DETERMINISTIC6G

CFM for FRER

János Farkas, Balázs Varga, Ferenc Fejes

janos.farkas@ericsson.com, balazs.a.varga@ericsson.com, ferenc.fejes@ericsson.com

DETERMINISTIC6G has received funding from the European Union’s Horizon Europe programme under grant agreement No 101096504.

Recap: Frame Replication and Elimination for Reliability (FRER)

- FRER provides great service protection
  - Stream survives as long as at least one of the redundant paths survive
  - Note: the maximally disjoint explicit paths are typically Traffic Engineered (TE) paths

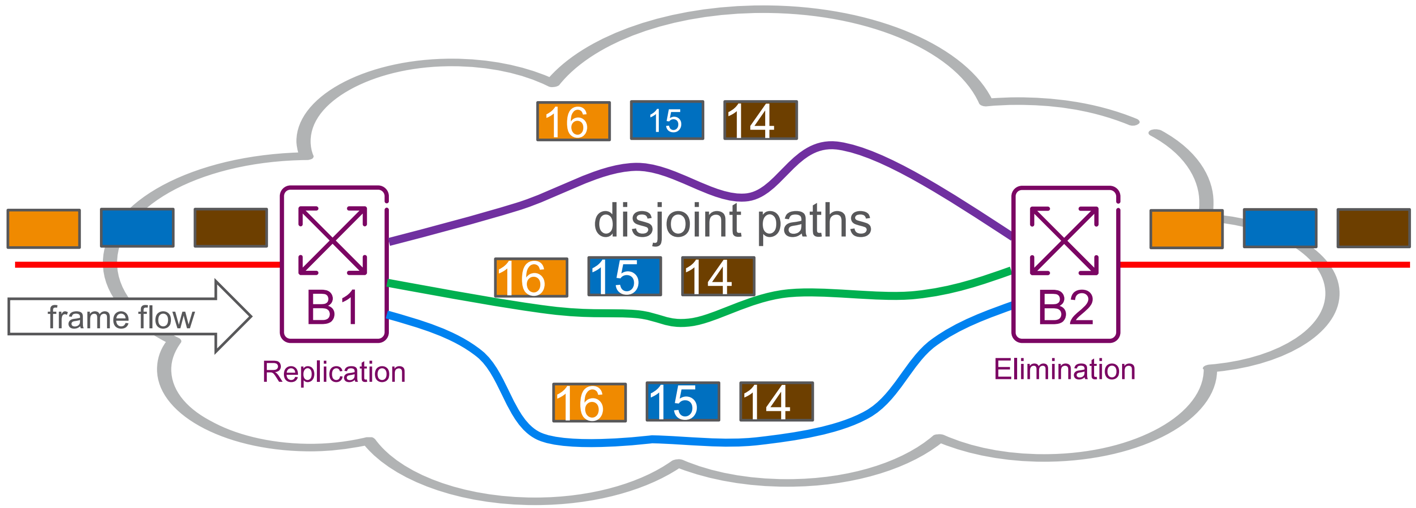

- Some are concerned that it is not known what is going on in the network as FRER hides failures
- FRER has in-built counters, which provide some node-local warnings
  - However, network-wide full picture cannot be derived from FRER counters

Recap: Connectivity Fault Management (CFM)

- Operations, Administration, and Maintenance (OAM):
  - The goal of OAM is to provide tools that make it possible to know what is going on in the network
  - The key task to be solved is fate sharing between data packets and OAM packets
- Based on https://www.ieee802.org/1/files/public/docs2021/60802-finn-intro-to-CFM-0721-v02.pdf:
- CFM provides OAM functions for VLAN bridged Ethernet networks
  - CFM was originally designed for service instances distinguished by their VID
- Two independent standards developed in collaboration
  - IEEE Std 802.1ag “Connectivity Fault Management” → Clauses 18-22 in IEEE Std 802.1Q
    - Specifies Fault Management (FM) functions
  - ITU G.8013/Y.1731 “OAM functions and mechanisms for Ethernet-based networks”
    - Specifies Performance Monitoring (PM) (e.g., loss measurement and delay measurement) on top of FM, i.e., ITU-T G.8013/Y.1731 is a superset
- Note: For simplicity, this presentation deliberately does not go into the details of Maintenance Domains (MDs), Maintenance Domain Levels (MDLs), and Maintenance Associations (MAs)

Recap: CFM Protocols

- CFM functions include path discovery, fault detection, fault notification, fault verification and isolation
- Maintenance Points (MPs), i.e., Maintenance End Points (MEPs) and Maintenance Intermediate Points (MIPs) can participate in the CFM protocols:
- Continuity Check protocol
  - Continuity Check Messages (CCMs) are transmitted periodically for fault detection, i.e., to detect both connectivity failures and unintended connectivity
  - A CCM does not generate a response
- Loopback protocol
  - The Loopback protocol includes the Loopback Message (LBM) and the Loopback Reply (LBR), which can be used to perform fault verification and isolation
- Linktrace protocol
  - The Loopback protocol includes the Linktrace Message (LTM) and the Linktrace Reply (LTR), which can be used to perform path discovery and fault isolation

Recap: CFM PDUs

- CFM PDUs can be identified by the EtherType value 89-02, see Table 21-1 in IEEE Std 802.1Q-2020
- Destination Address (DA)
  - For monitoring a service instance distinguished by its VID
    - CFM entities use the group MAC addresses for CCM and LTM PDUs listed in Table 8-18 and Table 8-19 in IEEE Std 802.1Q-2020
  - For monitoring a TE service
    - CFM entities use the individual MAC addresses or the group MAC addresses that are associated with the monitored services
- Source Address (SA)
  - The individual MAC address of the MP transmitting the CFM PDU

Recap: Traffic Engineering (TE) / Software Defined Networking (SDN)

- Traffic Engineering (TE) / Software Defined Networking (SDN) has been introduced to IEEE Std 802.1Q by IEEE Std 802.1Qay-2009 Provider Backbone Bridge Traffic Engineering (PBB-TE) (see also a paper on it)
- Introduces concepts and knobs for Traffic Engineering, e.g., TE-MSTID
- Extends CFM for Traffic Engineering
  - (CFM was originally designed for service instances distinguished by their VID)
- Although, the scope was Provider Backbone Bridge (PBB) networks, the introduced Traffic Engineering capabilities are generic
  - Some minor updates might be required for applicability in C-VLANs or S-VLANs

DETERMINISTIC6G

Recap: Some Definitions Introduced by 802.1Qay

- **Ethernet Switched Path (ESP)**: provisioned traffic engineered unidirectional connectivity path among two or more Customer Backbone Ports (CBPs) that extends over a Provider Backbone Bridged Network (PBBN). The path is identified by a 3-tuple <ESP-DA, ESP-SA, ESP-VID>, where ESP-DA and ESP-SA are Media Access Control (MAC) addresses and ESP-VID is a Virtual Local Area Network (VLAN) Identifier (VID) allocated to Traffic Engineering Multiple Spanning Tree Instance Identifier (TE-MSTID).
- **Ethernet Switched Path [ESP] Virtual Local Area Network [VLAN] Identifier [VID] (ESP-VID)**: A VLAN Identifier (VID) associated with a special (Traffic Engineering) value of the Multiple Spanning Tree Instance Identifier (MSTID) in the Multiple Spanning Tree (MST) Configuration Table, the TE-MSTID, indicating that the VID is under the control of an external agent responsible for setting up ESPs and that learning is disabled and forwarding is enabled.
- **Traffic Engineering service instance Identifier (TE-SID)**: An identifier of the Traffic Engineering service instance (TESI) that corresponds to a series of 3-tuples <ESP-DA, ESP-SA, ESP-VID>, each one identifying one of the TESI's Ethernet Switched Paths (ESPs).

Recap: CFM for TE Paths

- ESP: unidirectional traffic engineered path
  - Identified by: <DA, SA, VID> 3-tuple
- MEPs use the individual MAC address of the port on which they operate
- No requirements on using the same or different VIDs
- New TLV introduced for Response messages (LBR and LTR) due to the unidirectional paths (see 21.7.5 in IEEE Std 802.1Q-2020)

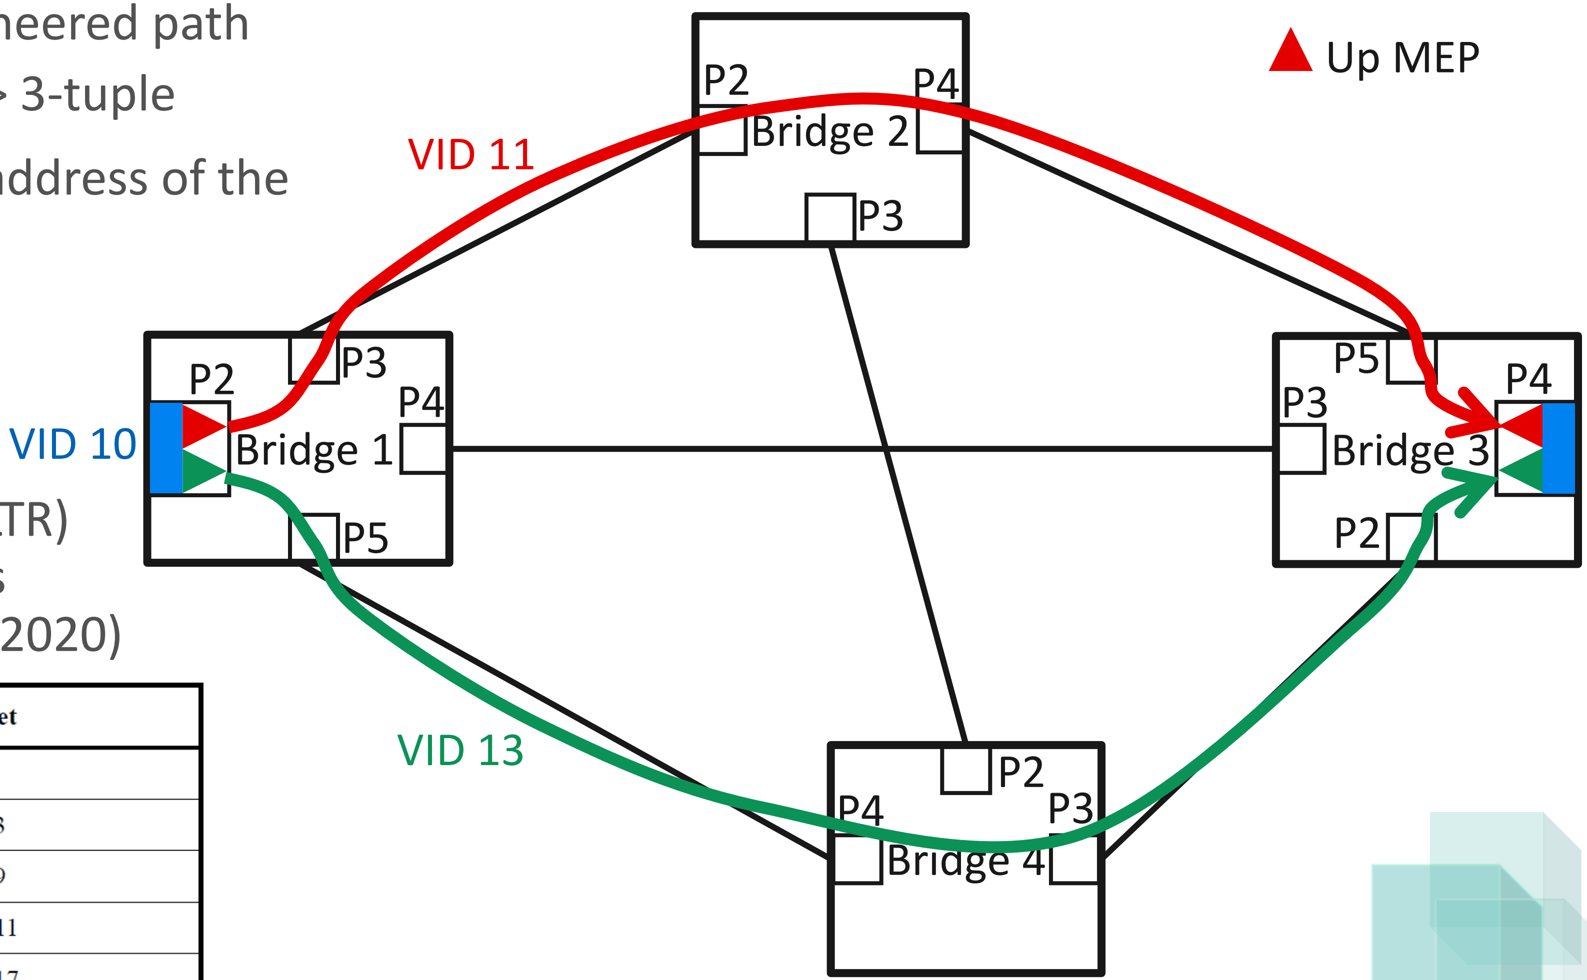

Field	Octet
Type = 9	1
Length	2–3
MIP MAC address	4–9
Reverse VID	10–11
Reverse MAC	12–17

What Is New in TSN?

- Some TSN tools operate on per-Stream basis, e.g., FRER
- Stream identification
  see Clause 6 in
  IEEE Std 802.1CB-2017
  as amended by
  IEEE Std 802.1CBdb-2021

Stream identification function	Active/passive	Examines	Overwrites	Reference
Null Stream identification	Passive	destination_address, vlan_identifier	None	6.4, 9.1.2
Source MAC and VLAN Stream identification	Passive	source_address, vlan_identifier	None	6.5, 9.1.3
Active Destination MAC and VLAN Stream identification	Active	destination_address, vlan_identifier	destination_address, vlan_identifier, priority	6.6, 9.1.4
IP Stream identification	Passive	destination_address, vlan_identifier, IP source address, IP destination address, DSCP, IP next protocol, source port, destination port	None	6.7, 9.1.5
Mask-and-match Stream identification function	Passive	destination_address, source_address, mac_service_data_unit	None	6.8, 9.1.6

- Note that traffic engineering, unidirectional communication, and the use of a different path for a response message are not new

Consequence of Per-Stream Operation

- Reminder: fate sharing between data packets and OAM packets is essential
- Consequently, the header fields used for packet processing have to be the same for OAM and data packets
- TE path (solid line) vs Stream (dashed line)
  - TE path identified by: <B3.P4, B1.P2, VID 11>
  - Stream identified by:
    - <L, VID 11>, or
    - <T, VID 11>,
    - etc.

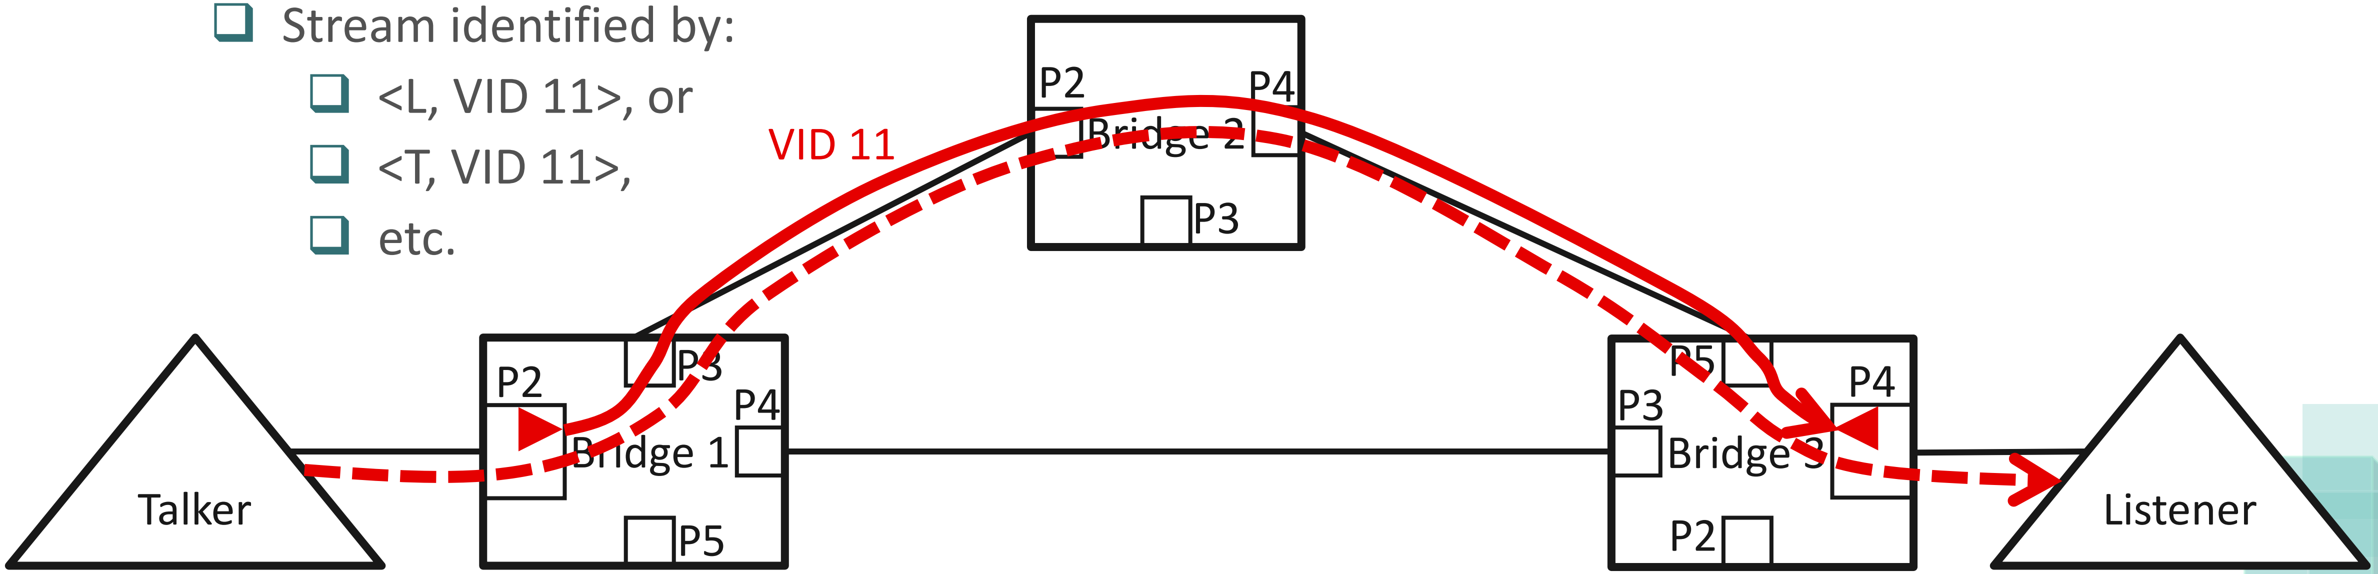

MEP Operation

- MEP operation has to be updated for CFM for Streams
- In case of Null Stream identification
  - The DA in CCM PDUs has to be the DA of the Listener to achieve fate sharing
  - MEPs have to be able use and operate on the DA of end stations
  - The DAs for the monitored Streams are to be configured at MEPs

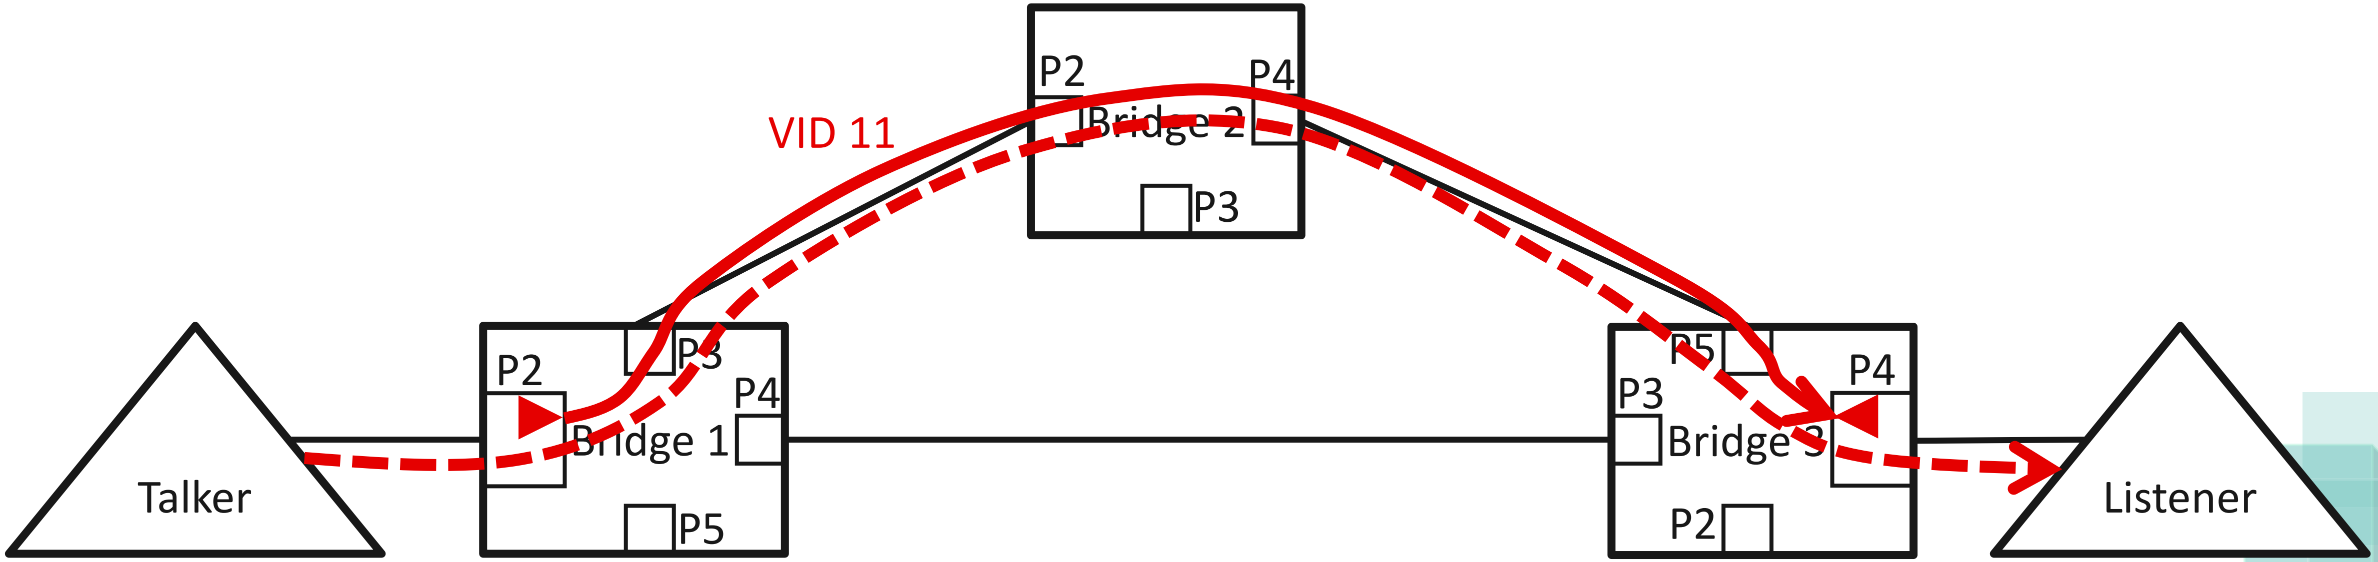

DETERMINISTIC6G

OAM for FRER?

- OAM covering FRER would provide the knowledge of what is going on in the network (despite of the great service protection provided by FRER)
- How to achieve fate sharing in case of FRER?
- Primarily, OAM packets have to be R-tagged and have to go through the Replication and Elimination functions just like data packets
  - OAM packets are distinguished by EtherType
  - OAM packets need to have their own sequence number space distinct from the sequence number space of data packets
  - Sequence recovery can be simple, i.e., the MatchRecoveryAlgorithm can be used for infrequent OAM packets
- No new header field needed for OAM
- However, we are not there yet; see next slides

CFM for Replication and Elimination Functions

- OAM has to cover Replication (R) and Elimination (E) functions
  - MEP before R
    - MEP generates CCMs
    - CCMs' DA and VID are identical to data packets
    - CCMs go through R
  - MEP after E
    - CCMs go through E before MEP receives them

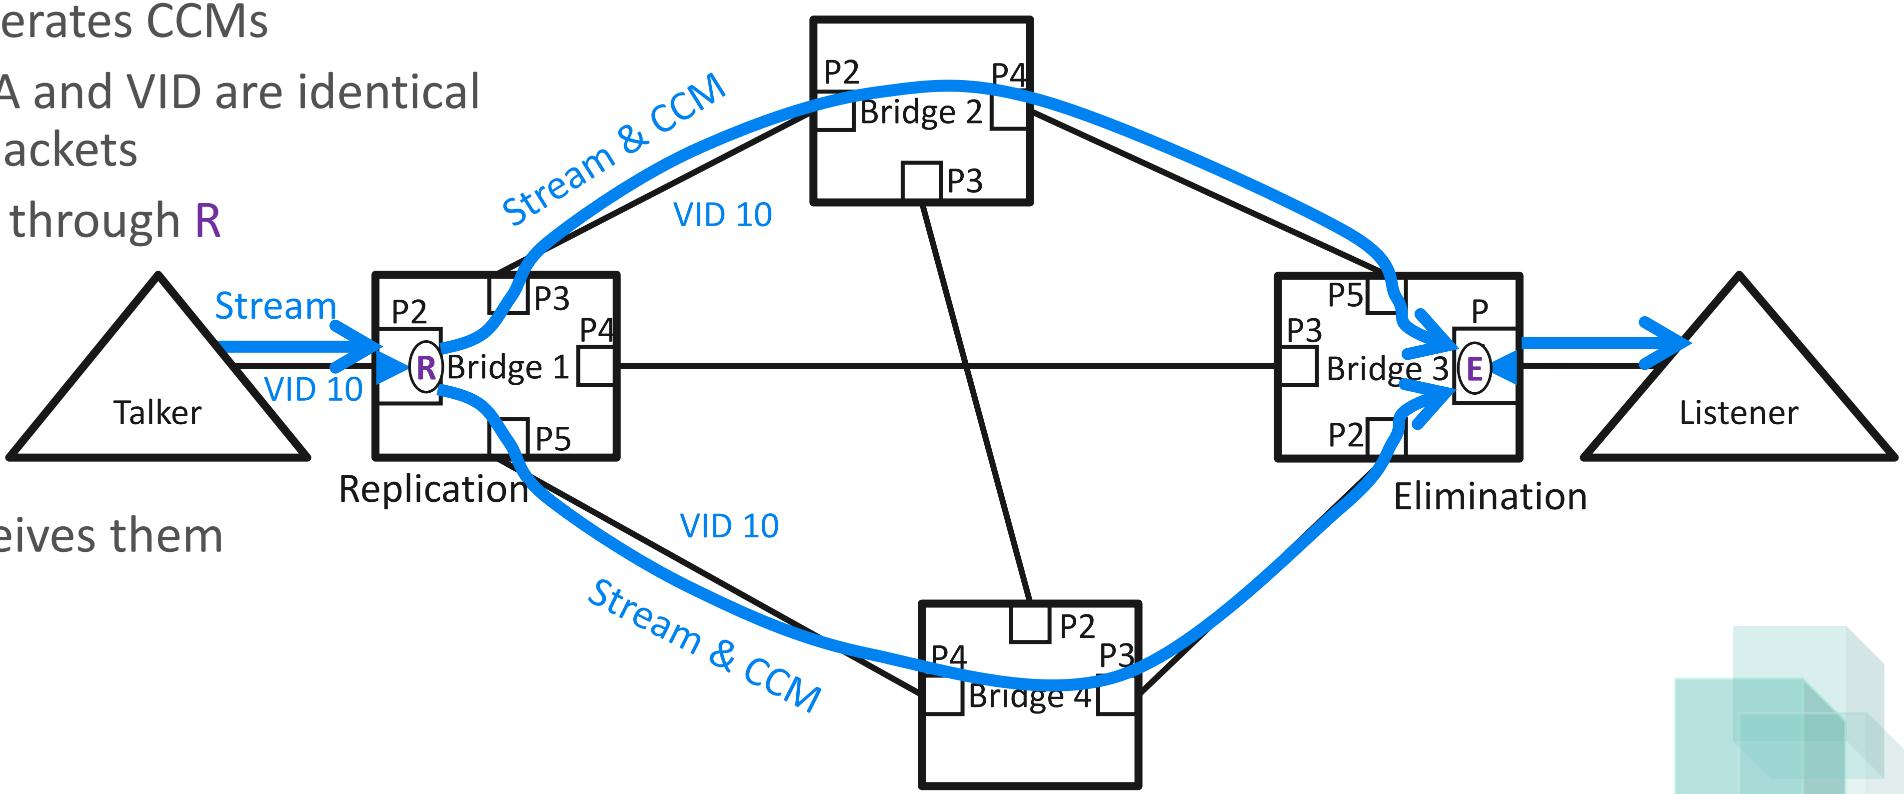

CFM for R+E Functions and TE Paths

- Blue MEPs cover the Replication (R) and Elimination (E) functions
- Red and Green MEPs cover the corresponding TE paths

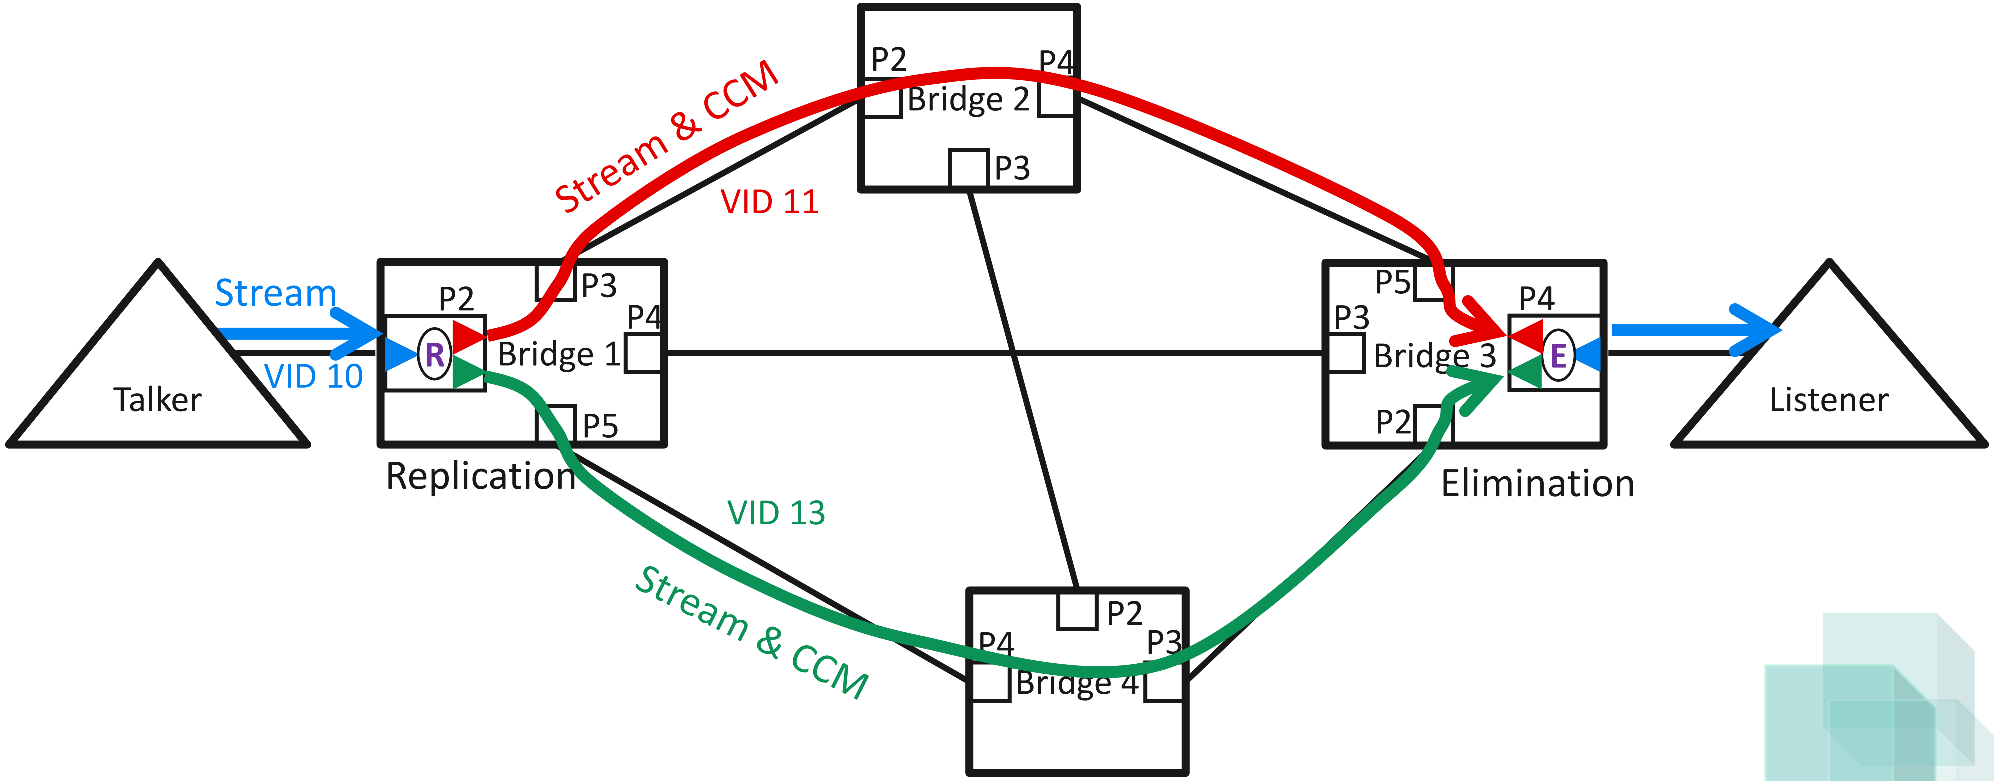

Where Are We at The Moment?

- No change to data plane, no new header fields needed to make CFM work for FRER
- MEP operations need to be extended for monitoring FRER
- Functionality introduced to 802.1Q for traffic engineering needs to be double checked and potential slight refinements might be needed
- Best practice guidelines could be provided, e.g.,
  - Use Null Stream identification
  - Add R-tag on edge bridges

- OAM is not that straightforward in case of IP Stream Identification and Mask & Match Stream Identification

Demo Network for Proof of Concept

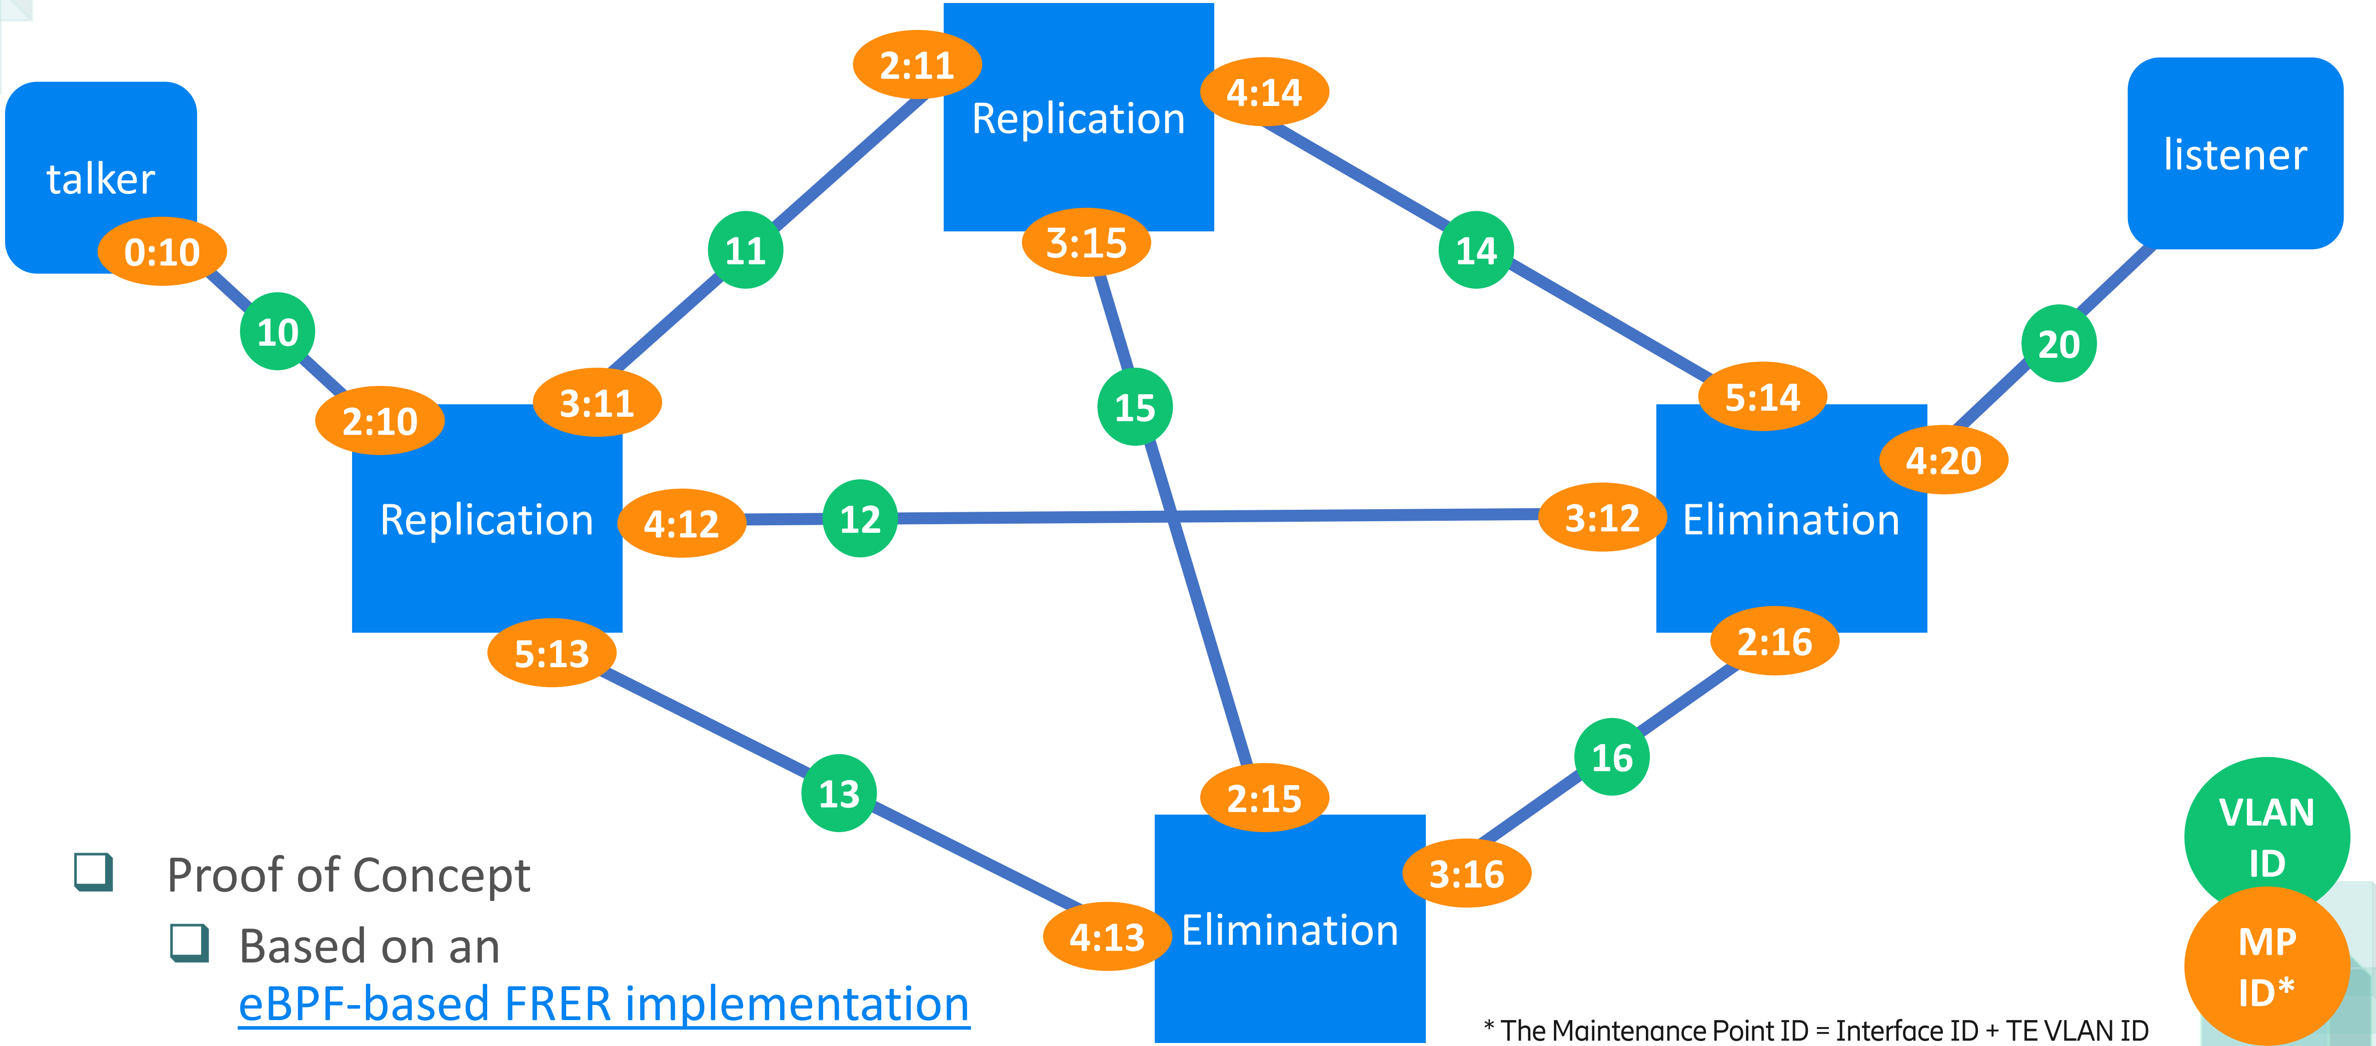

- Proof of Concept
  - Based on an [eBPF-based FRER implementation]

* The Maintenance Point ID = Interface ID + TE VLAN ID

Continuity Check

```
oamgen -c 5 -I tx-eth0 -t CC --mpid 3:16
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=0, oam_type=OAM_CC
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=1, oam_type=OAM_CC
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=2, oam_type=OAM_CC
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=3, oam_type=OAM_CC
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=4, oam_type=OAM_CC

--- MP-ID: 3:16 statistics ---
5 packets transmitted, 5 received, 0.0% packet loss
```

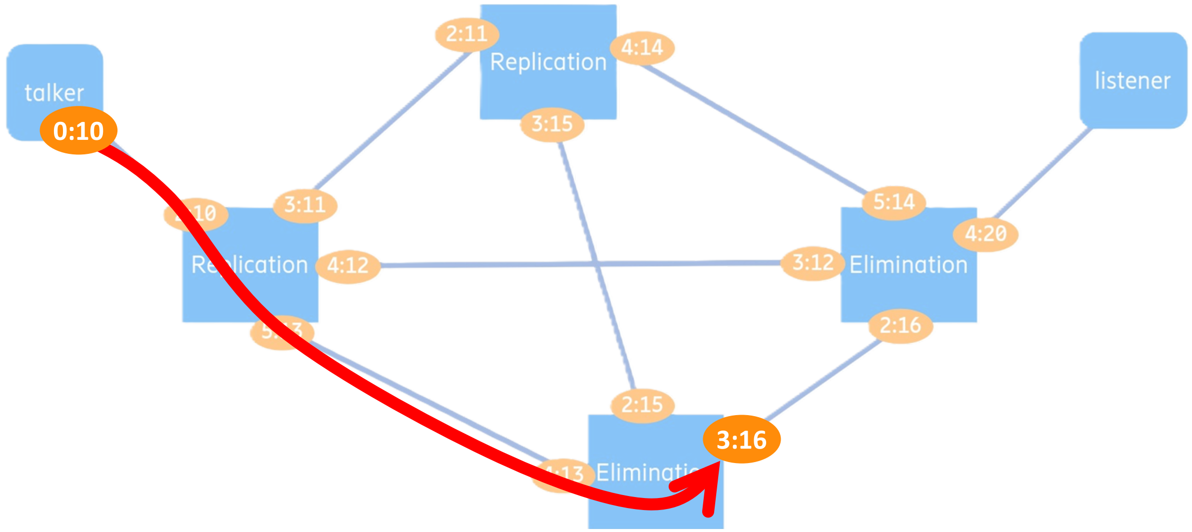

Continuity Check – cont'd

```
oamgen -c 5 -I tx-eth0 -t CC --mpid 3:16
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=0, oam_type=OAM_CC
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=1, oam_type=OAM_CC
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=2, oam_type=OAM_CC
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=3, oam_type=OAM_CC
34 bytes from de:c2:af:0a:d7:3f, rtag_seq=4, oam_type=OAM_CC

--- MP-ID: 3:16 statistics ---
5 packets transmitted, 5 received, 0.0% packet loss
```

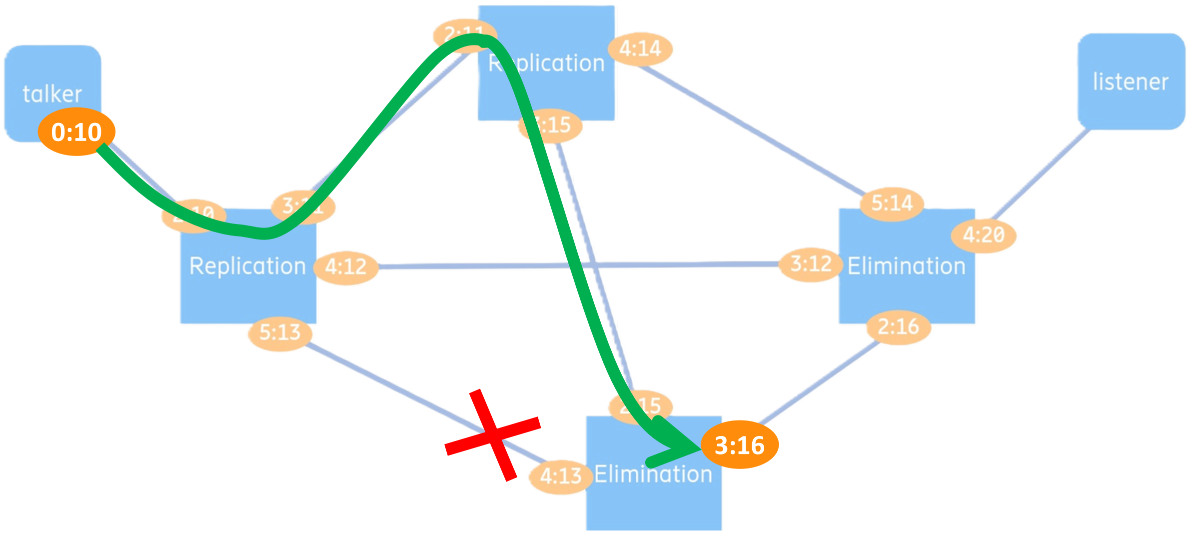

Record Route

- Record Route is a new FRER specific OAM function developed by the contributors as they found it useful

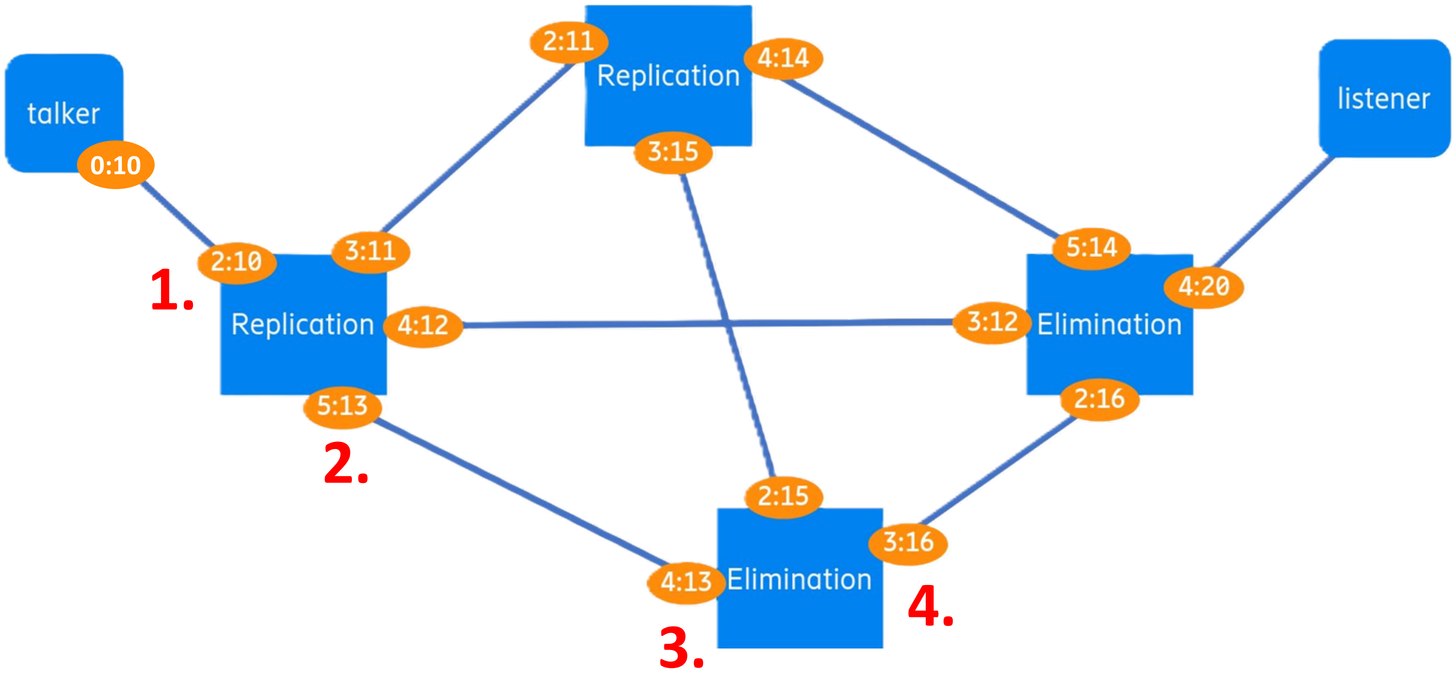


```
oamgen -c 5 -I tx-eth0 -t RR --mpid 3:16

50 bytes from de:c2:af:0a:d7:3f, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 5:13, 4:13, 3:16]
```

1. 2. 3. 4.

Record Route – cont’d

- Record Route is a new FRER specific OAM function developed by the contributors as they found it useful

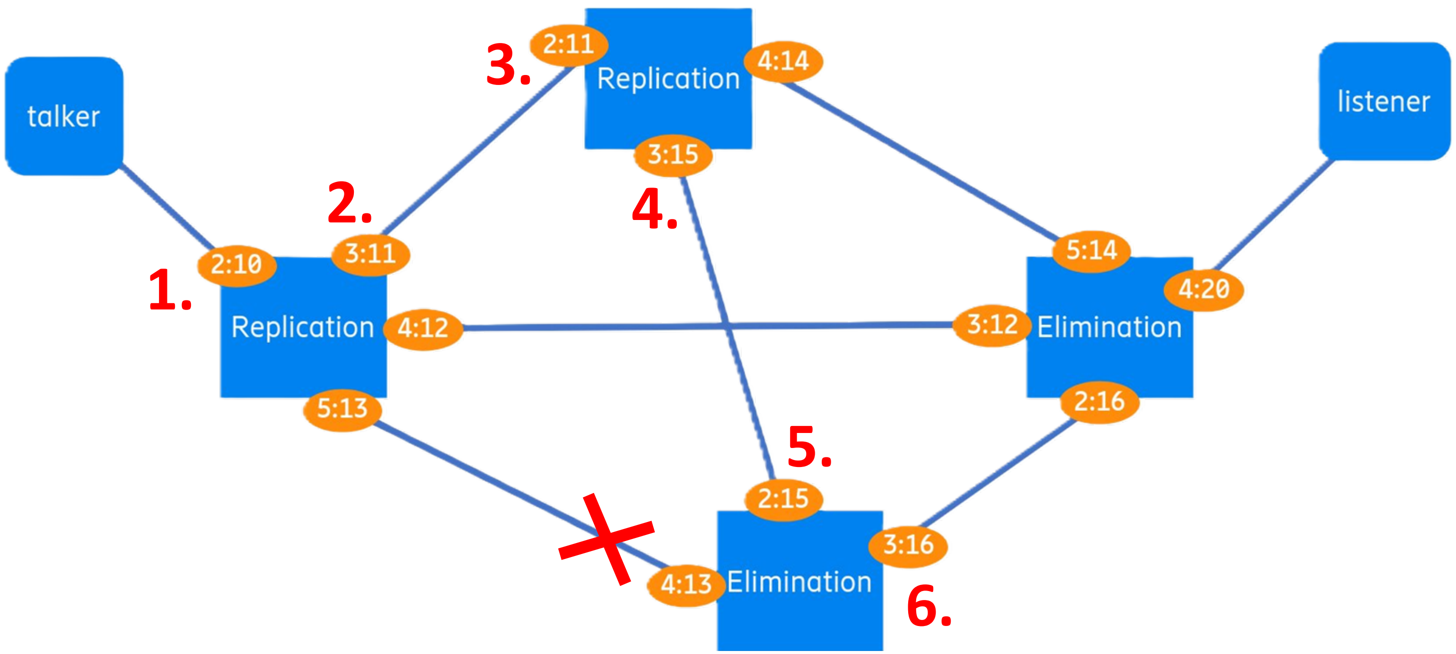


```
oamgen -c 1 -I tx-eth0 -t RR --mpid 3:16

58 bytes from de:c2:af:0a:d7:3f, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 3:11, 2:11, 3:15, 2:15, 3:16]
```

1. 2. 3. 4. 5. 6.

FRER Graph Discovery

- This is Record Route (see previous slides) for a special MP-ID "any"
- Every MP must reply

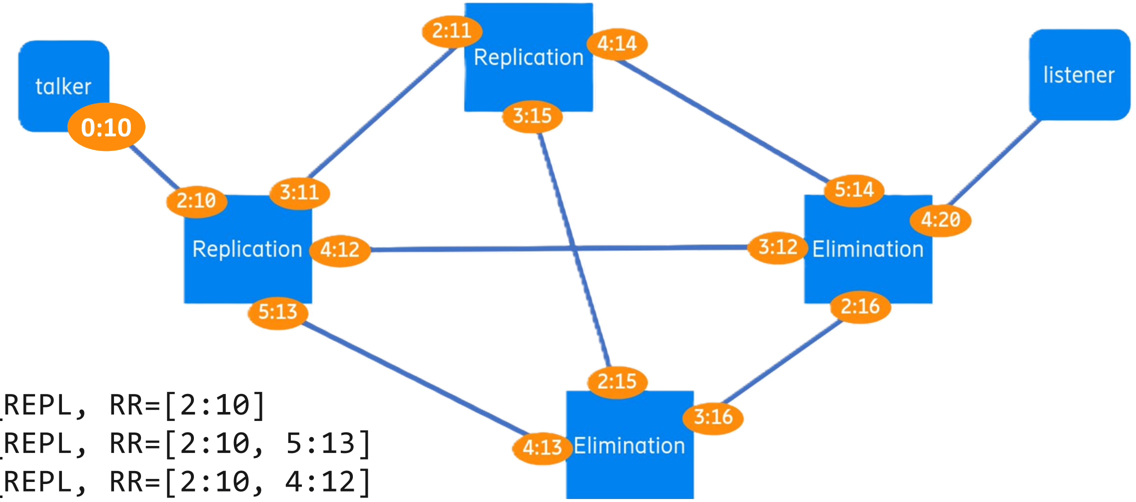


```
oamgen -c 1 -x 14 -I tx-eth0 -t RR --mpid any
38 bytes from 5a:12:03:9e:20:20, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10]
42 bytes from 5a:12:03:9e:20:20, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 5:13]
42 bytes from 5a:12:03:9e:20:20, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 4:12]
42 bytes from 5a:12:03:9e:20:20, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 3:11]
46 bytes from 6a:04:55:42:31:b2, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 4:12, 3:12]
50 bytes from 6a:04:55:42:31:b2, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 4:12, 3:12, 4:20]
54 bytes from 6a:04:55:42:31:b2, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 5:13, 4:13, 3:16, 2:16]
54 bytes from 6a:04:55:42:31:b2, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 3:11, 2:11, 4:14, 5:14]
46 bytes from fa:26:12:3a:68:3b, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 3:11, 2:11]
46 bytes from de:c2:af:0a:d7:3f, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 5:13, 4:13]
50 bytes from de:c2:af:0a:d7:3f, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 5:13, 4:13, 3:16]
50 bytes from fa:26:12:3a:68:3b, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 3:11, 2:11, 4:14]
50 bytes from fa:26:12:3a:68:3b, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 3:11, 2:11, 3:15]
54 bytes from de:c2:af:0a:d7:3f, rtag_seq=0, oam_type=OAM_RR_REPL, RR=[2:10, 3:11, 2:11, 3:15, 2:15]

--- MP-ID: any statistics ---
1 packet transmitted, 14 received, 0.0% packet loss
```

Summary

- The service protection provided by FRER hides network issues completely until a given Stream is fully broken
- FRER counters provide some level of warnings, however, do not reveal the network issues
- OAM applied to FRER can reveal network issues, provide information on what is going on in the network
- OAM can be provided for FRER by extending CFM
  - No change to the data plane
  - MEP operations need to be extended
  - Further refinements might be needed

DETERMINISTIC6G Grant Agreement No. 101096504

The DETERMINISTIC6G project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101096504.

If you need further information, please contact the coordinator:

János Harmatos, ERICSSON

E-Mail: coordinator@deterministic6g.eu

or visit: www.deterministic6g.eu

@DETERMINISTIC6G   DETERMINISTIC6G