

Hierarchical Service Function Chaining (hSFC)

draft-dolson-sfc-hierarchical-05

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)

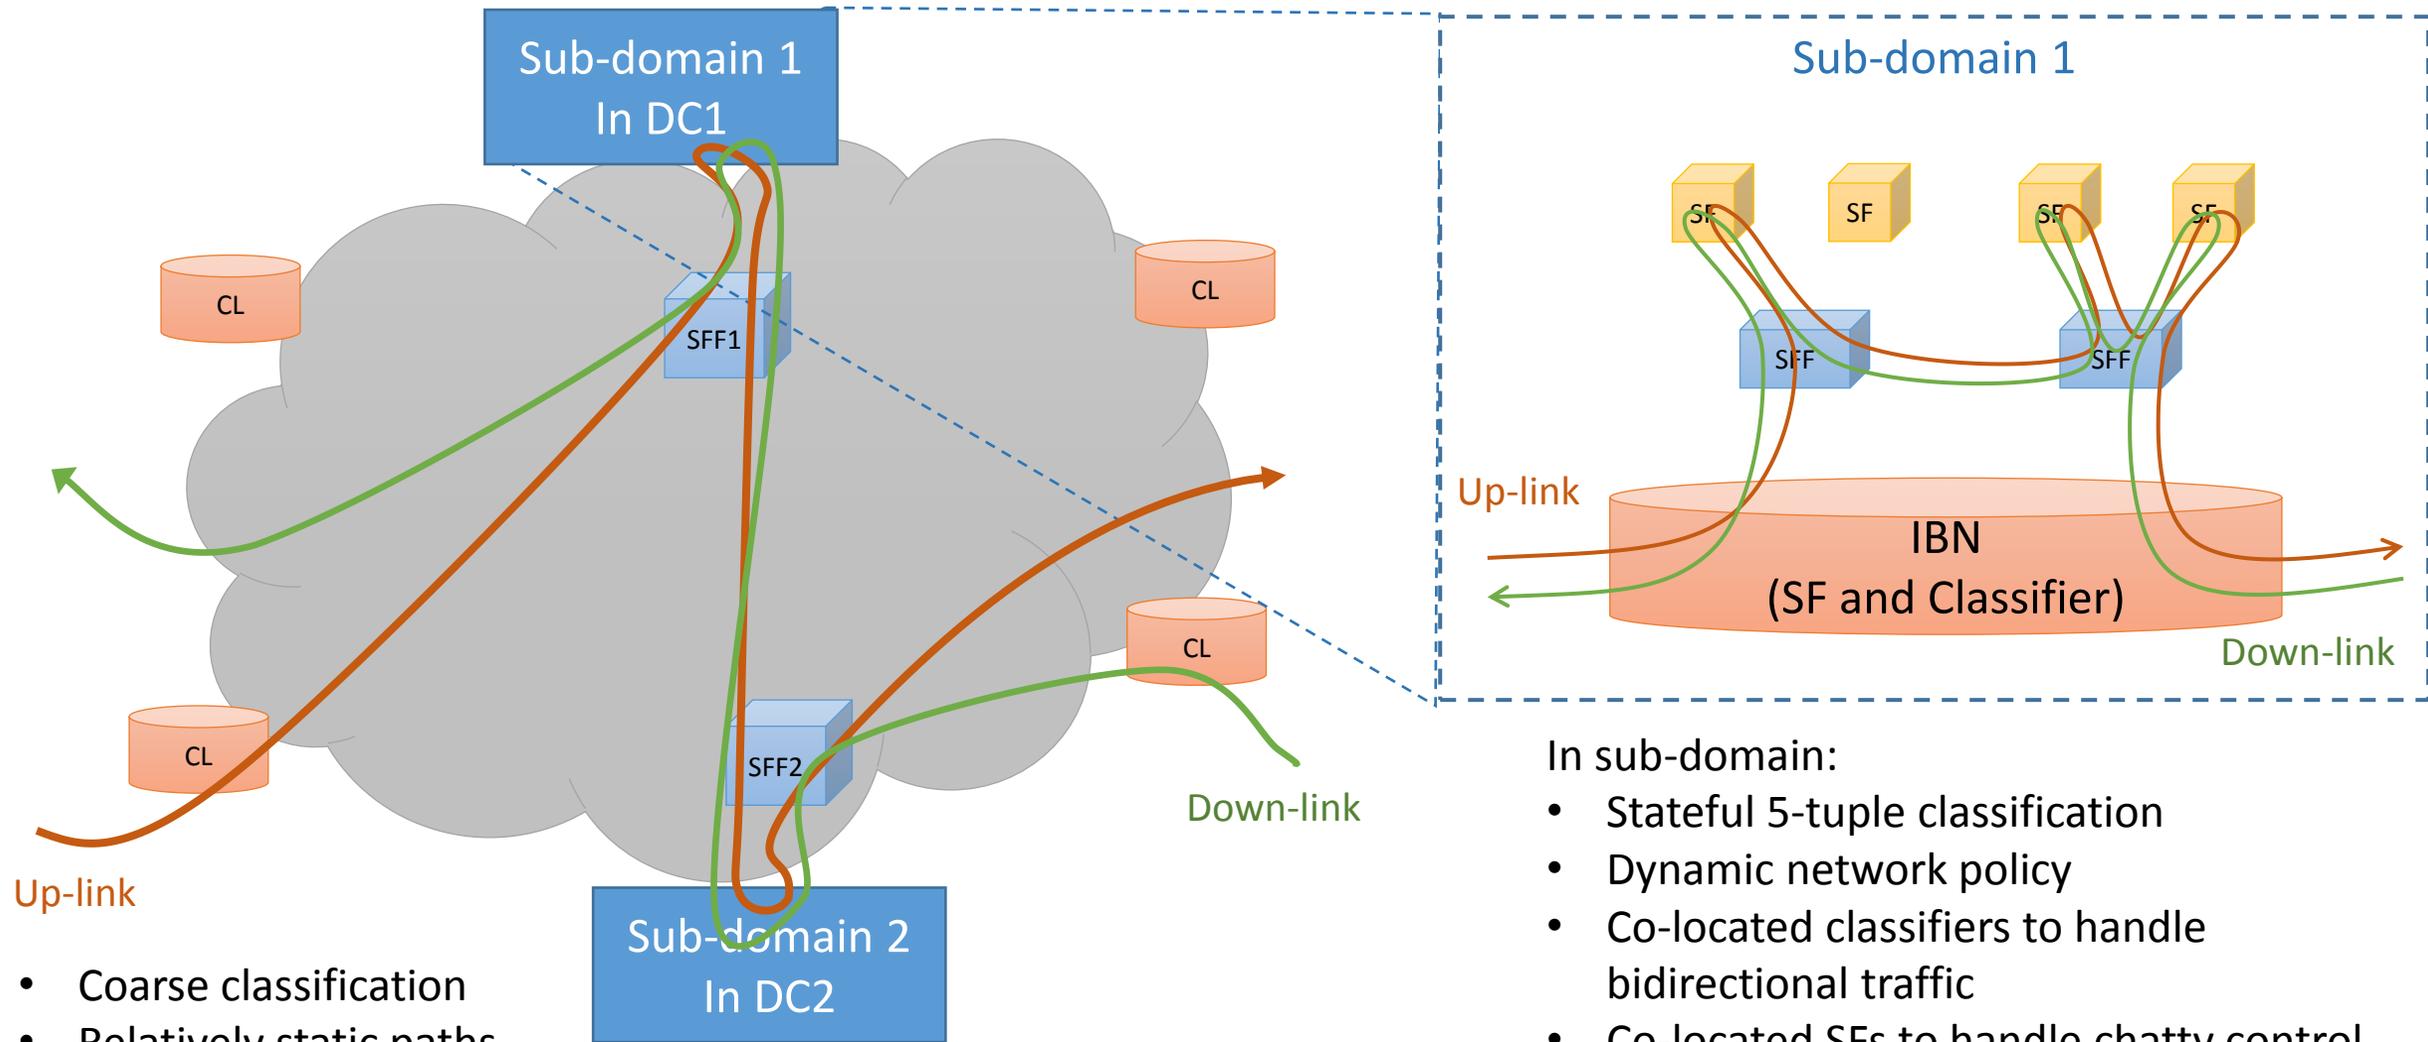
Dapeng Liu (max.ldp@alibaba-inc.com)

Ting Ao (ao.ting@zte.com.cn)

History

- Concept introduced in draft-homma-sfc-forwarding-methods-analysis
- At IETF 92 (Dallas), there was interest in creating a separate draft
- Initial Draft posted May 2015, proposing some mechanisms
- Version -02 presented at IETF 93 (Prague)
- Version -03 presented at IETF 94 (Yokohama)
- Incorporated contributions from
 - draft-liu-sfc-nesting-use-case-01
 - draft-ao-sfc-for-dc-interconnect-01
- Current version -05 posted March 7, 2016

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.

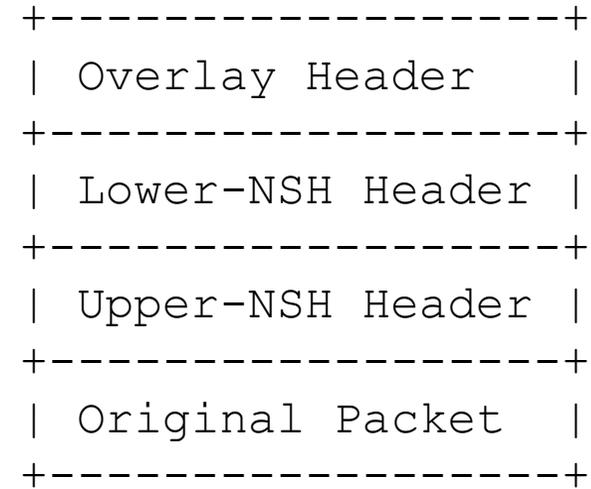
The Internal Boundary Node (IBN)

- We propose the IBN to bridge levels of hierarchy within a single administrative domain
 - A variant of “SFC Boundary Node” (per RFC7665)
 - 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
 - **Nesting upper-level NSH within lower-level NSH** ****New in -05****

Nesting NSH



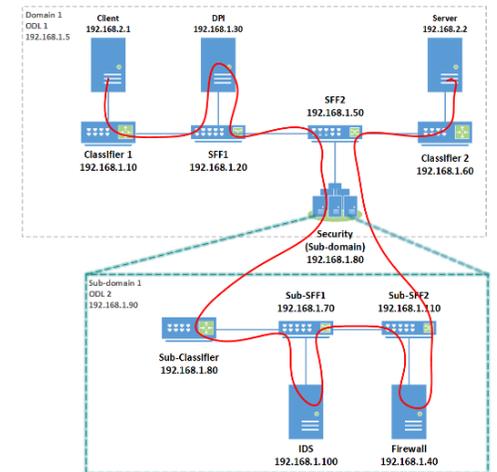
- Simplifies IBN responsibilities at the expense of complexity in the SF
- Requires a new allocation allowing NSH as “Next protocol” in the lower-NSH header.

Implementation

- Victor Vu has put together an implementation using OpenDaylight
- Method: encoding upper-level paths as context meta-data in the lower-level metadata
- And a presentation
 - <https://goo.gl/EfyOmE>

Demo

- Re-use sfc-103 demo topology and configurations
- Total 10 vagrant VMs. For clarification, 1 VM for each CF, SFF, SF



Control Plane Implications

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

- Control-plane standards should permit hSFC

Metadata Implications

- If in use, 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

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