

Packet Generation in Service Function Chains

Draft-penno-sfc-packet-03
IETF96 - Berlin

Reinaldo Penno (rpenno@cisco.com)
Carlos Pignataro (cpignata@cisco.com)
Chui-Tin Yen (tin@cisco.com)
Eric Wang (ejwang@cisco.com)
Kent Leung (kleung@cisco.com)
David Dolson (ddolson@sandvine.com)

Review of Problem Statement

Service Functions generate packets in the reverse flow direction to the source of the current in-process packet/flow

- Firewall
- NAT
- Proxies, caches, accelerators
- Intrusion Prevention Systems

Changes since -02

- Added an optimized method of encoding reverse path information in 32 bits of metadata (vs. 48 bits)
- Identified 6 classes of metadata

Methods of Path Reversal

1. SF has reverse path config -- configure SFs with reverse SPI/SI per forward SPI/SI
2. SF requests SFF to reverse using OAM protocol; SFF has reverse path config -- configure SFF with a reverse SPI/SI per forward SPI/SI
3. Classifier encodes reversal information in metadata
4. Algorithmic -- assuming systematic assignment of pairs of SPI/SI by control plane

Methods -- Scope of Impact

Method	Requires symmetric path	Control Classifier (C1)	Control SFF (C2)	Control SF (C3)	Requires OAM protocol	Special SPI/SI selection
1. SF has reverse path config				YES		
2. SF requests SFF to reverse; SFF has reverse path config			YES		YES	
3. Metadata encodes reversal	YES	YES				YES
4a. Same path, disjoint index (algorithmic)	YES					YES
4b. Flip path ID (algorithmic)	YES					YES

Document recommends Flip-Path-ID

- Solution contained to SF
- Assumes paths are assigned in pairs
- Requires symmetry in number of hops
 - Shows how “no-op” in SFF can create symmetry in some cases

“Same Path-ID with Disjoint Index Spaces” has similar properties and limitations

- Satisfying to use same SPI for both directions
- But more restrictions on SI selection

The No-Op Next-Hop in SFF

- The recommended approach does not work for asymmetric path layout
 - (The math assumes same number of SFs in each direction)
- Adding the no-op concept to the SFF simplifies things
- If the SFF finds “no-op” entry for a received SPI/SI, the SFF decrements the index and repeats the lookup.

Classes of Metadata

Type	Meaning	Example	Impact
Path-invariant	Same for all packets of a path	VPN identifier	Configure SF with the value for each path
Path default	Can vary packet by packet but there is a default value that works	QoS treatment	Configure SF with the default value globally or per path
Bidir cloneable	Represents both directions of a transport flow	Subscriber identifier	May be copied from one packet to another for the same transport flow.
Unidir cloneable	Represents one direction of a transport flow.	Ingress identifier	May be copied from one packet to another for same direction of same transport flow.
Created by SF	The SF is responsible for adding it	DPI output	The SF needs no special instructions
Reclassified	The SF cannot know what value to use	Policing color	Injected packet must be sent for reclassification

Questions for Audience

- Are SF implementers facing path reversal problem?
 - Yes? Which type of SF?
 - No? Which type of SF?
- What solutions have you chosen? Anything not in this draft?
 - Solutions for SPI/SI reversal?
 - Solutions for metadata of injected packets?

Adoption?

- Authors believe document is ready for adoption.
- Addresses critical issues for service chaining.
- Would chairs issue an Adoption poll?