


draft-litkowski-idr-bgp-timestamp

S. Litkowski, Orange

K. Patel, Cisco

J. Haas, Juniper

Problem statement



BGP infrastructure is used to transport more and more services.

Service Provider has to ensure that each AFI/SAFI requires guaranteed SLAs (path propagation time)


SLAs have to be monitored

Problem statement


- Service Providers require a solution for monitoring BGP path propagation time :
 - Single point of listening
 - No or very limited correlation need on collector
 - Bottlenecks identification
 - Accuracy of and between measurements (synchronization required)

Proposal

- Add a timestamp vector to BGP path to monitor propagation delay and track bottlenecks



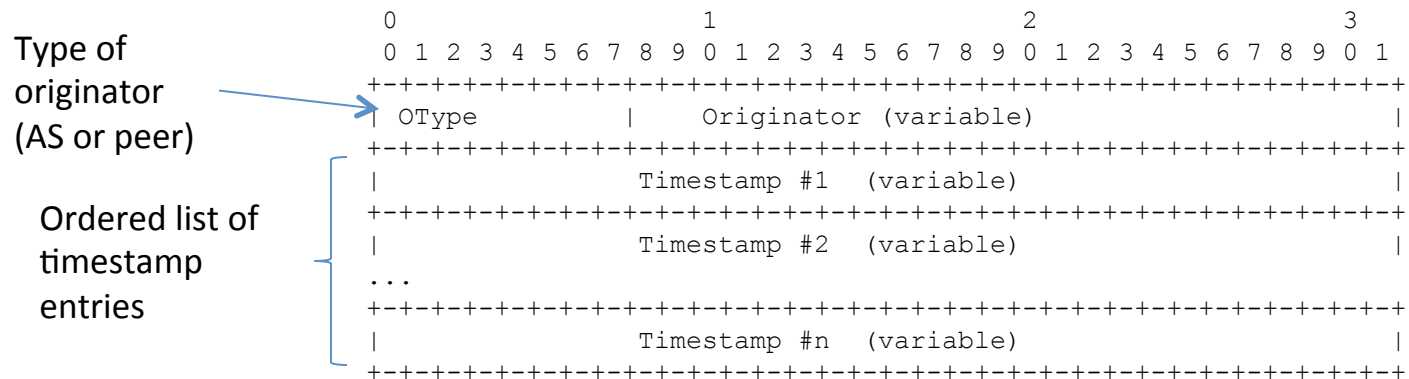
Architecture for propagation time measurement



- An external tool retrieves timestamp vectors from particular nodes
- Timestamp vector contains informations about the end to end propagation

Proposal

- Timestamp attribute (Optional transitive)



- Timestamp entry :

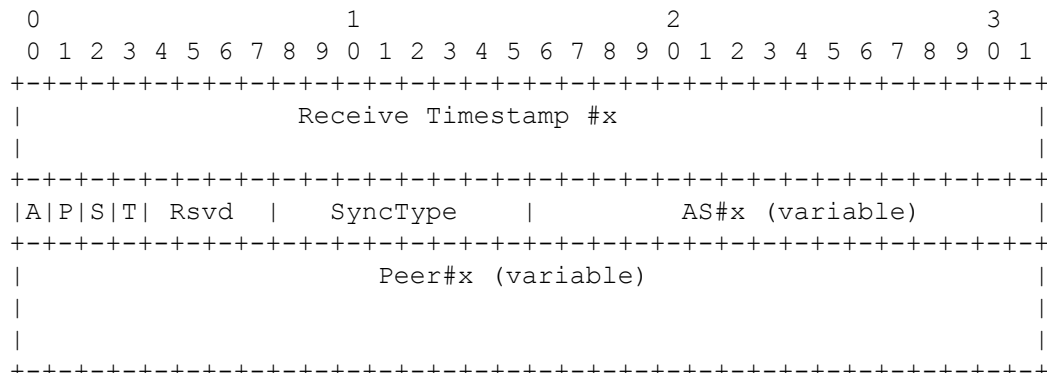
Flags :

A : AS type (AS2 or AS4)

P : Peer type (IPv4, IPv6)

S : Summary

T : Synchronized



Proposal

- Inspection list :
 - Timestamp are added only to a subset of BGP path matched by the inspection list (filter)
 - By default, do not timestamp
- Receiving a BGP path :
 - If BGP speaker supports BGP TS and path matches inspection list, timestamp is done :
 - If BGP-TS attr does not exist, it is created
 - If BGP-TS attr exists, new timestamp entry is added
 - Timestamp added is receive timestamp
 - If BGP speaker does not support BGP TS, it follows RFC4271 (transitive attribute)

Proposal

- Sending BGP path :
 - We suggest to send BGP TS attribute to only peers configured locally with a « send timestamp » option

Inter AS

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

Compared to BMP

- BMP does not mandate timestamps
- To retrieve timestamp vector, a BMP session would be required to each node including correlation in the external tool
- BMP does not provide information about synchronization state of the peer (is it free run or NTP ?)
- BMP basically dumps all received updates, strong filtering is required to catch interesting NLRIs

Next steps ...

- We need a solution to monitor BGP update propagation time
- First shot proposal and we have issues to solve with current draft :
 - NLRI with multiple originators and best path change over time (may lead to advertise stale timestamps)
 - Churn introduced by adding a new « variable » attribute
- WG feedback ?