

# Export of On-Path Delay in IPFIX

draft-tgraf-opsawg-ipfix-on-path-telemetry-00

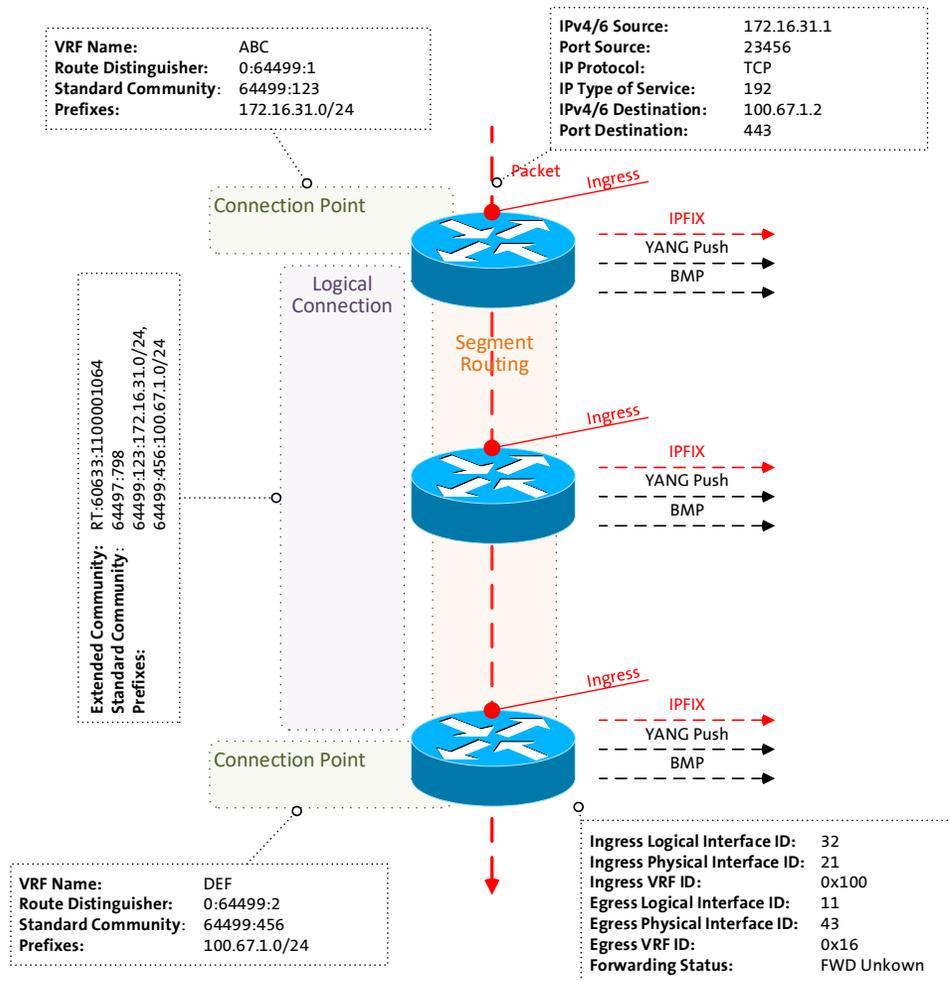
Enabling a statistical network delay view, giving insights  
where delay is being accumulated in the forwarding path

thomas.graf@swisscom.com  
benoit.claise@huawei.com  
alex.huang-feng@insa-lyon.fr

29. October 2022

# On-Path Delay with IPFIX Flow-Aggregation

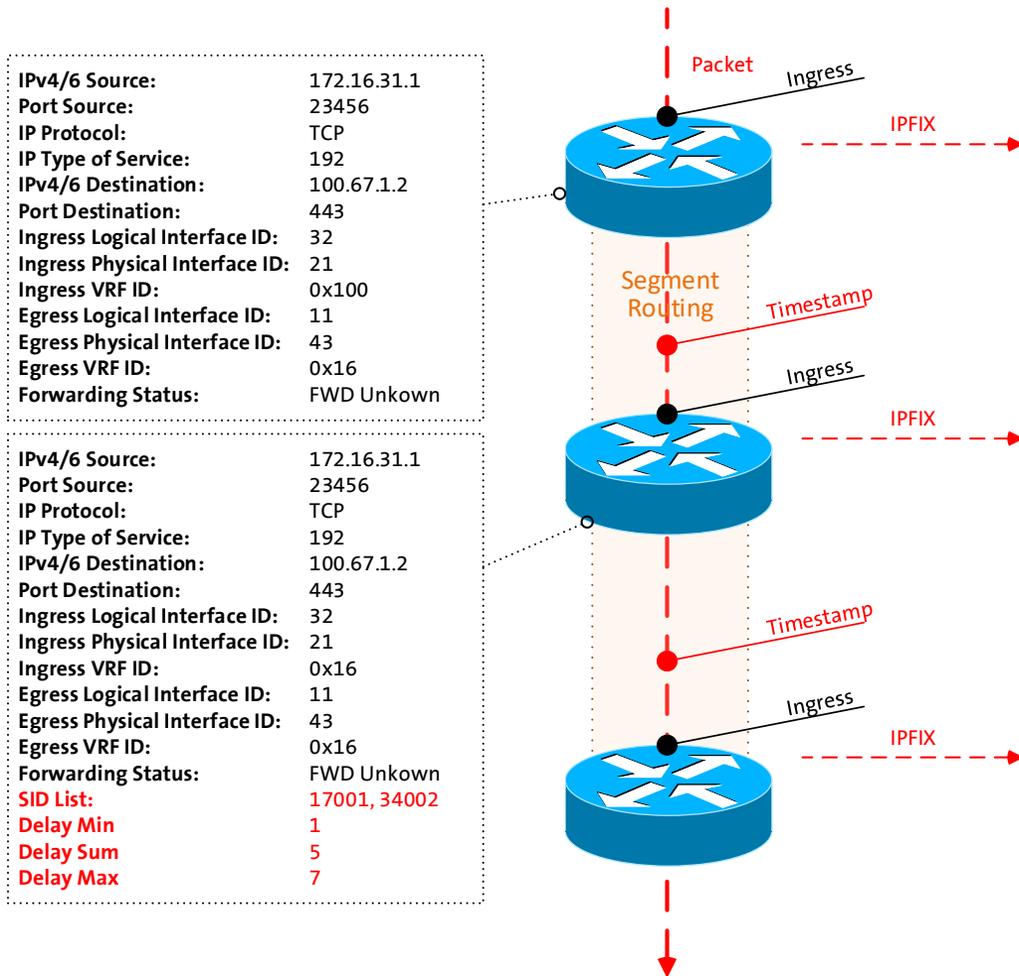
Aggregate and sample as early as possible – Chose your Cardinality



- IPFIX defines two key data engineering tools to reduce collected and exported amount of data. **Sampling and Aggregation**. Enabling a **statistical view from the network usage**. Also called **connectivity matrix**.
- IPFIX **measures packets and bytes** and give **device and control-plane context**.
- **With On-Path delay measurement**, IOAM, Path Tracing and iFIT, **delay can be measured** actively (probing) or passively. Metrics are exposed on every IOAM node, postcards or only at the last node (passport).
- **IPFIX lacks the ability to export delay**. A key element for monitoring Customer Service Level Agreements.
- **On-Path delay measurement lacks Flow Aggregation support** as defined in RFC 7015. Therefore, **scalability** in terms of data export and collection is **drastically limited** today.
- draft-tgraf-opsawg-ipfix-on-path-telemetry enables IPFIX to export delay while preserving the ability to aggregate and also **adds the On-Path delay metric definition** in the performance registry for proper delay definition.

# Measure delay and give network context

Enabling a statistical network delay view

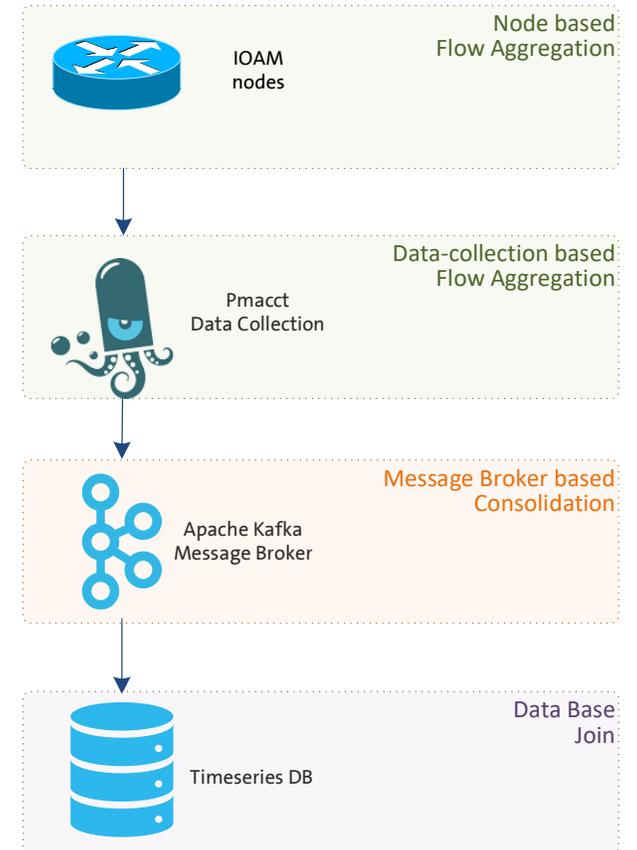


- Packets are captured ingress with an optional sampler, data-plane dimensions extracted, enriched with device and control-plane dimensions and **added with a unique flow ID to a flow cache on the node for aggregation.**
- The data-plane dimensions answers **which packet**. The control-plane **which service**. The device dimensions **where in the network**.
- **In case of On-Path delay measurement, a timestamp and optionally a direct export tag is added to the packet header when entering the On-Path delay measurement domain.**
- Each subsequent packet for the same flow increases byte and packet count. Each new flow creates a new flow ID in the flow cache.
- **In case of On-Path delay measurement, At each node in transit (postcard) or only at the last node (passport), the delay is calculated by comparing the timestamp in the packet and when packet is received on the node. Delay is populated into the flow cache besides packet and byte count.**

# On-Path Delay @ IPFIX

Draft Status

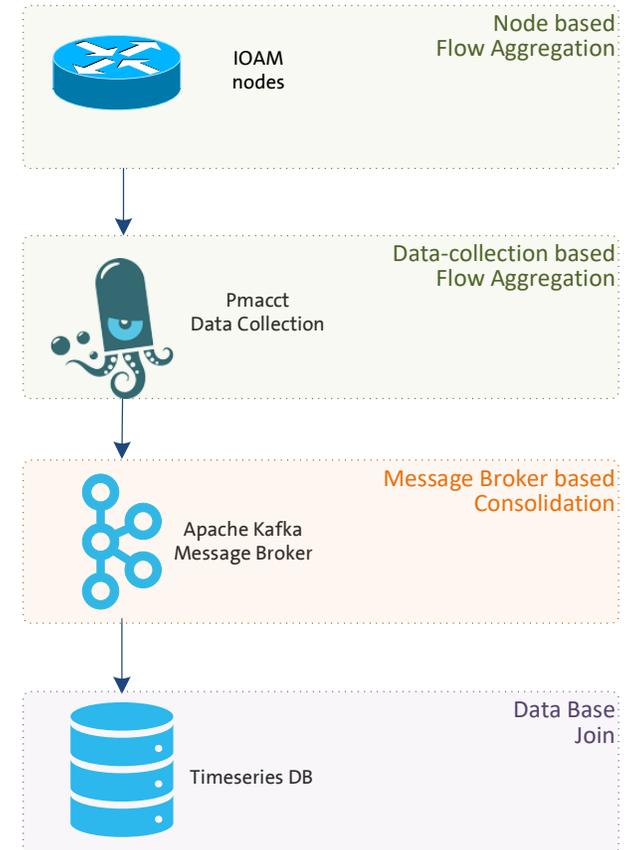
- Addressed comments from Greg to **change terminology from "Inband" to "On-Path" Telemetry.**
- Received comments from Tianran to change terminology from "IOAM transit and decapsulation nodes" to "transit and decapsulation nodes", **removing IOAM.** Authors believe that the **term IOAM applies to all On-Path Telemetry protocols.** -> Request feedback from the list.



# On-Path Delay @ IPFIX

## Next steps

- **Do you recognize the problem statement?**
- Network operators want to understand
  - **where delay with which network and device dimensions** is being accumulated
  - at highest scale **for a statistical network delay view.**
- IEs in document defined are independent from how the delay is being metered.
- Two vendors are validating technical feasibility. Others showing interest.
- INSA Lyon working on running open-source code in FD.io VPP.  
**Will be shown at IETF 116 hackathon.**
- Draft version -01 will contain data record and template examples.
- **-> Requesting adoption at OPSAWG.**



thomas.graf@swisscom.com  
benoit.claise@huawei.com  
alex.huang-feng@insa-lyon.fr

29. October 2022