

# RIFT Applicability

## draft-wei-rift-applicability-02

Yuehua Wei, Sandy Zhang , ZTE Corporation

Dmitry Afanasiev , Yandex

Tom Verhaeg , Interconnect Services B.V.

Jaroslawn Kowalczyk , Orange

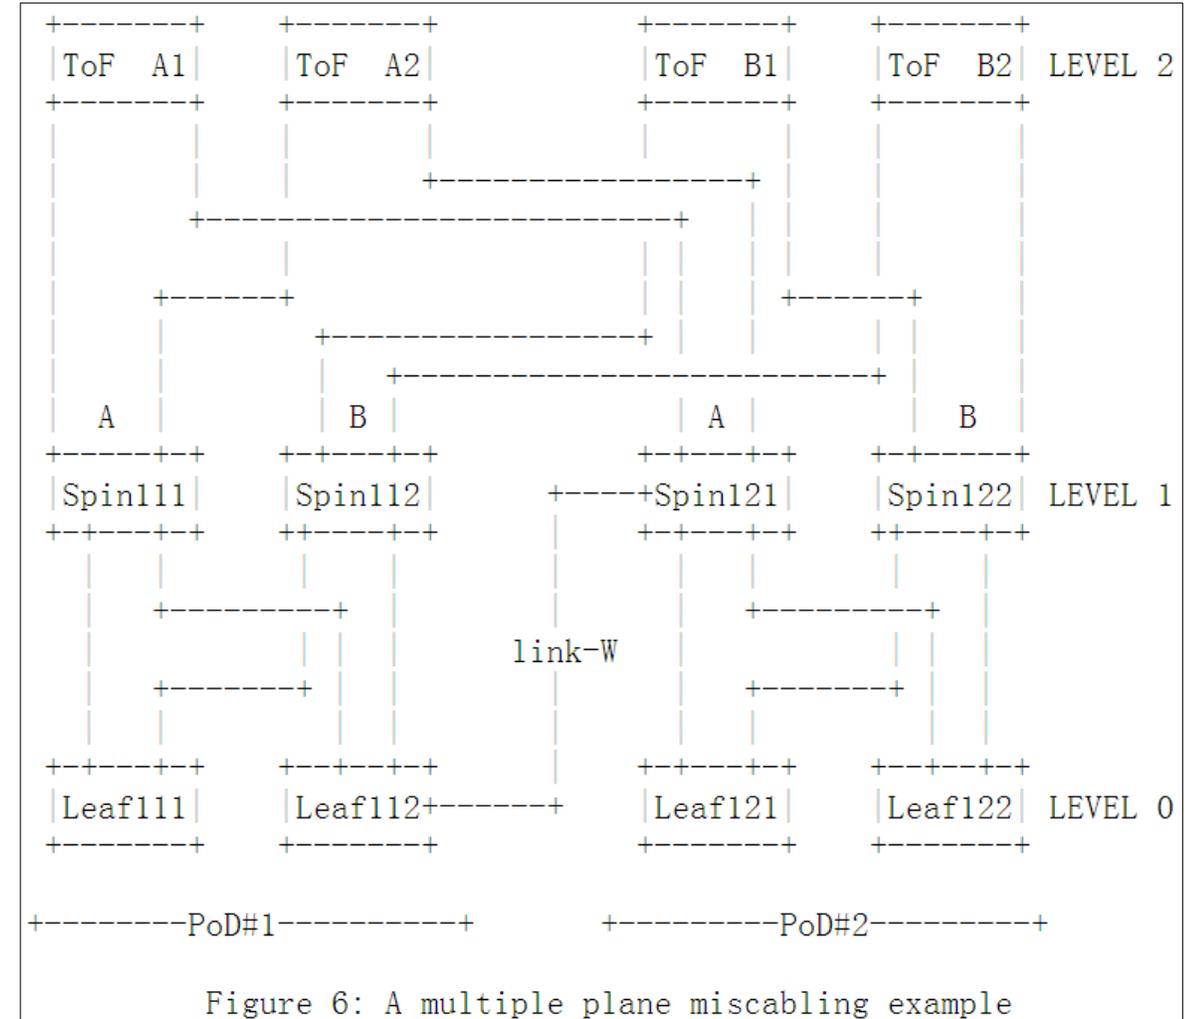
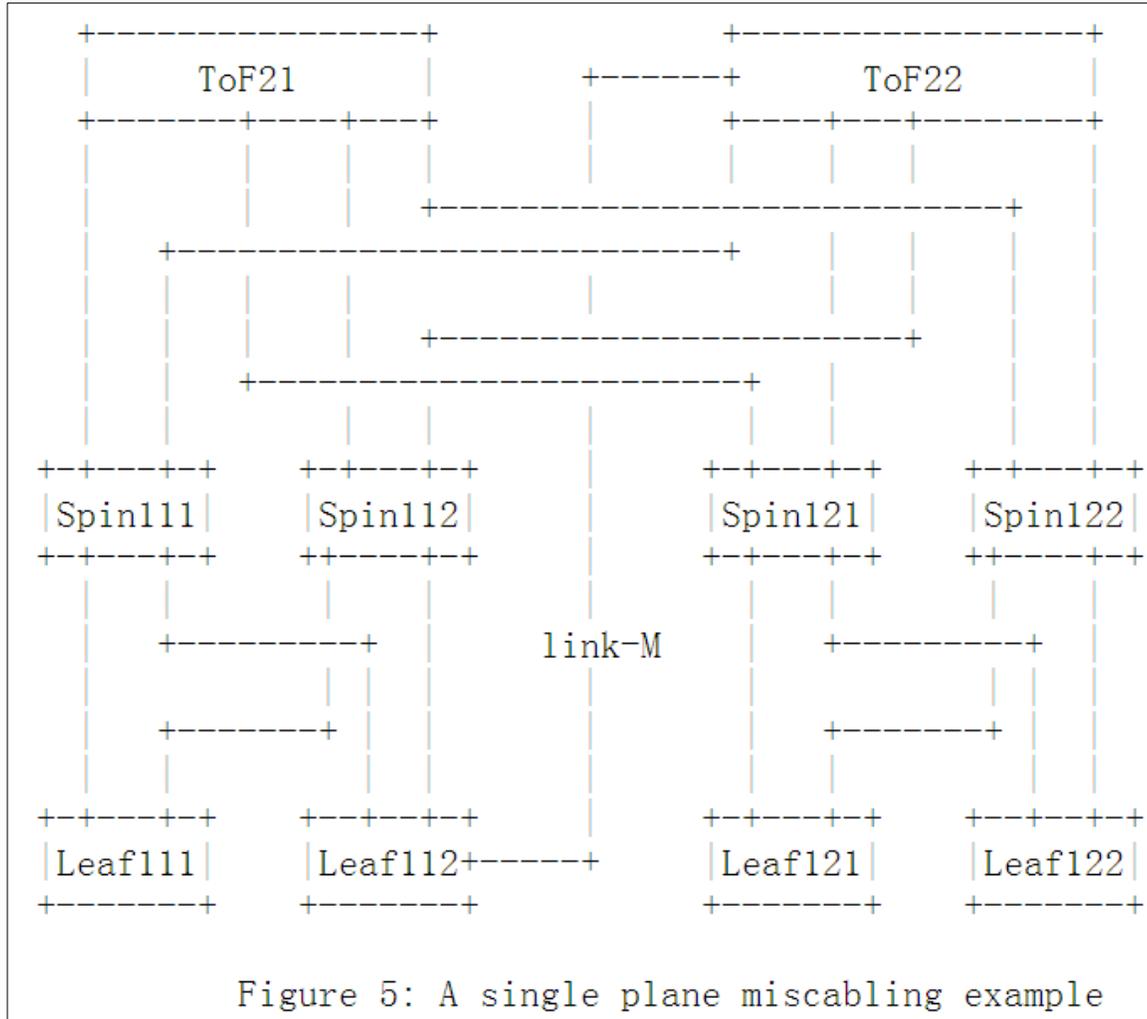
# Agenda

- Updates (since IETF-105)
- Next steps

# New paragraphs added to “4. Deployment Considerations

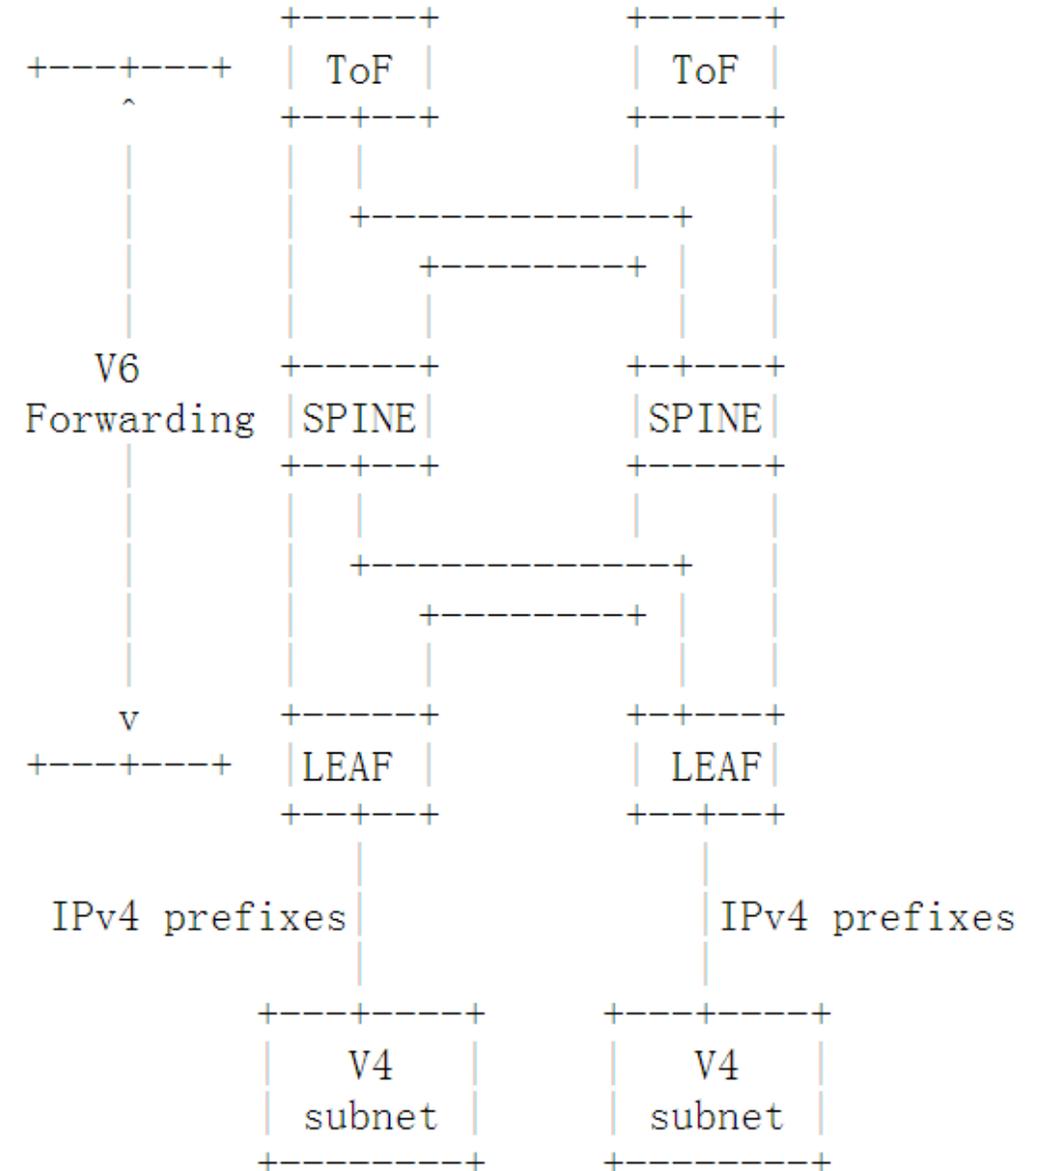
- Miscabling Examples
- IPv4 over IPv6
- Dual Homing Servers
- Fabric With A Controller
- Subnet Mismatch and Address Families
- Anycast Consideration

# Miscabling Examples

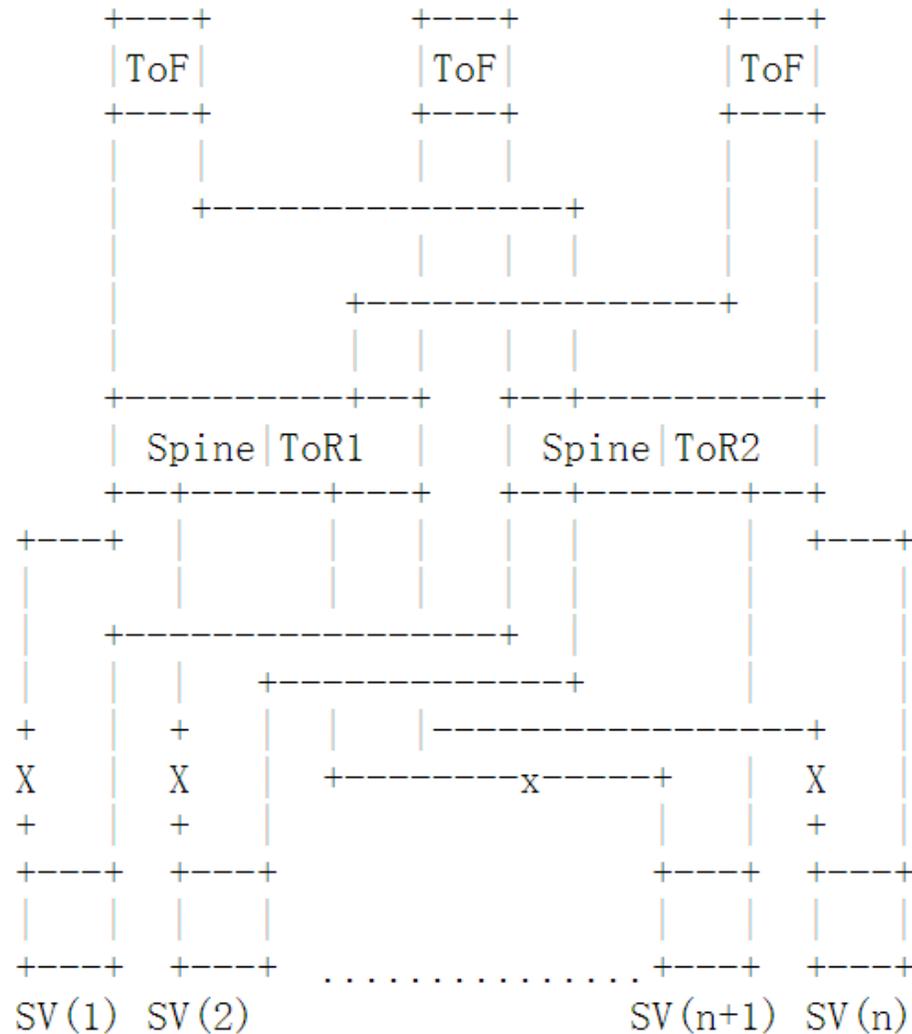


# IPv4 over IPv6

- RIFT allows advertising IPv4 prefixes over IPv6 RIFT network.
- It is expected that the whole fabric supports the same type of forwarding of address families on all the links.
- RIFT provides an indication whether a node is v4 forwarding capable and implementations are possible where different routing tables are computed per address family as long as the computation remains loop-free.

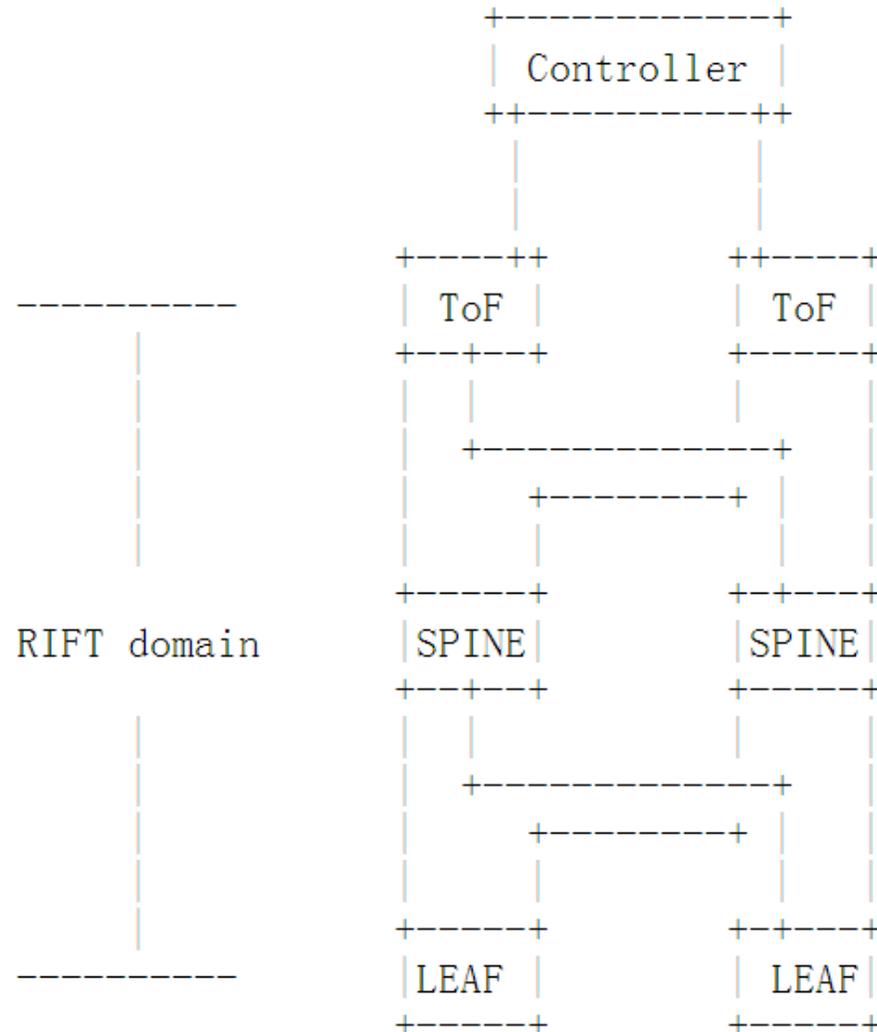


# Dual Homing Servers



- In the single plane, the worst condition is disaggregation of every other servers at the same level. All the servers' routes are disaggregated and the FIB of the servers will be expanded with  $n-1$  more specific routes.
- Support disaggregation from ToR to servers from start on

# Fabric With A Controller



- **Controller Attached to ToFs**

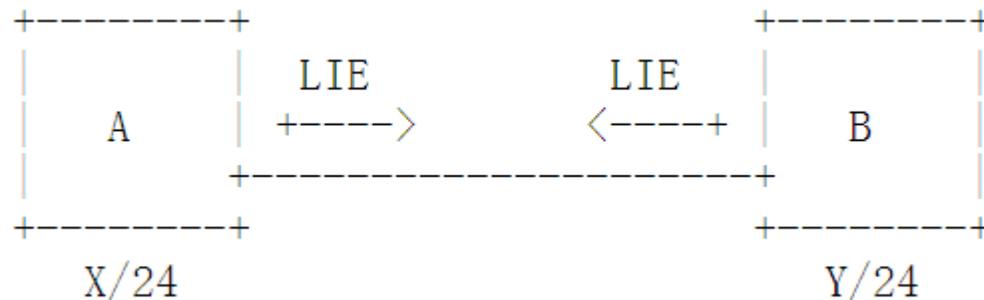
- Usually uses dual-homing connections. The loopback prefix of the controller should be advertised down by the ToF and spine to leaves.
- If the controller loses link to ToF, make sure the ToF withdraw the prefix of the controller

- **Controller Attached to Leaf**

- If the controller is attaching from a leaf to the fabric, no special provisions are needed.

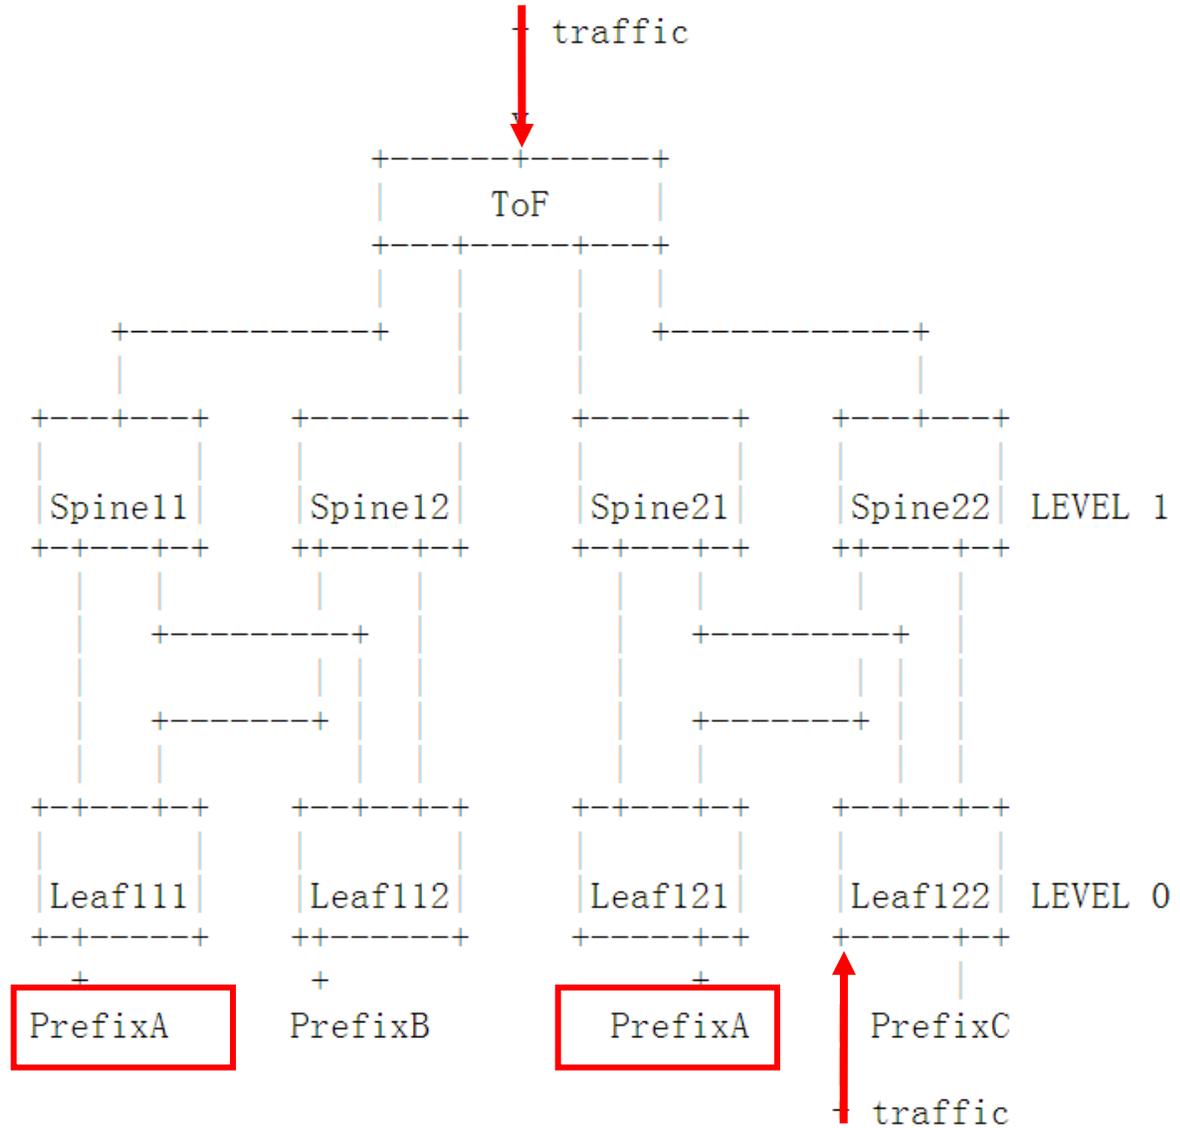
# Subnet Mismatch and Address Families

- Adjacency of node A and B may form
- Forwarding between node A and node B may fail because subnet X mismatches with subnet Y
- To prevent this a RIFT implementation should check for subnet mismatch just like e.g. ISIS does



# Anycast Consideration

- If the traffic comes from ToF to Leaf111 or Leaf121 which has anycast prefix PrefixA. RIFT can deal with this case well.
- But if the traffic comes from Leaf122, it will always get to Leaf121 and never get to Leaf111. If the intention is that the traffic should be offloaded to Leaf111, then use policy guided prefixes [PGP reference].



# Next steps

- Continue to seek comments
- Seek WG adoption

# Thank you!