Improving service assurance in packet transport networks

White paper

In carrier-grade packet transport networks, service operation, administration and maintenance (OAM) tools are necessary for Carrier Ethernet (CE) service assurance and for measuring compliance with a Service Level Specification (SLS). Frame Loss Ratio (FLR) is one of the most important performance metrics in the Ethernet Virtual Connection (EVC) SLS service attribute specified in MEF 10.4 [1]. FLR is measured using one of two Ethernet loss measurement tools specified in ITU-T Y.1731 [2] and following requirements specified in MEF 35.1 [3]. These two standards are used in the service OAM performance monitoring phase of the MEF 3.0 CE service lifecycle.

The Ethernet Synthetic Loss Measurement (ETH-SLM) tool has been used for more than a decade. It has the widest application scope and has been the tool available from most vendors. However, new transport services, such as those that support 5G networks or cloud interconnectivity services, require more accurate and faster FLR measurement capabilities. This need has led to a revival of interest in the original FLR measurement tool, the Ethernet Loss Measurement (ETH-LM tool, made possible by the availability of the necessary hardware functions. ETH-LM had many restrictions when it was first defined and was not commonly implemented by vendors.

This white paper explores the critical role that the ETH-LM and ETH-SLM tools will play in next-generation transport networks that offer CE services. It explains how ETH-LM works and provides a comparison between ETH-LM and ETH-SLM, specifically in terms of FLR accuracy and measurement time. Finally, it points out several additional considerations and conditions for using ETH-LM.

Contents

Introduction

In packet transport networks that carry IP-based services, the upper bound on the packet loss probability is 10-3, as defined in Clause 5.3 “Network QoS Classes” of ITU-T Y.1541 [4]. This corresponds to a typical FLR service-level agreement (SLA) for basic CE or IP VPN services that service providers offer to enterprises and cloud service providers. It also corresponds to the FLR ≤ 10-3 objective for point-to-point CE services with Class of Service (CoS) Label value “L” in MEF 23.2 [5].

For fixed services, businesses often require premium services with assurance of a much lower FLR, such as less than 10-5 for CoS Label value “H.” The development of 5G networks and emerging applications over 5G networks is also increasing assurance requirements for wireless services. For example, 3GPP Release 17 specifies stricter Quality of Service (QoS) requirements in Clause 6.1.7 “Standardized QoS Characteristics” of 3GPP TS 23.203 [6], where the upper bound for packet loss ratio is 10-8. In IEEE Std 802.1CM-2020 for fronthaul networks [7], the maximum FLR for user plane CPRI or eCPRI data flow is 10-7.

It is important to note that there is a difference between the FLR that a network can deliver and the accuracy of the service OAM FLR measurement capabilities. The former is ideally zero in a fault-free network that is properly traffic engineered. The FLR depends on many parameters, some of which service providers cannot control (e.g., FLR resulting from equipment or network failures). For service providers, the challenge is to be able to measure FLR with sufficient accuracy and speed so that they can detect performance degradation before end customers do and ensure SLS compliance.

Service providers need an efficient in-service OAM tool to monitor FLR in the network and achieve the primary goal of high accuracy over a reasonably short measurement interval. The ETH-SLM tool specified in ITU-T Y.1731 [2] is very popular and widely used as a flexible Ethernet service OAM tool for measuring frame loss for both point-to-point and multipoint services.

As service providers build more carrier-grade networks that offer greater capacity and resiliency and less oversubscription, the network FLR should decrease. This means the accuracy expected for FLR measurement is increasing. ETH-SLM counts only “synthetic” OAM frames specifically created for the purpose of measuring FLR. The lower the FLR to be measured, the longer it takes for ETH-SLM to measure it because of the statistical methodology it uses. This process could take weeks, months or even years, depending on the desired FLR accuracy. In comparison, the ETH-LM tool can measure frame loss very quickly and accurately because the loss of an individual frame can be detected without having to wait for the end of a long measurement interval.

Role of ETH-LM in service assurance

The lifecycle of MEF 3.0 CE services, shown in Figure 1, comprises three phases: Service Activation Testing, Service OAM Fault Management, and Service OAM Performance Monitoring. Each phase calls for the use of specific Ethernet service OAM protocols to validate and monitor the EVC. The ITU-T Y.1731 ETH-LM tool is one of the main tools for the Service OAM Performance Monitoring phase.

Figure 1. Ethernet Service OAM tools mapped to MEF 3.0 CE service lifecycle

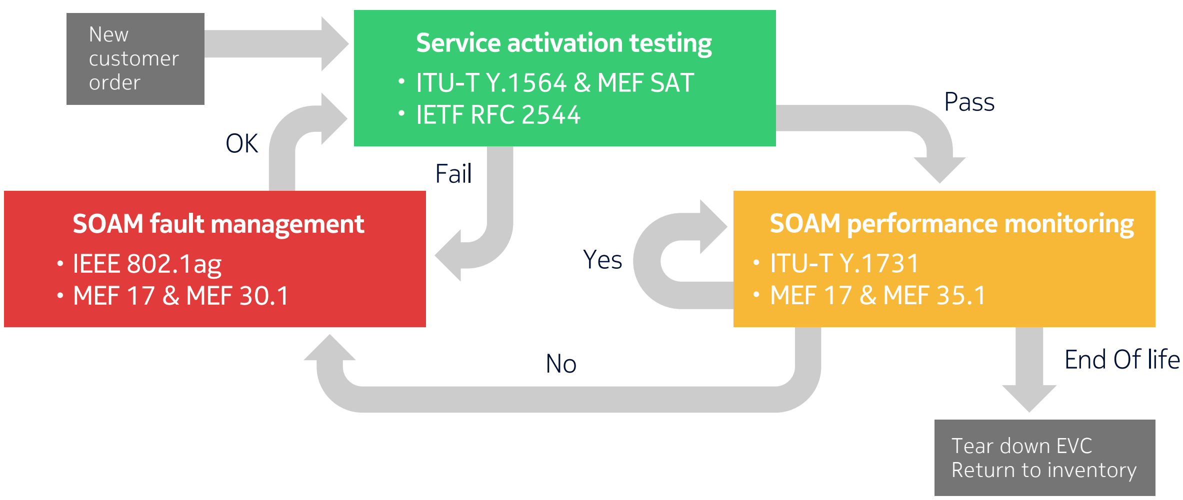

Service Activation Testing (SAT) is performed out of service before the provider delivers the service to its customer. It can be viewed as a process for validating that the services conform to the SLS. The SAT methodology verifies the configuration, validates the performance metrics of the SLS and provides a Key Performance Indicator (KPI) report (also known as a birth certificate) that the service provider can give to a customer when it declares that the service is up and running. The main test methodologies are as specified in ITU T Y.1564 [8] or in MEF 48.1 [9], a modern version of IETF RFC 2544 [10] that was not meant for production networks. These functions are implemented in test instruments or analyzers but can also be embedded as a test-head function in network elements (such as those from Nokia). The SAT phase involves active testing that is intrusive to the service because it injects test traffic at the User Network Interface (UNI).

Service OAM Fault Management is performed when an EVC is in service. The process aims to detect, diagnose, localize and correct network or service failures/fault conditions. The main tools are IEEE 802.1ag [11] or ITU-T Y.1731 OAM tools for fault management. Except for continuity and connectivity check functionality that runs proactively, the fault management tools (loopback and linktrace) run on demand because their use is limited to troubleshooting issues in the SAT or Service OAM Performance Monitoring phase.

Service OAM Performance Monitoring is performed once the EVC has been delivered to the customer, is in service, and may therefore be carrying customer traffic. This phase checks that the EVC is meeting the SLS throughout the period during which the SLS applies (not just at the moment of observation). Checking SLS conformance involves measuring the service performance metrics on an ongoing, proactive basis. The main tools are ITU-T Y.1731 OAM tools for performance monitoring (loss measurement and delay measurement). These tools run proactively (i.e., continuously) and perform passive testing (i.e., nonintrusive monitoring with no impact on the service). They do, however, inject OAM frames into the EVC, which consumes a small percentage of the EVC's bandwidth. Thus, service providers must take care to ensure that these frames do not compete with customer bandwidth.

For Service OAM Performance Monitoring, MEF 10.4 [1] specifies the following performance metrics in the EVC SLS service attribute:

- One-way Frame Delay Performance Metric
- One-way Mean Frame Delay Performance Metric
- One-way Frame Delay Range Performance Metric
- One-way Inter-Frame Delay Variation Performance Metric
- One-way Frame Loss Ratio Performance Metric
- One-way Availability Performance Metric
- One-way High Loss Intervals (HLI) Performance Metric
- One-way Consecutive High Loss Intervals (CHLI) Performance Metric
- One-way Composite Performance Metric
- One-way Group Availability Performance Metric

This white paper focuses on the use of the ITU-T Y.1731 ETH-LM in-service OAM tool to measure the last six of these ten metrics, all related to frame loss. The first four metrics are measured using the ITU-T Y.1731 ETH-DM delay measurement tool. Loss measurement could also be used to detect the Ethernet degraded signal defect (dDEG) specified in ITU-T G.8021 [12] which can be used as a signal degrade trigger for Ethernet linear or ring protection switching.

How ETH-LM works

Working principles

The one-way FLR, Availability, HLI and CHLI performance metrics specified by MEF require the measurement of the one-way frame loss between two monitoring points of an EVC. These monitoring points are called Up Maintenance Entity Group (MEG) End Points (MEPs), abbreviated as “Up MEPs.” Up MEPs are processing instances for initiating and terminating OAM frames over that EVC. As such, they are located at UNIs, each of which is a demarcation point between the service provider network and the end customer or subscriber network. Frame loss is the difference between the number of in-profile or “green” frames (i.e., those at or below the Committed Information Rate (CIR) so-called “green” frames of a given CoS) arriving at the ingress monitoring point and the number of green frames actually delivered by the EVC to the egress monitoring point, as exemplified in Figure 2. FLR is the ratio of green frames lost to the number of green frames arriving at the ingress monitoring point during the measurement interval.

Frame loss is measured unidirectionally per MEF 10.4 [1], providing a one-way (near-end to far-end) FLR performance metric. All the MEF performance metrics are one-way metrics. Their directionality provides a useful indication to the service provider about the nature of the degradation.

Figure 2. Example of Ethernet FLR measurement principles in CE networks

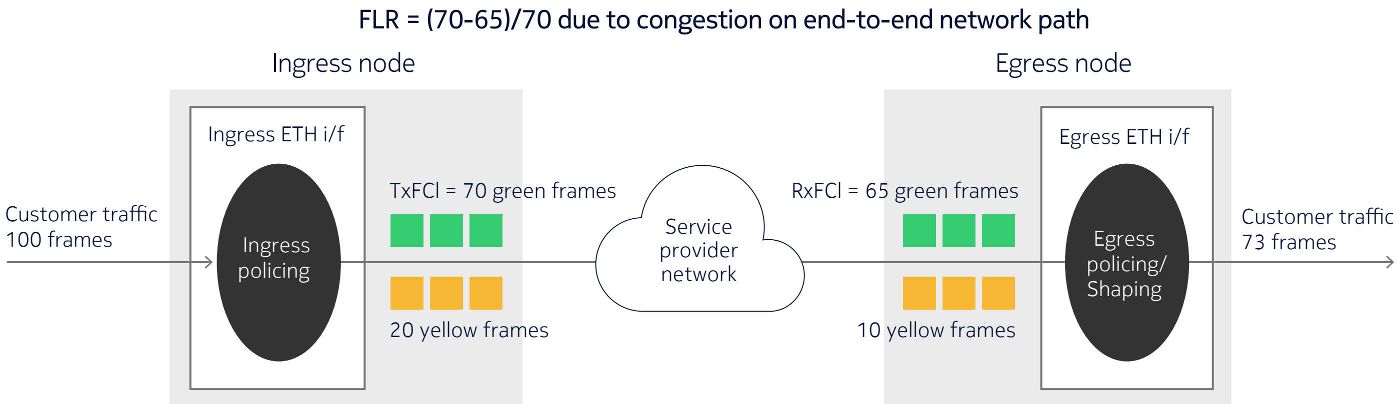

Recommendation ITU-T Y.1731 [2] specifies two in-service loss measurement methods that can be used on demand or for proactive measurements:

- **ETH-SLM**, specified in clause 8.4 of ITU-T Y.1731 [2], counts Synthetic Loss Message (SLM) frames injected in an EVC with a set of counts by a source (initiator) MEP and Synthetic Loss Reply (SLR) frames extracted from the EVC with incremented counts sent in response by the corresponding destination (responder) MEP (or MEPs). The source MEP retrieves each SLR frame and performs a frame loss computation. Figure 3 illustrates an example of this exchange. SLM/SLR frames are constructed with the destination Media Access Control (MAC) address set to the destination MEP's MAC address in both cases of a point-to-point or a multipoint EVC.

Figure 3. Working principles of single-ended ETH-SLM

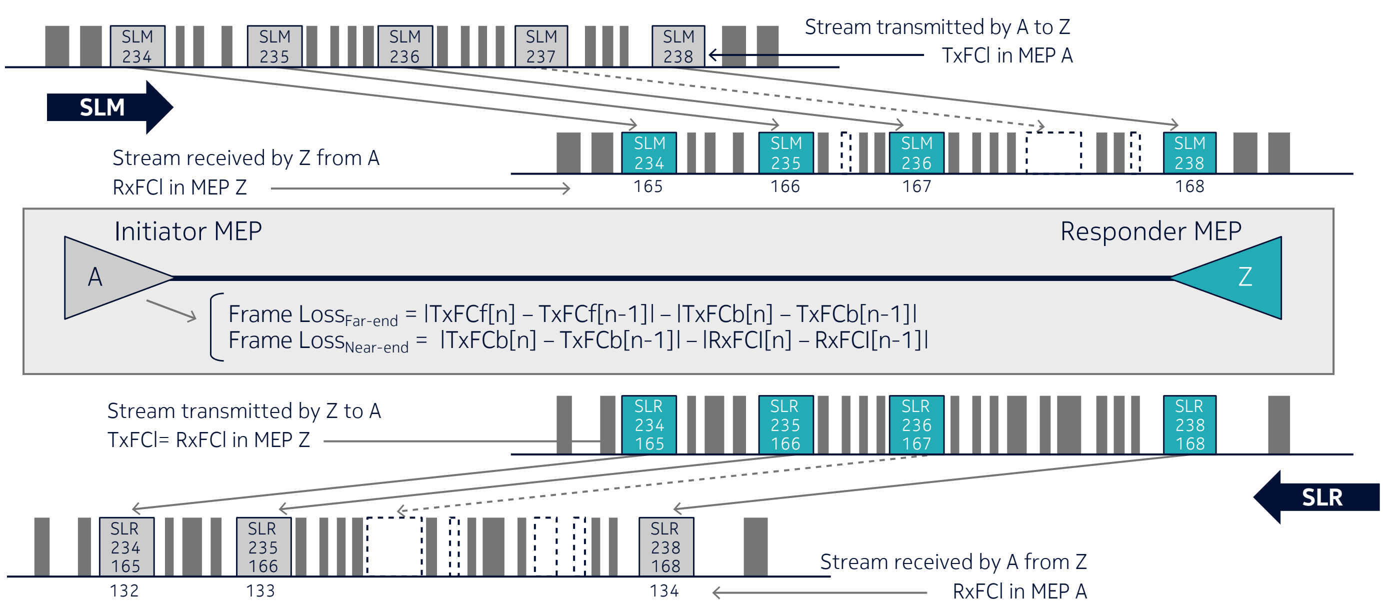

- **ETH-LM**, specified in clause 8.1 of ITU-T Y.1731 [2], counts actual service frames (unicast, unknown, multicast or broadcast) entering and exiting at the UNIs of the EVC. ETH-LM is limited to point-to-point Ethernet services because an ingress UNI may send service frames to multiple egress UNIs in a multipoint service, and an egress UNI may receive service frames from multiple ingress UNIs. As a result, the service frame counters at an ingress UNI and egress UNI do not necessarily match.

ITU-T Y.1731 defines two ETH-LM methods:

- Dual-ended ETH-LM uses Continuity Check Message (CCM) frames and is not specified in MEF 35.1 [3]. Consequently, it has had limited deployment (e.g., Nokia does not implement it).
- Single-ended ETH-LM uses Loss Measurement (LM) Message (LMM) and Reply (LMR) frames and is specified in MEF 35.1 [3] as part of its "PM-3 Solution." Single-ended ETH-LM is implemented by Nokia and described below.

Single-ended ETH-LM collects counts of frames per EVC or per <EVC, CoS> (depending on the hardware capabilities) at the ingress and egress monitoring points of an EVC and in both directions. The ETH-LM implementation creates the LMM and LMR Protocol Data Units (PDUs) including actual frame counts, injecting and extracting them at monitoring points to determine the unidirectional frame loss between point-to-point MEP peers. Figure 4 illustrates this exchange.

Figure 4. Working principles of single-ended ETH-LM

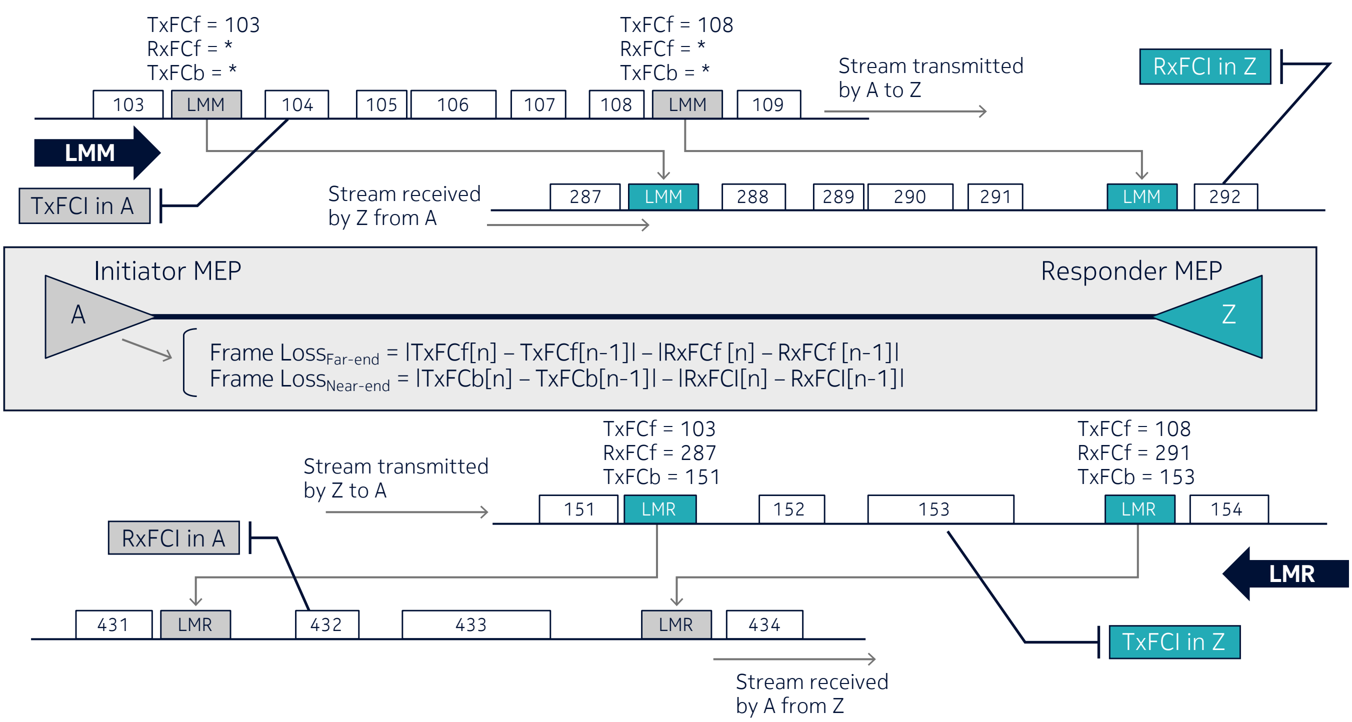

The single-ended ETH-LM protocol involves two measurement entities:

- The **LM initiator**, instantiated on a **local MEP**, initiates the loss measurement by constructing (as described in clause 9.12 of ITU-T Y.1731 [2]) and transmitting LMM PDUs to the peer MEP, and receiving LMR PDUs from the peer MEP. Based on a comparison of its counts and the counts collected from the LMR, the LM initiator determines the near-end and far-end frame loss.
- The **LM responder**, instantiated on a **remote MEP**, simply replies to LMM PDUs with LMR PDUs (as described in clause 9.13 of ITU-T Y.1731 [2]), incorporating counters as discussed below.

Each MEP (with LM initiator or LM responder role) maintains two local counters for each peer MEP and for each CoS:

- TxFCl: Counter for in-profile service frames transmitted to its peer MEP.
- RxFCl: Counter for in-profile service frames received from its peer MEP.

At regular intervals and for a given CoS, the LMM and LMR PDUs exchange these local counters between the LM initiator and LM responder to provide the LM initiator with the following four counts:

- TxFCf contains TxFCl at the moment of LMM PDU transmission from the LM initiator.
- RxFCf contains RxFCl at the moment of LMM PDU reception by the LM responder.
- TxFCb contains TxFCl at the moment of LMR PDU transmission from the LM responder.
- RxFCl is recorded at the moment of LMR PDU reception of the by the initiator MEP.

TxFCf, RxFCf and TxFCb fields are included in the LMM and LMR PDUs.

Calculating frame loss for on-demand and proactive ETH-LM

The real-time nature of the ETH-LM protocol requires hardware support to guarantee accuracy. LMM/LMR PDUs must be constructed and injected into the data stream immediately after reading the local counters, which count the last frames before LMM/LMR is inserted. Such immediacy can only be achieved in hardware. For example, with a software implementation, the locations in the data path where counters are incremented and where the LMM/LMR PDUs are constructed/processed may not be the same. Between the time local counters are read and the time LMM/LMR PDUs are constructed and injected, and depending on the bit rate, thousands of frames may have been received at the UNI, rendering the measurement completely inaccurate.

Upon receiving the first LMR, the LM initiator collects the four counts listed above and uses them as a starting point. Each LMR subsequently received for the same measurement during the desired measurement interval triggers the following frame loss calculations:

- Far-end Frame Loss = |TxFCf[n] - TxFCf[n-1]| - |RxFCf[n] - RxFCf[n-1]|
- Near-end Frame Loss = |TxFCb[n] - TxFCb[n-1]| - |RxFCl[n] - RxFCl[n-1]|

Note that the measurement is done at one end. The LM initiator MEP collects counts for both incoming (near-end) and outgoing (far-end) directions; the other end simply responds.

- **On-demand ETH-LM** consists of one or multiple simultaneous on-demand measurement sessions (also called “test” in certain graphical management interfaces). It is intended to measure and report the frame loss during a finite period of time, as a one-shot measurement. An on-demand ETH-LM test starts when enabled by the user, runs during the indicated time, and is then terminated. The user is provided with the results at the end of the test, via the different management interfaces (command-line interface (CLI), WebUI, Nokia WaveSuite NOC). Table 1 shows the counters in the measurement results.

Table 1. On-demand ETH-LM counters

On-demand counters	Counter acronym
Near-end Transmitted Data Frame	TN_TFCnt
Near-end Received Data Frame	TN_RFCnt
Far-end Transmitted Data Frame	TF_TFCnt
Far-end Received Data Frame	TF_RFCnt

- **Proactive ETH-LM** consists of one or multiple simultaneous proactive or continuous measurement sessions (also called “test” in certain graphical management interfaces). It is intended to continuously measure the frame loss performance once it has been started. Measurement results are collected into Performance MONitoring (PMON) bins over the regular ITU-T G.7710 [13] measurement intervals (15 minutes or one day). Table 2 shows the counters in the measurement results. Threshold Crossing Alerts (TCAs) are provided for selected counters, for example, average and maximum FLR, to alert the operator when service degradation begins and when the SLS is violated.

Table 2. Proactive ETH-LM PMON counters (15 minutes or 24 hours)

PMON counters	Counter acronym	TCA
Near-end Transmitted Data Frame	N_TF	No
Near-end Received Data Frame	N_RF	No
Far-end Transmitted Data Frame	F_TF	No
Far-end Received Data Frame	F_RF	No
Minimum Near-End FLR	mN_FLR	No
Average Near-End FLR	aN_FLR	Yes
Maximum Near-End FLR	xN_FLR	Yes
Minimum Far-End FLR	mF_FLR	No
Average Far-End FLR	aF_FLR	Yes
Maximum Far-End FLR	xF_FLR	Yes
Near-End Unavailability Interval (UAI)	N_UAI	No
Near-End High Loss Interval (HLI)	N_HLI	Yes
Far-End Unavailability Interval (UAI)	F_UAI	No
Far-End High Loss Interval (HLI)	F_HLI	Yes

Key benefits

The key benefits of ETH-LM are as follows:

- It provides the ultimate accuracy: The loss of an individual frame can be detected without having to rely on synthetic frames and a statistical measurement approach.
- FLR measurement is quick: Unlike ETH-SLM, there is no need to wait for a sufficiently long measurement period.
- The accuracy is completely independent of the LMM/LMR PDU interval, which does not need to be small as for ETH-SLM.
- It consumes less bandwidth for OAM.

Frames to be counted

To achieve an accurate frame count between the two monitoring points (MEPs), which might be on products from different vendors, it is important that the same frame types be counted in the TxFCl and RxFCl. If the MEPs at both ends of an ETH-LM session do not count the same frame types, there could be a systematic processing error. This error could easily be larger than the “normal” frame loss rate and in some cases lead to completely meaningless results.

As summarized in ITU-T G Suppl. 53 [15] and its Table 1, and formalized in ITU-T G.8021 [12], the rule is that the following frames shall be counted in the TxFCl and RxFCl:

- All unicast/multicast/broadcast service frames with a "green" level of bandwidth profile conformance (i.e., in-profile service frames) that arrive at the UNI and are mapped into the EVC after passing local IEEE 802.1Q [14] port filtering entities (i.e., locally filtered/dropped frames are not counted).
- Specific OAM PDUs:
  - OAM PDUs of higher MEG Level than the counting MEP are counted.
  - OAM PDUs of lower MEG Level are not counted.
  - For OAM PDUs of equal MEG Level, and other control PDUs, there are subtleties as summarized in Table 1 of ITU-T G Suppl. 53 [15].

Early on, there were some contradictions between the ITU-T Recommendations (ITU-T Y.1731 [2], ITU-T G.8021 [12]) and also differences with MEF 35.1 [3]. This led to liaison exchanges between ITU-T SG15 and MEF, which resulted in additional standardization work. MEF 35.1 mentions the discrepancies. The aligned ITU-T SG15 requirements are included in the latest versions of these ITU-T Recommendations and summarized in Table 1 of ITU-T G.Suppl. 53 [15].

The Nokia 1830 Photonic Service Switch (PSS) implementation complies with ITU-T G.8021 [12], which is different from what MEF 35.1 [3] specifies. The main difference is that ITU-T G.8021 counts proactive OAM frames in addition to service frames while MEF 35.1 counts service frames only. The advantages of following ITU-T G.8021 are:

- When no service frames are present in the EVC, the count of proactive OAM frames can still be used to monitor the EVC service FLR.
- Counting proactive OAM frames together with service frames increases the total sample size and helps to shorten the measurement time to achieve the desired FLR measurement accuracy.
- It overcomes the drawback that ETH-LM in PM-3 and Appendix J of MEF 35.1 [3] cannot calculate the Availability performance metric in an idle situation because it cannot measure loss when there are no service frames transmitted.

Additional considerations

Service providers should be aware of several constraints when using the ETH-LM tool:

- MEG and EVC applicability – ETH-LM is designed for a point-to-point MEG (i.e., a MEG with exactly two MEPs) and thus only applicable to point-to-point EVCs (E-Line). It is impossible to report frame loss results correctly if the point-to-point relationship is broken. For example, if a MEP is receiving service frames from multiple peer MEPs simultaneously, as is the case for multipoint EVCs, ETH-LM will not work in a reliable way. Some considerations for ETH-LM in point-to-point and multipoint EVCs are discussed in Appendix E of MEF 35.1 [3]. ETH-LM is well suited for the more than 70 percent of 5G transport connectivity in fronthaul, midhaul and backhaul networks that uses point-to-point connections. The exact percentage is operator dependent.
- ETH-LM over link aggregation – ETH-LM requires frame order preservation so that the frames between two adjacent LMM or LMR PDUs are the same as those between the pair of MEPs. The order of LMM and LMR PDUs relative to the frames that are counted is important. When the MEG traverses a link aggregation group, depending on its operational mode, link aggregation may not guarantee frame order preservation. This is discussed in Appendix VII "ETH-LM and Link Aggregation" of ITU-T Y.1731 [5]. As shown in Figure 5 below, reproduced from Figure VII.2 of ITU-T Y.1731 Appendix VII, if the LMM PDU shifts position relative to the service frames around it, the comparison will report an artificial frame loss or gain. This would reduce the FLR accuracy measured by ETH-LM.

Figure 5. Example of LMM PDU overtaking data frame causing artificial frame loss (or gain)

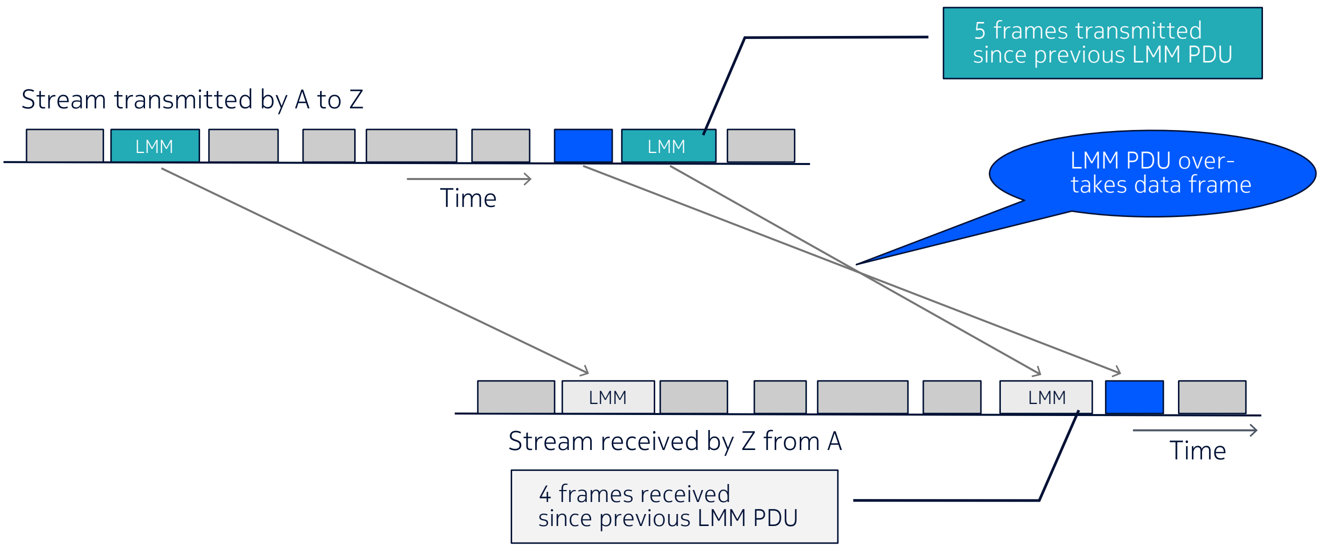

This issue is present when link aggregation operates in load-sharing mode. It can be avoided if link aggregation operates in active/standby mode (i.e., all traffic is forwarded over the active link of a link aggregation group with two aggregation links) or conversation-sensitive mode (i.e., assigning all the traffic in a conversation to a single aggregation link).

- ETH-LM QoS – To guarantee that the same service frames are counted at both MEPs, they must be configured for the same CoS and color (i.e., "green") of service frames as for the end-to-end EVC, otherwise the FLR accuracy measured by ETH-LM will be reduced.
- Frames to be counted – As explained in earlier, the MEPs at both ends of an ETH-LM session must count the same frame types to avoid a systematic FLR measurement error. When the MEPs at each EVC endpoint are from different vendors, the service provider should verify that they adhere to the same standards.

Conclusion

An efficient and accurate in-service loss measurement tool is essential for the service assurance of CE services. The ITU-T Y.1731 ETH-SLM tool has been the most popular and widely deployed service performance management tool over the past decade for both point-to-point and multipoint EVCs, but it has some significant drawbacks. With the emergence of more critical applications that demand lower latency and much lower FLR, an ETH-SLM measurement tool could take too long (from weeks to months or years) to achieve sufficient FLR accuracy with confidence. During this time, the service provider would be unaware of the real performance of the EVC and would be exposed in terms of SLS compliance.

Cloud infrastructure services and 5G applications in fronthaul, midhaul and backhaul transport networks mainly use point-to-point services. For these applications, the original, hardware-assisted, ITU-T Y.1731 ETH-LM has become a more relevant tool. Its main advantage is the much shorter measurement time required to achieve suitably high accuracy for a very low FLR. With the precision provided by ETH-LM, service providers have no blind spots in their network and can comfortably offer premium, higher revenue-generating CE services, with confidence in their ability to precisely measure SLS conformance for these services.

Abbreviations

CCM	Continuity Check Message
CE	Carrier Ethernet
CHLI	Consecutive High Loss Intervals
CIR	Committed Information Rate
CoS	Class of Service
CPRI	Common Public Radio Interface
ETH-LM	Ethernet Loss Measurement function
ETH-SLM	Ethernet Synthetic Loss Measurement function
EVC	Ethernet Virtual Connection
FLR	Frame Loss Ratio
HLI	High Loss Intervals
IEEE	Institute of Electrical and Electronics Engineers
ITU-T	International Telecommunication Union - Telecommunication Standardization Sector
KPI	Key Performance Indicator
LM	Loss Measurement
LMM	Loss Measurement Message
LMR	Loss Measurement Reply
MAC	Media Access Control
MEF	MEF Forum
MEG	Maintenance Entity Group
MEP	Maintenance Entity Group End Point
NOC	Network operations center
OAM	Operation, administration and maintenance
PDU	Protocol Data Unit
PMON	Performance MONitoring
QoS	Quality of Service
PSS	Photonic Service Switch
SAT	Service Activation Testing
SLA	Service level agreement
SLM	Synthetic Loss Message
SLR	Synthetic Loss Reply
SLS	Service Level Specification
TCA	Threshold Crossing Alert
UNI	User Network Interface

References

[1] MEF 10.4 Subscriber Ethernet Service Attributes

[2] Recommendation ITU-T G.8013/Y.1731 Operation, administration, and maintenance (OAM) functions and mechanisms for Ethernet-based networks

[3] MEF 35.1 Service OAM Performance Monitoring Implementation Agreement

[4] Recommendation ITU-T Y.1541 Network performance objectives for IP-based services

[5] MEF 23.2 Implementation Agreement MEF 23.2 Carrier Ethernet Class of Service – Phase 3

[6] 3GPP TS 23.203 Policy and charging control architecture (Release 17)

[7] IEEE Std 802.1CM-2018, IEEE Standard for Local and Metropolitan Area Networks — Time-Sensitive Networking for Fronthaul, as amended by IEEE Std 802.1CMde-2020

[8] Recommendation ITU-T Y.1564 Ethernet service activation test methodology

[9] MEF 48.1 Ethernet Service Activation Testing

[10] IETF RFC 2544 Benchmarking Methodology for Network Interconnect Devices

[11] IEEE Std 802.1ag, IEEE Standard for Local and Metropolitan Area Networks — Virtual Bridged Local Area Networks Amendment 5: Connectivity Fault Management

[12] Recommendation ITU-T G.8021/Y.1341 Characteristics of Ethernet transport network equipment functional blocks

[13] Recommendation ITU-T G.7710/Y.1701 Common equipment management function requirements

[14] IEEE Std 802.1Q, IEEE Standard for Local and Metropolitan Area Networks — Bridges and Bridged Networks

[15] ITU-T Series G Supplement 53, Guidance for Ethernet OAM performance monitoring

About Nokia

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering networks that sense, think and act by leveraging our work across mobile, fixed and cloud networks. In addition, we create value with intellectual property and long-term research, led by the award-winning Nokia Bell Labs.

Service providers, enterprises and partners worldwide trust Nokia to deliver secure, reliable and sustainable networks today – and work with us to create the digital services and applications of the future.

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: (April) CID212779