

Fine-Grained VLAN Labeling

draft-eastlake-trill-rbridge-fine-labeling-01.txt

Radia Perlman

Intel Labs

radiaperlman@gmail.com

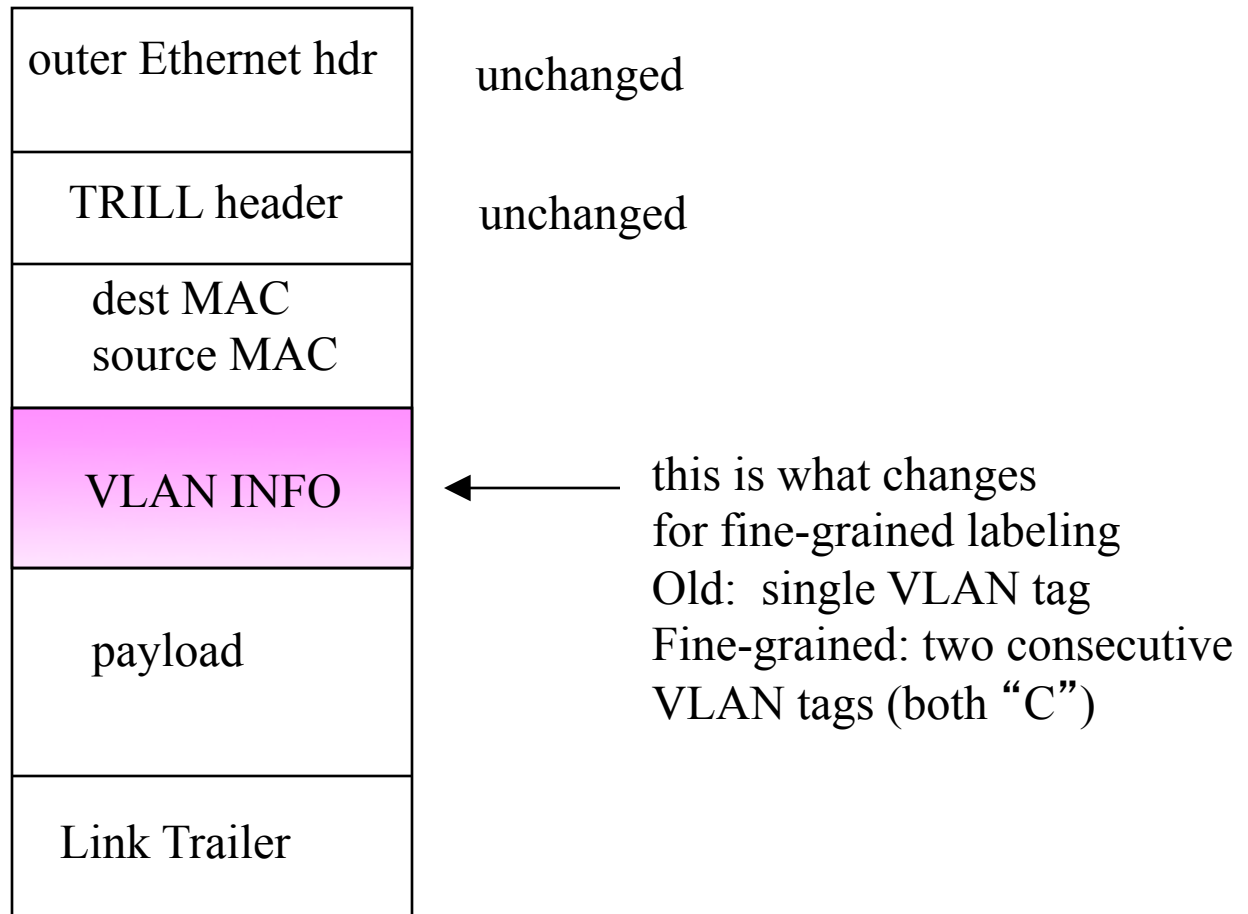
Goals

- More VLANs
- Work with existing silicon
- Allow existing RBridges to work in the core
- Make sure misconfigurations (having old RBridges on the edge) don't do anything bad
- Opportunistic opportunity: 2nd priority field

Basic Idea

- Use two VLANs (in the inner packet), to encode 24 bits of VLAN (16 million)

Double VLAN tag



Notation

- ST=“Single Tag” RBridge, i.e., not fine-grained labeling aware
- DT=“Double Tag” RBridge, i.e., fine-grained labeling aware
- (X.Y) = Fine-grained label with first 12 bits=X, and 2nd 12 bits=Y

Reporting DT-capability

- DT RBridges will report their DT-capability in LSPs
- So therefore DT RBridges will know which RBridges are ST, and which are DT

Reporting Attached fine-grained labels

- New sub-TLV in LSP for appointed forwarder to report attachment to a range of 24-bit labels
- Will be ignored by ST RBridges

Some issues with mixing ST and DT RBridges

- If ST RBridge in core, make sure it doesn't falsely filter multidestination (letting extra traffic through is OK, but blocking traffic because confusion with double tag **not OK**)
- If core ST RBridge exists, then DT R2 attached to X.Y must report attachment to ST VLAN "X"
- To allow some ST edge RBridges, we will allow some VLANs (say first k) to be announced with a single tag, whether ingressed by DT or ST

Illegal configuration

- Illegal if ALL of the following are true
 - Single tag VLAN “X” exists
 - Double tag VLAN “X.Y” exists
 - There is a ST RBridge on the edge reporting it is connected to X
- Detected through LSPs
- Mediated by having DT RBridge that is appointed forwarder for DT VLAN X.Y refrain from ingressing or egressing VLAN X.Y, **and complaining**

ST in core

- The only time core RBridges look at VLAN tag is for filtering multideestination frames
- It's OK for filtering not to be perfect (i.e., only filter on 12 of the 24 VLAN bits – it just wastes a bit of bandwidth, and filtering is a SHOULD anyway, not a MUST)
- But must not falsely filter

Avoiding false filtering with ST R1 in core

- To make sure that R1 doesn't falsely filter (X.Y), all edge RBridges attached to any fine-grained label with the first 12 bits=X must report connectivity to ST VLAN "X" in addition to DT VLAN (X.Y)

Unicast

- Basically no change, except for priority
- Given there are two VLAN tags, there's two priority fields
- Transit RBridges use the priority in the first VLAN tag

Possible Variant

- There might be reasons for R1 to send a multidestination frame as several unicasts instead
 - there might be a very small number of edge RBridges that need to see the frame (e.g., a VLAN that only exists at a few places)
 - As an alternative instead of disabling connectivity to X.Y, when illegal configuration
- It is optional for ingress to send, but mandatory for egress to receive a multidestination frame as a directed unicast