Exporting Aggregated Flow Data using IPFIX (draft-trammell-ipfix-a9n-02)

B. Trammell, E. Boschi,
A. Wagner, B. Claise

a9n in a nutshell

- Draft defines a general purpose architecture operational model for an Intermediate Aggregation Process (IAP), and support for aggregated flow export.
- Expands greatly on initial treatment given in Mediator problem statement and framework
- Much progress since Beijing and Maastricht to *generalize* a descriptive architecture for an Intermediate Aggregation Process

Contents

- 1. Introduction
- 2. Terminology
  - **Aggregated Flow**: *A Flow, as defined by [RFC5101], derived from a set of zero or more original Flows within a defined Aggregation Interval.*
- 3. Use Cases
  - Time series generation
  - Adaptive resolution of flow data
  - Anonymizing effects of aggregation
  - *This section requires some expansion, and harmonization with section 8 (Examples)*

More Contents

- 4. Architecture
  - How an Intermediate Aggregation Process fits into a Mediator, and with other IPFIX Architecture entities
  - A generalized, descriptive model for the internal arrangement of an Intermediate Aggregation Process
- 5. Operations
  - Detailed description of each of the operations outlined in the internal architecture
- 6. Additional Considerations
  - Exact versus Approximate Counting
  - Considerations for Aggregation of Sampled Data

Yet More Contents

- 7. Export
  - Guidelines, IEs, and options templates for exporting according to the model in sections 4-6
- 8. Examples
  - Traffic Time-Series per Source
  - Core Traffic Matrix
  - Distinct Source Count
  - Traffic Time-Series with Counter Distribution
  - *These are rough outlines only, and need completion*
- 9. Security
- 10. IANA

IAP Architecture

- Decomposition into iterative temporal and spatial steps
- Spatial aggregation implies temporal aggregation
  - interdependency due to special treatment of intervals in IPFIX

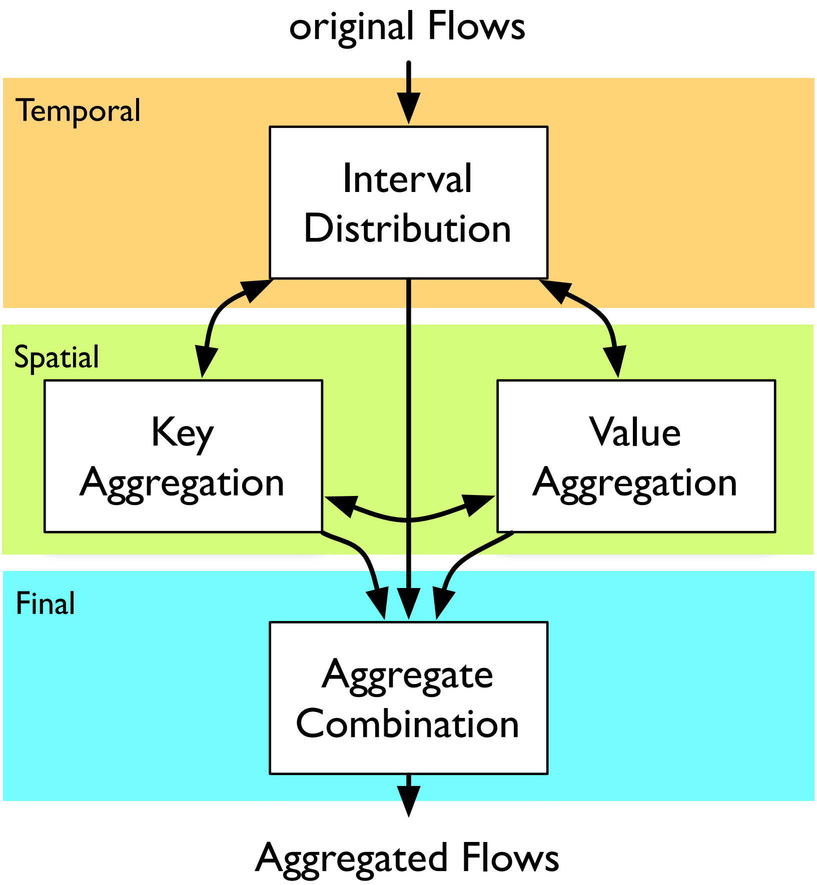

Interval Distibution

- Imposition of a time interval on partially aggregated Flows
- Time interval may be…
  - …regular (e.g. "5-minute bins") or irregular
  - …original-Flow-dependent (e.g. "in
  - …eternal (i.e., interval covering all packets in all the contributing flows)
- Interval distribution can be applied alone
  - e.g., to join long or low-active-timeout flows that were split by the original MP

Spatial Aggregation Operations

- Key aggregation: add or modify Flow Key fields in partially aggregated Flows
  - e.g. sourceIPv4Address masking or AS lookup
- Value aggregation: add or modify non-Key fields in partially aggregated Flows
  - e.g. counter averages or other descriptive statistics
- Aggregate combination: combine duplicate records for the same interval and keys into a single Aggregate Flow record.

Exporting Aggregates

- Guidelines for time interval export
- IEs for original flow counting
  - originalFlowsPresent: non-conservative
  - originalFlowsInitiated/Completed: conservative
  - originalFlows: conservative, distributable
- IEs for distinct host counting
  - distinctCountOf[Source|Destination]IPv[46]
- IE/Options for valueDistributionMethod
  - Applicable in specific cases where imposed interval shorter than intervals on original Flows

Next steps

- Better use cases and complete examples
- WG adoption for submission to IESG in late 2011
- Incorporate reviews and comments from IPFIX WG
  - Ensure widest possible applicability of the draft