Measuring Frame Loss Ratio for Carrier Ethernet Services

Sample sizes and measurement intervals for ETH-SLM and ETH-LM

White paper

Frame Loss Ratio (FLR) is a key performance metric that can help service providers determine whether their Carrier Ethernet (CE) services comply with a Service Level Specification (SLS) as specified in MEF 10.4 [1]. Service providers can measure adherence to an FLR performance objective in an SLS using one of two Ethernet frame loss measurement tools specified in ITU-T Y.1731 [2]: Ethernet Synthetic Loss Measurement (ETH-SLM) or Ethernet Loss Measurement (ETH-LM). As a companion document to the Nokia white paper “Improving service assurance in packet transport networks” [7], this technical white paper aims to compare the measurement time between ETH-SLM and ETH-LM.

To measure FLR with confidence and sufficient accuracy, the service provider must decide how many samples to collect and for how long. This white paper provides insights that can help service providers determine the minimum sample size they need to measure FLR as accurately as possible. It also compares the measurement intervals required to achieve a desired FLR accuracy using ETH-SLM and ETH-LM.

Contents

Introduction

With the emergence of new enterprise and cloud service provider applications, 5G services, Industry 4.0, and growing interest in augmented reality/virtual reality (AR/VR) applications for the metaverse, packet transport networks are required to provide ever-lower Frame Loss Ratio (FLR), down to 10^{-7} or 10^{-8}, which equates to one lost frame per 100 million frames transmitted [3][4][5][6]. To meet this FLR objective and comply with the Service Level Specification (SLS) agreed with subscribers, service providers need to engineer Carrier Ethernet (CE) services carefully and have an efficient tool for monitoring FLR performance. But as the FLR decreases, it takes more time to collect enough samples and ensure that the measured FLR is stable and credible.

As described in the white paper “Improving service assurance in packet transport networks” [7], Ethernet Loss Measurement (ETH-LM) counts actual service frames to measure the FLR, while Ethernet Synthetic Loss Measurement (ETH-SLM) counts synthetic frames to estimate the FLR. Before starting to compare the two tools, it is helpful to clarify the types of FLR and the relationships between them. This provides the basis for understanding the rationale behind engineering CE services and choosing a measurement interval that will generate a sufficient sample size.

Definitions of Network FLR, Service FLR and Synthetic Frame FLR

FLR can be categorized into three types:

Network FLR: The FLR that the network is built to provide, independent of whether it is carrying active services. This is a fundamental criterion of network design. It depends on the physical characteristics of the equipment and the design rules for network planning, Quality of Service (QoS) and traffic engineering (e.g., the amount of oversubscription). As a result, it is not easy to improve Network FLR once the network has been built.

EVC Service FLR: The FLR performance metric experienced by service frames in an Ethernet Virtual Connection (EVC), as defined in MEF 10.4 [1]. Since the user traffic in packet networks is usually bursty and will seldom run at line rate, EVC Service FLR is a sort of sampling of the Network FLR.

Synthetic Frame FLR: The FLR experienced by synthetic Operations, Administration and Maintenance (OAM) frames in the network, such as ETH-SLM frames. Synthetic OAM frames are usually generated at a constant rate, so Synthetic Frame FLR is also a sort of sampling of the Network FLR.

Figure 1 depicts the relationships between the Network FLR, EVC Service FLR and Synthetic Frame FLR, as described in Appendix D of MEF 35.1 [8].

Figure 1. Relationships between Network FLR, EVC Service FLR and Synthetic Frame FLR

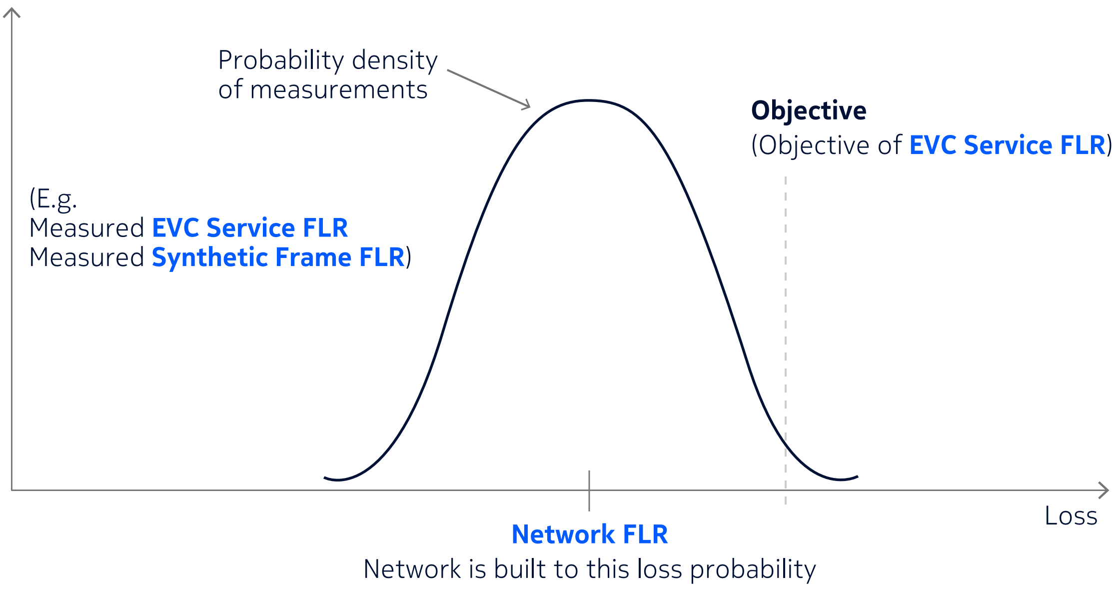

Per MEF 35.1 [8], FLR is a performance metric that requires a long observation period. Contrary to common belief, it cannot be measured instantaneously. As a result, FLR is not suitable for immediate action, in contrast to other one-way performance metrics such as for High Loss Intervals (HLI) and Consecutive High Loss Intervals (CHLI).

A service provider and its subscriber agree on the objective of the EVC Service FLR as part of the SLS. To meet this objective, the service provider and its network operator partners engineer the network to support the smallest possible FLR (i.e., the Network FLR), which is lower than the FLR objective of the SLS.

It is difficult to make the Network FLR extremely low because of restrictions in network components, network planning and traffic engineering, unless one is ready to significantly over-engineer a network, leading to higher CAPEX. When the service and/or synthetic OAM frames are transmitted over the network, the EVC Service FLR or Synthetic Frame FLR can start being measured. The measured FLR may be different in each case, with the percentages probabilistically distributed around the Network FLR. As the number of samples (service frames or synthetic OAM frames) increases, the variation of the measured FLR decreases. The measured FLR will be more concentrated around the Network FLR and will have a lower probability of exceeding the FLR objective.

To ensure that the measured FLR is lower than the EVC Service FLR objective without requiring the Network FLR to be unrealistically small, the service provider needs to use measurement intervals that are long enough to capture sufficient samples. Minimizing statistical variation depends heavily on selecting the correct sample size, which is tied to the measurement interval for a specific rate.

Consult Appendices D and J of MEF 35.1 [4] for more information on statistical considerations for ETH-SLM and frame loss count accuracy.

Number of samples required to measure FLR with confidence

As described in MEF 35.1 Appendix D [8], the Coefficient of Variation (CoV), i.e., the ratio of a probability density's standard deviation to its mean, is used to measure the variation of the measured FLRs. The smaller the CoV is, the more accurate the measurements will be.

$$\mathrm{CoV} = \frac{\sigma}{\mu} = \frac{\sqrt{n*p*q}}{n*p} = \sqrt{\frac{q}{np}} = \sqrt{\frac{q}{p}} * \frac{1}{\sqrt{n}}$$

Assumption: the event each frame is lost or not follows binominal distribution.
n – sample size
p – probability that a frame is lost
q – probability that a frame is not lost
σ – standard deviation to the expected measure FLR
μ – mean of the expected measured FLR

For example, consider a case where the Network FLR is built 30 percent smaller than the FLR objective (e.g., the FLR objective is 1×10^{-7}, Network FLR is built at 0.7×10^{-7}). The CoV should be less than or equal to 0.2 so that the service provider is confident that the measured FLR has a 95 percent probability of being less than the FLR objective. This is also known as FLR accuracy.

To achieve CoV <= 0.2, the number of required samples $n >\approx 25 / p$ when p is far less than 1. For example, to measure $p = 0.7 \times 10^{-7}$, the minimum sample size should be 35.7×10^{7}.

Figure 2 illustrates the probability and cumulative probability of the measured FLR results with different numbers of samples for a Network FLR fixed at 0.7×10^{-7}. In each graph, the Network FLR is illustrated by the vertical green dashed line, the FLR objective is illustrated by the vertical orange dotted line. The sample size is increased by a factor of ten between each row of two graphs. The graph on the left indicates the probability of each measured FLR result; the graph on the right indicates the cumulative probability of each measured FLR result.

These graphs show that the variation of the measured FLR result (also reflected in the CoV) becomes smaller with the increase in the number of samples. When CoV reaches 0.2, the cumulative probability of the measured FLR results becomes greater than 95 percent, which means that sufficient FLR accuracy is achieved.

Nokia has designed an ETH-SLM accuracy calculator that can be used to determine the minimum sample size and required ETH-SLM frame transmission period for a given target CoV. Contact your Nokia sales representative to obtain this tool.

Figure 2. Probability and cumulative probability of measured FLR results versus sample size

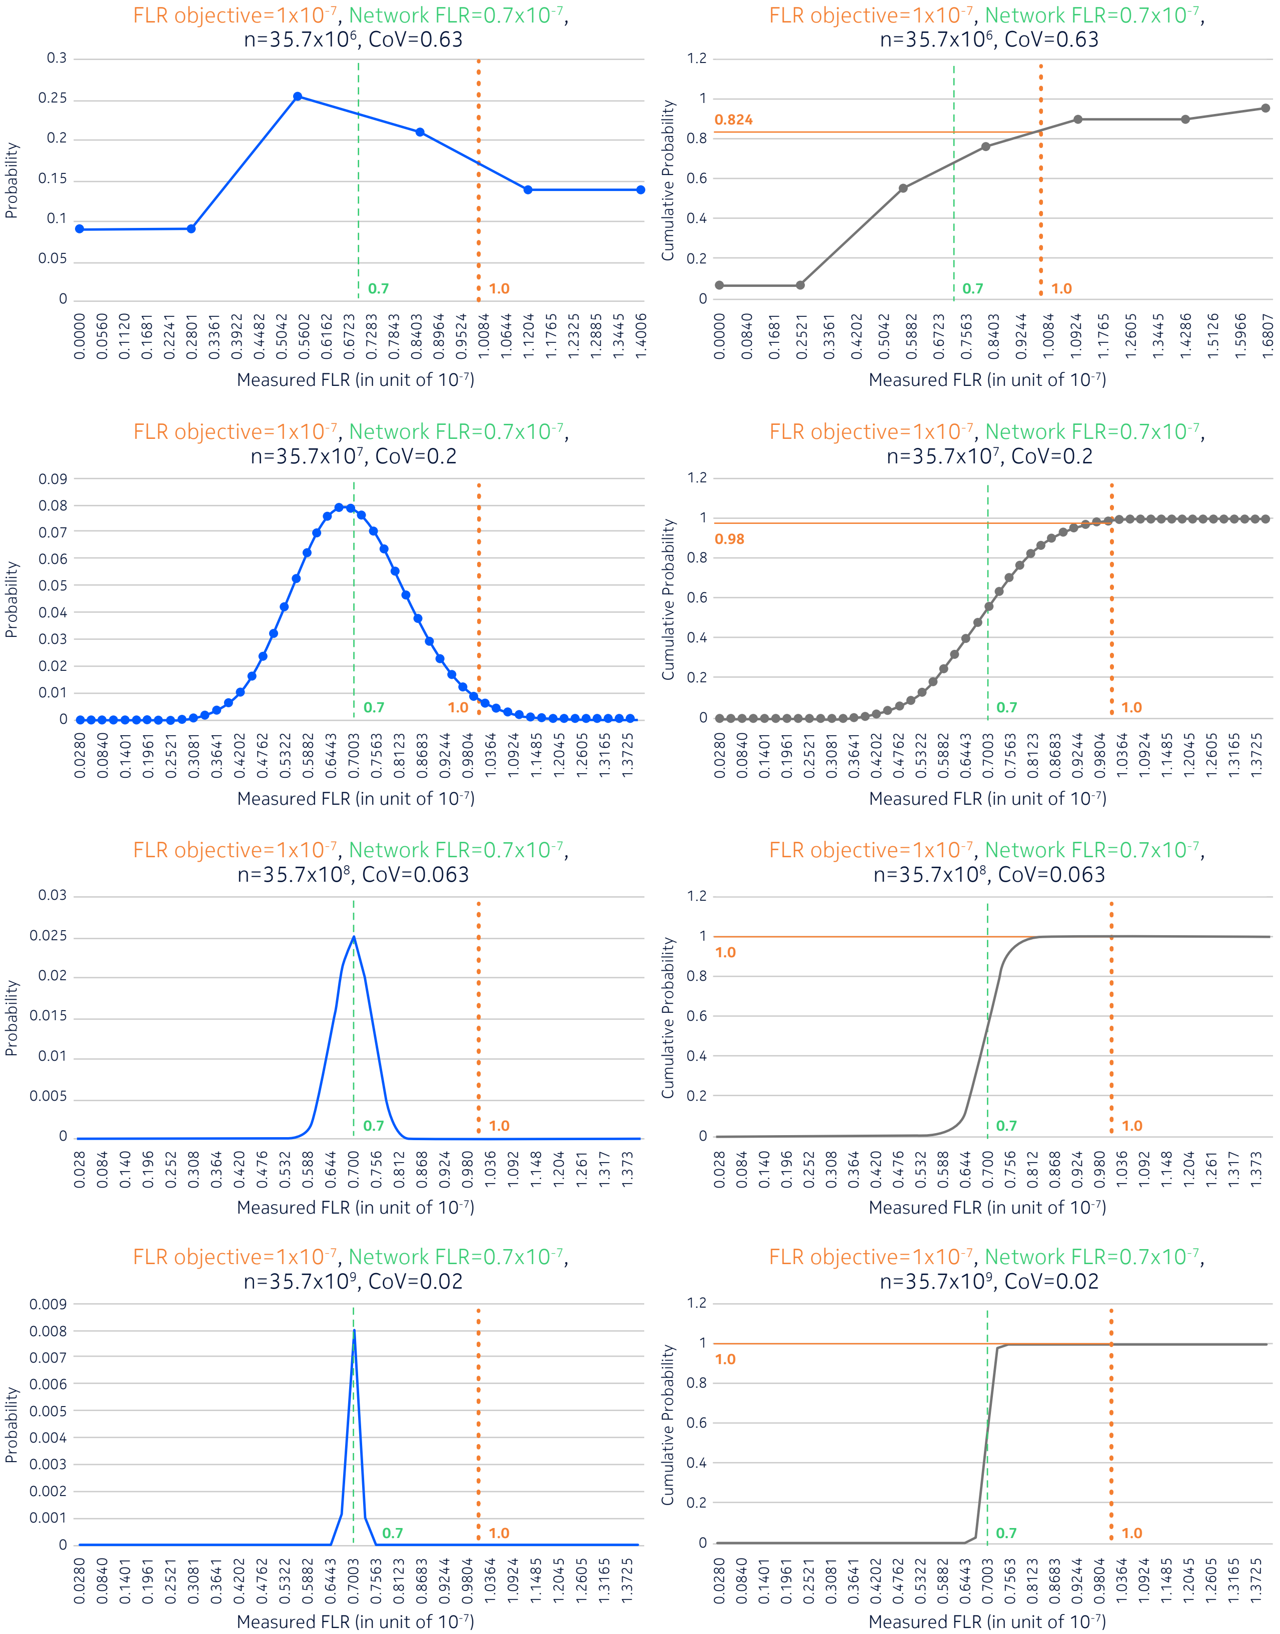

The right measurement interval for FLR

Once the service provider has determined the minimum number of samples required to measure the FLR with the right level of confidence, the next challenge is to determine how long it takes for the OAM tool (ETH-SLM or ETH-LM) to collect this minimum number of samples.

Figure 3 builds on the previous example of the required sample size (35.7 × 10^7 samples) to show that the required measurement time decreases as the sampling rate (number of samples per second) increases. For a 10,000 times faster sampling rate, the measurement time decreases by the same factor, as illustrated by the orange solid line compared with the blue dotted line.

Figure 3. Time for the measured FLR to be stable around the Network FLR and be within the EVC Service FLR objective

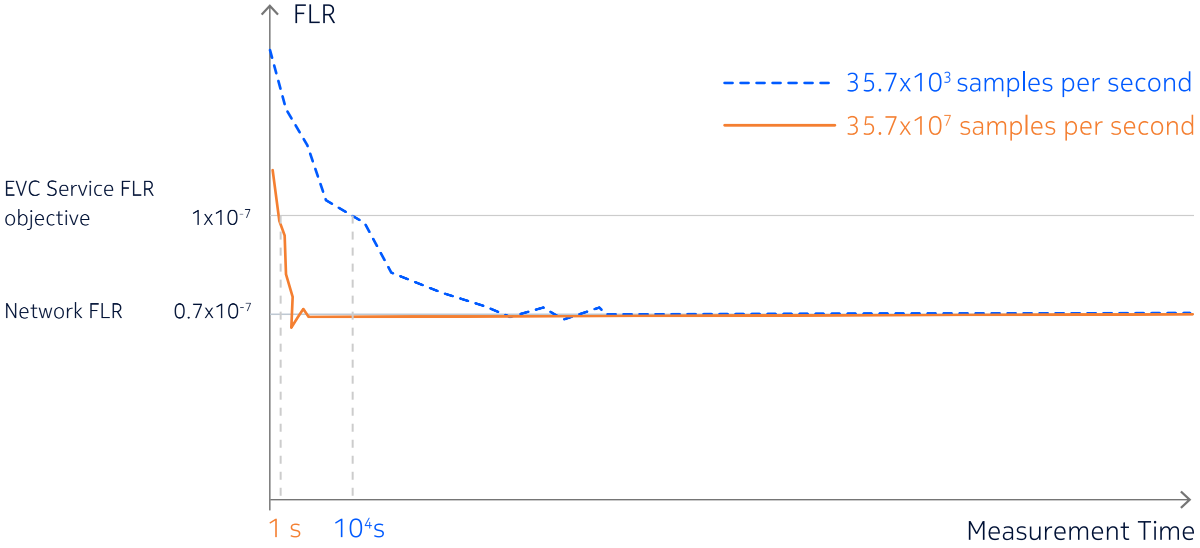

Different sampling rates can be used to analyze the measurement time needed for each loss measurement method.

For ETH-LM, which counts service frames, the following Ethernet port rates are used as examples: 100 Mbps, 1 Gbps, 10 Gbps, 25 Gbps or 100 Gbps, with a Media Access Control (MAC) frame size of 1,000 bytes. To achieve FLR accuracy with confidence, it takes at least 29,131 seconds (8.1 hours), 2,913 seconds (48.6 minutes), 291 seconds (4.86 minutes) or 29 seconds, respectively.

For ETH-SLM, which counts synthetic frames, the possible ETH-SLM frame transmission periods of 10 seconds, 1 second, 100 milliseconds, 10 milliseconds, 3.33 milliseconds are used as examples. To achieve FLR accuracy with confidence, it takes 3,570,000,000 seconds (113.2 years), 357,000,000 seconds (11.3 years), 35,700,000 seconds (1.13 years), 3,570,000 seconds (41.3 days) or 1,188,810 seconds (13.76 days), respectively.

In the above examples, achieving an FLR accuracy of 10^{-7} with confidence means that the measured FLR has a 95 percent probability of being lower than the FLR objective when the CoV is less than or equal to 0.2.

Figure 4 uses these example numbers to show the time that ETH-LM and ETH-SLM take to achieve an FLR accuracy of 10^{-7}. It clearly illustrates the advantage of using ETH-LM.

Figure 4. Comparison of the time taken by ETH-LM vs. ETH-SLM to achieve a FLR accuracy of 10^{-7}

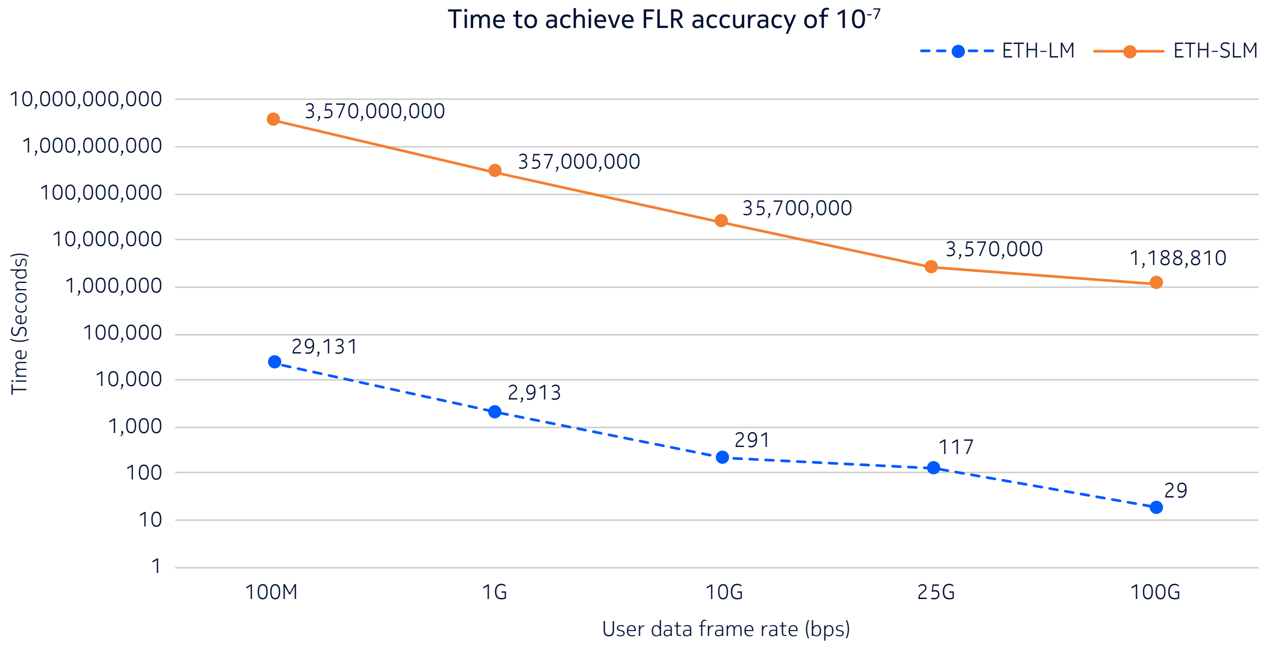

In general, it takes a long time to achieve a certain FLR accuracy. As indicated earlier, per MEF 35.1 [8], FLR is a performance metric that necessitates a long measurement time.

From the above comparison, it can be observed that ETH-LM takes much less time than ETH-SLM to achieve the desired FLR accuracy, i.e., to approach the Network FLR with confidence. For example, ETH-SLM would take 13.76 days to measure FLR performance with an FLR objective of 10^{-7} in the network at a bandwidth of 100 Gbps using the lowest ETH-SLM frame transmission period (3.33 milliseconds), while ETH-LM would take only 29 seconds. This means ETH-LM may take 1 / 10,000 of the measurement time of ETH-SLM. Therefore, the key benefit of ETH-LM is that the FLR measurement is quick – at least one order of magnitude faster than ETH-SLM.

Even though ETH-LM can detect the loss of a single frame and instantaneously measure the EVC Service FLR required to be monitored by MEF 10.4 [1], it still needs some time to stabilize and approach the Network FLR with the right confidence level. If the measurement time is too short, the measured FLR may sometimes violate the FLR objective with a probability of more than 5 percent. This would prevent a service provider from meeting the service-level agreement (SLA) contracted with its customers.

Moreover, with ETH-LM, the sample size is dependent on the service frame rate rather than the transmission rate of the ETH-LM frames, so the transmission period of ETH-LM frames need not be as small as that of ETH-SLM frames. For example, the transmission period for ETH-LM frames is 100 milliseconds, 1 second or 10 seconds, whereas that for ETH-SLM frames is 3.33 milliseconds, 10 milliseconds, 100 milliseconds, 1 second or 10 seconds. At a 100 millisecond transmission period, for example, ETH-LM consumes one-thirtieth of the bandwidth used by ETH-SLM, commensurately reducing the bandwidth consumed by OAM.

Conclusion

It is critical for service providers to have an accurate and efficient tool for measuring FLR performance with high confidence. Achieving more stringent FLR down to 10^{-7} or even less required by new applications brings new challenges with respect to the measurement time needed to achieve sufficient FLR accuracy. This white paper examined the rules needed to monitor high-performance CE services, taking into account the minimum sample size required to ensure that the measured FLR can meet the EVC Service FLR objective with high confidence. It also compared the measurement time intervals required to achieve a given FLR accuracy between ETH-LM and ETH-SLM, concluding that ETH-LM reaches this accuracy several orders of magnitude faster than ETH-SLM while consuming less bandwidth. The other advantages of ETH-LM over ETH-SLM are detailed in a companion Nokia white paper [7].

Abbreviations

AR	augmented reality
CE	Carrier Ethernet
CHLI	Consecutive High Loss Intervals
CoV	Coefficient of Variation
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
LM	Loss Measurement
MAC	Media Access Control
MEF	MEF Forum
OAM	Operations, Administration and Maintenance
PDU	Protocol Data Unit
QoS	Quality of Service
SLA	Service-Level Agreement
SLS	Service Level Specification
VR	virtual reality

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

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

[5] 3GPP TS 23.203 Policy and charging control architecture (Release 17)

[6] 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

[7] Nokia white paper “Improving service assurance in packet transport networks”

[8] MEF 35.1 Service OAM Performance Monitoring Implementation Agreement

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) CID213193