Data Label Based Tree Selection for Multi-destination

draft-yizhou-trill-tree-selection-04

liyizhou@huawei.com

March 2015

Background (1)

- Each distribution tree SHOULD be pruned per VLAN

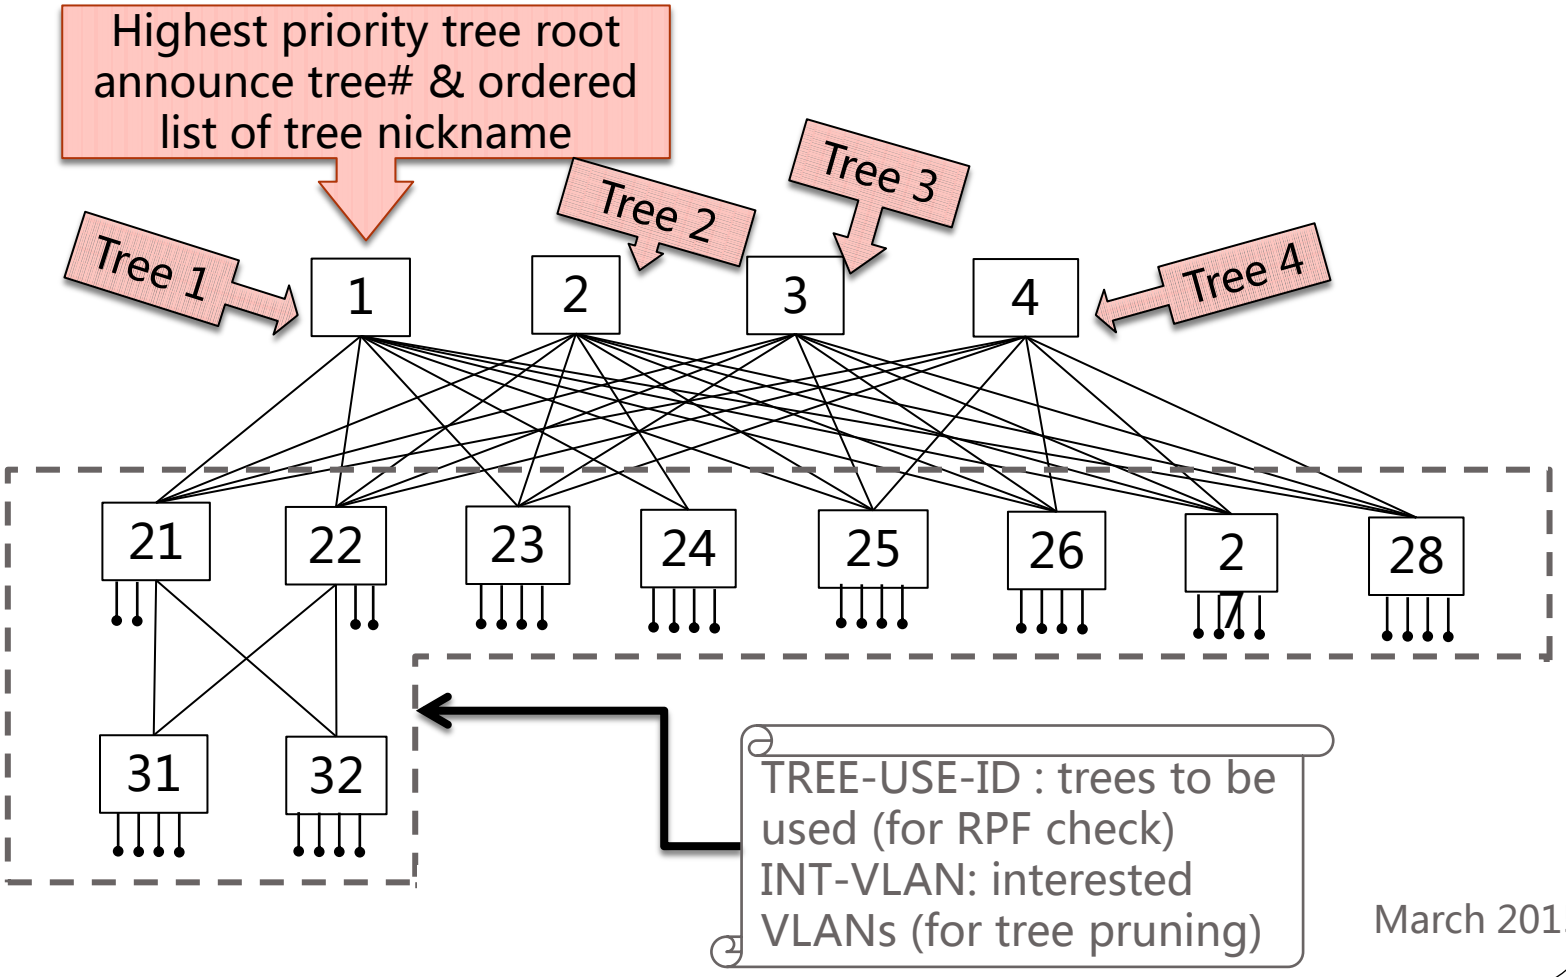

Background（2）

- 4 trees are built

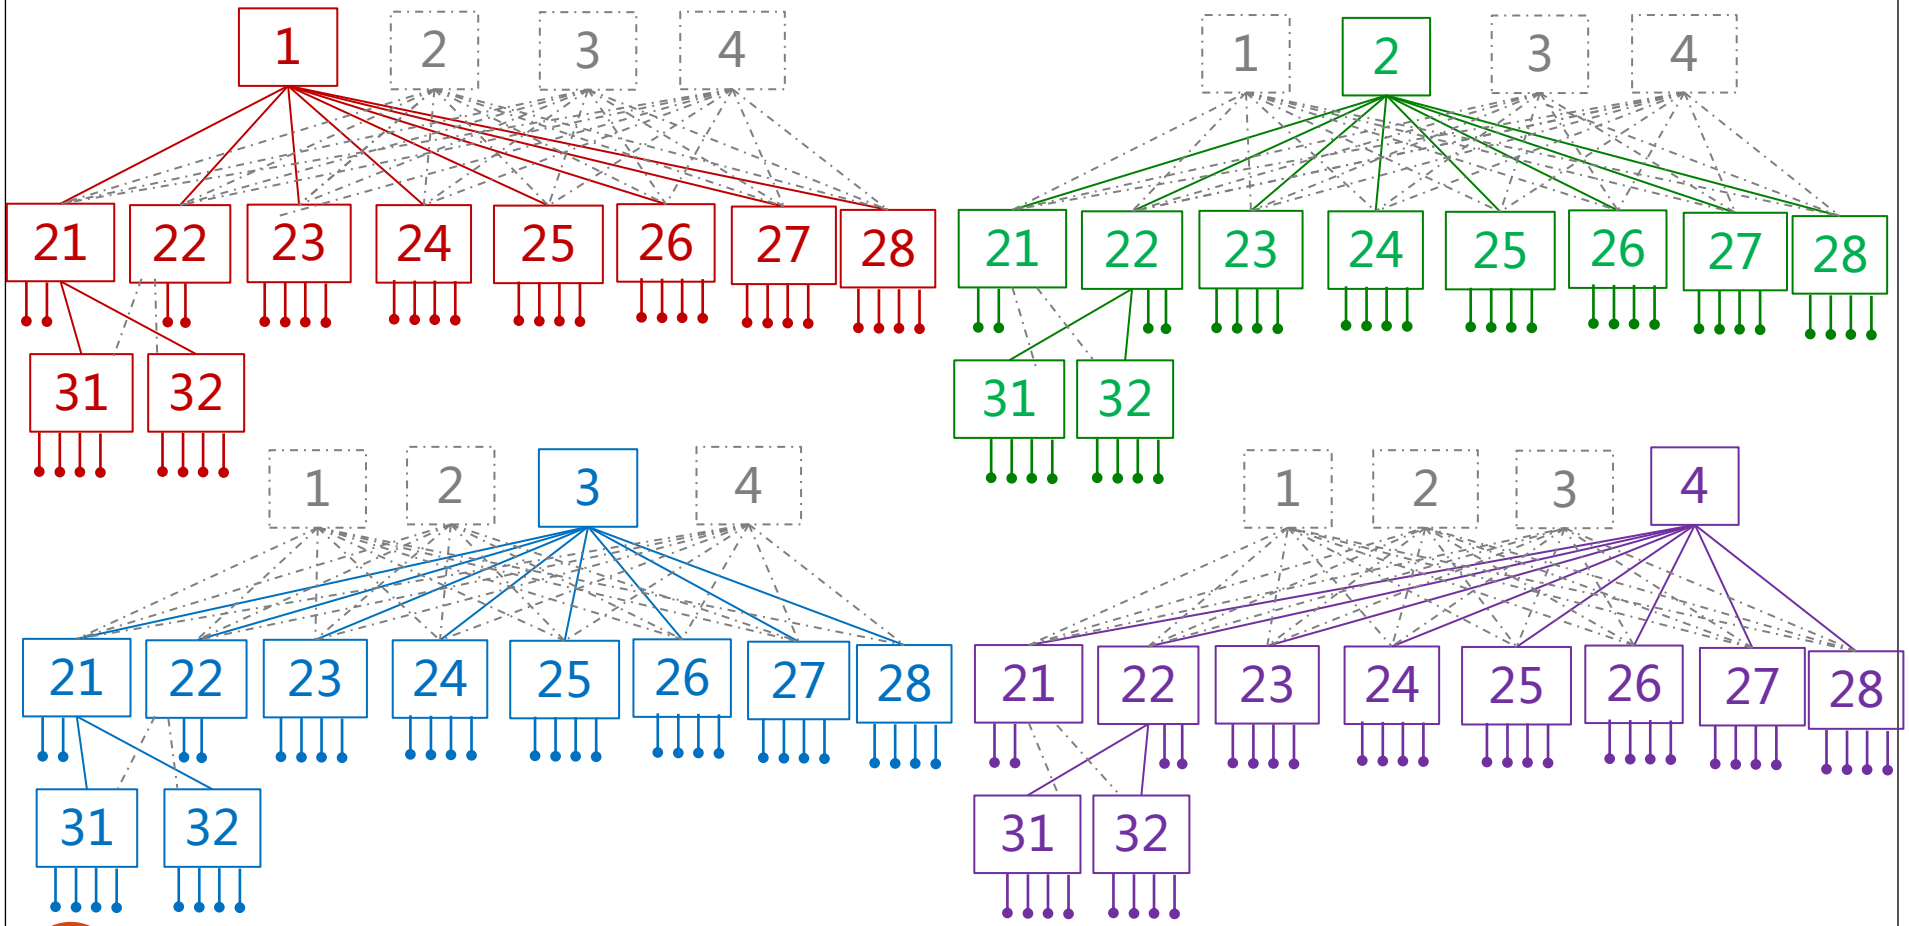

Motivations (1)

- Multicast forwarding table on RB21 has 16K entries.

Tree nickname	VLAN	Port list
tree1	1	1, 31，32, local access ports
tree1	…	…
tree1	4095	1, 31，32, local access ports
tree2	1	2, local access ports
tree2	…	…
tree2	4095	2, local access ports
tree3	1	3, 31，32, local access ports
tree3	…	…
tree3	4095	3, 31，32, local access ports
tree4	1	4, local access ports
tree4	…	…
tree4	4095	4, local access ports

4K entries for tree1

4K entries for tree2

4K entries for tree3

4K entries for tree4

Motivations (2)

- Table size = n*m entries. (n is #of trees, m is #of VLANs with downstream receivers).
- More entries required if L2/L3 multicast address is to be used for further pruning.
- Linearly increasing with # of trees.
- Table size is limited. May share a 8K/16K-entry table with IP multicast/VSI forwarding entries.
- Proposed: Data Label (VLAN or FGL) based tree selection to reduce the table size.
  - still allows the traffic sharing among trees

VLAN based Tree Selection

- Concept:
  - Highest priority tree root announces tree-VLAN correspondence which is the value pair of (tree id, VLANs allowed on this tree id).
  - Ingress RB selects the tree-VLAN correspondence it is interested in and wishes to use from the list.
  - It should not transmit VLAN x frame on tree y if the highest priority tree root does not say VLAN x is allowed on tree y.
  - Achieved VLAN based load balancing by selecting different trees.

VLAN based Tree Selection Example

- If we let the highest priority tree root announces:
  - (tree1, Vlan 1-1000)
  - (tree2, Vlan 1001-2000)
  - (tree3, Vlan 2001-3000)
  - (tree4, Vlan 3001-4095)
- Ingress selects and announces（tree id, interested vlan) from the announced tree-VLAN correspondence
- Multicast table entries are reduced to 4K（maximum).
- Table size shrunk:
  - n * m → m

Tree nickname	VLAN	Port list
tree1	1	
tree1	…	
tree1	1000	
tree2	1001	
tree2	…	
tree2	2000	
tree3	2001	
tree3	…	
tree3	3000	
tree4	3001	
tree4	…	
tree4	4095	

1K entries for tree1

1K entries for tree2

1K entries for tree3

1K entries for tree4

Updates from last revision

- Change from "VLAN" to "Data Label" to include FGL (Fine Grained Label (RFC 7172)) case.
- The (sub-)TLVs are put into E-L1FS (Extended Level 1 Flooding Scope) FS-LSPs [RFC7356].
- Re-use Group Sub-TLVs format as per defined in GADDR TLV [RFC7176] rather than a new format for multicast extension.
- Editorial changes.

Next Step

- Ready for Call for Adoption