Authors: G. Mirsky   J. Halpern   X. Min   A. Clemm
         Ericsson   Ericsson   ZTE Corp.   Futurewei
         J. Strassner   J. Francois   L. Han
         Futurewei   Inria   China Mobile

Precision Availability Metrics for SLO-Governed End-to-End Services

Abstract

   This document defines a set of metrics for networking services with
   performance requirements expressed as Service Level Objectives
   (SLO). These metrics, referred to as Precision Availability Metrics
   (PAM), can be used to assess the service levels that are being
   delivered. Specifically, PAM can be used to determine the degree of
   compliance with which service levels are being delivered relative to
   pre-defined SLOs. PAM can be used to provide a service according to
   its SLO as part of accounting records, to account for the actual
   quality with which services were delivered and whether or not any
   SLO violations occurred. Also, PAM can be used to continuously
   monitor the quality with which the service is 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 5 September 2022.

Copyright Notice

Table of Contents

1. Introduction

Network operators and network users often need to assess the quality
with which network services are being delivered. In particular in
cases where service level guarantees are given and service level
objectives (SLOs) are defined, it is essential to provide a measure
of the degree with which actual service levels that are delivered
comply with SLOs that were promised. Examples of service levels
include end-to-end latency and packet loss. Simple examples of SLOs
associated with such service levels would be target values for the
maximum end-to-end latency or maximum amount of loss that would be
deemed acceptable.

To express the quality of delivered networking services versus their
SLOs, corresponding metrics are needed that can be used to
characterize the quality of the service being provided. Of concern
is not so much the absolute service level (for example, actual
latency experienced), but whether the service is provided in
accordance with the contracted service levels. For instance, whether
the latency that is experienced falls within the acceptable range
that has been contracted for the service. The specific quality

depends on the SLO that is in effect. Different groups of
applications set forth requirements for varying sets of service
levels with different target values. Such applications range from
Augmented Reality/Virtual Reality to mission-critical controlling
industrial processes. A non-conformance to an SLO might result in
degradation of the quality of experience for gamers up to
jeopardizing the safety of a large area. However, as those
applications represent significant business opportunities, they
demand dependable technical solutions.

The same service level may be deemed perfectly acceptable for one
application, while unacceptable for another, depending on the needs
of the application. Hence it is not sufficient to simply measure
service levels per se over time, but to assess the quality of the
service being provided with the applicable SLO in mind. However, at
this point, there are no metrics in place that are able to account
for the quality with which services are delivered relative to their
SLOs, and whether their SLOs are being delivered on at all times.
Such metrics and the instrumentation to support them are essential
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
delivered, important for monetization of such services as well as
for triaging of problems).

The current state-of-the-art of metrics available today includes
(for example) interface metrics, useful to obtain data on traffic
volume and behavior that can be observed at an interface [RFC2863]
and [RFC8343] but agnostic of actual end-to-end service levels and
not specific to distinct flows. Flow records [RFC7011] and [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
to meet their targets, i.e., their associated SLOs.

This specification introduces a new set of metrics, Precision
Availability Metrics (PAM), 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. The term "availability" reflects
the fact that a service which is characterized by its SLOs is
considered unavailable whenever those SLOs are violated, even if
basic connectivity is still working. "Precision" refers to the fact
that services whose end-to-end service levels are governed by SLOs,
and which must therefore be precisely delivered according to the
associated quality and performance requirements. It should be noted
that "precision" refers to what is being assessed, not to the
mechanism used to measure it; in other words, it does not refer to
the precision of the mechanism with which actual service levels are
measured. The specification and implementation of methods that

provide for accurate measurements is a separate topic independent of the definition of the metrics in which the results of such measurements would be expressed.

[Ed.note: 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 as part of the Working Group discussions.]

2. Conventions used in this document

2.1. Terminology and Acronyms

[Ed.Note: needs updating.]

PAM Precision Availability Metric

OAM Operations, Administration, and Maintenance

EI Errored Interval

EIR Errored Interval Ratio

SEI Severely Errored Interval

SEIR Severely Errored Interval Ratio

EFI Error-Free Interval

3. Performance Availability Metrics

3.1. Preliminaries

When analyzing the availability metrics of a service flow between two nodes, we need to select a time interval as the unit of PAM. In [ITU.G.826], a time interval of one second is used. That is reasonable, but some services may require different granularity. For that reason, the time interval in PAM is viewed as a variable parameter though constant for a particular measurement session. Further, for the purpose of PAM, each time interval, e.g., second or decamillisecond, is classified either as Errored Interval (EI), Severely Errored Interval (SEI), or Error-Free Interval (EFI). These are defined as follows:

  *An EI is a time interval during which at least one of the performance parameters degraded below its pre-defined optimal level threshold or a defect was detected.

*An SEI is a time interval during which at least one the
    performance parameters degraded below its pre-defined critical
    threshold or a defect was detected.

   *Consequently, an EFI is a time interval during which all
    performance objectives are at or above their respective pre-
    defined optimal levels, and no defect has been detected.

The definition of a state of a defect in the network is also
necessary for understanding the PAM. In this document, the defect is
interpreted as the state of inability to communicate between a
particular set of nodes. It is important to note that it is being
defined as a state, and thus, it has conditions that define entry
into it and exit out of it. Also, the state of defect exists only in
connection to the particular group of nodes in the network, not the
network as a domain.

From these defitions, a set of basic metrics can be defined that
count the numbers of time intervals that fall into each category:

   *EI count.

   *SEI count.

   *EFI count.

3.2. Derived Performance Availability Metrics

A set of metrics can be created based on PAM introduced in Section
3. In this document, these metrics are referred to as derived PAM.
Some of these metrics are modeled after Mean Time Between Failure
(MTBF) metrics - a "failure" in this context referring to a failure
to deliver a packet according to its SLO.

   *Time since the last errored interval (e.g., since last errored
    ms, since last errored second). (This parameter is suitable for
    the monitoring of the current health.) [Ed. note: Need a
    definition of "current health". Is there an alternative to
    "current"? Past health?]

   *Packets since the last errored packet. (This parameter is
    suitable for the monitoring of the current health.)

   *Mean time between EIs (e.g., between errored milliseconds,
    errored seconds) is the arithmetic mean of time between
    consecutive EIs.

   *Mean packets between EIs is the arithmetic mean of the number of
    SLO-compliant packets between consecutive EIs. (Another variation
    of "MTBF" in a service setting.)

An analogous set of metrics can be produced for SEI:

   *Time since the last SEI (e.g., since last errored ms, since last
    errored second). (This parameter is suitable for the monitoring
    of the current health.)

   *Mean time between SEIs (e.g., between severely errored
    milliseconds, severely errored seconds) is the arithmetic mean of
    time between consecutive SEIs.

   *Mean packets between SEIs is the arithmetic mean of the number of
    SLO-compliant packets between consecutive SEIs. (Another
    variation of "MTBF" in a service setting.)

Determining the period in which the path is currently PAM-wise is
helpful. But because switching between periods requires ten
consecutive intervals, shorter conditions may not be adequately
reflected. Two additional PAMs can be used, and they are defined as
follows:

   *errored interval ratio (EIR) is the ratio of EI to the total
    number of time unit intervals in a time of the availability
    periods during a fixed measurement interval.

   *severely errored interval ratio (SESR) - is the ratio of SEIs to
    the total number of time unit intervals in a time of the
    availability periods during a fixed measurement interval.

3.3. Network Availability in Performance Availability Metrics

The definitions of EI, SEI, and EFI allow for characterization of
the communication between two nodes relative to the level of
required and acceptable performance and when performance degrades
below the acceptable level. The former condition in this document
referred to as network availability. The latter - network
unavailability. Based on the definitions, SEI is the one time
interval of network unavailability while EI and EFI present an
interval of network availability. But since the conditions of the
network are everchanging periods of network availability and
unavailability need to be defined with duration larger than one time
interval to reduce the number of state changes while correctly
reflecting the network condition. The method to determine the state
of the network in terms of PAM is described below:

   *If ten consecutive SEIs been detected, then the PAM state of the
    network is determined as unavailability, and the beginning of
    that period of unavailability state is at the start of the first
    SEI in the sequence of the consecutive SEIs.

*Similarly, ten consecutive non-SEIs, i.e., either EIs or EFIs,
     indicate that the network is in the availability period, i.e.,
     available. The start of that period is at the beginning of the
     first non-SEI.

    *Resulting from these two definitions, a sequence of less than ten
     consecutive SEIs or non-SEIs does not change the PAM state of the
     network. For example, if the PAM state is determined as
     unavailability, a sequence of seven EFIs is not viewed as an
     availability period.

4. Statistical SLO

   It should be noted that certain Service Level Agreements (SLA) may
   be statistical, requiring the service levels of packets in a flow to
   adhere to specific distributions. For example, an SLA might state
   that any given SLO applies only to a certain percentage of packets,
   allowing for a certain level of, for example, packet loss and/or
   exceeding packet delay threshold take place. Each such event, in
   that case, does not necessarily constitute an SLO violation.
   However, it is still useful to maintain those statistics, as the
   number of out-of-SLO 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 packets between 25 and 30ms (or merely 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. **Availability of Anything-as-a-Service**

   Anything as a service (XaaS) describes a general category of
   services related to cloud computing and remote access. These
   services include the vast number of products, tools, and
   technologies that are delivered to users as a service over the
   Internet. In this document, the availability of XaaS is viewed as
   the ability to access the service over a period of time with pre-
   defined performance objectives. Among the advantages of the XaaS
   model are:

     *Improving the expense model by purchasing services from providers
      on a subscription basis rather than buying individual products,
      e.g., software, hardware, servers, security, infrastructure, and
      install them on-site, and then link everything together to create
      networks.

     *Speeding new apps and business processes by quickly adapting to
      changing market conditions with new applications or solutions.

     *Shifting IT resources to specialized higher-value projects that
      use the core expertise of the company.

   But XaaS model also has potential challenges:

     *Possible downtime resulting from issues of internet reliability,
      resilience, provisioning, and managing the infrastructure
      resources.

     *Performance issues caused by depleted resources like bandwidth,
      computing power, inefficiencies of virtualized environments,
      ongoing management and security of multi-cloud services.

     *Complexity impacts enterprise IT team that must remain in the
      process of the continued learning of the provided services.

   The framework and metrics of the PAM defined in Section 3 allow a
   provider of XaaS and their customers to quantify, measure, monitor
   for conformance what is often referred to as an ephemeral -
   availability of the service to be delivered. There are other
   definitions and methods of expressing availability. For example,
   [HighAvailability-WP] uses the following equation:

       Availability Average = MTBF/(MTBF + MTRR),
       where:
       MTBF (Mean Time Between Failures) - mean time between individual
       component failures. For example, a hard drive malfunction or
       hypervisor reboot.
       MTTR (Mean Time To Repair) - refers to how long it takes to fix
       the broken component or the application to come back online,

While this approach estimates the expected availability of a XaaS,
the PAM reflects near-real-time availability of a service as
experienced by a user. It also provides valuable data for more
accurate and realistic MTBF and MTTR in the particular environment,
and simplifies comparison of different solutions that may use
redundant servers (web and database), load balancers.

In another field of communication, mobile voice and data services,
the definition of service availability is understood as "the
probability of successful service reception: a given area is
declared "in-coverage" if the service in that area is available with
a pre-specified minimum rate of success. Service availability has
the advantage of being more easily understandable for consumers and
is expressed as a percentage of the number of attempts to access a
given service." [BEREC-CP]. The definition of the availability used
in the PAM throughout this document is close to the quoted above. It
might be considered as the extension that allows regulators,
operators, and consumers to compare not only the rate of
successfully establishing a connection but the quality of the
connection during its lifetime.

6. Other PAM Benefits

PAM provides a number of important benefits with other, more
conventional performance metrics. Without PAM, it would be possible
to conduct ongoing measurements of service levels and maintain a
time-series of service level records, then assess compliance with
specific SLOs after the fact. However, doing so would require the
collection of vast amounts of data that would need to be generated,
exported, transmitted, collected, and stored. In addition, extensive
postprocessing would be required to compare that data against SLOs
and analyze its compliance. Being able to perform these tasks at
scale and in real-time would present significant additional
challenges.

Adding PAM allows for a more compact expression of service level
compliance. In that sense, PAM does not simply represent raw data
but expresses actionable information. In conjunction with proper
instrumentation, PAM can thus help avoid expensive postprocessing.

7. Discussion Items

The following items require further discussion:

  *Terminology - "Errored" vs. "Violated". The key metrics defined
   in this draft refer to intervals during which violations of
   objectives for service level parameters occur as "errored". The
   term "errored" was chosen in continuity with the concept of
   "errored seconds", often used in transmission systems. However,

"violated" may be a more accurate term, as the metrics defined
here are not "errors" in an absolute sense, but relative to a set
of defined objectives.

   *Metrics. The foundational metrics defined in this draft refer to
    errored/violated intervals. In addition, counts of errors/
    violations related to individual packets may also need to be
    maintained. Metrics referring to violated/errored packets, i.e.
    packets that on an individual basis miss a performance objective
    may be added in a later revision of this document.

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 "errored" and "severely
    errored" time interval.

   *Additional second-order metrics, such as "longest disruption of
    service time" (measuring consecutive time units with SEIs).

8. IANA Considerations

   TBA

9. Security Considerations

   Instrumentation for metrics that are used to assess compliance with
   SLOs constitute an attractive target for an attacker. By interfering
   with the maintaining of such metrics, services could be falsely
   identified as complying (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 adequately
   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 adequately secured.
   Likewise, where SLOs can be adjusted, the correlation between any
   metrics instance and a particular SLO must be clear. The same
   service levels that constitute SLO violations for one flow that
   should be maintained as part of the "errored time units" and related
   metrics, may be perfectly compliant for another flow. In cases when

it is impossible to tie together SLOs and PAM properly, it will be preferable to merely maintain statistics about service levels delivered (for example, overall histograms of end-to-end latency) without assessing which constitutes violations.

By the same token, where the definition of what constitutes a "severe" or a "significant" error depends on policy or context. The configuration of such policy or context needs to be specially secured. Also, the configuration of this policy must 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 can tamper with such policies will render the corresponding metrics useless (in the best case) or misleading (in the worst case).

10. Acknowledgments

TBA

11. References

11.1. Informative References

[BEREC-CP] Body of European Regulators for Electronic Communications, "BEREC Common Position on information to consumers on mobile coverage", Common Approaches/ Positions BoR (18) 237, June 2018, <https:// berec.europa.eu/eng/document_register/subject_matter/ berec/regulatory_best_practices/ common_approaches_positions/8315-berec-common-position- on-information-to-consumers-on-mobile-coverage>.

[HighAvailability-WP] Avi Freedman, Server Central, "High Availability in Cloud and Dedicated Infrastructure", <https://www.deft.com/wp-content/uploads/pdf/SCTG-High- Availability-White-Paper-Part-2.pdf>.

[ITU.G.826] ITU-T, "End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections", ITU-T G.826, December 2002.

[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>.

[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>.

Authors' Addresses

   Greg Mirsky
   Ericsson

   Email: gregimirsky@gmail.com

   Joel Halpern
   Ericsson

   Email: joel.halpern@ericsson.com

   Xiao Min
   ZTE Corp.

   Email: xiao.min2@zte.com.cn

   Alexander Clemm
   Futurewei
   2330 Central Expressway
   Santa Clara, CA 95050
   United States of America

   Email: ludwig@clemm.org

   John Strassner
   Futurewei
   2330 Central Expressway
   Santa Clara, CA 95050
   United States of America

   Email: strazpdj@gmail.com

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

   Email: jerome.francois@inria.fr

Liuyan Han
China Mobile
32 XuanWuMenXi Street
Beijing
100053
China

Email: hanliuyan@chinamobile.com