1. Services provided over EPON in MSO access networks

EPON is a truly converged network, capable of delivering different types of services, with varied QoS guarantees, while sharing a single physical OLT port. Thanks to the advanced DBA mechanisms in place in EPON, it is possible to mix residential best-effort as well as commercial bandwidth-guaranteed services on the same PON port.

In deployment models for EPON-based optical access adopted in MSO space by some of the operators, residential and commercial customers are typically separated and served using independent OLT platforms, sometimes co-located at the same hub. The main reason for this separation of residential and commercial customers is the way these services are provisioned, with residential platforms operating in the so-called DPoE mode with the full use of DOCSIS provisioning automation, while commercial platforms operated in the so-called EPON mode with manual configuration via CLI and very limited use of automation tools.

It is expected, though, that in the near future, both residential and commercial platforms merge and be served off the very same OLTs. At that time, it is likely that separation between residential and commercial customers still exists at the port level, i.e., a single OLT port will be used for either residential or commercial customers, but not both simultaneously. This may change in the emergence of best-effort commercial services, as discussed later on in this section.

1.1 Residential FTTH

In the residential FTTH service, the OLT operates in the routed mode. In this mode, IP connections from customer CPEs connected to the ONU are carried across the PON and then routed at the OLT into the Service Edge Router (SER) located northbound of the OLT and then into the public Internet across the converged transport network.

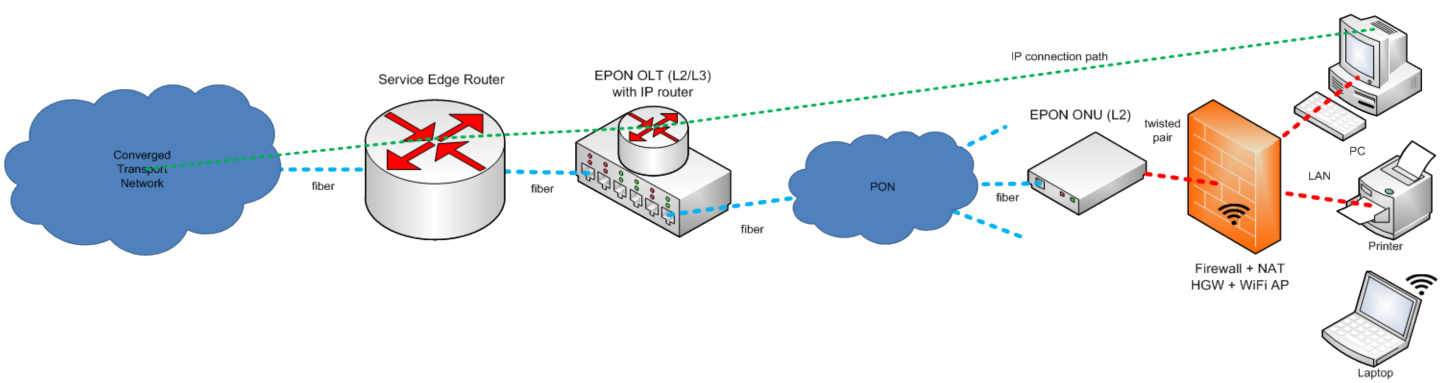

Figure 1: Architecture of a residential FTTH services

All the provisioning processes in the residential FTTH are derived from their DOCSIS counterparts, following the DPoE service and provisioning models. This means that all the existing tools and backoffice procedures developed over the years to deploy, manage, and bill DOCSIS-based residential customers are directly applicable to residential customers served over EPON. The only differences are obviously in the type of used media (fiber rather than coax) as well as bandwidth tiers (increased all the way to 1 Gbps symmetric over 1G-EPON).

1.2 Direct Internet Access

Direct Internet Access (DIA) is a special type of commercial services, where a commercial customer receives a dedicated L2 tunnel connected to public Internet across the converged transport network. In DIA, the OLT does not operate in the routed mode, and all Internet data between customer CPE and the DIA router is carried within a L2 tunnel based on a pre-configured combination of [802.1Q] VLAN tags.

DIA is typically provided as a secondary service to customers with exiting MEF service circuits, where MEF circuit uses primary UNI on the ONU, capable of delivering speeds up to 1 Gbps, while DIA is typically provided over the secondary UNI on the ONU, capable of delivering speeds up to 100 Mbps. The ONU forwards frames into the DIA tunnel using UNI port as the frame classifier: all frames exchanged over secondary UNI belong to DIA circuit.

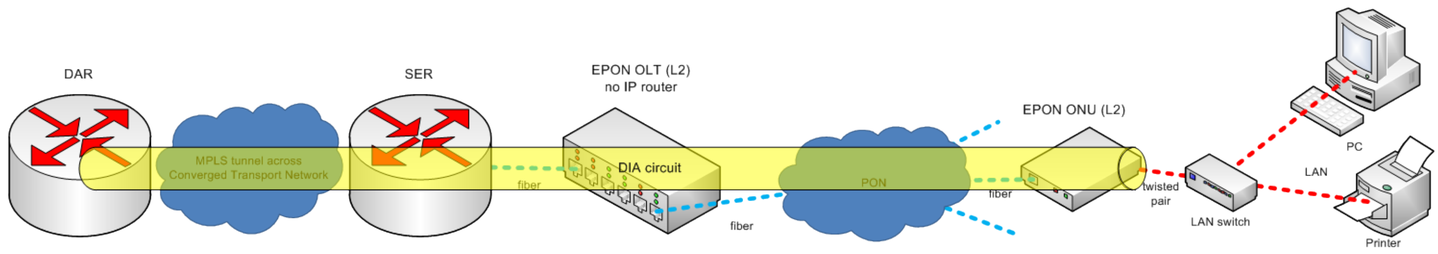

Figure 2: Architecture of DIA service

The OLT forwards frames that belong to the DIA circuit into a specific uplink port connected indirectly to the SER, which then builds MPLS tunnel across the converged transport network into a remote Dedicated Access Router (DAR) instance, connected directly to the public Internet.

In the downstream direction (from public Internet towards the customer CPE), all traffic policing is performed on DAR, making sure that the prescribed SLA and QoS profiles are adhered to. In the upstream direction (from the customer CPE towards public Internet), all traffic policing is performed on the EPON ONU and OLT, making sure that the DIA circuit is not overloaded northbound from the OLT.

1.3 MEF services

MEF services are used to interconnect at least two remote customer locations across the converged transport network. Similar to DIA, the OLT does not operate in the routed mode, and all data generated by the customer at one site is transported to the other location in a dedicated L2 tunnel, less locally significant L2CP traffic. The said L2 tunnel is designed based on a pre-configured combination of [802.1Q] VLAN tags.

MEF services are typically provided to commercial customers via the primary UNI on the ONU, capable of reaching the maximum throughput of 1 Gbps. In practice, MEF services have been limited to around 700 Mbps due to costs associated with providing a dedicated 1G-EPON OLT port to a single customer. Any MEF services above 700 Mbps are provided today with dedicated P2P CWDM Ethernet solution. It is expected that in the future such circuits will be migrated into 10G-EPON, once it becomes widely deployed.

The ONU forwards frames into the MEF tunnel using UNI port as the frame classifier: all frames exchanged over primary UNI belong to MEF circuit. Customer frames may be either forwarded as received (transport mode) without further modifications in the ONU and the OLT (used when customer frames are VLAN tagged in a reliable fashion), or encapsulated with additional VLAN tags (encapsulation mode), where the scope of the added outer tag is limited to the converged transport network.

Figure 3 shows an example of a simple MEF service, with a L2 tunnel interconnecting two customer sites: SITE1 and SITE2. Dedicated MEF tunnels are built at both sites across ONU and OLT operating both in L2 mode only. There is no routing involved within deployed MEF circuits. The OLT forwards frames that belong to the MEF circuit into a specific uplink port connected indirectly to the SER, which then builds MPLS tunnel across the converged transport network into a remote SER connected to an OLT and an ONU serving the remote customer site. Traffic is allowed to flow between both customer sites without the need to reach public Internet, providing inherent security for customer data and lowering transport costs. In some cases, especially when customer sites are geographically distributed, it is necessary to use third-party transport network, in which case an MPLS transport circuit is deployed across a combination of a local converged transport network and third-party transport network, reaching the remote SER instance. Functionally, for the end customer there is no difference in both cases.

Note also that it is possible to design and deploy more complex architectures than a P2P MEF circuit, interconnecting multiple customer sites. All the configuration differences are at interconnection between individual SERs, where P2MP MPLS tunnels need to be deployed, allowing traffic to flow to specific destination within the MPLS-based LAN. Given the level of complexity of such a drawing, this scenario is not shown in Figure 3.

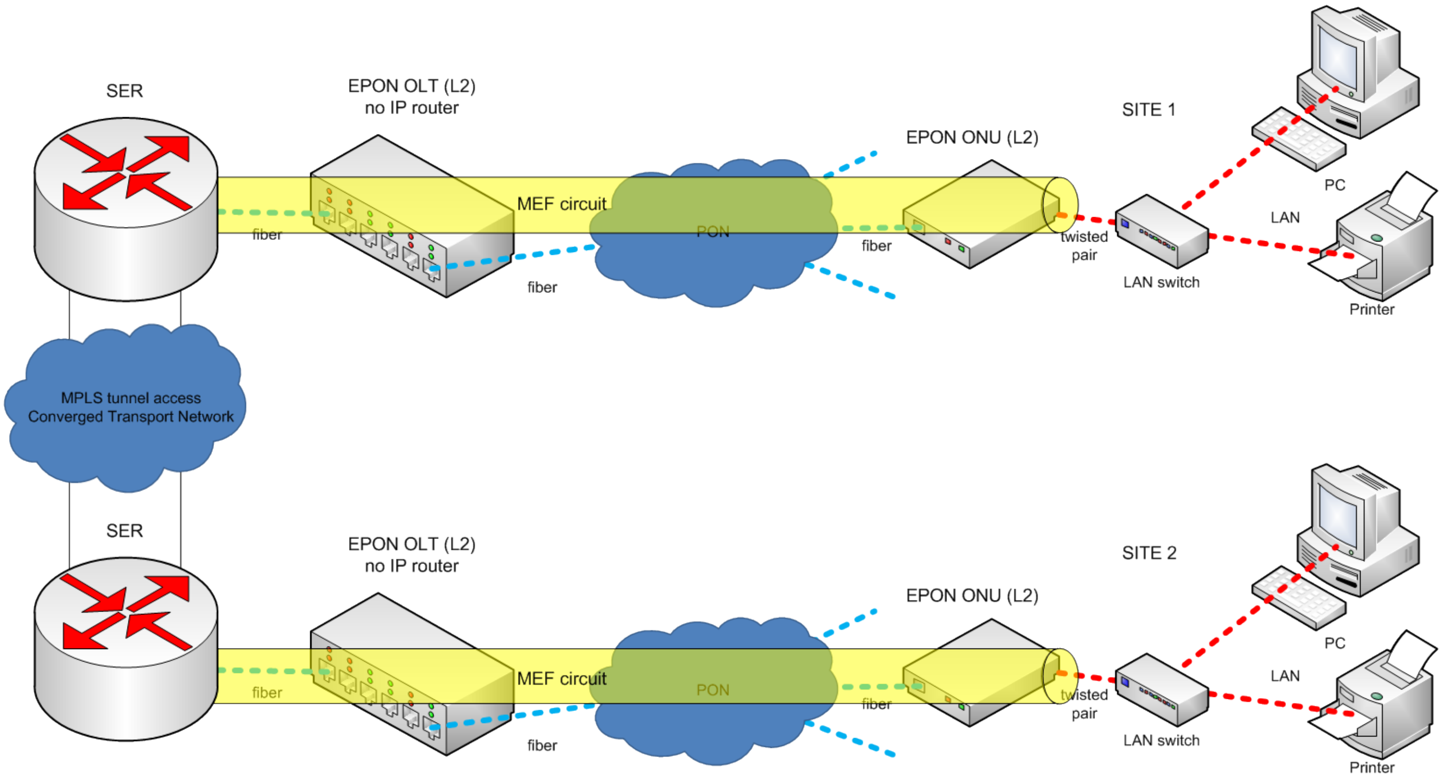

Figure 3: Architecture of MEF service

There are several types of MEF services that can be supported over the EPON, namely: E-LINE, E-TREE, and E-LAN. These individual MEF service types can be demonstrated using Figure 4 as the reference access network architecture.

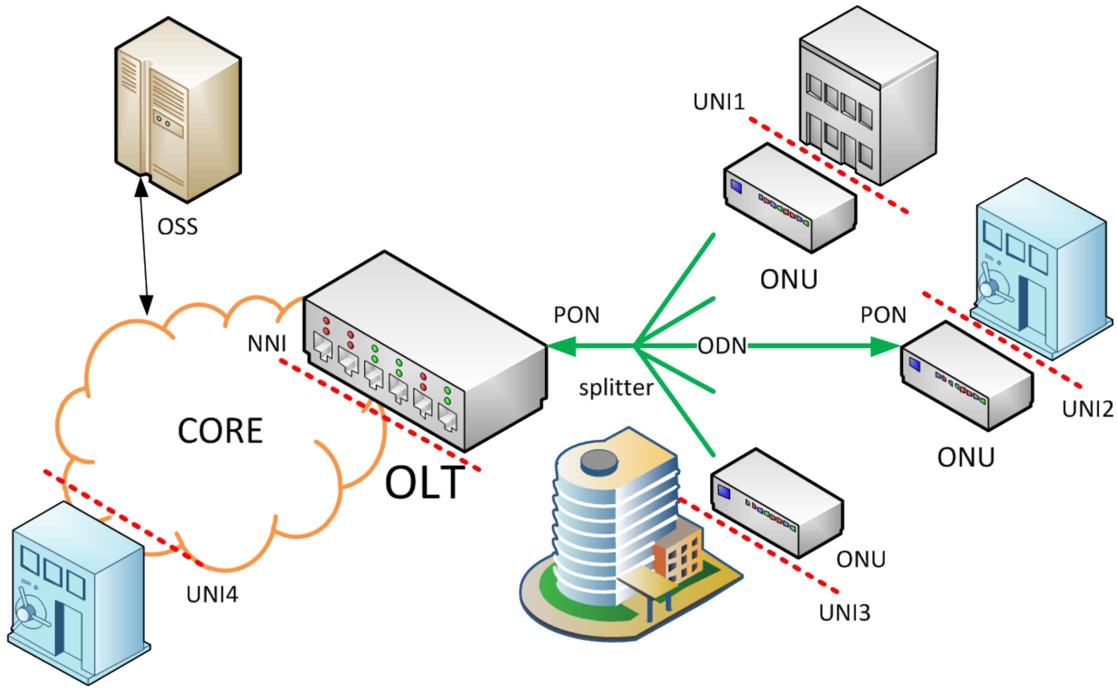

Figure 4: Reference scenario for description of MEF service types

An E-LINE service spans between one of UNIs (UNI1, UNI2, or UNI3) on any of the ONUs connected to the OLT, and the remote UNI4. Effectively, each E-LINE creates a simple P2P tunnel between selected MEF UNIs.

An E-TREE service creates a rooted multipoint service that connects a number of UNIs providing sites with hub and spoke multipoint connectivity. Each UNI is designated as either 'root' or 'leaf'. A root UNI can communicate with any leaf UNI, while a leaf UNI can communicate only with a root UNI. An E-TREE root could be located at the OLT, and individual leaf sites at UNI1, UNI2, or UNI3. Alternatively, the root could be located at the remote UNI4, and individual leaf sites at UNI1, UNI2, or UNI3.

An E-LAN service creates a multipoint to multipoint service that connects a number of UNIs (2 or more) providing full mesh connectivity for those sites. Each UNI can communicate with any other UNI that is connected to that Ethernet service. In the architecture shown in Figure 4, an E-LAN could be created among UNI1, UNI2, and remote UNI4.

Each MEF service type can be further sub-divided into Private and Virtual service. We have therefore:

- EP-LINE: Ethernet Private Line
- EV-LINE: Ethernet Virtual Line
- EP-LAN: Ethernet Private LAN
- EV-LAN: Ethernet Virtual LAN
- EP-TREE: Ethernet Private Tree
- EV-TREE: Ethernet Virtual Tree

In a Private service type, the end-user is provided with an exclusive dedicated, physical UNI port This means that an end-user is connected to a dedicated UNI port on an ONU.

In a Virtual service type, the end-user shares the selected physical UNI port with several other end-users, typically via means of an additional, external Ethernet switch connected to the demarcation UNI. This means that an end-user is connected to a port on an Ethernet switch with uplink connected to an UNI on an ONU.

From a provisioning perspective, the difference between private and virtual MEF services lies only in association of individual services (service flows) to individual UNIs on the ONU:

- multiple instances of virtual services are assigned to one and the same physical UNI, sharing its bandwidth;
- a single private service instance is always assigned to one dedicated physical UNI; no other service instances share this particular UNI

As of the time of writing, MSOs using legacy 1G-EPON platforms provide mostly Private Ethernet services, limited by the capabilities of the legacy EPON platforms. The introduction of second-generation EPON OLT into production is expected to enable larger deployment of Virtual Private Ethernet services, especially in locations where multiple commercial customers are co-located geographically and could be served with a single EPON ONU.

1.4 Public WiFi backhaul

The public WiFi backhaul service is very similar to the residential FTTH service (see Figure 5), in that a WiFi AP is treated as a CPE device connected to an EPON ONU. The ONU forwards all traffic to the OLT, which operates in the routed mode.

During the initialization phase, the WiFi AP retrieves an IP address from the assigned DHCP server and then uses the newly established L3 connectivity to communicate with the WiFi core controller. The said controller configures the WiFi AP with specific service parameters, including SSID, bandwidth profiles, etc. required for its proper operation. All the control data is transmitted in-band across public Internet. Once customer traffic begins to flow, all WiFi AP control traffic is intermixed with user data and routed accordingly by the SER.

The WiFi AP can be connected to the EPON ONU in several different ways, with two of them being used most commonly. In one arrangement, a stand-alone EPON ONU is used, and WiFi AP is then connected using standard twisted pair (at least CAT5e-class) cable. Alternatively, if the WiFi AP is equipped with an SFP port, an SFP ONU can be plugged directly into the WiFi AP, providing functionality of an EPON ONU in an SFP format. The second configuration is preferred for all new deployments, lowering the power consumption requirements, as well as minimizing the number of potential failure points in the network configuration.

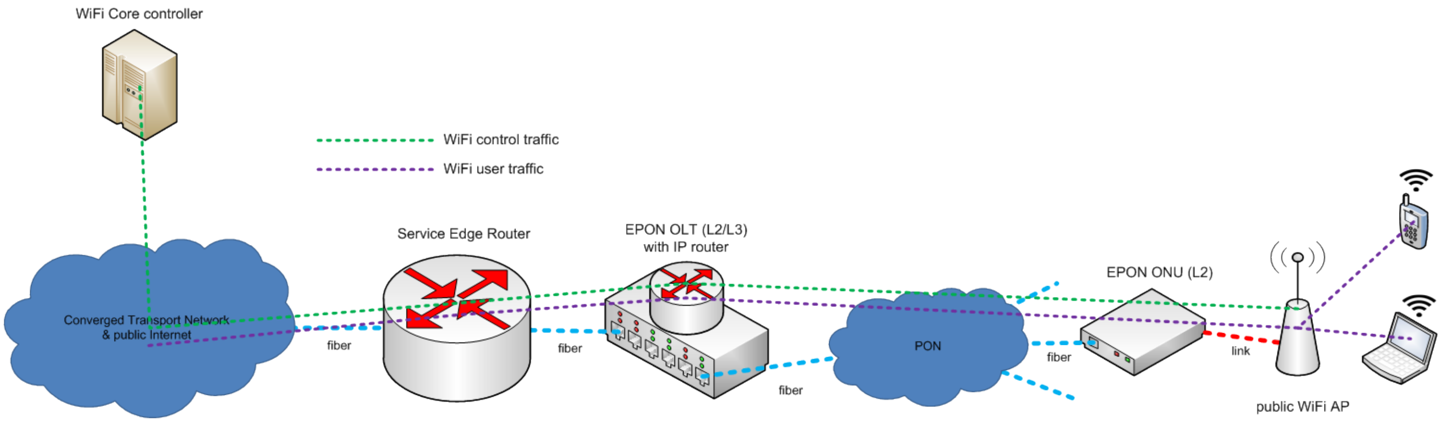

Figure 5: Architecture of public WiFi backhaul service

1.5 Cellular backhaul

The cellular backhaul is a very specific type of commercial service, where digital data generated by the radio interfaces on a cellular base station is then backhauled into a dedicated DWDM transport network within operator footprint. Apart from this one distinction, the service model is very similar to a P2P EV-LINE MEF circuit.

However, it is worth noting that for redundancy purposes, a NID located between the EPON ONU and the cellular tower (NID_A) creates two independent VLANs transported by the ONU and the OLT. One of the VLANs is injected into the primary DWDM transport ring, while the other VLAN is injected into the secondary DWDM transport ring. This arrangement creates effectively a sort of a redundancy north of the OLT, while not providing any protection against the failure of the ONU or the OLT itself. VLANs created by the NID_A are separated and injected into appropriate DWDM transport ring by NID_B, as shown in Figure 6.

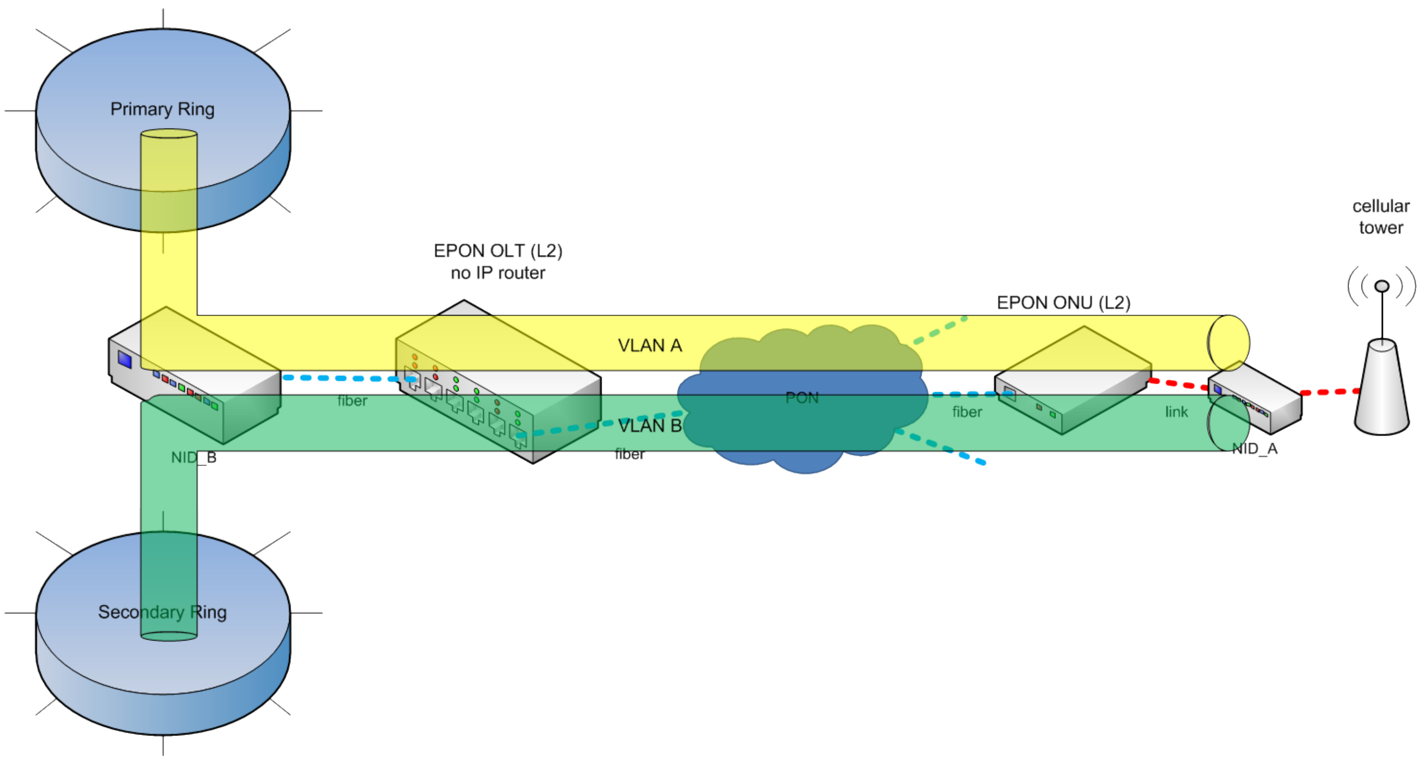

Figure 6: Architecture of cellular backhaul service

Along with the deployment of second-generation OLTs into production, this service model will be extended to take advantage of the redundancy features on the OLT itself. With some of the available OLT platforms, L2 protection mechanisms could be employed, providing additional protection for uplinks into individual DWDM transport rings. The selection of the target DWDM transport ring will still require the use of NIDs external to the ONU and the OLT, since an end-to-end (E2E) monitoring is required to decide which of the DWDM transport rings to use for CBH traffic. This decision cannot be taken autonomously by the OLT, since it does not have enough visibility into the condition of the E2E network. Only once the OLT and ONU introduce the support for [802.1ag] / [Y1731] service OAM mechanisms, NIDs will become effectively unnecessary, since the OLT and ONU will receive ability to monitor network status E2E and perform switching decisions based on the acquired network status.

1.6 Emerging applications

At the time of writing, emerging applications focus primarily on a more efficient use of access equipment, as well as new service types.

New service types are primarily associated with new business offerings for commercial customers, including remote management of their network infrastructure, managing edge routers, as well different applications for deployed MEF services. Network architecture-wise, all these applications take advantage of either P2P or MPLS-based P2MP services.

However, with the introduction of second-generation EPON OLTs into production, it will be possible to deploy P2MP MEF services in the form of EV-LAN or EVP-LAN directly on the OLT, without the need to reach out into the SER for switching between two separate P2P EV-LINE or EVP-LINE services. This is demonstrated in Figure 7.

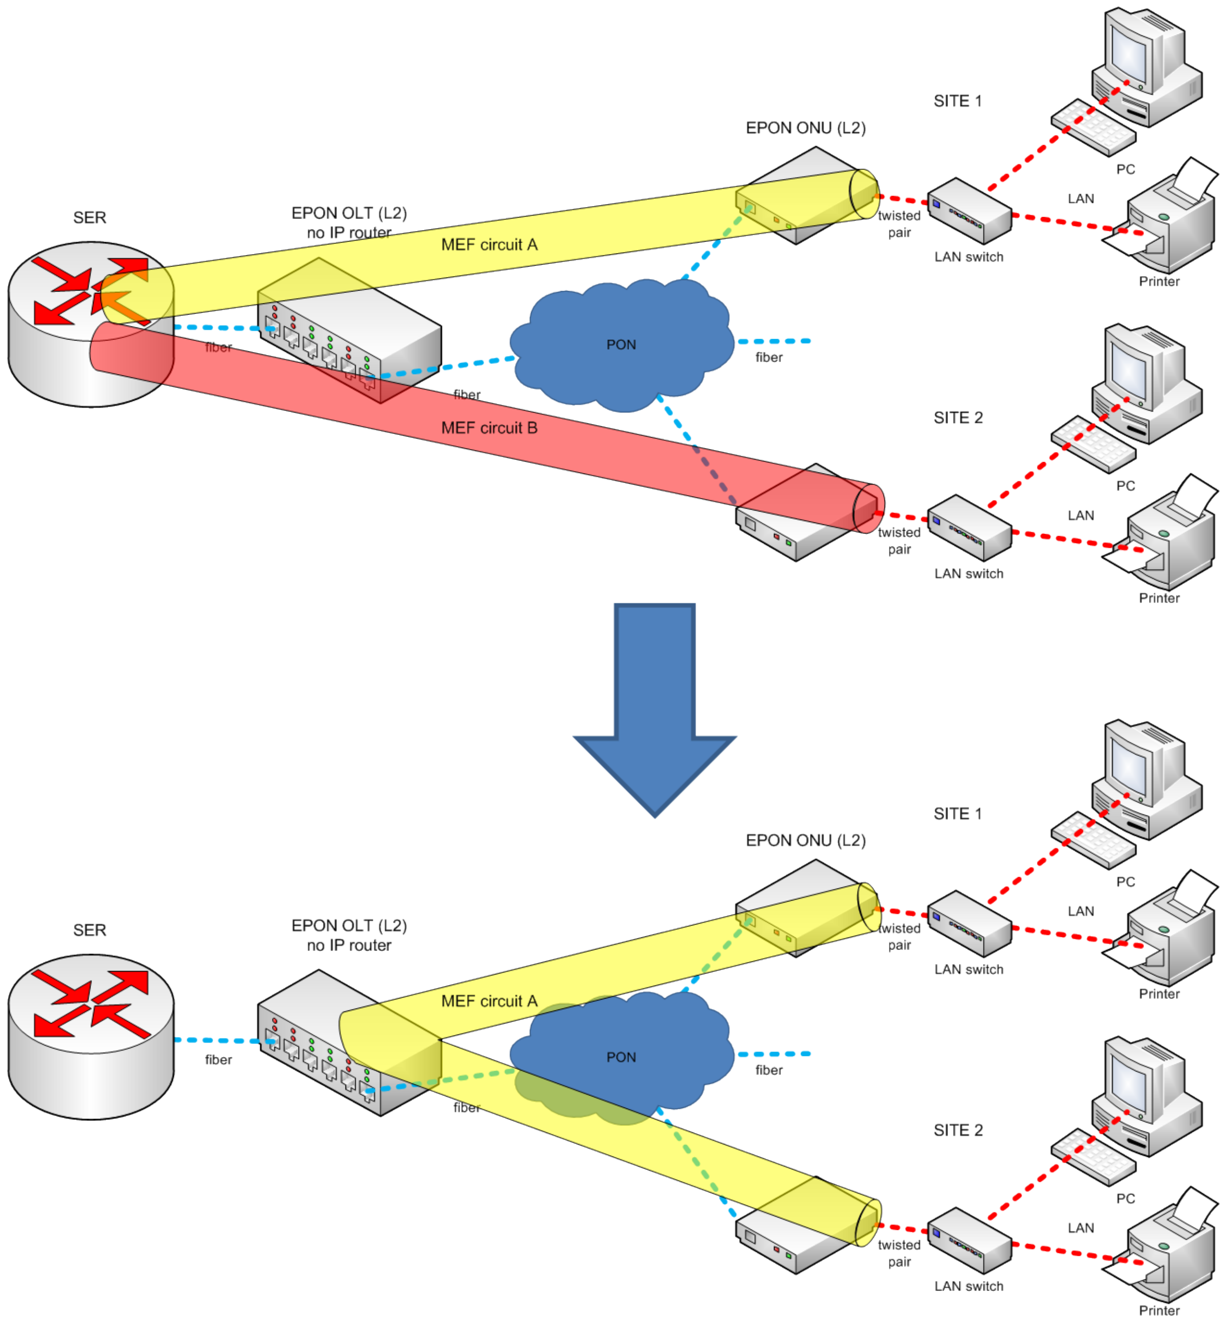

Figure 7: Two customer sites connected via two EV-LINE spokes (upper) and via EV-LAN on the OLT (lower)

Within the EV-LAN, customer traffic is still allowed to flow subject to specific QoS / SLA requirements. However, one EV-LAN is deployed in the access network, the local SER is completely offloaded from participating in data exchange between individual customer sites connected via one and the same OLT. SER interaction will still be needed when customer sites are geographically remote and communication between different OLT is required to build an EV-LAN, though substantial bandwidth and port load savings are expected with the introduction of local EV-LAN implementations.

In the near future, more complex MEF services may be also deployed, including EV-TREE architectures, though demand for such services at this time is very limited.