Overlay OAM Requirements

draft-ooamdt-rtgwg-ooam-requirement

Greg Mirsky gregory.mirsky@ericsson.com
Carlos Pignataro cpignata@cisco.com
Nagendra Kumar naikumar@cisco.com
Deepak Kumar dekumar@cisco.com

Erik Nordmark nordmark@acm.org
David Mozes davidm@mellanox.com
Mach Chen mach.chen@huawei.com
Santosh Pallagatti santosh.pallagatti@gmail.com

IETF-95 April, 2016

Motivation

- Work in NVO3, SFC and BIER WGs on OAM Framework and Requirements drafts
- Proposed adaptation, extension of existing OAM protocols (BFD in VXLAN)
- Proposed new mechanism (Transcending Traceroute)
- Common OAM presentation and discussion at IETF-94

Structure

- Fault Management
  - Proactive FM
  - On-demand FM
- Performance Management
- Alarm Indication Signal (Supression)
- Resiliency

Requirements

- OOAM independent from a transport layer
- Any node implicitly serves as MEP
- SDN-azation of Overlay OAM
- Proactive and on-demand OAM created equal
- Unidirectional Overlay OAM (CC and PM) optimization as services (multicast, SFC) are unidirectional
- OAM is about what is going in the transport layer and thus it must be in-band , i.e. fate sharing with data traffic
- Bi-directional OAM is important too, e.g. CC-CV and out-of-band Fault Management Signal

Fault Management

- Proactive
  - Continuity Check
  - Remote Defect Indication
  - Connectivity Verification
- On-demand
  - LoC defect localization
  - path tracing through overlay network
  - verification of mapping between overlay network and client layer services
  - ECMP discovery and verification
  - proxy ping/traceroute
- Fault Management Signals like Alarm Indication Signal to suppress client layer alarms when server layer fault detected
- Overlay network survivability may use protection switching and restoration

Performance Measurements

- Passive and Active Performance Measurement OAM are complimentary instruments in OOAM toolbox
  - One-way active and passive
  - Two-way active
- Support calculation of performance metrics:
  - packet delay
  - packet delay variation
  - packet loss
  - goodput (delivered throughput)

BFD in BIER

based on draft-ooamdt-rtgwg-gap-analysis-00

Greg Mirsky gregory.mirsky@ericsson.com
Erik Nordmark nordmark@acm.org
Carlos Pignataro cpignata@cisco.com
Nagendra Kumar Nainar naikumar@cisco.com

Deepak Kumar dekumar@cisco.com
Mach Chen mach.chen@huawei.com
Santosh Pallagatti santoshpalagatti@gmail.com
David Mozes davidm@mellanox.com

IETF-95 April, 2016, Buenos Aires

Motivation

- Overlay OAM Design Team is working on OOAM Requirements and OOAM Gap Analysis
- The gap analysis covers OAM tools to enable Fault Management and Performance Measurement:
  - Proactive continuity check
  - On-demand continuity check and connectivity verification
  - Active and passive performance measurement
  - Alarm Indication Signal (Suppression)

BFD with IP/UDP

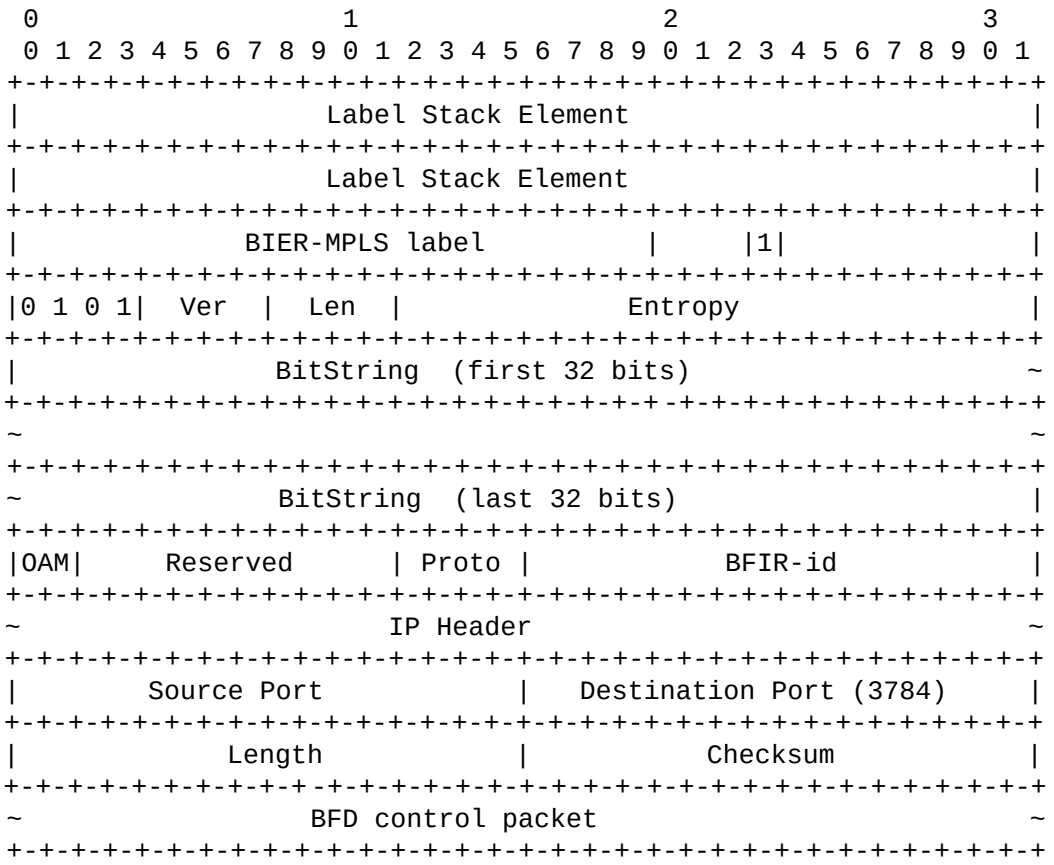

BFD in G-ACh

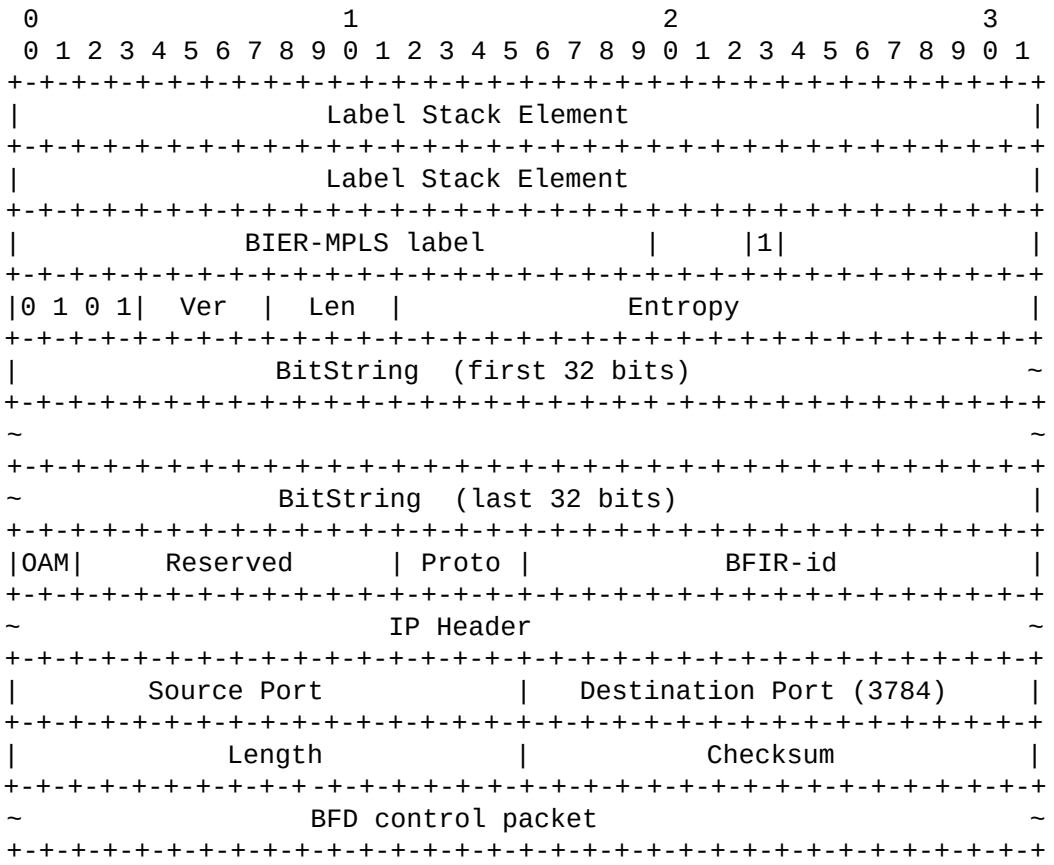
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Label Stack Element
Label Stack Element
BIER-MPLS label
1
0 1 0 1
Ver
Len
Entropy
BitString (first 32 bits)
BitString (last 32 bits)
OAM
Reserved
Proto
BFIR-id
IP Header
Source Port
Destination Port (3784)
Length
Checksum
BFD control packet

BFD for SFC

based on draft-ooamdt-rtgwg-gap-analysis-00

Greg Mirsky gregory.mirsky@ericsson.com
Erik Nordmark nordmark@acm.org
Carlos Pignataro cpignata@cisco.com
Nagendra Kumar Nainar naikumar@cisco.com
Deepak Nagendra dekumar@cisco.com
Mach Chen mach.chen@huawei.com
Santosh Pallagatti santoshpalagatti@gmail.com
David Mozes davidm@mellanox.com

IETF-95 April, 2016, Buenos Aires

Use Case

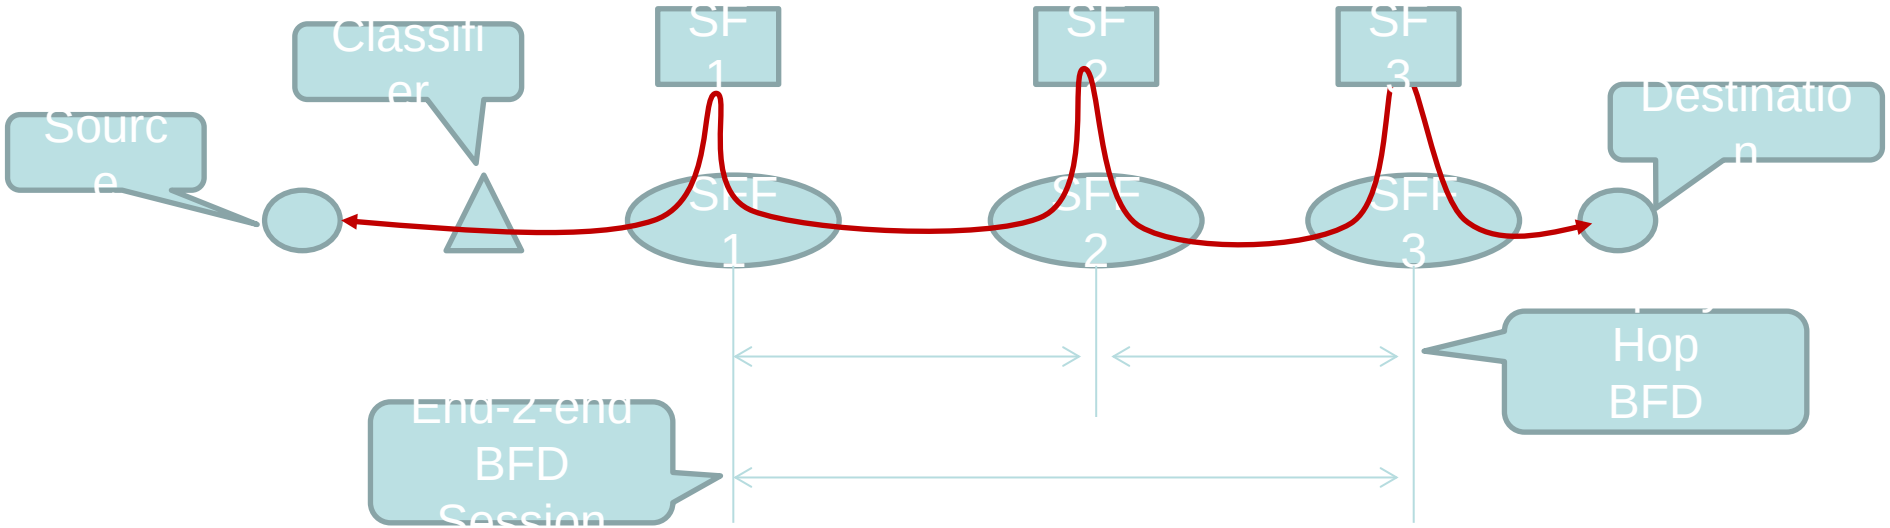

- Where to deploy the BFD sessions?
  - Between SFFs (the major case)
    - Hop-by-hop BFD session (e.g., SFF1<->SFF2, SFF2<->SFF3)
    - End-2-end BFD session (e.g., SFF1<->SFF3)
  - Other possibilities ?

Hop-by-hop Case

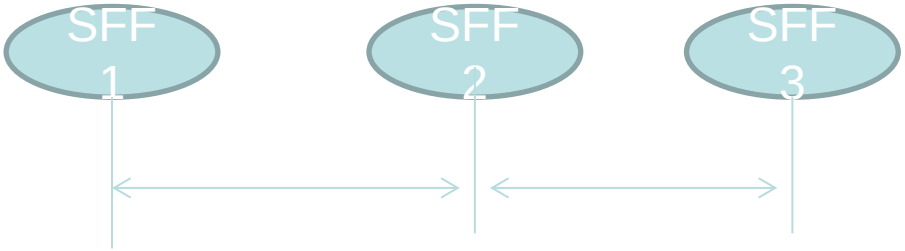

- An SFF should have capability to determine whether a packet should be delivered to an SF or terminated.
- Encapsulation dependent

End-2-End Case

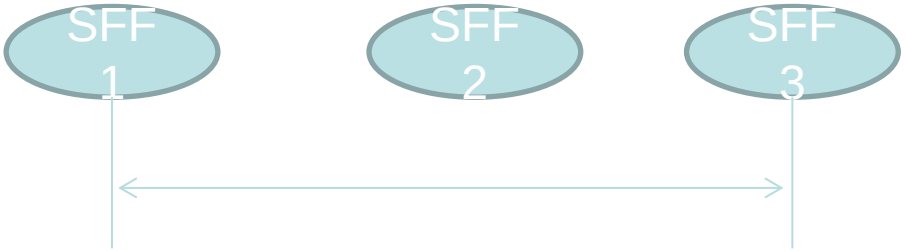

- An SFF (e.g., SFF2) should have capability to determine whether a packet should be delivered to an SF or the next hop SFF.
- Encapsulation dependent

Control Plane

- BFD session bootstrapping
  - In-band signaling
  - Out-of-band channel
  - Centralized controller

Encapsulations

- BFD with IP/UDP encapsulation
  - Same as RFC5881 and 5884
  - The source/destination addresses and UDP port are derived from the IP/UDP header
- BFD without IP/UDP encapsulation
  - Add source and destination addresses field
  - UDP port is not necessary, the “Next Protocol” and/or “type” fields can be used to indicate a BFD packet
- BFD with embedded Src/Dst Info
  - Source and destination address are embedded in the BFD control packet
  - Similar to RFC6428, e.g., Source MEP ID TLV

Next steps

- Update the draft to cover following aspects of the gaps
  - Use Case
  - Control plane
  - Data plane/Encapsulations