Fiber to the Room Architectural Developments

Frank Effenberger

Futurewei

Sep 2023

Simplest view of PON-based FTTR

- In its simplest terms, the PON-based FTTR network is a "PON of PONs"
  - The traditional access PON carries data to the house
  - The home PON carries data inside the home
- This raises the basic problem: How to interconnect the two PONs?

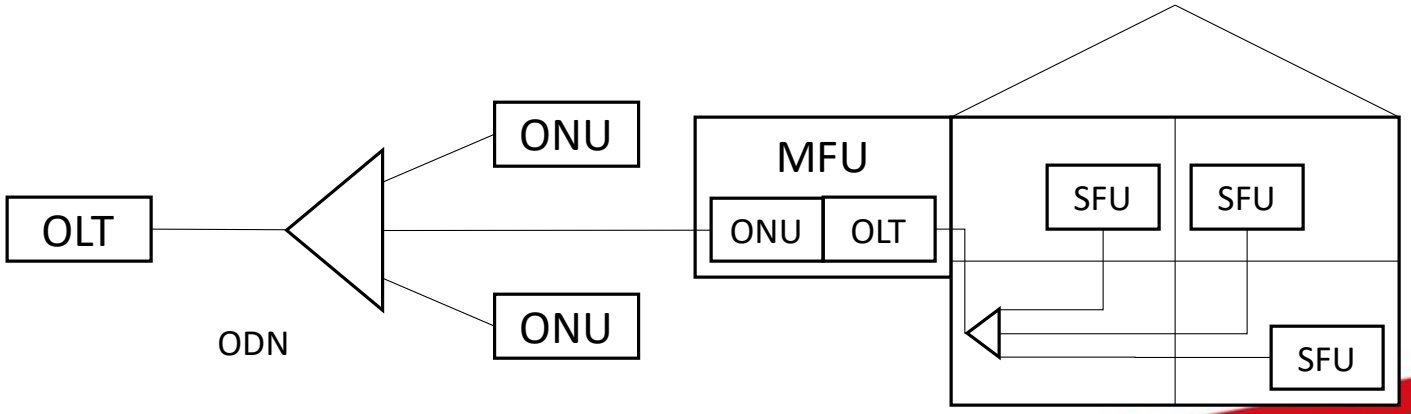

Interworking at different layers

Layer 0

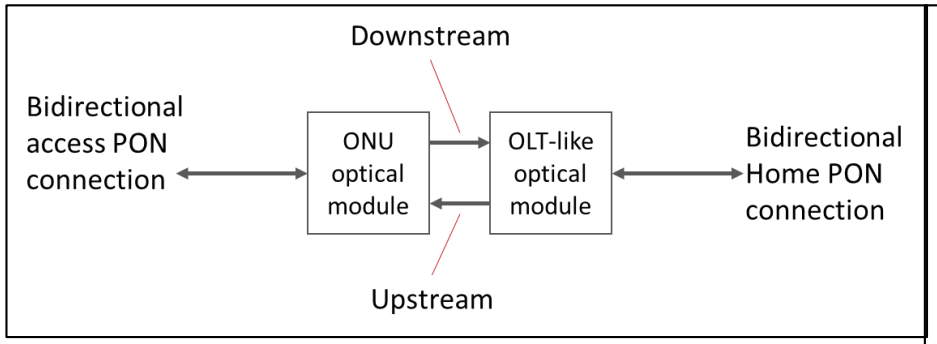

Layer 1

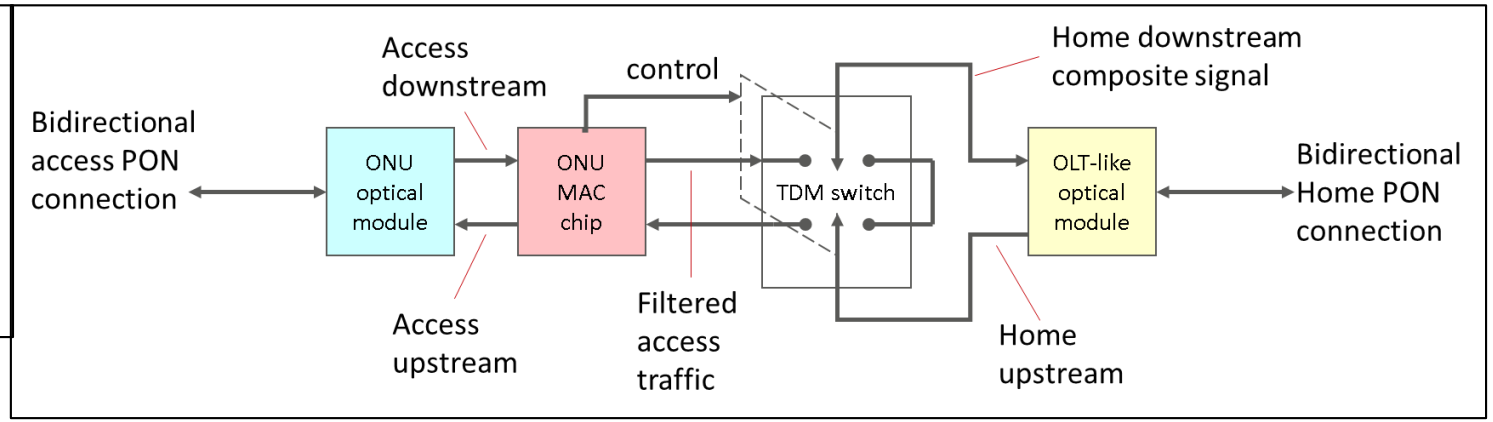

Layer 2

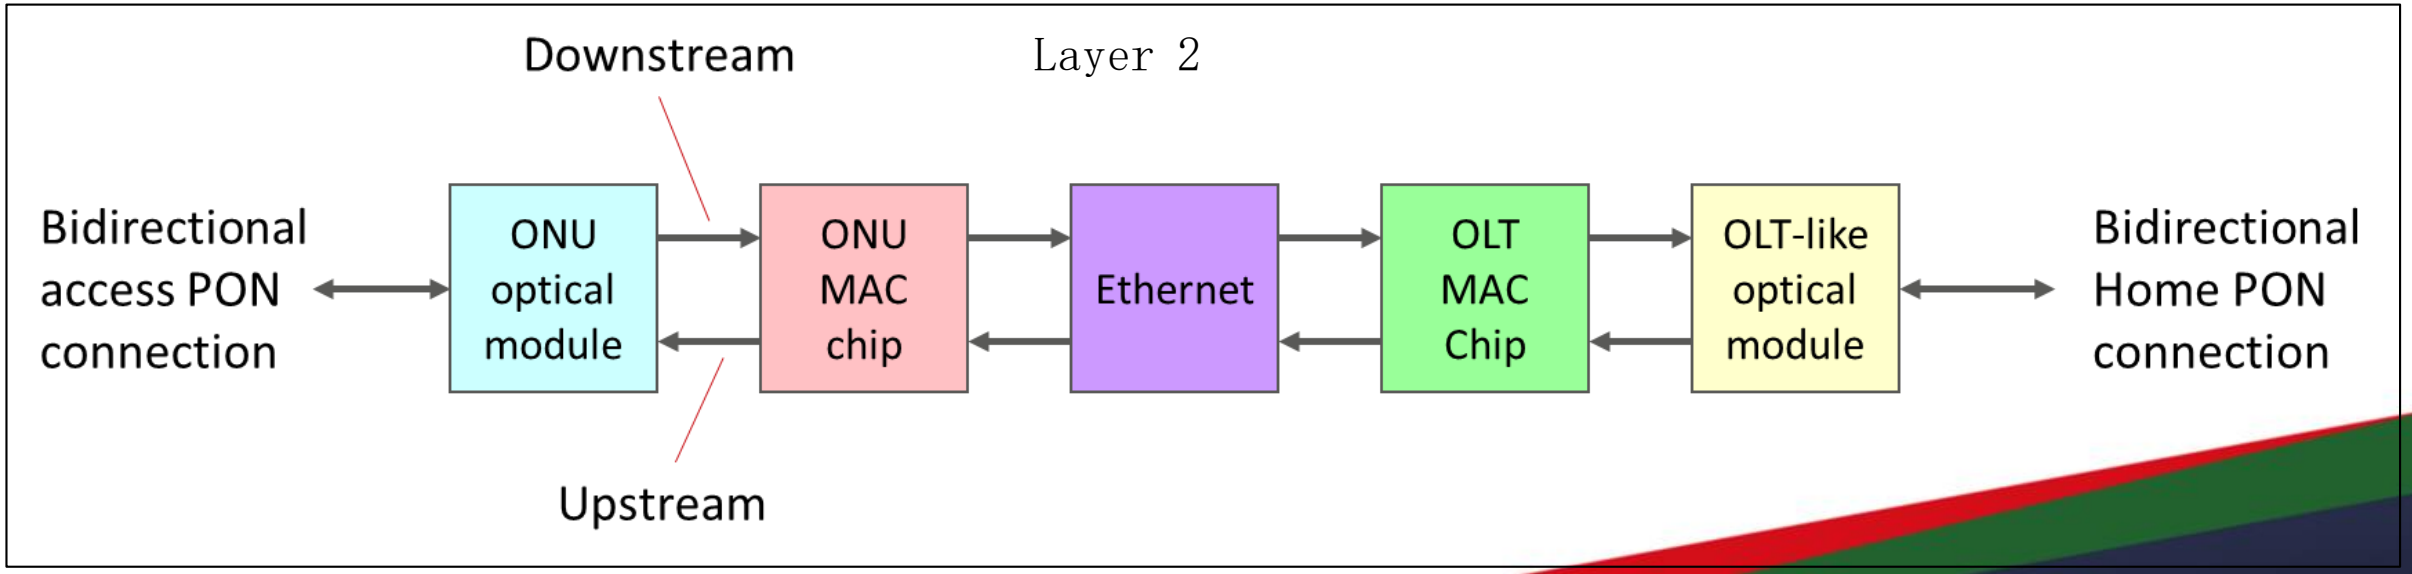

Practical commercial limitations

- While it is interesting to think about optimizations on the gateway architecture, we must remember that nearly all of these will require new ASIC development
  - ASICs are very expensive to design and fabricate, but then very cheap to mass produce
  - Because of this, developers won't pull the trigger on a new chip until the market is solidly realized
  - FTTR is making its way, with a few million units deployed, but it still isn't there yet
- Therefore, we need to figure out how to build a MFU gateway using existing ASIC devices

MFU gateway using existing ASICs

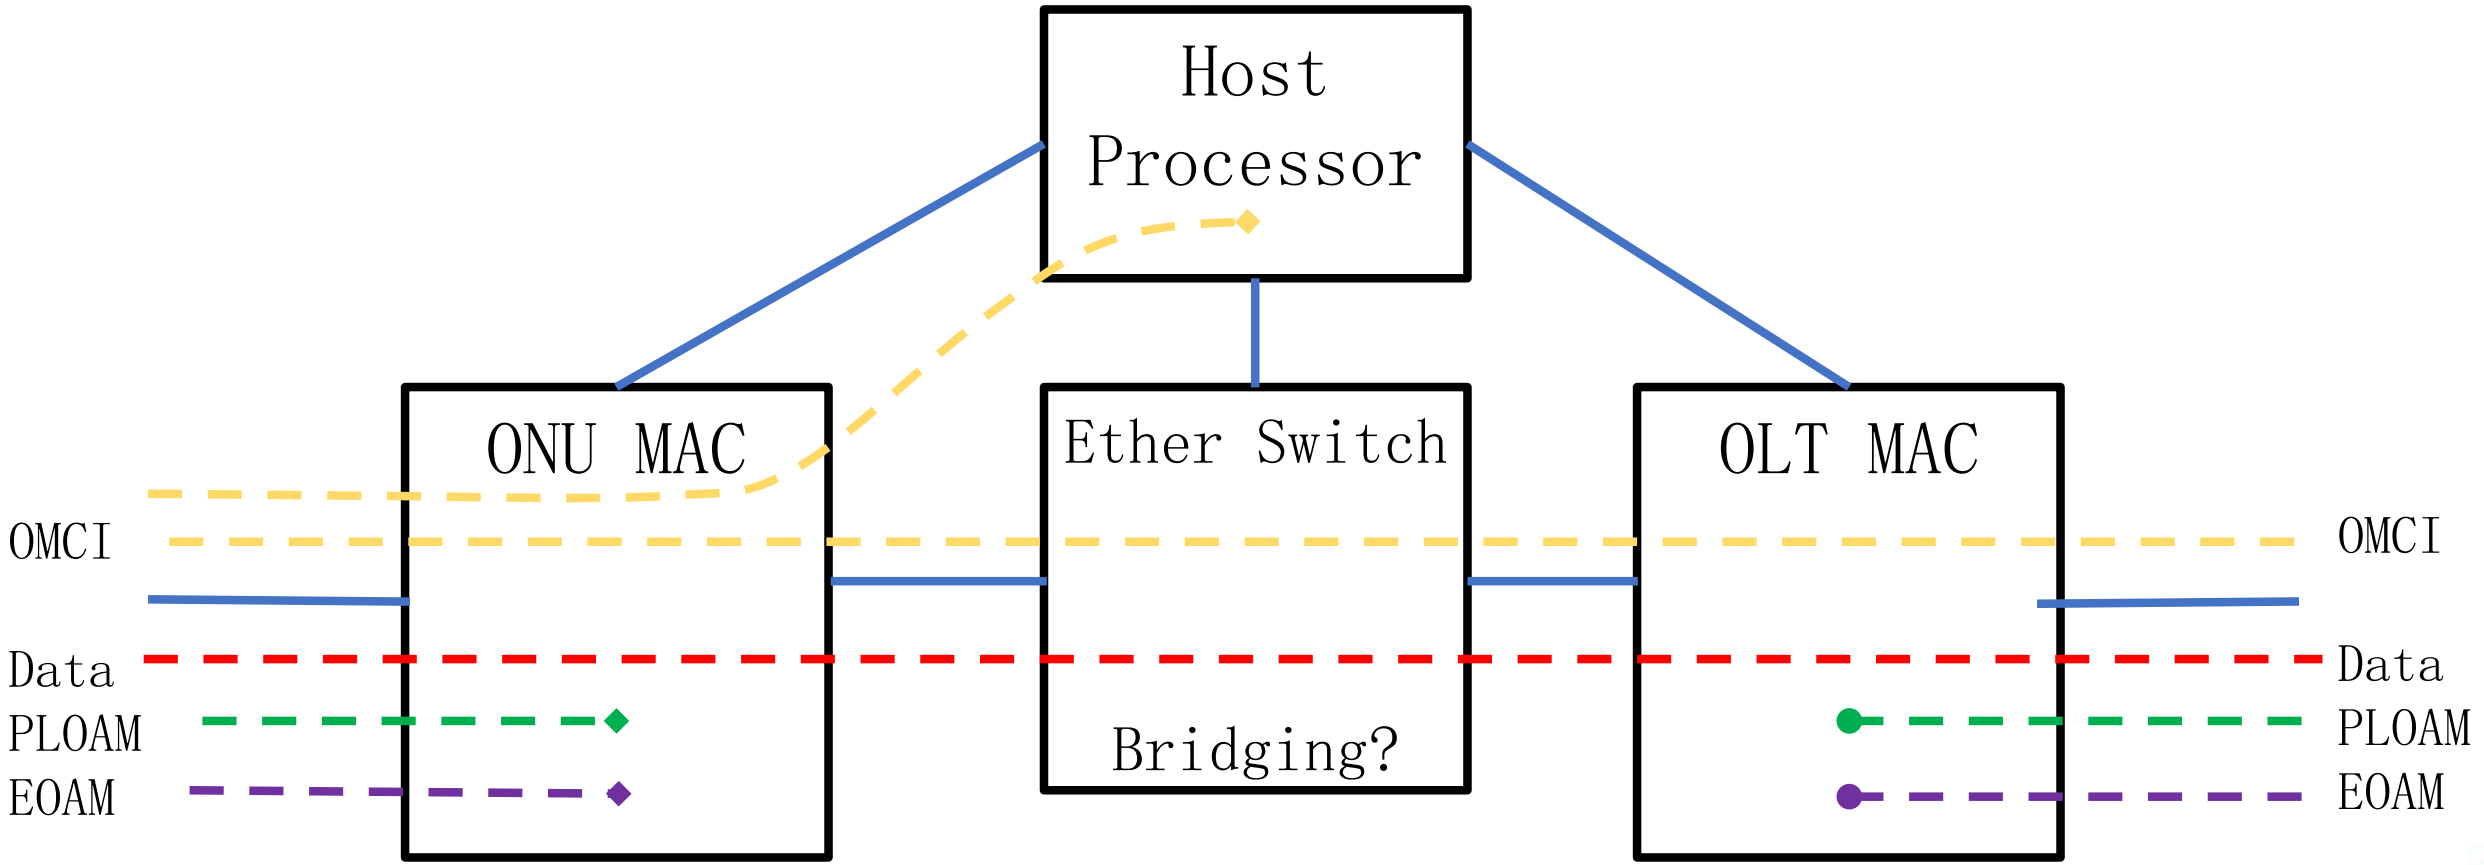

Embedded OAM

- The EOAM data is received by the ONU MAC from the access PON signal. There are two classes of EOAM, static data (the PSBd info) and dynamic (bandwidth map)
- The static data could simply be copied from the ONU MAC to the OLT MAC (via read/write from the host processor)
- The BWmap is too fast to be “copied”
  - Therefore, the MFU OLT MAC will have to make its own BWmaps
  - This requires traffic contracts to drive the DBA algorithm
  - Maybe the home PON will be undersubscribed, and the DBA could be simplified (e.g., a static assignment of BW)
  - The T-Conts must also be configured (PLOAM and OMCI channels)

Physical Layer OAM

- PLOAM messages come in several varieties
  - Timing non-critical vs. critical
  - Broadcast vs. unicast
  - Created in the ASIC or in the host processor
- The timing non-critical messages can be relayed by the MFU onto the home PON
  - Examples include system profiles, overhead, assign ONU-ID, deactivate SN, assign Alloc-ID, etc.
  - The access OLT sets up the Alloc-IDs, which must be used by the MFU-local BWmap generator
  - The MFU traffic contracts can be low level: a list of entries (Alloc-ID, minBW, guarBW, maxBW)
- The timing critical messages will likely be terminated in the ONU and OLT MACs
- This means that the access PON will have one PLOAM connection to the MFU, and then the MFU will need to have PLOAM connections to each of its subtending SFUs
  - Ranging is the most important example; however, the home PON physical extent is small, so the need for ranging could be reduced
- Note: if we have PLOAMs coming from the OLT and PLOAMs coming from the SFU, the sequence numbers will need to fit together. The SFU will not realize that it is actually talking to two OLTs

ONU Management and Control Interface

- The OMCI is the most complex part of the PON system
- It is desirable that the access OLT should manage the SFU's OMCI directly, and for the SFU OMCI MIB to match that of TR-156 (to leverage BBF's great work)
- The OMCI is carried over a GEM-port, just like every other data stream. Thus the OMCI connection can be handled just like the data
  - In that sense, the MFU is 'hands off' with the SFU's OMCI model
  - However, the MFU does need to interwork with the SFU using that OMCI model
  - This means that the MFU will need to know about the SFU's Alloc-IDs, Tconts, and portIDs
  - Of course, the MFU will also have its own OMCI channel, and its own model of traffic that it locally terminates
- In addition, the MFU will need to manage the UNI-side OLT and its optical interfaces

Data flows

- There are two types of flows
  - Access flows run from the Access OLT to the SFU
  - Local flows run from SFU to SFU, presumably through the Ethernet switch.
- For the access flows, the data flows need to be mapped across the MFU
  - For both directions, the simplest mode of operation is that the GEM fragments that come in on one side of the MFU will go out the other side with the same PortID
  - In the downstream, this should require just a simple 1-to-1 mapping. VLAN markings can be used to get across the Ethernet switch
  - In the case of no local traffic, there cannot be any contention in the downstream direction, so no QoS features are needed
  - In the upstream direction, the incoming PortID / Alloc-ID must be mapped into the upstream Tconts. This should be a simple 1-to-4 mapping, where the MFU will have the usual 4 Tconts, and then each home-side PortID will get assigned to one of the 4
- Local flows are a wholly new traffic pattern
- There are three ways this could be configured:
  - SFU unaware, partially aware, and SFU fully aware

Unaware SFU style

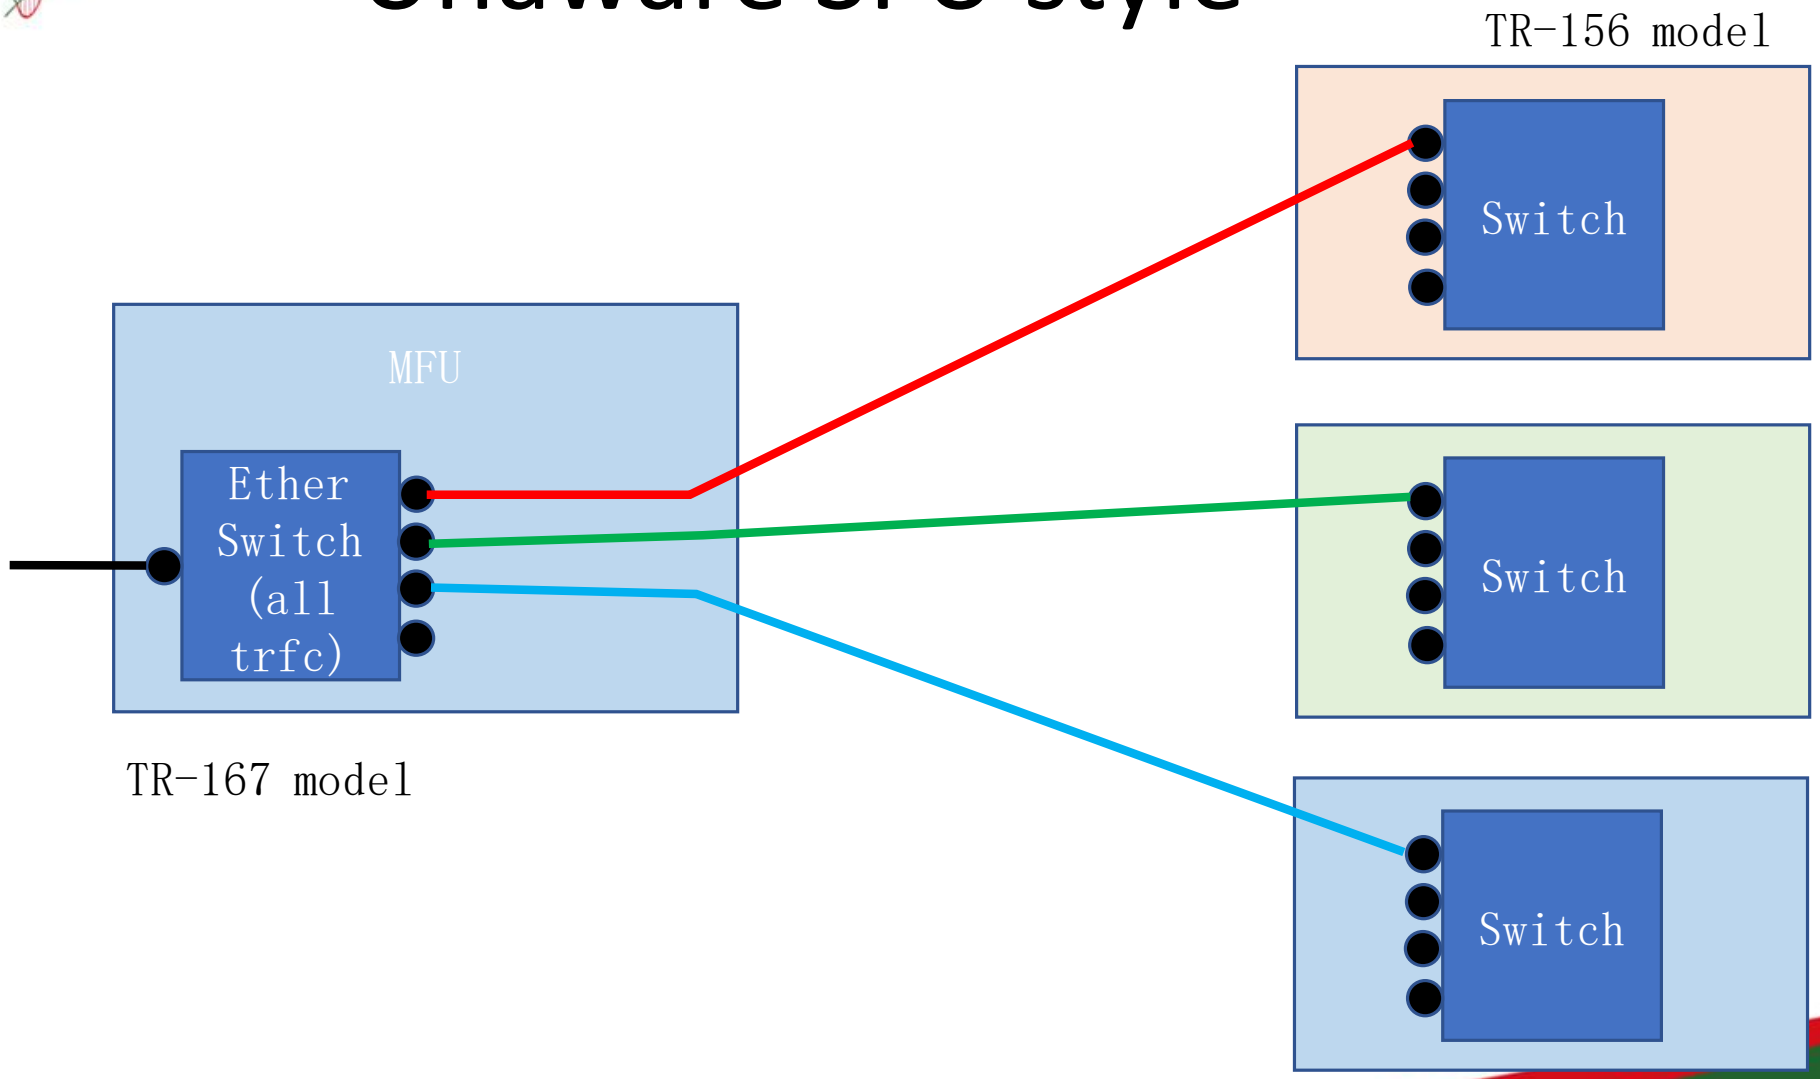

Semi-aware SFU style

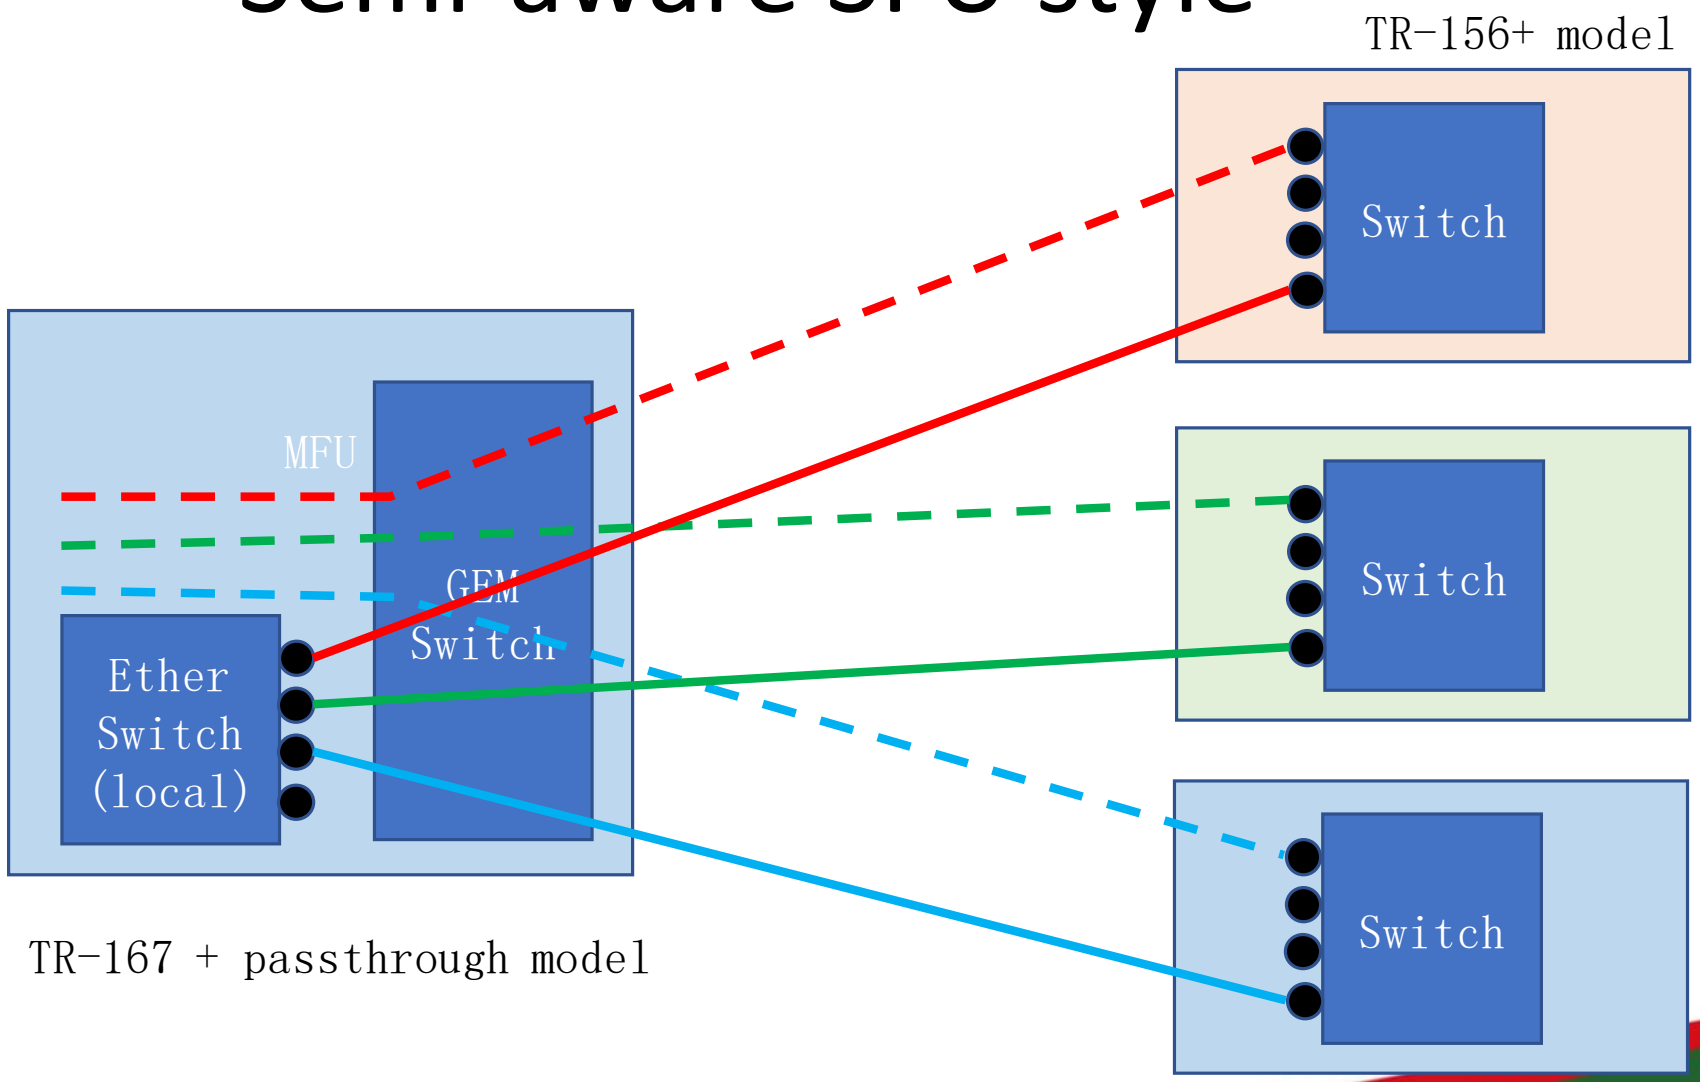

Fully aware SFU style

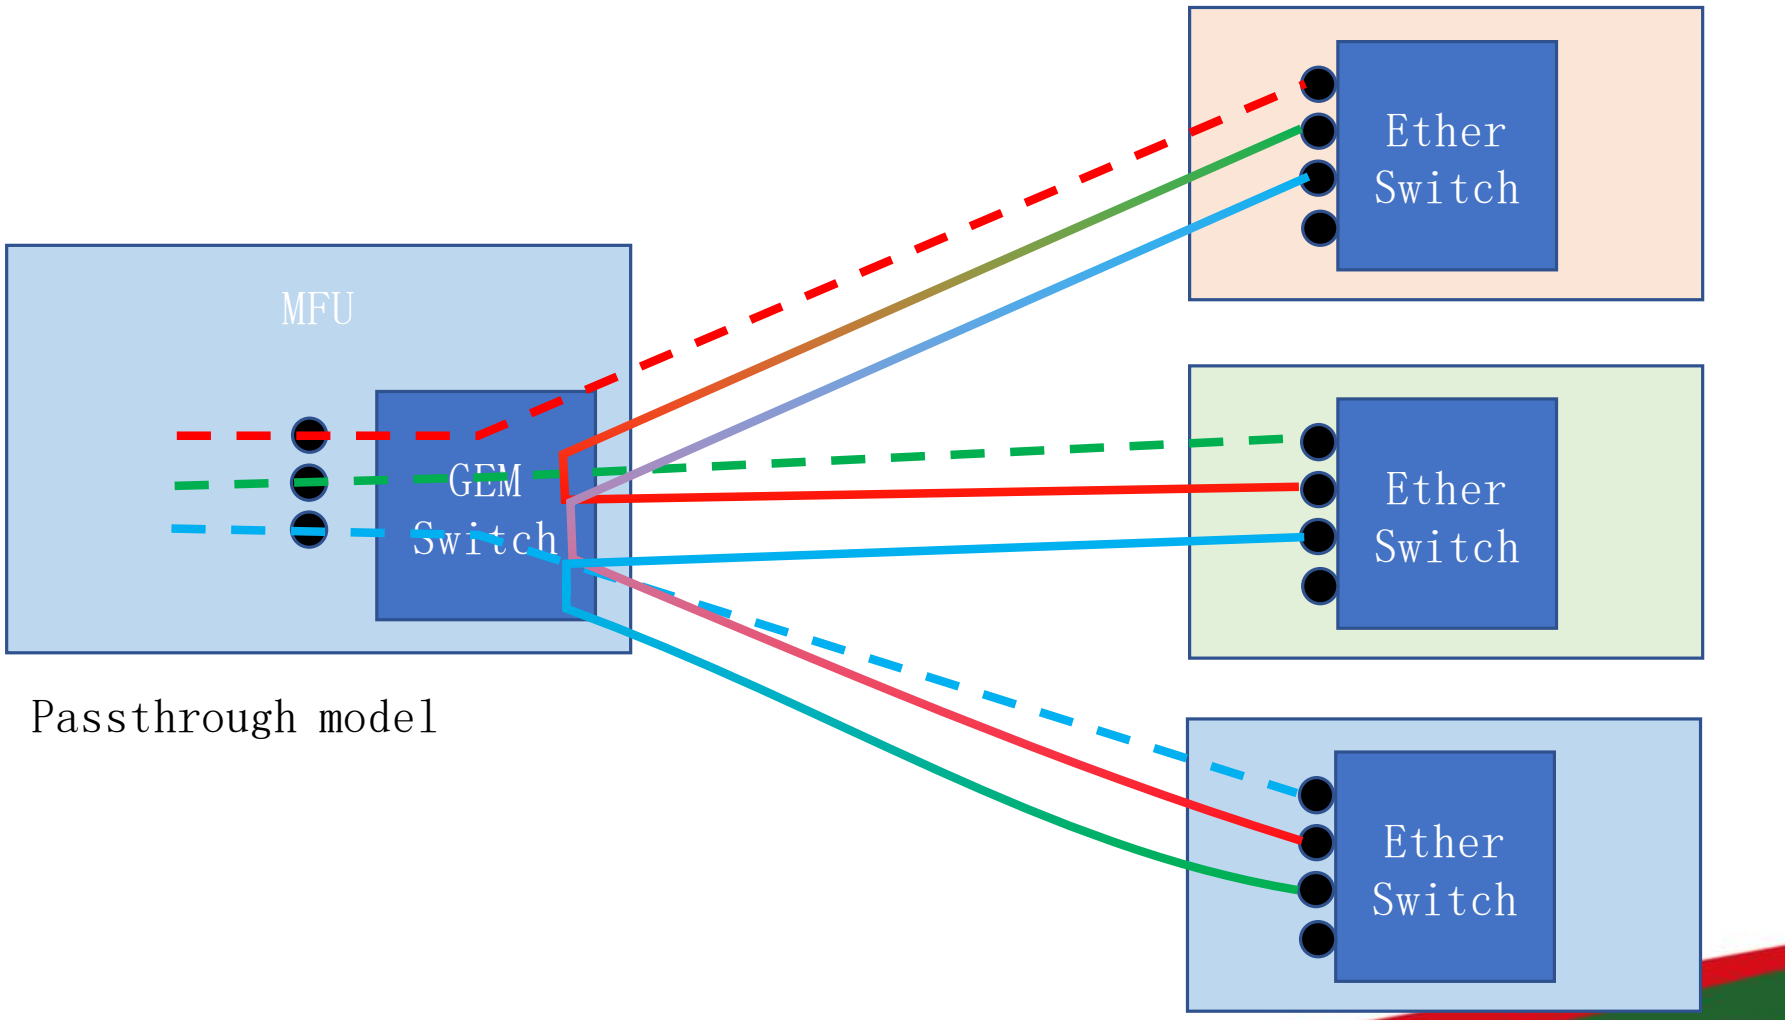

Summary

- Building an MFU using existing ASIC devices is possible
- Many low layer parameters can be directly copied over
  - However, the MFU will be creating its own Bwmap
- PLOAM and OMCI can be centralized to the access OLT
  - But work must be done to hide the complexities from the SFU
- Data flow can be handled with varied levels of complexity
  - SFU can be unaware, but it requires full switching at the MFU
  - As the SFU becomes more aware, the switching can be simplified