Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 23, 2022

A. Clemm
J. Strassner
Futurewei
J. Francois
Inria

October 20, 2021

High-Precision Service Metrics
draft-csfx-ippm-hipmetrics-00

Abstract

This document defines a set of metrics for high-precision networking services. These metrics can be used to assess the service levels that are being delivered for a networking flow. Specifically, they can be used to determine the degree of compliance with which service levels are being delivered relative to service level objectives that were defined for the flow. The metrics can be used as part of flow records and/or accounting records. They can also be used to continuously monitor the quality with which high-precision networking service are being delivered.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 23, 2022.

Copyright Notice

Table of Contents

1. Introduction

Many networking applications increasingly rely on high-precision networking services that have clearly defined service level objectives (SLOs), for example with regards to end-to-end latency. Applications requiring such services include industrial networks, for example cloud-based industrial controllers for precision machinery, vehicular applications, for example tele-driving in which a vehicle is remotely controlled by a human operators, or Augmented Reality / Virtual Reality (AR/VR) applications involving rendering of point clouds remotely. Many of those applications are not tolerant of degrading service levels. A slight miss in SLOs does not merely result in a slight deterioration of the Quality of Experience to end users, but may render the application inoperable. At the same time, many of those applications are mission critical, in which sudden failures can jeopardize safety or have other adverse consequences. However, clearly those applications represent significant business opportunity demanding dependable technical solutions.

Because of this, efforts such as Deterministic Networking (DetNet) [RFC8655] are attempting to create solutions in which clear bounds on parameters such as end-to-end latency and jitter can be defined in order to make service levels being delivered predictable and, ideally, deterministic. However, one area that has not kept pace concerns metrics that can account for service levels with which services are delivered, specifically the degree of precision for agreed-upon service level objectives. Such metrics, and the instrumentation to support them, are important for a number of purposes, including monitoring (to ensure that networking services

are performing according to their objectives) as well as accounting (to maintain a record of service levels actually delivered, important for monetization of such services as well as for triaging of problems).

The current state-of-the-art of such metrics includes (for example) interface metrics, useful to obtain data on traffic volume and behavior that can be observed at an interface [RFC2863] [RFC8343] but agnostic of actual end-to-end service levels and not specific to distinct flows. Flow records [RFC7011] [RFC7012] maintain statistics about flows, including flow volume and flow duration, but again contain very little information about end-to-end service levels, let alone whether the service levels delivered meet their targets, i.e. their associated SLOs.

This specification introduces a new set of metrics aimed at capturing end-to-end service levels for a flow, specifically the degree to which flows comply with the SLOs that are in effect.

It should be noted that at this point, the set of metrics proposed here is intended as a "starter set" that is intended to spark further discussion. Other metrics are certainly conceivable; we expect that the list of metrics will evolve over time as part of Working Group discussions.

2. Key Words

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Definitions and Acronyms

MTBF: Mean Time Between Failures

SL: Service Level

SLA: Service Level Agreement

SLO: Service Level Objective

4. Metrics

The following section proposes a set of accounting metrics focus on end-to-end latency objectives. They indicate whether any violations of end-to-end latency occurred at the packet level. These metrics

are intended to be applied on a per-flow basis and are intended to assess the degree to which a flow's end-to-end service levels comply with the SLO in effect for that flow.

While the focus in this document concerns end-to-end latency objectives, analogous metrics could also be defined for other end-to-end service level parameters, such as loss (which is distinct from loss occurring at any one given interface) or delay variation.

- Violated Packets. This indicates the number of packets for which a violation of a latency SLO occurred.

- Violated Time Units (e.g. violated seconds, violated milliseconds). This indicates the number of time units during which one or more violations of SLOs were observed, regardless of how many violations took place during the same interval. This measure is useful in scenarios where bursts of violations might suddenly occur (e.g. due to temporary network congestion, during route convergence etc.) and the count of violated packets by itself might paint a misleading picture.

The following additional set of metrics may be useful in certain scenarios as well. However, their precise definition may be subject to policy and further discussion is needed:

- Significantly Violated Packets. This indicates the number of packets for which a "significant" violation occurred, where "significant" implies an SLO that was not merely a near-miss but that missed the objective by a degree determined especially significant.

- Significantly Violated Time Units (e.g. significantly violated seconds, significantly violated milliseconds). This indicates the number of time units during which any significant violation occurred.

- Severely Violated Time Units (e.g. severely violated seconds, severely violated milliseconds). "Severe" here refers to the occurrence of multiple violations within the same time unit. The definition of "severe" may be subject to policy; it may also take into account the significance of the violations that occur.

Note that there is no definition of Severely Violated Packets. The term "severe" is used in conjunction with the occurrence of multiple violations related to multiple packets, not any one packet in isolation.

From these first-order metrics, second-order metrics can be defined that build on the first set of metrics. Some of these metrics are modeled after Mean Time Between Failure, or MTBF metrics - a "failure" in this context referring to a failure to deliver a packet according to its SLO.

- Time since last violated time unit (i.e., since last violated ms, since last violated second). (This parameter is particularly useful for the monitoring of the current health.)

- Packets since last violated packet. (This parameter is particularly useful for the monitoring of the current health.)

- Mean time between violated time units (i.e. between violated milliseconds, between violated seconds). This refers to the arithmetic mean of time between violations such as violated time units.

- Mean packets between violations. This refers to the arithmetic mean of the number of SLO-compliant packets between SLO violations. (Another variation of "MTBF" in a service setting.)

The same set of metrics can also be applied to significant violations, and to severe violations:

- Time since last significantly violated time unit (i.e., since last significantly violated ms, since last significantly violated second).

- Time since last severely violated time unit (i.e., since last severely violated ms, since last severely violated second).

- Packets since last significaty violated packet.

- Mean time between significantly violated time units (i.e. between significantly violated milliseconds, between significantly violated seconds).

- Mean time between severely violated time units (i.e. between severely violated milliseconds, between severely violated seconds).

- Mean packets between significant violations. This refers to the arithmetic mean of the number of SLO-compliant packets between significant SLO violations.

The next set of metrics puts the violations in relationship to non-violations. It is intended to provide an analogous measure to that

of availability, typically defined as the number of time units during which a system (or service) is unavailable divided by the total number of time units. In analogy, a time unit that is "violated" can be viewed as one in which a service is not available with the advertised precision:

- Precision availability (of milliseconds, of seconds): the ratio between violated time units (seconds, milliseconds) and the total time units for the duration of the service.

- Analogous metrics for precision availability re: severely violated time units, re: significantly violated time units.

It should be noted that certain Service Level Agreements may be statistical in nature, requiring the service levels of packets in a flow to adhere to certain distributions. For example, an SLA might state that any given SLO applies only to a certain percentage of packets, allowing for a certain amount of violations to take place. A "violated packet" in that case does not necessarily constitute an SLO violation. However, it is still useful to maintain those statistics, as the number of violated packets still matters when looked at in proportion to the total number of packets.

Along that vein, an SLA might establish an SLO of, say, end-to-end latency to not exceed 20ms for 99% of packets, to not exceed 25ms for 99.999% of packets, and to never exceed 30ms for anything beyond. In that case, any individual packet missing the 20 ms latency target cannot be considered an SLO violation in itself, but compliance with the SLO may need to be assessed after the fact.

To support statistical SLAs more directly, it is feasible to support additional metrics, such as metrics that represent histograms for service level parameters with buckets corresponding to individual service level objectives. For the example just given, a histogram for a given flow could be maintained with three buckets: one containing the count of packets within 20ms, a second with a count of packets between 20 and 25ms (or simply all within 25ms), a third with a count of packet between 25 and 30ms (or simply all packets within 30ms, and a fourth with a count of anything beyond (or simply a total count). Of course, the number of buckets and the boundaries between those buckets should correspond to the needs of the application respectively SLA, i.e. to the specific guarantees and SLOs that were provided. The definition of histogram metrics is for further study.

5. Discussion Items

The following is a list of items for which further discussion is needed as to whether they should be included in the scope of this specification:

- A YANG data model
- A set of IPFIX Information Elements
- Statistical metrics: e.g. histograms/buckets
- Policies regarding the definition of "significant" and "severe" violations
- Additional second-order metrics, such as "longest disruption of service time" (measuring consecutive time units with violations)

6. IANA Considerations

TBD

7. Security Considerations

Instrumentation for metrics that are used to assess compliance with SLOs consitute an interesting target for an attacker. By interfering with the maintaining of such metrics, services could be falsely identified as being in compliance (when they are not), or vice-versa flagged as being non-compliant (when indeed they are). While this document does not specify how networks should be instrumented to maintain the identified metrics, such instrumentation needs to be properly secured to ensure accurate measurements and prohibit tampering with metrics being kept.

Where metrics are being defined relative to an SLO, the configuration of those SLOs needs to be properly secured. Likewise, where SLOs can be adjusted, it needs to be clear which particular SLO any given metrics instance refers to. The same service levels that constitute SLO violations for one flow, and that should be maintained as part of the "violated time units", "violated packets", and related metrics, may be perfectly compliant for another flow. Where it is not possible to properly tie together SLOs and violation metrics, it will be preferrable to merely maintain statistics about sevice levels that were delivered (for example, overall histograms of end-to-end latency), without assessing which of these constitute violations.

By the same token, where the definition of what constitutes a "severe" violation or a "significant" violation depends on policy or

context, the configuration of such policy or context needs to be specially secured and the configuration of this policy be bound to the metrics being maintained. This way it will be clear which policy was in effect when those metrics were being assessed. An attacker that is able to tamper with such policies will render the corresponding metrics useless (in the best case) or misleading (in the worst case).

8. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2863] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", RFC 2863, DOI 10.17487/RFC2863, June 2000, <https://www.rfc-editor.org/info/rfc2863>.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <https://www.rfc-editor.org/info/rfc7011>.

[RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, DOI 10.17487/RFC7012, September 2013, <https://www.rfc-editor.org/info/rfc7012>.

[RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <https://www.rfc-editor.org/info/rfc7950>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8343] Bjorklund, M., "A YANG Data Model for Interface Management", RFC 8343, DOI 10.17487/RFC8343, March 2018, <https://www.rfc-editor.org/info/rfc8343>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/info/rfc8655>.

Authors' Addresses

Alexander Clemm
Futurewei
2330 Central Expressway
Santa Clara CA 95050
USA

Email: ludwig@clemm.org

John Strassner
Futurewei
2330 Central Expressway
Santa Clara CA 95050
USA

Email: strazpdj@gmail.com

Jerome Francois
Inria
615 Rue du Jardin Botanique
Villers-les-Nancy 54600
France

Email: jerome.francois@inria.fr