TRILL over IP

draft-ietf-trill-over-ip-01.txt

IETF 91, Honolulu

Margaret Wasserman <mrw@painless-security.com>
Donald Eastlake, Dacheng Zhang

Current Document Summary

- Defines a natural IP/UDP encapsulation for TRILL.
- Treats an IP network as a link connecting TRILL switch ports thus providing a method to connected remote TRILL sites into a single TRILL campus.
- Two Scenarios are described
  - Remote Office Scenario
  - IP Backbone Scenario

Encapsulation in Current Draft

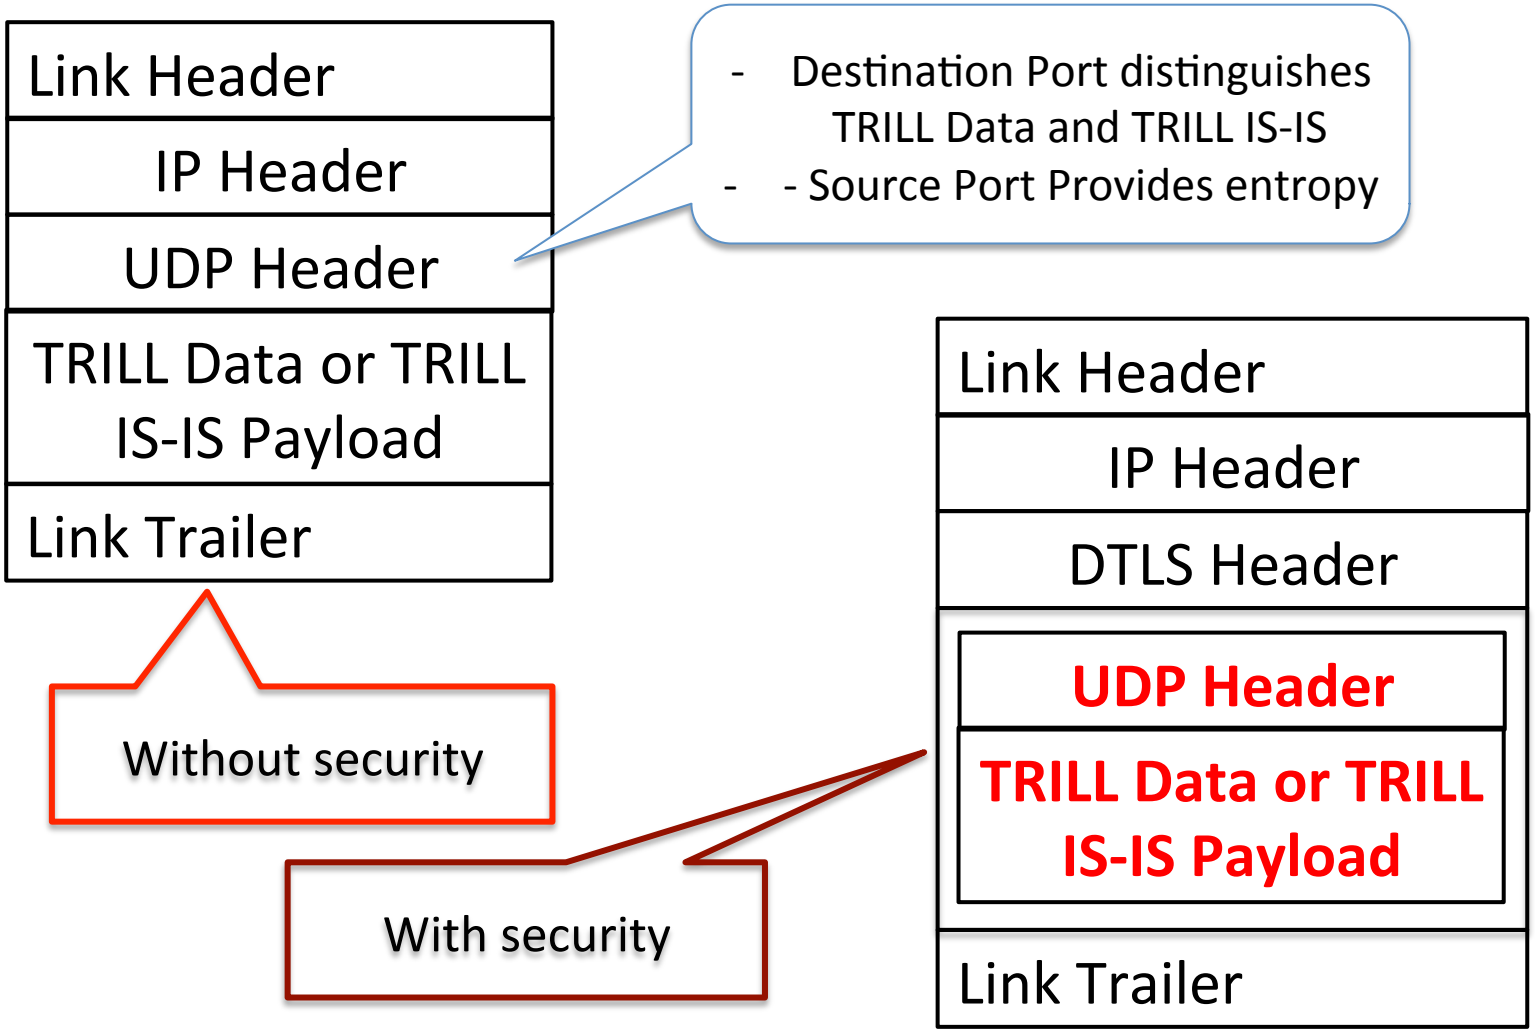

Current Document Summary (cont)

- Uses DTLS for security.
  - Does not interfere with IS-IS authentication of TRILL IS-IS packets.
  - If TRILL over IP switches support certificates, they MUST support :
    - TLS_RSA_WITH_AES_128_CBC_SHA256
  - If TRILL over IP switches support pre-shared keys, they MUST support:
    - TLS_PSK_WITH_AES_128_CBC_SHA256

Work Remaining

- Congestion Considerations section is inadequate
- Middle Box Considerations section is empty
- QoS Considerations are absent (DSCP)
- Needs a clear specification of the configuration associated with a TRILL over IP port
- Current draft did not take into account hardware support of encapsulation or security protocol:
  - Fast path support is important for demanding applications.

Question: Security Protocol

- Fast path hardware support is more common for IPSEC than for the currently mandated DTLS.
  - In either case, default keying can be derived from IS-IS keying so we are primarily talking about the data format, not necessarily the key exchange.
- It seems undesirable to have to support both.
- Should TRILL over IP change to using IPSEC as the mandatory to implement security?

Question: Encapsulation

- The current draft only supports natural UDP encapsulation. But there is more fast path hardware support and perhaps more flexibility with other encapsulations such as VxLAN.
- There was a consensus determination that the TRILL WG preferred UDP/IP over a new custom encapsulation (such as a new IP protocol type number) but we are not talking about either here.
- This encapsulation question is essentially independent of the security question.

Question: Encapsulation (cont.)

- Suggestion:
  - The default mode for a TRILL over IP port could be to exchange Hellos using natural encapsulation.
    - TRILL Hellos are sent at most once a second so this could be done in software.
  - The port capabilities sub-TLV in each Hello would indicate what encapsulations the sending port is willing to use.
    - Could vary between ports on the same switch due to port hardware.
  - Data connectivity is only established if TRILL switches have a common supported and enabled encapsulation.
  - A TRILL over IP port could also be configured to use one specified encapsulation for all TRILL communications.

Feedback? Questions?

Back up slides

THE TRILL ENCAPSULATION ARCHITECTURE

TRILL Link Encapsulations

- A TRILL link protocol encapsulation needs to:
  - Get a TRILL packet from one TRILL switch port to another TRILL switch port over the link.
  - Provide one mandatory to implement variation for interoperability.
  - Distinguish between TRILL Data packets and TRILL IS-IS packets.
  - If the link can have more than two ports on it, provide the address of the destination port(s).
  - Maybe other stuff depending on link technology.

In RFC 6325

- 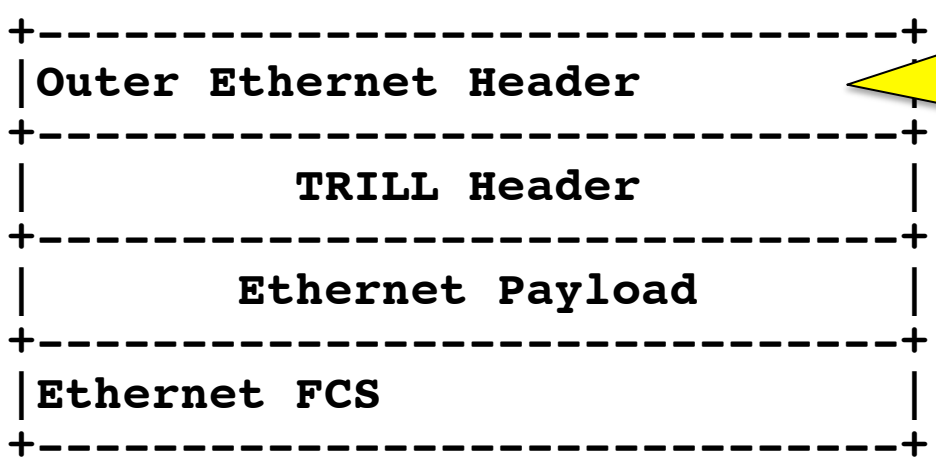

TRILL over Ethernet:
Ethernet Header before TRILL Header. Outer addresses needed because Ethernet link could be a bridged LAN with many stations on it.

- 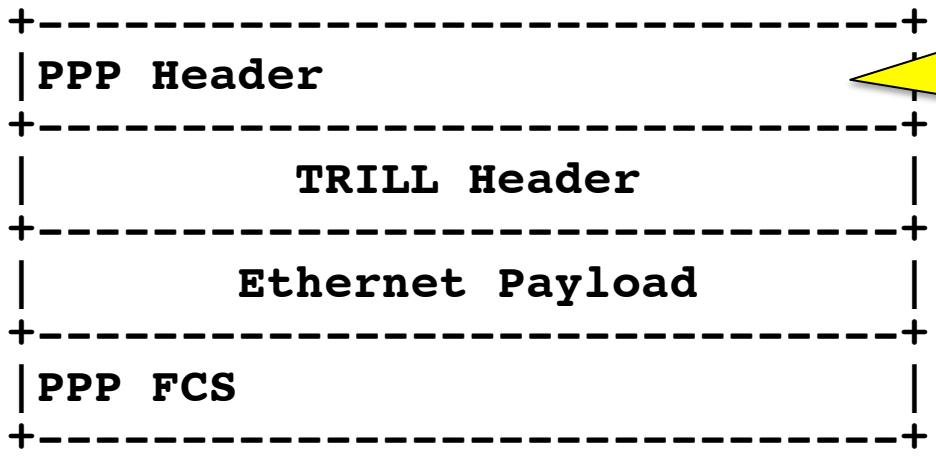

TRILL over PPP:
No addresses needed.
No Ethernet Header before TRILL Header

TRILL Link Encapsulaton

- In TRILL over Ethernet, Ethertypes indicate TRILL Data (0x22F3) or TRILL IS-IS (0x22F4). [RFC 6325]
- In TRILL over PPP, PPP code points indicate TRILL Data (0x005D) or TRILL IS-IS (0x405D). [RFC 6361]

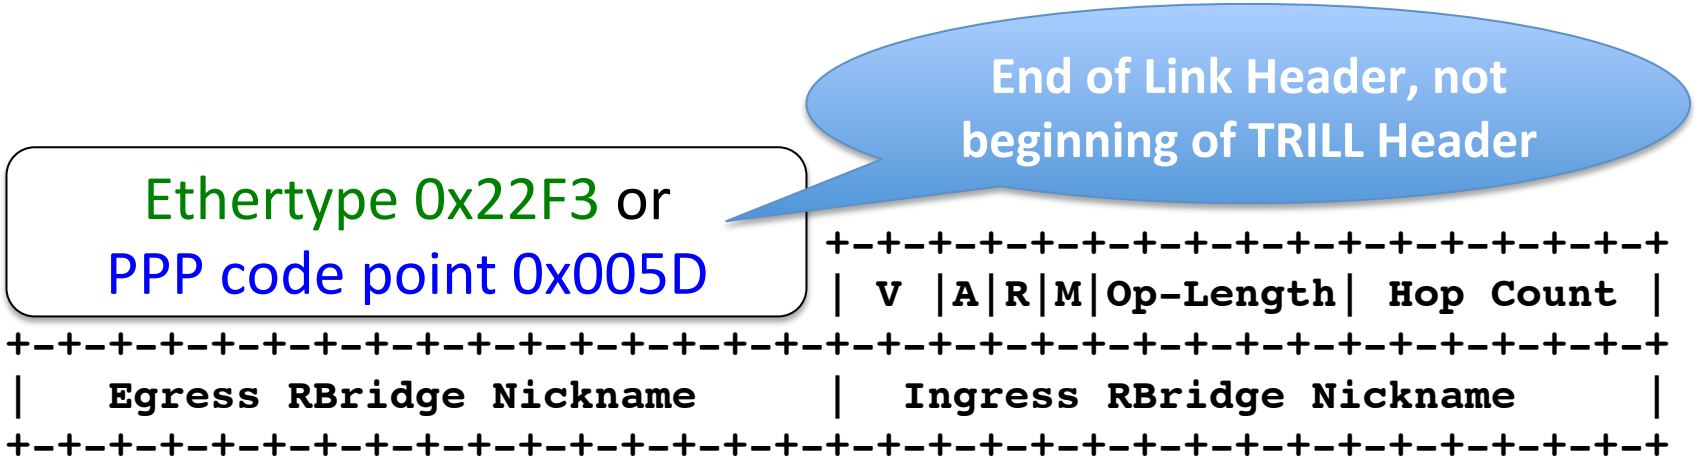

The 6-byte TRILL Data Header

The IP Link Protocol

- What about TRILL over IP?
  - (Use of IP does not necessarily imply long distance. You can have a local IP core and long distance carrier Ethernet, for example.)
- As with any other Link protocol, its purpose is to get a TRILL packet from one TRILL switch port to another and distinguish TRILL Data from TRILL IS-IS.
- The source TRILL switch IP port and the destination TRILL switch IP port have IP addresses which are provided by an IP Header.

The IP Link Protocol (cont.)

- An IP Link will be one TRILL hop but could be composed of multiple IP hops.

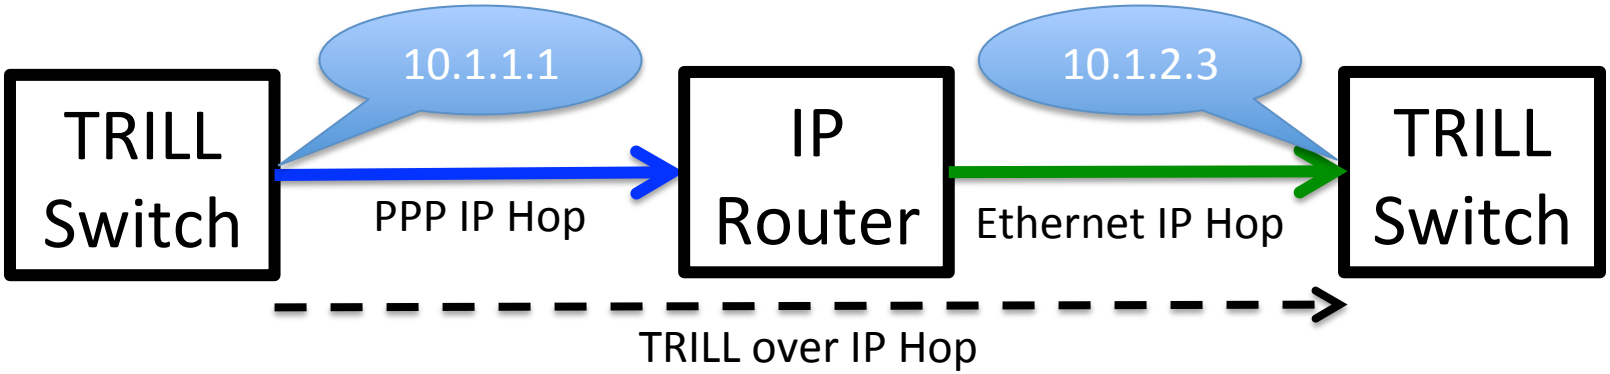

- Each IP hop composing the TRILL hop is over some lower layer, possibly different for each hop, and all irrelevant at the TRILL layer.

The IP Link Protocol (cont.)

- So you have an IP header and a TRILL header.

- You still need something in between to distinguish data from IS-IS (unless you use up two IP Protocol number and never care about problems with middle boxes due to unknown IP Protocol numbers) and provide entropy.

The IP Link Protocol (cont.)

- You could require TRILL over Ethernet over IP but:
  - You would be adding 12 bytes of useless “MAC addresses” that would be thrown away by the next TRILL switch in the path.
  - It would be inconsistent with the TRILL link encapsulation architecture in RFC 6325 and the standardized method of doing TRILL over PPP (RFC 6361) and TRILL over pseudowire (RFC 7174).