RBRIDGE MULTI-DESTINATION OAM

Yizhou Li
Weiguo Hao
David Michael Bond
Vishwas Manral

Motivations

- Multi-destination frame takes the different path from the unicast frame

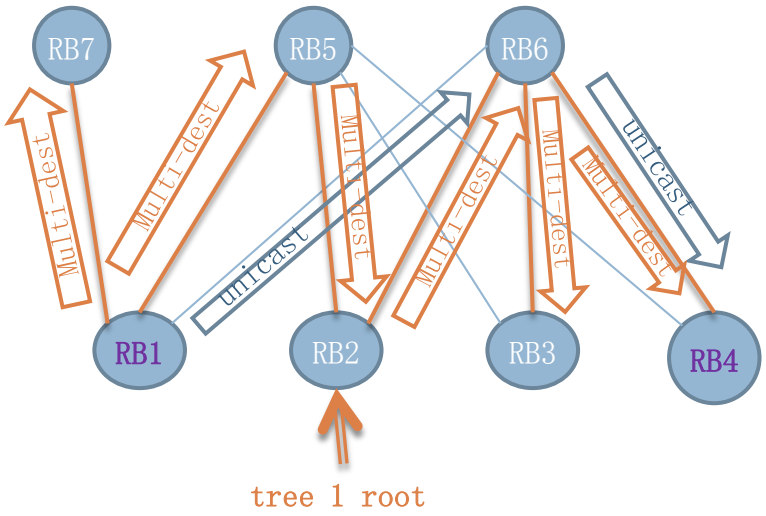

For diagnostic purpose, we may want to know the following:

1. Who are the leaf nodes of a tree in a VLAN? (leaf node: RB advertises its interest of the inner Vlan)
2. Check the connectivity to target(s) in a tree
3. Trace a target in a tree to find the failed hop

- We need OAM messages to support the checking on the multi-destination path in addition to the unicast path

New Channel Protocol

- Use RBridge Channel, define a new OAM channel protocol for Echo in the Long Format

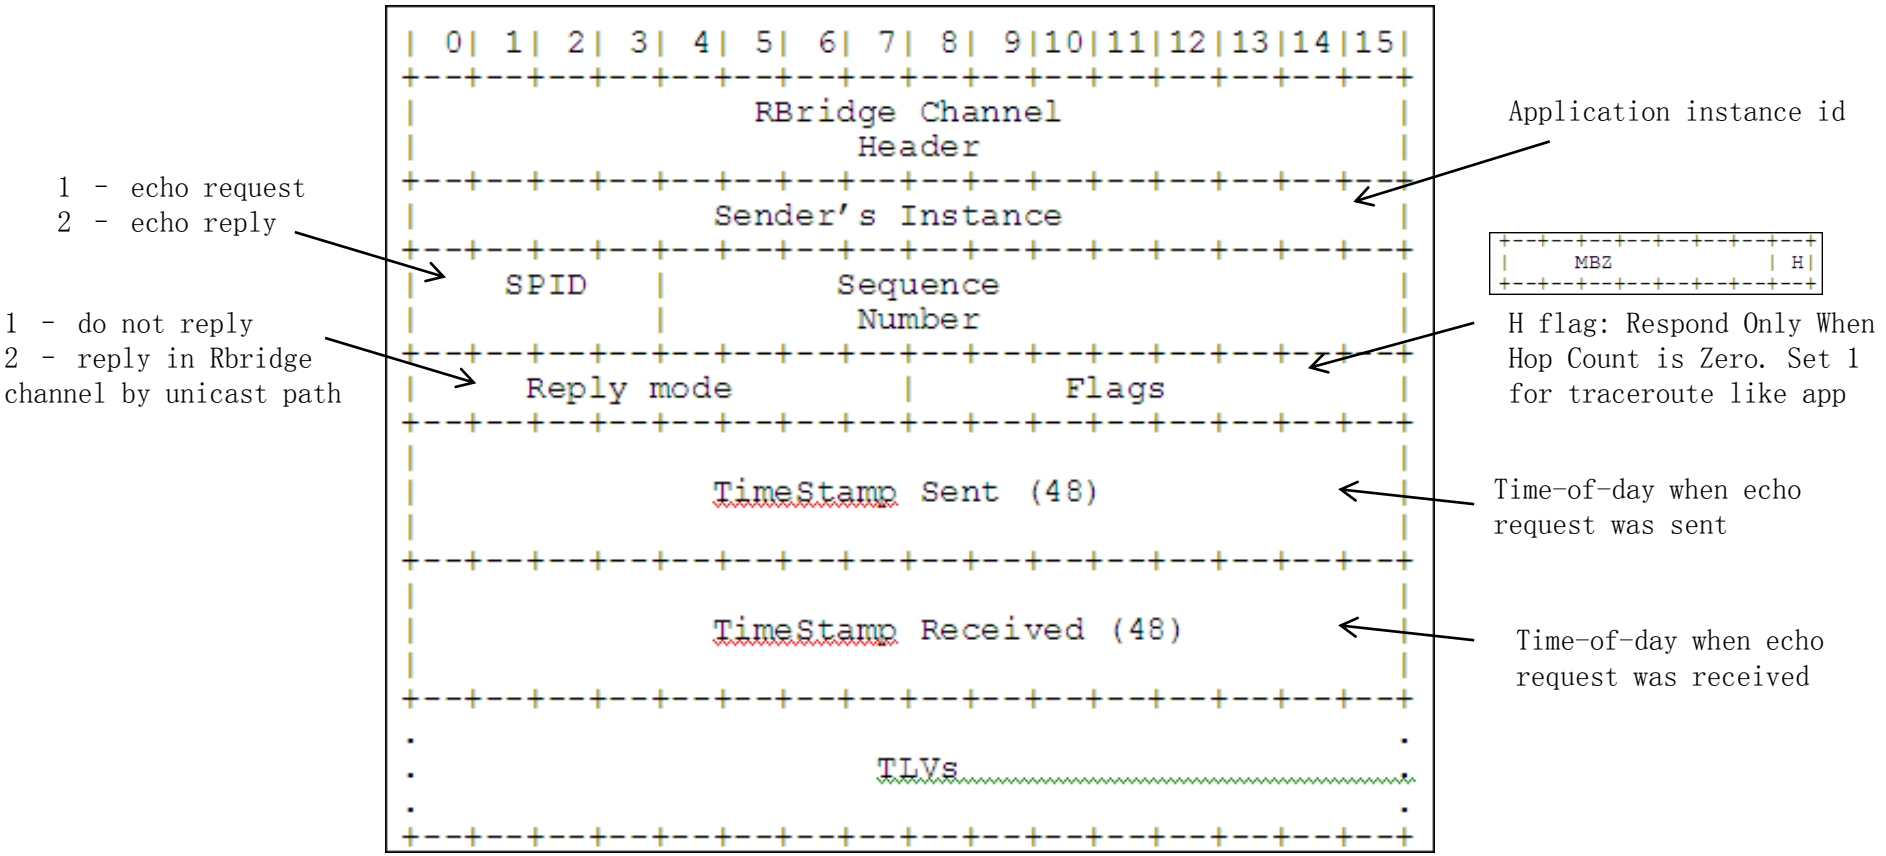

TLV

- Target Rbridge

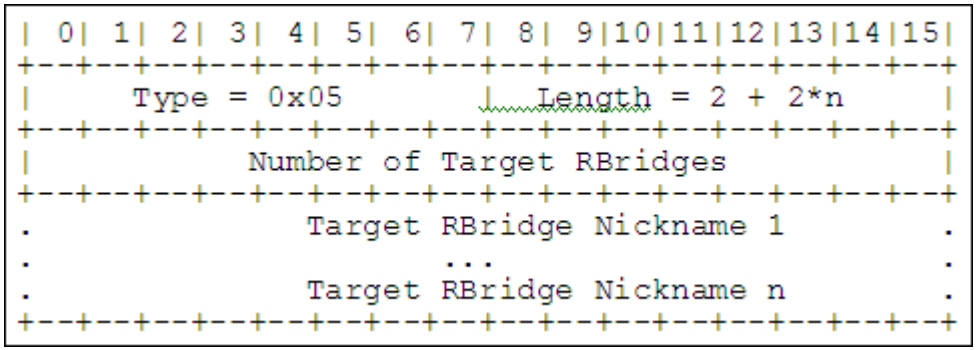

- Jitter

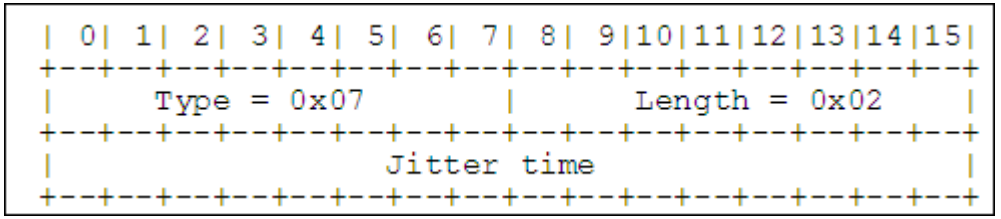

Operations

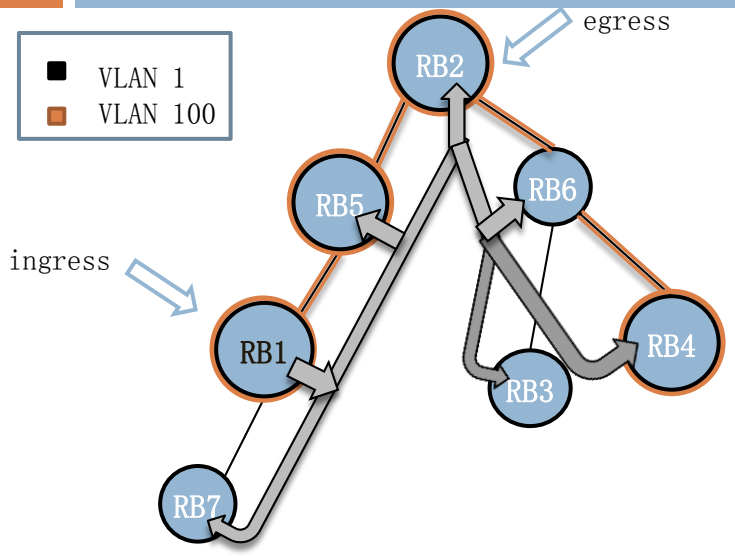

- Sending echo request
  - Inner.MacDA: All-Egress-Rbridges
  - Inner.VLAN: Defaults to 1. It can be any enabled VLAN
  - H(Respond Only When Hop Count is Zero) flag: 1 for traceroute like app, 0 for ping like app
  - Target TLV: optional. When not present, it means unspecified target.
  - Jitter TLV: optional.
- Receiving echo request
  - Replicate the frame to the control plane for processing (ethertype + Dmac), at the same time, do the normal multi-destination data forwarding
  - H flag is not set
    - Unspecified target: leaf node in the VLAN responds echo reply
    - Specified target: leaf node owning one of the targets in the VLAN responds echo reply
  - H flag is set: process only when hop count is 1 in the incoming frame
    - Unspecified target: send back error notification
    - Specified target: send back echo reply if it is the target; send back error notification if it is not the target (error notification can be suppressed if it is not in the path to the target according to LSDB)

Sample Application – ping

- [system] **ping trill-multicast** [**-c** *count* | **-h** *hop-count-value* | **-m** *interval* | **-t** *time-out*] * **root** <*root-nickname*> **inner-vlan** <*innervlan-value*> [**target** *leaf-nickname* *]

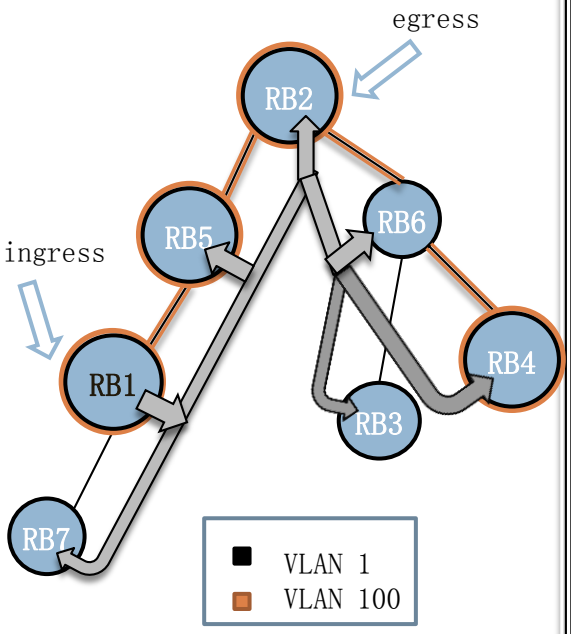


```
[~RB0x1111]ping trill-multicast -c 3  root 0x2222 inner-vlan 100
PING trill-multicast root 0x2222 vlan 100: 20 data bytes, press CTRL_C to
break
    Reply from 0x5555: bytes=20 sequence=1 hc=63 time=2 ms
    Reply from 0x2222: bytes=20 sequence=1 hc=62 time=2 ms
    Reply from 0x4444: bytes=20 sequence=1 hc=60 time=3 ms

    Reply from 0x5555: bytes=20 sequence=2 hc=63 time=2 ms
    Reply from 0x2222: bytes=20 sequence=2 hc=62 time=3 ms
    Reply from 0x4444: bytes=20 sequence=2 hc=60 time=3 ms

    Reply from 0x5555: bytes=20 sequence=3 hc=63 time=2 ms
    Reply from 0x2222: bytes=20 sequence=3 hc=62 time=4 ms
    Reply from 0x4444: bytes=20 sequence=3 hc=60 time=4 ms

--- 0x5555 ping statistics ---
   Packets: Sent = 3, Received = 3, Lost = 0 (0% loss)
   Round-trip min/avg/max = 2/2/2 ms
--- 0x2222 ping statistics ---
   Packets: Sent = 3, Received = 3, Lost = 0 (0% loss)
   Round-trip min/avg/max = 2/3/4 ms
--- 0x4444 ping statistics ---
   Packets: Sent = 3, Received = 3, Lost = 0 (0% loss)
   Round-trip min/avg/max = 3/3/4 ms
```

Sample Application – tracert

- [system]**tracert trill-multicast** [**-h** *hop-count-value* | **-t** *time-out*] * **root** <root-*nickname*> **inner-vlan** <*innervlan-value*> [**target** *leaf-nickname*]

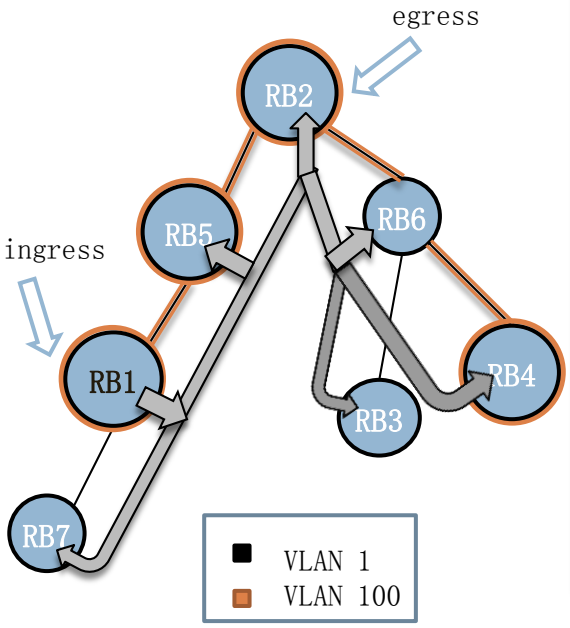


```
[~RB0x1111]tracert trill-multicast root 0x2222 inner-vlan 100 target
0x4444
Hop ReplyRBridge  Time (ms)  InPortId      PreRBNickname
0         0x1111          0    0xFFFF            0x1111
1         0x5555          2    0x0001            0x1111
2         0x2222          2    0x0002            0x5555
3         0x6666          4    0x0001            0x2222
4         0x4444          5    0x0003            0x6666
```

Next step?

- Do we want oam on multi-destination path?
- using a new channel protocol or the same protocol as that for unicast path oam but with diff SPID?
- Pruning:
  - How to achieve CAS(Channel associated signaling) pruning? (Borrowing CAS term here: refering to data-path associated OAM. make sure the OAM messages follow the exact data path, and are pruned in the exact way as real multi-destination data frame)
    - Make Dmac a real multi-destination data MAC, but Smac a special MAC?
    - Not strict CAS: use TLV to carry pruning info for control plain processing?
  - Pruning capability TLV: No prune/VLAN pruned/VLAN+MAC pruned/..?