Hierarchical Service Function Chaining (hSFC)

draft-dolson-sfc-hierarchical-03

David Dolson (ddolson@sandvine.com) -- presenting
Shunsuke Homma (homma.shunsuke@lab.ntt.co.jp)
Diego R. Lopez (diego.r.lopez@telefonica.com)
Mohamed Boucadair (mohamed.boucadair@orange.com)
Dapeng Liu (max.ldp@alibaba-inc.com)

History

- Concept introduced in draft-homma-sfc-forwarding-methods-analysis
- At IETF 92 (Dallas), there was interest in creating a separate draft
- Version -02 presented at IETF 93 (Prague)
- Initial Draft posted May 25, proposing some mechanisms
- 2 updates due to feedback
- Incorporated contributions from draft-liu-sfc-nesting-use-case-01
- Current version -03 posted Oct. 2.

The Problem

- SFC controller complexity in very large networks.
  - Millions of hosts
  - Thousands of forwarding elements
- Asymmetrical routing
  - But some SFs require bidirectional symmetry.
- Multiple operational teams

- How to avoid a “super controller” ?

The Solution:

- Make a big problem into several smaller problems

- The Key Idea:
- An SFC Sub-Domain can appear as a single SF to a high-level SFC domain

Hierarchical Service Paths

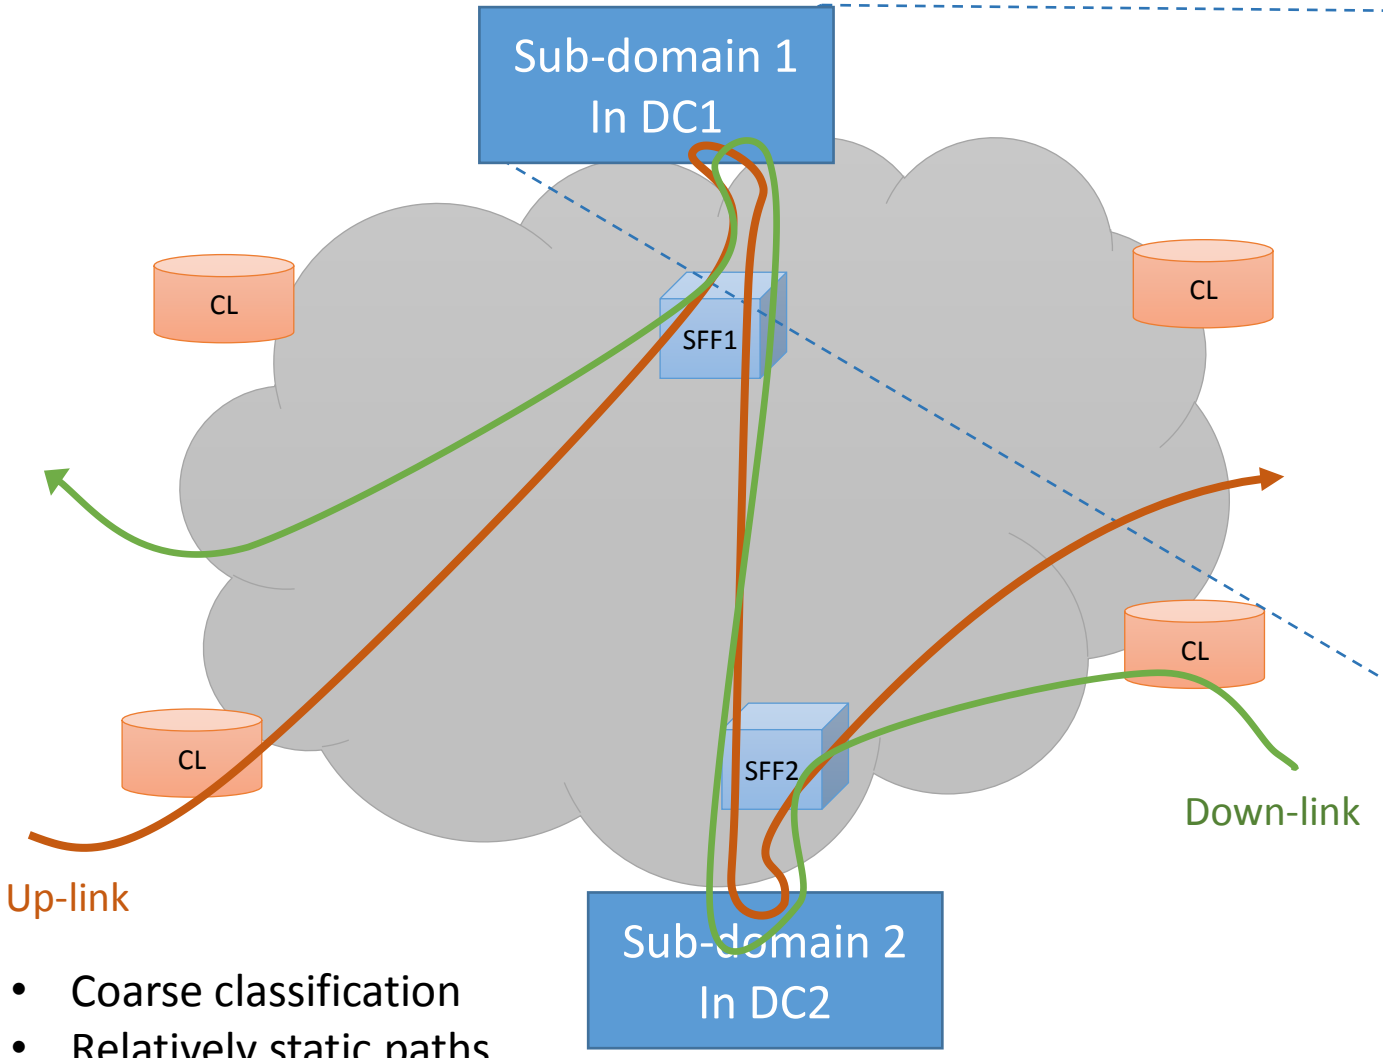

- Coarse classification
- Relatively static paths
- Geographically distributed classifiers

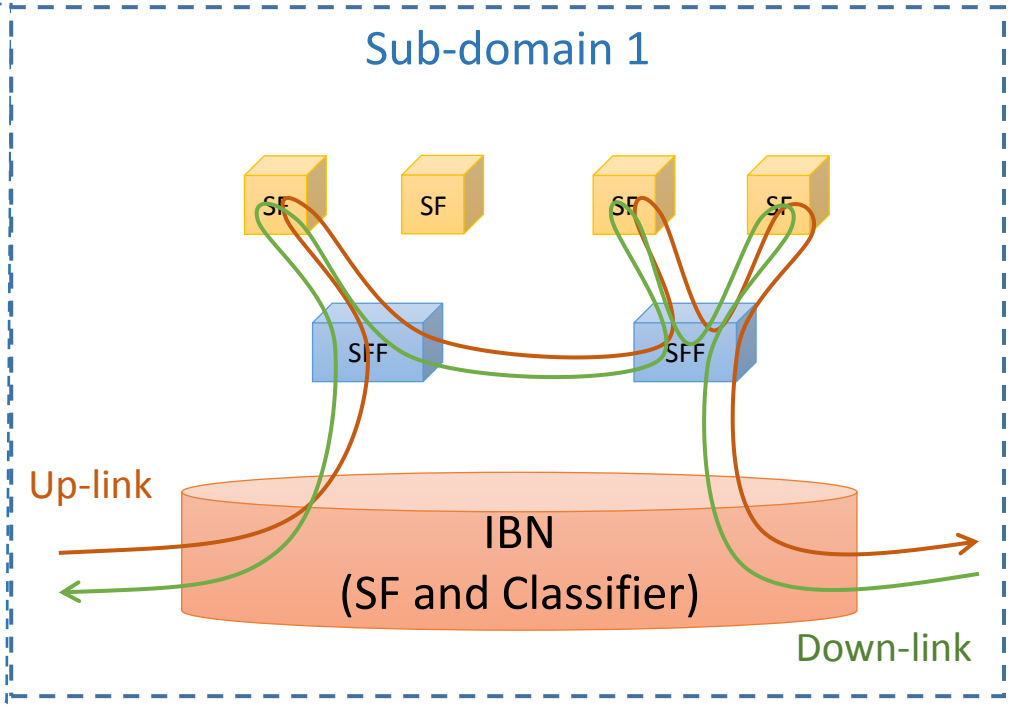

In sub-domain:

- Stateful 5-tuple classification
- Dynamic network policy
- Co-located classifiers to handle bidirectional traffic
- Co-located SFs to handle chatty control plane and NFV elasticity.

The Internal Boundary Node (IBN)

- SFC Architecture describes an “SFC Boundary Node”
  - [RFC7665, section 4.4]
  - Connecting SFC domains together
  - Does not specify details
- We propose the IBN to bridge levels of hierarchy within a single administrative domain
  - A variant of “SFC Boundary Node”
  - We have specified IBN behavior that is not described in RFC7665.
  - We identify IBN behavior to allow hSFC to be done safely.

Mechanisms

- Packets exiting lower-level domains are returned to paths in the higher levels. Challenge: which higher-level paths?
- Options:
  - Flow-stateful IBN – remember which path per 5-tuple
  - Encode upper-level paths as context metadata of lower-level
  - Unique lower-level paths per upper level path

Metadata Implications

- Metadata in the higher-level domain must be preserved when traversing the lower-level domain, by either:
  - Single metadata schema across domains
  - Pushing/popping/mapping mechanisms

Control Plane Implications

- IBN is an SF in the higher-level
- IBN is a Classifier in the lower-level
- Independence is desired

- Control-plane standards should permit hSFC

Examples in the draft

- Reducing the Number of Service Function Paths
  - Total number of managed paths is reduced

- Managing a Distributed Data-Center Network *new*
  - Avoid a super-controller across multiple DCs

Contributions

- Scalability to large networks
  - Can hide scaling considerations within a sub-domain
  - Avoid costly stateful classification in distributed classifiers
- Manageability of multiple domains
  - Simpler controllers
  - Easier to reason about
- Support multiple operational teams with local control
  - E.g., security team and optimization team
- IBN Function defined

Document Status

- Contributions from multiple authors
- Thorough review/contributions by several individuals
- All received comments have been addressed
- We are working to better describe mechanisms
- Would like the working group to adopt
  - To inform or standardize IBN behaviors