

BANANA
Internet Draft
Category: Proposed Standard

N. Leymann
C. Heidemann
Deutsche Telekom AG
L. Chen
M. Zhang
B. Sarikaya
Huawei
M. Cullen
Painless Security
December 21, 2017

Expires: June 24, 2018

BANdwidth Aggregation for interNet Access (BANANA)
Attributes for the Control Protocol of Bonding Tunnels
draft-leymann-banana-signaling-attributes-01.txt

Abstract

This memo specifies the attributes for the control protocol of BANdwidth Aggregation for interNet Access (BANANA).

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>

Copyright and License Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

INTERNET-DRAFT

BANANA Attributes

December 21, 2017

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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

1. Introduction	3
2. Acronyms and Terminology	3
3. GRE Tunnel Setup Request	4
3.1. Client Identification Name	4
3.2. Session ID	4
3.3. Synchronization Rate	5
4. GRE Tunnel Setup Accept	6
4.1. R IPv4 Address	6
4.2. R IPv6 Address	6
4.3. Session ID	7
4.4. RTT Difference Threshold	7
4.5. Bypass Bandwidth Check Interval	8
4.6. Active Hello Interval	8
4.7. Hello Retry Times	9
4.8. Idle Timeout	10
4.9. Bonding Key Value	10
4.10. Configured Link Upstream Bandwidth	11
4.11. Configured Link Downstream Bandwidth	11
4.12. RTT Difference Threshold Violation	12
4.13. RTT Difference Threshold Compliance	13
4.14. Idle Hello Interval	13
4.15. No Traffic Monitored Interval	14
5. GRE Tunnel Setup Deny	14
5.1. Error Code	15
6. GRE Tunnel Hello	16
6.1. Timestamp	16
6.2. IPv6 Prefix Assigned by the Remote	17
7. GRE Tunnel Tear Down	17
8. GRE Tunnel Notify	17
8.1. Bypass Traffic Rate	18

8.2.	Filter List Package	18
8.3.	Switching To Tunnel One	22
8.4.	Overflowing To Tunnel Two	22
8.5.	Link One Failure	23
8.6.	Link Two Failure	23

8.7.	IPv6 Prefix Assigned to Host	23
8.8.	Diagnostic Start: Bonding Tunnel	24
8.9.	Diagnostic Start: Tunnel One	24
8.10.	Diagnostic Start: Tunnel Two	25
8.11.	Diagnostic End	25
8.12.	Filter List Package ACK	26
8.13.	Switching To Active Hello State	26
8.14.	Switching To Idle Hello State	27
8.15.	Tunnel Verification	28
9.	Security Considerations	29
10.	IANA Considerations	29
11.	References	29
11.1.	Normative References	29
11.2.	Informative References	30
	Contributors	30
	Authors' Addresses	30

[1.](#) Introduction

The control plane of BANANA delivers configuration and control information between two peering BANANA boxes. Control messages such as GRE Tunnel Setup Request, GRE Tunnel Setup Accept, GRE Tunnel Setup Deny, GRE Tunnel Hello, GRE Tunnel Tear Down and GRE Tunnel Notify are specified in [[BANANA-signaling](#)]. This document further specifies the attributes to be carried as Attribute field in those control messages exchanged between the two peering BANANA boxes.

[2.](#) Acronyms and Terminology

GRE: Generic Routing Encapsulation [[RFC2784](#)] [[RFC2890](#)].

CIR: Committed Information Rate [[RFC2697](#)].

RTT: Round-Trip Time.

FQDN: Fully Qualified Domain Name. Generally, a host name with at

least one domain label under the top-level domain. For example, "dhcp.example.org" is an FQDN [[RFC7031](#)].

DSCP: The 6-bit codepoint (DSCP) of the Differentiated Services field (DS field) in the IPv4 and IPv6 headers [[RFC2724](#)].

DNS: Domain Name System. A hierarchical distributed naming system for computers, services, or any resource connected to the Internet or a private network.

DHCP: Dynamic Host Configuration Protocol. A standardized network protocol used on Internet Protocol (IP) networks for dynamically

distributing network configuration parameters, such as IP addresses for interfaces and services.

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.](#) GRE Tunnel Setup Request

The local BANANA box uses the GRE Tunnel Setup Request message to request that the remote BANANA box establishes the GRE tunnels. It is sent out from the local BANANA box's "F" and "S" interfaces (see Figure 3.1 of [[BANANA-signaling](#)]). Attributes that need to be included in this message are defined in the following subsections.

[3.1.](#) Client Identification Name

An operator uses the Client Identification Name (CIN) to identify the local BANANA box. The local BANANA box sends the CIN to the remote BANANA box for authentication and authorization as specified in [[TS23.401](#)]. It is REQUIRED that the GRE Tunnel Setup Request message sent out from the "S" interface contain the CIN attribute while the GRE Tunnel Setup Request message sent out from the "F" interface does not contain this attribute.

The CIN attribute has the following format:

+---+---+---+---+---+---+

|Attribute Type |

(1 byte)

```

+---+---+---+---+---+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+...-+
| Client Identification Name   (40 bytes) |
+---+---+---+---+---+---+---+---+---+---+...+

```

Attribute Type
CIN, set to 3.

Attribute Length
Set to 40.

Client Identification Name
This is a 40-byte string value encoded in UTF-8 and set by the operator. It is used as the identification of the local BANANA box in the operator's network.

[3.2.](#) Session ID

This Session ID is generated by the remote BANANA box when the first GRE Tunnel Setup Request message is received. The remote BANANA box announces the Session ID to the local BANANA box in the GRE Tunnel Setup Accept message. For the "F" and "S" interfaces that need to be bonded together, the local BANANA box MUST use the same Session ID. The local BANANA box MUST carry the Session ID attribute in each GRE Tunnel Setup Request message except the first time that the GRE Tunnel Setup Request message is sent to the remote BANANA box.

The Session ID attribute has the following format:

```

+---+---+---+---+---+
|Attribute Type | (1 byte)
+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+...-+
| Session ID                 (4 bytes) |
+---+---+---+---+---+---+---+---+---+---+...+

```

Attribute Type
Session ID, set to 4.

Attribute Length
Set to 4.

Session ID

An unsigned integer generated by the remote BANANA box. It is used as the identification of bonded GRE tunnels.

[3.3. Synchronization Rate](#)

If a GRE tunnel is using the Digital Subscriber Line (DSL), the local BANANA box uses the Synchronization Rate to notify the remote BANANA box about the downstream bandwidth of the DSL link. The GRE Tunnel Setup Request message sent on a GRE tunnel on a DSL link MUST include the Synchronization Rate attribute. GRE Tunnel Setup Request message sent on a GRE tunnel on other kinds of link SHOULD NOT include this attribute.

```
+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+...+
| Synchronization Rate       (4 bytes) |
+---+---+---+---+---+---+---+---+---+---+---+---+...+---
```

Attribute Type

Synchronization Rate, set to 7.

Attribute Length
Set to 4.

Synchronization Rate

An unsigned integer measured in kbps.

[4. GRE Tunnel Setup Accept](#)

The remote BANANA box uses the GRE Tunnel Setup Accept message as the response to the GRE Tunnel Setup Request message. This message indicates acceptance of the tunnel establishment and carries parameters of the GRE tunnels. Attributes that need to be included in this message are defined below.

[4.1.](#) R IPv4 Address

The remote BANANA box uses the "R" (see Figure 3.1 of [BANANA-signaling]) IPv4 Address attribute to inform the local BANANA box of the "R" IPv4 address. The local BANANA box uses the "R" IPv4 address as the destination endpoint IPv4 address of the GRE tunnels (the source endpoint IPv4 addresses of the GRE tunnels are the "F" interface IP address and the "S" interface IP address). The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the R IPv4 Address attribute.

```
+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+---+
| Attribute Length |            (2 bytes)
+---+---+---+---+---+---+---+---+---+...--+
| R IPv4 Address |            (4 bytes) |
+---+---+---+---+---+---+---+---+---+...++
```

Attribute Type
R IPv4 Address, set to 1.

Attribute Length
Set to 4.

R IPv4 Address
Set to the pre-configured IPv4 address (e.g., an IP address of a Line Card in the remote BANANA box), which is used as the endpoint IP address of GRE tunnels by the remote BANANA box.

[4.2.](#) R IPv6 Address

The remote BANANA box uses the "R" (see Figure 3.1 of [BANANA-signaling]) IPv6 Address attribute to inform the local BANANA box of the "R" IPv6 address. The local BANANA box uses the "R" IPv6 address as the destination endpoint IPv6 address of the GRE tunnels (the source endpoint IPv6 addresses of the GRE tunnels are the "F" interface IP address and the "S" interface IP address). The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the R IPv6 Address attribute.

```

+---+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length |       (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+...--+
| R IPv6 Address |       (16 bytes) |
+---+---+---+---+---+---+---+---+---+---+---+---+...--+

```

Attribute Type
R IPv6 Address, set to 2.

Attribute Length
Set to 16.

R IPv6 Address
Set to the pre-configured IPv6 address (e.g., an IP address of a Line Card in the remote BANANA box), which is used as the endpoint IP address of GRE tunnels by the remote BANANA box.

[4.3.](#) Session ID

The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the Session ID attribute as defined in [Section 3.2](#).

[4.4.](#) RTT Difference Threshold

The remote BANANA box uses the RTT Difference Threshold attribute to inform the local BANANA box of the acceptable threshold of the RTT difference between Tunnel 1 and Tunnel 2. If the measured RTT difference exceeds this threshold, the local BANANA box SHOULD stop offloading traffic to Tunnel 2. The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the RTT Difference Threshold attribute.

```

+---+---+---+---+---+

```



```

|Attribute Type | (1 byte)
+---+---+---+---+---+---+---+---+---+---+
| Attribute Length | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+...--+
| RTT Difference Threshold (4 bytes) |
+---+---+---+---+---+---+---+---+---+---+...--+

```

Attribute Type
RTT Difference Threshold, set to 9.

Attribute Length
Set to 4.

RTT Difference Threshold
An unsigned integer measured in milliseconds. This value can be chosen in the range 0 through 1000.

[4.5. Bypass Bandwidth Check Interval](#)

The remote BANANA box uses the Bypass Bandwidth Check Interval attribute to inform the local BANANA box of how frequently the bypass bandwidth should be checked. The local BANANA box should check the bypass bandwidth of the "F" interface in each time period indicated by this interval. The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the Bypass Bandwidth Check Interval attribute.

```

+---+---+---+---+---+---+---+---+---+---+
|Attribute Type | (1 byte)
+---+---+---+---+---+---+---+---+---+---+
| Attribute Length | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+...--+
| Bypass Bandwidth Check Interval (4 bytes) |
+---+---+---+---+---+---+---+---+---+---+...--+

```

Attribute Type
Bypass Bandwidth Check Interval, set to 10.

Attribute Length
Set to 4.

Bypass Bandwidth Check Interval
An unsigned integer measured in seconds. This value can be chosen in the range 10 through 300.

[4.6. Active Hello Interval](#)

The remote BANANA box uses the Active Hello Interval attribute to

inform the local BANANA box of the pre-configured interval for sending out GRE Tunnel Hellos. The local BANANA box should send out GRE Tunnel Hellos via both "F" and "S" interfaces in each time period as indicated by this interval. The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the Active Hello Interval attribute.

```

+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length           |      (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+...--+
| Active Hello Interval      |      (4 bytes)      |
+---+---+---+---+---+---+---+---+---+---+---+---+...++

```

Attribute Type
Active Hello Interval, set to 14.

Attribute Length
Set to 4.

Active Hello Interval
An unsigned integer measured in seconds. This value can be chosen in the range 1 through 100.

[4.7.](#) Hello Retry Times

The remote BANANA box uses the Hello Retry Times attribute to inform the local BANANA box of the retry times for sending GRE Tunnel Hellos. If the local BANANA box does not receive any acknowledgement from the remote BANANA box for the number of GRE Tunnel Hello attempts specified in this attribute, the local BANANA box will declare a failure of the GRE tunnel. The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the Hello Retry Times attribute.

```

+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length           |      (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+...--+
| Hello Retry Times          |      (4 bytes)      |
+---+---+---+---+---+---+---+---+---+---+---+---+...++

```

Attribute Type
Hello Retry Times, set to 15.

Attribute Length
Set to 4.

Hello Retry Times

An unsigned integer that takes values in the range 3 through 10.

[4.8.](#) Idle Timeout

The remote BANANA box uses the Idle Timeout attribute to inform the local BANANA box of the pre-configured timeout value to terminate Tunnel 1. When Tunnel 2 failure is detected, all traffic will be sent over Tunnel 1. If the failure of Tunnel 2 lasts longer than the Idle Timeout, subsequent traffic will be sent over the raw link 1 rather than over Tunnel 1, and Tunnel 1 SHOULD be terminated. The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the Idle Timeout attribute.

```
+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+---+
| Attribute Length |            (2 bytes)
+---+---+---+---+---+---+---+---+---+...--+
| Idle Timeout      |            (4 bytes)  |
+---+---+---+---+---+---+---+---+---+...+---
```

Attribute Type

Idle Timeout, set to 16.

Attribute Length

Set to 4.

Idle Timeout

An unsigned integer measured in seconds. It takes values in the range 0 through 172,800 with a granularity of 60. The default value is 86,400 (24 hours). The value 0 indicates that the idle timer never expires.

[4.9.](#) Bonding Key Value

The remote BANANA box uses the Bonding Key Value attribute to inform the local BANANA box of the number that is to be carried as the Key of the GRE header for subsequent control messages. The Bonding Key

Value is generated by the remote BANANA box and used for security purposes.

The method used to generate this number is left up to implementations. The pseudorandom number generator defined in ANSI X9.31, [Appendix A.2.4 \[ANSI-X9.31-1998\]](#) is RECOMMENDED. Note that random number generation "collisions" are allowed in the GRE Tunnel Bonding Protocol.

```

+-----+
|Attribute Type |                (1 byte)
+-----+
| Attribute Length          |    (2 bytes)
+-----+-----+...+
| Bonding Key Value          (4 bytes) |
+-----+-----+...+

```

Attribute Type
Bonding Key Value, set to 20.

Attribute Length
Set to 4.

Bonding Key Value
A 32-bit random number generated by the remote BANANA box.

[4.10](#). Configured Link Upstream Bandwidth

The remote BANANA box obtains the upstream bandwidth of a link from the management system and uses the Configured Link Upstream Bandwidth attribute to inform the local BANANA box. The local BANANA box uses the received upstream bandwidth as the CIR [\[RFC2697\]](#) for the link. GRE Tunnel Setup Accept messages sent on either Tunnel 1 or Tunnel 2 MAY include the Configured Link Upstream Bandwidth attribute. If a GRE tunnel is using the DSL, the GRE Tunnel Setup Accept message sent on this link MUST include the Configured Link Upstream Bandwidth attribute.

```

+-----+
|Attribute Type |                (1 byte)
+-----+

```


An unsigned integer measured in kbps.

[4.12.](#) RTT Difference Threshold Violation

The remote BANANA box uses the RTT Difference Threshold Violation attribute to inform the local BANANA box of the number of times in a row that the RTT Difference Threshold (see [Section 4.4](#)) may be violated before the local BANANA box MUST stop using Tunnel 2. If the RTT Difference Threshold is continuously violated for more than the indicated number of measurements, the local BANANA box MUST stop using Tunnel 2. The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the RTT Difference Threshold Violation attribute.

```
+---+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+...--+
| RTT Diff Threshold Violation (4 bytes) |
+---+---+---+---+---+---+---+---+---+---+---+---+...++
```

Attribute Type
RTT Difference Threshold Violation, set to 24.

Attribute Length
Set to 4.

RTT Difference Threshold Violation
An unsigned integer that takes values in the range 1 through 25.
A typical value is 3.

[4.13.](#) RTT Difference Threshold Compliance

The remote BANANA box uses the RTT Difference Threshold Compliance attribute to inform the local BANANA box of the number of times in a row that the RTT Difference Threshold (see [Section 4.4](#)) must be compliant before use of Tunnel 2 can be resumed. If the RTT Difference Threshold is continuously detected to be compliant across more than this number of measurements, the local BANANA box MAY resume using Tunnel 2. The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the RTT Difference Threshold Compliance attribute.

```

+---+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+...--+
| RTT Diff Threshold Compliance (4 bytes) |
+---+---+---+---+---+---+---+---+---+---+---+...--+

```

Attribute Type
RTT Difference Threshold Compliance, set to 25.

Attribute Length
Set to 4.

RTT Difference Threshold Compliance
An unsigned integer that takes values in the range 1 through 25.
A typical value is 3.

[4.14.](#) Idle Hello Interval

The remote BANANA box uses the Idle Hello Interval attribute to inform the local BANANA box of the pre-configured interval for sending out GRE Tunnel Hellos when the subscriber is detected to be idle. The local BANANA box SHOULD begin to send out GRE Tunnel Hellos via both "F" and "S" interfaces in each time period as indicated by this interval, if the bonded tunnels have seen no traffic for a period longer than the "No Traffic Monitored Interval" (see [Section 4.15](#)). The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the Idle Hello Interval attribute.

```

+---+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+...--+
| Idle Hello Interval         (4 bytes) |
+---+---+---+---+---+---+---+---+---+---+---+...--+

```

Attribute Type
Idle Hello Interval, set to 31.

Attribute Length
Set to 4.

Idle Hello Interval

An unsigned integer measured in seconds. This value can be chosen in the range 100 through 86,400 (24 hours) with a granularity of 100. The default value is 1800 (30 minutes).

[4.15.](#) No Traffic Monitored Interval

The remote BANANA box uses the No Traffic Monitored Interval attribute to inform the local BANANA box of the pre-configured interval for switching the GRE Tunnel Hello mode. If traffic is detected on the bonded GRE tunnels before this interval expires, the local BANANA box SHOULD switch to the Active Hello Interval. The GRE Tunnel Setup Accept message sent over Tunnel 2 MUST include the No Traffic Monitored Interval attribute.

```
+---+---+---+---+
|Attribute Type |                               (1 byte)
+---+---+---+---+---+---+---+---+---+---+
| Attribute Length |                               (2 bytes)
+---+---+---+---+---+---+---+---+---+---+...--+
| No Traffic Monitored Interval (4 bytes) |
+---+---+---+---+---+---+---+---+---+---+...+---
```

Attribute Type

No Traffic Monitored Interval, set to 32.

Attribute Length
Set to 4.

No Traffic Monitored Interval

An unsigned integer measured in seconds. This value is in the range 30 through 86,400 (24 hours). The default value is 60.

[5.](#) GRE Tunnel Setup Deny

The remote BANANA box MUST send the GRE Tunnel Setup Deny message to the local BANANA box if the GRE Tunnel Setup Request from this local BANANA box is denied. The local BANANA box MUST terminate the GRE tunnel setup process as soon as it receives the GRE Tunnel Setup Deny

message.

5.1. Error Code

The remote BANANA box uses the Error Code attribute to inform the local BANANA box of the error code. The error code depicts why the GRE Tunnel Setup Request is denied. GRE Tunnel Setup Deny messages sent over both Tunnel 1 and Tunnel 2 MUST include the Error Code attribute.

```
+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length                |    (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+...-+
| Error Code                      (4 bytes)  |
+---+---+---+---+---+---+---+---+---+---+---+---+...+-+
```

Attribute Type
Error Code, set to 17.

Attribute Length
Set to 4.

Error Code
An unsigned integer. The list of codes is as follows:

- 1: The remote BANANA box was not reachable via link 2 during the GRE Tunnel Setup Request.
- 2: The remote BANANA box was not reachable via link 1 during the GRE Tunnel Setup Request.
- 3: The Tunnel 2 to the remote BANANA box failed.
- 4: The Tunnel 1 to the remote BANANA box failed.
- 5: The given User ID of link 1 is not allowed to use the GRE Tunnel Bonding service.
- 6: The given User Alias / User ID (Globally Unique Identifier (GUID)) is not allowed to use the GRE Tunnel Bonding service.
- 7: The User IDs of link 1 and link 2 do not match.

- 8: The remote BANANA box denied the GRE Tunnel Setup Request because a bonding session with the same User ID already exists.
- 9: The remote BANANA box denied the GRE Tunnel Setup Request because the user's CIN is not permitted.
- 10: The remote BANANA box terminated a GRE Tunnel Bonding session for maintenance reasons.
- 11: There was a communication error between the remote BANANA box and the management system when the GRE Tunnel Setup Request message is being sent over Tunnel 2.
- 12: There was a communication error between the remote BANANA box and the management system when the GRE Tunnel Setup Request message is being sent over Tunnel 1.

6. GRE Tunnel Hello

After the Tunnel 1 / Tunnel 2 is established, the local BANANA box begins to periodically send out GRE Tunnel Hello messages via the tunnel; the remote BANANA box acknowledges the local BANANA box's messages by returning GRE Tunnel Hello messages to the local BANANA box. This continues until the tunnel is terminated.

6.1. Timestamp

The remote BANANA box uses the Timestamp attribute to inform the local BANANA box of the timestamp value that is used for RTT calculation. GRE Tunnel Hello messages sent over both Tunnel 1 and Tunnel 2 MUST include the Timestamp attribute.

```

+---+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length |                (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+...--+
| Timestamp                (8 bytes) |
+---+---+---+---+---+---+---+---+---+---+---+---+...--+

```

Attribute Type
Timestamp, set to 5.

Attribute Length
Set to 8.

The time since the system restarted. The high-order 4 bytes indicate an unsigned integer in units of 1 second; the low-order 4 bytes indicate an unsigned integer in units of 1 millisecond.

[6.2.](#) IPv6 Prefix Assigned by the Remote

The remote BANANA box uses the IPv6 Prefix Assigned by the Remote attribute to inform the local BANANA box of the assigned IPv6 prefix. This IPv6 prefix is to be captured via lawful intercept. GRE Tunnel Hello messages sent over both Tunnel 1 and Tunnel 2 MUST include the IPv6 Prefix Assigned by the Remote attribute.

```
+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length |                (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+...--+
| IPv6 Prefix Assigned by the Remote(16 bytes) |
+---+---+---+---+---+---+---+---+---+---+---+...+---
```

Attribute Type

IPv6 Prefix Assigned by the Remote, set to 13.

Attribute Length

Set to 17.

IPv6 Prefix Assigned by the Remote

The highest-order 16 bytes encode an IPv6 address. The lowest-order 1 byte encodes the prefix length. These two values are put together to represent an IPv6 prefix.

[7.](#) GRE Tunnel Tear Down

The remote BANANA box can terminate Tunnel 1 / Tunnel 2 by sending the GRE Tunnel Tear Down message to the local BANANA box via the tunnel. The Error Code attribute as defined in [Section 5.1](#) MUST be included in this message. After receiving the GRE Tunnel Tear Down message, the local BANANA box removes the IP address of R, which is the destination IP addresses of Tunnel 1 and Tunnel 2.

[8. GRE Tunnel Notify](#)

The local BANANA box and the remote BANANA box use the GRE Tunnel Notify message, which is transmitted through either Tunnel 1 or Tunnel 2, to notify each other about their status regarding the two GRE tunnels, the information for the bonded tunnels, the actions that need to be taken, etc.

Usually, the receiver just sends the received attributes back as the acknowledgement for each GRE Tunnel Notify message. However, there is an exception for the Filter List Package: since the size of the Filter List Package attribute can be very large, a special attribute -- the Filter List Package ACK attribute -- is used as the acknowledgement (see [Section 8.12](#)).

Attributes that need to be included in the GRE Tunnel Notify message are defined below.

[8.1. Bypass Traffic Rate](#)

There are a few types of traffic that need to be transmitted over the raw "F" interface (see Figure 3.1 of [[BANANA-signaling](#)]) rather than the bonded GRE tunnels. The local BANANA box has to set aside bypass bandwidth on the "F" interface for these traffic types. Therefore, the available bandwidth of Tunnel 1 is the entire "F" interface bandwidth minus the occupied bypass bandwidth.

The local BANANA box uses the Bypass Traffic Rate attribute to inform the remote BANANA box of the downstream bypass bandwidth for the "F" interface. The Bypass Traffic Rate attribute will be included in the GRE Tunnel Notify message sent over Tunnel 1. The remote BANANA box calculates the available downstream bandwidth for Tunnel 1 as the Configured Link Downstream Bandwidth minus the bypass bandwidth provided by the local BANANA box. The available downstream bandwidth will be used as the CIR of the coloring system [[RFC2697](#)].

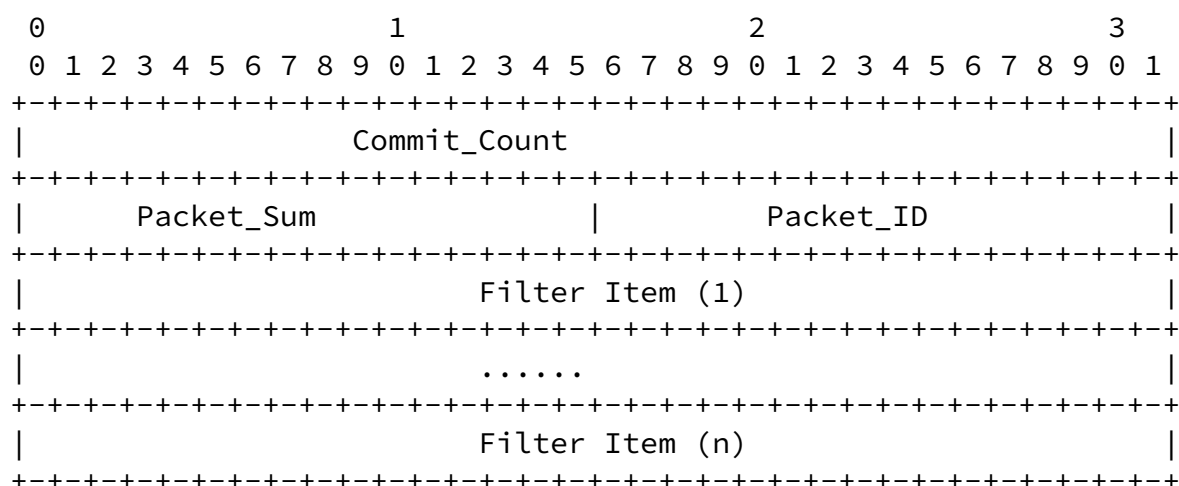
```
+---+---+---+---+
|Attribute Type |                               (1 byte)
+---+---+---+---+---+---+---+---+---+---+
| Attribute Length           |               (2 bytes)
+---+---+---+---+---+---+---+---+---+---+...
```


Attribute Length


The total length of the Filter List TLV. The maximum allowed length is 969 bytes.

Filter List TLV

The Filter List TLV occurs one time in a Filter List Package attribute. It has the following format:



where each Filter Item is of the following format:



Commit_Count

An unsigned integer that identifies the version of the Filter Item list. The version is shared by all Filter List Packages and increases monotonically by one for each new Filter Item list. The local BANANA box MUST refresh its Filter Item list when a new Commit_Count is received.

Packet_Sum

If a single Filter List Package attribute might make the control message larger than the MTU, fragmentation is used. The Packet_Sum indicates the total number of fragments.

Packet_ID

The fragmentation index for this Filter List Package attribute. Each fragment is numbered starting at 1 and increasing by one up to Packet_Sum.

Type

The Type of the Filter Item. Currently, the following types are supported:

Filter Item	Type
=====	=====
FQDN [RFC7031]	1
DSCP [RFC2724]	2
Destination Port	3
Destination IP	4
Destination IP & Port	5
Source Port	6
Source IP	7
Source IP & Port	8
Source MAC	9
Protocol	10

Source IP Range	11
Destination IP Range	12
Source IP Range & Port	13
Destination IP Range & Port	14

Other values are reserved for future use and MUST be ignored on receipt.

Length

The length of the Filter Item in bytes. Type and Length are excluded.

Enable

An integer that indicates whether or not the Filter Item is enabled. A value of 1 means "enabled", and a value of 0 means "disabled". Other possible values are reserved and MUST be ignored on receipt.

Description Length

The length of the Description Value in bytes.

Description Value

A variable-length string value encoded in UTF-8 that describes the Filter List TLV (e.g., "FQDN").

Value

A variable-length string encoded in UTF-8 that specifies the value of the Filter Item (e.g., "www.yahoo.com"). As an example, Type = 1 and Value = "www.yahoo.com" mean that packets whose FQDN field equals "www.yahoo.com" match the Filter Item. "Source MAC" (source Media Access Control address) values are specified using hexadecimal numbers. Port numbers are decimals as assigned by IANA in [[Port-NO](#)]. For the "Protocol" type, the value could be either a decimal or a keyword specified by IANA in [[Pro-NO](#)]. The formats for IP addresses and IP address ranges are defined in [[RFC4632](#)] and [[RFC4291](#)] for IPv4 and

IPv6, respectively. A Filter Item of Type 5, 8, 13, or 14 is a combination of two parameters; values for the two parameters are separated by a colon (":").

[8.3](#). Switching To Tunnel One

If the RTT difference between Tunnel 1 and Tunnel 2 is continuously detected to be in violation of the RTT Difference Threshold (see [Section 4.4](#)) more than the number of times specified in the RTT Difference Threshold Violation attribute (see [Section 4.12](#)), the local BANANA box uses the Switching To Tunnel One attribute to inform the remote BANANA box to use Tunnel 1 only. When the remote BANANA box receives this attribute, it MUST begin to transmit downstream traffic to this local BANANA box solely over Tunnel 1. The GRE Tunnel Notify message sent over Tunnel 1 MAY include the Switching To Tunnel One attribute.

```
+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+
| Attribute Length |            (2 bytes)
+---+---+---+---+---+---+---+---+
```

Attribute Type
Switching To Tunnel One, set to 11.

Attribute Length
Set to 0.

[8.4](#). Overflowing To Tunnel Two

If the RTT difference between Tunnel 1 and Tunnel 2 is continuously detected to not be in violation of the RTT Difference Threshold (see [Section 4.4](#)) more than the number of times specified in the RTT Difference Threshold Compliance attribute (see [Section 4.13](#)), the local BANANA box uses the Overflowing To Tunnel Two attribute to inform the remote BANANA box that Tunnel 2 can be used again. The GRE Tunnel Notify message sent over Tunnel 1 MAY include the Overflowing To Tunnel Two attribute.

```
+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+
| Attribute Length |            (2 bytes)
+---+---+---+---+---+---+---+---+
```

Attribute Type
Overflowing To Tunnel Two, set to 12.

Attribute Length
Set to 0.

[8.5.](#) Link One Failure

When the local BANANA box detects that the "F" interface status is "down", it MUST tear down Tunnel 1. It informs the remote BANANA box about the failure by using the Link One Failure attribute. The remote BANANA box MUST tear down Tunnel 1 upon receipt of the Link One Failure attribute. The Link One Failure attribute SHOULD be carried in the GRE Tunnel Notify message sent over Tunnel 2.

```
+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+
```

Attribute Type
Link One Failure, set to 18.

Attribute Length
Set to 0.

[8.6.](#) Link Two Failure

When the local BANANA box detects that the "S" interface status is "down", it MUST tear down Tunnel 2. It informs the remote BANANA box about the failure by using the Link Two Failure attribute. The remote BANANA box MUST tear down Tunnel 2 upon receipt of the Link Two Failure attribute. The Link Two Failure attribute SHOULD be carried in the GRE Tunnel Notify message sent over Tunnel 1.

```
+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+
```

Attribute Type
Link Two Failure, set to 19.

Attribute Length
Set to 0.

[8.7.](#) IPv6 Prefix Assigned to Host

If the local BANANA box changes the IPv6 prefix assigned to the host,

INTERNET-DRAFT

BANANA Attributes

December 21, 2017

it uses the IPv6 Prefix Assigned to Host attribute to inform the remote BANANA box. GRE Tunnel Notify messages sent over both Tunnel 1 and Tunnel 2 MAY include the IPv6 Prefix Assigned to Host attribute.

```

+---+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length                |      (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+...--+
| IPv6 Prefix Assigned to Host      (16 bytes) |
+---+---+---+---+---+---+---+---+---+---+---+...+---+

```

Attribute Type

IPv6 Prefix Assigned to Host, set to 21.

Attribute Length

Set to 17.

IPv6 Prefix Assigned to Host

The highest-order 16 bytes encode an IPv6 address. The lowest-order 1 byte encodes the prefix length. These two values are put together to represent an IPv6 prefix.

[8.8.](#) Diagnostic Start: Bonding Tunnel

The local BANANA box uses the Diagnostic Start: Bonding Tunnel attribute to inform the remote BANANA box to switch to diagnostic mode to test the performance of the entire bonding tunnel. The Diagnostic Start: Bonding Tunnel attribute SHOULD be carried in the GRE Tunnel Notify message sent over Tunnel 1.

```

+---+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length                |      (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+

```

Attribute Type

Diagnostic Start: Bonding Tunnel, set to 26.

Attribute Length

Set to 0.

[8.9.](#) Diagnostic Start: Tunnel One

The local BANANA box uses the Diagnostic Start: Tunnel One attribute to inform the remote BANANA box to switch to diagnostic mode to test

the performance of Tunnel 1. The Diagnostic Start: Tunnel One attribute SHOULD be carried in the GRE Tunnel Notify message sent over Tunnel 1.

```
+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+
| Attribute Length           |   (2 bytes)
+---+---+---+---+---+---+---+---+---+
```

Attribute Type
Diagnostic Start: Tunnel One, set to 27.

Attribute Length
Set to 0.

[8.10.](#) Diagnostic Start: Tunnel Two

The local BANANA box uses the Diagnostic Start: Tunnel Two attribute to inform the remote BANANA box to switch to diagnostic mode to test the performance of Tunnel 2. The Diagnostic Start: Tunnel Two attribute SHOULD be carried in the GRE Tunnel Notify message sent over Tunnel 1.

```
+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+
| Attribute Length           |   (2 bytes)
+---+---+---+---+---+---+---+---+---+
```

Attribute Type
Diagnostic Start: Tunnel Two, set to 28.

Attribute Length
Set to 0.

[8.11.](#) Diagnostic End

The local BANANA box uses the Diagnostic End attribute to inform the remote BANANA box to stop operating in diagnostic mode. The Diagnostic End attribute SHOULD be carried in the GRE Tunnel Notify message sent over Tunnel 1.

```
+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+
```

Leymann, et al.

Expires June 24, 2018

[Page 25]

INTERNET-DRAFT

BANANA Attributes

December 21, 2017

Attribute Type
Diagnostic End, set to 29.

Attribute Length
Set to 0.

[8.12.](#) Filter List Package ACK

The local BANANA box uses the Filter List Package ACK attribute to acknowledge the Filter List Package sent by the remote BANANA box. GRE Tunnel Notify messages sent over both Tunnel 1 and Tunnel 2 MAY include the Filter List Package ACK attribute.

```
+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+
| Attribute Length           | (2 bytes)
+---+---+---+---+---+---+---+---+---+...-+
| Filter List Package ACK      (5 bytes) |
+---+---+---+---+---+---+---+---+---+...+---
```

Attribute Type
Filter List Package ACK, set to 30.

Attribute Length
Set to 5.

Filter List Package ACK

The highest-order 4 bytes are the Commit_Count as defined in [Section 8.2](#). The lowest-order 1 byte encodes the following error codes:

0: The Filter List Package is acknowledged.

1: The Filter List Package is not acknowledged. The local BANANA box is a new subscriber and has not ever received a Filter List Package. In this case, the remote BANANA box SHOULD tear down the bonding tunnels and force the local BANANA box to re-establish the GRE tunnels.

2: The Filter List Package is not acknowledged. The local BANANA box has already gotten a valid Filter List Package. The filter list on the local BANANA box will continue to be used, while the remote BANANA box need not do anything.

[8.13](#). Switching To Active Hello State

If traffic is being sent/received over the bonding GRE tunnels before

the "No Traffic Monitored Interval" expires (see [Section 4.15](#)), the local BANANA box sends the remote BANANA box a GRE Tunnel Notify message containing the Switching To Active Hello State attribute.

The remote BANANA box will switch to Active Hello State and send the local BANANA box a GRE Tunnel Notify message carrying the Switching To Active Hello State attribute as the ACK.

When the local BANANA box receives the ACK, it will switch to Active Hello State, start RTT detection, and start sending GRE Tunnel Hello messages with the Active Hello Interval (see [Section 4.6](#)).

```
+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+
| Attribute Length |       (2 bytes)
+---+---+---+---+---+---+---+---+---+
```

Attribute Type

Switching To Active Hello State, set to 33.

Attribute Length
Set to 0.

[8.14.](#) Switching To Idle Hello State

The local BANANA box initiates switching To Idle Hello State when the bonding of GRE tunnels is successfully established and the GRE Tunnel Setup Accept message sent over Tunnel 2 carrying the Idle Hello Interval attribute (see [Section 4.14](#)) is received. The local BANANA box sends the remote BANANA box a GRE Tunnel Notify message containing the Switching To Idle Hello State attribute.

The remote BANANA box will switch to Idle Hello State, clear RTT state, and send the local BANANA box a GRE Tunnel Notify message carrying the Switching to Idle Hello State attribute as the ACK.

When the local BANANA box receives the ACK, it will (1) switch to Idle Hello State, (2) stop RTT detection and clear RTT state, and (3) start sending GRE Tunnel Hello messages with the Idle Hello Interval (see [Section 4.14](#)).

```
+---+---+---+---+
|Attribute Type |           (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Attribute Length           |           (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+
```

Attribute Type
Switching To Idle Hello State, set to 34.

Attribute Length
Set to 0.

[8.15.](#) Tunnel Verification

The remote BANANA box uses the Tunnel Verification attribute to inform the local BANANA box to verify whether an existing Tunnel 2 is still functioning. The Tunnel Verification attribute SHOULD be carried in the GRE Tunnel Notify message sent over Tunnel 2. It provides a means to detect the tunnel faster than the GRE Tunnel Hello, especially when Tunnel 2 is in the Idle Hello State and it

takes a much longer time to detect this tunnel.

When the remote BANANA box receives a GRE Tunnel Setup Request sent over Tunnel 2 and finds that the requested tunnel conflicts with an existing tunnel, the remote BANANA box initiates tunnel verification.

The remote BANANA box drops all conflicting GRE Tunnel Setup Request messages sent over Tunnel 2 and sends GRE Tunnel Notify messages carrying the Tunnel Verification attribute until the verification ends. The local BANANA box MUST respond to the remote BANANA box with the same Tunnel Verification attribute as the ACK if the tunnel is still functioning.

If the ACK of the Tunnel Verification attribute is received from the local BANANA box, the remote BANANA box determines that the existing tunnel is still functioning. A Tunnel 2 Deny message (with Error Code = 8) will be sent to the local BANANA box. The local BANANA box SHOULD terminate the GRE Tunnel Setup Request process immediately.

If the remote BANANA box does not receive a tunnel verification ACK message after three attempts (one initial attempt and two retries), it will regard the existing tunnel as failed, and the GRE Tunnel Setup Request sent over Tunnel 2 will be accepted.

```
+---+---+---+---+
|Attribute Type |                (1 byte)
+---+---+---+---+---+---+---+---+
| Attribute Length |                (2 bytes)
+---+---+---+---+---+---+---+---+
```

Attribute Type
Tunnel Verification, set to 35.

Attribute Length
Set to 0.

[9](#). Security Considerations

See the Security Considerations of [[BANANA-signaling](#)].

[10](#). IANA Considerations

IANA need not assign anything for this memo. RFC editor: please

remove this section before publication.

11. References

11.1. Normative References

- [Port-NO] IANA, "Service Name and Transport Protocol Port Number Registry", <<http://www.iana.org/assignments/service-names-port-numbers>>.
- [Pro-NO] IANA, "Assigned Internet Protocol Numbers", <<http://www.iana.org/assignments/protocol-numbers>>.
- [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>>.
- [RFC2697] Heinanen, J. and R. Guerin, "A Single Rate Three Color Marker", [RFC 2697](#), DOI 10.17487/RFC2697, September 1999, <<http://www.rfc-editor.org/info/rfc2697>>.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", [RFC 2784](#), DOI 10.17487/RFC2784, March 2000, <<http://www.rfc-editor.org/info/rfc2784>>.
- [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>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", [BCP 122](#), [RFC 4632](#), DOI 10.17487/RFC4632, August 2006, <<http://www.rfc-editor.org/info/rfc4632>>.

- [TR-069] Broadband Forum, "CPE WAN Management Protocol", Issue: 1 Amendment 5, November 2013, <https://www.broadband-forum.org/technical/download/TR-069_Amendment-5.pdf>.
- [TS23.401] 3GPP TS23.401, "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access", v11.7.0, September 2013.
- [BANANA-signaling]
N. Leymann, C. Heidemann, et al, "BANDwidth Aggregation for interNet Access (BANANA) The Control Protocol of Bonding Tunnels", [draft-leymann-banana-signaling](#), work in progress.

11.2. Informative References

- [802Type] IANA, "IEEE 802 Numbers",
<<http://www.iana.org/assignments/ieee-802-numbers>>.
- [ANSI-X9.31-1998]
ANSI Standard X9.31-1998, "Digital Signatures Using Reversible Public Key Cryptography for the Financial Services Industry (rDSA)", 1998.
- [RFC2724] Handelman, S., Stibler, S., Brownlee, N., and G. Ruth, "RTFM: New Attributes for Traffic Flow Measurement", [RFC 2724](#), DOI 10.17487/RFC2724, October 1999, <<http://www.rfc-editor.org/info/rfc2724>>.
- [RFC7031] Mrugalski, T. and K. Kinnear, "DHCPv6 Failover Requirements", [RFC 7031](#), DOI 10.17487/RFC7031, September 2013, <<http://www.rfc-editor.org/info/rfc7031>>.

Contributors

Li Xue
Individual
Email: xueli_jas@163.com

Zhongwen Jiang
Huawei Technologies
Email: jiangzhongwen@huawei.com

Authors' Addresses

INTERNET-DRAFT

BANANA Attributes

December 21, 2017

Nicolai Leymann
Deutsche Telekom AG
Winterfeldtstrasse 21-27
Berlin 10781
Germany
Phone: +49-170-2275345
Email: n.leymann@telekom.de

Cornelius Heidemann
Deutsche Telekom AG
Heinrich-Hertz-Strasse 3-7
Darmstadt 64295
Germany
Phone: +49-6151-5812721
Email: heidemannc@telekom.de

Lihao Chen
Huawei Technologies
No.156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China
EMail: lihao.chen@huawei.com

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
United States of America
Email: sarikaya@ieee.org

Margaret Cullen

Painless Security
14 Summer St. Suite 202
Malden, MA 02148
United States of America
Email: margaret@painless-security.com

Leymann, et al.

Expires June 24, 2018

[Page 31]

INTERNET-DRAFT

BANANA Attributes

December 21, 2017

Leymann, et al.

Expires June 24, 2018

[Page 32]