9 Service availability

9.1 Device and transceiver status monitoring and diagnostic functions

9.2 Definitions of events

9.3 Optical link protection

9.3.1 Introduction

This subclause defines optical link protection mechanisms, their functional description, and the associated OLT and ONU requirements. Two types of optical link protection are introduced, namely, a trunk protection (see 9.3.3) and a tree protection (see 9.3.4), each addressing a different application space and providing different types of functionality.

9.3.1.1 Terminology

In the remainder of this subclause, the terms *primary* and *backup* are used to describe the physical modules involved in the protection scheme whereas the terms *working* and *standby* describe the state of the physical modules. The working module refers to the module currently carrying subscriber traffic, and the standby module is not carrying subscriber traffic. During the actual switch event, both the primary and backup modules may be carrying active traffic, depending on the ~~actual~~ implementation; however, this condition is transient.

The switching time between the working OLT and the standby OLT is defined as the time between the last bit of the last envelope transmitted on the working OLT_MDI and the first bit of the first envelope transmitted on the standby OLT_MDI, assuming continuous flow of data to a single connected ONU. The time taken by the switching condition detection process is accounted for in the switching time. Note that the switching time measurement may not be accurate with multiple connected ONUs.

The switching time between the working L-ONU and the standby L-ONU is defined as the maximum time interval among the following:

- Time interval between reception of the last bit of the control message (*PON Interface Administrate* TLV, defined in 14.4.9.3) by the working L-ONU, requesting the ONU to perform switchover, and the first bit of the PLID envelope transmitted by the standby L-ONU and containing a *REPORT* MPCPDU reflecting the nonzero queue length.
- Time interval between the detection of loss of signal by the working L-ONU and the first bit of the PLID envelope transmitted by the standby L-ONU and containing a *REPORT* MPCPDU reflecting the nonzero queue length.
- Time interval between the reception of the first bit of a data frame by the standby L-ONU and the first bit of the PLID envelope transmitted by the standby L-ONU and containing a *REPORT* MPCPDU reflecting the nonzero queue length.

The above time intervals are measured under continuous flow of data to a single connected ONU.

9.3.2 Device architecture and requirements

EPON devices should support optical link protection. If optical link protection is ~~supporter~~supported, the EPON devices incorporate a protection switch function in specific locations in the MAC Client allowed for by the model defined in IEEE Std 1904.1, Clause 6 and instantiate the appropriate number of OAM and MAC Control Clients.

Optical link protection mechanisms are defined in 9.3.3 and 9.3.4. Specific requirements for the ONU and OLT devices are different, depending on the type of supported protection mechanism.

9.3.2.1 Line and Client protection

This subclause describes Line ONU/OLT protection and Client ONU/OLT protection schemes and their representation in the MAC Client reference model. In both cases, the operation of the protection function is controlled and provisioned via mechanisms specified in the MAC Client reference model.

Functional blocks within the MAC Client reference model may be categorized into two groups, based on their logical behavior:

- Functional blocks with *combinatorial logic*, where the output of a functional block only depends on the input. Such functional blocks are marked with white boxes in Figure 9-2 and Figure 9-3. Input/Output, Modifier, and CrossConnect implement combinatorial logic.
- Functional blocks with *sequential logic*, where the status of the output of a functional block depends on the status of the input, past history, or internal state of the block. Such functional blocks are shown as shaded boxes in Figure 9-2 and Figure 9-3. Classifier, Queues, Policer/Shaper, and Scheduler implement sequential logic.

Device behavior during the protection switchover event and its impact on data traffic are different depending on whether the sequential logic blocks are duplicated or shared among the primary and backup ESPs. These behavioral differences are detailed below. In the Line ONU/OLT protection scheme, all the sequential logic blocks are shared, while in the Client ONU/OLT protection scheme all the sequential logic blocks are duplicated.

9.3.2.1.1 Line device protection

In the case of Line ONU/OLT protection, the protection switch function is located between the Input/Output ports connected to the `MADI`/`MADR` primitives, as shown in Figure 9-2. The Classifier, Modifier, Policer/Shaper, Queues, and Shaper blocks are shared among the primary and backup ESPs, providing the required redundancy for the Line device elements only.

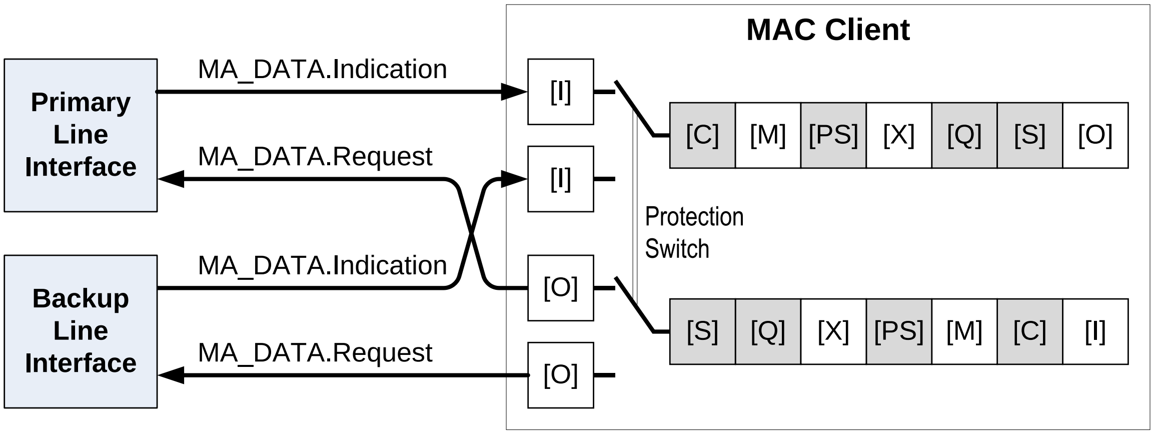

Figure 9-2—Line device protection architecture

After the switchover event, data stored in the Queues block of the primary path is available for the backup path as well, preventing data loss. Likewise, since the Policer/Shaper and Scheduler blocks are shared between the primary and backup paths, the size of the generated data burst (in the case of ONUs) does not exceed the maximum burst size allowed by Policer/Shaper parameters provisioned for the given device. In this way, the MAC Client maintains the state of individual functional blocks before and after the protection switching takes place.

9.3.2.1.2 Client device protection

In the case of Client ONU/OLT protection, the protection switch function is instantiated between the Input/Output ports connected to the UNI ports of the ONUs and NNI ports of the OLT, as shown in Figure 9-3. The Classifier, Modifier, Policer/Shaper, Queues, and Shaper blocks are duplicated, providing the required redundancy for the Client device elements.

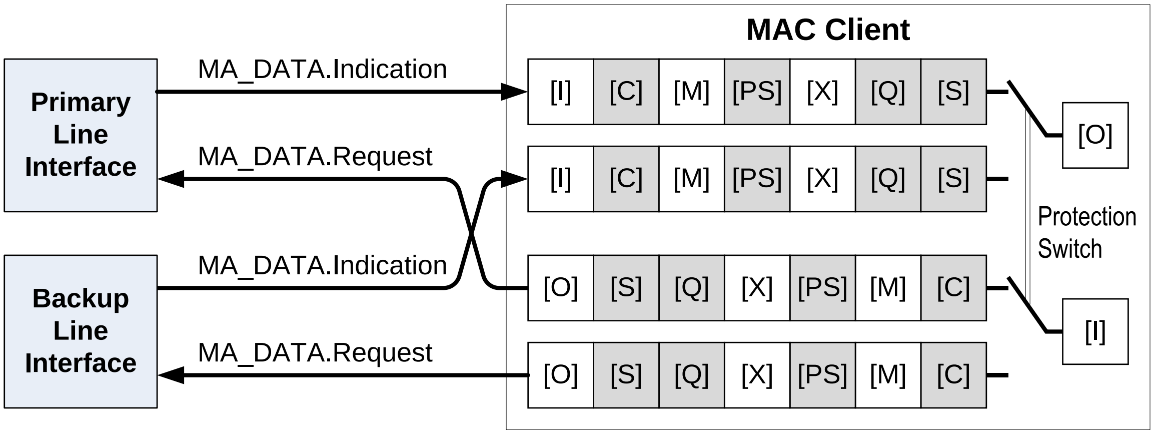

Figure 9-3—Client device protection architecture

After the switchover event, data stored in the Queues block of the primary path is not available for the backup path, allowing some data loss for frames already stored in the Queues block for the primary path. Likewise, since the Policer/Shaper blocks are duplicated, the size of the generated data burst can be double what was provisioned (the primary and backup Policer/Shaper each may admit maximum-size burst, right before and right after the protection switchover event).

9.3.2.2 Line fault detection

Both the working C-OLT and the working C-ONU observe the status of the working optical link, using available mechanisms for detection of the link failure in the upstream and downstream directions. The link failure detection may take place at both ends of the link or on one end of the link only. In either case, the side detecting link failure notifies the link peer about this event using the messages and mechanism specific to a given protection scheme.

9.3.2.2.1 OLT detection conditions

The working OLT shall be able to detect the fault condition on the working optical line using each of the mechanisms defined below. Each of these mechanisms is sufficient to indicate a fault condition.

a) Optical LoS: the working OLT fails to receive valid optical signal from multiple granted ONUs within $T_{LoS\_Optical}$ (two milliseconds by default), as identified by the Signal Detect Threshold value in IEEE Std 802.3, Table 141-17 and Table 141-18.

b) MAC LoS: the working OLT fails to receive ~~a *REPORT* MPCPDU or any other~~any MAC frame from ~~the~~ any ~~L~~ ONU ~~in~~ within a T_{LoS_MAC} window (50 ms by default).~~eight (see the MISSED_REPORT_LIMIT constant defined in IEEE Std 802.3, 144.3.7.1) consecutive grants for which the `ForceReport` flag in the corresponding *GATE* MPCPDU was set to `1`.~~ To avoid a false MAC LoS detection due to all ONUs being idle, the working OLT is expected to request at least eight *REPORT* MPCPDUs from every registered ONU within each T_{LoS_MAC} window. To request a *REPORT* MPCPDU, the OLT generates a *GATE* MPCPDU with the `ForceReport` flag associated with the PLID grant set to 1 (see 8.4.1.5.3). ~~Note that the time required to detect MAC LoS condition depends on the frequency of *REPORT* MPCPDUs. If it is desirable to detect MAC LoS within a T_{LoS_MAC} window (50 ms by default), the OLT is expected to generate at least one *GATE* MPCPDU (with `ForceReport` flag set to `1`) per ONU every 0.125 × T_{LoS_MAC} ms (6.25 ms by default).~~

The OLT may also use the signal quality degradation metrics, which identify whether the power of the received optical signal exceeds the lowest/highest threshold or the BER of the received signal exceeds a certain operator-defined threshold. The values of $T_{LoS\_Optical}$, T_{LoS_MAC}, the time thresholds of optical signal loss, and BER on the OLT side are configured via NMS and are outside the scope of this standard.

9.3.2.2.2 ONU detection conditions

The ONU shall detect the fault condition on the working optical line using any of the mechanisms defined below:

a) Optical LoS: the ONU fails to receive any valid optical signal within $T_{LoS\_Optical}$ (two milliseconds by default), as identified by the Signal Detect Threshold value in IEEE Std 802.3, Table 141-21 and Table 141-22.

b) MAC LoS: the working ONU fails to receive a *GATE* MPCPDU or any other frame from the OLT within T_{LoS_MAC} (50 ms by default). Note that to avoid a situation where a single lost *GATE* MPCPDU causes a protection switchover, the OLT is expected to generate at least one *GATE* MPCPDU to the ONU every 0.125 × T_{LoS_MAC} ms (6.25 ms by default).

The values for $T_{LoS\_Optical}$ and T_{LoS_MAC} parameters ~~should be~~are configured using the eOAM attribute *aOnuConfigProtection* (see 14.4.9.2)~~messages specific to a given protection scheme~~.

9.3.3 Trunk protection scheme

In the trunk protection scheme, the ODN span between the C-OLT and the 2:*N* optical splitter, used to join the two trunk segments, is protected. The C-ONU and the branch fiber (ODN span between the splitter and the ONU) are not protected. There are two types of trunk protection schemes, as shown in Figure 9-4 and Figure 9-5.

Figure 9-4 presents a trunk protection scheme with redundant L-OLT and trunk segments. In this scheme, the MAC, MAC Control, and OAM Clients in the C-OLT are shared by the primary and the backup L-OLTs and are not protected against failures. This trunk protection scheme reduces the implementation cost and targets protection only against the failures having highest potential impact: OLT optical transceiver failures and trunk fiber cuts. In this scheme, the OLT uses a line protection architecture (see 9.3.2.1.1).

The trunk protection with redundant L-OLT scheme supports only the *intra-chassis* protection scheme, where the primary L-OLT and backup L-OLT are located within the same chassis (either on the same line card or on separate line cards).

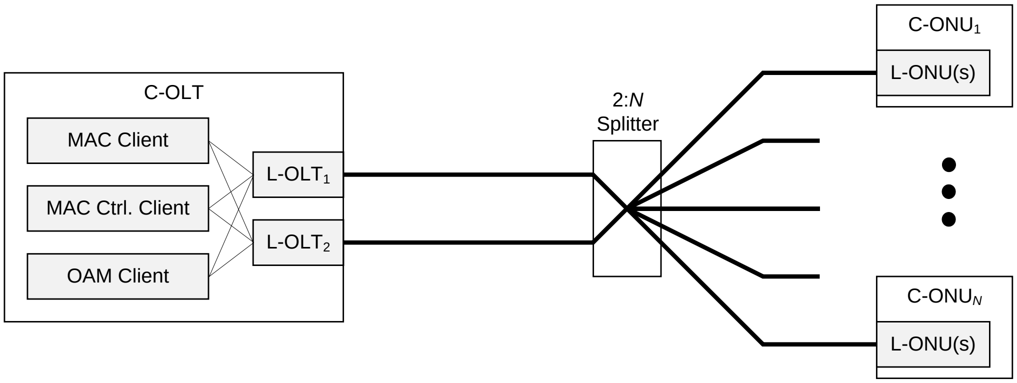

Figure 9-4—Trunk protection with redundant L-OLT

An alternative configuration of the trunk protection scheme is shown in Figure 9-5. This scheme provides added robustness as the whole C-OLT is duplicated, including the L-OLT and all MAC Clients.

In addition to intra-chassis protection, the trunk protection with redundant C-OLT scheme supports the protection, where the primary C-OLT and backup C-OLT are located in different chassis (either within the same central office or at geographically different locations). The inter-chassis protection scheme requires coordination of the protection states and functions among the primary and backup C-OLTs comprising the trunk protection group and may require communication over LANs and/or wide area networks (WANs) using public or proprietary protocols. The nature of information, data formats, and communications protocols used to coordinate protection functions among the primary and backup C-OLTs are outside the scope of this standard.

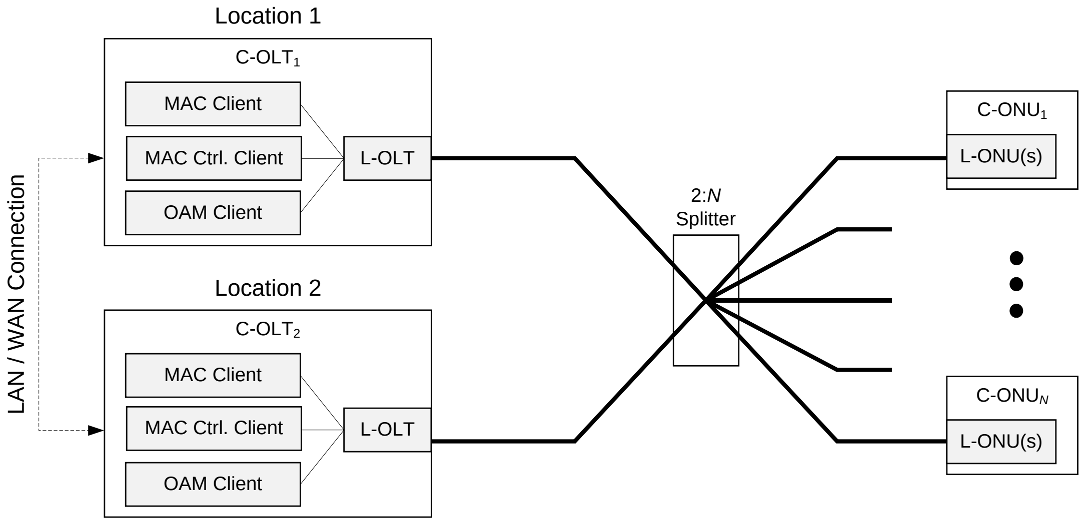

Figure 9-5—Trunk protection with redundant C-OLT

~~In the trunk protection scheme, the backup C-OLT acquires the round-trip time (RTT) values for individual ONUs without executing the MPCP discovery and registration process. Two possible approaches to acquire RTT are covered in Annex 9A; however, their selection and implementation details are outside the scope of this standard.~~

~~There are no ONU configuration differences between the trunk protection schemes shown in Figure 9-4 and Figure 9-5. Connected C-ONUs are configured in precisely the same manner.~~

~~The following subclauses provide technical requirements for the C-ONU and C-OLT devices participating in this protection scheme.~~

9.3.3.1 Trunk protection switching procedure

The protection switching procedure in the trunk protection scheme may be triggered in the following ways:

- Automatically, when both the OLT and the ONU detect the fault condition on the working optical line using any of the mechanisms specified in the following subclauses; or
- On-demand, when the OLT is requested by the NMS to switch to the standby path. This protection switch is executed typically for operational reasons, e.g., fiber repairs, maintenance of OLT cards.

In a 50G-EPON, a link failure detected on any 25Gb/s channel causes both channels to switch from the primary to the backup OLT.

The primary trunk fiber and the backup trunk fiber are assumed to follow disjoint paths and therefore to be of different lengths. Consequently, the ONU round-trip times (RTTs) observed by the primary L-OLT would be different from the RTTs observed by the backup L-OLT.

The encryption method specified in 11.3 relies on synchronization of *CipherClock* in the OLT with the *TxCipherClock* and *RxCipherClock* in an ONU (see 11.3.5.4.1). This synchronization, in turn, depends on the RTT of a given ONU, therefore, the encryption established between the primary L-OLT and an ONU cannot continue to operate between the backup L-OLT and the same ONU. When a protection switching event is triggered, the OLT shall disable the downstream encryption. If the ONUs have not detected the trunk failure condition independently, OLT's disabling of the downstream encryption causes ONUs to disable the upstream encryption and also to suspend all user traffic, as explained in 11.3.7.2.

9.3.3.1.1 Default procedure

In the event of trunk failure, the default trunk protection switching behavior is for all the ONU to execute the full start-up sequence, as illustrated in Figure 11-1, including the MPCP and OAM discoveries, ONU authentication, and establishment of encryption.

To ensure that all ONUs respond to the MPCP discovery performed by the backup OLT, the backup OLT should explicitly deregister all ONUs by issuing the *REGISTER* MPCPDU(s) with the *Flag* field equal to NACK (see IEEE 802.3, 144.3.6.4). An individual MPCPDU may be sent to each registered ONU or a single MPCPDU may be broadcast to all ONUs by sending it in an envelope addressed to the broadcast PLID (see *BCAST_PLID* in IEEE 802.3, 144.3.5) and with the MAC destination address 01–80–C2–00–00–01.

9.3.3.1.2 Optimized procedure bypassing MPCP and OAM discovery

Some OLT implementations may improve upon the default behavior described above by being able to exchange ONU-related non-security information, e.g., ONU's identities, measured RTTs, and ONU configurations. Such capabilities may allow a faster protection switching procedure that bypasses the MPCP and OAM discoveries (see Step 1 in Figure 11-1), and possibly the OAM provisioning (see Step 6 in Figure 11-1).

In case of trunk protection scheme with redundant L-OLT (see Figure 9-4), the shared MAC Control client ensures that the ONU identities and the RTT values measured by the primary L-OLT are also available to the backup L-OLT. Since only the trunk lengths are different, the backup L-OLT needs to apply a fixed offset to each ONU's RTT measured by the primary L-OLT. This fixed offset may be administratively provisioned onto the backup L-OLT or it can be measured after the switchover by re-ranging (i.e., direct reregistration, see IEEE Std 802.3, 144.3.7) of a single ONU and subtracting the ONU's RTT measured by the primary L-OLT from the RTT measured by the backup L-OLT.

Once the new RTT values are determined, the OLT may skip the MPCP registration and OAM/eOAM discoveries and proceed directly to ONU authentication (see 11.2), re-synchronizing of the *CipherClock* in the OLT with the *TxCipherClock* and the *RxCipherClock* in an ONU (see 11.3.5.4.1.3), and re-establishment of the initial encryption key (see 11.3.2).

Note that the above optimization may also be used with the trunk protection scheme with redundant C-OLTs. However, in this case, the identities and RTTs of all the ONUs are not automatically available to the backup C-OLT and need to be explicitly shared. The LAN/WAN communication channel used to share this information between the primary and backup C-OLTs is outside the scope of this standard.

This optimized protection switching procedure requires the ONUs to be provisioned with a sufficient holdover period to ensure that the ONUs do not auto-deregister while the standby L-OLT or C-OLT is being activated. The holdover period parameter is configured using the eOAM attribute *aOnuConfigHoldoverPeriod* (see 14.4.9.4).

9.3.3.1.3 Optimized procedure bypassing ONU authentication

This method of trunk protection switching optimization bypasses the ONU authentication in addition to bypassing the MPCP and OAM discoveries, as described in 9.3.3.1.2.

In the trunk protection scheme with redundant L-OLT, the instances of MAC, MAC Control, and OAM Clients are shared among the primary and the backup L-OLTs. This implies that the ONU authentication data is readily available to both the primary L-OLT and the backup L-OLT.

In addition to authenticating the ONUs, the EAP_TLS1.3 authentication method allows both the OLT and each ONU to derive a shared secret that is used to establish the initial encryption key, as described in 11.3.2. Bypassing the ONU authentication during the protection switching event also implies that the establishment of the initial key is also bypassed. This optimized trunk protection switching method assumes that the session encryption keys for every encryption entity are available to both the primary and the backup L-OLTs. This allows the standby L-OLT to restart the encryption/decryption processes by simply re-activating the encryption keys that where distributed to the ONUs by the working L-OLT prior to the protection switching event.

The protection switching optimization method that bypasses the ONU authentication is not recommended for the trunk protection scheme with redundant C-OLT, as it would require transfer of ONU-related security information across a LAN or a WAN.

~~9.3.3.1 Functional requirements~~

~~Trunk protection in EPON requires the following basic functionalities:~~

- ~~Ability to measure the standby path RTT (bRTT) in a way that does not affect live services. While the measurement of bRTT through re-registration of the affected ONUs is certainly possible, it would have a high impact on services (i.e., longer interruption) and is, therefore, not recommended. Annex 9A presents examples of dynamic bRTT measurement mechanisms.~~
- ~~Ability to switch the ONU between working and standby paths in under 150 ms, thus minimizing the impact on the operating services.~~

~~The process of protection switching in the trunk protection scheme may be executed in the following ways:~~

- ~~Automatically, when both the OLT and the ONU detect the fault condition on the working optical line using any of the mechanisms specified in the following subclauses; or~~

~~On-demand, when the OLT is requested by the NMS to switch to the standby path. This protection switch is executed typically for operational reasons, e.g., fiber repairs, maintenance of OLT cards.~~

~~9.3.3.1.1 C-OLT requirements~~

~~In the trunk protection mechanism, as defined in 9.3.3, the OLT is connected to two optical links, the primary and backup link, from which only one is active at any time, carrying OAM, CCP, eOAM, and MPCP control frames together with subscriber traffic.~~

~~The protection function present in the Operation, Administration, and Management block is responsible for switching between the primary and backup paths for subscriber and control frames.Both the Line ONU/OLT protection and Client ONU/OLT protection schemes can be supported by the OLT in the case of trunk protection, as defined in 9.3.2.1.1 and 9.3.2.1.2 respectively.~~

~~The protection function additionally instantiates the state diagram per Figure 9-6, controlling the operation of the MAC Client and L-OLTs.~~

~~The standby OLT and the working OLT participating in the trunk protection group exchange configuration details continuously, i.e., the standby OLT is informed by the working OLT of any changes in its configuration, registration, and authentication status of individual ONUs, etc. The specific method of communication between the standby OLT and the working OLT is outside the scope of this standard.~~

~~The primary OLT and backup OLT may be provisioned by the NMS in exactly the same manner, using the same configuration parameters, except that one OLT is provisioned as working and the other OLT is provisioned as standby to minimize the preparation time for the backup OLT to come online after the protection switchover.~~

~~The standby OLT may be in a cold standby mode or in a warm standby mode. In the first mode, the standby OLT remains powered off until protection switching is requested. In the second mode, the standby OLT remains powered on with minimum functions enabled and operational, i.e., the OLT has the capability of receiving and parsing upstream transmissions from the PON branch it is connected to, but does not send any data downstream. It is recommended that the standby OLT operates in the warm standby mode to facilitate fast response times to the optical protection switchover events:~~

- ~~Electronic subsystems are fully operational.~~
- ~~Optical subsystem is partially active, i.e., transmitter is powered down (no downstream transmission is needed), while receiver is powered up (receiving upstream transmissions from connected ONUs).~~

~~Assume the initial state of the network is such that the primary OLT is in the working state and the backup OLT is in the standby state in the following discussion. The working OLT monitors the status of the optical line according to 9.3.2.2.1, and once any of the line fault conditions are detected, the working OLT disables its optical transmitter and stops any downstream transmission. The working OLT then causes the standby OLT to enter into the full operating mode. The switching time between the working OLT and the standby OLT shall be lower than or equal to 150 ms. The definition of the switching time can be found in 9.3.1.1.~~

~~To complete the switchover process, the working OLT informs the NMS about the switchover event. It is recommended that the time between the events of the working OLT's switching its laser off and the standby OLT's switching its laser on be at least equal to the largest value of $T_{LoS\_Optical}$ for all connected ONUs, as defined in 9.3.2.2.2, to guarantee that connected ONUs can properly detect the line fault condition.~~

~~Information, notification, and alarms/warnings delivered by the working OLT to the NMS in the switchover condition, as well as message formats, are outside the scope of this standard.~~

~~The new working OLT, once the switchover process is complete, shall send one or more *GATE* MPCPDUs to force each registered ONU to resynchronize to the MPCP clock. The transmission of such *GATE*~~

~~MPCPDU is recommended to take place as soon after the end of the switchover event as possible to minimize frame loss during the switchover event.~~

~~9.3.3.1.2 C-ONU requirements~~

~~In the trunk protection mechanism, as defined in 9.3.3, the ONU is connected to a single optical link. In this case, the C-ONU does not contain primary and backup ESPs and typically remains registered throughout the switchover event. All the necessary changes take place on the OLT side, and the ONU is required only to suspend upstream transmissions for a specific period of time and remain in the HOLD_OVER_START state (per Figure 9-7) until a *GATE* MPCPDU is received.~~

~~As a result, only one instance of the OAM Client and MAC Control Client is needed on the ONU side, and the protection function present in the Operation, Administration, and Management block instantiates the state diagram per Figure 9-7, controlling the operation of the MAC Client and L-ONU(s).~~

~~Upon detection of a line fault, the C-ONU enters the HOLD_OVER_START state per Figure 9-7, where all currently stored upstream envelope descriptors are purged and the transmission of data from the ONU to the OLT is suspended. All incoming subscriber upstream data frames are queued. Frame loss is allowed in trunk protection when the local ONU queues overflow.~~

~~The C-ONU leaves the HOLD_OVER_START state upon the reception of the first *GATE* MPCPDU after entering the HOLD_OVER_START state. The upstream transmission is resumed using the newly allocated upstream transmission slots.~~

~~If the C-ONU fails to receive the *GATE* MPCPDU within the provisioned duration of the HOLD_OVER_START state (expressed by the `periodHoldOver` variable), the ONU enters the local deregistration state by sending the `MACR( REGISTER_REQ, status)` primitive with the `status` parameter set to NACK, per Figure 9-7. The OLT deregisters the ONU independently, based on the observed link status.~~

9.3.3.2 Trunk ~~switching~~ protection process

The ONU trunk protection process is defined via a state diagram shown in Figure 9-6. This process supports the default switching procedure (see 9.3.3.1.1), as well as the optimized switching procedures that bypass the MPCP and OAM discovery (see 9.3.3.1.2) or that bypass the ONU authentication (see 9.3.3.1.3).

~~9.3.3.1.3~~9.3.3.2.1 Variables

~~`backupLoS`~~

~~TYPE: Boolean~~

~~This variable indicates whether the MAC LoS or optical LoS condition is observed by the backup L-OLT, as defined in 9.3.2.2.1, or by the backup L-ONU (only in tree protection case), as defined in 9.3.2.2.2. The value of `true` indicates that the LoS condition is observed, and `false` indicates that the LoS condition is not observed. By default, this variable has the value of `false`.~~

`primaryLoS`

TYPE: Boolean

This variable indicates whether the MAC LoS or optical LoS condition is observed ~~by the primary L-OLT, as defined in 9.3.2.2.1, or~~ by the primary L-ONU, as defined in 9.3.2.2.2. The value of `true` indicates that the LoS condition is observed, and `false` indicates that the LoS condition is not observed. By default, this variable has the value of `false`.

periodHoldo~~O~~ver

TYPE: 32-bit unsigned integer

This variable represents the maximum period of time that the ONU may remain in the HOLD_OVER_START state. If the ONU does not receive at least one *resynchronization GATE* MPDPDU within the periodHoldo~~O~~ver, it deregisters. This variable is expressed in units of milliseconds, and its value is provisioned using the *aOnuConfigHoldoverPeriod* attribute (see 14.4.9.4).

registered

TYPE: Boolean

This variable holds the ONU's current ~~result of the discovery process~~registration status. ~~It is set to `true` once the discovery process is completed and registration is acknowledged.~~ This variable maps to the variable `Registered` defined in IEEE Std 802.3, 144.3.7.3~~.~~, and its value is controlled by the ONU Registration state diagram (see IEEE Std 802.3, 144.3.7.8).

~~9.3.3.1.4~~9.3.3.2.2 Timers

timerHoldo~~O~~ver

This timer is used to force the ONU leave the HOLD_OVER_START state if the period of time spent in this ~~the HOLD_OVER_START~~ state is longer than the provisioned value of periodHoldo~~O~~ver. Once this timer expires, the ONU deregisters.

~~9.3.3.1.5~~9.3.3.2.3 Functions

`clearCommittedEnvelopes()`

This function deletes all the envelope descriptors stored in `EnvList[0]` and `EnvList[1]` (see IEEE Std 802.3, 144.3.8). These envelopes have been scheduled by the working L-OLT before the protection switching event and are not valid or anticipated by the standby OLT. This function code is equivalent to `{EnvList[0].Clear(); EnvList[1].Clear()}`.

`deregisterRequest()`

This function causes the ONU MPCP client to issue an auto-deregistration request toward the ONU Registration state diagram (see IEEE Std 802.3, 144.3.7.8). This is equivalent to generating the primitive `MCSR(MsgRegisterReq)` with `MsgRegisterReq.Flag = NACK`. Upon the reception of this primitive from the MPCP client, the ONU Registration state diagram transitions from REGISTERED to LOCAL_DEREGISTER state and then unconditionally to the UNREGISTERED state.~~activateDataPath(portId)~~

~~This function controls the flow of subscriber data frames egressing the port identified by portId parameter, which can take the following values:~~

- ~~primaryPonIF identifies the primary physical port associated with OLT_MDI or ONU_MDI.~~
- ~~backupPonIF identifies the backup physical port associated with OLT_MDI or ONU_MDI.~~
- ~~primaryMacPort identifies the primary virtual port associated with OLT_LI. This port is identified by an LLID.~~
- ~~backupMacPort identifies the backup virtual port associated with OLT_LI. This port is identified by an LLID.~~

~~When the function is called with the argument set to `primaryPonIF` or `primaryMacPort`, the identified primary port becomes the working port, and the corresponding backup port becomes the standby port. Similarly, when the function is called with the argument set to `backupPonIF` or `backupMacPort`, the identified backup port becomes the working port, and the corresponding primary port becomes the standby port. Implementations may choose to accomplish this switching by modifying the rules or reconfiguring the association between the CrossConnect entries and the queues.~~

~~`opticalTX( portId, param )`~~

~~This function controls the status of the optical transmitter associated with the port identified by `portId` parameter. The `portId` parameter can take values as defined in the `activateDataPath( portId )` function. When the `param` variable has the value of `enable`, the optical transmitter is enabled, allowing the data transmission across the OLT_MDI or ONU_MDI. When the `param` variable has the value of `disable`, the optical transmitter is disabled (either powered down or disabled administratively), resulting in no frames being transmitted across the OLT_MDI or ONU_MDI.~~

~~`purgeGrants()`~~

~~This function causes an L-ONU to discard all stored (pending) grants.~~

~~`sendResyncGates( portId )`~~

~~This function is responsible for transmission of refresh *GATE* MPCPDUs to all L-ONUs connected to a port identified by `portId` parameter. The `portId` parameter can take values as defined in the `activateDataPath( portId )` function. Reception of these *GATE* MPCPDUs forces the ONUs to leave the HOLD_OVER_START state, as defined in Figure 9-7.~~

~~9.3.3.1.6~~9.3.3.2.4 Primitives

`MACI(` ~~`GATE`~~ `resyncGate )`

The acronym `MACI` is defined in 3.4. This primitive represents a reception of a ~~non-discovery~~resynchronization *GATE* MPCPDU at the ONU. and is equivalent to the `MCSI(MsgGate)` primitive defined in IEEE Std 802.3, 144.1.4.1 and 144.3.6.1. The resynchronization GATE is generated by the standby L-OLT and its purpose is to resynchronize the ONU’s MPCP clock to the standby L-OLT MPCP clock. The resynchronization GATE is sent in an envelope addressed to the ONU’s unicast PLID and has its timestamp precompensated by the ONU’s RTT, as measured by (or otherwise known to) the standby L-OLT. The MPCP clock synchronization is performed by the ONU’s Control Parser defined in IEEE Std 802.3, 144.2.1.5.

`MACI( switchGate )`

The acronym `MACI` is defined in 3.4. This primitive represents the reception of a special *GATE* MPCPDU that indicates to the ONU that the OLT has activated the standby L-OLT (i.e., the previously working L-OLT became the standby, and the L-OLT previously in standby mode became the working L-OLT). This primitive is equivalent to the `MCSI(MsgGate)` primitive defined in IEEE Std 802.3, 144.1.4.1 and 144.3.6.1, however the switch-indicating GATE does not contain any `EnvAlloc` structures. The `switchGate` message is sent in an envelope addressed to the broadcast PLID (*BCAST_PLID*, see IEEE Std 802.3, 144.3.5) and as such, it does not trigger the MPCP time adjustment by the ONU.

~~`MACR(REGISTER_REQ, status )`~~

~~The `MACR` acronym is defined in 3.4. This primitive represents the transmission of a *REGISTER_REQ* MPCPDU by the ONU and is equivalent to the `MCSR(MsgRegisterReq)` primitive defined in IEEE Std 802.3, 144.1.4.1 and 144.3.6.3. The `status` parameter sets the value of the `MsgRegisterReq.Flags` field.~~

~~`NMSI( messageId, failureCode )`~~

~~This primitive is used to inform the NMS about the protection switching event, during which the previously working and standby L-OLTs exchange their functions. It uses the following parameters:~~

- ~~`messageId` identifies whether the switching event was initiated by the OLT or the ONU and what the new working port is. The following messages are defined:~~
- ~~`MSG1`: The switching event was initiated at the OLT, and the `primaryPonIF` is in the working state.~~
- ~~`MSG2`: The switching event was initiated at the OLT, and the `backupPonIF` is in the working state.~~
- ~~`failureCode` identifies the reason for the protection switching, per TBD.~~

~~`NMSR( protection, switch )`~~

~~This primitive is used by the NMS to request the working OLT to initiate a protection switch, during which the previously working and standby OLTs exchange their functions.~~

~~9.3.3.1.7~~9.3.3.2.5 State diagram~~s~~

~~The C-OLT shall instantiate the switching process state diagram as defined in Figure 9-6. In case Client protection is implemented at the OLT (i.e., when two C-OLTs are used as shown in Figure 9-9), the combined operation of both C-OLTs shall be as defined in Figure 9-6.~~ The ~~C-~~ONU shall instantiate the ~~switching~~ trunk protection process state diagram as defined in Figure 9-6.

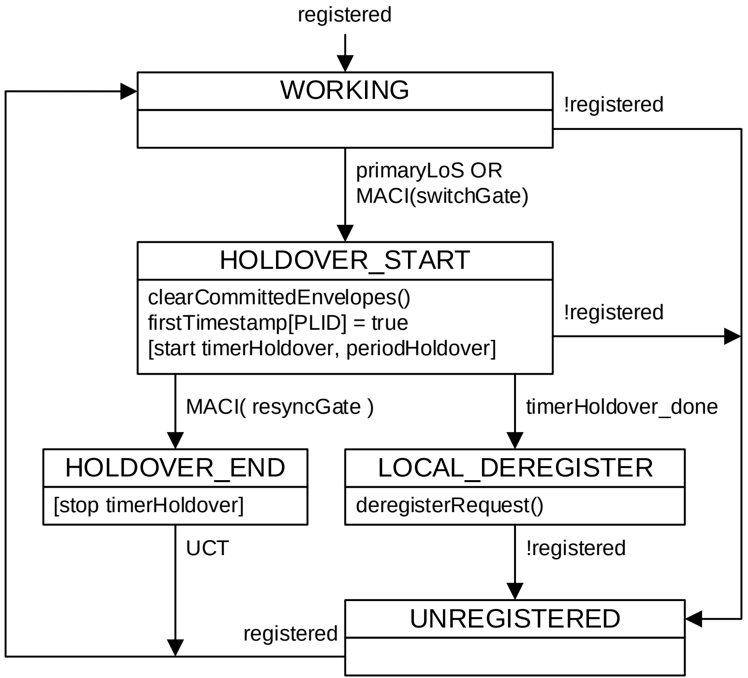

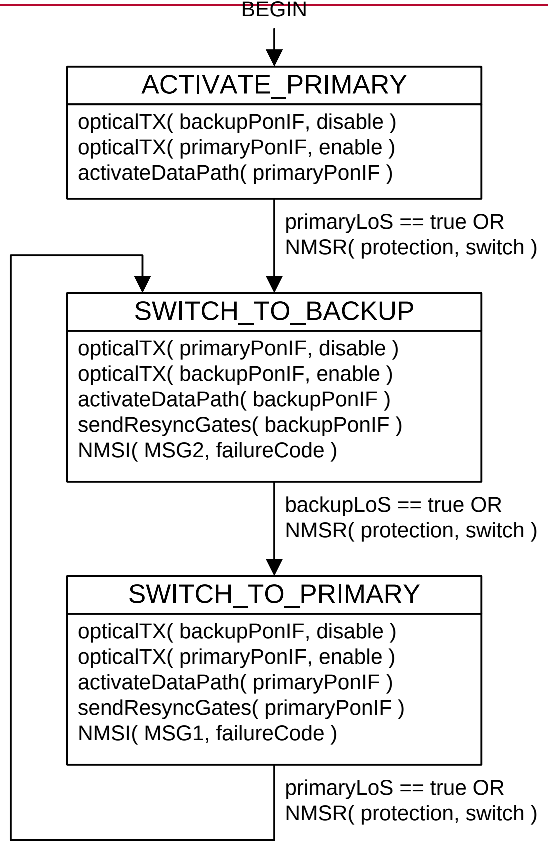

~~***Figure 9-6—Trunk protection process operating on the OLT***~~

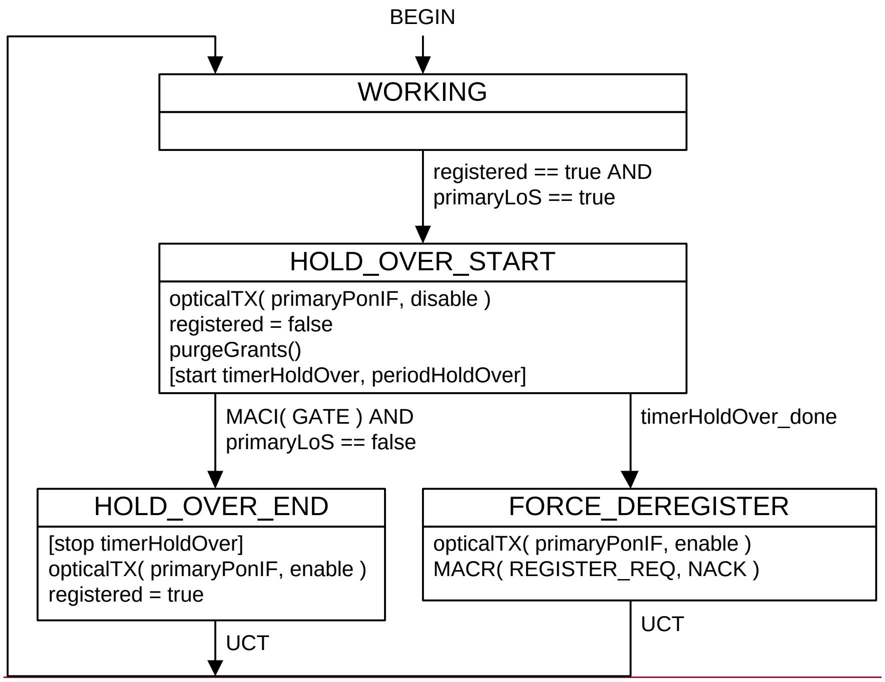

Figure ~~-7~~9-6—Trunk protection process operating on the ONU

Under normal operating conditions, the process remains in the WORKING state. When ONU detects the loss of signal (`primaryLoS == true`), it transitions to a HOLDOVER_START state.

Some kinds of failures are not readily detectable by ONUs, e.g., a failure of the OLT receiver sub-unit. For that reason, upon the protection switching event, the OLT generates a broadcast MPCPDU to all ONUs to indicate such event to the ONU and to solicit the necessary action form the ONUs. The exact message being broadcast depends on the protection switching procedure desired by the OLT.

For the default protection switching procedure, the OLT issues a broadcast request for ONUs to deregister (see 9.3.3.1.1). This MPCPDU is not intercepted by the ONU trunk switching process. Instead, it is handled by the ONU Registration process (see IEEE Std 802.3, 144.2.1.5), which results in ONU becoming unregistered. Whether the ONU has independently detected the loss of signal and transitioned into the HOLDOVER_START state or it has not setected it and remained in the WORKING state, the ONU trunk protection process reacts to the `registered` value becoming false and transitions into the UNREGISTERED state.

To perform the optimized protection switching procedure, the OLT issues a broadcast indication that the switch event took place (see `MACI(switchGate)` in 9.3.3.2.4), but does not request ONUs to deregister. If the ONU has not detected the loss of signal and remains in the WORKING state, the reception of the `switchGate` MPCPDU causes the ONU trunk switching process to enter the HOLDOVER_START state.

Upon entering the HOLDOVER_START state, the process performs the following actions:

a) Clears all committed envelope descriptors that have been scheduled by the working L-OLT before the protection switching event and that are not valid or anticipated by the new working (previously standby) L-OLT (see `cleadCommittedEnvelopes()` function definition in 9.3.3.2.3).

b) Sets the `firstTimestamp[PLID]` variable to true. This change ensures that the next MPCPDU received by the ONU Control Parser process (see IEEE Std 802.3, 144.2.1.5) resynchronizes the MPCP clock and does not cause ONU deregistration due to the timestamp drift exceeding the maximum allowed threshold (see `ProcessTimestamp(…)` function definition in IEEE Std 802.3, 144.2.1.4).

c) Starts the `timerHoldover` timer.

While in the HOLDOVER_START state, the following events may take place:

d) The `timerHoldover` timer expires. If the optimized trunk protection switching procedure fails to complete within the allocated holdover period, the ONU self-deregisters (in LOCAL_DEREGISTER state) . This causes the ONU to follow the default switching procedure, i.e., to re-register with the newly working L-OLT.

e) The ONU becomes unregistered due to an explicit OLT request. This also causes ONU to follow the default switching procedure, i.e., to re-register with the newly working L-OLT.

f) The ONU trunk protection process receives a resynchronization GATE MPCPDU (see `MACI(resyncGate)` in 9.3.3.2.4).

Before the ONU trunk protection process received the resynchronization GATE MPCPDU, that same MPCPDU was processed by the ONU Control Parser (see IEEE Std 802.3, 144.2.1.5), where its `Timestamp` value was used to resynchronize the local MPCP clock.

Upon reception of the resynchronization GATE MPCPDU, the ONU trunk switching process ends the holdover period (in HOLDOVER_END state) and unconditionally transitions back into WORKING state. The ONU is now synchronized to the new L-OLT clock and can be allocated PLID and MLID envelopes for upstream transmission. The user traffic is not yet enabled as the data encryption is yet to be reactivated.

To reactivate the encryption, the *CipherClocks* at the OLT and the ONU need to be resynchronized (see 11.3.5.4.1.3). This is achieved via the exchange of *Sync Cipher Clock* TLV (see 14.6.5.2). This exchange is not part of the ONU trunk protection process.

If the OLT uses the optimized protection switching procedure that bypasses the MPCP and OAM discovery (see 9.3.3.1.2), the next step is to initiate the ONU authentication (see 11.2.2), which also derives the new initial encryption key at the OLT and the ONU.

If the OLT uses the optimized protection switching procedure that bypasses the MPCP and OAM discovery and the ONU authentication (see 9.3.3.1.3), the OLT simply reactivates the same encryption key that was active before the protection switching event.