Exporting Aggregated Flow Data using IPFIX

draft-trammell-ipfix-a9n-01

Brian Trammell Elisa Boschi Arno Wagner

IETF 79, November 6-13, 2010 Beijing, China

Outline

Introduction

- Flow aggregation is one of the most important, widely applied mediator operations
- One document per specific Intermediate Process $\rightarrow$ we should have one for aggregation.
- Draft defines a general, interoperable, implementation-independent model for IPFIX Aggregation
- Built into the Mediator framework.
- a9n = “a*ggregatio*n”, aggregated.

What do you mean, “aggregation”?

- Temporal? Spatial?
- Combining short-lived flows into long-lived flows?
- Imposing time intervals on flows for time-series data generation?
- Combining flows from multiple Observation Domains?
- Yes.
- Aim of the document is to cover *all* commonly used aggregation operations.

What do you mean, “aggregation”?

- Temporal? Spatial?
- Combining short-lived flows into long-lived flows?
- Imposing time intervals on flows for time-series data generation?
- Combining flows from multiple Observation Domains?
- Yes.
- Aim of the document is to cover *all* commonly used aggregation operations.

Temporal Aggregation

- Temporal aggregation is defined in Mediator framework as
  - "[m]erging a set of Data Records within a certain time period into one Flow Record by summing up the counters where appropriate," and
  - composition, wherein "multiple consecutive Flow Records with identical Flow Key values are merged into a single Flow Record of longer Flow duration if they arrive within a certain time interval."
- Definition does not handle externally imposed intervals, important for time-series reporting, so requires expansion.

Spatial Aggregation

- Spatial aggregation defined in Mediator framework as an operation wherein “Data Records sharing common properties are merged into one Flow Record within a certain time period.”
- Definition does not cover flow key reduction (one major reason that Data Records would share common properties during aggregation), so requires expansion.

Spatiotemporal Interdependence

- IPFIX Flows do not represent events with a single point in time, rather events over an interval.
- Spatial aggregation therefore has an unavoidable temporal component
  - The interval of the aggregated flows is the minimum covering interval; or
  - (more generally) intervals are externally imposed.
- Better to model aggregation as a series of operations of effects, than to separate aggregation in space and time completely.

Definition of Aggregated Flow

- A Flow, as defined by 5101, derived from a set of zero or more original Flows within a defined time interval.
- 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).
- The defined time interval is externally imposed (but may be derived from other flows part of the same Aggregated Flow)

Aggregation Operations

- Defined to be implementation-independent (much like the IPFIX Architecture, or the anonymisation draft).
- *Interval distribution* modifies an input flow's time interval, optionally creating multiple Flows.
- *Key aggregation* modifies flow keys by reduction or replacement.
- *Combination* pulls together flows resulting from these two steps into a single Flow for each key and time interval, applying *counter distribution* to distribute counters split by inverval distribution.
  - **TODO:** Not yet clear these are in the right order

Aggregation and the Mediator Framework

- Framework presents a generalized, not-quite-adequate definition of spatial and temporal composition in discussing Intermediate Aggregation Processes
  - chosen for a subset of aggregation operations covered here
  - does not address spatiotemporal interdependence
- Terminology in a9n operational model defined with reference to framework, but named differently, chosen to avoid collision.
- a9n is *more specific*, does not update/obsolete Framework.

Aggregation and the Mediator Protocol

- a9n handles data level aggregation, applicable to
  - mediators,
  - direct export of Aggregated Flows,
  - processing of files...
- Some issues in aggregation are actually mediator-general:
  - architectural issues in many-to-one aggregation across observation points,
  - template and observation domain management across an aggregating mediator
  - etc.
- These are handled in the Mediator Protocol draft.

Time Interval export

- Time Interval export: each flow SHOULD contain begin and end timestamps
  - maximizes interoperability (principle: a Flow is a Flow)
- MAY omit end timestamp IFF intervals are regular for a given Observation Domain within a Transport Session.

Flow Count export

- New information elements for counting original Flows contributing to an Aggregated Flow
- *Conservative* counts are preserved across re-aggregation, *non-conservative* are not.
- `originalFlowsPresent`: non-conservative
- `originalFlowsInitiated`: conservative, flows with start time within interval
- `originalFlowsCompleted`: conservative, flows with end time within interval
- `originalFlows`: conservative, general

Counter Distribution export

- When intervals are shorter than the longest flow, counters must be distributed across multiple intervals
- `valueDistributionMethod` in an Options record exports the method used to do this, on a per-Template basis:
  - Simple: start interval, end interval, mid interval
  - Linear: simple uniform, proportional uniform
  - Nonlinear: simulated process, direct
- This draft does *not* specify that aggregation MUST support exotic distributions of counters, or distribution export.

Counter Distribution export

- When intervals are shorter than the longest flow, counters must be distributed across multiple intervals
- `valueDistributionMethod` in an Options record exports the method used to do this, on a per-Template basis:
  - Simple: start interval, end interval, mid interval
  - Linear: simple uniform, proportional uniform
  - Nonlinear: simulated process, direct
- This draft does *not* specify that aggregation MUST support exotic distributions of counters, or distribution export.

Open Issue: Distinct Count

- Often useful to count distinct keys reduced away during Key Aggregation
  - e.g., unique destination addresses per source address
- How to export these?
- Current suggestion: `distinctCountOf`*InformationElementName* Information Elements be registered with IANA as needed.
- If chosen: need to specify this a bit more precisely.
- Other possibility: 5103-style PEN meaning “unique count of”
  - but this seems inordinately hackish...

Open Issue: Distinct Count

- Often useful to count distinct keys reduced away during Key Aggregation
  - e.g., unique destination addresses per source address
- How to export these?
- Current suggestion: `distinctCountOf`*InformationElementName* Information Elements be registered with IANA as needed.
- If chosen: need to specify this a bit more precisely.
- Other possibility: 5103-style PEN meaning “unique count of”
  - but this seems inordinately hackish...

Open Issue: Distinct Count

- Often useful to count distinct keys reduced away during Key Aggregation
  - e.g., unique destination addresses per source address
- How to export these?
- Current suggestion: `distinctCountOf`*InformationElementName* Information Elements be registered with IANA as needed.
- If chosen: need to specify this a bit more precisely.
- Other possibility: 5103-style PEN meaning “unique count of”
  - but this seems inordinately hackish...

Comments

- Comments from Benoit Claise (thanks!) have been very helpful in generalizing the draft
- Aim of the draft: completely cover aggregation
- Aim of the draft: remain implementation independent
- Reviews and comments, especially from implementors, help us do this!

Adopt as WG item?

- Draft in much better shape than Maastricht
  - Key insight: aggregation means different things to different people $\rightarrow$ need to cover them all
  - Key insight: spatial and temporal flow aggregation interdependent in general case $\rightarrow$ need to handle them together
- Continued improvements planned over the winter
- WG-mature by Prague