

INTERNET-DRAFT  
Intended Status: Informational

M. Cullen  
Painless Security  
N. Leymann  
C. Heidemann  
Deutsche Telekom AG  
M. Boucadair  
France Telecom  
H. Deng  
China Mobile  
M. Zhang  
B. Sarikaya  
Huawei  
October 31, 2016

Expires: May 4, 2017

Problem Statement: Bandwidth Aggregation for Internet Access  
draft-zhang-banana-problem-statement-03.txt

## Abstract

Bandwidth aggregation capabilities for Internet access services can significantly improve end user's Quality of Experience. Such capabilities are especially attractive in environments where multi-interfaced boxes become commonplace and can technically connect to various access networks, both wired and wireless.

This document describes the problems with bandwidth aggregation for Internet access. A set of requirements are derived from the said problems.

## Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at  
<http://www.ietf.org/shadow.html>

INTERNET-DRAFT

Problem Statement

October 31, 2016

## Copyright and License Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1. Introduction</a>	<a href="#">3</a>
<a href="#">2. Acronyms and Terminology</a>	<a href="#">3</a>
<a href="#">3. Generic Reference Model</a>	<a href="#">4</a>
<a href="#">4. Problem Areas</a>	<a href="#">4</a>
<a href="#">4.1. Addressing</a>	<a href="#">4</a>
<a href="#">4.2. Traffic Classification</a>	<a href="#">5</a>
<a href="#">4.3. Traffic Distribution</a>	<a href="#">5</a>
<a href="#">4.4. Traffic Recombination</a>	<a href="#">6</a>
<a href="#">4.4.1. Reordering Buffer</a>	<a href="#">6</a>
<a href="#">4.5. Bypass</a>	<a href="#">7</a>
<a href="#">4.6. Measurement</a>	<a href="#">8</a>
<a href="#">4.7. Policy Control</a>	<a href="#">8</a>
<a href="#">5. Requirements</a>	<a href="#">9</a>
<a href="#">6. Related IETF Work</a>	<a href="#">10</a>
<a href="#">6.1. GRE Tunnel Bonding</a>	<a href="#">10</a>
<a href="#">6.2. LISP</a>	<a href="#">11</a>
<a href="#">6.3. Mobile IP</a>	<a href="#">11</a>
<a href="#">6.4. Multipath TCP Proxy</a>	<a href="#">11</a>
<a href="#">7. Security Considerations</a>	<a href="#">11</a>
<a href="#">8. IANA Considerations</a>	<a href="#">12</a>
<a href="#">9. Acknowledgements</a>	<a href="#">12</a>
<a href="#">10. References</a>	<a href="#">12</a>
<a href="#">10.1. Normative References</a>	<a href="#">12</a>
<a href="#">10.2. Informative References</a>	<a href="#">13</a>
<a href="#">Appendix A. Additional Requirements</a>	<a href="#">14</a>

## [1](#). Introduction

Use cases of BANDwidth Aggregation for interNet Access (BANANA, a.k.a., Hybrid Access) are described in the Technical Report [[TR-348](#)] published by Broadband Forum: by providing Hybrid Access, Service Providers can provide customers with increased access bandwidth and higher access reliability; Service delivery may also be fostered to access the Internet by means of providing a LTE (Long Term Evolution) connection while the wired line is being constructed.

Although host-based Hybrid Access is possible, the scope of this document is restricted to be network-based only. Host-based might be standardized in other places, such as the MIF Working Group.

Design techniques that are being investigated, developed and sometimes deployed to facilitate bandwidth aggregation and improve the resiliency of access conditions raise several problems from various standpoints: traffic routing and forwarding, congestion control, security, etc. This document aims at presenting these problems regardless of the nature of the design technique. It also intends to derive requirements accordingly, and which should be addressed by any bandwidth aggregation technique. Typically, this is one of the inputs that are expected from the IETF community.

A common framework will be sketched, including required functional modules and interactions. The various solution proposals (e.g., GRE, LISP, MIP, MPTCP) can be viewed as applicability assessments of the framework. To support BANANA, the problems to be addressed includes: addressing, traffic classification, distribution and recombination, bypassing, measurement and policy control. To address these problems, we may work as a group to

- specify the encapsulation format;
- develop a common control plane;
- and define or suggest approaches to address BANANA problems developed in this document.

## 2. Acronyms and Terminology

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

CPE: Customer Premises Equipment. An equipment which is the property of the network operator and located on the customer premises.

HG: Home Gateway. A CPE device that is enhanced to support the simultaneous use of both fixed broadband and 3GPP access connections.

BANANA

Expires May 4, 2017

[Page 3]

---

INTERNET-DRAFT

Problem Statement

October 31, 2016

HAAP: Hybrid Access Aggregation Point. A logical function in Operator's network, terminating bonded connections while offering high speed Internet.

PDP: Packet Data Protocol. A packet transfer protocol used in wireless GPRS (General Packet Radio Service)/HSDPA (High Speed Downlink Packet Access) networks.

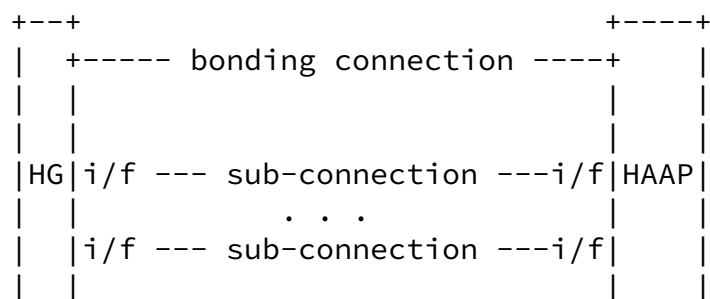
PPPoE: Point-to-Point Protocol over Ethernet is a network protocol for encapsulating PPP frames inside Ethernet frames.

DHCP: Dynamic Host Configuration Protocol [[RFC2131](#)].

DNS: Domain Name System [[RFC1035](#)].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

## 3. Generic Reference Model



+--+

+-----+

Figure 3.1: Reference model of the Hybrid Access

Customers access the Internet using the Hybrid Access which comprises of several key component functions as shown in Figure 3.1: the Home Gateway (HG) as one peer, the Hybrid Access Aggregation Point (HAAP) as the other peer, the bonding connection between the two peers and the sub-connections that logically make up the bonding connection.

## [4. Problem Areas](#)

### [4.1. Addressing](#)

At the HG side, interface addresses of sub-connections are locally acquired upon the bootstrap of the system by means of certain existing protocols such as Point-to-Point Protocol over Ethernet (PPPoE) [[RFC2516](#)] and Packet Data Protocol (PDP). At the HAAP side,

BANANA

Expires May 4, 2017

[Page 4]

---

INTERNET-DRAFT

Problem Statement

October 31, 2016

interface addresses are usually pre-configured by operators. HG and HAAP will rely on the control protocol that is to be developed to exchange these addresses. Afterwards, sub-connections are de-multiplexed by their interface addresses. Both IPv4 and IPv6 should be supported.

End users behind the HG box will regard the bonding connection as a traditional connection to the Internet. With the established sub-connections, connectivity between the HG and HAAP has been built up, therefore endpoint addresses for this bonding connection can be obtained from existing protocols, e.g., DHCP and DNS.

### [4.2. Traffic Classification](#)

Traffic classification occurs before the flows or packets are distributed among the connections. HG and HAAP should support the classification function that marks a flow or packets which are to be further processed by the traffic distribution function or bypass the Hybrid Access (See [Section 4.5](#)). Classification criteria include IP addresses, port numbers, etc. Traffic classification policies can be defined by end users and service providers and must be enforced by the HG and HAAP equipments.

### [4.3.](#) Traffic Distribution

For traffic that is to be distributed across multiple sub-connections, equal load balancing generally applies, possibly inferred by the bandwidth that is available in each link that supports sub-connection. Unequal load balancing should be supported as well. Traffic may be distributed across sub-connections as a function of their available bandwidth. Traffic may also be split in such a way that whenever one sub-connection is saturated, then traffic is forwarded over a secondary sub-connection.

There are two kinds of traffic distribution methods for the Hybrid Access: per-flow load balancing and per-packet load sharing. The per-flow load balancing method is used to be widely adopted in core IP networks. It is suitable for the scenario where there are a large number of flows to be distributed. For end users, usually there are few number of applications to be transmitted over the bonded sub-connections. Per-flow load balance techniques might not be able to achieve a fine grained load distribution, so that the per-packet load sharing is necessary.

For the per-flow load balancing, the load can be distributed using hashing methods. For the per-packet load splitting, the coloring mechanism specified in [[RFC2698](#)] can be used to classify customer's IP packets, both upstream and downstream, and packets will then be

forwarded over the appropriate sub-connections. For example, packets colored as green are forwarded over one sub-connection and packets colored as yellow are forwarded over another sub-connection. For scenarios that rely upon more than two sub-connections, multiple levels of coloring mechanism could be implemented.

### [4.4.](#) Traffic Recombination

For the packet-based traffic distribution, the recombination function at the receiver sides must reorder packets to preserve the integrity of the communication. The sender needs to mark each packet with a sequence number. The sequence number are set per the whole bonding connection rather than per sub-connection so that all packets forwarded over several sub-connections actually share the same reordering buffer.

#### 4.4.1. Reordering Buffer

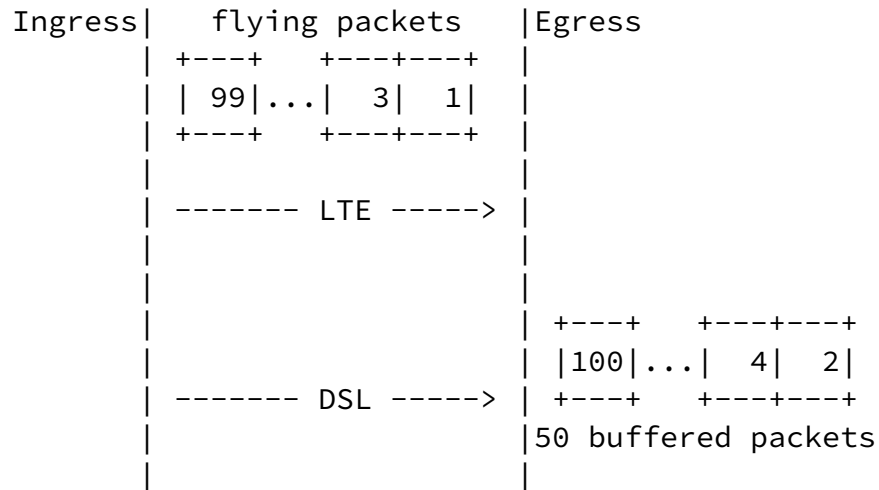


Figure 4.1: Minimizing the reordering buffer

Deployment experiences show that a secondary sub-connection might suffer from large latency, jitter and high packet loss rate. For packet-based traffic distribution, packets are distributed onto those sub-connections at the ingress and then recombined again in a buffer at the egress. If the secondary sub-connection suffers, the entire bonding connection will suffer as well due to the recombination function. For example, assume packet 1,3,...,99 are distributed onto the secondary sub-connection while packet 2,4,...,100 are distributed onto the primary sub-connection. If packet 1 is delayed by 100 ms, even packet 2 arrives at the recombination buffer at the first millisecond, it has to remain in the buffer awaiting for packet 1 as long as 99 ms. Packets distributed onto the primary sub-connection,

which arrive after packet 2, have to be buffered. This can easily lead to the overflow of the reordering buffer and the user's TCP throughput of the bonding connection might be greatly reduced.

Latency of each sub-connection will be monitored. For example, HG and HAAP may calculate the Adaptive Acknowledgment Time-Out of each sub-connection as specified in [RFC2637]; HG and HAAP may periodically exchange control messages to detect the RTT of each sub-connection

[[FLARE](#)]; Packet loss rate of each sub-connection may be monitored as well [[BondL3](#)]. If the difference of the monitored latency exceeds a predefined threshold or the secondary sub-connection exhibits a too high packet loss rate, attached HG and HAAP will stop distributing traffic onto this sub-connection.

Even if the latency of the two sub-connections are comparable and the packet loss rate of the secondary sub-connection is fine so that the reordering buffer does not overflow, it's still worthy to design solutions to minimize the usage of the reordering buffer. In order to realize this goal, the traffic distribution at the ingress should be manipulated. For example, the idea of [[FLARE](#)] might be borrowed: basically, a traffic flow would be split into "flowlets" by the gaps between the arriving packets. Packets of a specific flowlet is solely distributed onto one sub-connection. In this way, reordering is avoided or minimized. The load-balancing method of MPTCP [[RFC6824](#)] could be used as well: packets are always distributed to the sub-connection with the least congestion level and/or latency [[MPScheduler](#)].

#### [4.5](#). Bypass

Service Providers may require some traffic to bypass the Hybrid Access. For example, some delay sensitive applications such as live TV broadcasting carried over a lossy sub-connection would impair customers' Quality of Experience. Service providers could then make sure that such traffic is forwarded over a set of wired sub-connections only, thereby disregarding low-rate mobile connections, for example.

There are two types of bypass: the bypassing traffic are transmitted on a sub-connection out of all the sub-connections between HG and HAAP; the bypassing traffic is still transmitted on a sub-connection between HG and HAAP, just that the occupied bandwidth of the bypassing traffic on this sub-connection can not be used for bandwidth aggregation. In either case, the bypassing traffic would not be under control of the Hybrid Access scheme.

HG and HAAP needs to exchange information about bypassing through the control protocol, such as the application types that need to bypass

the Hybrid Access and the bandwidth occupied by the bypassing traffic



(See also [Section 4.6](#)).

#### [4.6](#). Measurement

HG and HAAP need to measure and exchange performance parameters of the Hybrid Access, including performance parameters that pertain to each sub-connection that belongs to the same connection. Such parameters include (but are not necessarily limited to):

- Operating state: The operating state has to be measured by control messages. When a sub-connection fails, this sub-connection has to be removed from the bonding connection.
- Round Trip Time (RTT): The measurement of this parameter is used by flow and congestion control algorithms for per-flow and per-packet distribution purposes. The probing packet could be piggy-backed by data packets or could be carried by control messages.
- Maximum sub-connection bandwidth: This parameter may be used to determine the amount of traffic that can be distributed across all or a subset of sub-connections.
- Bypassing bandwidth: This parameter has to be periodically monitored. Usually, this is measured for the opposite end to figure out the available sending bandwidth. For example, the HG reports the downloading bypassing bandwidth used in a sub-connection so that the HAAP can calculate the remaining downloading bandwidth of this sub-connection.

#### [4.7](#). Policy Control

Operators and customers may control the Hybrid Access with policies. These policies will be instantiated into parameters or actions that impact traffic classification, distribution, combination, measurement and bypassing. Such policies may consist in:

- Defining traffic filter lists used by the traffic classification function.
- Per-flow or per-packet forwarding policies.
- Operators may specify maximum value of the difference between two measured one-way transit delays. Operators may also specify the maximum allowed packet loss rate of a sub-connection.
- Operators may determine the maximum allowed size (See MAX\_PERFLOW\_BUFFER in [\[RFC2890\]](#)) of the buffer for reordering.

Operators may also define the maximum time (See `OUTOFORDER_TIMER` in [[RFC2890](#)]) that a packet can stay in the buffer for reordering. These parameters may pact traffic recombination.

- Operators and customers may specify the service types to bypass the Hybrid Access.
- Operators may specify the frequency for detecting a sub-connection and the detection retry times before a sub-connection can be declared as "failed".

## 5. Requirements

Requirements for the Hybrid Access are described in this section. Also, some additional requirements are listed for discussion in the Appendix.

The solution **MUST** apply for both IPv4 and IPv6 traffic.

The solution **MUST NOT** require any new capability to support Hybrid Access from the host.

In the Hybrid Access, forwarding traffic flows over various interfaces may have a negative impact on customers' experience (e.g., hazardous log outs, broken HTTPS associations, etc.). The solution should be carefully designed, so that

- activating the solution **MUST NOT** impact the stability, availability of the services delivered to the customer compared to customers who access the same service whose traffic is forwarded along a single path.

"Roles" (primary or backup) should be assigned to each supported network interface type (e.g., fixed or mobile access). This role is assigned by the network operator (e.g., fixed access can be set as the "primary"). Note that there may be more than two sub-connections to support primary and backup roles. A default setting can be considered.

- Setting of the role of the sub-connections **SHOULD NOT** be changed by the user.

The solution should provide Service Providers with means to enforce policy control of the Hybrid Access. For example,

- the solution **MUST** allow to rate limit the traffic on a per-

interface basis.

- the solution **MUST** support means to enable/disable the activation of multiple interfaces at the HG.
- the solution **MUST** support means to instruct the HG about the logic for mounting interfaces.
- the solution **MUST** support means to bind a given traffic to a subset of interfaces.

For the sake of policy enforcement or analytics at the network side,

- the solution **MAY** ease correlating flows, conveyed over multiple access networks, and which belong to the same customer.

Some services such as VoIP may be available over all active interfaces, but distinct logins and credentials may be used.

- The HG **SHOULD** be provided with clear instructions about which account to use to place outgoing sessions. For the sake of simplicity, it is **RECOMMENDED** to use the login/credentials that are independent of the underlying access network used to access the service.

## [6.](#) Related IETF Work

Hybrid Access designs can rely upon several solutions. The following subsections recap the work that has been or is being conducted by the IETF in this area. Note that solutions are listed in an alphabetic order. No preference order should be assumed by the reader.

### [6.1.](#) GRE Tunnel Bonding

GRE Tunnel Bonding [[GRE-HA](#)] uses per-packet traffic distribution to achieve a fine-grained load sharing among the sub-connections. Out-of-sequence packets may be received at the egress so that reordering function is provided. IP packets are encapsulated in the GRE header which is in turn encapsulated in an outer IP header and forwarded over the sub-connections. The Sequence Number field of the GRE header is used to number the packets at the sender side, while the receiver

uses of this sequence number to reorder the packets.

A new control plane is defined. Control messages are transported in the same GRE tunnels that are used to transport data packets. The control messages and data packets are distinguished by the GRE Protocol Type 0xB7EA.

GRE tunnel bonding has been deployed by Deutsche Telekom AG and Austria Telekom.

## [6.2.](#) LISP

LISP has the basic capability to support the Hybrid Access [[LISP-HA](#)] [[ILNP](#)]. LISP can be used to enforce traffic engineering based upon static load balancing that is not agnostic to link characteristics.

Packet-based traffic distribution has been considered in [[LISP-HA](#)] as well. The detail specification of the reordering mechanism based on a "Reorder" flag is left for future work.

## [6.3.](#) Mobile IP

Mobile IP [[RFC3775](#)] and Network Mobility (NEMO; [[RFC3963](#)]) used to handle multiple L3 connectivity to the Internet via multiple ISPs for a multi-homed end user [[RFC4908](#)]. By treating Hybrid Access as a special scenario, some existing capabilities of Mobile IP and NEMO could be reused to realize Hybrid Access. Take [[MIP-HA](#)] as an example, rely on the multiple Care-of Addresses (CoAs) capability [[RFC5648](#)] [[RFC6275](#)], the "addressing" problem of BANANA could be settled. Currently, per-flow traffic distribution has already been supported by Mobile IP and NEMO ([[RFC6088](#)], [[RFC6089](#)]) while packet-based traffic distribution is left for future work [[MIP-HA](#)].

## [6.4.](#) Multipath TCP Proxy

MPTCP provides the ability to establish a communication over multiple paths, by means of sub-flow establishment and operation [[RFC6824](#)]. However, the traditional MPTCP is a host-based technology therefore it's out the scope of this document. What is considered as a candidate technology to support the Hybrid Access is the MPTCP proxy mechanism. There are some implementations and deployments.

The MPTCP proxy operates at the transport layer and locates in the operator's network. A transparent MPTCP mode is proposed in [MPTCP-trans]: a MPTCP proxy terminates a user's TCP flow and reinitiates MPTCP sub-flows towards the other MPTCP proxy; The other MPTCP proxy will terminate the MPTCP sub-flows and restore the user's TCP flow; The MPTCP protocol suite provides features to manage the state of sub-flows between the two proxies. [MPTCP-plain] discusses MPTCP proxy (i.e., transparent MPTCP mode) deployment concerns and also specifies an extension to transport UDP datagrams in MPTCP packets. UDP traffic can therefore be forwarded over a MPTCP connection.

## 7. Security Considerations

Hybrid Access might introduce new threats to the network. For example, traffic sent on unsecured sub-connections would be easily intercepted by an attacker who performs man-in-the-middle attack. The

BANANA

Expires May 4, 2017

[Page 11]

---

INTERNET-DRAFT

Problem Statement

October 31, 2016

multi-path communication may be leveraged to perform Denial of Service (DoS) attack or Distributed Denial of Service (DDoS) attack (e.g., based upon flooding traffic) that may jeopardize the aggregation gateway as well as the access equipment and end station operation.

These kind of new security issues should be carefully considered in designing solutions that aim to address the problems of Bandwidth Aggregation for Internet Access.

## 8. IANA Considerations

No IANA action is required in this document.

## 9. Acknowledgements

Authors would like to thank the comments and suggestions from Christian Jacquenet and Li Xue.

## 10. References

### 10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI

10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

- [TR-348] Broadband Forum, "Technical Report on Hybrid Access Broadband Network Architecture", July, 2016, <<https://www.broadband-forum.org/technical/download/TR-348.pdf>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), DOI 10.17487/RFC2131, March 1997, <<http://www.rfc-editor.org/info/rfc2131>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), DOI 10.17487/RFC1035, November 1987, <<http://www.rfc-editor.org/info/rfc1035>>.
- [RFC2516] Mamakos, L., Lidl, K., Evarts, J., Carrel, D., Simone, D., and R. Wheeler, "A Method for Transmitting PPP Over Ethernet (PPPoE)", [RFC 2516](#), DOI 10.17487/RFC2516, February 1999, <<http://www.rfc-editor.org/info/rfc2516>>.
- [RFC2689] Bormann, C., "Providing Integrated Services over Low-bitrate Links", [RFC 2689](#), DOI 10.17487/RFC2689, September

BANANA

Expires May 4, 2017

[Page 12]

---

INTERNET-DRAFT

Problem Statement

October 31, 2016

1999, <<http://www.rfc-editor.org/info/rfc2689>>.

- [RFC2890] Dommety, G., "Key and Sequence Number Extensions to GRE", [RFC 2890](#), DOI 10.17487/RFC2890, September 2000, <<http://www.rfc-editor.org/info/rfc2890>>.

## [10.2](#). Informative References

- [RFC2637] Hamzeh, K., Pall, G., Verthein, W., Taarud, J., Little, W., and G. Zorn, "Point-to-Point Tunneling Protocol (PPTP)", [RFC 2637](#), DOI 10.17487/RFC2637, July 1999, <<http://www.rfc-editor.org/info/rfc2637>>.
- [BondL3] Maciej Bednarek, Guillermo Barrenetxea, Mirja Mirja Kuehlewind and Brian Trammell, "Multipath bonding at Layer 3", Applied Networking Research Workshop, July 16, 2016, Berlin, Germany

- [FLARE] Srikanth Kandula, Dina Katabi, Shantanu Sinha, and Arthur Berger, "Dynamic Load Balancing Without Packet Reordering", ACM SIGCOMM Computer Communication Review, April 2007.
- [MPscheduler] Hyunwoo Nam, Doru Calin and Henning Schulzrinne, "Towards Dynamic MPTCP Path Control Using SDN", IEEE NetSoft Conference and Workshops (NetSoft), June 2016.
- [GRE-HA] N. Leymann, C. Heidemann, M. Zhang, et al, "GRE Tunnel Bonding", [draft-zhang-gre-tunnel-bonding](#), work in progress.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", [RFC 6824](#), DOI 10.17487/RFC6824, January 2013, <<http://www.rfc-editor.org/info/rfc6824>>.
- [MPTCP-tans] B. Peirens, G. Detal, S. Barre and O. Bonaventure, "Link bonding with transparent Multipath TCP", [draft-peirens-mptcp-transparent](#), work in progress.
- [MPTCP-plain] M. Boucadair and C. Jacquenet, "An MPTCP Option for Network-Assisted MPTCP Deployments: Plain Transport Mode", [draft-boucadair-mptcp-plain-mode](#), work in progress.
- [MIP-HA] P. Seite, A. Yegin and S. Gundavelli, "Multihoming support for Residential Gateways", [draft-seite-dmm-rg-multihoming](#), work in progress.
- [RFC5213] Gundavelli, S., Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", [RFC 5213](#), DOI

10.17487/RFC5213, August 2008, <<http://www.rfc-editor.org/info/rfc5213>>.

- [RFC6088] Tsirtsis, G., Giarreta, G., Soliman, H., and N. Montavont, "Traffic Selectors for Flow Bindings", [RFC 6088](#), DOI 10.17487/RFC6088, January 2011, <<http://www.rfc-editor.org/info/rfc6088>>.
- [RFC6089] Tsirtsis, G., Soliman, H., Montavont, N., Giaretta, G., and K. Kuladinithi, "Flow Bindings in Mobile IPv6 and Network Mobility (NEMO) Basic Support", [RFC 6089](#), DOI

10.17487/RFC6089, January 2011, <<http://www.rfc-editor.org/info/rfc6089>>.

[802.1AX] IEEE, "IEEE Standard for Local and Metropolitan Area Networks - Link Aggregation", 802.1AX-2014, 24 December 2014.

[LISP-HA] M. Menth, A. Stockmayer and M. Schmidt, "LISP Hybrid Access", [draft-menth-lisp-ha](#), work in progress.

[ILNP] "ILNP - Identifier-Locator Network Protocol", online available: <http://ilnp.cs.st-andrews.ac.uk/>

## Appendix A. Additional Requirements

The following requirements are listed as record and may subject to change.

- The solution MUST be valid for any kinds of interfaces that need to be aggregated. No dependency to the underlying media should be assumed.
- The solution MUST comply with servers policy regarding IP addresses that are bound to (HTTP session) cookies.
- The solution MUST NOT break TLS associations.
- Activating the solution MUST NOT have negative impacts on the service usability (including the HG management).
- Service degradation MUST NOT be observed when enabling the solution.
- Enabling the solution MUST increase the serviceability of the HG. In particular, the solution MUST allow for the HG to always establish a network attachment when the primary connectivity is out of service.

- The solution SHOULD NOT alter any mechanism, to aggregate available resources or to ensure a service continuity among multiple access points, that is supported by end-devices connected to the HG.



- The HG MUST bind the DNS server(s) discovered during the network attachment phase to the interface from which the information was received.
- The HG MUST bind the service information (e.g., SIP Proxy Server) discovered during the network attachment phase to the interface from which the information was received.
- When sending the traffic via a given interface, the HG MUST use as source address an address (or an address from a prefix) that was assigned for that interface.
- For protocols such as RTP/RTCP, the same IP address MUST be used for both RTP and RTCP sessions.
- Dedicated tools SHOULD be provided to the customer to assess the aggregated capacity (instead of link-specific). This can be included as part of the HG UI, a dedicated portal, etc.

## Author's Addresses

Margaret Cullen  
Painless Security  
14 Summer St. Suite 202  
Malden, MA 02148 USA

EMail: [margaret@painless-security.com](mailto:margaret@painless-security.com)

Nicolai Leymann  
Deutsche Telekom AG  
Winterfeldtstrasse 21-27  
Berlin 10781  
Germany

Phone: +49-170-2275345  
EMail: [n.leymann@telekom.de](mailto:n.leymann@telekom.de)

Cornelius Heidemann  
Deutsche Telekom AG  
Heinrich-Hertz-Strasse 3-7  
Darmstadt 64295  
Germany

Phone: +4961515812721  
EMail: [heidemannc@telekom.de](mailto:heidemannc@telekom.de)

Mohamed Boucadair  
France Telecom  
Rennes 35000  
France

EMail: [mohamed.boucadair@orange.com](mailto:mohamed.boucadair@orange.com)

Hui Deng  
China Mobile  
53A,Xibianmennei Ave.,  
Xuanwu District,  
Beijing 100053  
China

EMail: [denghui@chinamobile.com](mailto:denghui@chinamobile.com)

---

INTERNET-DRAFT

Problem Statement

October 31, 2016

Mingui Zhang  
Huawei Technologies  
No.156 Beiqing Rd. Haidian District,  
Beijing 100095 P.R. China

EMail: zhangmingui@huawei.com

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

EMail: sarikaya@ieee.org

BANANA

Expires May 4, 2017

[Page 17]