Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 22, 2015

M. Boucadair, Ed.
D. Binet
S. Durel
B. Chatras
France Telecom
T. Reddy
Cisco Systems
B. Williams
Akamai, Inc.
B. Sarikaya
L. Xue
Huawei
R. Wheeldon
Cisco Systems

July 21, 2014

Scenarios with Host Identification Complications
draft-boucadair-intarea-host-identifier-scenarios-07

Abstract

This document describes a set of scenarios in which complications to identify which policy to apply for a host are encountered. This problem is abstracted as "host identification". The document does not include any solution-specific discussion.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 22, 2015.

Table of Contents

1. Introduction

The goal of this document is to enumerate scenarios which encounter the issue of uniquely identifying a host among those sharing the same IP address. An exhaustive list of encountered issues for the Carrier Grade NAT, A+P, and Application Proxies scenarios are documented in [RFC6269]. In addition to those issues, some of the scenarios described in this document suffer from additional issues such as:

- Identify which policy to enforce for a subscriber/UE (e.g., limit access to the service based on some counters such as volume-based service offering); enforcing the policy will have impact on all hosts sharing the same IP address.
- Need to correlate between the internal address:port and external address:port to generate and therefore to enforce policies.
- Query a location server for the location of an emergency caller based on the source IP address.

The analysis of the scenarios listed in this document indicates several root causes for the host identification issue:

1. Presence of address sharing (NAT, A+P, application proxies, etc.).
2. Use of tunnels between two administrative domains.
3. Combination of address sharing and presence of tunnels in the path.

2. **What is in?** What is out?

The goal of this document is to identify scenarios the authors are aware of and which share the same complications to identify which policy to apply for a host. This problem is abstracted as host identification problem.

This document can be used as a tool to design solution(s) mitigating the encountered issues. Describing the scenario allows to identify what is common between the scenarios and then would help during the solution design phase. Note, [RFC6967] focuses only on the CGN, A+P, and application proxies cases. The analysis in [RFC6967] may not be accurate for some of the scenarios that do not span multiple administrative domains (e.g., Section 10.1).

This document does not target means that would lead to expose a host (or a user) beyond what the original packet, issued from that host, would have already exposed. Such means are not desirable, nor required to solve the issues encountered in the scenarios discussed in this document. The focus is exclusively on means to restore the information conveyed in the original packet issued by a given host. These means are intended to help in identifying which policy to apply for a given flow. These means rely on some bits of the source IP address and/or port number(s) used by the host to issue packets. To prevent side effects and mis-uses of such means on privacy, solution specification document(s) should explain, in addition to what is already documented in [RFC6967], the following:

- To what extent the solution can be used to nullify the effect of using address sharing to preserve privacy (see for example

[EFFOpenWireless]). Note, this concern can be mitigated if the address sharing platform is under the responsibility of the host's owner and the host does not leak information that would interfere with its privacy protection tool.

- To what extent the solution can be used to expose privacy information beyond what the original packet would have already exposed. Note, the solutions discussed in [RFC6967] do not allow to reveal extra information than what is conveyed in the original packet.

This document does not include any solution-specific discussion. In particular, the document does not elaborate whether explicit authentication is enabled or not. Moreover, this document does not discuss whether specific information is needed to be leaked in packets, whether this is achieved out-of-band, etc. Those considerations are out of scope.

3. Scenario 1: CGN

Several flavors of stateful CGN have been defined. A non-exhaustive list is provided below:

1. NAT44 ([RFC6888], [I-D.tsou-stateless-nat44])
2. DS-Lite NAT44 [RFC6333]
3. NAT64 [RFC6146]
4. NPTv6 [RFC6296]

As discussed in [RFC6967], remote servers are not able to distinguish between hosts sharing the same IP address (Figure 1). As a reminder, remote servers rely on the source IP address for various purposes such as access control or abuse management. The loss of the host identification will lead to issues discussed in [RFC6269].

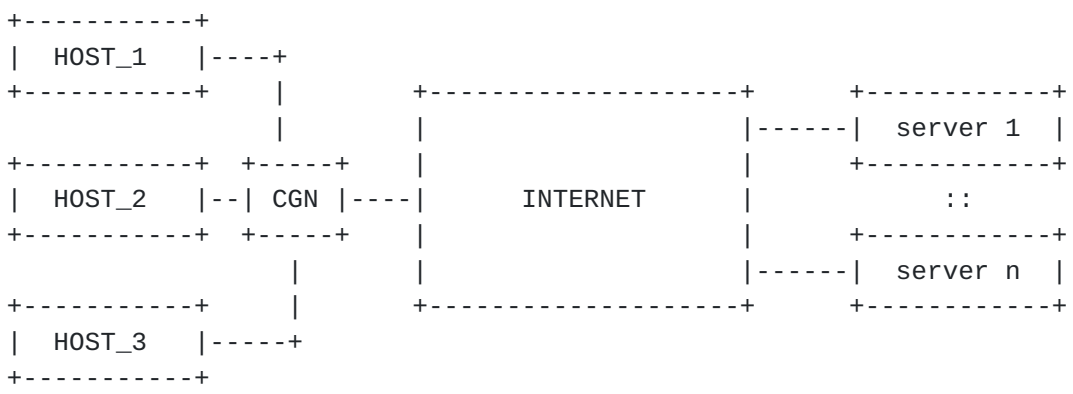

Figure 1: CGN Reference Architecture

Some of the above referenced CGN scenarios will be satisfied by eventual completion of the transition to IPv6 across the Internet (e.g., NAT64), but this is not true of all CGN scenarios (e.g. NPTv6 [RFC6296]) for which some of the issues discussed in [RFC6269] will be encountered (e.g., impact on geolocation).

Privacy-related considerations discussed in [RFC6967] apply for this scenario.

4. Scenario 2: A+P

A+P [RFC6346][I-D.ietf-softwire-map][I-D.ietf-softwire-lw4over6] denotes a flavor of address sharing solutions which does not require any additional NAT function be enabled in the service provider's network. A+P assumes subscribers are assigned with the same IPv4 address together with a port set. Subscribers assigned with the same IPv4 address should be assigned non overlapping port sets. Devices connected to an A+P-enabled network should be able to restrict the IPv4 source port to be within a configured range of ports. To forward incoming packets to the appropriate host, a dedicated entity called PRR (Port Range Router, [RFC6346]) is needed (Figure 2).

Similar to the CGN case, remote servers rely on the source IP address for various purposes such as access control or abuse management. The loss of the host identification will lead to the issues discussed in [RFC6269]. In particular, it will be impossible to identify hosts sharing the same IP address by remote servers.

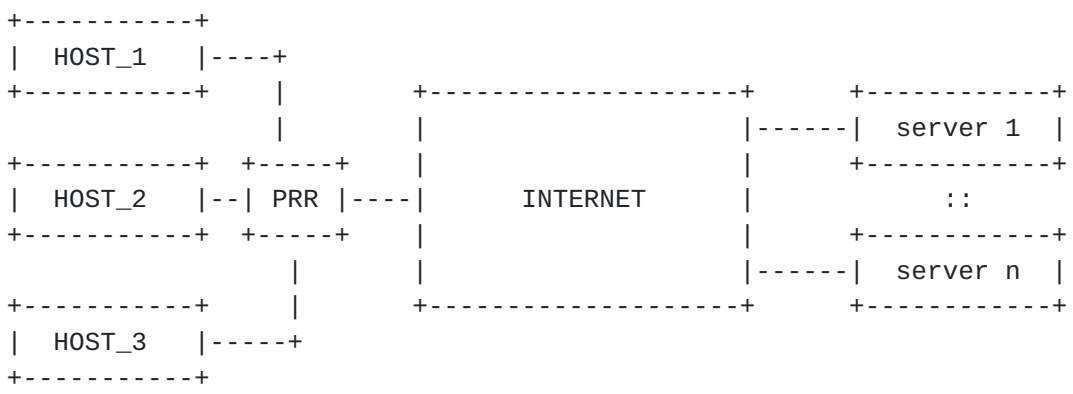

Figure 2: A+P Reference Architecture

Privacy-related considerations discussed in [RFC6967] apply for this scenario.

5. Scenario 3: On-Premise Application Proxy Deployment

This scenario is similar to the CGN scenario. Remote servers are not able to distinguish hosts located behind the PROXY. Applying policies on the perceived external IP address as received from the PROXY will impact all hosts connected to that PROXY.

Figure 3 illustrates a simple configuration involving a proxy. Note several (per-application) proxies may be deployed. This scenario is a typical deployment approach used within enterprise networks.

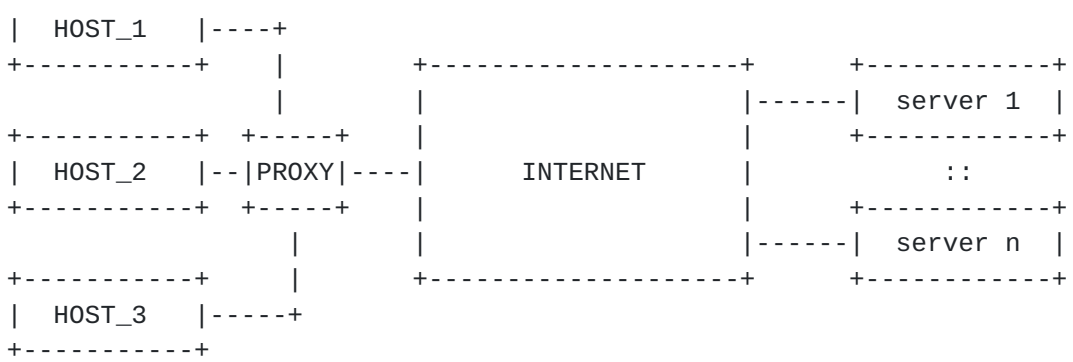

Figure 3: Proxy Reference Architecture

The administrator of the proxy may have many reasons for wanting to proxy traffic - including caching, policy enforcement, malware scanning, reporting on network or user behavior for compliance or security monitoring. The same administrator may also wish to

selectively hide or expose the internal host (or user) identity to servers. He/she may wish to hide the identity to protect end-user privacy or to reduce the ability of a rogue agent to learn the internal structure of the network. He/she may wish to allow upstream servers to identify hosts (and/or users) to enforce access policies (for example on documents or online databases), to enable account identification (on subscription-based services) or to prevent spurious misidentification of high traffic patterns as a DoS attack. Application-specific protocols exist for enabling such forwarding on some plain-text protocols (e.g., Forwarded headers on HTTP [RFC7239] or time-stamp-line headers in SMTP [RFC5321]).

Servers not receiving such notifications but wishing to perform host or user-specific processing are obliged to use other application-specific means of identification (e.g. Cookies [RFC6265]).

Packets/connections must be received by the proxy regardless of the IP address family in use. The requirements of this scenario are not satisfied by eventual completion of the transition to IPv6 across the Internet. Complications will arise for both IPv4 and IPv6.

Privacy-related considerations discussed in [RFC6967] apply for this scenario.

6. Scenario 4: Distributed Proxy Deployment

This scenario is similar to the proxy deployment scenario (Section 5) with the same use-cases. However, in this instance part of the functionality of the application proxy is located in a remote site. This may be desirable to reduce infrastructure and administration costs or because the hosts in question are mobile or roaming hosts tied to a particular administrative zone of control but not to a particular network.

In some cases, a distributed proxy is required to identify an account holder on whose behalf it is performing the caching, filtering or other desired service - for example to know which policies to enforce. Typically, IP addresses are used as a surrogate. However, in the presence of CGN, this identification becomes difficult. Alternative solutions include the use of cookies, which only work for HTTP traffic, tunnels or proprietary extensions to existing protocols.

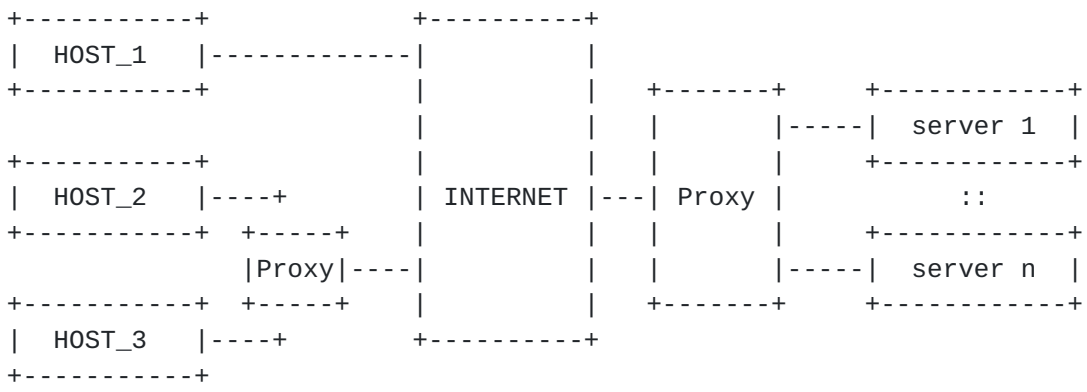

Figure 4: Distributed Proxy Reference Architecture (1)

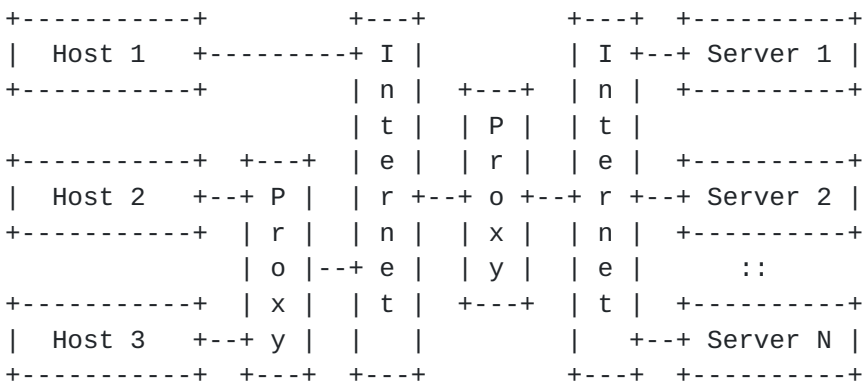

Figure 5: Distributed Proxy Reference Architecture (2)

Packets/connections must be received by the proxy regardless of the IP address family in use. The requirements of this scenario are not satisfied by eventual completion of the transition to IPv6 across the Internet. Complications will arise for both IPv4 and IPv6.

If the proxy and the servers are under the responsibility of the same administrative entity (Figure 4), no privacy concerns are raised. Nevertheless, privacy-related considerations discussed in [RFC6967] apply if the proxy and the servers are not managed by the same administrative entity (Figure 5).

7. Scenario 5: Overlay Network

An overlay network is a network of machines distributed throughout multiple autonomous systems within the public Internet that can be used to improve the performance of data transport (see Figure 6). IP packets from the sender are delivered first to one of the machines that make up the overlay network. That machine then relays the IP

packets to the receiver via one or more machines in the overlay network, applying various performance enhancement methods.

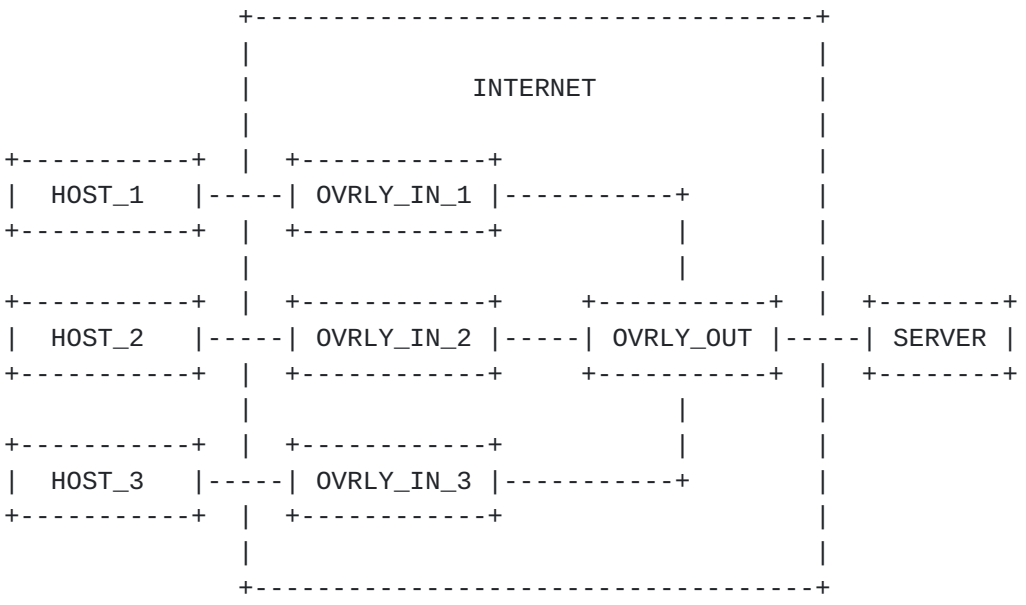

Figure 6: Overlay Network Reference Architecture

Such overlay networks are used to improve the performance of content delivery [IEEE1344002]. Overlay networks are also used for peer-to-peer data transport [RFC5694], and they have been suggested for use in both improved scalability for the Internet routing infrastructure [RFC6179] and provisioning of security services (intrusion detection, anti-virus software, etc.) over the public Internet [IEEE101109].

In order for an overlay network to intercept packets and/or connections transparently via base Internet connectivity infrastructure, the overlay ingress and egress hosts (OVERLAY_IN and OVERLAY_OUT) must be reliably in-path in both directions between the connection-initiating HOST and the SERVER. When this is not the case, packets may be routed around the overlay and sent directly to the receiving host.

For public overlay networks, where the ingress and/or egress hosts are on the public Internet, packet interception commonly uses network address translation for the source (SNAT) or destination (DNAT) addresses in such a way that the public IP addresses of the true endpoint hosts involved in the data transport are invisible to each other (see Figure 7). For example, the actual sender and receiver may use two completely different pairs of source and destination addresses to identify the connection on the sending and receiving networks in cases where both the ingress and egress hosts are on the public Internet.

```
           ip hdr contains:                ip hdr contains:
SENDER -> src = sender   --> OVERLAY --> src = overlay2  --> RECEIVER
          dst = overlay1                  dst = receiver
```

Figure 7: NAT operations in an Overlay Network

In this scenario, the remote server is not able to distinguish among hosts using the overlay for transport. In addition, the remote server is not able to determine the overlay ingress point being used by the host, which can be useful for diagnosing host connectivity issues.

In some of the above referenced scenarios, IP packets traverse the overlay network fundamentally unchanged, with the overlay network functioning much like a CGN (Section 3). In other cases, connection-oriented data flows (e.g. TCP) are terminated by the overlay in order to perform object caching and other such transport and application layer optimizations, similar to the proxy scenario (Section 5). In both cases, address sharing is a requirement for packet/connection interception, which means that the requirements for this scenario are not satisfied by the eventual completion of the transition to IPv6 across the Internet.

More details about this scenario are provided in [I-D.williams-overlaypath-ip-tcp-rfc].

This scenario does not introduce privacy concerns since the identification of the host is local to a single administrative domain (i.e., CDN overlay Network) or passed to a remote server to help forwarding back the response to the appropriate host.

8. Scenario 6: Policy and Charging Control Architecture

This issue is related to the framework defined in [TS23.203] when a NAT is located between the PCEF (Policy and Charging Enforcement Function) and the AF (Application Function) as shown in Figure 8.

The main issue is: PCEF, PCRF and AF all receive information bound to the same UE(User Equipment) but without being able to correlate between the piece of data visible for each entity. Concretely,

- PCEF is aware of the IMSI (International Mobile Subscriber Identity) and an internal IP address assigned to the UE.

- AF receives an external IP address and port as assigned by the NAT function.

- PCRF is not able to correlate between the external IP address/port assigned by the NAT (received from the AF) and the internal IP address and IMSI of the UE (received from the PCEF).

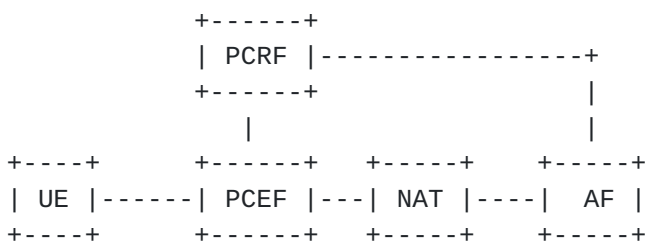

Figure 8: NAT located between AF and PCEF

This scenario can be generalized as follows (Figure 9):

- Policy Enforcement Point (PEP, [RFC2753])

- Policy Decision Point (PDP, [RFC2753])

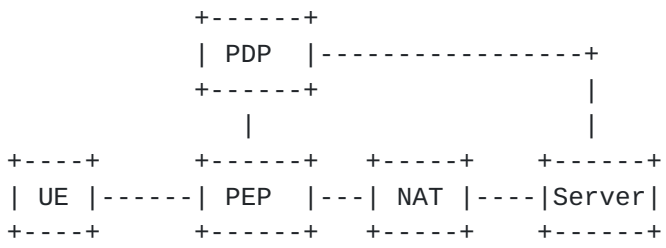

Figure 9: NAT located between PEP and Server

Note that an issue is encountered to enforce per-UE policies when the NAT is located before the PEP function (see Figure 10):

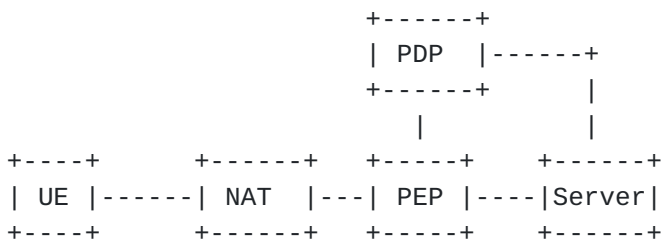

Figure 10: NAT located before PEP

This scenario does not introduce privacy concerns since the identification of the host is local to a single administrative domain and is meant to help identifying which policy to select for a UE.

9. Scenario 7: Emergency Calls

Voice service providers (VSPs) operating under certain jurisdictions are required to route emergency calls from their subscribers and have to include information about the caller's location in signaling messages they send towards PSAPs (Public Safety Answering Points, [RFC6443]), via an Emergency Service Routing Proxy (ESRP, [RFC6443]). This information is used both for the determination of the correct PSAP and to reveal the caller's location to the selected PSAP.

In many countries, regulation bodies require that this information be provided by the network rather than the user equipment, in which case the VSP needs to retrieve this information (by reference or by value) from the access network where the caller is attached.

This requires the VSP call server receiving an emergency call request to identify the relevant access network and to query a Location Information Server (LIS) in this network using a suitable look-up key. In the simplest case, the source IP address of the IP packet carrying the call request is used both for identifying the access network (thanks to a reverse DNS query) and as a look-up key to query the LIS. Obviously the user-id as known by the VSP (e.g., telephone number, or email-formatted URI) can't be used as it is not known by the access network.

The above mechanism is broken when there is a NAT between the user and the VSP and/or if the emergency call is established over a VPN tunnel (e.g., an employee remotely connected to a company VoIP server through a tunnel wishes to make an emergency call). In such cases, the source IP address received by the VSP call server will identity the NAT or the address assigned to the caller equipment by the VSP (i.e., the address inside the tunnel). This is similar to the CGN case (Section 3) and overlay network case (Section 7) and applies irrespective of the IP versions used on both sides of the NAT and/or inside and outside the tunnel.

Therefore, the VSP needs to receive an additional piece of information that can be used to both identify the access network where the caller is attached and query the LIS for his/her location. This would require the NAT or the Tunnel Endpoint to insert this extra information in the call requests delivered to the VSP call servers. For example, this extra information could be a combination of the local IP address assigned by the access network to the caller's equipment with some form of identification of this access network.

However, because it shall be possible to setup an emergency call regardless of the actual call control protocol used between the user

and the VSP (e.g., SIP [RFC3261], IAX [RFC5456], tunneled over HTTP, or proprietary protocol, possibly encrypted), this extra information has to be conveyed outside the call request, in the header of lower layers protocols.

Privacy-related considerations discussed in [RFC6967] apply for this scenario.

10. Other Deployment Scenarios

10.1. Scenario 8: Open WLAN or Provider WLAN

In the context of Provider WLAN, a dedicated SSID can be configured and advertised by the RG (Residential Gateway) for visiting terminals. These visiting terminals can be mobile terminals, PCs, etc.

Several deployment scenarios are envisaged:

1. Deploy a dedicated node in the service provider's network which will be responsible to intercept all the traffic issued from visiting terminals (see Figure 11). This node may be co-located with a CGN function if private IPv4 addresses are assigned to visiting terminals. Similar to the CGN case discussed in Section 3, remote servers may not be able to distinguish visiting hosts sharing the same IP address (see [RFC6269]).

2. Unlike the previous deployment scenario, IPv4 addresses are managed by the RG without requiring any additional NAT to be deployed in the service provider's network for handling traffic issued from visiting terminals. Concretely, a visiting terminal is assigned with a private IPv4 address from the IPv4 address pool managed by the RG. Packets issued form a visiting terminal are translated using the public IP address assigned to the RG (see Figure 12). This deployment scenario induces the following identification concerns:

   * The provider is not able to distinguish the traffic belonging to the visiting terminal from the traffic of the subscriber owning the RG. This is needed to identify which policies are to be enforced such as: accounting, DSCP remarking, black list, etc.

   * Similar to the CGN case Section 3, a misbehaving visiting terminal is likely to have some impact on the experienced service by the subscriber owning the RG (e.g., some of the issues are discussed in [RFC6269]).

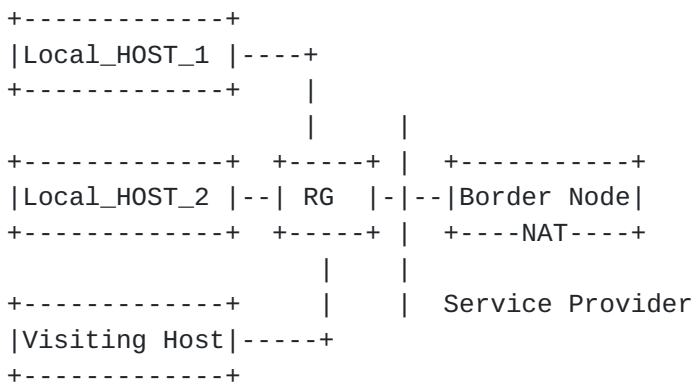

Figure 11: NAT enforced in a Service Provider's Node

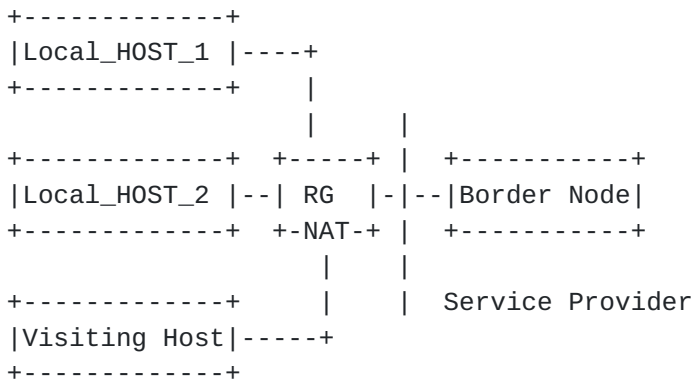

Figure 12: NAT located in the RG

This scenario does not introduce privacy concerns since the identification of the host is local to a single administrative domain and is meant to help identifying which policy to select for a visiting UE.

10.2. Scenario 9: Cellular Networks

Cellular operators allocate private IPv4 addresses to mobile terminals and deploy NAT44 function, generally co-located with firewalls, to access to public IP services. The NAT function is located at the boundaries of the PLMN (Public Land Mobile Network). IPv6-only strategy, consisting in allocating IPv6 prefixes only to mobile terminals, is considered by various operators. A NAT64 function is also considered in order to preserve IPv4 service continuity for these customers.

These NAT44 and NAT64 functions bring some issues very similar to those mentioned in Figure 1 and Section 8. This issue is

particularly encountered if policies are to be applied on the Gi interface: a private IP address is assigned to the mobile terminals, there is no correlation between the internal IP address and the external address:port assigned by the NAT function, etc.

Privacy-related considerations discussed in [RFC6967] apply for this scenario.

10.3. Scenario 10: Femtocells

This scenario can be seen as a combination of the scenarios described in Section 10.1 and Section 8.

The reference architecture is shown in Figure 8.

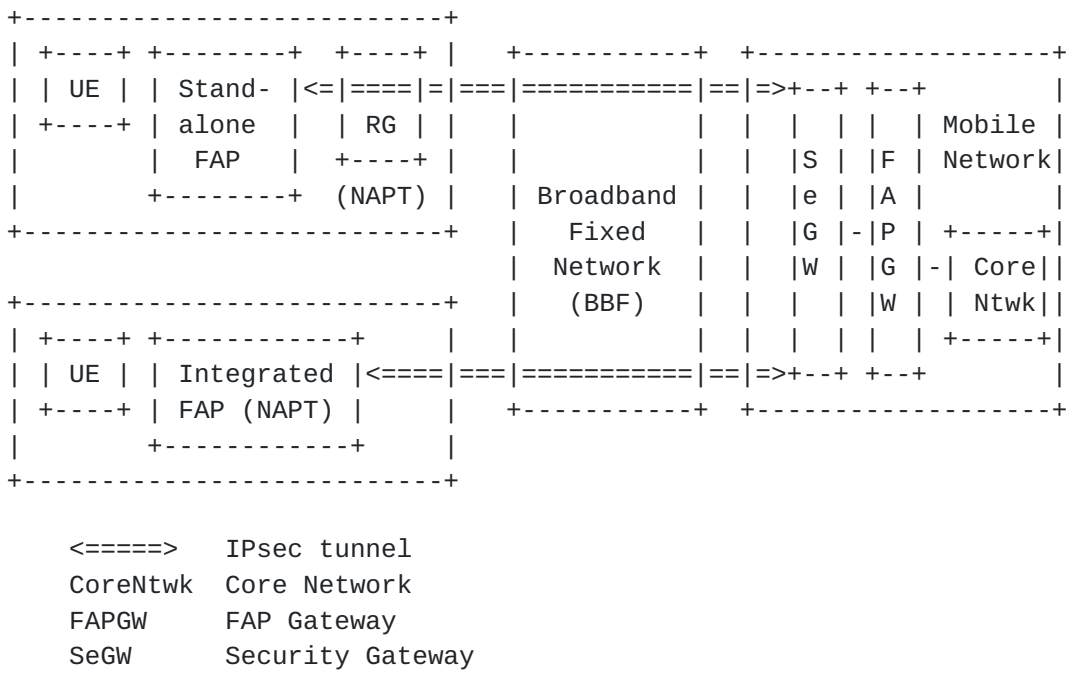

Figure 13: Femtocell Reference Architecture

UE is connected to the FAP at the residential gateway (RG), routed back to 3GPP Evolved Packet Core (EPC). UE is assigned IPv4 address by the Mobile Network. Mobile operator's FAP leverages the IPsec IKEv2 to interconnect FAP with the SeGW over the BBF network. Both the FAP and the SeGW are managed by the mobile operator which may be a different operator for the BBF network.

An investigated scenario is the mobile operator to pass on its mobile subscriber's policies to the BBF to support traffic policy control . But most of today's broadband fixed networks are relying on the private IPv4 addressing plan (+NAPT) to support its attached devices

including the mobile operator's FAP. In this scenario, the mobile network needs to:

- determine the FAP's public IPv4 address to identify the location of the FAP to ensure its legitimacy to operate on the license spectrum for a given mobile operator prior to the FAP be ready to serve its mobile devices.

- determine the FAP's pubic IPv4 address together with the translated port number of the UDP header of the encapsulated IPsec tunnel for identifying the UE's traffic at the fixed broadband network.

- determine the corresponding FAP's public IPv4 address associated with the UE's inner-IPv4 address which is assigned by the mobile network to identify the mobile UE to allow the PCRF to retrieve the special UE's policy (e.g., QoS) to be passed onto the Broadband Policy Control Function (BPCF) at the BBF network.

SeGW would have the complete knowledge of such mapping, but the reasons for unable to use SeGW for this purpose is explained in "Problem Statements" (section 2 of [I-D.so-ipsecme-ikev2-cpext]).

This scenario involves PCRF/BPCF but it is valid in other deployment scenarios making use of AAA servers.

The issue of correlating the internal IP address and the public IP address is valid even if there is no NAT in the path.

This scenario does not introduce privacy concerns since the identification of the host is local to a single administrative domain and is meant to help identifying which policy to select for a UE.

10.4. Scenario 11: Traffic Detection Function

Operators expect that the traffic subject to the packet inspection is routed via the Traffic Detection Function (TDF) function as requirement specified in [TS29.212], otherwise, the traffic may bypass the TDF. This assumption only holds if it is possible to identify individual UEs behind NA(P)T which may be deployed into the RG in fixed broadband network, shown in Figure 14. As a result, additional mechanisms are needed to enable this requirement.

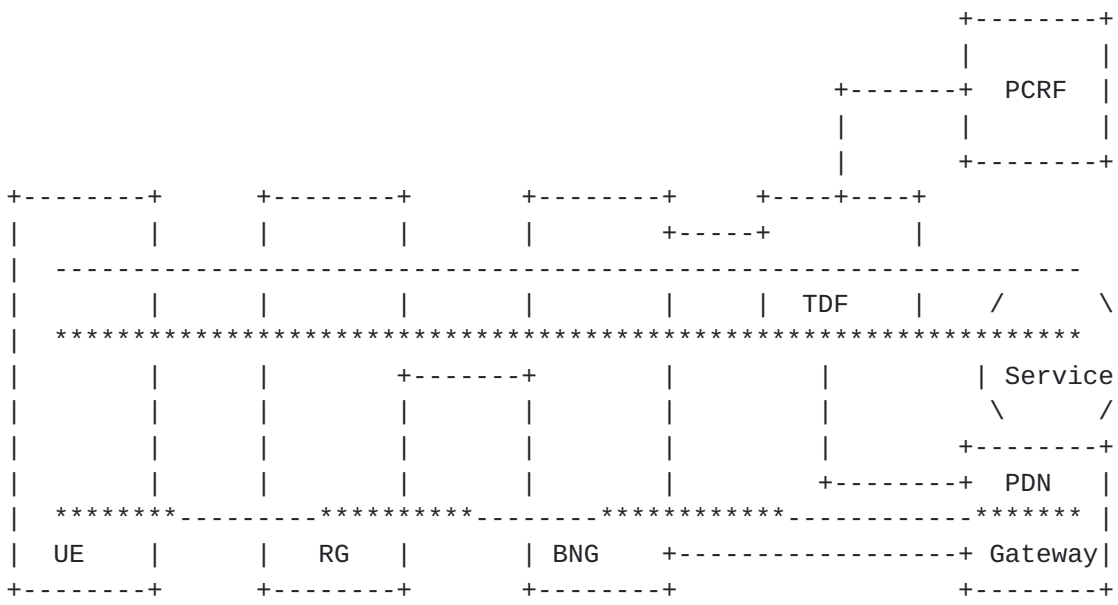

Legends:

---------	3GPP UE User Plane Traffic Offloaded subject to packet inspection
*********	3GPP UE User Plane Traffic Offloaded not subject to packet inspection
*****----	3GPP UE User Plane Traffic Home Routed

Figure 14: UE's Traffic Routed with TDF

This scenario does not introduce privacy concerns since the identification of the host is local to a single administrative domain and is meant to help identifying which policy to select for a UE.

10.5. Scenario 12: Fixed and Mobile Network Convergence

In the Policy for Convergence of Fixed Mobile Convergence (FMC) scenario, the fixed broadband network must partner with the mobile network to acquire the policies for the terminals or hosts attaching to the fixed broadband network, shown in Figure 15 so that host-specific QoS and accounting policies can be applied.

A UE is connected to the RG, routed back to the mobile network. The mobile operator's PCRF needs to maintain the interconnect with the Broadband Policy Control Function (BPCF) in the BBF network for PCC (Section 8). The hosts (i.e., UEs) attaching to fixed broadband network with a NA(P)T deployed should be identified. Based on the UE identification, the BPCF to deploy policy rules in the fixed broadband network can acquire the associated policy rules of the identified UE from the PCRF in the mobile network. But in the fixed

broadband network, private IPv4 address is supported. The similar requirements are raised in this scenario as Section 10.3.

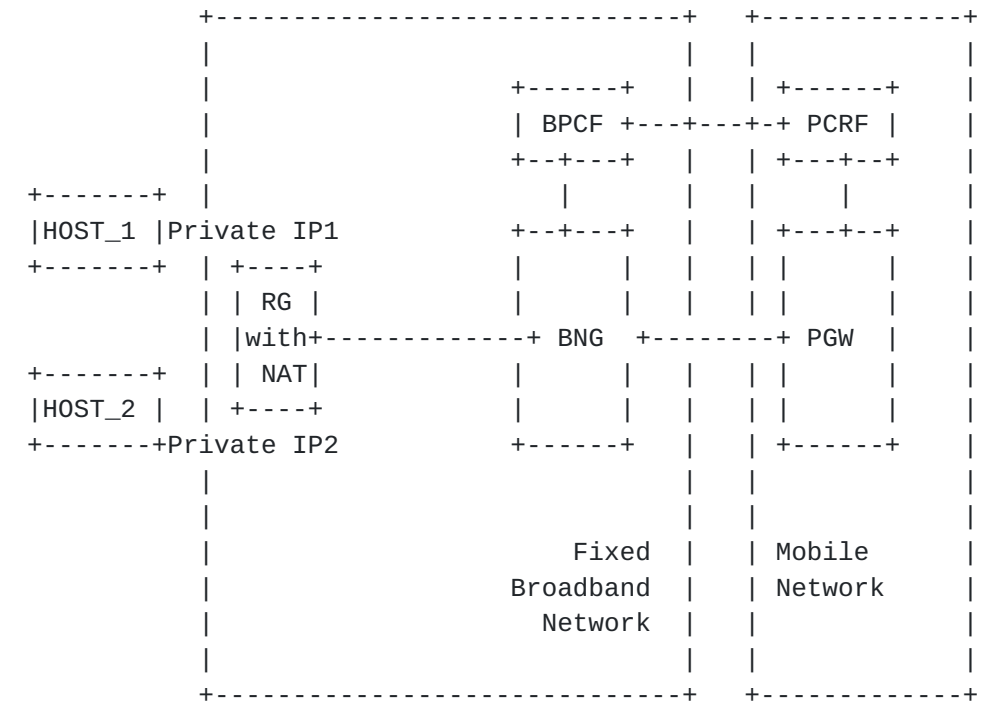

Figure 15: Reference Architecture for Policy for Convergence in Fixed and Mobile Network Convergence (1)

In IPv6 network, the similar issues exists when the IPv6 prefix is sharing between multiple UEs attaching to the RG (see Figure 16). The case applies when RG is assigned a single prefix, the home network prefix, e.g. using DHCPv6 Prefix Delegation [RFC3633] with the edge router, BNG acting as the Delegating Router (DR). RG uses the home network prefix in the address configuration using stateful (DHCPv6) or stateless address assignment (SLAAC) techniques.

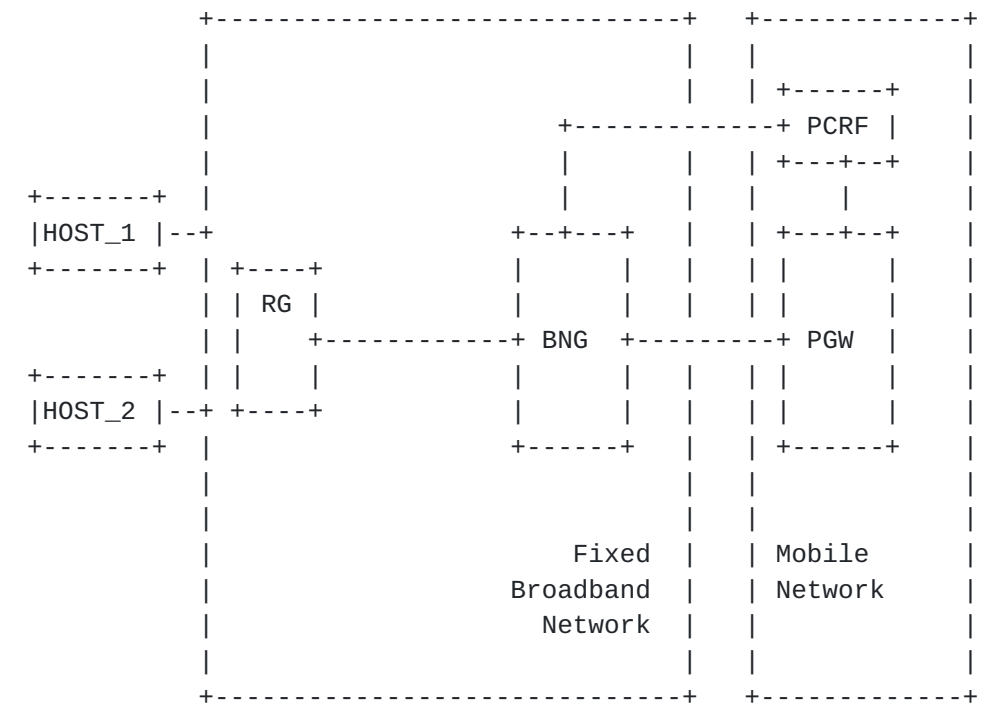

Figure 16: Reference Architecture for Policy for Convergence in Fixed and Mobile Network Convergence (2)

BNG acting as PCEF initiates an IP Connectivity Access Network (IP-CAN) session with the policy server, a.k.a. Policy and Charging Rules Function (PCRF), to receive the Quality of Service (QoS) parameters and Charging rules. BNG provides to the PCRF the IPv6 Prefix assigned to the host, in this case the home network prefix and an ID which in this case has to be equal to the RG specific home network line ID.

HOST_1 in Figure 16 creates an 128-bit IPv6 address using this prefix and adding its interface id. Having completed the address configuration, the host can start communication with a remote hosts over Internet. However, no specific IP-CAN session can be assigned to HOST_1, and consequently the QoS and accounting performed will be based on RG subscription.

Another host, e.g. HOST_2, attaches to RG and also establishes an IPv6 address using the home network prefix. Edge router, the BNG, is not involved with this and all other such address assignments.

This leads to the case where no specific IP-CAN session/sub-session can be assigned to the hosts, HOST_1, HOST_2, etc., and consequently the QoS and accounting performed can only be based on RG subscription and not host specific. Therefore IPv6 prefix sharing in Policy for

Convergence scenario leads to similar issues as the address sharing as it has been explained in the previous scenarios in this document.

11. Synthesis

The following table shows whether each scenario is valid for IPv4/ IPv6 and if it is within one single administrative domain or span multiple domains.

scenario	IPv4	IPv6		Single Administrative Domain
		Client	Server	
CGN	Yes	Yes(1)	No	No
A+P	Yes	No	No	No
Application Proxy	Yes	Yes	Yes	No
Distributed Proxy	Yes	Yes	Yes	Yes/No
Overlay Networks	Yes	Yes(3)	Yes(3)	No
PCC	Yes	Yes(1)	No	Yes
Emergency Calls	Yes	Yes	Yes	No
Provider WLAN	Yes	No	No	Yes
Cellular Networks	Yes	Yes(1)	No	Yes
Femtocells	Yes	No	No	No
TDF	Yes	Yes	No	Yes
FMC	Yes	Yes(1)	No	No

Notes:

(1) e.g., NAT64

(2) A proxy can use IPv6 for the communication leg with the server or the application client.

(3) This scenario is a combination of CGN and Application Proxies.

12. Privacy Considerations

Privacy-related considerations that apply to means to reveal a host identifiers are discussed in [RFC6967]. This document does not introduce additional privacy issues than those discussed in [RFC6967].

13. Security Considerations

This document does not define an architecture nor a protocol; as such it does not raise any security concern. Host identifier related security considerations are discussed in [RFC6967].

14. IANA Considerations

This document does not require any action from IANA.

15. Acknowledgments

Many thanks to F. Klamm, D. Wing, and D. von Hugo for their review.

J. Touch, S. Farrel, and S. Moonesamy provided useful comments in the intarea mailing list.

Figure 8 and part of the text in Section 10.3 are inspired from [I-D.so-ipsecme-ikev2-cpext].

16. Informative References

[EFFOpenWireless]
EFF, , "Open Wireless, https://www.eff.org/issues/open-wireless", 2014.

[I-D.ietf-softwire-lw4over6]
Cui, Y., Qiong, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture", draft-ietf-softwire-lw4over6-10 (work in progress), June 2014.

[I-D.ietf-softwire-map]
Troan, O., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, "Mapping of Address and Port with Encapsulation (MAP)", draft-ietf-softwire-map-10 (work in progress), January 2014.

[I-D.so-ipsecme-ikev2-cpext]
So, T., "IKEv2 Configuration Payload Extension for Private IPv4 Support for Fixed Mobile Convergence", draft-so-ipsecme-ikev2-cpext-02 (work in progress), June 2012.

[I-D.tsou-stateless-nat44]
Tsou, T., Liu, W., Perreault, S., Penno, R., and M. Chen, "Stateless IPv4 Network Address Translation", draft-tsou-stateless-nat44-02 (work in progress), October 2012.

[I-D.williams-overlaypath-ip-tcp-rfc]
Williams, B., "Overlay Path Option for IP and TCP", draft-williams-overlaypath-ip-tcp-rfc-04 (work in progress), June 2013.

[IEEE101109]
Salah, K., Calero, J., Zeadally, S., Almulla, S., and M. ZAaabi, "Using Cloud Computing to Implement a Security Overlay Network, IEEE Security & Privacy, 21 June 2012. IEEE Computer Society Digital Library.", June 2012.

[IEEE1344002]
Byers, J., Considine, J., Mitzenmacher, M., and S. Rost, "Informed content delivery across adaptive overlay networks: IEEE/ACM Transactions on Networking, Vol 12, Issue 5, ppg 767-780", October 2004.

[RFC2753] Yavatkar, R., Pendarakis, D., and R. Guerin, "A Framework for Policy-based Admission Control", RFC 2753, January 2000.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.

[RFC5456] Spencer, M., Capouch, B., Guy, E., Miller, F., and K. Shumard, "IAX: Inter-Asterisk eXchange Version 2", RFC 5456, February 2010.

[RFC5694] Camarillo, G. and IAB, "Peer-to-Peer (P2P) Architecture: Definition, Taxonomies, Examples, and Applicability", RFC 5694, November 2009.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

[RFC6179] Templin, F., "The Internet Routing Overlay Network (IRON)", RFC 6179, March 2011.

[RFC6265] Barth, A., "HTTP State Management Mechanism", RFC 6265, April 2011.

[RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, June 2011.

[RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.

[RFC6346] Bush, R., "The Address plus Port (A+P) Approach to the IPv4 Address Shortage", RFC 6346, August 2011.

[RFC6443] Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, "Framework for Emergency Calling Using Internet Multimedia", RFC 6443, December 2011.

[RFC6888] Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, April 2013.

[RFC6967] Boucadair, M., Touch, J., Levis, P., and R. Penno, "Analysis of Potential Solutions for Revealing a Host Identifier (HOST_ID) in Shared Address Deployments", RFC 6967, June 2013.

[RFC7239] Petersson, A. and M. Nilsson, "Forwarded HTTP Extension", RFC 7239, June 2014.

[TS23.203] 3GPP, , "Policy and charging control architecture", September 2012.

[TS29.212] 3GPP, , "Policy and Charging Control (PCC); Reference Points", December 2013.

Authors' Addresses

Mohamed Boucadair (editor)
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

David Binet
France Telecom
Rennes
France

Email: david.binet@orange.com

Sophie Durel
France Telecom
Rennes
France

Email: sophie.durel@orange.com

Bruno Chatras
France Telecom
Paris
France

Email: bruno.chatras@orange.com

Tirumaleswar Reddy
Cisco Systems
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tireddy@cisco.com

Brandon Williams
Akamai, Inc.
Cambridge MA
USA

Email: brandon.williams@akamai.com

Behcet Sarikaya
Huawei
5340 Legacy Dr. Building 3,
Plano, TX 75024
USA

Email: sarikaya@ieee.org

Li Xue
Huawei
Beijing
China

Email: xueli@huawei.com

Richard Stewart Wheeldon
Cisco Systems
Qube, 90 Whitfield Street
London W1T 4EZ
UK

Email: rwheeldo@cisco.com