

draft-litkowski-idr-bgptimestamp-02

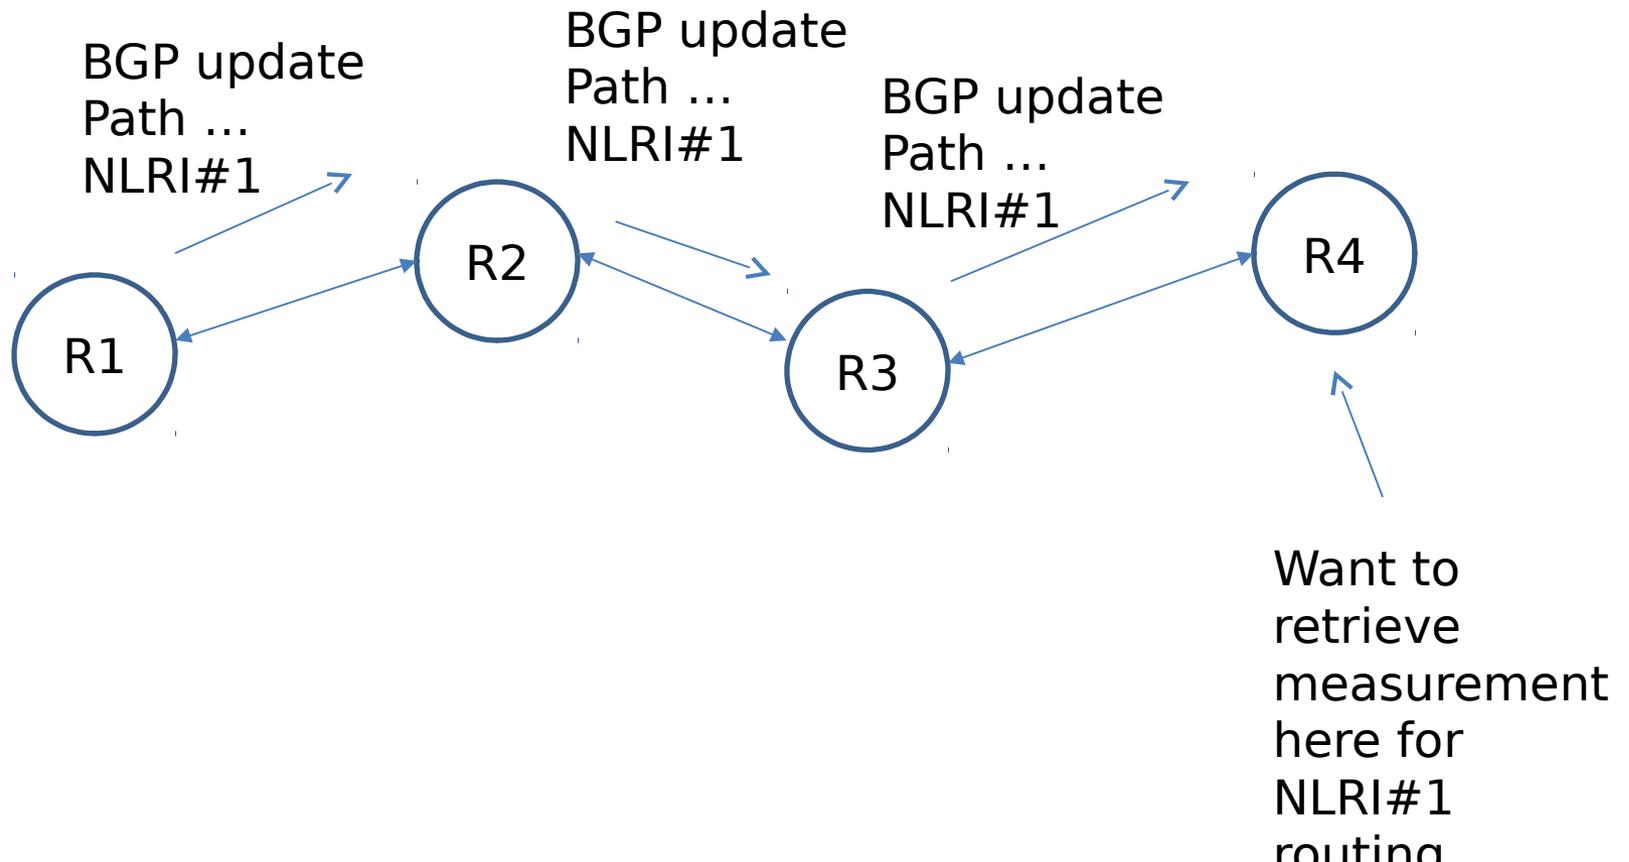
S. Litkowski

J. Haas

K. Patel

Problem statement

- How much time does it take to propagate BGP path ?



Problem statement

- Why is it important ?
 - Path propagation time involved in restoration of traffic in case of failure
 - Path propagation time involved in time to receive multicast flow
 - We already faced a lot of issues linked to BGP propagation
- Why it becomes more critical now ?
 - More and more AFI/SAFIs
 - Some AFI/SAFI may impact other
 - Some AFI/SAFI are more critical
 - Some « non critical » churn in a single AFI/SAFI will prevent important information to be propagated in a fast way

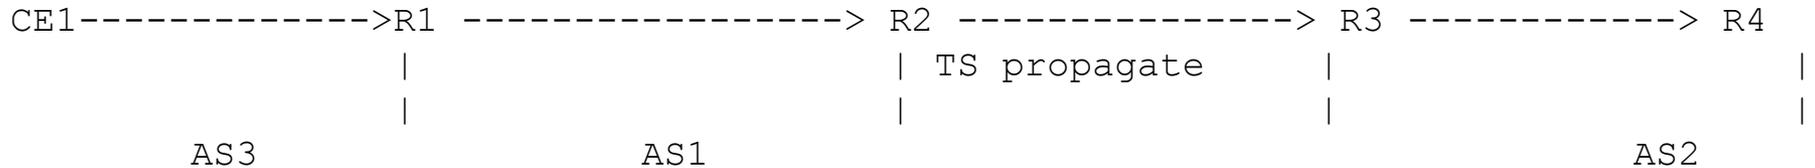
Problem statement

- Requirements :
 - Need to monitor performance of controlplane
 - Single point of listening, no/limited correlation
 - Need to be able to identify bottlenecks
 - Need accuracy

Proposed solution

- New BGP attribute : BGP-TS
- Use a timestamp vector that will be updated by each BGP Speaker in the path vector.

BGP update	BGP update	BGP update	BGP update
10.0.0.0/8	10.0.0.0/8	10.0.0.0/8	10.0.0.0/8
TS:	TS:	TS:	TS :
AS3;CE1:rT1,sT2	AS3;CE1:rT1,sT2	AS3;CE1:rT1,sT2	AS3;CE1:rT1,sT2
	AS1;R1:rT3,sT4	AS1;R1:rT3,sT4	AS1;R1:rT3,sT4
		AS1;R2:rT5,sT6	AS1;R2,rT5,sT6
			AS2;R3,rT6,sT7



Proposed solution

- Scope :
 - Do the timestamping only on for specific NLRIs to be monitored (real or beacons)
 - Not on all ! (notion of inspection list)
- Sampling will help to provide good accuracy
 - We already have such sampling system today which already helps to identify blackholes, or high delays in controlplane

InterAS

- Service provider may not want to expose its timestamp information to external peers
- Four options available :
 - Propagate : propagate TS vector as for internal peers (all details provided)
 - Drop : strip BGP-TS attribute
 - Drop AS : drop all TS entries for the local AS before sending
 - Summary : modify TS vector by aggregating local AS entries into a single summary AS entry

Next steps

- As a SP, we think that there is a need for such BGP monitoring solution
- WG feedback ?
- Is GROW a home for the problem statement and solution And let just encoding at IDR ? Or do we keep the work in IDR ?