

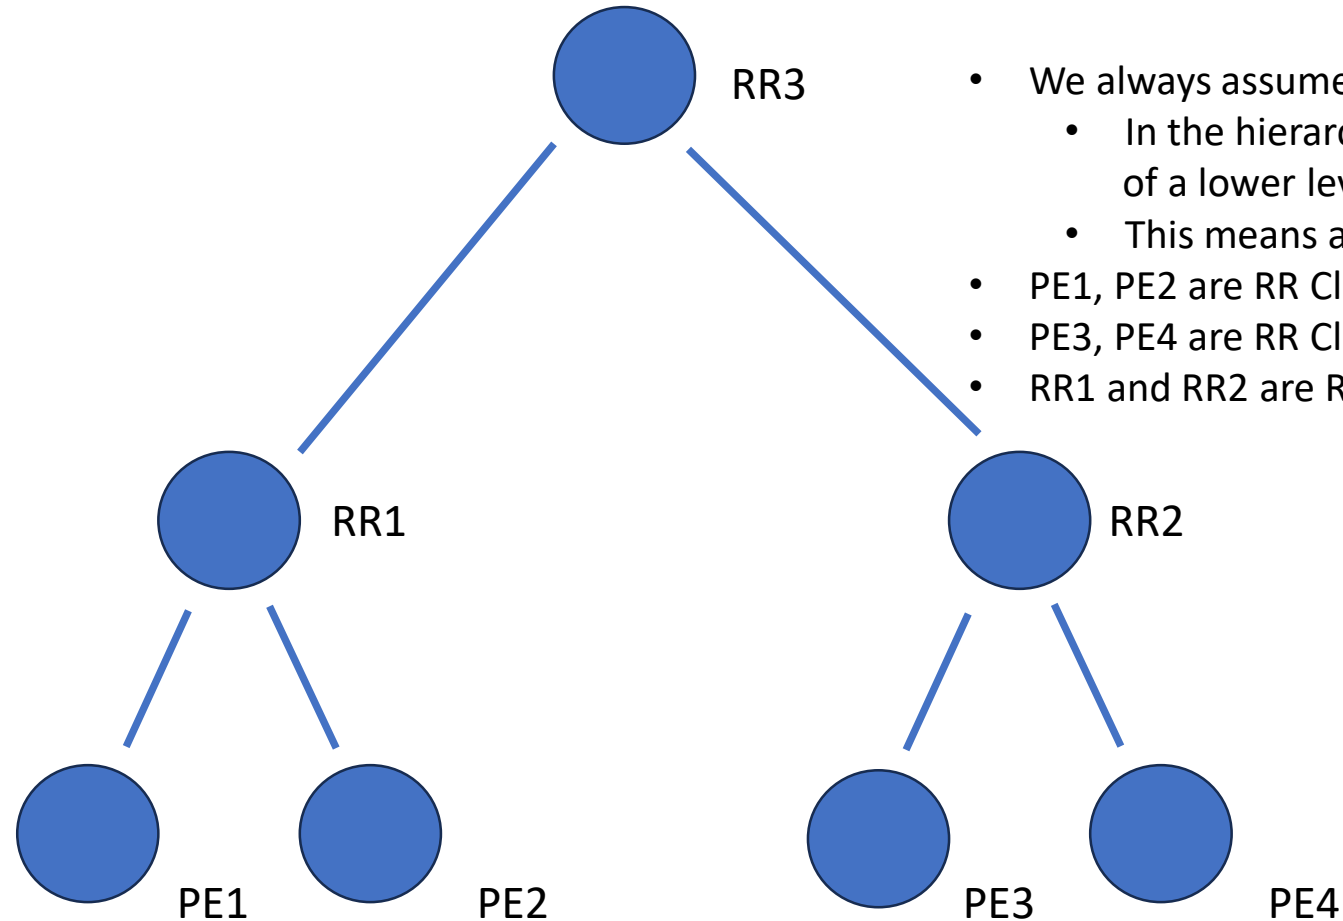
A Solution to the Hierarchical Route Reflection Issue in RT Constraints

Satya R Mohanty

Juan Alcaide

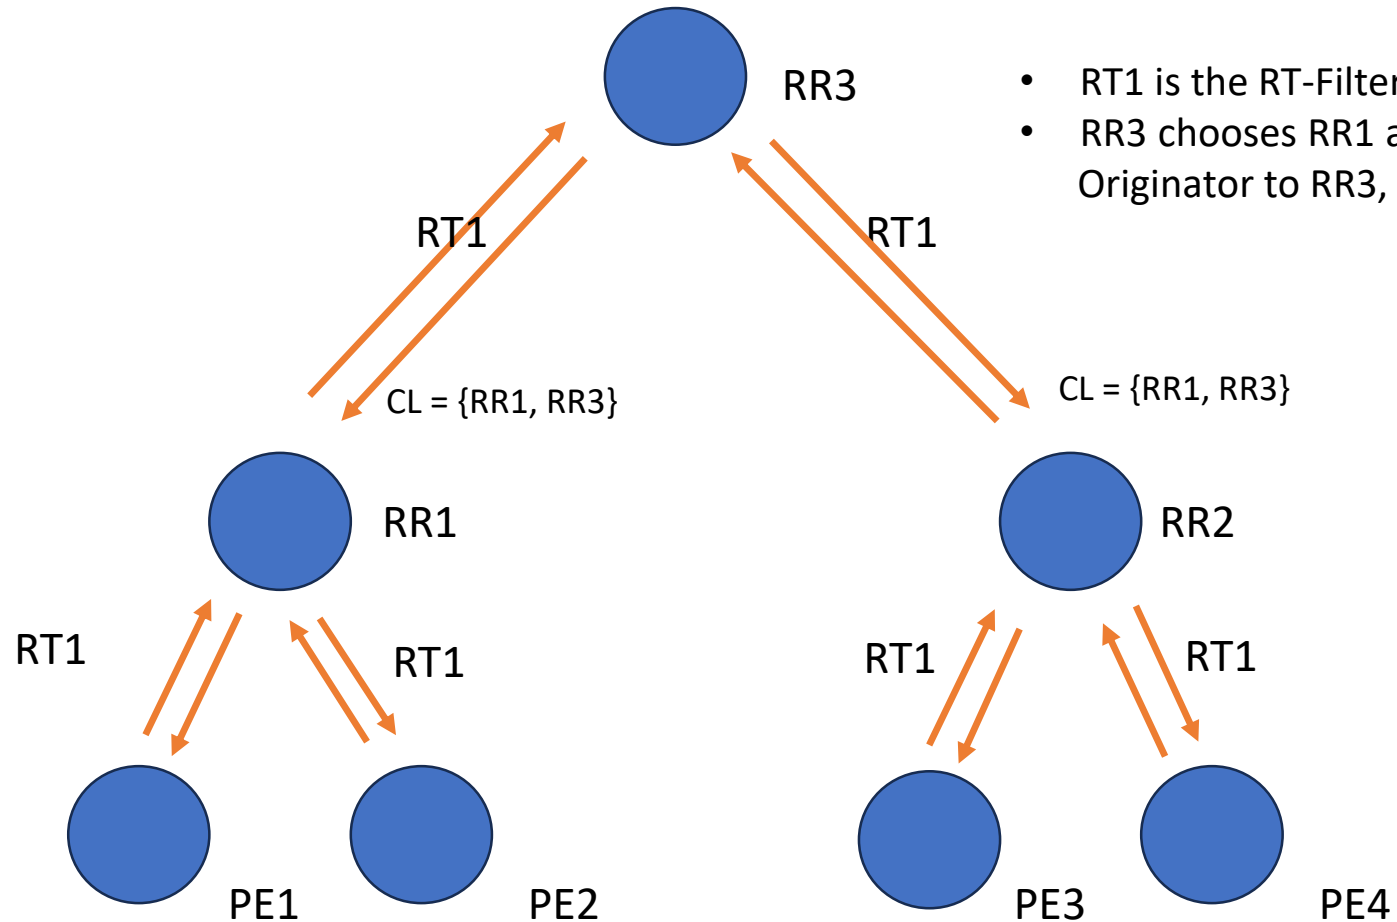
Mrinmoy Ghosh

Definition: RR Canonical Network



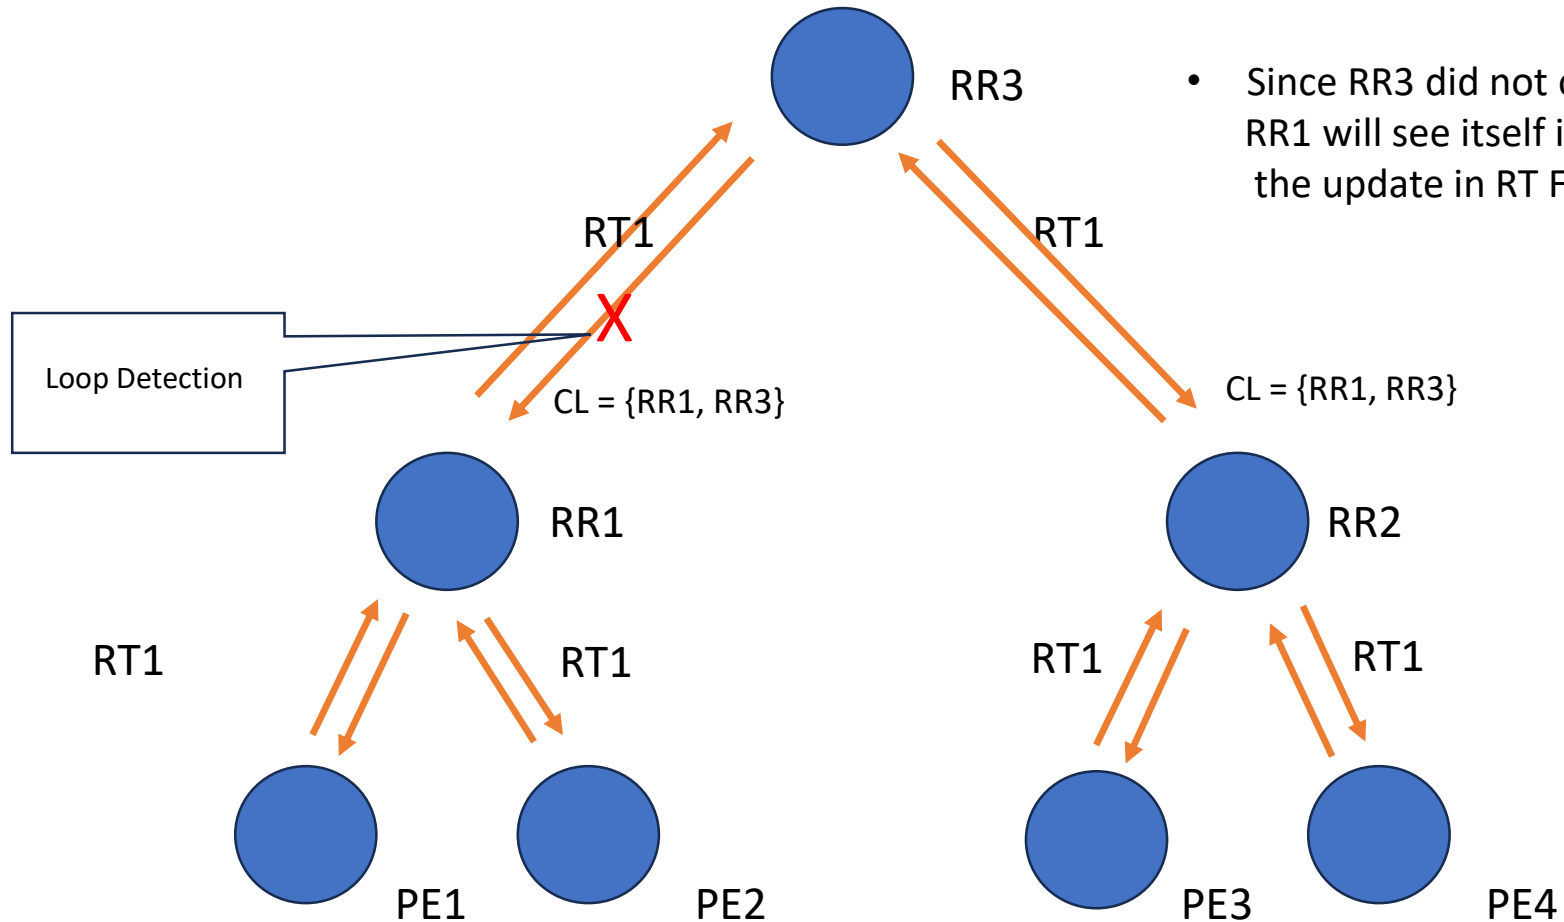
- We always assume a RR Canonical Network
 - In the hierarchy, we have as client only speaker of a lower level
 - This means a reflected Route never comes back
- PE1, PE2 are RR Clients of RR1
- PE3, PE4 are RR Clients of RR2
- RR1 and RR2 are RR Clients of RR3

Problem Description




- RT1 is the RT-Filter Route.
- RR3 chooses RR1 as best, rewrites the NH and Originator to RR3, and advertises it to RR1 and RR2

Problem Description



- Since RR3 did not overwrite the cluster, RR1 will see itself in the cluster-list and will drop the update in RT Filter Address-Family

Propagation of VPN Routes



- RR1 will not advertise VPN routes to RR3 as the RT-Filter, RT1, was rejected
- RR1 will not advertise its client's routes to RR3
- RR2 and its clients do not see RD1:n1

Previous Work

Solutions presented in [I-D.ietf-idr-rtc-hierarchical-rr] are based on

- Addpath, RR-1 will advertise both paths from RR-2 and RR-3 to RR-2 and RR-3 so that each first level RR will accept at least one and install filter
- When RR-1 will advertise the best-path to a client or non-client speaker, and that speaker is the one whose path is the best, the advertising router will use the most "diverse" path (different next-hop and ORIGINATOR_ID than the best-path) to accomplish the same goal, i.e. the path will be accepted at the receiving speaker.

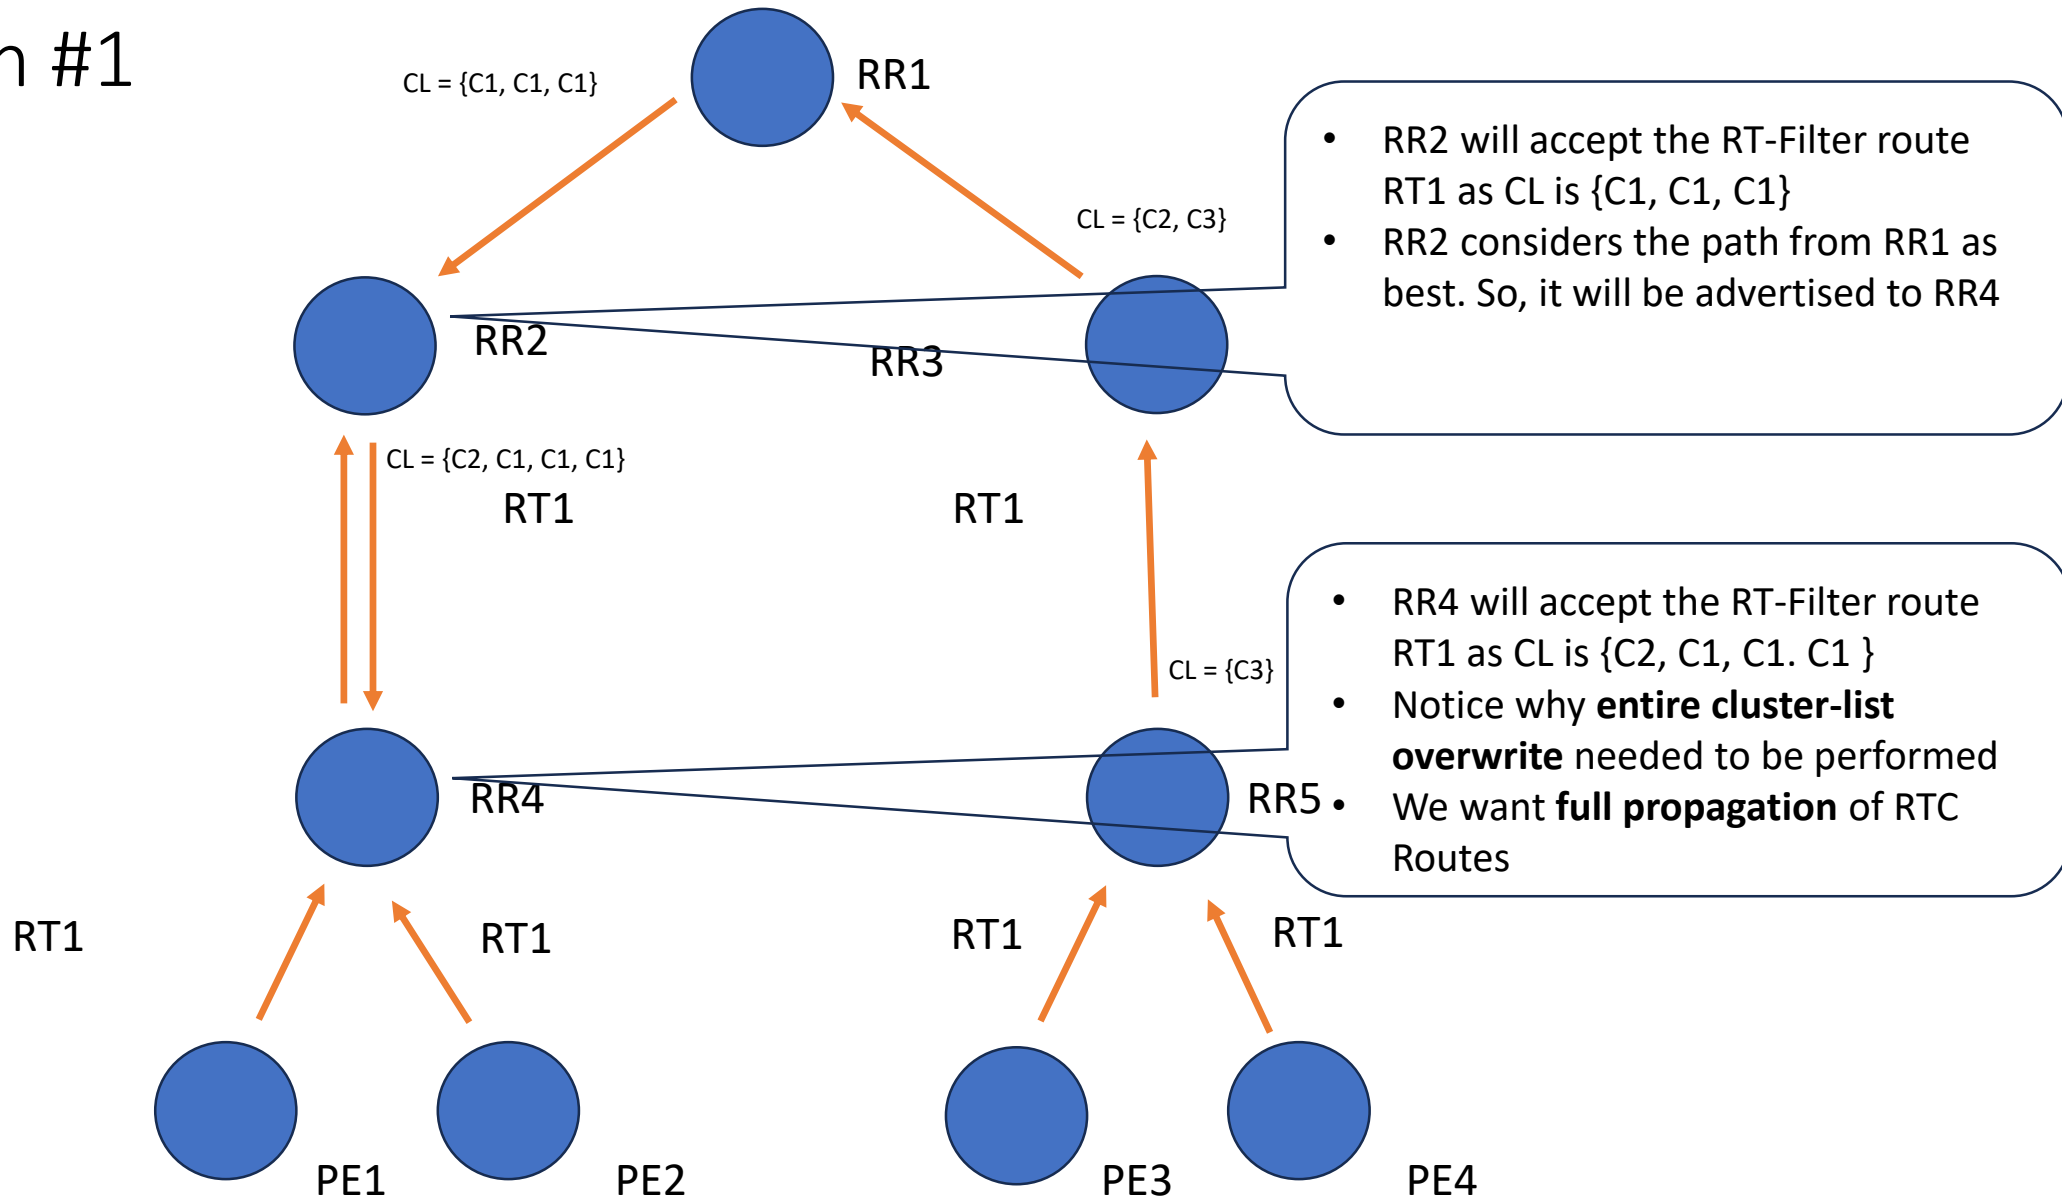
Previous Work (Disadvantages)

- One of the problems of solution 1 are a higher management burden (higher level RR need to be identified, add-paths need to be configured) and therefore an increased number of paths to be advertised. The decision on what paths to be advertised also increases the management burden
- For solution 2, a measure of how disjoint are the paths is not well defined. But suffers of the same problems than solution 1. In addition, the new requirement is sending a different update for every client. This effectively breaks the shared peer update-formatting implementation than most vendors use.

Hierarchical RR Rule 1 (for senders)

- When reflecting the (best-path) RTC route from RR client to RR client, the following rules will apply:
 - When RTC route has CLUSTER_LIST, overwrite all CLUSTER_ID of CLUSTER_LIST to local CLUSTER_ID. Note that when advertising that RTC route, the local CLUSTER_LIST will still be prepending per usual RR rules.
 - ORIGINATOR_ID is set or overwritten with local router-id.
 - NEXT_HOP is overwritten with local peering address (next-hop-self).

Solution #1



Final Hierarchical RR Rule 1 (for senders)

The following recommended (NEXT_HOP-IGNORE) rules can be implemented:

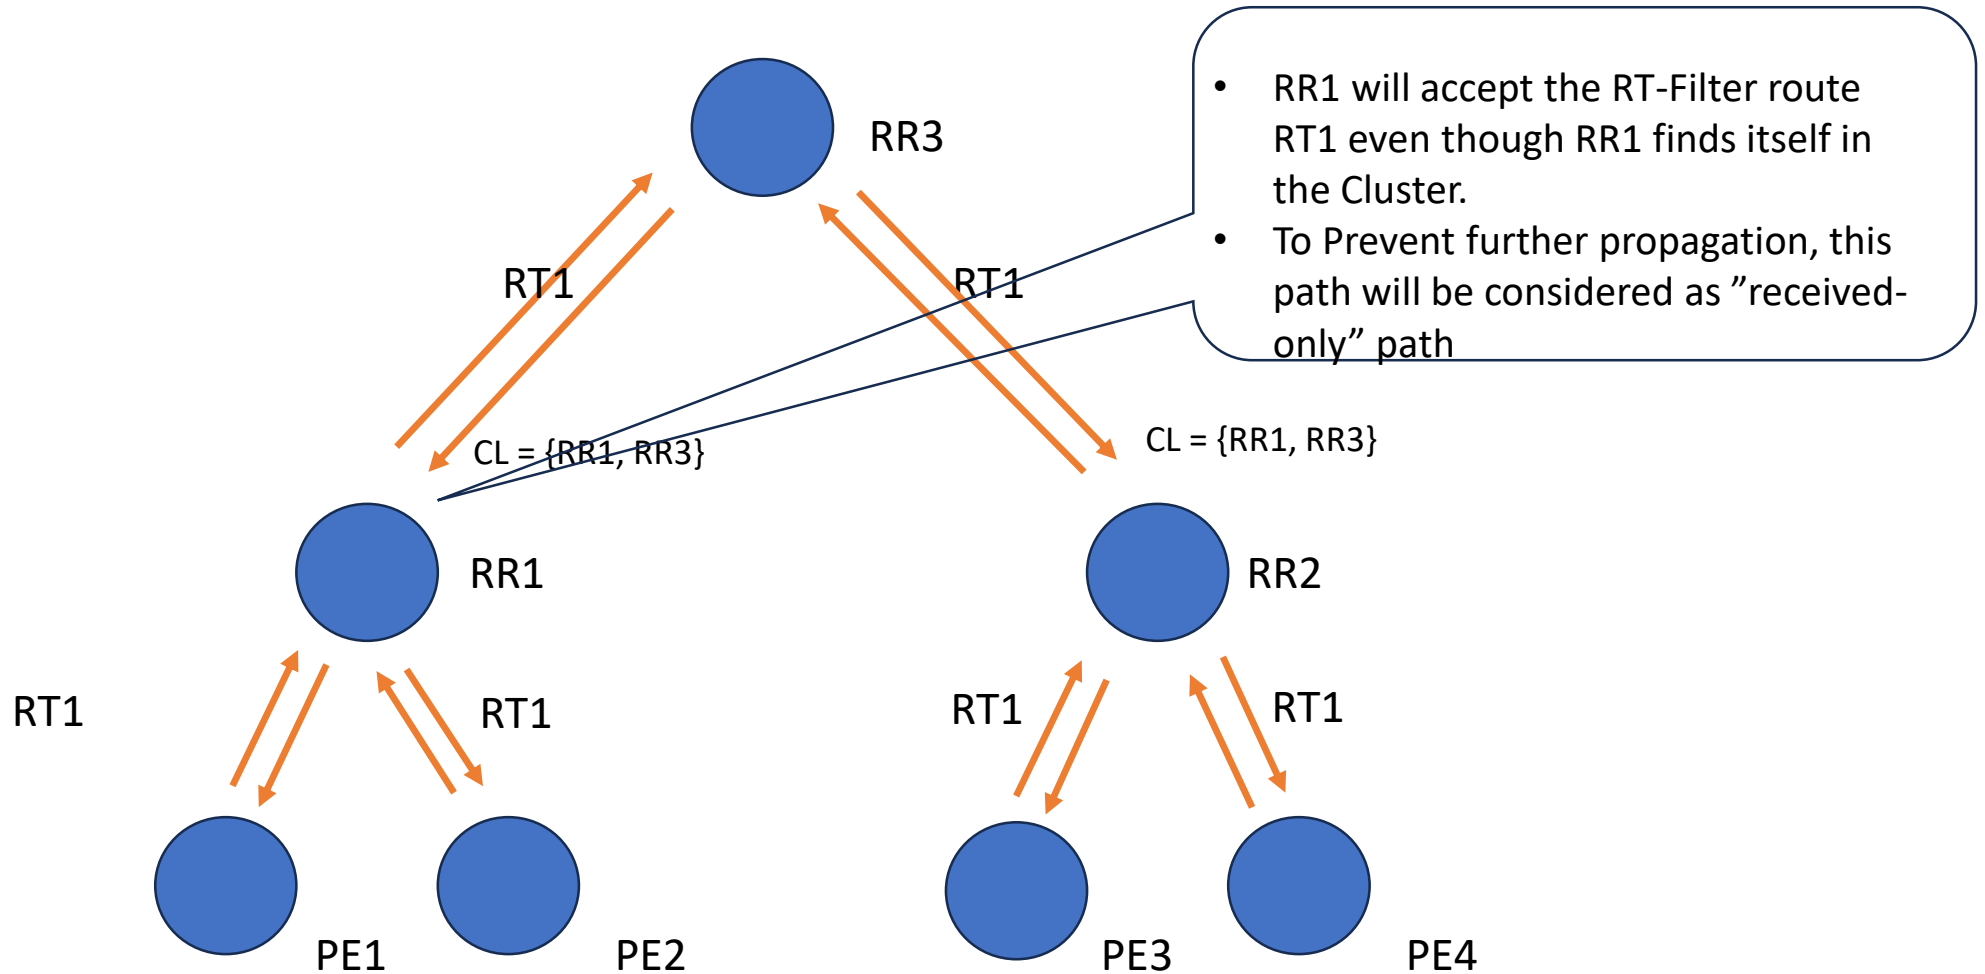
- When reflecting a RTC route, NEXT_HOP overwrite is disabled.
- When receiving A RTC route, it is not discarded even if the received NEXT_HOP is one of the IP addresses of the speaker.
- Change with respect to RFC4684 even for a single level of RR. Note that disabling NEXT_HOP check doesn't create any more loop conditions in a canonical network.
- With these rules the selection of best-path RTC route is now determined by the IGP cost to the original next-hop, and not the IGP costs of each individual peering address.
 - Best-path selection and propagation of RTC routes will be more stable.

Hierarchical RR Rule 2 (for Receivers)

This rule is to be used by the receiver of RTC routes.

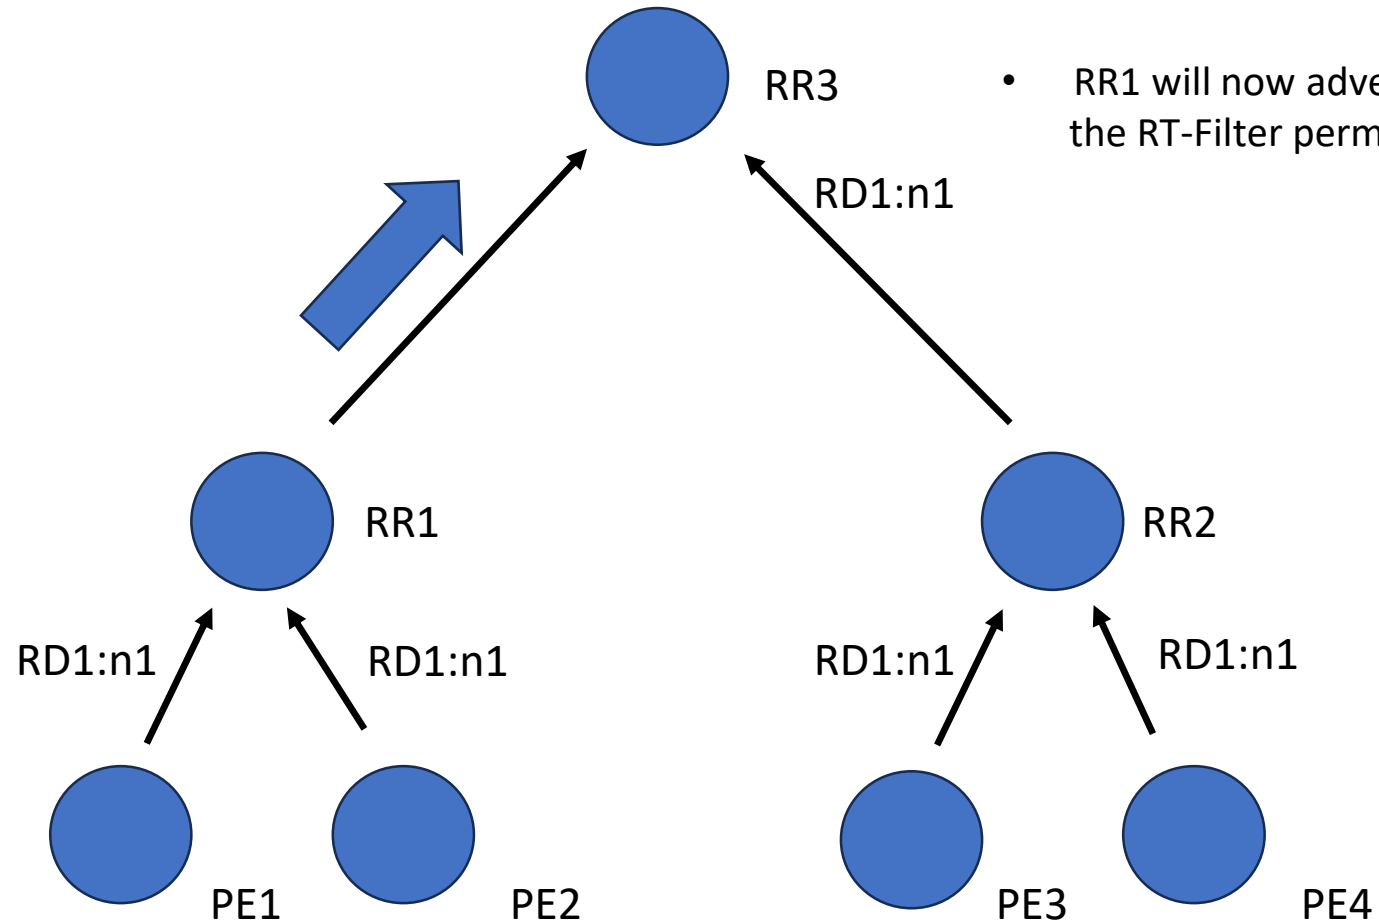
- When receiving a RTC route, the following rules will apply:
 - CLUSTER_ID, ORIGINATOR_ID and NEXT_HOP checks will be considered, but instead of discarding the routes, the route will be kept in Adj-RIB-IN as a Received-only route.
 - A route in Received-only state will not be considered for best -path nor advertised to any peer.
 - A route in Received-only state will be considered to install a VPN filter.
- The rules above could be used also when there is just one level of RR. This solution is not considered in [RFC468]
- This can interact with Rule 1 (for senders)
 - Mixed Deployments where Rule 1 is not supported
 - Rule 1 will effectively have priority

Solution #2



Solution #2

Propagation of VPN Routes



- RR1 will now advertise VPN routes to RR3 as the RT-Filter permits it.

Thanks

Questions?