Motivation for Passive Packet Sampling [Why psamp?]

Albert Greenberg

AT&T Labs

agreenberg@att.com

psamp@ops.ietf.org

Outline

- Goals
- Why standardize passive packet sampling?
- ipfix and psamp
- Example applications
- Summary

Goals

- Aim to greatly assist a very wide range of applications, which benefit from detailed passive measurements of traffic demands
  - e.g., traffic engineering, DoS attack detection, data for capacity planning and billing
- Aim for simplicity
  - call on a very simple set of primitive capabilities, which can be implemented ubiquitously at maximal line rate with minimal additional state, to support reliable, detailed, direct, timely measurements
- Allow for flexibility in implementation
  - allow simple configuration of sampling and export parameters
  - tuneable control over volume of measurement data
  - stay clear of discussion of integration with packet control actions (policing, marking, shaping, queuing).
  - attempt to decrease the burden of export of router state needed to interpret exported usage information
  - full packet capture not in psamp scope (RFC 2804)

Why Passive Packet Sampling?

- Why passive?
  - To measure traffic across all edges
- Why packet?
  - To obtain information immediately beyond what we get from passive SNMP coarse-grained counters and active performance probe data
- Why sampling?
  - To scale to high rate, and enable implementation across all network edges, while trading off some statistical accuracy

Why Standardize Passive Packet Sampling?

- To create standard with consistent and well-defined interfaces to support a broad spectrum of applications
  - Provide specifications that vendors can build to
- To reach agreement among network vendors, software developers, xSPs on simple traffic measurement capabilities for operational management tasks
  - Some of the related products/solutions now on: INmon, Juniper, Foundry, Cisco (raw sampled netflow)
- To help drive towards obtaining these capabilities in every monitor, every router, every line card, every measurement ASIC, …
  - Just like SNMP usage statistics (which are simple!)

psamp and ipfix

- ipfix is concerned with standardizing passive flow measurements
  - A very good thing. See http://www.ipfix.doit.wisc.edu
  - Focus on export of aggregations providing summaries of packet trains
- psamp is concerned with standardizing passive, packet sampling capabilities
  - Offers packet-level measurements to higher level applications, which might be “on-board” or “off-board”
  - Allows for low-latency between measurement and reporting, which will be particularly useful
  - Aims for parallel measurement
    - e.g., 1 in N continuous sampling for baselining
    - e.g., access-control-list-like filters with associated counters for billing
- Aim
  - Listen and learn from ipfix. There is potential to use ipfix solutions for data export and information model, where requirements line up
  - Don’t slow either effort down

Idea

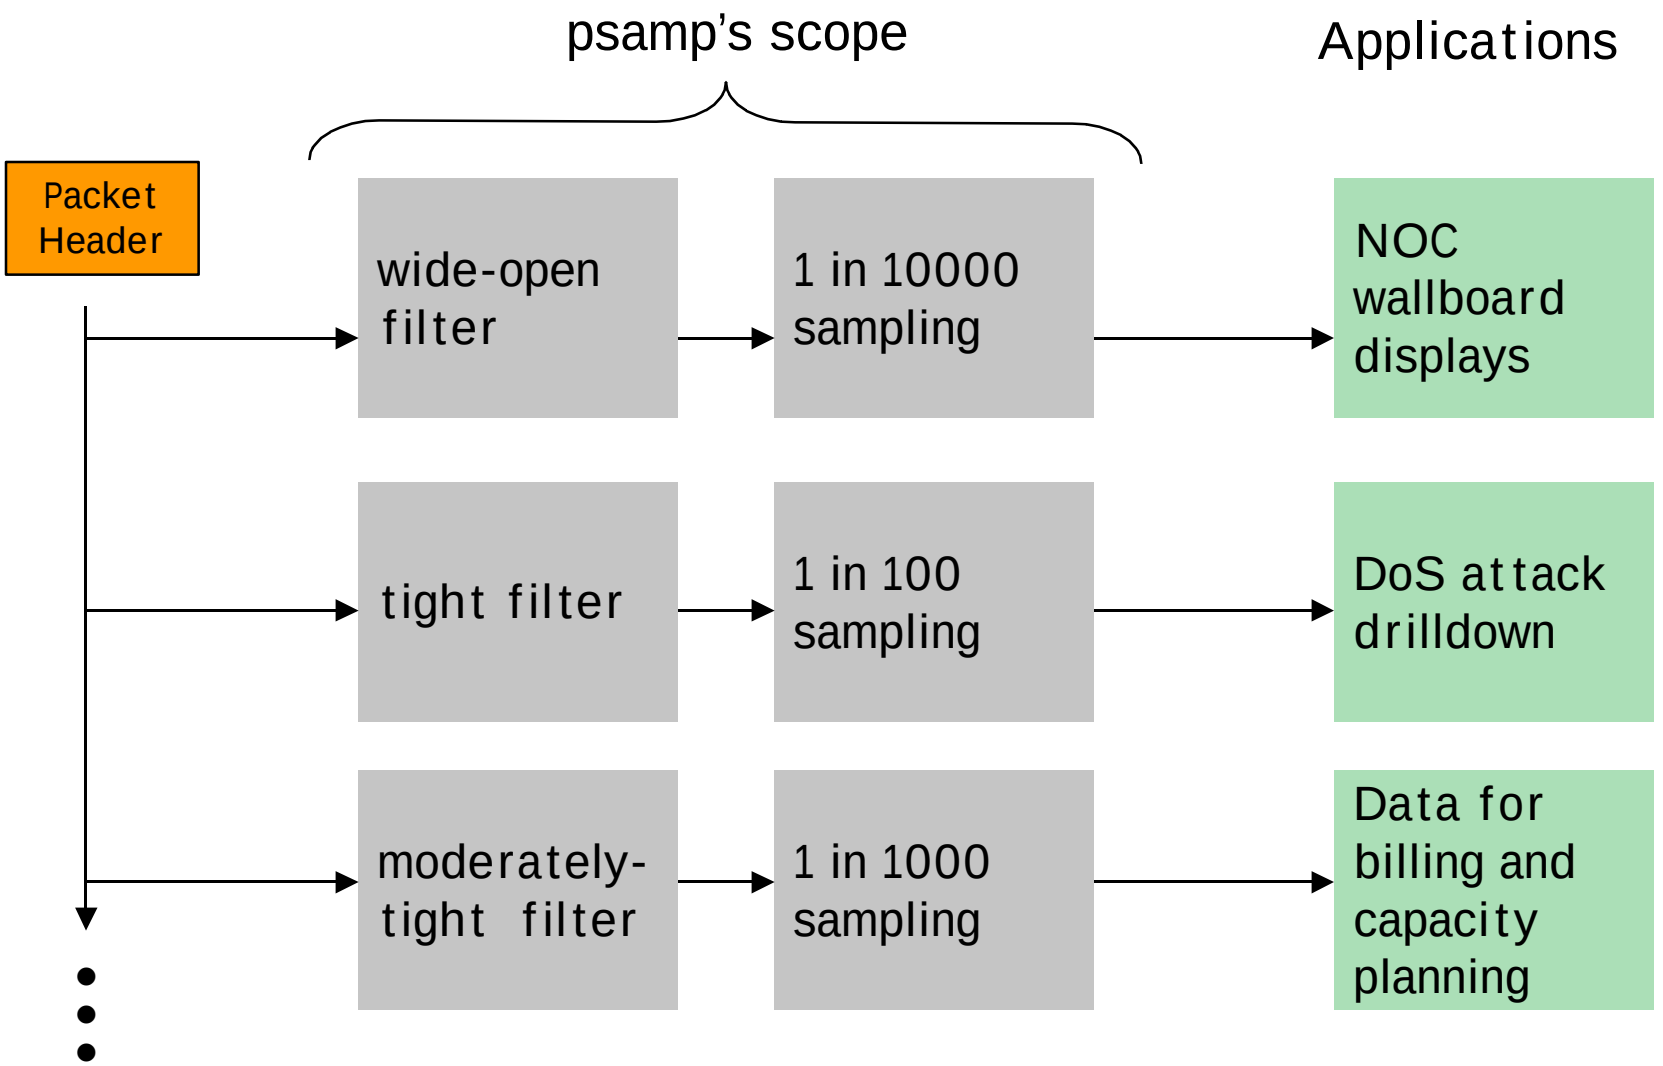
psamp's scope
Applications
Packet Header
wide-open filter
1 in 10000 sampling
NOC wallboard displays
tight filter
1 in 100 sampling
DoS attack drilldown
moderately-tight filter
1 in 1000 sampling
Data for billing and capacity planning

psamp Primitives

- Which packets to select
  - filter: e.g., match/mask on source/destination prefix, port numbers, protocol, … + tags to indicate the associated (sub)interface
  - sample: e.g., 1 in N deterministic, random or hash-based
- What info to export
  - selected packet header fields
  - timestamp
  - certain associated router state (in/out interface, matching routing table entries for source/destination prefix and source/destination AS), if available
- Simple primitives are powerful
  - enable a very wide range of measurement applications
  - above suggestions just examples – remains for the working group to decide

Example Application: Troubleshooting

- Problem
  - On receiving congestion alert (e.g., high SNMP utilization, or large probe delay), identify which services, peers, customers impacted
- Measurement Solution
  - Use unfiltered sampling for coarse-grained view of the traffic demands. Identify interesting subset of traffic (e.g., a service type, or a source address prefix corresponding to some customer)
  - Refine filters to zoom in on this traffic, and boost the sampling rate correspondingly.

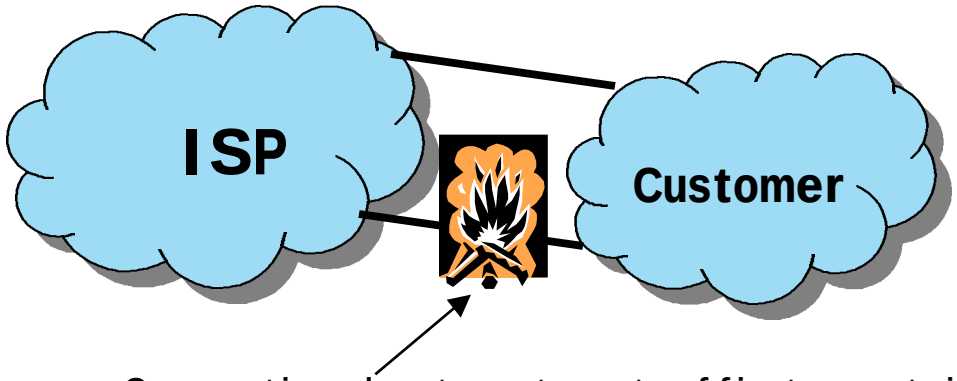

Congestion due to return traffic to certain customer prefixes

Examples: Traffic Engineering, Capacity Planning, Managing Peering Relationships

- Problems
  - Traffic engineering: improve service quality and asset utilization, via network-wide control of routing
    - valuable input: traffic matrix (e.g., volumes per ingress-egress pair)
  - Network engineering: improve design, capacity planning, where to attach new customers
    - valuable input: traffic matrices, over longer time scales
  - Manage peering relationships: adjust who to peer with and where
    - valuable input: AS-level level traffic matrices, over long time scales
- Measurement Solution
  - Sample packets across the network edge, looking for trends as well as significant shifts or anomalies in traffic.
  - Use wide-open, low rate sampling to identify heavy-hitters, and potentially use more narrow filters to drill down

Direct Observation of Network Behavior

- Problem
  - Capture information about the current network state and behavior
    - Identify the precise set of paths packets traversing an overloaded link
    - Trace the paths of traffic to a given prefix, for a multi-homed customer seeing congestion on one access link for that prefix
  - Today, this is hard
    - Involves scheduling unreliable downloads of voluminous routing and forwarding table, joins of data sets, and working with stale data
  - Need
    - A method essentially equivalent to selecting and marking packets at the edge and then selecting and measuring marked packets at every hop
- Measurement Solution
  - Sample a given packet at every hop in a domain, or not at all. Construct trajectories from the sampled packets.
  - Hash-based sampling. (Discussed later in the BOF)
    - N.B. AT&T may own intellectual property applicable to this contribution

Need to Control Measurement Overhead

- Need configurable maximum export rate
  - Want capabilities for high speed links
  - Can be problematic to predict the volume of measurement data
    - e.g., packets matching a filter associated with a DoS attack
  - Measurement infrastructure will be engineered to accept up to a particular rate of measurements
    - don't want to overload it
    - really about reliable engineering mechanisms ⇒ cap the rate that packets are supplied to transport
- Need information about missing data (e.g., sequence numbers)
  - Data can get lost inside the network or inside the router
  - Want to have sequence numbers and indications of number of packets that matched the filter that have not been exported
- Info on configuration state of sampling
  - E.g., sampling rate, filter type – finesse the operational headache of joining usage with the associated sampling configuration

Summary

- Application needs
  - network-wide measurements: e.g., routing policy optimization for traffic engineering
  - timely information: e.g., DoS attack detection
  - controllable accuracy: e.g., data for capacity planning
  - guidance for what-if's: e.g., what services to offer, whether to deploy caches, what billing model to use
- Implies capabilities that are reliable, detailed, direct, timely and available ubiquitously
- Goals
  - To reach agreement among community on simple traffic measurement capabilities for operational management tasks
  - To create standard with consistent and well-defined interfaces to support a broad spectrum of applications
  - As a 1st step, to focus discussion on charter of for a working group!
  - psamp@ops.ietf.org