

PSAMP Framework Document

draft-ietf-psamp-framework-02.txt

Duffield, Greenberg, Grossglauser, Rexford: AT&T

Chiou: Avici


Claise, Marimuthu, Sadasivan: Cisco

Summary of Changes

- Architecture
- Redefinition of "Measurement Process"
- Selection Process
- Reporting Process
- Export Process

Architecture

- ❑ Change definition of "Measurement process"
- ❑ Selection Process -> Reporting Process -> Export Process



- ❑ PSAMP measurement proc. analogous to IPFIX metering process
- ❑ Multiple parallel measurement processes
 - ✦ Can feed single export process
 - ✦ E.g. router with multiple line cards
 - Per line card measurement processes
 - Single export process on router

Selection Operations

- ❑ Output of selection operation is selected packet
 - ✦ Previously: output was binary selection decision
 - ✦ Now: easier to express ordered composite selection operations
- ❑ Selector Sequence Number
 - ✦ Each selector keeps counter of input packets
 - ✦ Counter value reported as sequence number for selected packets
 - ✦ Used at collector to infer attained sampling rate (c.f. sFlow)
 - attained sampling rate needed to infer actual traffic rate
 - robust with respect to loss of reports after sampling

Selection Operations

- ❑ Count based vs. Timer based
 - ✦ Timer based: simpler to implement? Useful for IPPM support
 - ✦ Count based: more accurate for single packet statistics (Claffy 93)
- ❑ Simple random sampling:
 - ✦ generalization: n from N random sampling
- ❑ Systematic 1 in N sampling
 - ✦ generalization: n from N periodic
- ❑ Hash-based
- ❑ Stratified, non-uniform probability
 - ✦ Probably too complex to specify at first cut
 - ✦ No big demand from applications

Filtering

- ❑ Selection of packets based on packet fields, packet treatment
 - ✦ Demand from applications, e.g. drill-down
- ❑ Don't expect all PSAMP devices to support filtering
 - ✦ filter must first parse fields, then filter on them
- ❑ Feasible for many existing devices
 - ✦ routers already parsing fields and filtering for ACLs
 - ✦ packet treatment also available at sufficiently low rate
- ❑ Filtering for measurement can be simpler than filtering for ACL
- ❑ Simple Proposal
 - ✦ Filter on each of set of fields:
 - Single match/mask (IP addresses, TCP flags, ...)
 - Single range (TCP/UDP port numbers, AS numbers,)
 - ✦ Select packet if it passes all field filters
 - ✦ Simple to configure in MIB
 - No attempt to reproduce complexity of general ACL specification

Composite Selection Operations

- Application: drill down, e.g.,
 - ✦ Baseline 1 in 10,000 sampling: notice "interesting" traffic
 - ✦ Configure filter onto interesting traffic, 1 in 100 sampling
- Proposal:
 - ✦ Allow composition of filtering with sampling, either order
 - Filtering -> Sampling
 - Sampling -> Filtering
 - ✦ Advantageous to put first the operation that thins traffic most
 - Allowing either order extends domain of utility

Multiple Parallel Measurement Processes

- ❑ Multiple measurement processes acting on same traffic stream
- ❑ Application: drill down, e.g.,
 - ✦ Baseline 1 in 10,000 sampling: notice "interesting" traffic
 - ✦ Configure filter onto interesting traffic, 1 in 100 sampling
- ❑ Want to be able to drill down while continuing base measurements

Which sampling operations?

□ Capability Model

- ✦ Standard specifies each sampling method
- ✦ Implementers decide which to support
- ✦ Marketplace decides which are important

□ Conformance Levels

- ✦ MUST/SHOULD/MAY
- ✦ Standards decide minimum PSAMP capabilities
- ✦ Clearer understanding of minimum PSAMP capabilities in practice
- ✦ What are the criteria to decide?
- ✦ Difficult to place newer sampling methods in correct level

Current Draft Proposal: 2 Conformance Levels

□ MUST

- ✦ one of 1 in N systematic, or 1/N simple random
 - both are currently available from vendors

□ SHOULD

- ✦ both options above
- ✦ n from N systematic sampling
- ✦ hash-based selection
- ✦ filtering (see slide 6)
- ✦ composite selection operations (see slide 7)
- ✦ at least 2 parallel selection processes

Packet Reports

❑ Mandatory Reports:

- ✦ Report first n bytes beyond link level header
 - No protocol and field parsing required
 - Burden of interpretation falls in collector
- ✦ Report sequence numbers from selection operation(s)
- ✦ Report PSAMP device interfaces used by packet
- ✦ Any additional fields calculated during sampling e.g. hash, timestamps

❑ Optional Reports:

- ✦ Report configurable combination of selected fields instead of n bytes
- ✦ Saves bandwidth, less burden of interpretation for collector
- ✦ Should not be hard for a device that already filters on fields

❑ Either should be compatible with IPFIX, suitably tweaked

Export Process

- ❑ No substantive changes, yet
- ❑ Requirements
 - ✦ Congestion avoiding
 - ✦ Not onerous on PSAMP device
 - ✦ Reliability not required, avoid overhead (buffers, ack processing)
- ❑ Candidate export protocols
 - ✦ Collector based rate renegotiation
 - ✦ Protocols in development (DCCP, PR-SCTP?)
 - ✦ Whatever IPFIX decides
 - NetFlow v 9 basis, + TCP, unreliable transport TBD (PS-SCTP?)
- ❑ Other proposals?