Service Function Chaining: Framework & Architecture

draft-boucadair-sfc-framework

London, March 2014

M. Boucadair (mohamed.boucadair@orange.com)
C. Jacquenet (christian.jacquenet@orange.com)
R. Parker (Ron_Parker@affirmednetworks.com)
D. Lopez (diego@tid.es)
J. Guichard (jguichar@cisco.com)
C. Pignataro (cpignata@cisco.com)

Context

- IP networks rely more and more on the combination of advanced functions
  - Besides basic routing and forwarding functions
- The goal is to enforce service-inferred forwarding for traffic traversing a given domain
  - Differentiated by the set of Service Functions to be invoked
  - Service-inferred forwarding is policy-based. Policies may be:
    - subscriber-aware
    - based on flow characteristics
    - TE-oriented (e.g., optimize network resource usage)
    - combination thereof

Main Principles

- SFC provides a new forwarding paradigm to process traffic according to an ordered set of Service Functions
- Dynamic SFC provisioning is a different operation than packet processing
- Service Functions are viewed as logical instances
  - The set of SFs is defined by the service to be provided and according to the networking environment
  - There is no global nor standard list of SFs enabled
- SF chaining adapts to the service and the traffic directionality
  - Transparent to communication endpoints
  - Policy-driven
  - Network transport agnostic
  - Supports sharing of information between SFs

All is about Policies

- SFC logic is domain-specific and policy-driven
  - No global or standard SF chaining logic
  - The ordered set of activated SFs is specific to each administrative domain
  - The chaining of SFs and the criteria to invoke some of them are local to a domain
- Preserve current engineering practices
  - Topology hiding
  - SF chaining logic and related policies should not be leaked outside of the SFC domain
- Several SF chains can be simultaneously enforced within a domain
  - To meet various business requirements
- How to bind the traffic to a given SF chaining is policy-based

Overall Approach

- SFC operations are abstracted
- No implementation-specific details are included
  - No assumptions on how policies are enforced
    - How SF-specific policies are enforced is out of scope
  - Traffic forwarding between SFC-aware SFs can rely on IP-based tunnels, MAC-in-MAC, etc.
- Fragmentation handling
  - A stateless approach is described
  - MTU tweaking can be envisaged
- Compliance with DiffServ and ECN

Supported Features

- Core functionalities
  - SFC structuring
  - SFC-related policy enforcement
- Additional functionalities
  - Control plane
    - SF discovery
    - SF/SFC diagnostic
    - SF liveliness detection
  - SF loop detection
  - Overhead optimization

Functional Elements

- SFC Boundary Node: A node that connects one SFC-enabled domain to a another SFC-enabled domain or an is SFC-unaware domain
  - SFC Ingress Node: A SFC Boundary Node that handles traffic which enters the SFC-enabled domain
  - SFC Egress Node: A SFC Boundary Node that handles traffic which leaves the SFC-enabled domain
- SF Node: Any node within an SFC-enabled domain that embeds one or multiple SFs
- SFC Classifier: An entity which classifies packets based on (some of) their contents, and optionally local policies (i.e., subscriber aware and flow aware)
- PDP (Policy Decision Point): An entity which is responsible for structuring SFCs and communicate policies to other involved functional elements

Co-locating SFC Functional Elements

- It is deployment-specific
- Operators can make their own design choices, e.g.,
  1. A BN can act as Egress BN and Ingress BN for the same flow
  2. Distinct Ingress BN and Egress BN may be crossed by a packet
  3. Distinct Ingress BNs for upstream and downstream
  4. An Ingress BN can embed a Classifier
  5. An Ingress BN may not embed a Classifier, but be responsible for dispatching flows among a set of Classifiers
  6. Ingress BN, Egress BN, and Classifier may be co-located
  7. A PDP may be co-located with a Classifier
  8. An SF Node may be co-located with a Classifier
  9. Many Network Elements may behave as xBNs

The Intelligence Resides In The PDP

- PDP makes decisions according to policies documented in SFC Policy Tables
  - PDP decisions are applied by SF Nodes, Classifier, and Boundary Nodes which process traffic accordingly
  - A PDP may manage multiple SFC domains

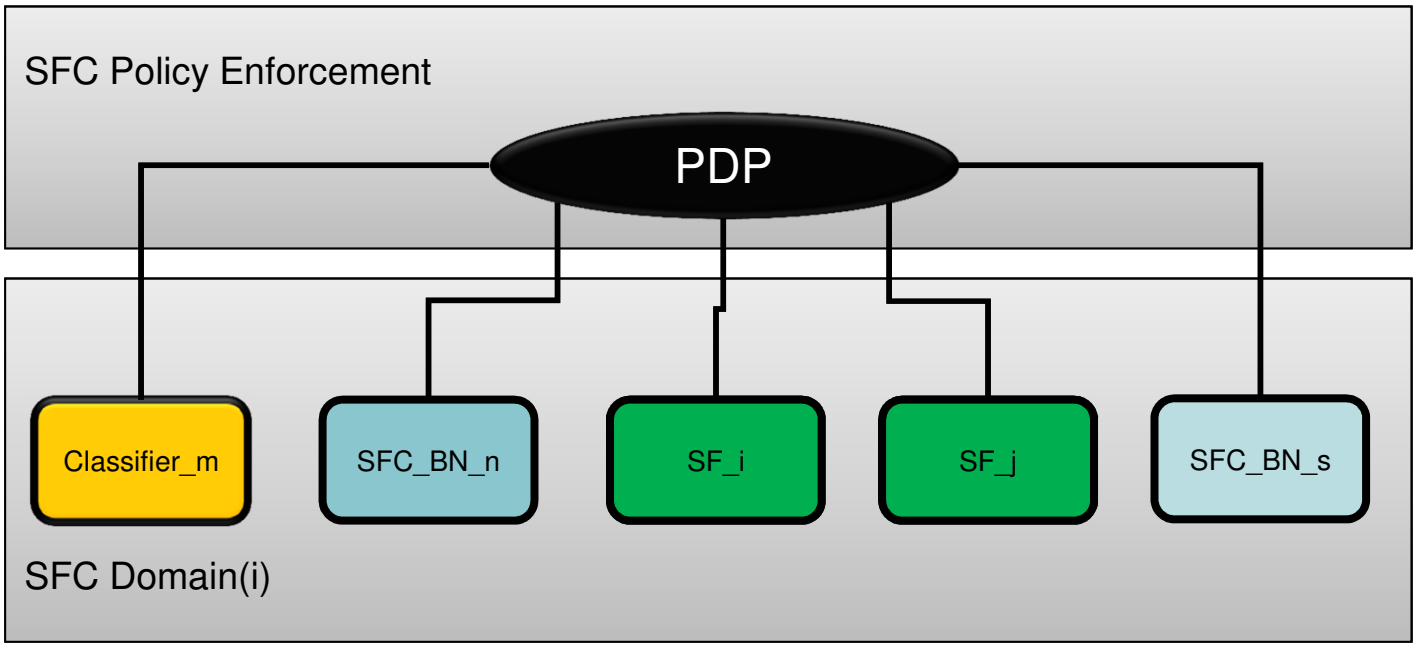

Main SFC Operations

- Assign SF Identifiers
  - SFs are listed and identified in a repository maintained by the SFC administrative entity (*e.g.*, ISP)
- Assign SF Locators
  - Meant to locate a SF which can be supported by several devices
    - Locator is typically an IP address (could be a FQDN, MAC @, etc.)
    - One or multiple Locators can be configured for each SF (*e.g.*, 1.2.3.4, 2001:db8::1)
- Build SF Chains
  - Detail the list of SFs to be invoked in a specific order
  - Abstracted as “SF Map”
    - SF Map is an ordered list of “SF Identifiers”
    - SF Maps are specific to traffic directionality
- Coordination between PDP, Classifiers, and SFs

Deployment Considerations

- The document identifies and discusses deployment models where:
  1. A marking-based scheme is used together with an encapsulation mode
  2. No encapsulation nor tagging mechanisms are needed
  3. Marking can be used but no encapsulation is required
  4. No protocol extension is required

Focus on the Marking-Based Scheme

- Chaining is described by an information processed by devices that participate to the delivery of a given service
- Such information needs to be signaled in data packets themselves
- Some open issues are discussed
  - What marking information is required to be signaled?
  - Where to inject such marking information?
  - How to steer traffic forwarding to cross specific SF Nodes?
- A preliminary analysis is included in the I-D, while a detailed analysis is recorded in I-D.boucadair-sfc-design-analysis

Next Step?

- Adopt as a Working Group document to fulfill this charter item:

"*Architecture: This document will provide a description of the SFC architectural building blocks and their relationships including interconnection, placement of SFC specific capabilities, management, diagnostics, design analysis, and security models, as well as requirements on the protocol mechanisms. The initial scope is restricted to a single administrative domain, keeping in mind that architectural decisions made for the intra-domain case may have implications for the inter-domain case.*"

- Incorporate concepts from other proposed I-Ds on architectural aspects
- Comments and contributions are welcome