MTU UP !

Marek Hajduczenia, PhD

marek.hajduczenia@mybrighthouse.com

MTU in EPON (1)

o When 1G-EPON standard was first published, 802.3-2005 supported MTU up to 1500B. Life was simple.

o 802.3-2008 added differentiation for "basic frames", "Q-tagged frames", and "envelope frames". Operators moved to require support of envelope frames in EPON.

3.2.7 MAC Client Data field

The MAC Client Data field contains a sequence of octets. Full data transparency is provided in the sense that any arbitrary sequence of octet values may appear in the MAC Client Data field up to a maximum field length determined by the particular implementation.

Ethernet implementations shall support at least one of three maximum MAC Client Data field sizes defined as follows:

a) 1500 decimal—basic frames (see 1.4.73)
b) 1504 decimal—Q-tagged frames (see 1.4.291)
c) 1982 decimal—envelope frames (see 1.4.151)

Excerpt from 802.3-2008, 3.2.7

MTU in EPON (2)

o A large share of existing 1G-EPON implementations support up to 2kB MTU today, allowing to transport Ethernet frames with payload larger than 1500 bytes.

o Similarly, all existing 10G-EPON implementations support 2kB MTU.

o All of 10G (and above) P2P interfaces on routers, switches, etc. deployed today support at least 9kB MTU.

o A growing class of applications puts increased MTU requirements on operators, demanding 4kB and 9kB MTU support in the access space.

o Transition to larger MTU in P2P links is relatively simple, given that only two link peers are involved. In EPON, situation is not that simple, especially in burst-mode upstream direction.

MTU Drivers

o Today, there are three major applications driving MTU demand towards 4/9kB:

    o cellular backhaul (driven by protocols, such as CPRI, etc),

    o business customers running distributed storage / cloud applications between remote sites (driven by end application),

    o customer demand (no specific technical reason, apart from the perceived "need" for larger MTU and better link efficiency)

o In these cases, moving towards larger frames improves link utilization and minimizes "Ethernet tax" (framing overhead).

o Today, an operator typically sets off the "Ethernet tax" by increasing the provisioned service rate by certain fixed amount.

    o This works for most applications, especially when IP hosts are involved and source data can be fragmented without any problems.

    o In cellular backhaul or distributed storage, source hosts may require specific frame size above 2kB and fragmentation would require additional NIDs to be placed before access equipment, driving cost and complexity

MTU Solution

o The MTU quandary has two possible solutions:

    a) Support fragmentation between the ONU and the OLT;

    OR

    b) Increase the MTU size, supporting larger Ethernet frame sizes.

o Individual options are discussed in the following two slides

o In the scope of NG-EPON effort, it is recommended that the support for MTU >2kB becomes one of the official requirements to support existing and future applications, driving MTU size up.

o Without ability to carry larger MTU over NG-EPON, some applications will have to be migrated into CWDM P2P links, taking away a large share of potential NG-EPON market.

MTU Solution (a): Fragmentation

o P802.3br (IET) Task Force is currently working to define extended MAC architecture to support interspersing express traffic. It will also add framework for MAC frame fragmentation.

o The same framework could be reused in NG-EPON for packet fragmentation: large incoming frames would be fragmented into 2kB chunks, transmitted to OLT and reassembled.

o The drawback of this solution is the increased jitter, delay, and increased cost of resulting equipment due to fragmentation / reassembly at high data rates.

o Further analysis of this solution, its impact on jitter budget in NG-EPON, limitations, transmission overhead, etc. would be required when NG-EPON becomes a Task Force.

MTU Solution (b): increase MTU

o Alternatively, the size of Ethernet frame could be again increased beyond 2kB envelope frames.

o Since the last time this topic was discussed, line rates increased from 10 Gb/s to 100 Gb/s and beyond (400 Gb/s is under active development). With a fixed maximum frame size, frame rates increase linearly with the increase in the data rate.

o If MTU is increased to 9kB, the operator is responsible for making sure that all devices support this MTU and then configure them accordingly.

o Legacy devices would not be expected to support such frames (no need to solve the legacy-equipment-in-the-middle type of problems that IEEE 802.1 usually has to deal with)

THANKS !