PON Protection Holdover Timers for Package A

Modifying the ONU Holdover Timer Range to Support PON Protection

Michael “Mike” Emmendorfer
Vice President, Systems Engineering and Architecture
ARRIS Group

Supporters

- ....

Proposal Summary

1. This proposal seeks to expand the configurable time range that ONUs will wait before proceeding to deregister during an OLT failover event, this range is known as the holdover period.

2. Specifically this submission changes a single configurable parameter timer range or duration that an ONU will remain in a holdover state to allow a standby OLT (e.g. line card) to takeover from the working OLT that failed.

3. This proposal specifically expands the configurable timer range limit from 1 second to 4.5 seconds.

4. Important Note: this is a configurable parameter range thus service providers that want 1 second or even less may continue to do so and this proposal simply enables service providers the ability to use a higher holdover period if desired.

5. Allowing a configurable range up to 4.5 seconds will not impact other timers on the ONU or OLT, additionally this will remain in alignment with all other aspects of SIEPON and IEEE specifications.

Background (1 of 4)
December 2014 Proposal for PON Protection for Package A

- **December 2014 Marek Hajduczenia submitted and received approval to add PON protection for Package A**

- **Drivers from Marek's Proposal to add PON Protection for Package A included:**
  - Package A had no PON protection support defined, either Trunk or Tree
  - MSOs need the option to have redundant hardware and hitless failover where appropriate
  - The main focus is on intra-chassis protection
  - Specifically a PON line card failure could be moved to the protecting redundant PON line card

> Motion#5: Move to accept rmtf_1412_hajduczenia_1a.pdf as the baseline proposal for the Optical Link Protection for package A.
> – (Required ≥ 2/3)
> – Moved: Marek Hajduczenia
> – Seconded: Mark Laubach
> – Y:7 N:0 A:0
> http://www.ieee1904.org/meeting_archive/2014/12/anwg_1412_closing.pdf

Source: http://www.ieee1904.org/revision/meeting_archive/2014/12/rmtf_1412_hajduczenia_2.pdf

Background (2 of 4)
December 2014 Proposal for PON Protection for Package A

Marek Hajduczenia submitted this information in his proposal to add PON Protection to Package A rmtf_1412_hajduczenia_2.pdf (Slide 3 below)

Redundancy and Failover

- For high-availability of services, MSOs need the option to have redundant hardware and hitless failover where appropriate.
  - The main focus is on intra-chassis protection at this time.
  - Inter-chassis protection will be addressed in the future.
- On PON line card failure, the services could be moved to the protecting redundant PON line card
  - 1+1 protection is required
  - N+1 protection would be welcome in the future
- Trunk protection is the primary focus at this time
- Support for tree protection will be added to specification as well
- Implementation demand is on trunk protection today, with tree protection likely becoming more important in residential FTTH in the future

We agree on all points and especially the 1:N protection would be welcome in the future.

The MSOs are interested in 1:N protection and making a single modification to this previously approved proposal could provide additional support for 1:N protection

Source: http://www.ieee1904.org/revision/meeting_archive/2014/12/rmtf_1412_hajduczenia_2.pdf

Background (3 of 4)
December 2014 Proposal for PON Protection for Package A

- **Sourced Directly from rmtf_1412_hajduczenia_1a.pdf**
  - **Page 5 Section 9.3.3 Trunk protection scheme**

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).

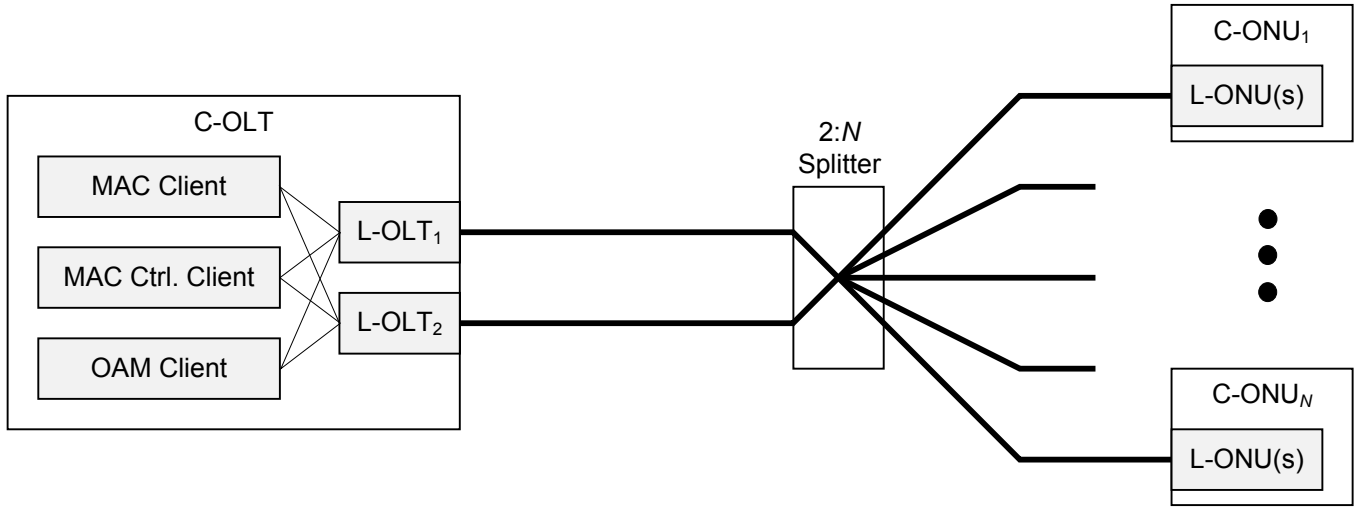

Figure 9 8—Trunk protection with redundant L-OLT

Source: http://www.ieee1904.org/revision/meeting_archive/2014/12/rmtf_1412_hajduczenia_1a.pdf

Background (4 of 4)
December 2014 Proposal for PON Protection for Package A

- **Sourced Directly from rmtf_1412_hajduczenia_1a.pdf**
  - **Page 6 Section 9.3.3 Trunk protection scheme**

In addition to intra-chassis protection, the trunk protection with redundant C-OLT scheme supports the inter-chassis protection, where the primary C-OLT and backup C-OLT are located in different chassis (either within the same central office or geographically different locations).

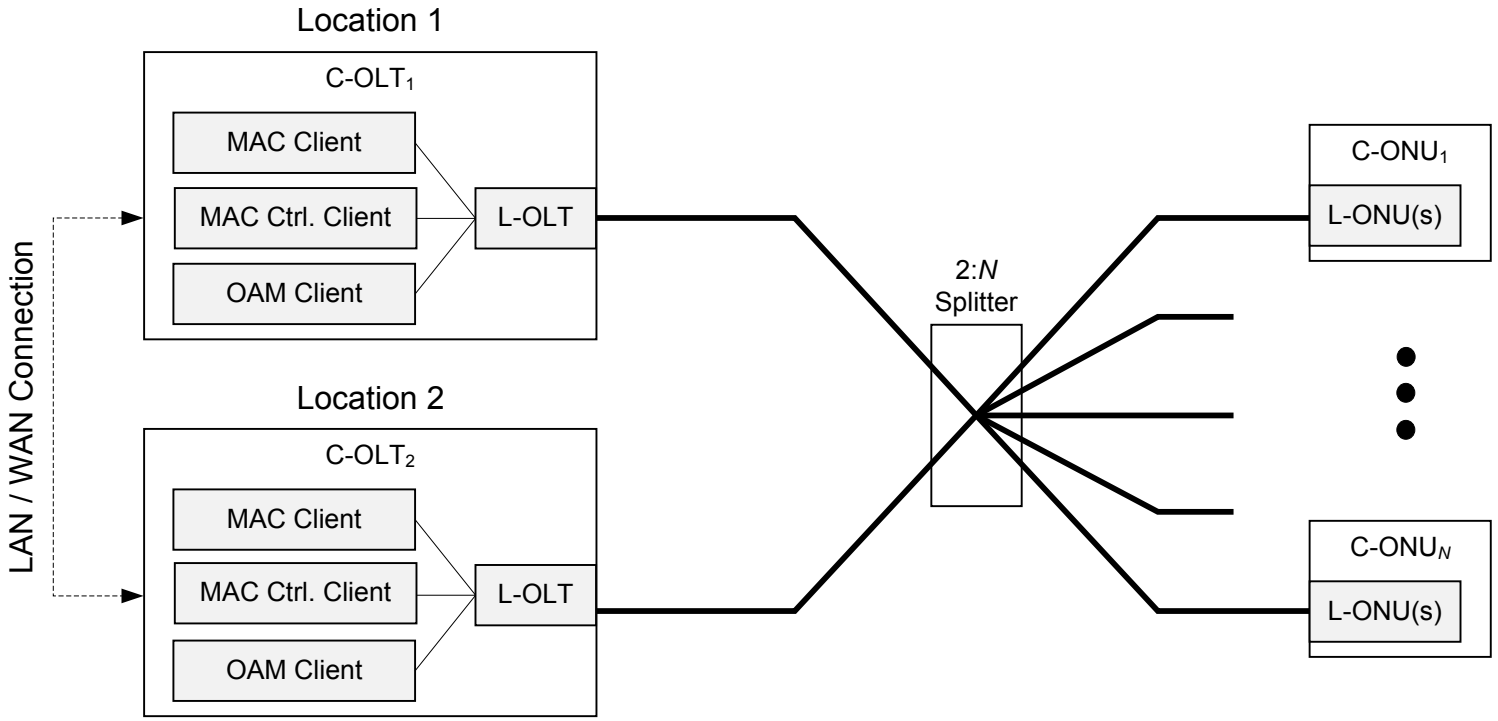

Figure 9-9—Trunk protection with redundant C-OLT

Source: http://www.ieee1904.org/revision/meeting_archive/2014/12/rmtf_1412_hajduczenia_1a.pdf

Leveraging Lessons Learned by the Cable Industry's Use of 1:N Line Card Protection Schemes

- **Cable Operators have extensive use of 1:N line card protection schemes when deploying DOCSIS technology and many welcome the future 1:N line card protection for PON**

- **Point of Reference and Lessons Learned with MSO's wide scale 1:N line card protection**
  - CableLabs DOCSIS 1.0 and 1.1 Specifications
    - Downstream Lost Sync Interval is 600 ms
    - Cable Operators observation is that cable modems would re-register after 600 ms ← Lesson Learned
  - CableLabs DOCSIS 2.0 Specifications ← CableLabs updated the specifications to increase the timer settings to support line card protection
    - MSOs / CableLabs corrected the short duration timer problem with timers with D1.0/D1.1
    - MSO wanted to predictably support 1:N Line Card Protection and expanded the timer range
      - Upstream
        - T4 (ranging interval) Timer Minimum Value is 30 seconds (this is also known as station maintenance timeout value)
      - Downstream
        - QAM/FEC Lock / Sync Message / Modem Timing will rerange based on T4 timing

Key Takeaway: MSO experience with 1:N protection is that 1 second Timers are not enough

Specific Change to rmtf_1412_hajduczenia_1a.pdf

14.4.1.9.4 Attribute aOnuConfigHoldoverPeriod (0xD7/0x09-03)

Sub-attribute aOnuConfigHoldoverPeriod.sHoldOverPeriod:

Syntax:	Unsigned integer
Range:	0x00-00 to 0x11-94 (4.5 seconds) ← Changed from: 0x03-E8 (1 second)
Remote access:	Read/Write
Unit:	1 ms
Default value:	0x00-C8
Description:	This sub-attribute represents the value loaded into the timerHoldOver timer, as defined in 9.3.3.

So what does this change mean?
The "configurable range" for ONU holdover period moves from 1 second to 4.5 seconds for the ONU to wait for an MPCP or Holdover end message before beginning the deregister process

So why a HoldOverPeriod configurable range up to 4.5 seconds?

- We want to stay within the specification of the IEEE 802.3
- While in HoldOverPeriod there are some functions on the ONU that run normally even in a holdover state
  - While MPCP is on hold the OAM keep-alive continues to run normally in an ONU (important point)
- **So while in holdover SIEPON needs to stay within the bounds of the IEEE 802.3 2008_section5 (EPON) 57.3.1.5 specification for local_lost_link_timer (the OAM timers)**
  - **this is defined as 5 Seconds ± 10% and this would mean the time range should not exceed 4.5 seconds**
- **A 4.5 second holdover period is "available" technically but it is just that the SIEPON specification are just not using it!!!!**
  - **So why not have the option to use up to 4.5 seconds if desired to avoid ONU deregistration if 1 second proves not to be enough time?**

Why the change to ONU Holdoverperiod?
Answer: We may need the additional time to support 1:N protection

- Keeping the ONUs registered is critical
  - the re-registration process is time consuming and avoiding re-registration of ONUs and vCMs is beneficial
- If using this additional and “available” holdover time of up to 4.5 seconds avoids the ONU from deregistering this is a very good thing
- Increasing Densities in terms of ONUs and LLIDs per port, per line card, and ultimately in a PON protection group may require more time
- Service providers may want to allow extra time to make sure all ONUs remain registered while the OLT standby line card becomes the working line card
- We are proposing a configurable range change from 1 second to 4.5 seconds and while some service providers may continue use of 1 second, others may want up to 4.5 seconds

Comments and/or Questions

- Is there a technical reason this should not be add to the SIEPON specs?

- Is there any reason that SIEPON should not support at least the option for a “configurable range” up to 4.5 seconds for the holdover period?

- Other comments / questions?

Motion

Move to adopt Sub-attribute aOnuConfigHoldoverPeriod.sHoldOverPeriod Range from 1 second to 4.5 seconds as described in rmtf_1508_emmendorfer_1.pdf (page 37) and in rmtf_1508_emmendorfer_2 (page 9).

- Motion (≥ 2/3)
- Moved by: Mike Emmendorfer, ARRIS Group
- Seconded by:
- Yes (#), No (#), Abstain (#)
- Motion passes / fails

Thank You

Backup Materials from 1904.1 SIEPON for Explanation Purposes

SIEPON 9.3.3.3 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 9-11) 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 9-11, 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 9-11, where all currently stored upstream transmission grants 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(DA , REGISTER_REQ , status = deregister) primitive to the underlying MPCP sublayer, per Figure 9-11. The OLT deregisters the ONU independently, based on the observed link status.

SIEPON 9.3.5.3.1 Automated trunk protection switching

9.3.5.3.1 Automated trunk protection switching

Figure 9-20 shows the scenario of an automated trunk protection. Upon detection of a link fault, ONUs move from the WORKING state to the HOLD_OVER_START state (see Figure 9-23). While in this state, the ONU blocks deregistration due to timestamp drift.

Simultaneously, the working OLT detects the link fault for all connected ONUs. The working OLT performs the switchover, during which the backup OLT becomes the working OLT and the primary OLT becomes the standby OLT. Once this process is completed, the working OLT sends unicast resync GATE MPCPDUs to synchronize ONUs' MPCP clocks to the new working OLT MPCP clock. Then the working OLT broadcasts the HOLDOVER with Flag set to end to all ONUs. On reception of the HOLDOVER message with Flag set to end, the ONUs move from the HOLD_OVER_START state to the HOLD_OVER_END state and then unconditionally to the WORKING state. If the ONU does not receive a GATE MPCPDU and the HOLDOVER message within periodHoldOver time interval, it deregisters.

SIEPON 9.3.5.3.2 NMS-driven trunk protection switching

9.3.5.3.2 NMS-driven trunk protection switching

Figure 9-21 shows the scenario of an NMS-driven trunk protection. After the working OLT receives the protection switching request from the NMS, the working OLT broadcasts the HOLDOVER message with Flag set to start to all ONUs. On reception of the HOLDOVER message with Flag set to start, the ONUs move from the WORKING state to the HOLD_OVER_START state (see Figure 9-23). While in this state, the ONU blocks deregistration due to timestamp drift. The working OLT performs the switchover, during which the backup OLT is activated and the working OLT is deactivated. Once this process is completed, the working OLT sends one or more resync GATE MPCPDUs to synchronize all ONUs' MPCP clocks to the new working OLT MPCP clock. Then the working OLT broadcasts the HOLDOVER message with Flag set to end to all ONUs. On reception of the HOLDOVER message with Flag set to end, the ONUs move from the HOLD_OVER_START state to the HOLD_OVER_END state and then unconditionally to the WORKING state. If the ONU does not receive a GATE MPCPDU and the HOLDOVER message within periodHoldOver time interval, it deregisters.