Pseudonode Nickname

draft-hu-trill-pseudonode-nickname-00

Hongjun Zhai (ZTE)
Fangwei Hu (ZTE)
Radia Perlman (Intel Labs)
Donald Eastlake 3rd (Huawei)

Contents

- Problem Statement
  - Traffic failure caused by AF change
- Pseudonode Nickname
- Control Plane of Pseudonode Nickname
- Data Plane of Pseudonode Nickname
- TLV Extension for Pseudonone Nickname

Problem Statement

- Only AF(Appointed Forwarder) can ingress/egress native frames into or from TRILL network on a multi-access link
- VLAN-x AF changes for some reasons:
  - DRB changed, and the new DRB maybe designate another RBridge as VLAN-x AF in the link
  - AF down or overloaded, DRB designates another RBridge as VLAN-x AF

Problem Statement

- Issues caused by VLAN-x AF changing:
  - End station can not perceive the AF change
  - MAC entries in remote RBridge's cache can not be updated by self-learning in time
  - Remote traffic to local end station will be destined to the stale VLAN-x AF, and be discarded
  - ESADI can alleviate the issues, but it is only an optional protocol, not running on all ingress/egress RBridges
  - It is even worse for the Swap Nickname Field approach in multi-level TRILL

Problem Statement

- Issues caused by VLAN-x AF changing(examples):

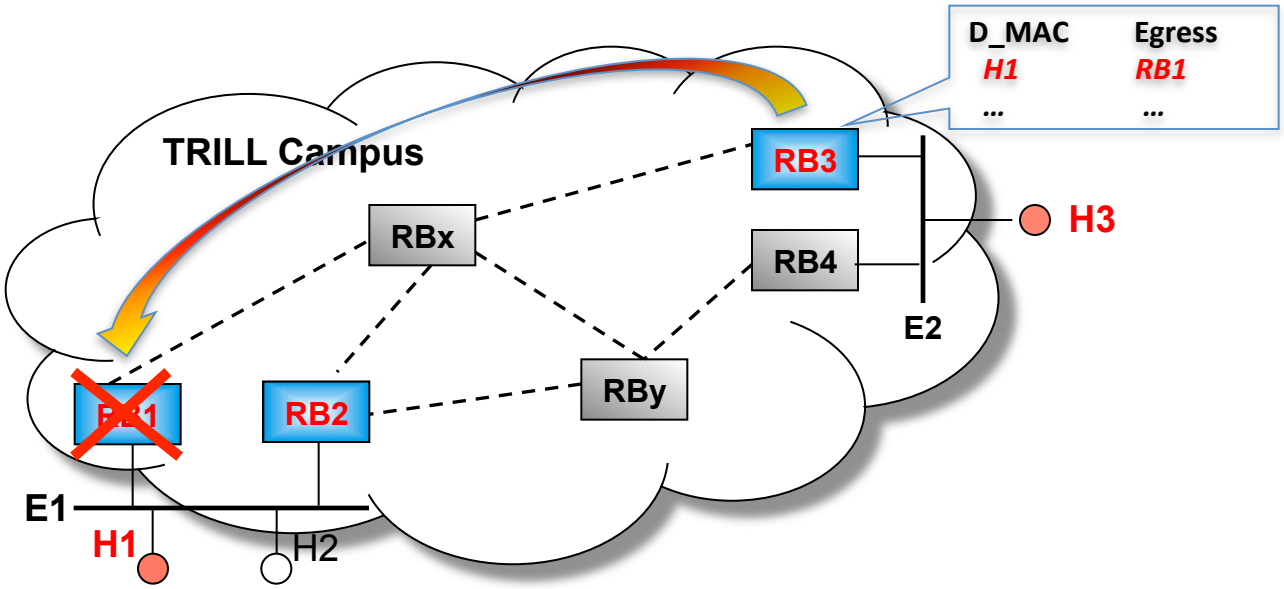

In the RB3's forwarding table, egress is an RBridge, not a link or a group of RBridges, which causes the issue

Contents

- Problem Statement
  - Traffic fails caused by AF change
- Pseudonode Nickname
- Control Plane of Pseudonode Nickname
- Data Plane of Pseudonode Nickname
- TLV Extension for Pseudonone Nickname

Pseudonode Nickname

- Issues is solved by pseudonode nickname:

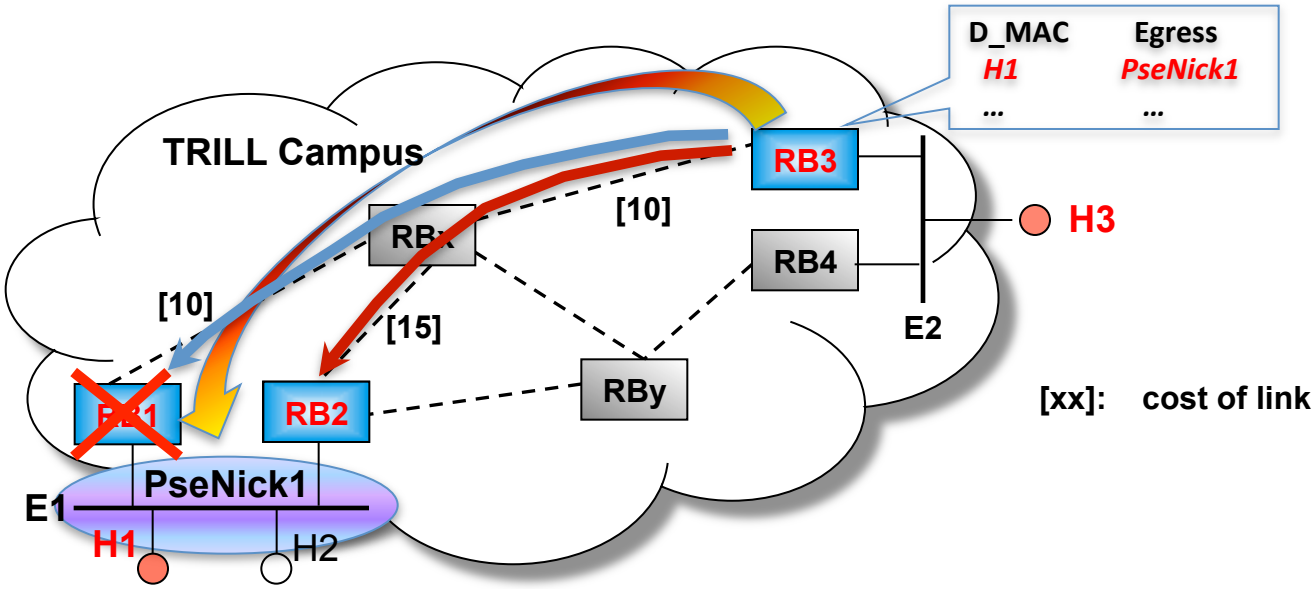

In the RB3's forwarding table, egress is the pseudonode nickname of E2 link, not an RBridge, so the issue is solved

Pseudonode Nickname

- Pseudonode nickname is used to identify a link or a group of Rbridges
- It is assigned by DRB on the link:
  - If there is not a pseudonode nickname on a link, the DRB will contend for one (just as an RBridge nickname) and announce it in its hellos
  - Else, the DRB should reuse the existing pseudonode nickname on this link

Contents

- Problem Statement
  - Traffic fails caused by AF change
- Pseudonode Nickname
- Control Plane of Pseudonode Nickname
- Data Plane of Pseudonode Nickname
- TLV Extension for Pseudonone Nickname

Control Plane of Pseudonode Nickname

- Pseudonode nickname function
  - Each RBridge on the link announces its pseudonode nickname capability in its TRILL Hellos
  - Only if DRB confirms that all the adjacencies in Report state support and enable this capability, this function can be enabled on the link

Control Plane of Pseudonode Nickname

- Pseudonode nickname in Hellos:
  - Each RBridge announces the pseudonode nickname used on this link if it knows such one nickname
  - From the pseudonode nickname contained in adjacencies' hellos, a new DRB can obtain the pseudonode nickname used on this link and reuse it

Control Plane of Pseudonode Nickname

- Pseudonode nickname in LSPs:
  - Pseudonode nickname is announced in DRB's pseudonode LSPs across TRILL campus
  - With the help of pseudonode LSPs and the contained pseudonode nickname, other Rbridges can calculate a shortest path to the pseudonode nickname

Contents

- Problem Statement
  - Traffic fails caused by AF change
- Pseudonode Nickname
- Control Plane of Pseudonode Nickname
- Data Plane of Pseudonode Nickname
- TLV Extension for Pseudonone Nickname

Data Plane of Pseudonode Nickname

- Pseudonode nickname is not used for multicast frames, but for unicast data frames:
  - Multicast TRILL data frames is forwarded along the given Distribution Tree, which contains all the RBridges in TRILL campus
  - So multicast traffic can not fail even if AF has changed;
- Pseudonode nickname influences only ingress and egress RBridges, not transit Rbridges in data plane

Data Plane of Pseudonode Nickname

- Ingress processing:
  - Ingress nickname is pseudonode nickname instead of RBridge's nickname in TRILL header

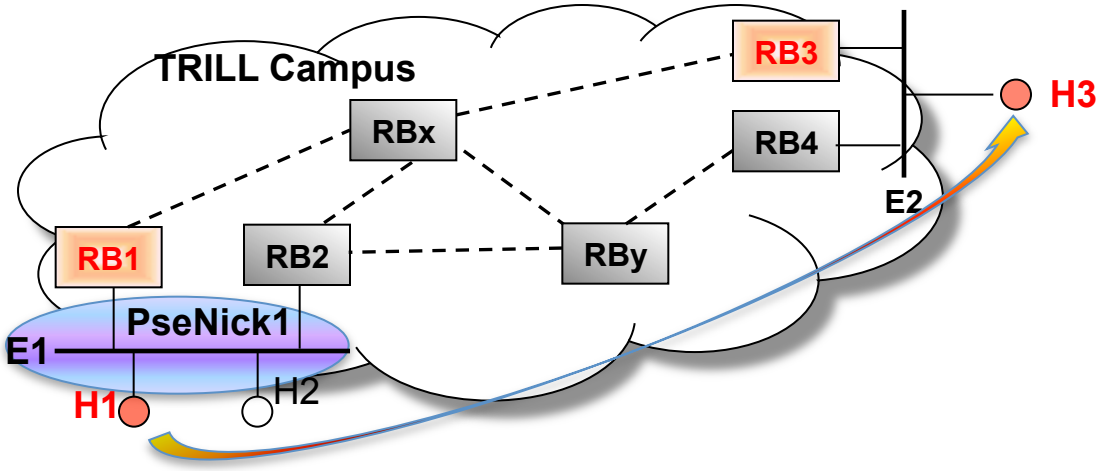

When RB1 encapsulates H1's native data frame into TRILL form, the ingress nickname in TRILL header is PseNick1 not RB1 if the pseudonode nickname function is enabled on link E1

Data Plane of Pseudonode Nickname

- Egress processing:
  - Only VLAN-x forwarder on the destination link can decapsulate the TRILL data frame to native form, and do self-learning
  - In the case that the egress nickname in TRILL header is the pseudonode nickname of one connected link, the receiving RBridge should forward it to the VLAN-x forwarder if it is not
  - Two methods are given for receiving RBridge to forward the frame:
    - Unicasting to VLAN-x AF
    - Multicasting to VLAN-x AF

Data Plane of Pseudonode Nickname

- Egress processing(Unicasting to AF):
  - TRILL header unchanged, the Outer.MacDA is replaced with one unicast MAC of VLAN-x AF on the destination link.

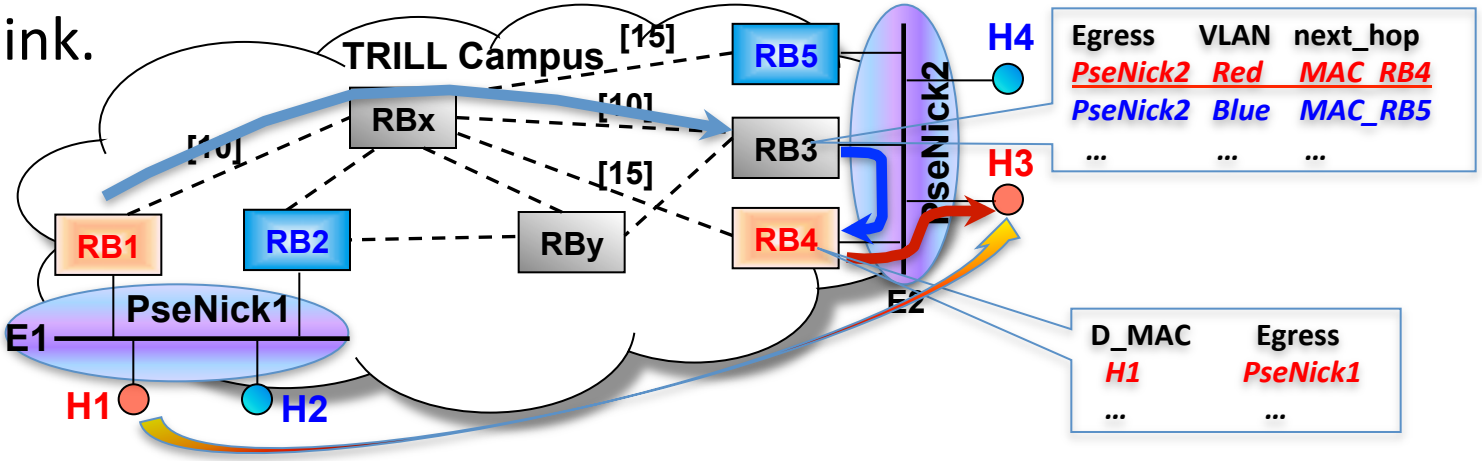

In this approach, the forwarding table has to be based on {nickname, VLAN}, instead of {nickname}, because RB3 must know which Rbridges on E2 is VLAN-Red AF just from the view of data plane

Data Plane of Pseudonode Nickname

- Egress processing(Multicasting to AF):
  - TRILL header unchanged, the Outer.MacDA is replaced with a multicast MAC, i.e., “AF RBridges on this link”

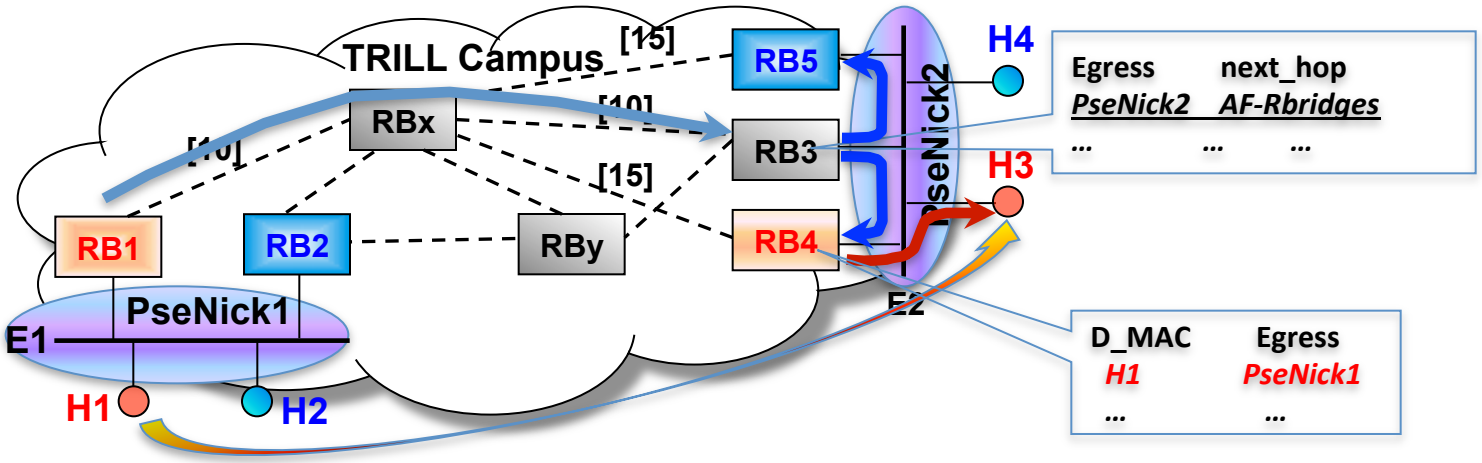

The forwarding remains unchanged in form, but a multicast MAC, i.e., “AF RBridges on this link” is required
Only RB4, the VLAN-Red AF, will decapsulate this TRILL data frame to native frame and forward it to H3

Contents

- Problem Statement
  - Traffic fails caused by AF change
- Pseudonode Nickname
- Control Plane of Pseudonode Nickname
- Data Plane of Pseudonode Nickname
- TLV Extension for Pseudonone Nickname

TLV Extension for Pseudonone Nickname

- Pseudonode Nickname Capability in Hellos

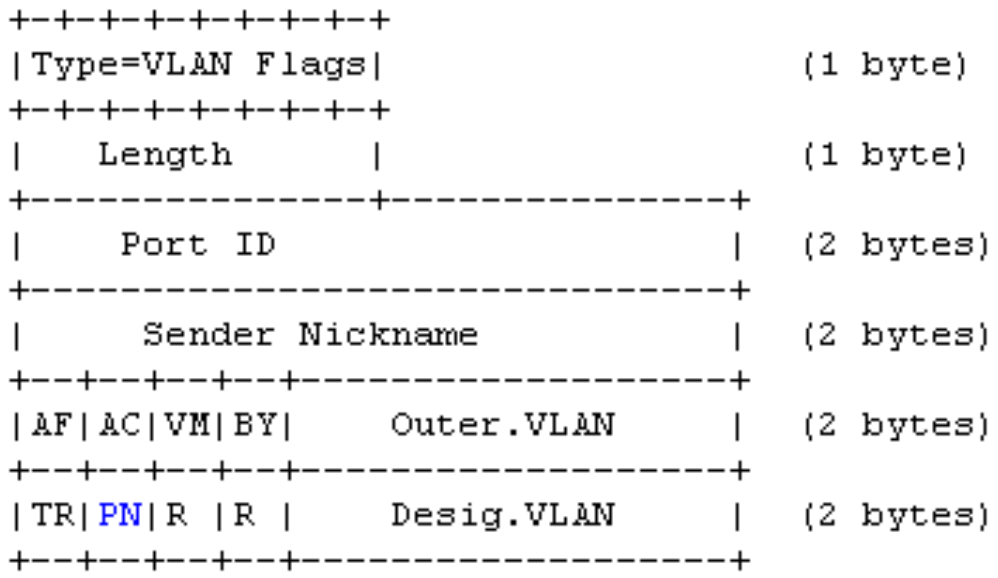

The PN bit is used to indicate whether a RBridge supports and enables the pseudonode nickname

TLV Extension for Pseudonone Nickname

- Pseudonode Nickname TLV

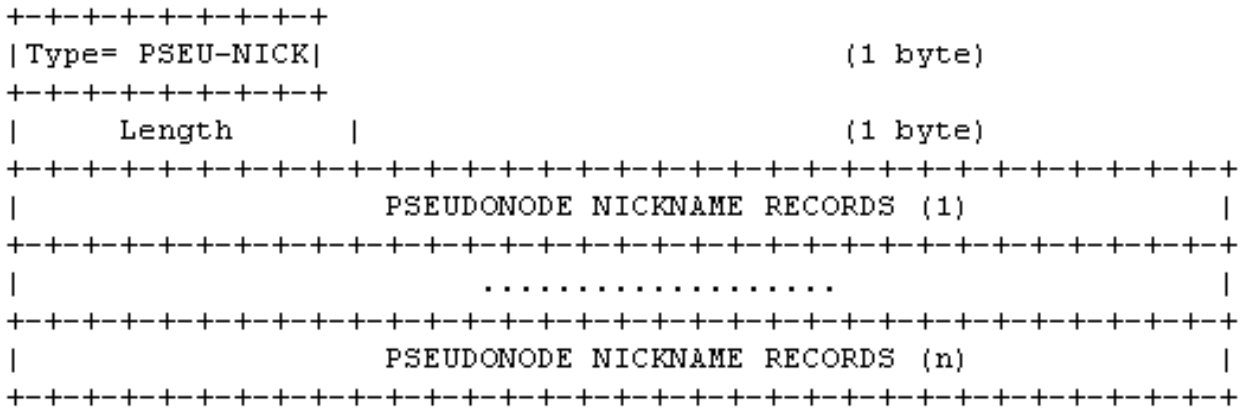

Pseudonode nickname record:

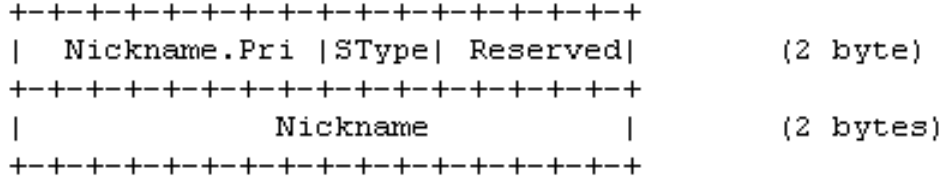

This TLV is carried in DRB's pseudonode LSPs, and the RBridges' Hellos

END

Hongjun Zhai (ZTE)
Fangwei Hu (ZTE)
Radia Perlman (Intel Labs)
Donald Eastlake 3rd (Huawei)