


Practical Implementation of EPON ONU Management using OAM over VLC

How does IEEE 1904.2 D2.0 scale?

Background

- IEEE 1904.2 was originally chartered because there was a need to allow "remote" management of EPON ONUs or EPOC CNU's.
- In any practical implementation, this would naturally be a one manager to many ONUs relationship, where the manager is a single server (bare metal or virtual machine)


Desired Topology



- Single Management Server should be able to manage multiple EPON ONUs using OAM over VLC
- VLC is implemented as a software program running on the server

Implementation Idea 1

- VLC Sublayer is implemented as a software program running on the server
- From a software and resource efficiency perspective, a single instance of VLC in the manager would be good
- VLC Rules have no way to distinguish between OAM AB and OAM AZ as sources, so there is no way to steer OAM into the correct tunnel



Egress Rule 1 in VLC SL A:

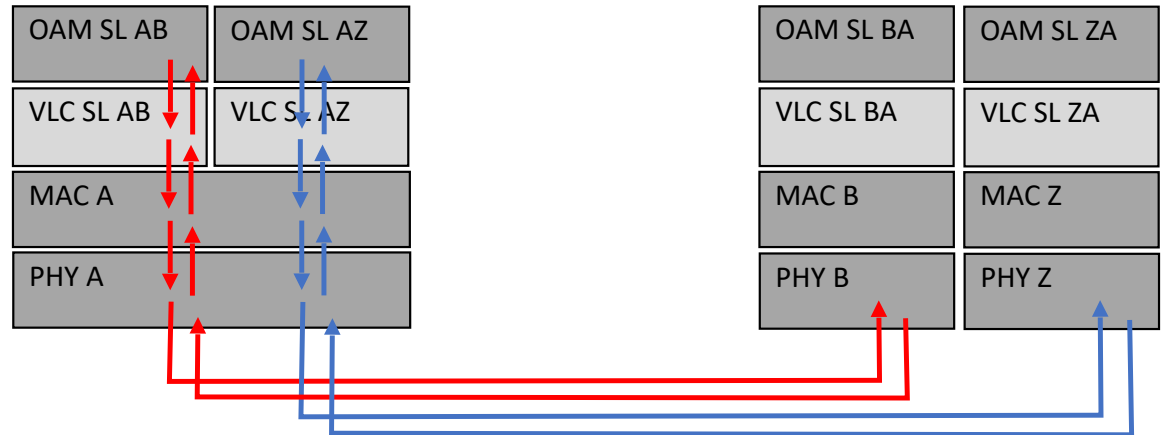
```
IF FID_DST_ADDR==SP_ADDR AND FID_LEN_TYPE==ETHERTYPE_SP  
AND FID_SUBTYPE==SUBTYPE_OAM  
THEN REPLACE(FID_DST_ADDR, Z) AND REPLACE(FID_LEN_TYPE,  
ETHERTYPE_VLC)
```

Egress Rule 2 in VLC SL A:

```
IF FID_DST_ADDR==SP_ADDR AND FID_LEN_TYPE==ETHERTYPE_SP  
AND FID_SUBTYPE==SUBTYPE_OAM  
THEN REPLACE(FID_DST_ADDR, B) AND REPLACE(FID_LEN_TYPE,  
ETHERTYPE_VLC)
```

Implementation Idea 2

- Alternative: Multiple VLC instances running in software (separate processes, separate threads, or simply separate sockets)
- Egress from Manager - Works
- Ingress to Manager - How does MAC A know which VLC SL receives incoming VLC PDU?
- This approach doesn't work



Egress Rule 1 in VLC SL AZ:

IF FID_DST_ADDR==SP_ADDR AND FID_LEN_TYPE==ETHERTYPE_SP AND FID_SUBTYPE==SUBTYPE_OAM

THEN REPLACE(FID_DST_ADDR, Z) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_VLC)

Ingress Rule 1 in VLC SL AZ:

IF FID_SRC_ADDR==Z AND FID_LEN_TYPE==ETHERTYPE_VLC AND FID_SUBTYPE==SUBTYPE_OAM

THEN REPLACE(FID_DST_ADDR, SP_ADDR) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_SP)

Egress Rule 1 in VLC SL AB:

IF FID_DST_ADDR==SP_ADDR AND FID_LEN_TYPE==ETHERTYPE_SP AND FID_SUBTYPE==SUBTYPE_OAM

THEN REPLACE(FID_DST_ADDR, B) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_VLC)

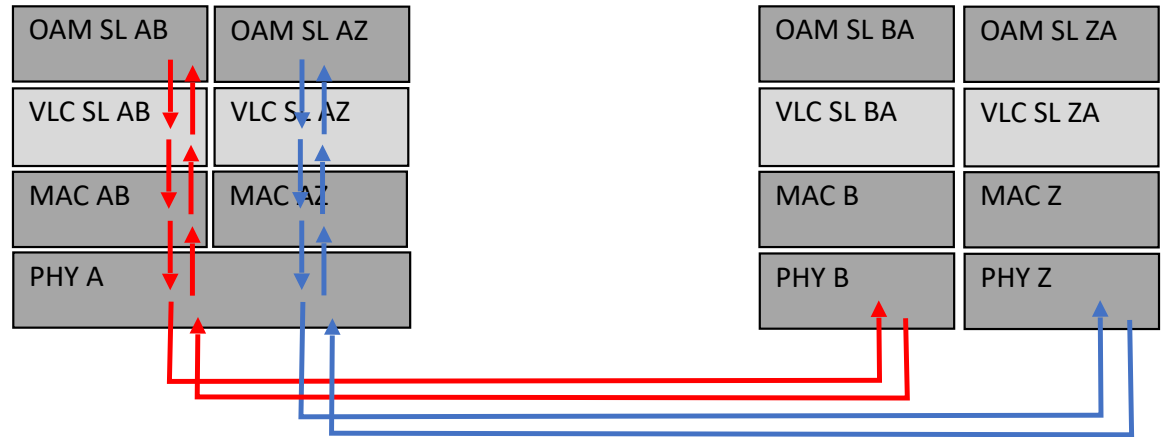
Ingress Rule 1 in VLC SL AB:

IF FID_SRC_ADDR==Z AND FID_LEN_TYPE==ETHERTYPE_VLC AND FID_SUBTYPE==SUBTYPE_OAM

THEN REPLACE(FID_DST_ADDR, SP_ADDR) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_SP)

Implementation Idea 3

- Alternative: Multiple VLC instances running in software (separate processes, separate threads, or simply separate sockets) with multiple “virtual” MACs.
- This approach works...



Egress Rule 1 in VLC SL AZ:

IF FID_DST_ADDR==SP_ADDR AND FID_LEN_TYPE==ETHERTYPE_SP
AND FID_SUBTYPE==SUBTYPE_OAM

THEN REPLACE(FID_DST_ADDR, **Z**) AND REPLACE(FID_LEN_TYPE,
ETHERTYPE_VLC)

Ingress Rule 1 in VLC SL AZ:

IF FID_SRC_ADDR==**Z** AND FID_LEN_TYPE==ETHERTYPE_VLC AND
FID_SUBTYPE==SUBTYPE_OAM

THEN REPLACE(FID_DST_ADDR, **SP_ADDR**) AND
REPLACE(FID_LEN_TYPE, ETHERTYPE_SP)

Egress Rule 1 in VLC SL AB:

IF FID_DST_ADDR==SP_ADDR AND FID_LEN_TYPE==ETHERTYPE_SP
AND FID_SUBTYPE==SUBTYPE_OAM


THEN REPLACE(FID_DST_ADDR, **B**) AND REPLACE(FID_LEN_TYPE,
ETHERTYPE_VLC)

Ingress Rule 1 in VLC SL AB:

IF FID_SRC_ADDR==**Z** AND FID_LEN_TYPE==ETHERTYPE_VLC AND
FID_SUBTYPE==SUBTYPE_OAM

THEN REPLACE(FID_DST_ADDR, **SP_ADDR**) AND
REPLACE(FID_LEN_TYPE, ETHERTYPE_SP)

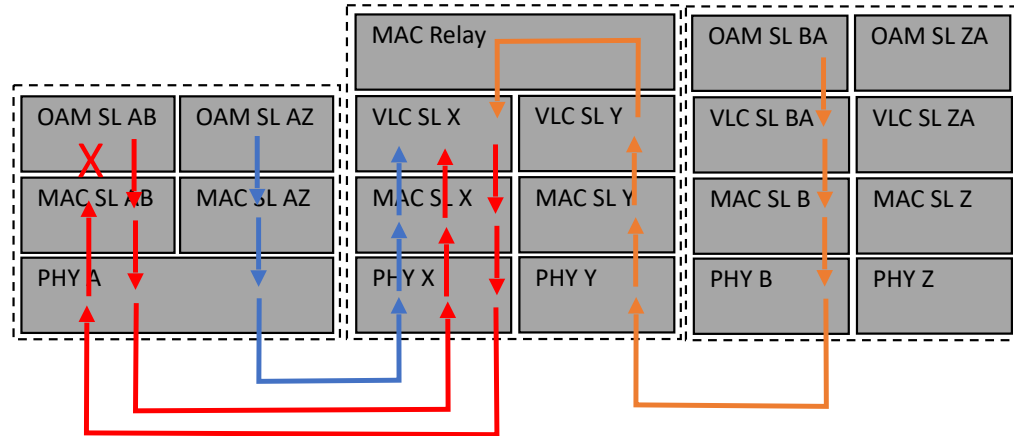
Implementation Idea 3 (continued)



- This approach works, but does it scale?
 - One server managing a single PON with 128 ONUs requires 128 MACs, each with its own MAC address
 - A single OLT system (S-OLT) might have >64 PON ports
 - A single server serving a typical S-OLT would need $128 \times 64 = 8192$ MACs, each with its own MAC address
 - This is very tedious to manage making it impractical

Implementation Idea 4

- What if VLC is in the bridge?
 - Scaling on the server is still impractical, but that's an OAM problem and no longer a VLC problem.
 - Rules that can be applied on the bridge are able to handle this situation better.
- There are no "hardware" bridges capable of this today
- Requires that the bridge-to-manager OAMPDUs be addressed to the individual MAC address



Ingress Rule 1 in VLC SL X:

IF FID_DST_ADDR==SP_ADDR AND FID_SRC_ADDR==**AB** AND FID_LEN_TYPE==ETHERTYPE_SP AND FID_SUBTYPE==SUBTYPE_OAM
 THEN REPLACE(FID_DST_ADDR, **B**) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_VLC)

Ingress Rule 2 in VLC SL X:

IF FID_DST_ADDR==SP_ADDR AND FID_SRC_ADDR==**AZ** AND FID_LEN_TYPE==ETHERTYPE_SP AND FID_SUBTYPE==SUBTYPE_OAM
 THEN REPLACE(FID_DST_ADDR, **Z**) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_VLC)

Egress Rule 1 in VLC SL X:

IF FID_SRC_ADDR==**B** AND FID_LEN_TYPE== ETHERTYPE_VLC AND FID_SUBTYPE==SUBTYPE_OAM
 THEN REPLACE(FID_DST_ADDR, **AB**) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_OAM)

Egress Rule 1 in VLC SL X:

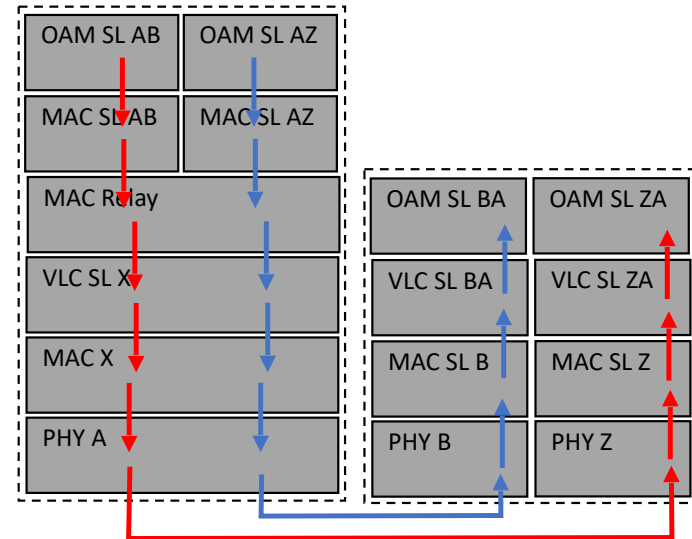
IF FID_SRC_ADDR==**Z** AND FID_LEN_TYPE== ETHERTYPE_VLC AND FID_SUBTYPE==SUBTYPE_OAM
 THEN REPLACE(FID_DST_ADDR, **AZ**) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_OAM)

802.3CL57 → NO

IS THIS LEGAL? Will the receiving OAM sublayer accept the OAMPDU without DST_MAC_ADDR == SP_MAC_ADDR?

Implementation Idea 5

- What if VLC is in the bridge and the bridge is in the server?
- Bridge is implemented in software, so no dependencies on availability of a hardware-based bridge
- Still requires that the OAMPDU VLC SL X to OAM SL AB or OAM SL AZ use individual MAC address



Egress Rule 1 in VLC SL X:

IF FID_DST_ADDR==SP_ADDR AND FID_SRC_ADDR==**AB** AND FID_LEN_TYPE==ETHERTYPE_SP AND FID_SUBTYPE==SUBTYPE_OAM
 THEN REPLACE(FID_DST_ADDR, **B**) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_VLC)

Egress Rule 2 in VLC SL X:

IF FID_DST_ADDR==SP_ADDR AND FID_SRC_ADDR==**AZ** AND FID_LEN_TYPE==ETHERTYPE_SP AND FID_SUBTYPE==SUBTYPE_OAM
 THEN REPLACE(FID_DST_ADDR, **Z**) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_VLC)

Ingress Rule 1 in VLC SL X:

IF FID_SRC_ADDR== **B** AND FID_LEN_TYPE== ETHERTYPE_VLC AND FID_SUBTYPE==SUBTYPE_OAM
 THEN REPLACE(FID_DST_ADDR, **AB**) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_OAM)

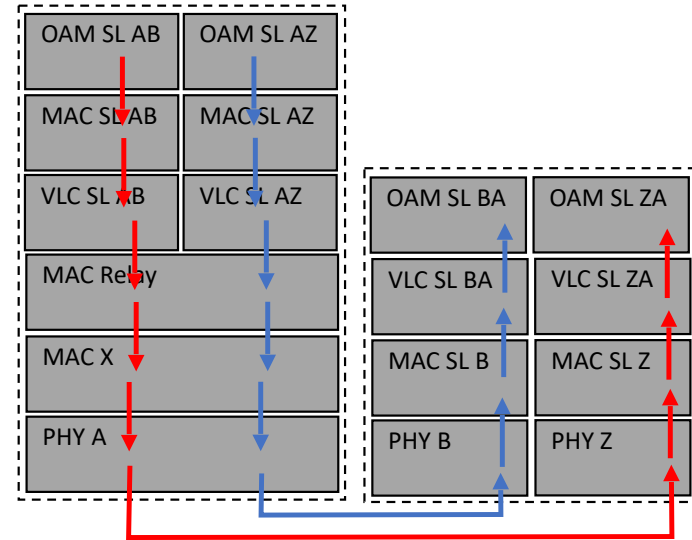
Ingress Rule 1 in VLC SL X:

IF FID_SRC_ADDR== **Z** AND FID_LEN_TYPE== ETHERTYPE_VLC AND FID_SUBTYPE==SUBTYPE_OAM
 THEN REPLACE(FID_DST_ADDR, **AZ**) AND REPLACE(FID_LEN_TYPE, ETHERTYPE_OAM)

Still a problem: OAMPDU doesn't have
 DST_MAC_ADDR == SP_MAC_ADDR?

Implementation Idea 6

- VLC instance on each “virtual” MAC
- This is effectively the same as Idea 3 with the same outstanding question:
 - Is there anything that we can do in VLC to enable an implementation that is more easily scaled to support management of many CPE devices?



Summary

- The architecture and rules defined in the D2.0 text describe a one-to-one relationship which aligns with IEEE 802.3.
- This alignment is understandable but raises a question of scalability.
- Applying the draft as written, it is necessary to have a one-to-one relationship (one manager MAC to one ONU MAC).
- If an implementer wishes to manage multiple ONUs from a single "server", which is a completely reasonable expectation, then the "server" would need to have multiple MACs.
- It is not uncommon for a single S-OLT to have 64 or more PON ports, each serving 64 to 128 ONUs. This scale would require the "manager" to have up to 8192 MACs.
- It would not be unreasonable to expect a single "manager" server to support multiple S-OLTs, requiring multiple 10's of thousands of MACs

Summary (continued)

- In this presentation we've established that IEEE 1904.2 D2.0 supports only two models
 - “Hardware Bridge” (Implementation Idea 4)
 - Multiple MACs in “Manager” (Implementation Idea 3 and Implementation Idea 6)
- “Hardware bridge” is not a practical near-term option
- Is there anything that we can do in VLC to enable an implementation that is more easily scaled to support management of many CPE devices?