol Encapsulation over ATM AAL5*	IETF	2006
[18] RFC 2865	*Remote Authentication Dial-in User Service*	IETF	2000
[19] RFC 3118	*Authentication for DHCP Messages*	IETF	2001
[20] RFC 3315	*Dynamic Host Configuration for IPv6*	IETF	2003

[21] RFC 3633	*IPv6 Prefix Options for DHCP version 6*	IETF	2003
[22] RFC 4649	*DHCPv6 Relay Agent Remote ID Option*	IETF	2006
[23] RFC 4861	*Neighbor Discovery for IPv6*	IETF	2007
[24] RFC 4862	*IPv6 Stateless Address Autoconfiguration*	IETF	2007
[25] RFC 5176	*Dynamic Authorization Extensions to RADIUS*	IETF	2008
[26] RFC 5880	*Bidirectional Forwarding Detection*	IETF	2010
[27] RFC 5881	*Bidirectional Forwarding Detection*	IETF	2010
[28] RFC 6088	*Traffic Selectors for Flow Binding*	IETF	2011
[29] RFC 6221	*Lightweight DHCPv6 Relay Agent*	IETF	2011
[30] RFC 6704	*Forcerenew Nonce Authentication*	IETF	2012
[31] RFC 6788	*Line-Identification Option*	IETF	2012

2.3 **Definitions**

The following terminology is used throughout this Technical Report.

Accounting Record An accounting record represents a summary of information collected for the Subscriber Session construct. A Service Edge may create an accounting record, according to its accounting policy, at Subscriber Session start, stop and at interim intervals during the existence of the Subscriber Session. Each accounting record contains an identifier that uniquely identifies the Subscriber Session for which it is generated.

Flow A grouping of traffic identified by a set of header information and port information including, but not limited to: L3 header, L2 Header, Virtual and/or Physical interface Port, and/or Agent Circuit ID information for a remote port in the access network.

IP Edge A service provider controlled "Service Edge" device capable of hosting an IP Session. This can be a BNG in the context of TR-101[4] or BRAS in the context of TR-59 [1].

IP Flow An IP Flow is identified by a 5-tuple IP parameter traffic classifier. An IP Flow identifier forms the classification element of a traffic policy that is applied to a Subscriber Session. The 5-tuples is made up of following header fields: source IP address, source port, destination IP address, destination port and protocol.

IP Session An IP Session is a grouping of traffic according to one or more classifiers visible at a control point, called the IP Edge, in the broadband network. The classifier is composed of, at a minimum, a Subscriber's IP address (v4 or v6), IPv4 subnet or IPv6 prefix. Additional Layer1 and Layer 2 parameters may be part of the IP Session where appropriate.

Service Edge A device capable of hosting a Subscriber Session. This can be a BNG in the context of TR-101 [4] or a BRAS in the context of TR-59 [1], or it can also include an Ethernet edge device or an Access Node in the context of an Ethernet service.

Session A Session is a logical construct intended to represent a network connectivity service instance at a network node. Data and control plane policies are associated with Sessions. Sessions are initiated and configured dynamically or statically. A Session may have associated state.

Subscriber Session	A Subscriber Session is either a PPP Session, an IP Session, or an Ethernet Session. Subscriber sessions are used to represent all traffic that is associated with that subscriber by a given service provider in order to provide a context for policy enforcement.
Ethernet Session	A Ethernet sessions represents all traffic from a subscriber that is associated with a Subscriber's Ethernet address and/or access port. This is also described as a layer 2 Active Line Access Session.
Access Session	An Access Session where the access link comes up and is available for data transmission. In the DSL case, this starts when the DSL modem has trained up with the DSLAM, and via ANCP the DSLAM would then transmit a Port Up message to BNG
Traffic Rule session	A Traffic Rule session is an abstraction of a set of policy rules. This would be used with an identifier to allow an operator to know if a particular "set of Traffic rules" is enabled or not without needing to know the underlying rule details. For example, an http redirect service would be a set of rules that allows DNS traffic to be transmitted, http traffic redirected to a web portal, and all other traffic dropped.
DHCP Proxy:	In IPv4, a device that acts as a DHCP Server to downstream clients while acting as a DHCP client to upstream DHCP servers.
Delegating router	In IPv6, a router that acts as a DHCP server or relay and is responding to the prefix requests (RFC 3633 [21])
Requesting router	The requesting router is the router that acts as a DHCP client. It is requesting the prefix(es) to be assigned (RFC 3633 [21])

2.4 Abbreviations

This Technical Report uses the following abbreviations:

3GPP	Third Generation Partnership Project
AAA	Authentication Authorization Accounting
AN	Access Node
ANCP	Access Node Control Protocol
API	Application Programming Interface

ARP	Address Resolution Protocol
ATM	Asynchronous Transfer Mode
BFD	Bidirectional Forwarding Detection
BNG	Broadband Network Gateway
BPCF	Broadband Policy Control Framework
BRAS	Broadband Remote Access Server
C-VID	Customer VLAN ID
CoA	Change of Authorization
COPS	Common Open Policy Service
CPE	Customer Premise Equipment
DAD	Duplicate Address Detection
DEI	Drop Eligible Indicator
DHCP	Dynamic Host Configuration Protocol
DoS	Denial of Service
DSCP	Differentiated Service Code Point
DSL	Digital Subscriber Line
DS-Lite	Dual Stack Lite
DSLAM	Digital Subscriber Line Access Multiplexer
FQDN	Fully Qualified Domain Name
HMAC	Host-based Message Authentication Code
IA_PD	Identity Association for Prefix Delegation
ICMP	Internet Message Control Protocol
IP	Internet Protocol
IPoE	IP over Ethernet
L1	Layer 1
L2	Layer 2
L3	Layer 3
MAC	Media Access Control
ME	Maintenance Entity
NAI	Network Address Identifier
NAT	Network Address Translation
ND	Neighbor Discovery
NUD	Neighbor Unreachability Detection
OAM	Operations And Maintenance
PD	Prefix Delegation
PDF	Policy Decision Function

PIO	Prefix Information Option
PPP	Point-to-Point Protocol
PPPoA	PPP over ATM
PPPoE	PPP over Ethernet
PSTN	Public Switched Telephone Network
RA	Router Advertisement
RADIUS	Remote Authentication Dial-In User Service
RG	Residential Gateway
RIP	Routing Information Protocol
RS	Router Solicitation
S-VID	Service VLAN ID
SLAAC	StateLess Address AutoConfiguration
SP	Service Provider
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VCI	Virtual Circuit ID
VLAN	Virtual LAN
VPI	Virtual Path ID
WAN	Wireless Area Network

3 Technical Report Impact

3.1 Energy Efficiency

TR-146 has no impact on Energy Efficiency.

3.2 IPv6

IPv6 is in the scope of this document. IPv6 requirements in this TR should be considered along with requirements found in TR-187 [11] and TR-177[10].

3.3 Security

TR-146 introduces security requirements and an authentication framework to authenticate a fixed line access device such as an RG and a method to authenticate visiting devices such as 3GPP UEs (User Equipments) connected to the RG.
An IPv6 network can allow each end device in the customer network to have its own address which can be directly visible to, and accessible from, the Internet, as opposed to being hidden behind the NAT in an RG. This raises general security and privacy issues which are not uniquely related to Policy, but still apply in the Policy context. This document explores the implications of the use of IPv6 on IP Subscriber Sessions but doesn't attempt to address the broader security issues related to the removal of NAT from an IPv6 customer premises network.

3.4 Privacy

TR-146 does not introduce any new privacy-related issues.

4 Introduction

TR-146 describes Subscriber Sessions as an evolution of broadband access networks from PPP-based to Ethernet and IPoE-based networks. TR-146 describes Subscriber Session Management for Ethernet, IPv4 and IPv6 Subscriber Sessions based on TR-101 [4], TR-145 [7], TR-156 [8], TR-167 [9], TR-177 [10] and TR-187 [11] architectures. TR-146 leverages AAA and IP address/prefix allocation of IPoE connected endpoints in order to create Subscriber Session and provide specific profile management for them.

TR-146 defines Subscriber Sessions and Flow classifiers, along with Session authentication and management applicable to a broadband access environment. This allows service providers to provide a diversified set of Ethernet and IP services, whilst still having the network tools to control and account for them.

TR-146 presents basic network element requirements to ensure overall functionality and vendor interoperability.

5 Subscriber Sessions

This section defines Subscriber Session of different types and provides requirements for network elements to support their establishment and termination.

5.1 Subscriber Session Life Cycle

A Subscriber Session life cycle can be decomposed into the following stages:

- Subscriber Session Creation
- Determination and Execution of applicable Subscriber Session Policies/Profiles (including authentication)
- Subscriber Session Termination

The Subscriber Session can be triggered by a data plane event or a control plane event. Such a data plane event could be the reception of a particular packet. Once initiated, authentication occurs and policies may be triggered by the control plane.

Once a Subscriber Session is created on a Service Edge, it is necessary to monitor the Subscriber Session state via control or data plane events. For example, this could be when the user has explicitly disconnected, an inactivity timeout has been reached, or when a traffic volume quota has been exceeded.

While the Subscriber Session is active, it is also desirable to be able to change its state, or some associated parameters.

The actions that are taken following Subscriber Session creation, i.e. traffic detection monitoring and change events, constitute the Subscriber Session policy. This is split into control policies, which deal with events from control protocols and processes associated with the Subscriber Session, and Session traffic policies. The latter requires a means of classifying the traffic, and the concept of IP Flow has associated with it such a classifier.

Each Subscriber Session also needs to have a means of being terminated, and removing the Subscriber's context and policies, when appropriate. Such termination can be achieved through data plane or control plane actions governed by the Subscriber policy/termination rules.

The failure of a "keep-alive" protocol can form a Subscriber Session Termination rule, or the Subscriber Session State derived from a control protocol, e.g. DHCP or 802.1X, implies that Subscriber Sessions must persist as long as the related control protocol state is valid, potentially even following a keep-alive failure event. Both perspectives are in fact legitimate and underscore different potential uses for Subscriber Sessions.

5.2 Types of Subscriber Sessions

In broadband environments, Subscriber connectivity may be permanent or ephemeral. Similarly, IP addresses and prefixes are assigned either statically or dynamically, or, in the case of an Ethernet service, not at all.

This Technical Report specifies two types of Subscriber Sessions, based on their provisioning method and overall purpose.

Fixed Sessions*:* A fixed Session is instantiated and terminated by management provisioning action rather than by data or control plane events.

Certain services, for example those offered to business Subscribers, are nailed-up. They can support an Ethernet service, or an IP service with fixed IP address(es) or prefix(es). In this case, the service is always on. Fixed Subscriber Sessions allow external parties or systems to have continuous connectivity. For fixed Subscriber Sessions, authentication is inferred from the facility.

Dynamic Sessions: A dynamic Session is instantiated and terminated by data plane or control plane events.

In typical Subscriber services, connectivity to the service provider's network is initiated by the Subscriber via an attachment procedure that may include authentication.

In the context of TR-59 [1] and TR-101 [4], both of the above IP Session types will be typically hosted on a BRAS or BNG. However this document is generally agnostic to these architectures and so the generic term Service Edge is used instead. Moreover a Service Edge can also support the native Ethernet access to metro Ethernet business services as well as wholesale partners for Subscriber Internet access.

R-01 The Service Edge MUST support the creation and termination of Fixed and Dynamic Subscriber Sessions.

5.3 IP Sessions and IP Flows

Extending the generic Session concept to the IP protocol yields the concept of an IP Session, which represents a Subscriber's IP Session Traffic or a portion thereof, and is associated with a Subscriber's IP address or prefix. IP Sessions are abstractions from a service provider's perspective, of the connections and resources associated with a Subscriber IP end-point. This abstraction allows applying policies for traffic residing on a service provider's IP Edge.

As mentioned previously, the application of traffic policies requires a traffic classification mechanism. In the context of an IP Session, this is the IP Flow classifier.

A basic IP Edge requirement can be derived based on the overview so far:

R-02 The IP Edge MUST be able to support the application of traffic policies to an IP Session with configurable IP Flow classifiers.

5.4 The Ethernet Session

Extending the generic Session concept to the Ethernet protocol yields the concept of an Ethernet Session, which represents a Subscriber's Ethernet traffic, but typically not a portion thereof, and is associated with a Subscriber's Ethernet address and/or access port. Ethernet Sessions are an abstraction for a service provider of the connections and resources of a Subscriber end-point and may include multiple IP stacks and tunneling protocols, or may be just plain Ethernet. This abstraction allows for both Ethernet and IP policies to be applied to the Subscriber Session residing on a service provider's Service Edge.

A base Service Edge functional requirement can be derived based on the overview so far:

R-03 The Ethernet Edge MUST be able to support traffic policies with configurable Ethernet Flow classifiers.

5.5 Relationships among Subscriber Sessions

There can be relationships between IP flows, IP Sessions and Ethernet Sessions. Generally speaking, one Session contains another when that Session encapsulates the other in the protocol stack. So for instance, an IP flow can be carried in an IP Session that is carried within a PPP Session, and that PPP Session can be carried within an Ethernet Session. If the Ethernet Session terminates, it implies that the PPP Session and IP Sessions also terminate, and the IP Flow is removed. It's also possible to Group Sessions independent of their encapsulations, as described in section 6.5. So there are 2 types of relationships:

1. an encapsulation relationship,
2. a grouping/binding relationship.

5.6 Subscriber Session Creation

This section deals with the mechanisms for detecting events triggering Subscriber Sessions creation. Both Ethernet and IP Sessions are combined in this section because of their similarities. The section is organized into interface specific considerations followed by Subscriber Session creation sub sections. The PPP and IP-specific functional requirements can be separated, allowing operators to determine the applicable element requirements based on their planned use of Ethernet, IP, PPPoE or mixed access mode architectures.

The concept of a Subscriber Session is independent of the physical interface that communicates the Session traffic. However the appropriate Subscriber Session creation mechanisms need to be supported on physical and logical interfaces connecting the Service Edge to the aggregation network. In particular it is envisaged that Subscriber Session traffic will be arriving on the Service Edge via an interface corresponding to reference point in the aggregation network architecture.

At a minimum, a Service Edge needs to support Ethernet, IP and where required PPP Subscriber Session creation on the following protocol stacks:

- 802.1ad [13]
- IP over 802.1ad
- PPP over E over 802.1ad
- 802.1Q[14]
- IP over 802.1Q
- PPP over E over 802.1Q
- RFC2684 [17] bridged
- IP over RFC2684 [17] bridged
- IP over RFC2684 [17] routed
- PPP over E over RFC2684 [17] bridged
- PPP over ATM

In order to deal with mixed protocol deployments, whereby a single Subscriber may use both the PPPoE and IPoE protocols (e.g. PPPoE for internet access and IPoE for Set-Top-Box connectivity); it is desirable for the Service Edge to support multiple protocols on a given interface.

5.6.1 Fixed Subscriber Session Provisioning

Subscribers that have static configuration which are bound to a logical interface but not shared with other Subscribers can be represented by a static Subscriber Session corresponding to a logical interface. This is typically the case when there is a 1:1 relationship between the logical interface and the Subscriber. In this case, the act of configuring a static Subscriber Session by the operator is deemed to be the Subscriber Session start event.

R-04 The Service Edge MUST be able to support provisioning Subscriber Sessions on dedicated logical interfaces.

5.6.2 Triggers for Dynamic Subscriber Session creation

Neither Ethernet nor IP has a tightly bound or native control protocol to signal the start of a Subscriber Session before the Subscriber transmits data packets.

The following triggers as observed at the Service Edge are the most useful in defining the "first sign of life", or Subscriber Session Start.

- An Ethernet 802.1ad or 802.1Q packet received by the Service Edge.
  This event can be the first sign of life for a new Subscriber Session when the access network opens a bridging gate through 802.1X [12].

- A DHCP packet received by the IP Edge.

- An IPv4 packet that is received by the IP Edge from a source IP address that has no association with an established Subscriber Session

- An IPv6 RS packet with a local loop identifier

Note 1: When IPv6 prefixes are delegated to the home, traffic that carries these prefixes needs to be associated with this Subscriber Session.

Note 2: PPPoE includes a tightly bound control protocol from which the Session could be directly inferred.

In the case of DHCPv4 initiated IP Sessions, it is necessary for the IP Edge to monitor the most relevant value with regard to the IP Edge of the DHCPv4 lease that is associated with the IP Session and update the lease information during periodic DHCPv4 renewals for active IP Sessions.

However the DHCP lease obtained through the DHCP server does not override the IP Session lifetime parameter (e.g. RADIUS Session-Timeout) provided by the AAA server at IP Session initiation or during subsequent updates.

IPv6 has two methods for address assignment, SLAAC ([24] and DHCPv6 [20]. When an IPv6 device attaches to a link for the first time or re-attaches, it initializes its interface, sending a Router Solicitation (RS) upstream and a Neighbor Solicitation for Duplicate Address Detection (DAD) to validate its addresses. The IP Edge will receive the RS and responds by sending a solicited Router Advertisement (RA).

For networks that use native IPv6 over Ethernet in the access, it is desirable to control both the SLAAC address assignment procedure as well as the DHCPv6 procedure on a per Subscriber basis, while limiting the number of exchanges with the AAA subsystem. In order to do so it is desirable to trigger the IP Session start process only once the Service Edge has received sufficient credentials for authenticating and authorizing the Subscriber and only then continue with the address assignment procedures. Prior to providing prefix(es) via DHCPv6, link local address assignment is performed first.

The following requirements therefore arise:

5.6.2.1 **Service Edge requirements**

R-05 The IP Edge MUST support IPv4 IP Session initiation triggered by the reception of a unique Subscriber originated DHCPv4 DHCPDISCOVER message (RFC 2131 [16]).

R-06 The IP Edge MUST support IPv6 IP Session initiation triggered by the reception of a unique Subscriber originated DHCPv6 SOLICIT message (RFC 3315 [20]).

R-07 The IP Edge MUST support IPv6 IP Session initiation triggered by the reception of a unique Subscriber originated RS message (RFC 4861 [23]) that specify the Subscriber line ID.

R-08 The IP Edge MUST support IP Session initiation triggered by the reception of any Subscriber originated IPv4 packet whose source IP address is not associated with an already-established Session.

R-09 The IP Edge MUST support a per logical interface configuration of the IP Session initiation event method (e.g. DHCPv4 -triggered or IPv4 packet-triggered).

R-10 The IP Edge MUST support the configuration of a permissible IP (subnet) range(s) when using the IPv4 packet-triggered IP Session initiation method.

R-11 The IP Edge MUST be able to create valid routing or forwarding table entries corresponding to the address(es) or prefix(es) associated with the IP Session.

R-12 The IP Edge MUST support updating an IP Session DHCPv4 lease time by reading the lease duration specified in DHCPACK messages.

R-13 The IP Edge MUST be configurable so as to maintain DHCPv4 initiated IP Sessions for the duration of the DHCPv4 lease.

R-14 When SLAAC is used for Subscriber upstream address configuration, the IP Edge MUST be configurable so as to maintain IP Sessions as long as at least one of the prefixes in IA_PD option in the received REPLY message, is associated a valid lifetime.

R-15 When DHCPv6 is used for Subscriber upstream address configuration, the IP Edge MUST be configurable so as to maintain IP Sessions as long as at least one of the addresses in the IA_NA option in the REPLY message, is associated a valid lifetime.

R-16 The IP Edge MUST support IP Session initiation following the reception of a valid Subscriber originated DHCPv6 (RFC3315 [20]) packet.

R-17 The IP Edge MUST support DHCPv6 address assignment to RG as per RFC3315 [20].

R-18 The IP Edge MUST be able to be configured as a delegating router, receiving IPv6 Prefix(es) on a per Subscriber basis via AAA, and advertising them in ICMPv6 RAs and DHCPv6 IA_PD option (RFC3633 [21]) respectively.

R-19 The IP Edge SHOULD be capable of sending unicast RA messages listing the AAA derived prefix in the Prefix Information Option and indicating this prefix to be eligible for SLAAC. The unicast RAs MUST be sent to the host MAC address which is derived from the incoming RS messages .

R-20 The IP Edge MUST support sending of unsolicited multicast RA messages that do not contain a Prefix Information Option (PIO).

R-21 The IP Edge MUST support sending of solicited unicast RA messages that do not contain a Prefix Information Option (PIO)

R-22 For each IP Session the IP Edge MUST create and maintain the appropriate forwarding entries, this will be one of the following :

- the interface prefix and delegated prefix(es),
- The IPv4 address (of the RG or host)
- The IPv6 address (of the RG or host)

R-23 On the Subscriber facing interface the IP Edge MUST be capable of associating traffic with an IP Session on the basis of the link assigned and/or delegated IPv6 prefix(es).

R-24 The Ethernet Edge MUST be able to support Ethernet Session initiation following the reception of an Ethernet frame whose source MAC address and/or S-VID and/or C-VID is not associated with an already-established Session.

R-25 The Ethernet Edge MUST support installing one or more MAC address filters at an access port based on an 802.1X [12]supplicant and an authentication interaction.

5.6.2.2 **RG Requirements**

R-26 The RG MUST support Neighbor Discovery and IPv6 Stateless Auto Configuration as specified in RFC4861 [23] and RFC4862 [24] respectively.

R-27 The RG MUST support Prefix Delegation as per TR-124 [5], section on WAN.IPv6.

5.6.3 PPP Session detection

The mechanisms for detecting PPP sessions are well known and described in TR-092 [3]and in TR-101 [4].

5.7 Subscriber Session Termination

The Subscriber Session Termination mechanism is closely coupled to the Subscriber Session traffic detection method, the use of a "keep-alive" protocol, and whether the Session is static or dynamic. The termination of a Subscriber Session may be triggered in one of the following ways:

- **The reception of DHCPRELEASE message from the Subscriber's DHCPv4 client.**

  When DHCP is used as a means to initiate an IPv4 Session and no DHCPRELEASE message is coming from the Subscriber's DHCP client, the Session termination will be triggered by the DHCP lease expiry.

- **The reception of RELEASE message from the Subscriber's DHCPv6 client.**

  When DHCP is used as a means to initiate an IPv6 Session and no RELEASE message is coming from the Subscriber's DHCP client, the Session termination will be triggered by the valid lifetime expiry of the allocated addresses.

- **A Session keep-alive protocol failure event**

  This is covered in Section 6.2.

- **A Session Termination Command**

  This method refers to the case when an automatically generated or operator issued Session termination command is received by the Service Edge. Such a command may be due to a volume or time quota being exceeded as determined by the user policy, the actions of the Subscribers themselves, e.g. as a result of a sign-off at the portal, or the actions of an operator. The command could be generated internally (an API) or passed via an operator command line interface, or received externally from an AAA or Application sub-system.

- **Termination of the underlying interface construct**

  The termination or de-activation of an underlying construct, e.g. a link failure or link-down condition, should trigger the termination of the Sessions relying on this construct.

As specified in RFC 3633 [21], prefix delegation can be done either by the IP Edge acting as a DHCP server, or by the IP Edge acting as a relay between subscribers and a more centralized DHCP server.

The following requirements may be derived:

R-28 The IP Edge MUST be configurable to act as a DHCP Server, and conform to RFC 3315 [20] and RFC 2131 [16].

R-29 The IP Edge MUST be configurable to act as a DHCP Relay, and conform to RFC 3315 [20] and RFC 2131 [16].

R-30 The IP Edge MUST terminate the corresponding IP Session when detecting a DHCPRELEASE from a Subscriber's DHCPv4 client.

R-31 The IP Edge MUST terminate the corresponding IP Session when detecting an IPv6 address/prefix lease's expiry.

R-32 When SLAAC is used for Subscriber upstream address configuration, the IP Edge MUST terminate the corresponding IP Session when observing an IA_PD option in DHCPv6 message signifying that all the delegated prefixes have a zero valid lifetime.

R-33 When DHCP is used for Subscriber upstream address configuration, the IP Edge MUST terminate the corresponding IP Session when observing a DHCPv6 message signifying that all its addresses have a zero valid lifetime.

R-34 The IP Edge MUST be configurable to terminate IP Sessions when loss of L1, L2 or L3 connectivity to Subscriber's premises is detected.

R-35 The Ethernet Edge MUST be able to be configured to terminate Ethernet Sessions when loss of L1 or L2 connectivity to Subscriber's premises is detected.

6 Subscriber Session Policies

In order to support certain types of network services, such as pre-paid services, or service selection portals, it is desirable to have the capability to change policy on an existing IP or Ethernet Session.

The Service Edge should provide a management or policy interface that would allow at least the following operations to be performed on a Subscriber Session:

- Terminating a Subscriber Session,
- Forcing a Subscriber Session re-authentication and reauthorization,
- Changing the Subscriber Session's service policy,
  (e.g. a change of QoS, redirection, etc)
- Allowing a Subscriber Session status query.

The protocol to achieve this should have the following characteristics:

- Open, standards-based and capable of both customization and future extensions,
- Able to disassociate the transport protocol from the carried information, i.e. does not impose an informational model,
- Supports message level security for authentication, replay protection and message integrity,
- Stateless or able to easily recover state (this improves scalability and resiliency),
- Provides a clear way to address or reference Subscriber Session contexts,
- Supports external Subscriber Session queries, alongside queries for specific Subscriber Session characteristics.

The selection of a particular protocol is outside the scope of this document.

R-36 The Service Edge MUST support the configuration and enforcement of per- Subscriber Session policies triggered by Subscriber Session control plane events.

R-37 The Service Edge MUST support the configuration and enforcement of traffic policies within the context of any given Subscriber Session.

6.1 Subscriber Session Authentication, Authorization and Accounting (AAA)

This section describes how AAA interacts with the Subscriber Sessions defined in this document. The AAA process is typically handled via Radius AAA, but can also be handled by other relevant protocols e.g. Diameter.

Subscriber Sessions introduce some new challenges related to the fact that Ethernet and IP Sessions generally lack the means to obtain Subscriber credentials and perform authentication. In

order to define a means for collecting Subscriber credentials, a Subscriber Session authentication model needs to be enhanced. This will be the topic of future work.

Service Providers choosing to offer services without requiring credentials, e.g. through service authorization based on Subscriber line-id, can use Subscriber Session identifier(s) derived directly from the packet that triggers the generic Session. This can include one or more of the following identifiers:

1. DHCP Options contained in the respective v4 and v6 DHCP message (e.g. the Agent Circuit ID and/or Agent Remote ID sub-options of DHCPv4 Option 82 "Relay Agent Information option") or DHCPv4 Option-60 "Vendor Class identifier",
2. Line-identification" contained in RS messages sent by a Subscriber's RG (e.g. as in RFC 6788 [31]),
3. Source MAC address contained in the DHCPv4 packet,
4. DUID contained in a Subscriber-originated DHCPv6 message,
5. The packet's source IP address,
6. The packet's incoming interface identifier, e.g. VLAN or VCI/VPI,
7. The prefix used by the Subscriber, e.g. delegated via DHCP IA_PD.

In some cases the packet's source MAC address may be a useful identifier. However, given that it is not specific to IP packets; it can be used both for Ethernet services as well as groupings of IP services. For single IP services, it is best seen as a supplemental identifier to 1, 2 and 3 above. The source MAC address and an associated IP address may be bound as part of security mechanisms defined in Section 5.7 of TR-101 [4], and this binding can be further cemented through use of port attachment authentication (e.g. 802.1X [12]).

The use of DHCP for address assignment and as trigger for IP Session creation, require some special considerations:

- It is necessary for the IP Edge to act as a:
  - DHCP relay,
  - or a DHCPv4 proxy,
  - or a DHCP server.

- The IP Edge needs to monitor the DHCP negotiation process and be able to instantiate appropriate IP and ND or ARP forwarding entries once an address is assigned. Similarly, appropriate client MAC to DHCP address bindings should be instantiated as a security mechanism against IP Session spoofing or hijacking and DoS attacks. This functionality is already defined in TR-092 [3] and TR-101 [4].

- In order to address the need for supporting different IP address assignment ranges, as commonly found in Layer 3 wholesale scenarios, the IP Edge needs to have a mechanism for assigning each DHCP initiated IP Session to a DHCP address pool or directing the

DHCP message to a specific DHCP server. Such DHCP handling may be dictated by Subscriber profile information stored in a AAA back-end server.

- The IP Edge may need to retain stateful information regarding the DHCP lease, given that this is intrinsically linked to the state of the IP Session.

- It is desirable to synchronize DHCP lease time and IP Session timeout provided by the AAA server, in order to avoid long connectivity loss, especially when no keep-alive protocol is activated between the RG and the IP Edge.

The presence of statically addressed IP Subscribers also has some implications for the creation of a IP Session.

R-38 The IP Edge MUST be able to create a binding between the IP address or prefix and the source MAC address following successful IP Session authorization, to prevent Session spoofing or hijacking and DoS attacks.

Upon Subscriber Session authorization, the Service Edge should have the ability to retrieve the Subscriber Session policies, including any IP Flow classifiers.

Subscriber Session traffic is expected to be accounted for using well established start/interim/stop mechanisms conveying information on Subscriber Session Traffic to the AAA Server which counts the traffic during the Session duration.

R-39 It MUST be possible to use RADIUS as the AAA protocol.

R-40 The Service Edge MUST support a configurable mechanism for AAA Authorization of Subscriber Sessions.

R-41 The IP Edge SHOULD be able to function as a DHCPv4 proxy.

R-42 The IP Edge MUST support the forwarding of the SOLICIT or DHCPDISCOVER message to a DHCP server whose address is returned from the Radius server as part of the AAA authentication/authorization.

R-43 The Service Edge MUST support configurable identifiers, and SHOULD support combinations thereof used for Subscriber Session from the following list:

- Agent Circuit ID and/or Agent Remote ID sub-option(s) contained in DHCPv4 Option 82 of the DHCPDISCOVER message,
- The DHCPv4 Option 60 value contained in the DHCPDISCOVER message,
- The Source MAC address of the DHCPDISCOVER message,
- The packet's source IP address,
- The packet's source MAC address,

- The DHCPv6 DUID in Subscriber-originated DHCPv6 messages,
- The DHCPv6 Interface-id and/or Remote-id options contained in the RELAY-FORWARD or RELAY-REPLY messages (RFC 4649 [22]),
- An option containing the derivation of access line id. in received RS messages,
- The prefix(es) used by a Subscriber, e.g. prefix(es) delegated via DHCPv6 IA_PD,
- PPP identifiers (agent circuit_id, agent remote_id),
- The Service Edge source port of an 802.1X [12] packet exchange,
  Note: the AAA exchange for 802.1X [12] includes credentials.

R-44 The Service Edge MUST be able to trigger the authorization of a Subscriber Session based on the identifiers listed in R-43.

R-45 For DHCP initiated IP Sessions, the IP Edge when acting as a DHCP server or relay, MUST be capable of assigning to each such IP Session an address coming from an address pool determined from parameters returned in a Session AAA authentication or authorization messages.

R-46 For DHCP initiated IP Sessions, the IP Edge when acting as a DHCP relay MUST be capable of forwarding the DHCP request to a DHCP server indicated by parameters returned in the AAA Session authentication or authorization message

R-47 For DHCPv4-initiated Subscriber Sessions, the IP Edge MUST be configurable to request authorization from the AAA server prior to DHCP lease extension.

R-48 The Service Edge SHOULD support a method for retrieving flow classifiers from a AAA repository.

R-49 The Service Edge MUST be able to use flow classifiers in a traffic policy applied to a Subscriber Session during authentication/authorization.

R-50 For Dynamic Subscriber Session update triggered by the AAA server, the IP Edge MUST support RADIUS Change-of-Authorization (CoA) (RFC5176 [25]) and Disconnect Message Types.

R-51 The IP Edge MUST support RADIUS Accounting Start, Stop and Interim Update messages for both IP and PPP Sessions (as per R-5-15/TR-092 [3] and R-5-16/TR-092 [3]).

R-52 The Ethernet Edge MUST support authenticating 802.1X [12] supplicants using RADIUS.

6.2 **Subscriber Session Monitoring**

Session monitoring can provide an accurate picture of the state of a Subscriber Session. This is not the same as OAM, but can be regarded as an addition to the OAM functionality of the transport network.

Given that IP does not have a universal, tightly bound control-protocol that is able to monitor for the continued existence of Subscriber Sessions some other mechanism is needed for this purpose. The below high-level requirements apply to any such mechanism :

R-53 The IP Edge MUST allow the service provider to retrieve the status of individual IP Sessions.

R-54 The frequency of keep-alive messages SHOULD be configurable on a per Subscriber Session basis.

R-55 The keep-alive response timeout and count SHOULD be configurable on a per Subscriber Session basis.

R-56 The Service Edge MUST be configurable to terminate Subscriber Sessions upon keep-alive failure.

R-57 The Service Edge MUST be configurable to provide to the operator the connectivity status of each Subscriber Session that it supports.

6.2.1 Ethernet OAM

All of the architecture and requirements to test and maintain the connectivity of an Ethernet Session are contained in the OAM sections of TR-101 [4], TR-156 [8], and TR-167 [9]. This TR incorporates that architecture and those requirements by reference. These capabilities span unicast and multicast and also interwork with ATM OAM when that technology is used in the last mile. Ethernet OAM includes continuity checks (CCM), loopback tests (LBM, LBR), and path determination tests (LTR, LTM).

It's possible to use one or more of the Ethernet OAM tools to provide the link availability of an IP Session or many IP Sessions when the L2 and L3 topologies are congruent. That is, Ethernet OAM can test and determine link availability for an Ethernet portion of the network that is used to support one or more IP Sessions. It will not be able to test IP availability though a router or the availability of the IP stack co-located at a Service Edge. Because of these limitations, service providers may choose to detect failures at the Ethernet layer, IP layer, or both. Thus, the following RG requirement applies.

R-58 The RG MUST issue a DHCP renewal message following a random delay between 1 and 30 seconds after it detects a restoration of Ethernet continuity at the customer ME level.

6.2.2 Unidirectional IP Session Monitoring using BFD Echo

A common problem in deployments is IPoE recovering from loss of DHCP and/or ICMPv6 derived state in devices such as Access Nodes and IP Edges. Typically, an RG initiates such state following a link-up event on its WAN interface by sending DHCP and/or ICMPv6 packets. Any event that involves the loss of state on the said upstream nodes but does not translate to a link reset on the RG, results in a situation where traffic sent by the RG in the upstream direction will be dropped, until the state is re-initialized by means of a DHCP and/or ICMPv6 exchange. In the absence of link layer triggers, such re-initialization is dictated by the DHCP and/or ICMPv6 protocol timers, which when based on default values are likely to result in lengthy interruptions of service. An efficient remedy to this problem is to equip the RG with a mechanism that allows it to detect this condition, and treat it as a link flap event that leads to a re-initialization via DHCP and/or ICMPv6. In addition, it is highly desirable for such a mechanism to not require specific support at the IP Edge, nor tax its control plane.

BFD Echo (RFC 5880 [26]) can be used in a unidirectional mode and meets all of the above criteria. This uses a subset of the full BFD protocol. It allows the RG to detect failures in IP connectivity based on the periodic sending of BFD packets on its WAN interface addressed to one of the RG's dynamically assigned IP addresses, or the IPv6 subnet router address. The process of sending a BFD packet that is intended to be sent back to the sender is known as BFD Echo. It should be noted that the only expectation of an IP Edge is for it to route the packet, which will naturally result in routing back towards the sender. Because the IP Edge sees all BFD Echo traffic as user IP traffic, no additional load is placed on the access node or IP Edge's control plane.

It should be noted that the BFD Echo mechanism is only suitable to detect and recover from the loss of IP forwarding connectivity between the RG and IP Edge, as observed and detected by the RG. The full BFD keep-alive mechanism described in Section 6.2.3 is necessary if additional operational requirements need to be met, e.g. the generation of accounting messages at the IP Edge relating to a detected connectivity failure.

Moreover, there are cases where connectivity checking has to occur on a node other than the IP Edge. This may be the case for DS-Lite-based deployment, where connectivity checking has to occur on the CGN node, using basic BFD Echo request/response messages.

R-59 The RG MUST support the configuration of the BFD Echo functionality, as per RFC5881 [27].

R-60 The RG MUST support sending BFD Echo packet(s) on its WAN interface at configurable regular interval with a default value of 30s. The destination IP address of such packets MUST be taken from the list of IP addresses assigned to or via the WAN interface, including the Subnet-Router address of a DHCPv6 delegated prefix.

R-61 The RG MUST support receiving self originated BFD echo packets addressed to its assigned address or the Subnet-Router IPv6 delegated prefix.

R-62 The RG MUST issue a DHCP renewal message after a failure to receive a configurable number of successive BFD Echos. This renewal MUST be delayed by a random time between 1 and 30 seconds.

6.2.3 Bidirectional IP Session Monitoring using BFD Echo Keep-alive

This subsection presents IP Edge and RG requirements relating to the use of Bidirectional Forwarding Detection (BFD) as an IP Session keep-alive. BFD offers an IP Session monitoring mechanism that can detect failure on the path between the IP Edge and a remote IP device.

BFD over IP has the following properties:

- Variable granularity (time between successive ping packets) is part of the BFD protocol; it can be configured dynamically and adjusted based on traffic pattern.
- If desired, periodic polls in one or both directions can be disabled with polling becoming purely event driven (Demand mode).
- If the remote IP device is not interested in monitoring connectivity, the BFD Echo function allows only one side to keep the state of the BFD Session. The processing on the remote peer is limited to turning BFD Echo packets around which could be implemented in the Forwarding plane.
- BFD is designed for very frequent keep-alive messages and exhibits good scaling properties.
- BFD can operate in the context of both IPv4 and IPv6 IP Sessions.
- BFD supports optional authentication that can provide protection from DOS attacks.
- BFD is independent of DHCP, so the solution applies to non-DHCP based situations (e.g. business-DSL with a static IP address, or RIP-based setting, etc).

However, the BFD protocol stack is not likely to be found on most end-hosts.

When an IP Session is established, the IP Edge will create a partner BFD Session if it is configured to do so by the Subscriber keep-alive authorization data. The BFD Session will be initiated in a "Down" state and the IP Edge will assume an active role. If the remote IP device supports BFD, and a successful initial protocol exchange takes place, the BFD Session will transition into INIT and then into UP state, as specified in IETF RFC 5880 [26].

For RGs using addresses assigned by DHCP, BFD initiation must be coordinated with RG configuration through DHCP. As the RG's IP address is not available until the DHCPv4 DHCPACK or DHCPv6 REPLY message is received, BFD is only initiated when DHCP configuration is completed.

6.2.3.1 **IP Edge BFD Keep-alive Requirements**

R-63 The IP Edge MUST support the BFD protocol as per RFC5880 [26].

R-64 The IP Edge MUST take on the Active role during BFD Session initiation.

Note that the Active role is defined in section 6.1 of RFC5880 [26].

R-65 The IP Edge MUST be able to operate in Demand mode.

R-66 The IP Edge MUST support the BFD echo function.

R-67 The IP Edge MUST support either a periodic timer or an inactivity timer tracking incoming Subscriber traffic for each IP Session. This timer MUST be used as the trigger for BFD polls in Demand mode.

R-68 The IP Edge MUST allow IP Sessions to be created and forward traffic prior to successful negotiation of BFD.

R-69 The IP Edge MUST support a configurable upstream inactivity timer as a fallback mechanism in the event of unsuccessful BFD negotiation.

R-70 The IP Edge MUST support a configurable timeout for a BFD session on a given IP Session to establish connectivity with its remote peer. If a BFD session fails to transition to the UP state within that time BFD MUST be automatically disabled on the IP Session.

R-71 The IP Edge MUST be able to configure the DSCP bits of BFD packets.

R-72 The IP Edge MUST be able to configure the Ethernet Priority bits of BFD packets.

R-73 BFD session parameters (e.g., Enable/Disable, Mode, Control / Echo packet transmission intervals, connection failure definition) MUST be able to be provisioned on the IP Edge or communicated from the AAA .

R-74 The IP Edge MUST be able to set the BFD state to AdminDown.

R-75 When a BFD session is in AdminDown state, the IP Edge MUST set Diagnostic code in BFD poll messages to “Path Down”.

6.2.3.2 **RG BFD Keep-alive Requirements**

R-76 RG MUST support the BFD protocol for IP Session Keep-alive. The BFD implementation MUST be compliant with the BFD standard as described in the RFC5880 [26].

R-77 BFD MUST be initiated after both the RG and the IP Edge's IP addresses are available on the RG.

R-78 The RG MUST take on the Passive role during BFD session initiation.

R-79 The RG MUST support BFD Demand mode

R-80 The RG MUST support BFD Asynchronous mode.

R-81 The RG MUST be able to process BFD echo packets in the data plane as specified in RFC5881.

R-82 The RG MUST be able to configure the DSCP bits of BFD packets.

R-83 The RG MUST be able to configure the Ethernet Priority bits of BFD packets.

R-84 The RG SHOULD respond to IP Edge initiated BFD polls using the same DSCP and Ethernet Priority values received in the packet.

R-85 The RG MUST ignore IP packets arriving on the BFD UDP port other than those originating on the IP Edge.

R-86 The BFD configuration on the RG MUST be configurable using TR-069 [2] mechanism.

R-87 When using BFD Demand mode, the RG MUST run an inactivity timer based on the Detect Interval negotiated with the IP Edge.

R-88 When a BFD session on the RG receives a poll with a Diag code set to "Path Down" it MUST perform the following actions:

- Transition into the Down state;
- Respond to the poll with the Diag code set to 3 ("Neighbor Signaled BFD Session Down ")
- Prompt the DHCP client to transition into the Init-Reboot state for DHCPv4 initiated IP Sessions .
- Prompt the DHCP client to send a CONFIRM message for DHCPv6 initiated IP Sessions.

R-89 The RG DHCP client MUST be able to enter DHCPv4 Init-Reboot state or DHCPv6 Confirm state upon detecting that BFD has transitioned into "Down" state.

For correct operation, the BFD session needs to be configured and activated on the RG. TR-069 [2] provides a means of setting up or altering such a BFD configuration on the RG. However since TR-069 uses IP as transport protocol it will only be possible for an operator to configure BFD on an RG only after an IP session is established. Since any subsequent TR-069 configuration may

involve a non-deterministic lag between the IP session establishment and the RG being provisioned for BFD, a Broadband Forum default RG configuration needs to be recommended. Such a default RG configuration will prevent the IP Edge from misinterpreting the situation as lack of BFD support on RG and falling back to another keep-alive.

R-90 The RG MUST use the IP Edge address as the destination for BFD Control packets.

R-91 The RG MUST be able to be pre-provisioned with the following Broadband Forum specified default configuration.

- Version (1)
- Control Plane Independent (0)
- Authentication Present (0)
- Demand (1)
- Detect Multiplier (3)
- Local Discriminator (a random 32-bit value)
- Desired Minimum Transmit Interval (1,000,000)
- Required Minimum Receive Interval (1,000,000)
- Required Minimum Echo Receive Interval (0)
- State (Down)

In summary, this default configuration specifies that the RG supports BFD Demand mode, does not support Echo packets and does not support BFD Authentication. The minimum supported intervals between polls to and from IP Edge are both set to 1 second. The number of retries is set to 3. The Local Discriminator is a unique identifier of a BFD session used by IP Edge to de-multiplex BFD packets from multiple hosts. It can be pre-provisioned to a random value or generated on demand using some random number generator on CPE. The initial state of the BFD session is "Down". RFC 5880 [26] can be consulted for more detailed definitions of these parameters.

6.2.4 ARP Keep-alive

R-92 The IP Edge SHOULD support a configurable ARP ping IP session keep-alive mechanism for active monitoring of an IP RG associated with the IP session.

6.2.5 IP Session Monitoring with IPv6 Neighbor Unreachability Detection

IPv6 has a native mechanism called Neighbor Unreachability Detection (NUD) which enables an IPv6 node to track the status of neighbor nodes using the ND protocol. This method is used by IPv6 hosts to maintain adjacency state with its neighbors by updating their neighbor cache tables. Since this is passive monitoring, when a host/node is not sending any traffic, a link failure cannot be detected by this method.

Neighbor Unreachability Detection is used by Residential Gateway and the IPv6 IP Edges they are connected to.

An RG keeps the state of its adjacent IP Edge by sending unicast Neighbor Solicitation messages to verify that the path between itself and the IP Edge node is still available and the neighbor is "Reachable". The receiving node (IP Edge) responds by sending a Neighbor Advertisement to the sender (RG) which, upon receiving the Neighbor Advertisement, confirms that the forwarding path is valid. The IP Edge node will also send a Neighbor Solicitation message to the RG which will respond by sending a Neighbor Advertisement message. This ensures that both neighbors maintain state and that both one-way paths are valid. Each neighbor maintains state from its own perspective.

R-93 The IP Edge MUST support IPv6 Network Unreachability Detection for IP Sessions (RFC 4861 [23]).

R-94 The RG MUST support IPv6 Network Unreachability Detection for IP Sessions (RFC 4861 [23]).

6.3 Change of Policies using RECONFIGURE or DHCPFORCERENEW messages

DHCPv4 DHCPFORCERENEW mechanism, defined in RFC 6704 [30], provides a method for performing change of policies on IP Sessions. These changes are triggered by the DHCP server and that makes this mechanism especially useful with long DHCP lease times.

DHCPFORCERENEW messages allow an operator to force an IP Session to renew its IP lease. At the same time policy changes can be made, including the following:

- Changing the IP address of an IP Session,
- Terminating an IP Session,
- Forcing an IP Session re-authentication or reauthorization,
- Changing the IP Session's service policy (e.g. a change of QoS, redirection, etc)

The application of DHCPFORCERENEW is limited by a requirement that DHCPFORCERENEW message is always authenticated using procedures as described in RFC3118 [19]. Authentication for DHCP is mandatory for DHCPFORCERENEW, however as currently defined in Section 5 of RFC 3118[19] it requires distribution of token or shared-secret out-of-band to DHCP clients. In addition, some protection is provided by the split-horizon architecture in TR-059 [1] and TR-101 [3] – essentially it is highly improbable to snoop or spoof neighbor DHCP messages, and for some implementations, this protection may suffice.

Additional protection can be provided by an authentication protocol for DHCPv4. The DHCPFORCERENEW mechanism is based on a DHCPFORCERENEW Key that is exchanged between the DHCP server and DHCP client in the initial DHCPACK and is used for verification of any

subsequent DHCP DHCPFORCERENEW messages. The server then includes an HMAC computed from the DHCPFORCERENEW Key in subsequent DHCPFORCERENEW messages.

Both the DHCPFORCERENEW Key sent from the server to the client and the HMAC in subsequent DHCPFORCERENEW messages are carried as the Authentication information in a DHCP Authentication option. DHCPFORCERENEW Key Authentication follows the model set forward in DHCPv6 (RFC3315 [20]) as the Reconfigure Key Authentication Protocol.

6.3.1 Requirements for IPv4

The following requirements are apply to the RG:

R-95 The RG MUST support the use of DHCPFORCERENEW (RFC 6704 [30]) for changing the configuration parameters or the IP address associated with an IP Session.

R-96 The RG MUST indicate support for RFC 6704 [30] by sending the DHCPFORCERENEW_NONCE_CAPABLE. option in the DHCPDISCOVER and in the DHCPREQUEST messages.

R-97 The RG MUST support using the DHCPFORCERENEW Key for validating DHCPFORCERENEW messages received from the DHCP Server, as per RFC 6704 [30].

The following requirements apply to the IP Edge:

R-98 The DHCP Server or the DHCP Proxy on the IP Edge MUST indicate support of RFC 6704 [30] by sending the DHCPFORCERENEW_KEY_CAPABLE option in the DHCPOFFER.

R-99 The DHCP Server on the IP Edge MUST support selecting a DHCPFORCERENEW Key and sending it to the DHCP client in the DHCP Authentication option in the DHCPACK message, as per RFC 6704 [30].

R-100 The DHCP Server on the IP Edge MUST support computing an HMAC-MD5 of the DHCPFORCERENEW message using the DHCPFORCERENEW Key and inserting the result in the authentication information field in the Authentication option included in the DHCPFORCERENEW message sent to the client, as per RFC 6704 [30].

6.3.2 Requirements for IPv6

The following requirements are apply to the RG:

R-101 The RG MUST support the use of DHCPv6 reconfigure as per RFC 3315 [20] (for changing the configuration parameters or the IP address associated with an IP Session).

R-102 The RG MUST support using the reconfigure Key for validating DHCP RECONFIGURE messages received from the DHCP Server, as per RFC 3315 [20].

The following requirements apply to the IP Edge when acting as DHCP server

R-103 The DHCP Server MUST support selecting a RECONFIGURE Key and sending it to the DHCP client in the DHCP Authentication option in the DHCP REPLY message, as per RFC 3315 [20].

R-104 The DHCP Server MUST support computing an HMAC-MD5 of the RECONFIGURE message using the RECONFIGURE Key and inserting the result in the authentication information field in an AUTHENTICATION option included in the RECONFIGURE message sent to the client, as per RFC 3315 [20].

6.4 Traffic Policies for Subscriber Sessions

Traffic policies can be applied to Subscriber Sessions. A simple exemplary policy would be to drop all traffic on un-authorized Subscriber Sessions and to forward all traffic on authorized Sessions. Service Edges however need to support the definition of much more detailed traffic policies.

A Flow is defined by a traffic classifier and is used as the conditional element of a traffic policy activated on a Subscriber Session. A traffic policy could be activated/instantiated for a Subscriber Session based on the overall network service associated with the Subscriber Session, e.g. "UDP packets to destination XYZ are to be policed to 64kbps".

Multiple Subscriber Sessions may share the same traffic policy definitions, however each such policy will be instantiated only in the context of its Subscriber Session, i.e. traffic policies are not aggregated across Subscriber Sessions that share an IP flow traffic classifier definition, or common actions. The exception to this is when one Subscriber Session contains another, as is the case for IP Sessions contained in an Ethernet Session, or a 6rd Session contained in an IPv4 Session. In these cases, any policies applied to the container Subscriber Session can be applied to all contained Subscriber Sessions.

6.4.1 Ethernet Traffic Classifier

A traffic policy is constructed using the Ethernet Flow classifier as its condition match, and a set of actions. The policy is applied in the context of the Subscriber Session either directly at Session activation or dynamically following a Subscriber or operator network service activation event, e.g. following the activation of a turbo button for some specific traffic category. Ethernet Flow classifier definitions are expected to be either predefined or downloaded to the Service Edge from

a repository during Subscriber logon, network service activation, or both. The Flow classifier parameter definition can be passed directly, or referenced to a named template.

A network element providing this functionality needs to support the definition of traffic classifiers allowing the use of one or more of the following:

1. The destination MAC address
2. The source MAC address
3. The C-VID
4. The S-VID
5. The Ethertype – which MUST include PPPoE (0x8863, 0x8864), IPv4oE (0x0800), IPv6oE (0x086DD), and ARP (0x0806).
6. The Ethernet Priority bits
7. The Ethernet DEI bit

6.4.2 IP Traffic Classifier

A traffic policy is constructed using the IP Flow as its condition match and a set of actions. The policy is applied in the context of the Subscriber Session either directly at session activation or dynamically following a Subscriber or operator network service activation event, e.g. following the activation of a turbo button for some specific traffic category. IP Flow definitions are expected to be either predefined or downloaded to the Service Edge from a repository during Subscriber logon, network service activation, or both. The IP Flow parameter definition can be passed directly, or referenced to a named template.
A network element aiming to provide this functionality needs to support the definition of traffic classifiers allowing the use of one or more of the following traffic parameters (RFC 6088 [28]):

1. The destination IP address
2. The source IP address
3. The transport protocol
4. The transport Protocol destination port
5. The transport protocol source port

A situation may arise when multiple traffic policies are applied to a Subscriber Session, with one or more of the policies' IP Flow classifier definitions overlapping each other. In order to deal with this it is desirable to have a method of organizing the traffic policy classifiers so as to resolve potential overlaps. This can be achieved using a simple priority scheme with traffic matching the highest priority policy.

R-105 The Service Edge MUST be able to associate traffic policies with a Subscriber Session.

R-106 The IP Flow classifiers of a traffic policy MUST be configurable and allow classification based on any combination of:

- DSCP for IPv4 and IPv6,
- Incoming port/interface PPP Session using the FQDN/NAI,
- Source IP address,
- Destination IP address,
- IP Protocol,
- Source TCP/UDP port,
- Destination TCP/UDP port,
- Packet length.

Although not part of the scope of TR-146, higher protocol layer traffic classifiers can be conceived, and could be used in a similar manner.

As with Subscriber Session accounting, it is desirable to have a means of accounting for traffic corresponding to a given Flow.

R-107 The Service Edge MUST support collecting accounting statistics on a per Flow basis.

R-108 The Service Edge MUST be able to activate, deactivate or modify traffic policies of an existing Subscriber Session.

6.5 **Subscriber Session Grouping**

It is desirable to have a mechanism for Grouping Subscriber Sessions which are logically linked to each other into a Group which contains all the traffic delivered to that Group; this allows the application of a common set of policies. This is the case if a single Subscriber access line had several identifiable IP hosts connected to it (e.g. via a Layer 2 RG), each host represented by a Subscriber Session, yet all these Subscriber Sessions can be bound together logically by the access line id.

There are two categories of Grouped Subscriber Sessions:

a) All of the Subscriber Session traffic is presented on the Service Edge via a common physical or logical interface that corresponds directly to the desired Subscriber Session Group, e.g. a 1:1 VLAN. In this case the logical interface provides the Subscriber Session Group construct.

b) Subscriber Sessions from multiple Subscribers are presented over a shared logical interface, e.g. an N:1 VLAN. In this case the Subscriber Sessions can be grouped according to their RG MAC address or access-line circuit-id parameter, passed in a DHCPv4 Option-

82, or DHCPv6 Option 18 "interface_id" (i.e. the RG upstream interface ID) or a PPPoE circuit-id Tag.

c) Combinations of a) and b) may also exist.

Given these two categories, the following requirements are needed:

R-109 The IP Edge MUST support the logical grouping of IP Sessions based on DHCPv4 Option-82, sub-option "Agent Circuit ID".

R-110 The IP Edge MUST support logical grouping of IP Sessions based on PPPoE circuit-id tag.

R-111 The IP Edge MUST support logical grouping of IP Sessions based on DHCPv6 option 18 "interface_id".

R-112 The IP Edge MUST support the logical grouping of Subscriber Sessions based on a combination of logical interfaces.

R-113 The IP Edge MUST support the logical grouping of Subscriber Sessions based on a combination of authenticated client information.

R-114 The Ethernet Edge MUST support the logical grouping of Ethernet Sessions based on a combination of physical interfaces.

R-115 The Ethernet Edge MUST support the logical grouping of Ethernet Sessions based on a combination of VLANs.

R-116 The Ethernet Edge MUST support the logical grouping of Ethernet Sessions based on a combination of authenticated client MAC address

R-117 The Service Edge MUST support the application of Subscriber Session control and data plane policies to a Subscriber Session Group.

Introducing IPv6 increases the number of possible scenarios where Subscriber Session Grouping capability may be desired. In particular, in a dual stack IPv4/IPv6 environment there are scenarios where different QoS policies may be required for different logical Groups of Subscriber Sessions. In this case the following logical Groups of Subscriber Sessions can be identified:

1. a Group of IPv4-only IP Sessions;
2. a Group of IPv6-only IP Sessions (e.g. as a consequence of Prefix Delegation);
3. a Group of all IPv4 and IPv6 IP Sessions sharing the same access line.

A group of Subscriber Sessions can fate share with an anchor Subscriber Session, which is typically instantiated in the underlying connectivity.
It should be noted that the network topology for individual Subscriber Sessions may differ, and Session Groups may also differ, depending on the point in the network where policies or profiles are applied. For example, the Access Node may see Ethernet and IP Sessions, but may not "see"

6rd or DSLite Sessions. A BNG behind A-10 may not see the Ethernet or IP Sessions, but may see the 6rd or DSLite Sessions. So, in every case, the Subscriber Session Groups are within the context of visible traffic at a specific point.

A mechanism to identify and create the three different types of Subscriber Session Groups is required.
For IP Sessions, DHCPv4 option 82 sub-option "Agent Circuit ID" and sub-option "Agent Remote ID" inserted by Access Node can be used to bundle together IP Sessions belonging to the same access line at the BNG (see TR-101 [4] sections 3.8.1, 3.9.1 and Appendix B). IP Sessions can be bundled using one of the following DHCPv6 options:

- option 18 "Interface-id" (RFC 3315 [20]) can be added by a Lightweight DHCPv6 Relay Agent (see RFC 6221 [29]) to identify the interface on which the client message was received;
- option 37 "Relay Agent Remote-ID" (RFC 4649 [22]) can be added by a Lightweight DHCPv6 Relay Agent to identify the remote host end of the circuit.

In order to be able to bundle together IPv4 and IPv6 IP Sessions, a mechanism that recognizes and Groups together access line information provided by DHCPv4 and DHCPv6 must be provided on the IP Edge.

R-118 The IP Edge MUST support the creation of an IP Session Group based on DHCPv6 Interface-ID.

R-119 The Service Edge MUST support the creation of a Subscriber Session Group based on a common logical and/or physical interface(e.g. 1:1 VLAN on a given physical port).

R-120 The IP Edge MUST support the creation of a Subscriber Session Group based on a Remote-ID Options carried in RELAY-FORW messages.

R-121 The IP Edge MUST support the logical Grouping of IPv4 and IPv6 IP Sessions based on a common DHCPv4 / DHCPv6 derived remote-id.

R-122 The Service Edge MUST support the concurrent enforcement of traffic policies (e.g. QoS, filtering, bandwidth throttling) on a session group basis and individual Subscriber Sessions within Session Group.

R-123 The group policy MUST take precedence.

6.5.1 Grouped Authentication and Authorization

When an IP Subscriber Session is grouped with other IP Sessions, it may be advantageous for optimization reasons, to skip authentication exchanges with the AAA server once one of the Sessions in that Group has already been authenticated. For instance, consider the case where a

Subscriber establishes an IPv4 IP Session and authentication is successfully performed. When the same Subscriber establishes an IPv6 IP Session that is grouped with the previous one (e.g. according to the line-id) then this IPv6 IP Session would be implicitly authenticated (i.e. locally authenticated by the Service Edge) thanks to the previous authentication of the IPv4 IP Session. The model for such authentication is to identify the IP Session that is to be authenticated (the "master session").

If no such IP session is identified, then each and every IP session will be authenticated and accounted for individually. This allows service providers to choose the binding among IP Sessions, and specify the behavior when one or more Sessions within a Group are established or terminated. For example, when an Ethernet Session is authenticated using 802.1X [12], additional IPv4 and IPv6 IP Sessions can be grouped without any additional authentication. If either IP Session is dropped, then the other continues. If the Ethernet Session is dropped, then both IP Sessions are dropped. Similarly, if a service provider prefers to authenticate using DHCP, it is still possible to authenticate the Ethernet Session using that method, and this choice would allow independent establishment and termination of IPv4 and IPv6 IP Sessions once that authentication has been completed and until the Ethernet Session drops.

Alternately, if a service provider wished to ensure that IPv4 and IPv6 IP Sessions were established and terminated in concert, they could bind IPv6 to IPv4 (or the reverse) and whenever the authenticated IP Session drops, those grouped IP Sessions that share that authentication are also terminated. Finally, when a Subscriber Session is contained within another, then it is terminated when the container Subscriber Session terminates. For example, a 6rd Session will terminate when the IPv4 IP Session it relies on terminates and that IPv4 IP Session will terminate when the Ethernet Session it relies on terminates.

R-124 The Service Edge MUST implicitly authenticate and authorize any individual Subscriber Session within a given group if the master Subscriber Session in that Group has already been authenticated and authorized.

R-125 The Service Edge MUST NOT authenticate and authorize any individual Subscriber Session within a given group if the master Subscriber Session in that Group has not been authenticated and authorized.

R-126 If no master Subscriber Session has been designated for a given Subscriber Session Group, the Service Edge MUST authenticate and authorize each and every Session individually.

R-127 When a master Subscriber Session terminates, the Service Edge MUST terminate all the other Subscriber Sessions in that Group.

R-128 Any given Subscriber Session, MUST only be assigned to a single Group.

6.5.2 Grouped Accounting and Reporting

Accounting and reporting can be performed on a Subscriber Session Group basis. This is useful, for instance, to perform volume-based accounting related to the total volume of all Subscriber Sessions established by a Subscriber.

R-129 The Service Edge MUST be able to perform accounting and reporting to the AAA server on a per Subscriber Session Group basis.

R-130 The volume information sent to the AAA server MUST be the total volume of all the Subscriber Sessions within a given Group

R-131 An Accounting-Stop message pertaining to a Subscriber Session Group MUST only be sent when the Subscriber Session to which the authentication is bound is terminated.

6.5.3 Traffic Policies for Grouped Subscriber Sessions

R-132 The IP Edge MUST support traffic policies for a given Group with classification rules that can be IPv4-specific (e.g. IPv4 subnet), IPv6-specific (e.g. IPv6 prefix) or apply to both IPv4 and IPv6.

R-133 When receiving authorization information from the AAA server, the Service Edge MUST support the activation of traffic policies for the authorized Subscriber Session and all the Subscriber Sessions that will be Grouped with it.

End of Broadband Forum Technical Report TR-146