

Hierarchical Service Function Chaining (hSFC)

draft-dolson-sfc-hierarchical-02

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 expanding to a separate draft
- Initial Draft posted May 25, proposing some mechanisms
- 2 updates due to feedback
- Use cases contributed in draft-liu-sfc-nesting-use-case-01, moving here

The Problem

- SFC controller complexity in very large networks.
 - Millions of hosts
 - Thousands of forwarding elements
 - Asymmetrical routing
 - Multiple operational teams
- How to avoid a “super controller” ?

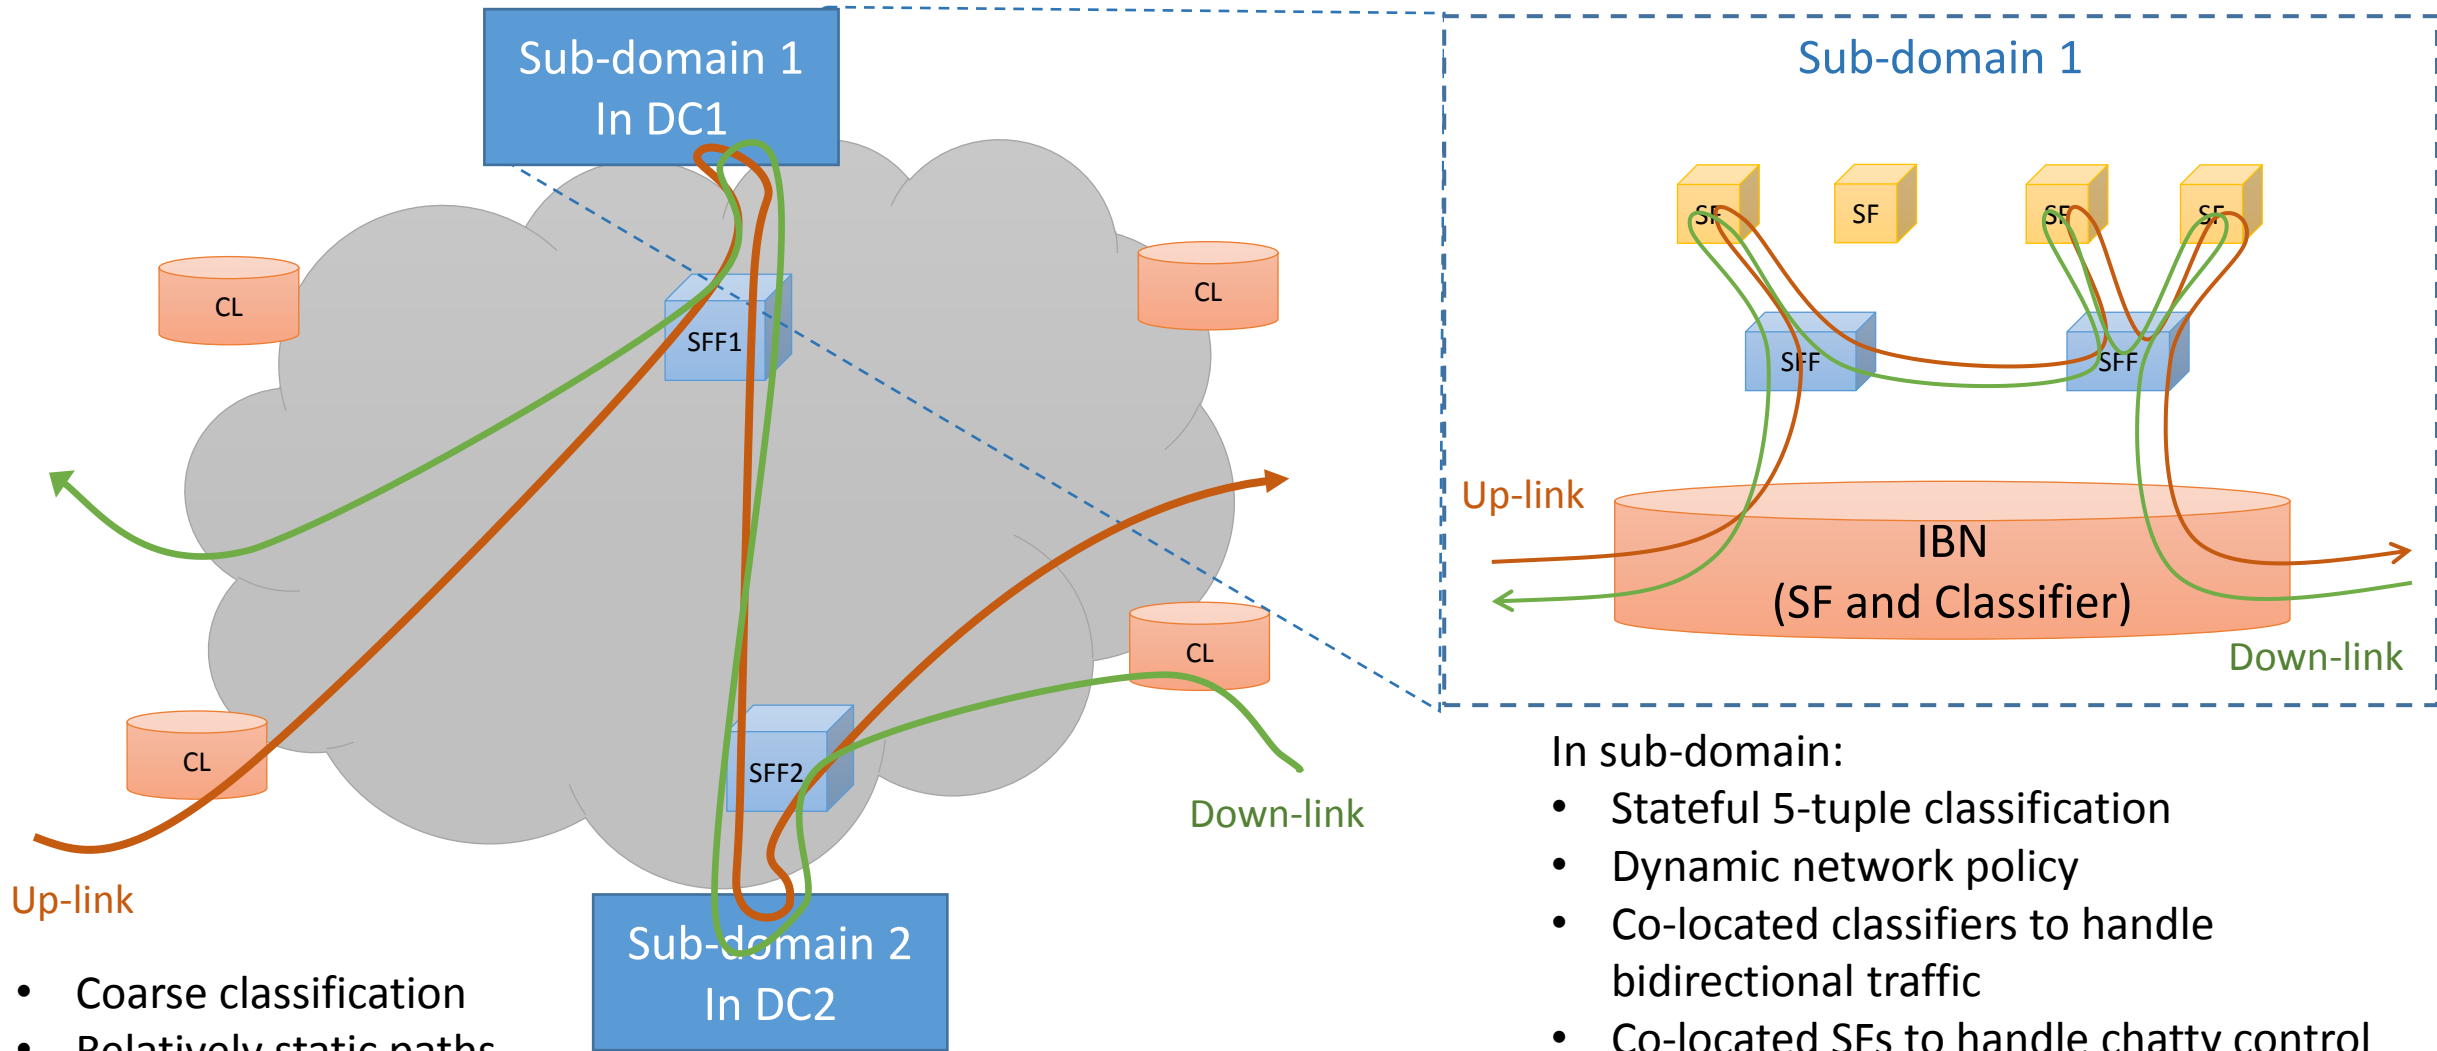
The Solution:

- Hierarchy of control

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.

Internal Boundary Node (IBN)

- (Name changes “SF Domain Proxy” → “SF Domain Gateway” → “Internal Boundary Node”)
- This is the module that links the domains
- Looks like an SF to top level
- Looks like classifier and end-of-chain to low level

Benefits:

- Avoid costly stateful classification at distributed classifiers
 - Scales to very large networks
 - Supports specialized sub-domains with local control (e.g., per tenant)
- Inclusion of IBN in SFC Architecture?

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

Control Plane Implications

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

- Very precise control interfaces for each function would benefit hSFC
 - (vs., for example, a control element wanting to own the entire node)

Metadata Implications

- What happens to metadata in the higher-level domain when packets are forwarded into the lower-level domain?
 - Share a single metadata schema across domains?
 - Pushing/popping/mapping?

Next Steps

- We would like to see working group adoption, to inform discussions about meta-data and control-plane.
- Enhance section on examples that clearly benefit from hSFC.