ROLL Interim Meeting

2015-02-10
slide set version 5

Note Well

Any submission to the IETF intended by the Contributor for publication as all or part of an IETF Internet-Draft or RFC and any statement made within the context of an IETF activity is considered an "IETF Contribution". Such statements include oral statements in IETF sessions, as well as written and electronic communications made at any time or place, which are addressed to:

- The IETF plenary session
- The IESG, or any member thereof on behalf of the IESG
- Any IETF mailing list, including the IETF list itself, any working group or design team list, or any other list functioning under IETF auspices
- Any IETF working group or portion thereof
- Any Birds of a Feather (BOF) session
- The IAB or any member thereof on behalf of the IAB
- The RFC Editor or the Internet-Drafts function

All IETF Contributions are subject to the rules of RFC 5378 and RFC 3979 (updated by RFC 4879).

Statements made outside of an IETF session, mailing list or other function, that are clearly not intended to be input to an IETF activity, group or function, are not IETF Contributions in the context of this notice. Please consult RFC 5378 and RFC 3979 for details.

A participant in any IETF activity is deemed to accept all IETF rules of process, as documented in Best Current Practices RFCs and IESG Statements.

A participant in any IETF activity acknowledges that written, audio and video records of meetings may be made and may be available to the public.

Source: https://www.ietf.org/about/note-well.html

Material

Agenda: http://www.ietf.org/proceedings/interim/2015/02/10/roll/agenda/agenda-interim-2015-roll-1

Presenter: Ralph Droms

Additional Inf: http://www.ietf.org/proceedings/interim/2015/02/10/roll/proceedings.html

Agenda for meeting

1. Agenda bashing
2. Goals and desired outcomes for this meeting
3. Cover origin of problem: how did we get here
   a. 6tisch requirements for minimal
   b. 6man review of flow label
   c. 6lo review of NHC
   d. publication of dispatch header solution
4. Architectural deep dive of proposals on the table

Goal:

- Arrive at consensus on the scope of the compression of 6553/6554 problem.
- How much effort to spend and where shall the problem be solved.
- Explore the three alternatives and the constraints on the solution space.

Outcome

- Which WG should solve this problem,
- Scale and timing of solution.

[will the solution be 6lowpan-specific? shall it employ lossless compression?]

RPL Artifacts compression

by Pascal Thubert

draft-ietf-6tisch-minimal

- Last Call completed
- 6TiSCH supports classical PHY
  - Uncompressed RPL artifacts are show stoppers
  - Minimal needs a compression mechanism
  - Initially flow label, currently NHC
- Open issue on NHC reference
  - What of we do not reference anything
  - Else what should default be
  - Status on appeal to ADs
  - See Brian’s answer next slide

Talk with Brian Haberman

> All in all, if you agree, we can progress the draft in IESG review leaving that particular reference open to be changed till later in the process?

Why would you expect that to be acceptable? If a piece of the spec is needed for interoperability, it needs to be there during the review.

Why send it to the IESG if there are still questions on how that piece will be done? If it is left out, how will anyone know if there is consensus on the spec as a whole?

It seems to me that:

1. There have been several proposals put forth
2. Issues have been raised on each one

So, it seems prudent (to me) to look at each of the alternatives and determine which warts everyone concerned is willing to live with.

Running code is a good way to do that.

Making that determination needs to be done before the spec is sent to the IESG.

Problem statement

- RPL adds 3 artifacts in data packets.
  - These are an IP in IP encapsulation,
  - an option in HbH header (called RPI) and
  - an RH (RH type 3).
- These artifacts are only partially and sub-optimally compressed in 6LoWPAN.
- The RPI is not compressed at all, consumes 8 octets per packet, and will cause IP in IP encapsulation in some packets.
- Any combination of artifacts is possible though RPI usually omitted if RH3 is present.

More on RPL option [RFC6553] (RPI)

- 8-octets overhead for HbH header in data pkt
  - Almost innocuous with G-PHY (ZigbeeIP, CG-Mesh)
  - May cause fragmentation with classical PHY (127 octets/Frame in 2.4 GHz band)
  - Not compressed by 6LoWPAN HC
  - Wasted Energy in constrained devices
- Additional IP-in-IP encapsulation
  - Deeply aggravating factor for energy consumption and fragmentation
  - Needed when the packet flows outside the LLN to be able to add or remove the RPI (6MAN)

When does the RPI force IPIP encapsulation?

- When egressing the LLN -- When ingressing the LLN

The RPI can be placed in a RPL Option for use in an IPv6 Hop-by-Hop header...Since the Hop-by-Hop header must not flow outside of the RPL domain, it must be removed from packets that leave the domain, and be inserted in packets entering the domain. In both cases, this operation implies an IP-in-IP encapsulation [dispatch-03]

For reasons such as security and the capability to send ICMP errors back to the source, an original packet must not be tampered with, and any information that must be inserted in or removed from an IPv6 packet must be placed in an extra IP-in-IP encapsulation[dispatch-03]

In the case of the Non-Storing Mode of Operation (MOP), RPL requires a Routing Header type 3 (RH3) for all packets that are routed down a RPL graph. With Non-Storing RPL, even if the source is a node in the same LLN, the packet must first reach up the graph to the root so that the root can insert the RH3 to go down the graph. In any fashion, whether the packet was originated in a node in the LLN or outside the LLN, and regardless of whether the packet stays within the LLN or not, as long as the source of the packet is not the root itself, the source-routing operation also implies an IP-in-IP encapsulation at the root to insert the RH3 [dispatch-03]

- Can the RH3 be added by an originating LLN node for intra-LLN traffic?

Problem with RH3 [RFC6554]

- Only partial compression
  - Almost innocuous with G-PHY (ZigbeeIP, CG-Mesh)
    - [large MTU, but appropriate MTUs need to be advertised]
  - Elides common prefix in addresses
  - Still wasted Energy in constrained devices
- Additional IP-in-IP encapsulation
  - **Deeply aggravating factor for energy consumption and fragmentation**
  - **Needed for packets incoming the LLN (6MAN)**
- 6TiSCH supports classical PHY
  - Frame (128 bytes) size limits hops

History

- ROLL agreed in Toronto to place the RPI data in the IPv6 flow label, which would reduce the per-packet overhead from 8 bytes into 20 bits and would avoid additional overhead from IP in IP encapsulation. Implementation was demonstrated at the 6TiSCH plug fest.
- Brian Carpenter reviewed the work and approved it, provided that the LLN operations would not leak outside the LLN.
- ROLL submitted the case to 6MAN for validation (that's the flow label draft)
- During the 6MAN review, it was suggested that an alternate approach to compress the RPI data could be found at 6lo, and evaluation work started in that direction (that's the NHC draft). 6MAN chairs confirmed recommendation not to use the flow label in chairs-ads exchanges
- 6TiSCH supported that work and asked for help from 6lo
- Reviewing the NHC draft in Honolulu, 6lo decided that we would need to understand how RH3 would be compressed in that approach to fully evaluate it.
- The authors found that extending the NHC approach to IP in IP and RH3 lead to too much implementation complexity, and that placing the RPL artifacts in a separate dispatch from the original packet lead to a more efficient processing at source, destination and intermediate nodes a better reuse of existing 6LoWPAN compression code. They produced a new draft called the Routing Dispatch draft.

Proposals on the table

3 proposals to compress RPL artifacts in data not all of them necessarily in all packets, which makes the compression operation delicate:

- The most efficient is based on Flow Label (less bits and no IP in IP) but is incomplete since it does not cover the RH => now dismissed

- A 6LoWPAN alternate, the NHC draft, which also compresses the RPI, but requires IP in IP and does not cover RH either.

- Another 6LoWPAN alternate called the Routing Dispatch draft. This one addresses the whole story (IP in IP, RPI, and RH3). This is the response from the NHC authors to the 6lo concerns on the NHC approach.

On the 6MAN side

3 possible outcomes on the IoT/FL question to 6MAN:

- No change to existing specs. In that case, even resetting the Flow Label inside the LLN is not acceptable. There are little chances that IoT devices conform.
- **Only (re)setting the FL to zero inside the LLN is acceptable, in which case the FL cannot be used for RPI. This ask is critical for all LLNs and ensures continued conformance of existing standards such as ISA100.11a.**
- LLNs are free to use the FL as seem fit, as long as this activity does not leak outside the LLN and the setting of outgoing packets conforms 6437 (this is the current question of the 6MAN FL draft that Brian Carpenter reviewed)

How does ISA100 use Flow Label

1. sets the flow-label at the end systems, it is unchanged throughout the LLN
2. the flow label is unchanged when the packet egresses
3. the flow label is usually set properly by corresponding nodes that sending into the LLN, but the gateway may, according to ACL, reset the Flow Label to a valid value.

(Bob hints that even the ISA100.11a approach is not acceptable)

Current base solution

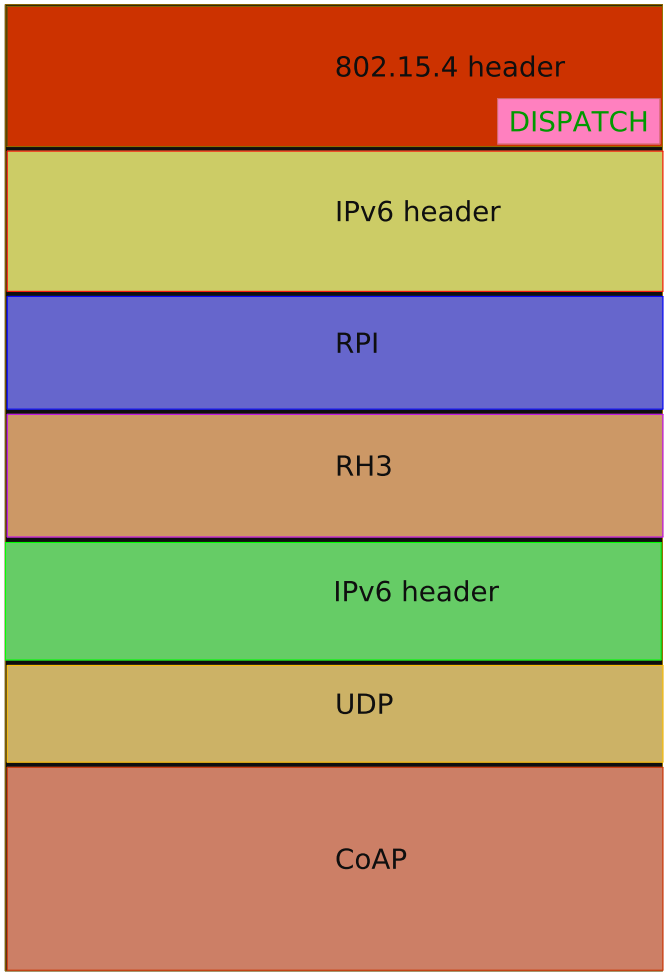

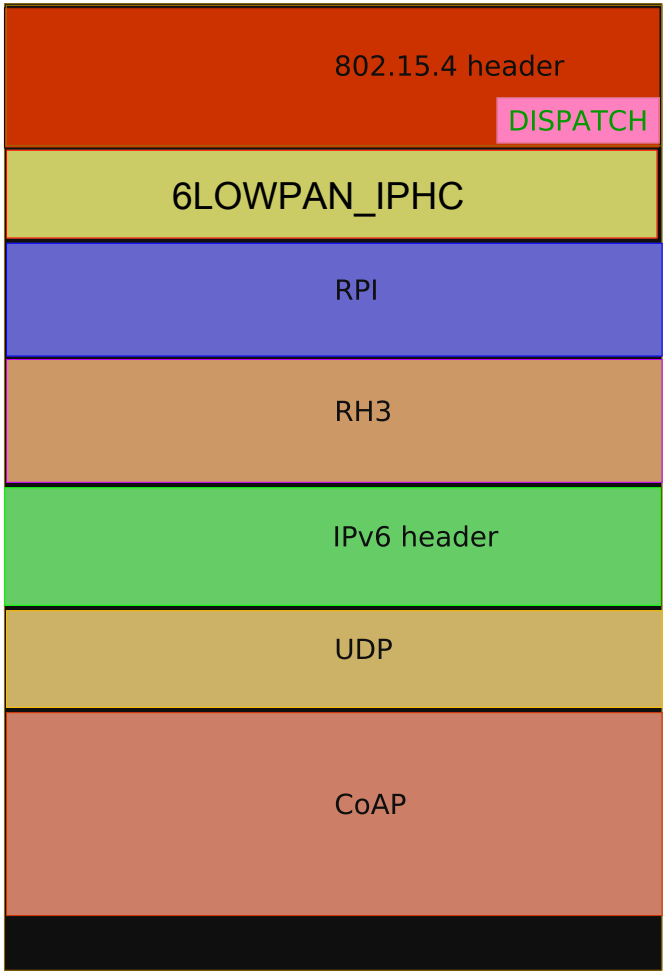

Note: IGNORES FRAGMENTATION HEADERS

Flow Label Proposal

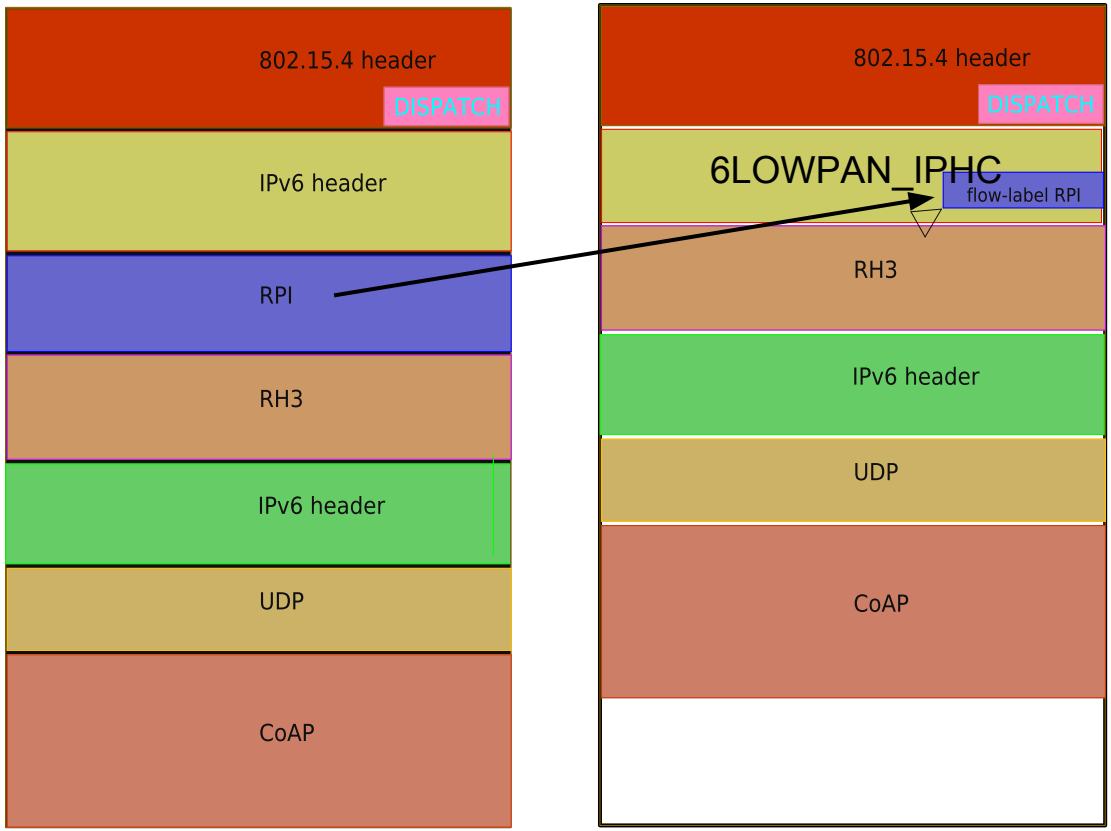

NHC solution: architectural view

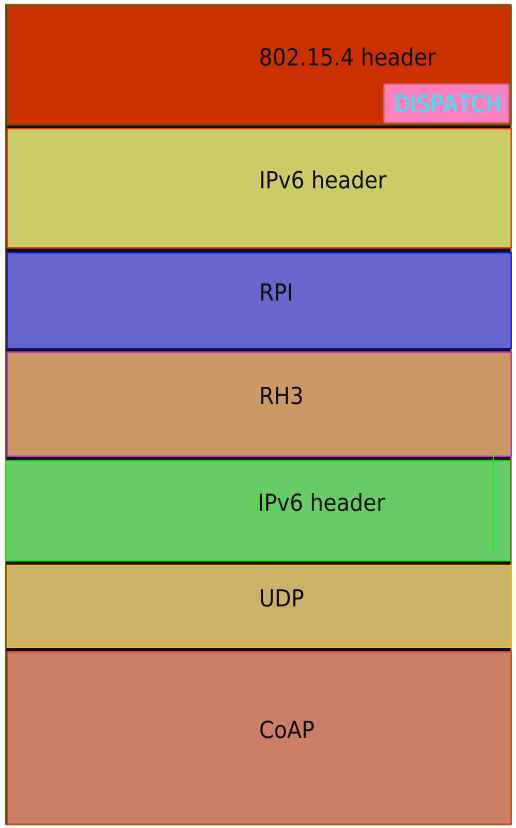

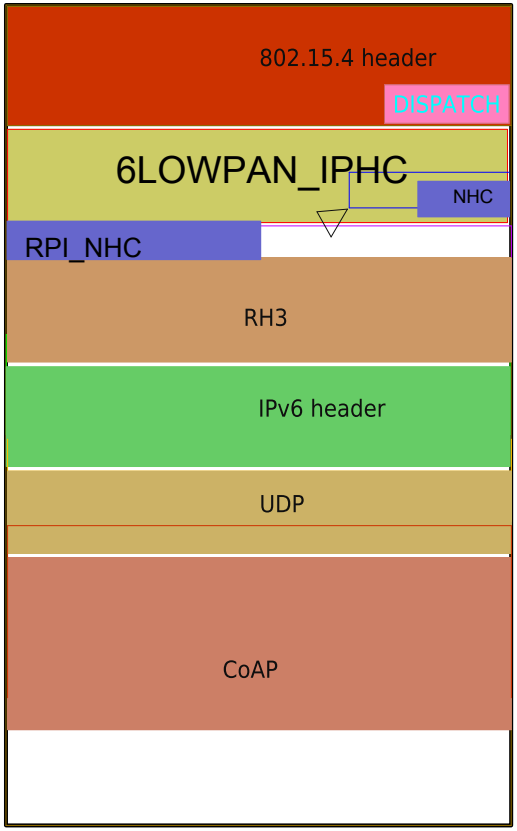

Dispatch: architectural view

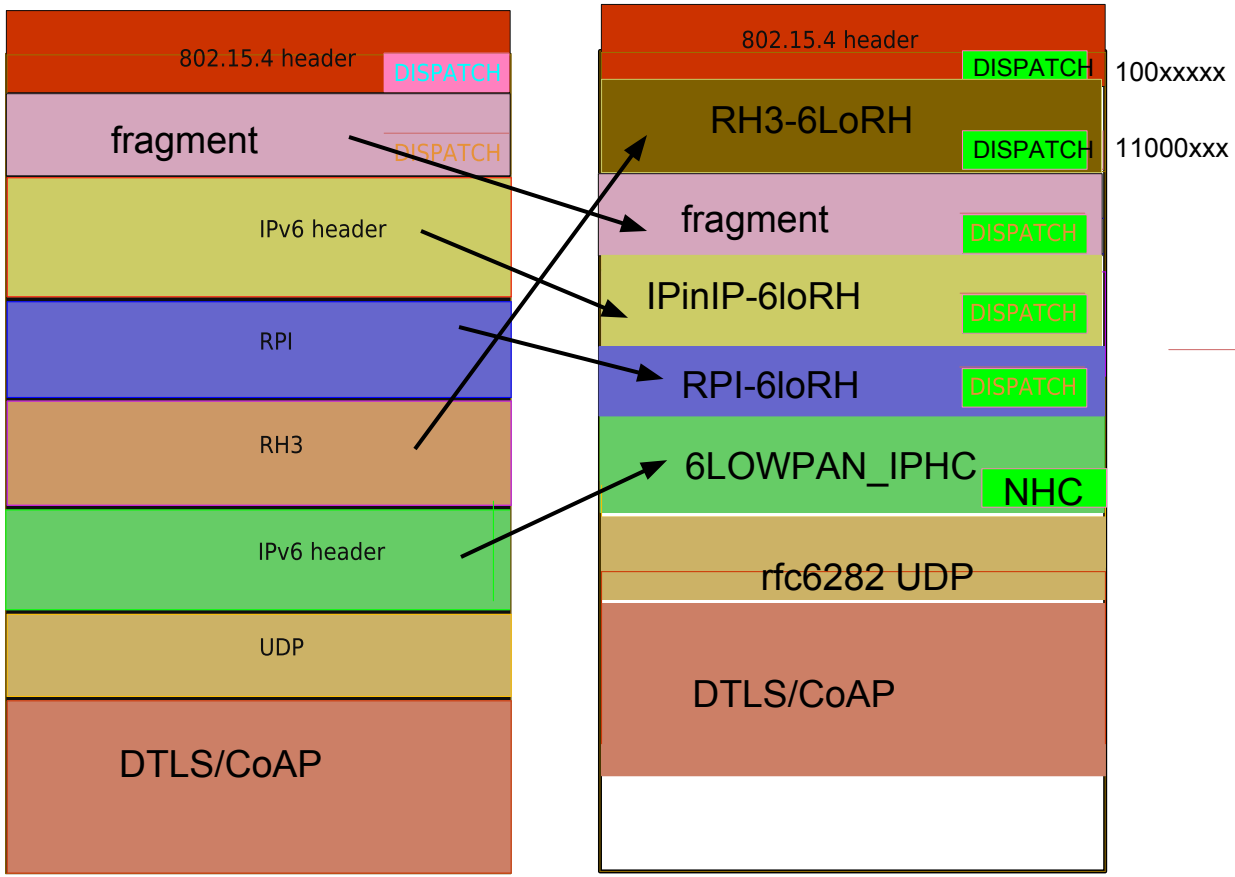

BYTE SIZES ARE NOT TO SCALE

draft-thubert-6lo-rpl-nhc

draft-thubert-6lo-rpl-nhc

Status

- 02 published
- Priority work due to 6TiSCH dependency

Includes

- 6lo next header compression (NHC) for the RPI*
- 3 proposals for the exact format to debate

Does not include

- RH3 compression
- IP-in-IP compression

* Really: a compressed form for an IPV6 HbH header with a RPL Option that transports the RPL Packet Information (RPI)

Compressing the RPLInstanceId

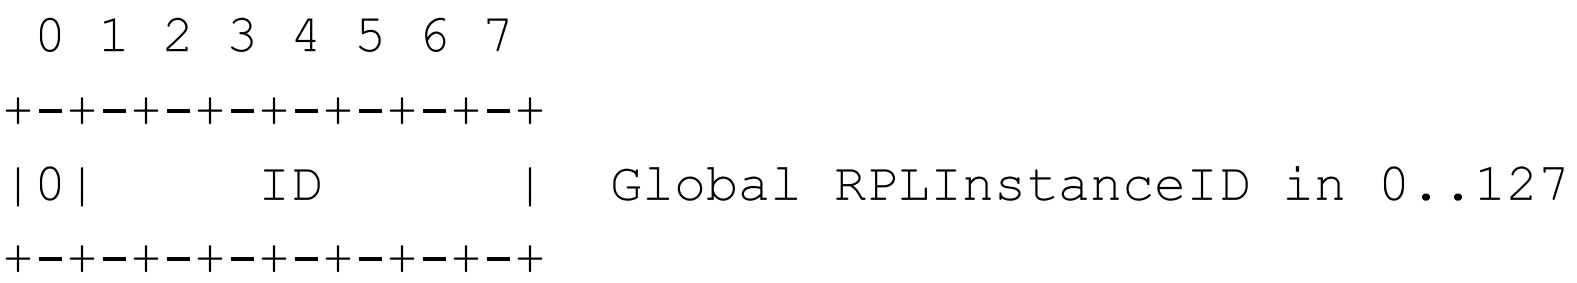

For the global Instance 0, the RPLInstanceID field is all zeroes

Define ‘I’ flag that, when set, signals a global Instance 0, elided.

Compressing the SenderRank

```
DAGRank(rank) = floor(rank/MinHopRankIncrease)
```

If MinHopRankIncrease is set to a multiple of 256, the least significant 8 bits of the SenderRank will be all zeroes; by eliding those, the SenderRank can be compressed into a single byte.

Define a 'K' flag that, when set, signals that a single byte is used.

Compressing the RPLInstanceId

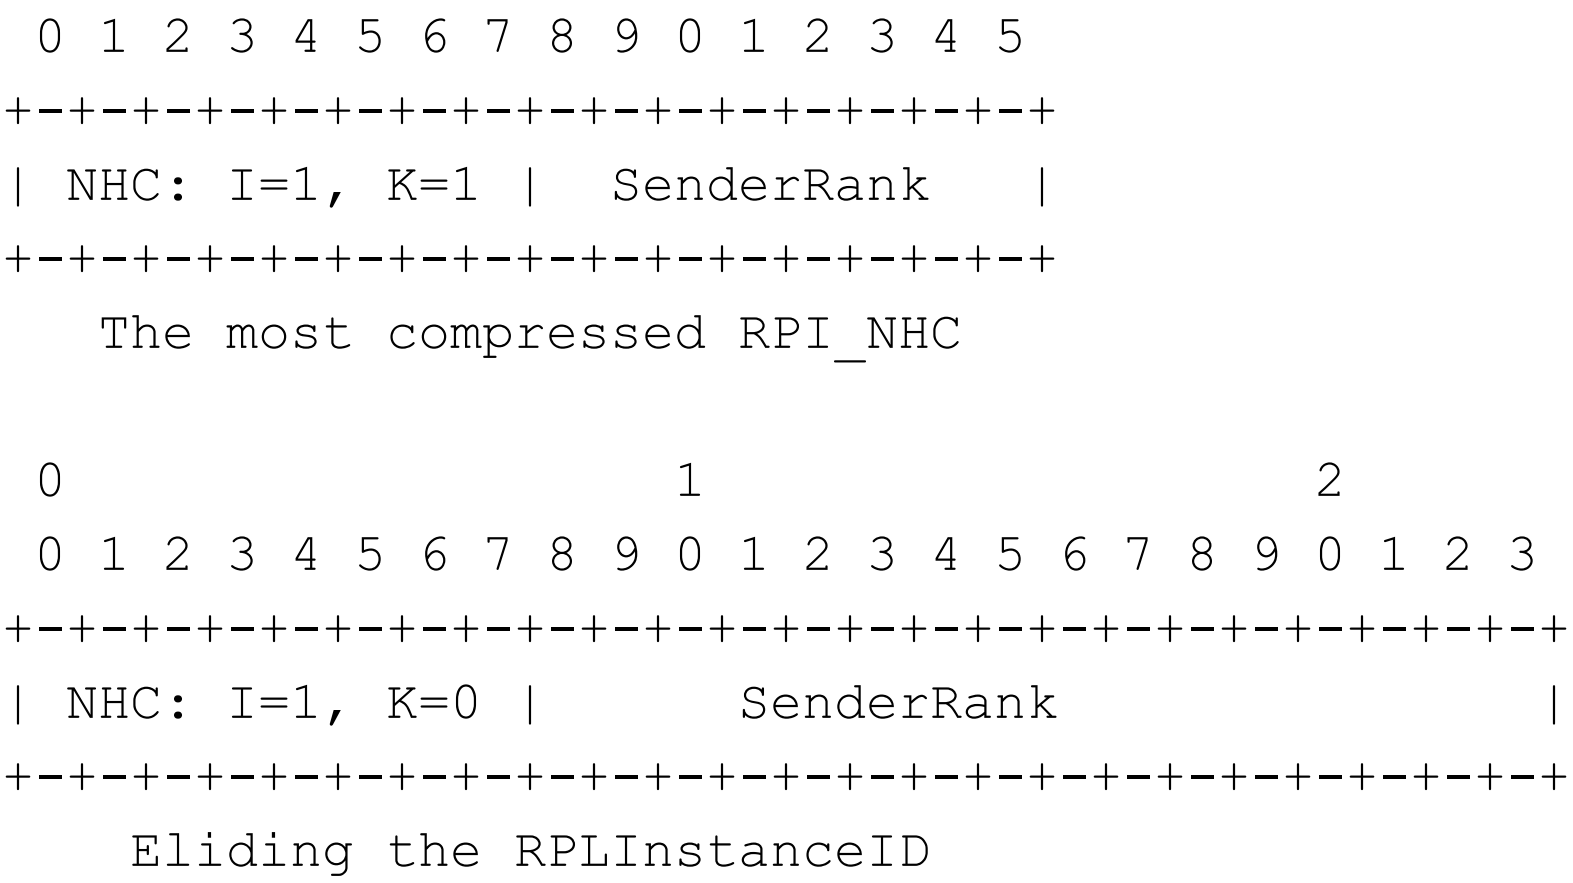

Compressing the RPLInstanceId

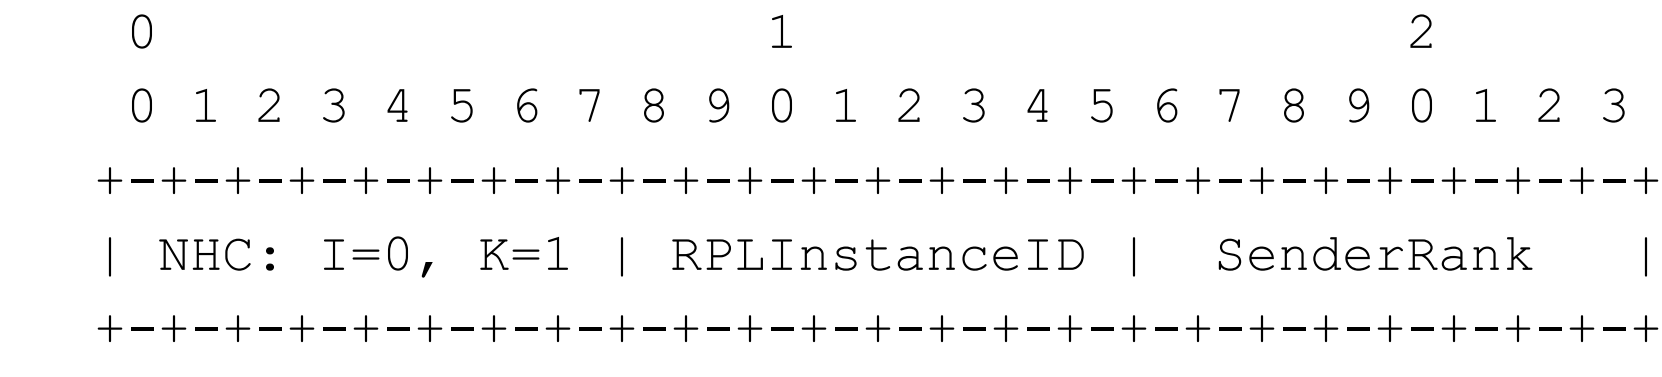

Compressing SenderRank

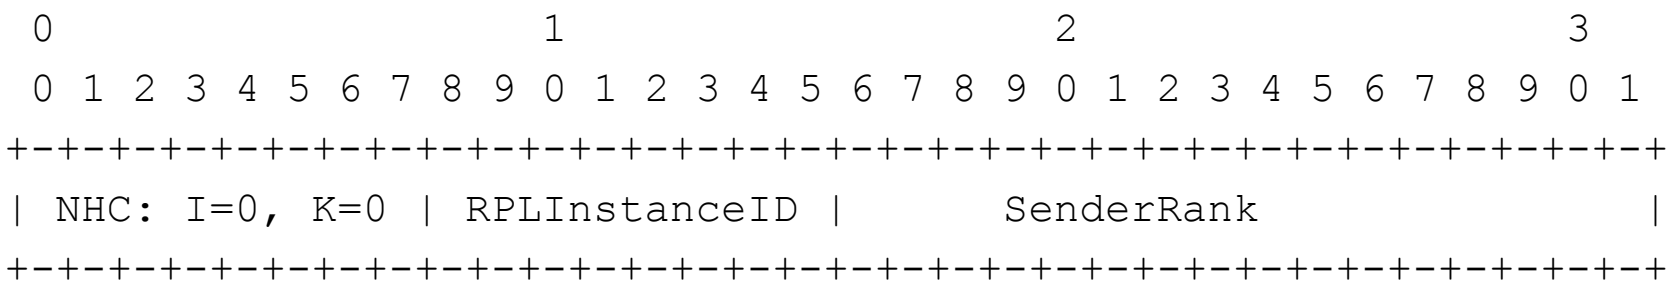

Least compressed form of RPI_NHC

Possible NHC encodings

- “Greedy” encoding option
  - Uses NHC 1:0:x:x:x:x:x:x for RPL opt. flags + NH flag
- “Conservative” encoding option
  - Adds to Ext. ID enum. of RFC 6282 LOWPAN_NHC
  - An additional byte for RPI flags (vs. Greedy)
- “Efficient” encoding option
  - Trades Complexity for NHC footprint
  - Compression as good as “Greedy” in normal case
  - Breaks from simple RFC 6282 encoding model
  - Defines new escape mechanism
  - Inserts byte upon forwarding exception

draft-thubert-6lo-rpl-nhc

Status

- 02 published
- Priority work due to 6TiSCH dependency

Includes

- 6lo next header compression (NHC) for the RPI*
- 3 proposals for the exact format to debate

Does not include

- RH3 compression
- IP-in-IP compression

* Really: a compressed form for an IPV6 HbH header with a RPL Option that transports the RPL Packet Information (RPI)

draft-thubert-6lo-routing-dispatch

draft-thubert-6lo-routing-dispatch

Status

- 03 published
- Federates authors of previous proposals

Features

- Compression of all RPL artifacts
- Extensible through TLV
- Reuse of the mesh header in non-mesh-under
  - | 10 xxxxxx | MESH => 1/3 of 6LoWPAN addressable
  - Not backward compatible,implies different networks
  - Proposes a mesh header compression for mixed mode

Why not NHC++ ?

RFC 6282 Code Point starvation

- NHC “greedy” discussion => conservative use

RFC 6282 Code Base

- Backward compatible code on IP
- Separate routing operation from endpoint

RPL or not RPL?

- TLV space for other routing and tagging purposes

Legacy hosts

- Allows decaps. of 6LoRH at penultimate node

Generic TLV format

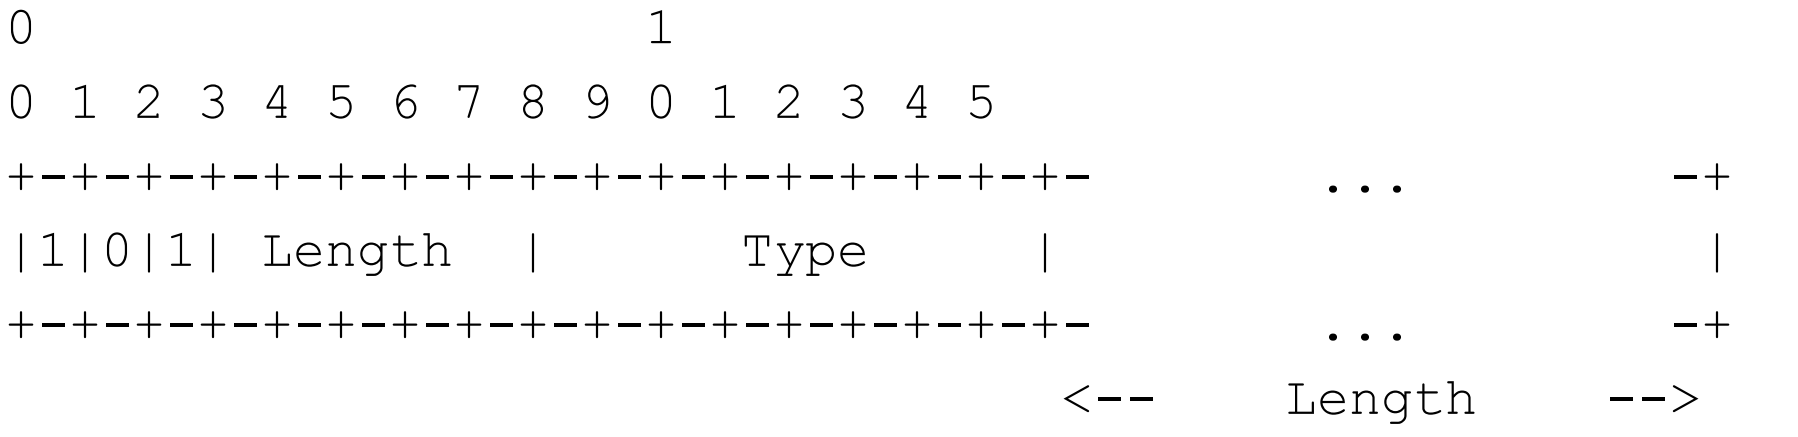

Elective format (above) is skippable thanks to standard length field whereas Critical format is not. Length is derived from Type/ Type Specific Extension

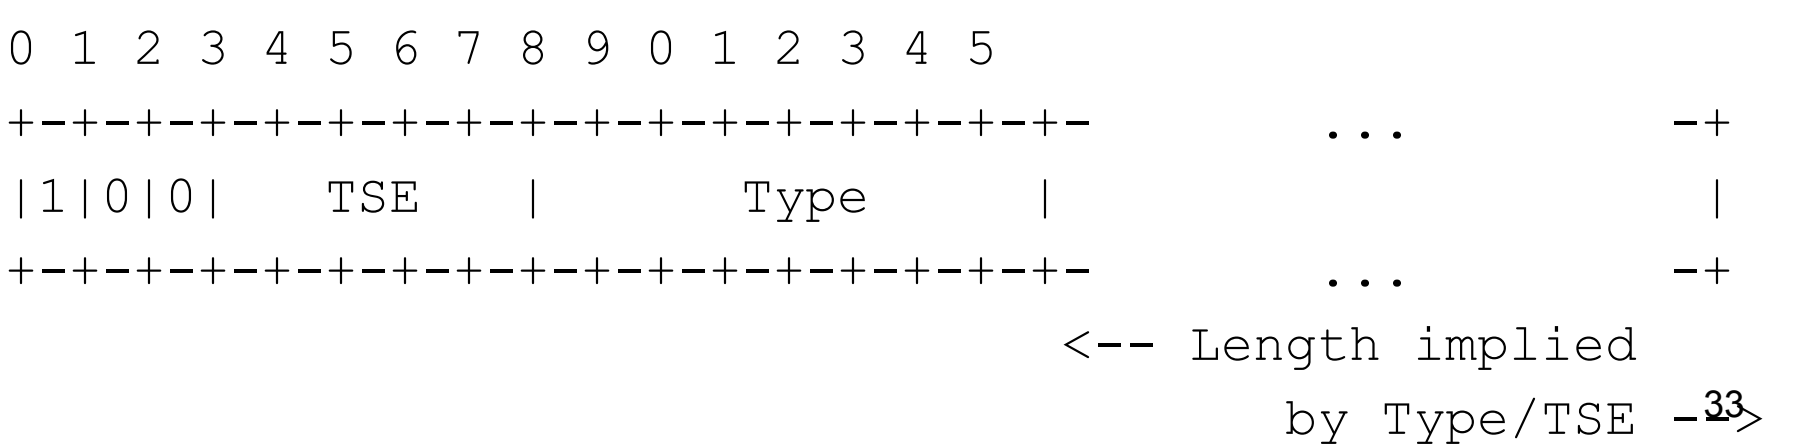

RPI-6LoRH

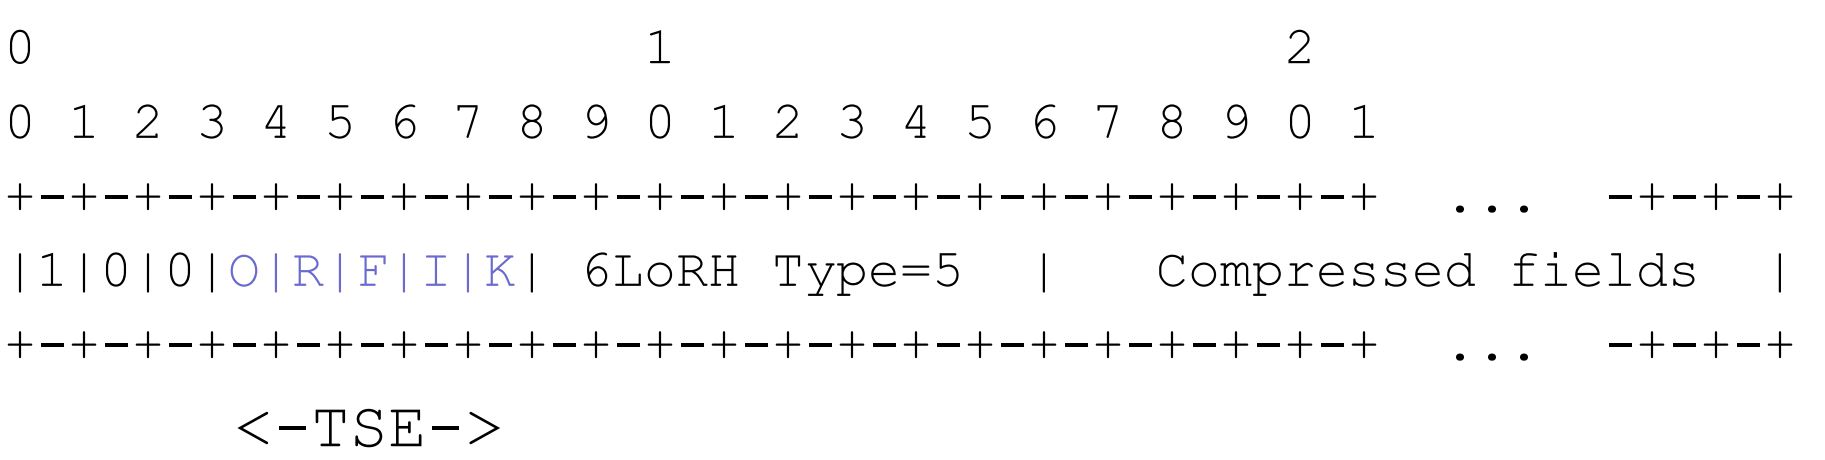

Critical format,

O,R,F bits from RFC 6553 in the TSE

Same I, K flags as in NHC proposal, placed in TSE as well

Compressed fields inferred from I and K flags, 1 to 3 bytes

RH3-6LoRH

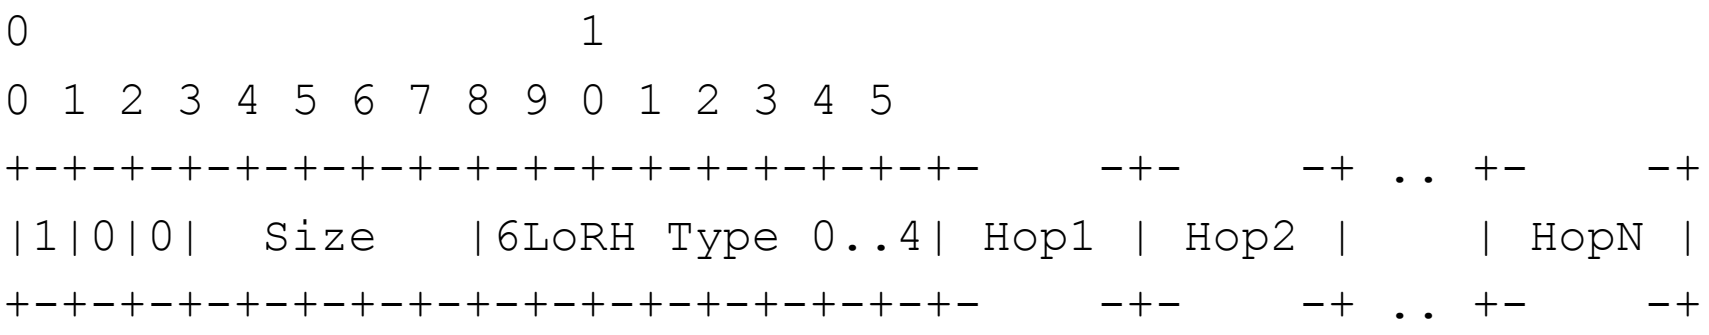

Size indicates the number of compressed addresses (minus1)

The form of compression is indicated by the Type as follows:

Type	Size Unit
0	1
1	2
2	4
3	8
4	16

< More details on compression needed in next version >

Example use of RH3-6LoRH

Root sends

```
0x80 0x04 <2001:DB8::ABCD>       -> 1 * 16-bytes address
0x82 0x01 0xCAFE 0xBEEF 0xCA5A   -> 3 * 2-bytes addresses
```

2001:DB8::ABCD forwards

```
0x80 0x04 <2001:DB8::CAFE>       -> 1 * 16-bytes address
0x81 0x01 0xBEEF 0xCA5A          -> 2 * 2-bytes addresses
```

2001:DB8::CAFE forwards

```
0x80 0x04 <2001:DB8::BEEF>       -> 1 * 16-bytes address
0x80 0x01 0xCA5A                 -> 1 * 2-bytes addresses
```

2001:DB8::BEEF forwards

```
0x80 0x04 <2001:DB8::CA5A>       -> 1 * 16-bytes address
```

2001:DB8::CA5A removes the 6LoRH and routes to dest. in IPHC

RPI-6LoRH

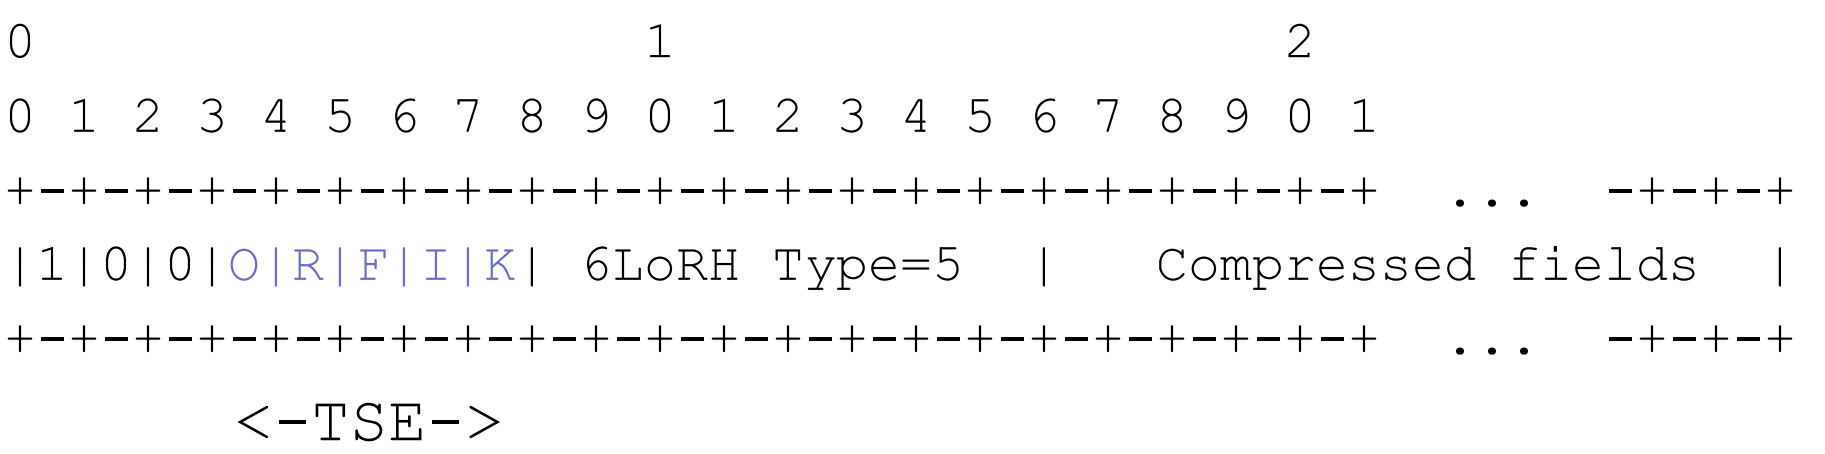

Critical format,

<need to indicate whether to remove on penultimate>

O,R,F bits from RFC 6553 in the TSE

Same I, K flags as in NHC proposal, placed in TSE as well

Compressed fields inferred from I and K flags, 1 to 3 bytes

Optimized IP in IP

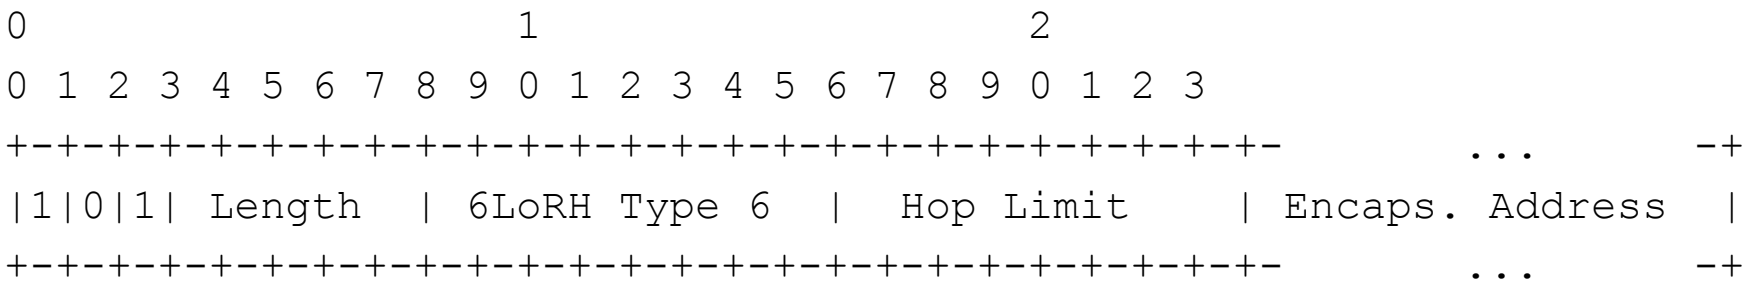

Elective format, size of 1 indicates encapsulator is root.

<need to provide compression when not>

Usually, only one byte of data for Hop Limit

Inner (IPHC) hop limit can be compressed elided

If destination downwards is same as that of IPHC,

or root on upward packets => can be elided too

Use RH when destination cannot be elided, examples given

IETF 93 - 6tisch Interop

Which approach to use for that?

Conclusions:

to be filled in during meeting

Acknowledgment

Thank you very much to Pascal Thubert for his help on the slides