

Multi Level/Topology TRILL

Mingui Zhang, Donald Eastlake, Radia Perlman,
Margaret Cullen, Hongjun Zhai

zhangmingui@huawei.com, d3e3e3@gmail.com

MULTI LEVEL/TOPOLOGY TRILL DRAFT STATUS

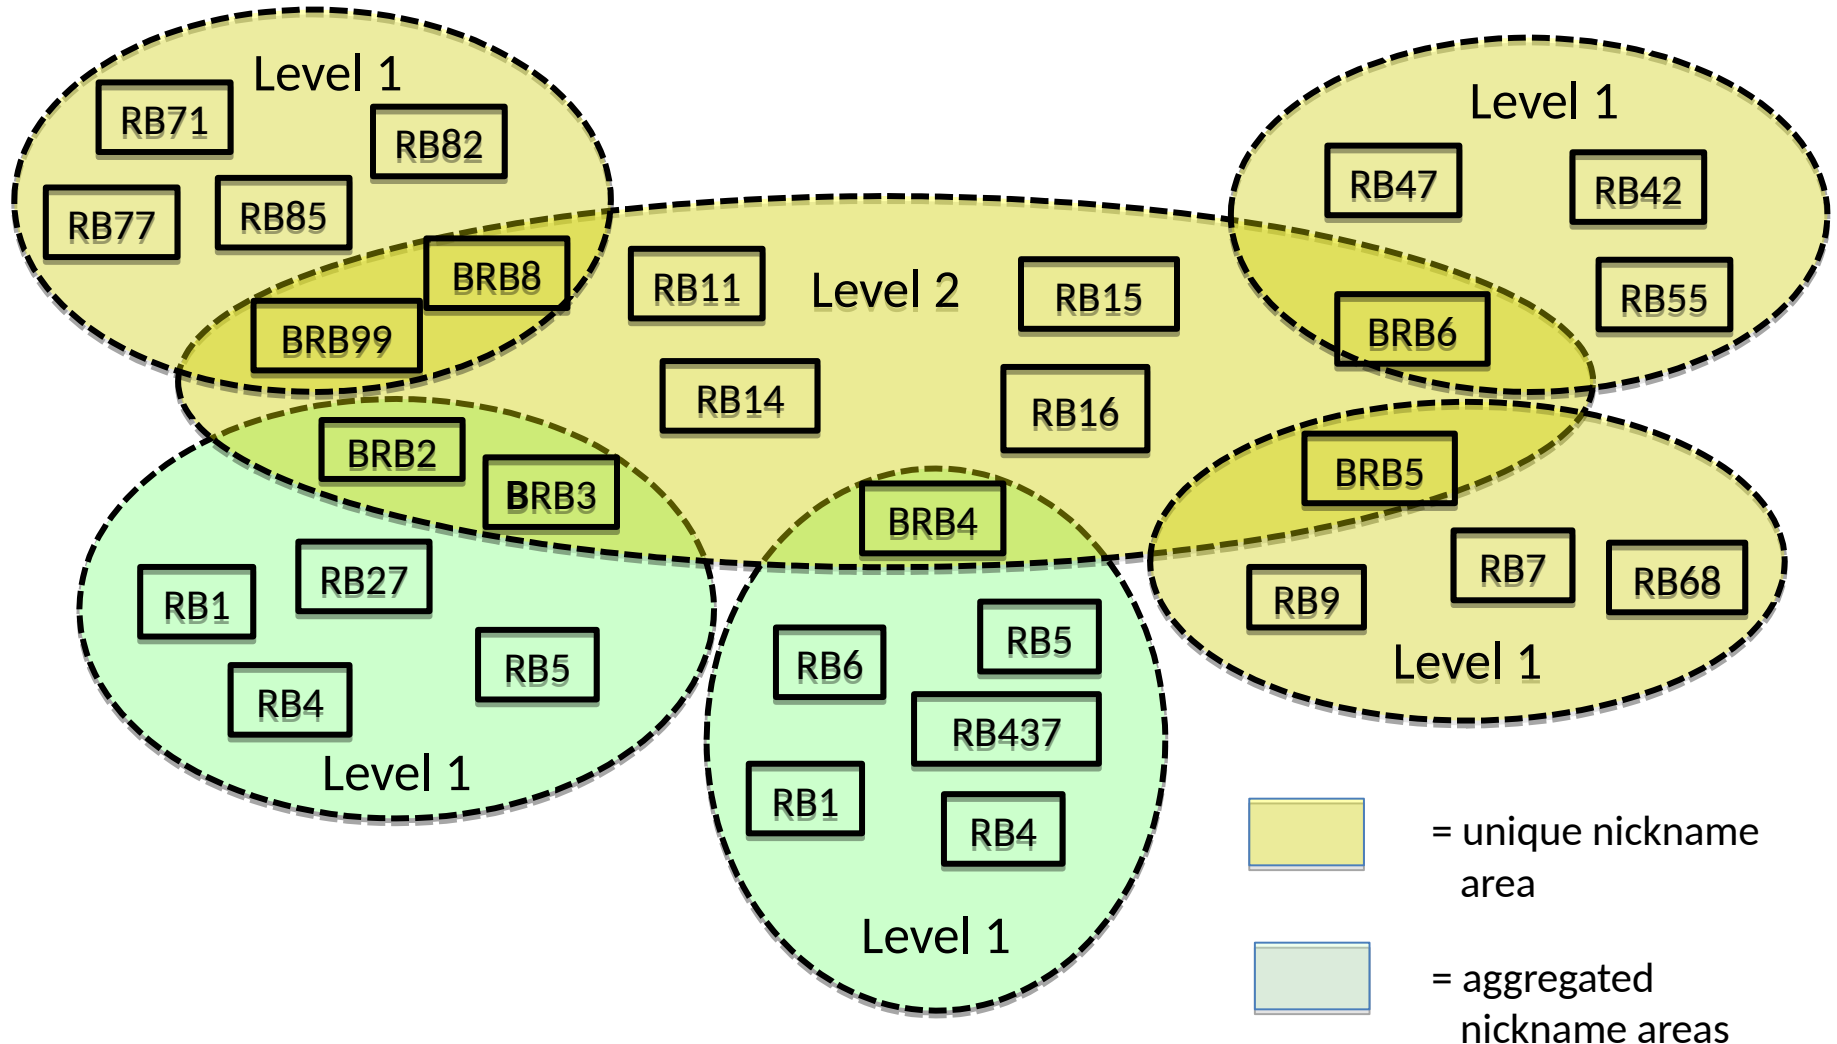
Draft Status

- Past WG Last Call - informational
 - draft-ietf-trill-rbridge-multilevel-01
- WG Drafts
 - draft-ietf-trill-rbridge-multi-topology-01
 - draft-ietf-trill-multilevel-single-nickname-01
- New Personal Draft
 - draft-zhang-trill-multilevel-unique-nickname-00

Multi-Level Drafts

- draft-ietf-trill-rbridge-multilevel-01
 - Informational: Advantages of and options for TRILL support of multi-level (IS-IS) routing.
 - Main option is unique RBridge nicknames across the campus versus aggregated nicknames for Level 1 areas.
 - Unique nicknames simpler for border RBridges. OK for smaller campuses but can exhaust nicknames in a huge campus
 - Aggregated nicknames solve nickname exhaustion but probably requires fast path changes in border Rbridges.
 - This draft recommends supporting hybrid use of both unique and aggregated nicknames in a campus.

Hybrid Multi-level Campus



Standards Track Drafts

- draft-ietf-trill-multilevel-single-nickname-01
 - Proposed Standard: Specifies a method of aggregation with Level 1 areas identified by the set of border RBridges.
- draft-zhang-trill-multilevel-unique-nickname-00
 - Proposed Standard: Specifies a multilevel method with unique nicknames across Level 1 areas and the level 2 area.
- draft-ietf-trill-rbridge-multi-topology-01
 - Proposed Standard. Specifies a method of applying IS-IS multi-topology to TRILL.

Standards Track Drafts


- draft-ietf-trill-multilevel-**single**-nickname-01
 - Proposed Standard: Specifies a method of **aggregation** with Level 1 areas identified by the set of border R Bridges.
- draft-zhang-trill-multilevel-**unique**-nickname-00
 - Proposed Standard: Specifies a multilevel method with **unique** nicknames across Level 1 areas and the level 2 area.
- draft-ietf-trill-rbridge-multi-topology-01
 - Proposed Standard. Specifies a method of applying IS-IS multi-topology to TRILL.

MULTI LEVEL/TOPOLOGY TRILL

DRAFT-ZHANG-TRILL-MULTILEVEL-UNIQUE-NICKNAME

TRILL Multilevel

- The major issue of TRILL multilevel is how to manage RBridge nicknames.
 - Unique nickname approach
 - No nickname duplication across L1/L2.



Nickname announcement


- The border RBridge needs to announce nicknames
 - In L1, as if it owns all nicknames outside this area
 - In L2, as if it owns all nicknames inside the area
- A route to a nickname in another area will be calculated “segment by segment”.

[L1 Segment] RB27->RBx->RBz->RB2

[L2 Segment] RB2->RBb->RBc->RBd->RBe->RB3

[L1 Segment] RB3->RBk->RB44

Unitcast routing



- RB27
 - Learned D is attached to RB44
 - Encapsulates the packet [ingress=27, egress=44]
 - Routes to RB2 since RB2 announced in Area X it also owns RB44
- RB2
 - Routes to RB3 in L2 since RB3 announced in L2 it owns RB44
- RB3
 - Routes to RB44 in Area Y
- RB44
 - Decapsulates and learn S is attached to RB27


Multicast routing

- Local distribution tree
 - With an L1 root nickname, advertised by the border RBridge with the highest root priority.
 - Multicast routing as specified in RFC 6325
 - A multi-destination packet on a local tree **MUST NOT** be leaked into Level 2


- Global distribution tree
 - With an L2 root nickname, also advertised by the border RBridge with the highest root priority.
 - The global distribution is calculated “segment by segment”.

Global distribution tree, different views

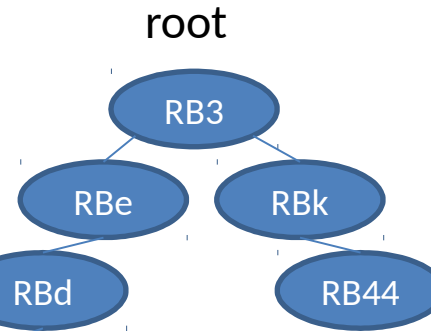
RB27's view



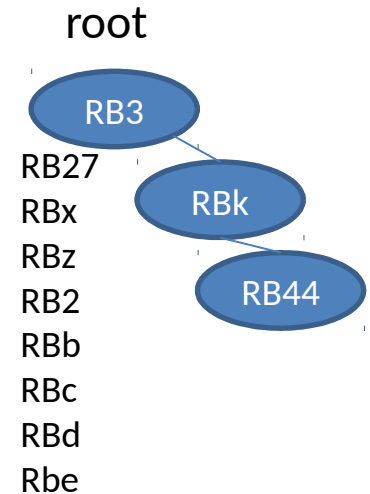
RB2's view



RB3's view



RB44's view

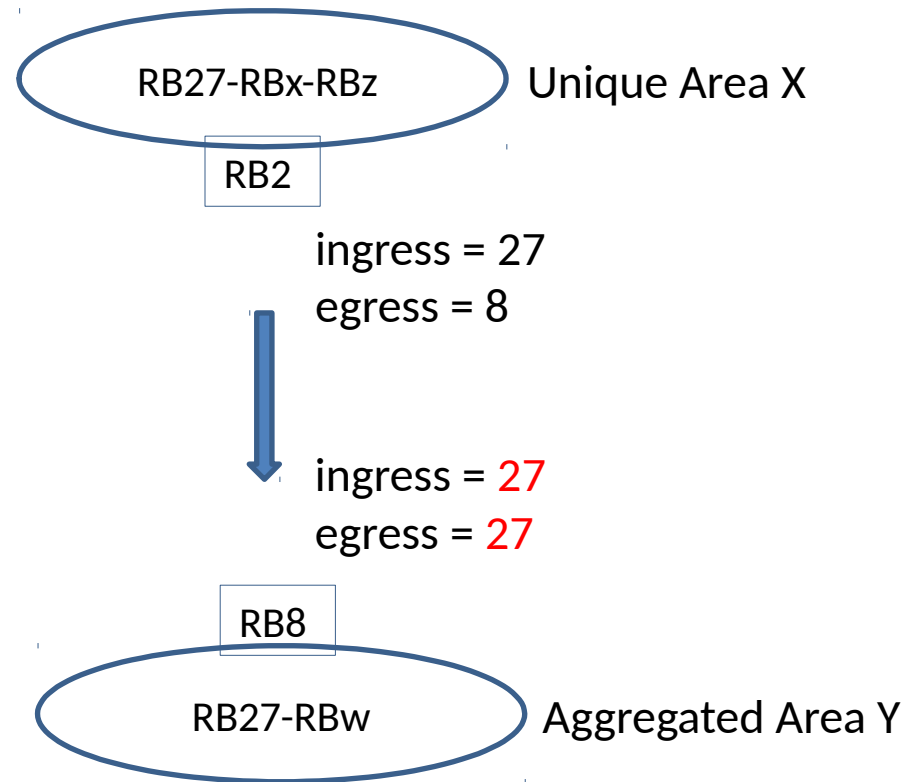


Multicast forwarding

- RB27
 - Produces a multi-destination packet [M = 1, ingress=27, egress=3]
 - Floods the packet using the segment in Area X.
- RB2
 - Floods the packet using the segment in Level 2.
- RB3
 - Floods the packet using the segment in Area Y.
- RB44
 - As a multicast listener, RB44 decapsulates the multi-destination packet.
 - Learns that S is attached to nickname 27.

Mix Unique/Aggregated Areas

- We know
 - Nickname reuse is allowed between aggregated areas.
 - Nickname reuse is NOT allowed between unique areas.
- Question
 - Whether nickname reuse is allowed between unique and aggregated areas?
 - The answer is NO!!



Next steps

- Ask for WG adoption of draft-zhang-trill-multilevel-unique-nicknam
- Incorporate comments we receive.

END

Mingui Zhang, Donald Eastlake, Radia Perlman,
Margaret Cullen, Hongjun Zhai

zhangmingui@huawei.com, d3e3e3@gmail.com