Exporting Aggregated Flow Data using IPFIX

~~(draft-trammell-ipfix-a8n-00)~~

(draft-trammell-ipfix-a9n-0x)

B. Trammell, E.Boschi, A. Wagner

IETF 78 Maastricht – 29 July 2010

Scope of work

- Provide terminology and architecture for export of aggregated flows using IPFIX
- Define requirements for aggregation within IPFIX
- Outline issues specific to aggregation, and define implementation-independent approaches to solving them
- Define new information elements specific to aggregated flows
- Define metadata export for aggregated flows

Terminology

- Aggregated Flow: A Flow, as defined by [RFC5101], derived from a set of zero or more original Flows within a defined time interval. The two primary differences between a Flow and an Aggregated Flow are
  - (1) that the time interval of a Flow is generally derived from information about the timing of the packets comprising the Flow, while the time interval of an Aggregated Flow are generally externally imposed; and
  - (2) that an Aggregated Flow may represent zero packets (i.e., an assertion that no packets were seen for a given Flow Key in a given time interval).
- Intermediate Aggregation Function: mapping[1] from a set of zero or more original Flows, that separates the original Flows into a set of one or more given time intervals[2].
  - [1]This isn't *technically* a function since aggregation may be probabilistic (sampling within aggregation).
  - [2]Time intervals are practically almost always regular (e.g. every five minutes, every hour, every day), but *need not be*.

Architecture

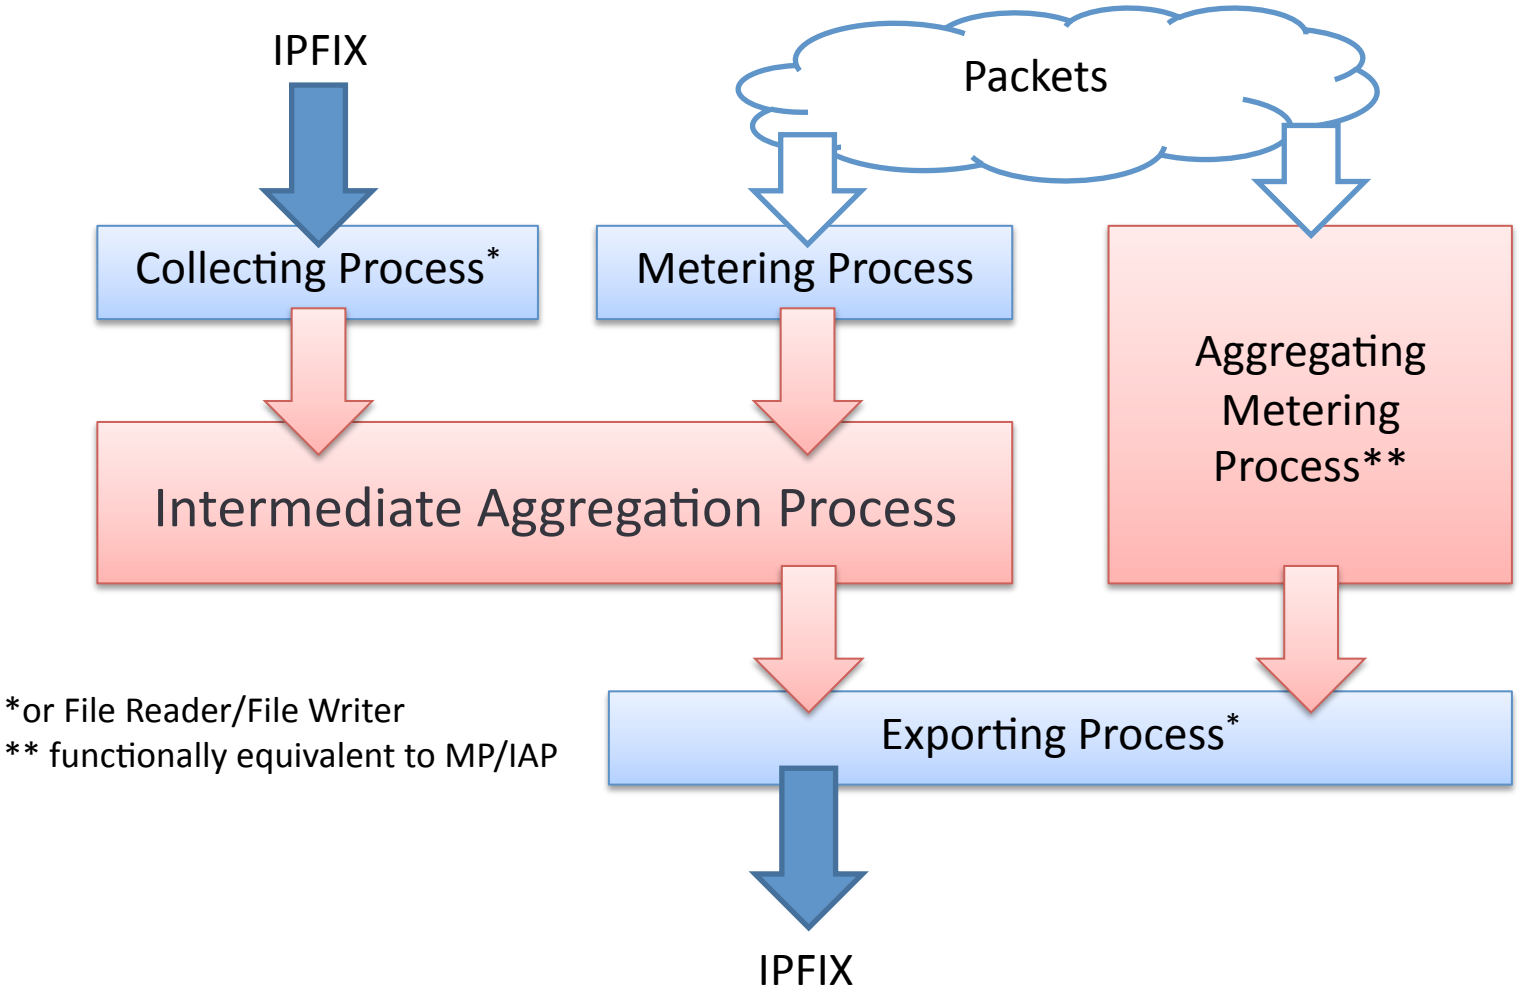

*or File Reader/File Writer
** functionally equivalent to MP/IAP

Requirements

- Interoperability/backward compatibility
  - Since an Aggregated Flow is essentially a Flow, any solution *must* allow *unmodified* 5101-compliant CPs to receive and correctly interpret Aggregated Flows.
- Implementation independence
  - Like the IPFIX Architecture, aggregation architecture is *descriptive*, not proscriptive.
  - Exported metadata describes properties of the data, *not* operations/algorithms on the data.

Aggregation-Specific Issues

- Counter distribution: how to derive counters for aggregated Flows when taken from an original Flow spanning multiple time intervals?
  - First, last, mid-interval
  - Simple and proportional uniform distribution
  - Simulated and direct distribution
- How to count original Flows?
  - Need both conservative and non-conservative flow counters
- How to export time intervals?
  - Attached to each flow, start and end
  - Metadata support for regular intervals (common case)

Distribution to Aggregated Flows

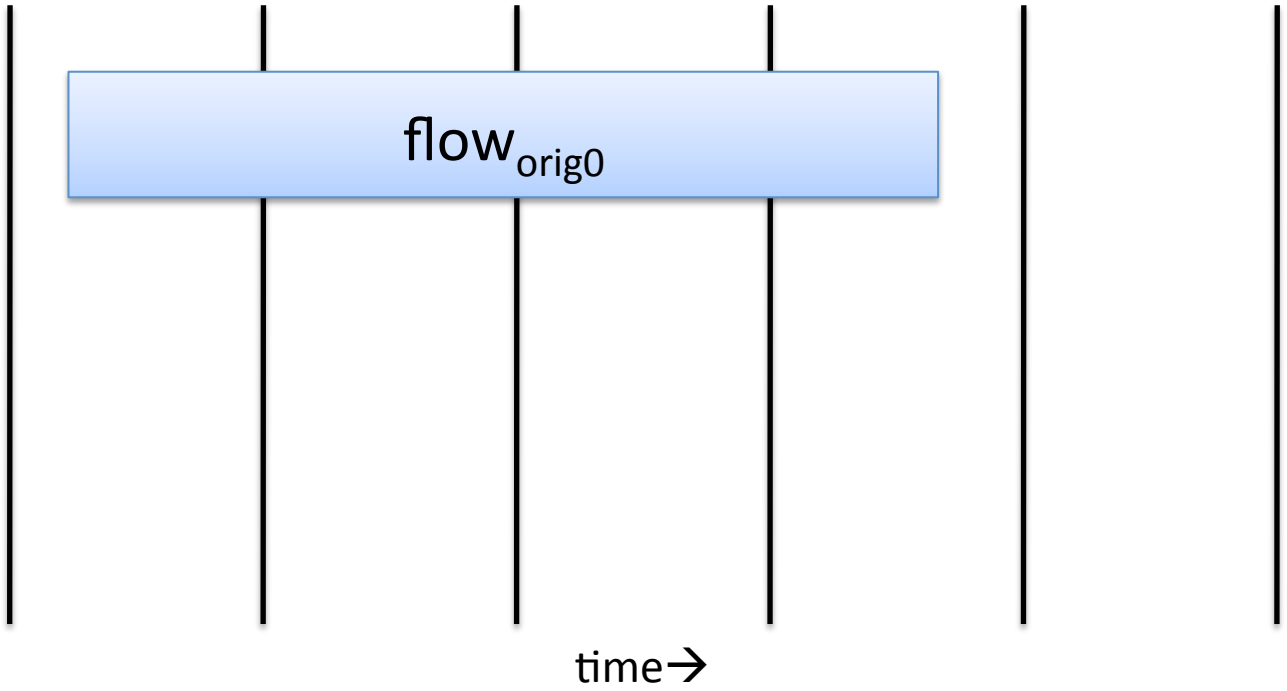

Distribution to Aggregated Flows

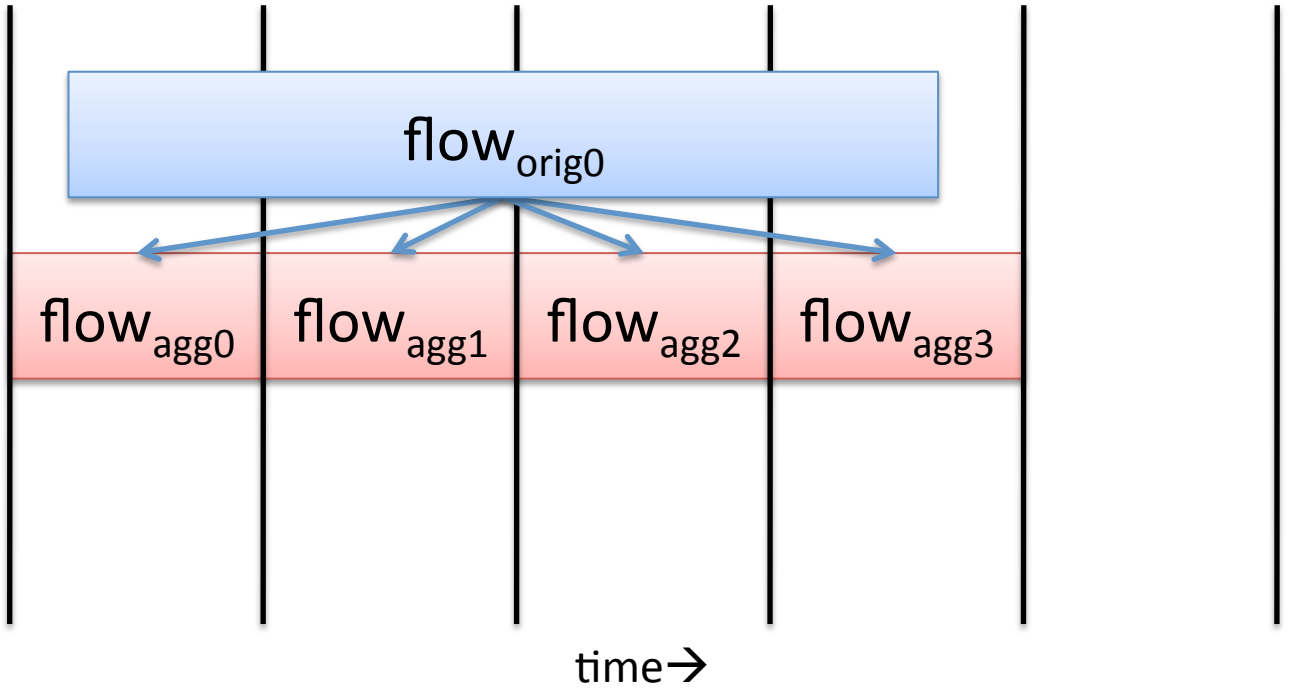

Distribution to Aggregated Flows

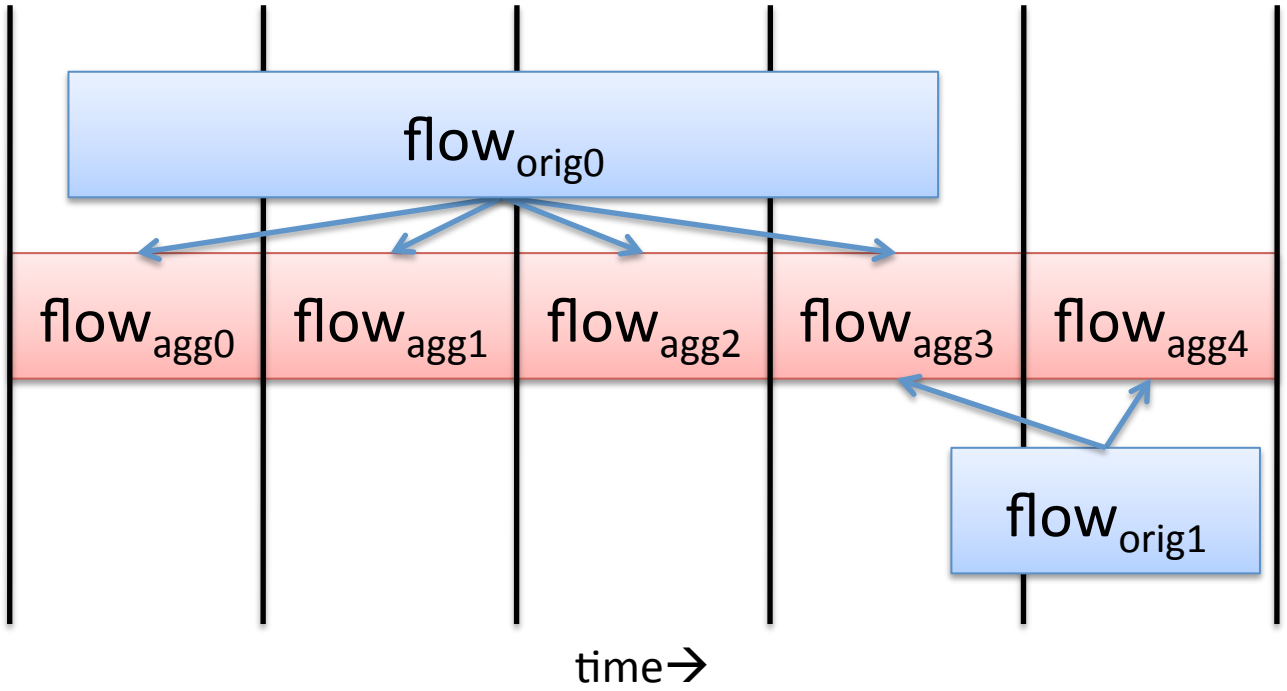

Flow Counters

- originalFlowsPresent: non-conservative counter
- originalFlowsInitiated: conservative counter, start interval distribution
- originalFlowsCompleted: conservative counter, end interval distribution
- originalFlows: conservative counter, any distribution, float64 representation

Next steps

- Complete the draft
  - Define metadata export for counter distribution
  - Define metadata export for time interval length
  - Fix the name (oops)
- Adoption as WG item within the Mediator activity.