TRILL over IP

draft-ietf-trill-over-ip-02.txt

IETF 92, Dallas

Margaret Wasserman mrw@painless-security.com

Dacheng Zhang, Donald Eastlake,

Document Summary

- "TRILL over IP" treats an IP network as a link connecting TRILL switch ports, thus providing a method to connected TRILL sites into a single TRILL campus.
- Two Scenarios are described in the draft
  - Remote Office Scenario
  - IP Backbone Scenario
- Specifies encapsulation, security, and transport considerations including congestion, MTU, fat flows, recursive ingress, …

Changes from -01

- Changes primarily motivated by the hardware support required for high data rates:
  - Security: Use of IPsec instead of DTLS due to better hardware support available for IPsec. This change is in the current Version -02.
  - Encapsulation: Use of alternative encapsulations with better hardware support, planned for next version -03.
- Also Section 6 on Port Configuration added in -02.

Security

- Draft now specifies IPsec ESP (Encapsulating Security Protocol) in Tunnel Mode.
  - Some details needs to be filled in such as
    - mandatory to implement crypto algorithms
    - details of default keying and key negotiation.
  - Use of ESP Tunnel Mode supports use of IPsec appliances separate from the actual RBridge port hardware.

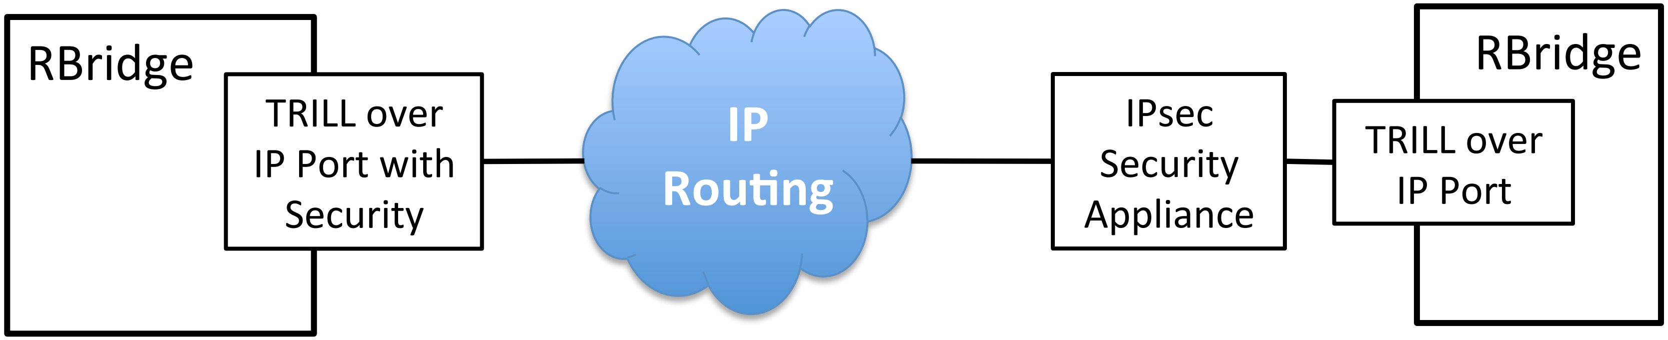

IPsec ESP in Tunnel Mode

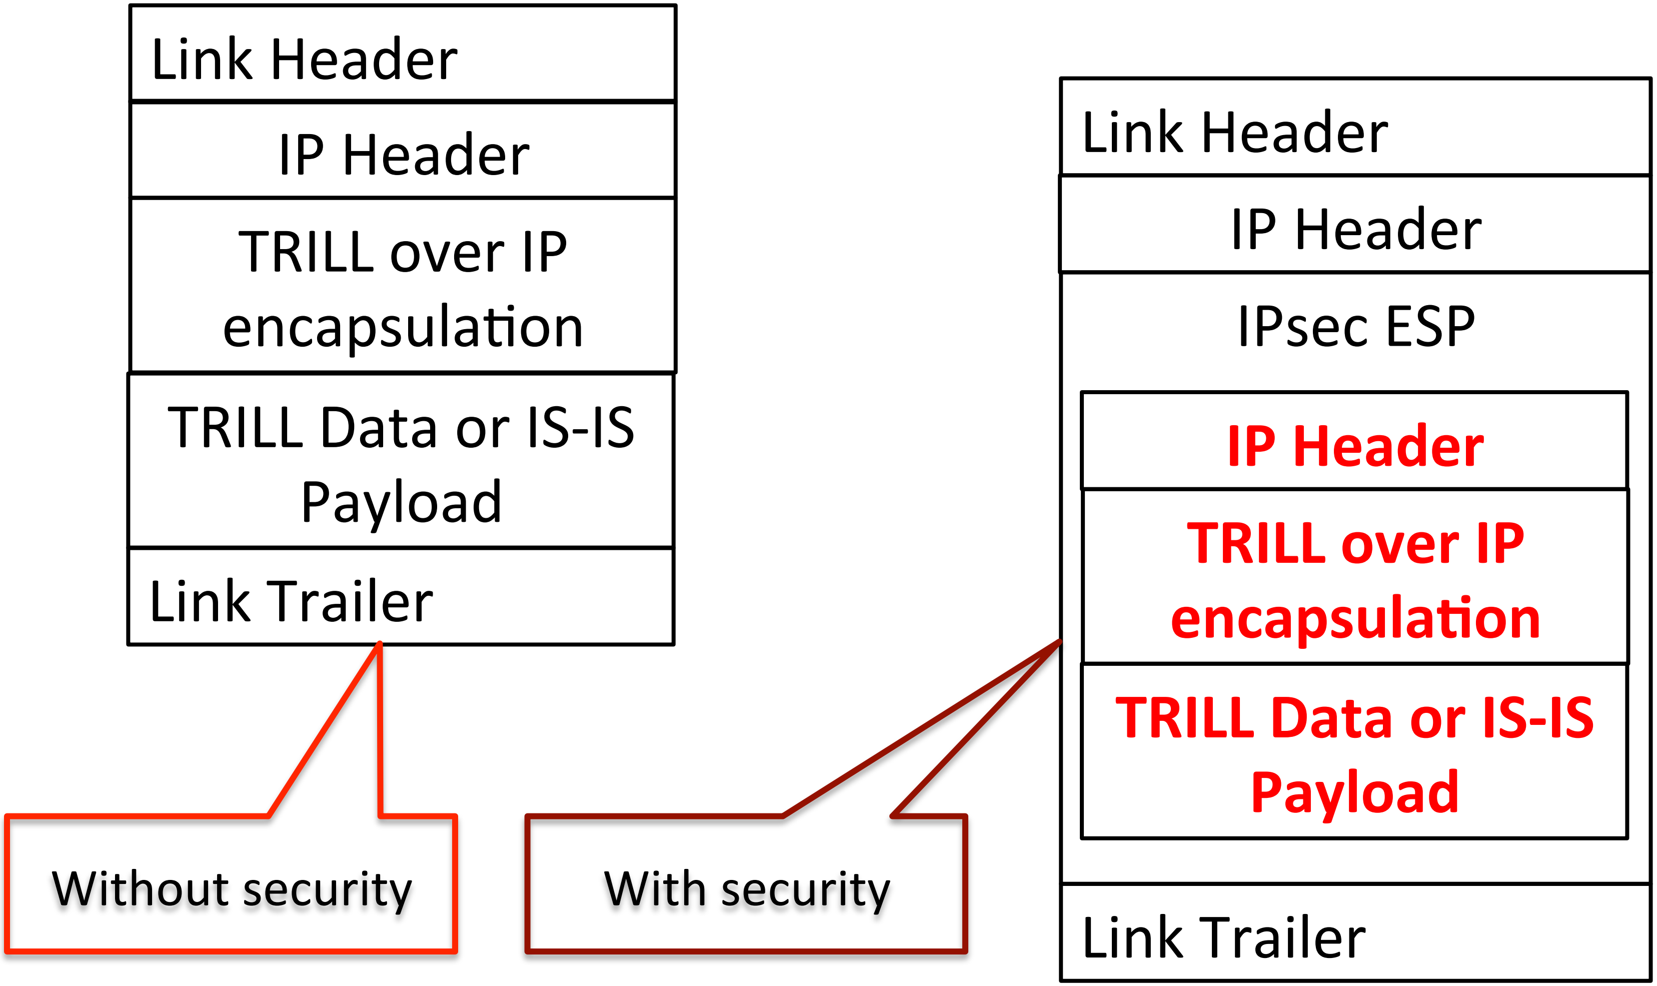

Link Header

IP Header

TRILL over IP encapsulation

TRILL Data or IS-IS Payload

Link Trailer

Link Header

IP Header

IPsec ESP

IP Header

TRILL over IP encapsulation

TRILL Data or IS-IS Payload

Link Trailer

Without security

With security

Encapsulation

- The current draft only specifies direct UDP encapsulation. But there is better fast path hardware support and more flexibility with other encapsulations such as VxLAN.
  - "UDP encapsulation" is really TRILL over UDP over IP. TRILL Data versus IS-IS is indicated by destination UDP socket.
  - "VxLAN encapsulation" with current VxLAN [RFC7348] is really TRILL over Ethernet over VxLAN over UDP over IP. TRILL Data versus IS-IS is indicated by EtherType but the Ethernet DA&SA are 12 bytes of wasted space.
  - Other encapsulations are being developed in other working groups. We might optionally use those but there is no proposal to develop an encapsulation in the TRILL WG

Encapsulation

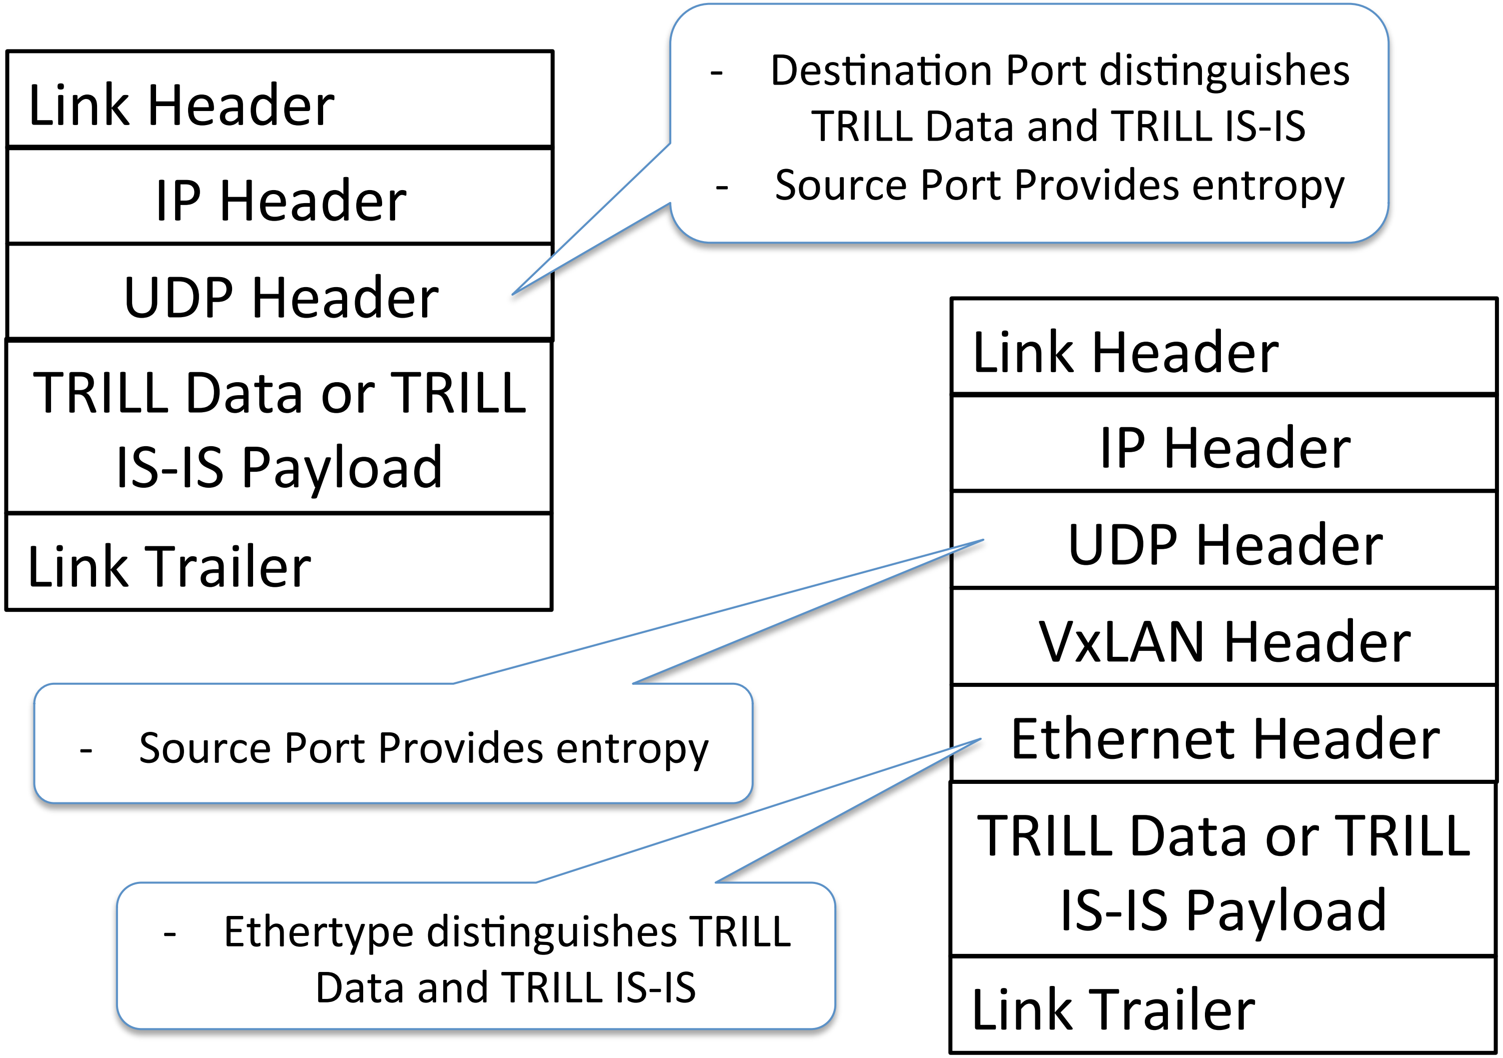

Link Header

IP Header

UDP Header

TRILL Data or TRILL IS-IS Payload

Link Trailer

- Destination Port distinguishes TRILL Data and TRILL IS-IS
- Source Port Provides entropy

Link Header

IP Header

UDP Header

VxLAN Header

Ethernet Header

TRILL Data or TRILL IS-IS Payload

Link Trailer

- Source Port Provides entropy

- Ethertype distinguishes TRILL Data and TRILL IS-IS

Encapsulation

- Proposal:
  - The initial mode for a TRILL over IP port would be to exchange Hellos and E-L1CS LSPs using UDP encapsulation.
    - This is a small enough amount of traffic it can be done in software.
  - What data encapsulations a port is willing to use, in priority order, can be advertised in Hellos or E-L1CS LSPs. Can vary between ports due to port hardware.
  - Data connectivity (adjacency) is established if TRILL switches have a common supported and enabled encapsulation.
  - A TRILL over IP port could also be configured to always use a specified encapsulation for all TRILL communications.

Other Work Remaining

- Other work remaining includes:
  - QoS Considerations are absent (how to map TRILL packet priority to IP)
  - Middle Box Considerations section is empty.

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 TRILL Base RFC 6325

-
  ```
  +------------------------------+
  |Outer Ethernet Header         |
  +------------------------------+
  |         TRILL Header         |
  +------------------------------+
  |       Ethernet Payload       |
  +------------------------------+
  |Ethernet FCS                  |
  +------------------------------+
  ```

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

-
  ```
  +------------------------------+
  |PPP Header                    |
  +------------------------------+
  |         TRILL Header         |
  +------------------------------+
  |       Ethernet Payload       |
  +------------------------------+
  |PPP FCS                       |
  +------------------------------+
  ```

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]

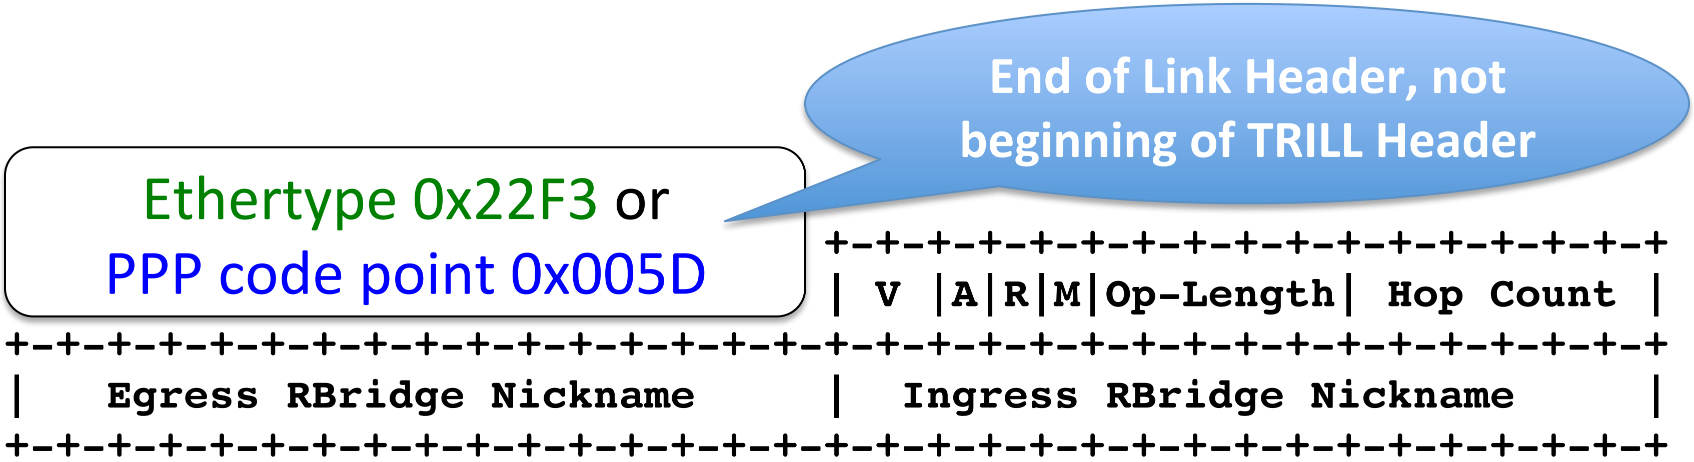

Ethertype 0x22F3 or
PPP code point 0x005D

End of Link Header, not
beginning of TRILL Header

```
                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                               | V |A|R|M|Op-Length| Hop Count |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Egress RBridge Nickname    |  Ingress RBridge Nickname     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The 6-byte TRILL Data Header

The IP Link Protocol

- What about TRILL over IP?
  - (Use of IP instead of Ethernet 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.

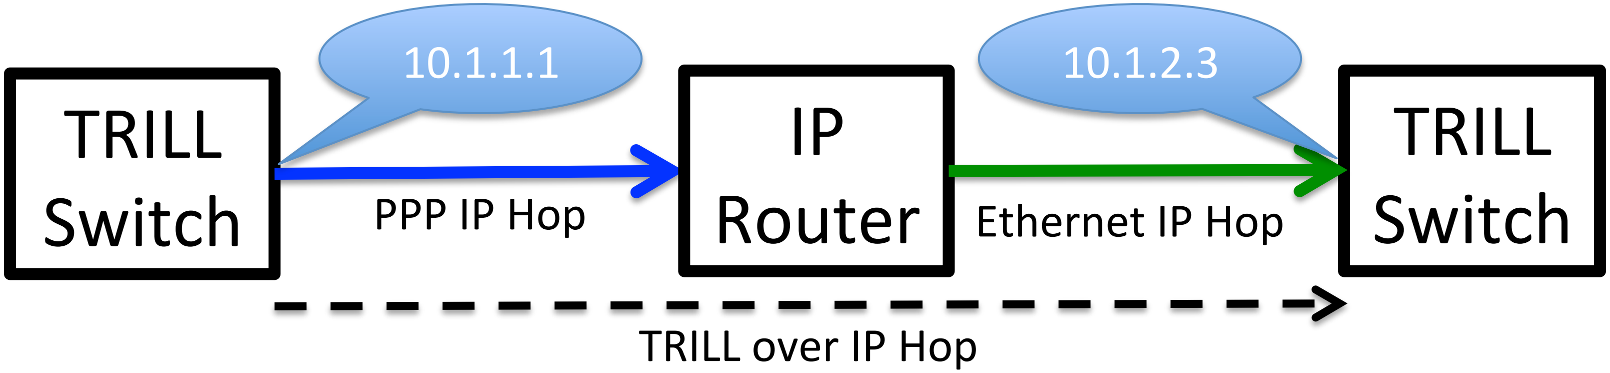

- Each IP hop composing the TRILL hop is over some lower layer, possibly different for each hop, and all irrelevant at the TRILL layer.

The IP Link Protocol (cont.)

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

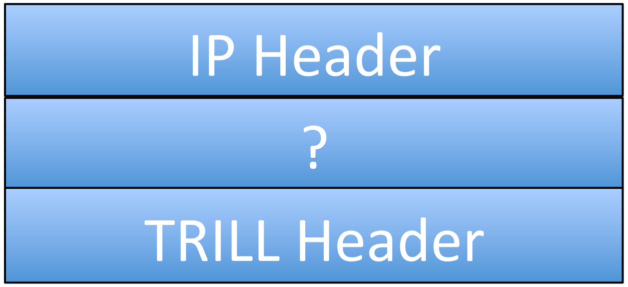

IP Header

?

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 always require TRILL over Ethernet [over x] 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).