PCE
Internet-Draft
Intended status: Informational
Expires: 16 August 2024
L. M. Contreras
Telefonica
F. Agraz
S. Spadaro
Universitat Politecnica de Catalunya
13 February 2024

Path Computation Based on Precision Availability Metrics
draft-contreras-pce-pam-02

Abstract

The Path Computation Element (PCE) is able of determining paths according to constraints expressed in the form of metrics. The value of the metric can be signaled as a bound or maximum, meaning that path metric must be less than or equal such value. While this can be sufficient for certain services, some others can require the utilization of Precision Availability Metrics (PAM). This document defines a new object, namely the PRECISION METRIC object, to be used for path calculation or selection for networking services with performance requirements expressed as Service Level Objectives (SLO) using PAM.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The Path Computation Element (PCE) [RFC4655] is able of determining paths according to constraints expressed in the form of metrics. For that purpose, the METRIC object is defined in [RFC5440]. The value of the metric included in the METRIC object can be signaled as a bound or maximum, meaning that path metric must be less than or equal such value.

While this can be sufficient for certain services, some others can require the utilization of Precision Availability Metrics (PAM) [I-D.ietf-ippm-pam]. That is the case of services like Network Slice [I-D.ietf-teas-ietf-network-slices] or deterministic [RFC5878] [RFC8655] services. These networking services express their performance requirements by means of Service Level Objectives (SLO) with target values for certain metrics.

At the time of calculating a path by the PCE, the METRIC object [RFC5440] serves for the purposes of indicating either the metric that MUST be optimized by the path computation algorithm, or a bound on the path cost that MUST NOT be exceeded for the path to be considered as acceptable. The value of the metric refers to the

instantaneous observed behavior of that parameter, without a notion of behavior along the preceding time. This cannot be sufficient for certain networking services which require to experience stable behavior along the time according to their SLOs.

The precision availability metrics indicate whether or not a given service has been available according to expectations along the time, for whatever SLO considered as constraint. Thus, at the time of computing a path for networking services described by means of SLOs, it is convenient to express the applicable metric constraints according to the definition of precision availability metrics. This permits the PCE to calculate paths showing a behavior compatible to the desired SLOs over a period. This document defines new object, namely the PRECISION METRIC object, using PAM for that purpose.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT","SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

In addition, the terms defined in [I-D.ietf-ippm-pam] are also used in this document.

3. Rationale of the usage of PAM for path calculation

[I-D.ietf-ippm-pam] introduced the concept of intervals for measuring the behavior of measurable performance parameters against some predefined thresholds. Those intervals consider a given time window. Thus, it is possible to define a Violated Interval (VI) as the time interval during which at least one of the performance parameters presents degradation respect to a predefined optimal level threshold. Similarly, when the threshold is defined as critical, the degradation of the performance parameter in a time window generates a Severe Violated Interval (SVI).

Taking into account the VIs and SVIs it is feasible to generate availability metrics showing some degree of historic behavior in the form of the following ratios:

* Violated Interval Ratio (VIR), defined as the ratio of the summed numbers of VIs and SVIs to the total number of time unit intervals along a predefined availability period.

* Severely Violated Interval Ratio (SVIR), defined as the ratio of SVIs to the total number of time unit intervals along a predefined availability period.

At the time of provisioning a networking service which requires stable SLOs along the time, it is important to ensure that the selected path has shown such stable behavior in the past. Despite the fact that the past behavior is not a guarantee of future behavior, it can be presumed that those paths with lower VIR and SVIR will better satisfy the SLOs of the networking service. Alternatively, PAM can be used by the path computation entity for fine-grained path computation. Then PAM is a useful criteria for calculating and selecting paths.

4. Definition of the PRECISION METRIC Object

The PRECISION METRIC object is defined according to the following structure.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Flags  |C|S|     Type      | Stat Function |     Tiers     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    AvPeriod   |    TI_Units   |          TI_Value             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   Violated Interval Ratio                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               Severely Violated Interval Ratio                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
~                          Thresholds                           ~
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The following fields are defined.

* From the Flags field, two flags are defined in this document.

  o S flag (Statistical - 1 bit): determines if the metric follows a statistical distribution function. When S=0, it means that the metric will be assessed against an optimal (for VI) and a critical (for SVI) thresholds. When S=1, it means that the metric will be assessed against a multi-tiered SLO, presenting different thresholds per tier. In case the SLO is defined in N tiers, each tier is associated with a threshold. Following the example in [I-D.ietf-ippm-pam], a latency metric defined in this way could be expressed in the form of

    + not to exceed 30 ms for any packet;
    + to not exceed 25 ms for 99.999% of packets;
    + to not exceed 20 ms for 99% of packets.

o C (Computed Metric - 1 bit), with similar meaning and implications to the C flag defined on the METRIC object in [RFC5440]. That is, when C=1 in a PCReq message it indicates that the PCE MUST provide the computed path precision metric value in the PCRep message.

o Unassigned flags MUST be set to zero on transmission and MUST be ignored on receipt.

* Type (8 bits): specifies the metric type. The valid metric type values are those allocated by IANA for the original METRIC object T field.

  (Note. To check with PCE WG if this is the correct approach, or if alternatively it is convenient to allocate specific values for the PRECISION METRIC object).

* Stat Function field (0 bits): in case S=1, this field determines the statistical function for describing the SLO. The following functions are considered:

  - 0x0: this is a reserved value.
  - 0x1: histogram
  - 0x2: cumulative distribution function
  - 0x3 - 0x255: these are reserved for future use.

  When S=0, this field SHOULD be ignored.

* Tiers (8 bits): determines the number of tiers in which the statistical distribution of the SLO is defined. The following values are considered:

  - 0x0-0x1: these are invalid values.
  - 0x2: two tiers, valid for the case S=0.
  - 0x3- 0x255: multiple tiers, valid for the case S=1.

* AvPeriod (Availability Period - 8 bits): specifies the total number of of time unit intervals to be considered for the calculation of VIR and SVIR shown by the path.

* TI_Units (Time Interval Units - 8 bits): specifies the units for the definition of the time window of the interval. The following units are considered:

- 0x0: this is a reserved value.
- 0x1: microsecond
- 0x2: millisecond
- 0x3: second
- 0x4: minute
- 0x5: hour
- 0x6: day
- 0x7: week
- 0x8: month
- 0x9: year
- 0x10 - 0x255: these are reserved for future use.

A PRECISION METRIC Object with values 0x0 or 0x1 SHOULD be discarded. A PRECISION METRIC Object with S=0 and Tiers field different than 0x2 SHOULD be discarded. This value implies that the Threshold field will be composed by an Optimal Threshold (for VI) and a Critical Threshold (for SVI). Finally, a PRECISION METRIC Object with S=1 and Tiers field lower than 0x3 SHOULD be discarded. When a generic value of N is provided in this field, it implies that the Threshold field will be compose by N-1 thresholds (for VI per tier) and a Critical Threshold (for SVI corresponding to the highest tier).

* TI_Value (Time Interval Value - 8 bits): specifies the numerical value for the definition of the time window of the interval.

* Violated Interval Ratio (32 bits): specifies the expected VIR for the path, encoded in 32 bits in IEEE floating point format [IEEE.754.2019]. The VIR of the path calculated by the PCE SHOULD be lower or equal than this value. The way in which the PCE calculates the VIR is out of scope of this document.

* Severely Violated Interval Ratio (32 bits): specifies the expected SVIR for the path, encoded in 32 bits in IEEE floating point format [IEEE.754.2019]. The SVIR of the path calculated by the PCE SHOULD be lower or equal than this value. The way in which the PCE calculates the SVIR is out of scope of this document.

Regarding the Thresholds field, this will be variable in size depending on the statistical nature of the precision metric. When the metric is defined only according to an optimal and critical thresholds (S=0 case), then only those thresholds are included in the field. However, when the SLO is defined by means of a multi-tiered statistical distribution (S=1 case), then one threshold field is included per tier. In summary, this would be the different possible situations for the Thresholds field:

* S=0, meaning that only an optimal and critical thresholds are considered. In this case, the Thresholds field follows the following structure:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Optimal Threshold Tier Boundary                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Optimal Threshold                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Critical  Threshold                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Optimal Threshold Tier Boundary, the Optimal Threshold and the Critical Threshold fields are encoded in 32 bits in IEEE floating point format [IEEE.754.2019].

* S=1, meaning that only an optimal and critical thresholds are considered. In this case, the Thresholds field follows the following structure:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Tier 1 Boundary                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   Threshold for Tier 1                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                             ...                               ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Tier N-1 Boundary                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   Threshold for Tier N-1                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Critical  Threshold                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

All the Threshold fields are encoded in 32 bits in IEEE floating point format [IEEE.754.2019].

The way in which the PCE calculates the different thresholds is out of scope of this document.

4.1. Summary of the PRECISION METRIC Object

The PRECISION METRIC Object is extended to take into consideration PAMs. The PRECISION METRIC object is defined to accommodate the expression of constraints following the PAM proposition in [I-D.ietf-ippm-pam].

According to the definition before, and depending on the statistical description of the SLO, two different messages can be found.

When S=0 the SLO or metric is defined against an optimal and a critical thresholds. In consequence, the message format is as follows:

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |    Flags    |C|0|      Type     | Stat Function |      0x2      |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |    AvPeriod   |    TI_Units   |           TI_Value            |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                   Violated Interval Ratio                     |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |               Severely Violated Interval Ratio                |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |               Optimal Threshold Tier Boundary                 |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                     Optimal Threshold                         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                    Critical  Threshold                        |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

In this case, the message has a fixed size of 28 bytes.

When S=1 the SLO or metric is defined following an statistical distribution with N tiers, representing a total of N-1 optimal thresholds plus a critical one. In consequence, the message format is as follows:

```
        0                   1                   2                   3
        0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |    Flags  |C|1|      Type      | Stat Function |      0xN      |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |    AvPeriod   |    TI_Units    |          TI_Value             |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                   Violated Interval Ratio                     |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |               Severely Violated Interval Ratio                |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                      Tier 1 Boundary                          |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                   Threshold for Tier 1                        |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       ~                             ...                               ~
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                     Tier N-1 Boundary                         |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                  Threshold for Tier N-1                       |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                    Critical  Threshold                        |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

In this case, the message has a variable size determined by (4+(2N-1))*4 bytes, being N the number of tiers of the SLO statistical distribution.

4.2. Examples on the usage of the PRECISION METRIC Object.

4.2.1. PRECISION METRIC coding examples

The following are examples of usage of the PRECISION METRIC Object. Path Delay metric type is used as precision metric in these examples.

The first example assumes a a networking service characterized by a SLO defined by means of two tiers with optimal threshold of 20 ms for 99,9% of the packet latency samples, and critical threshold of 25 ms. The availability expectation for this service is to show a VIR of 5% and a SVIR of 0,2%. The availability period considered is one day, while the time interval is considered 1 hour. In these conditions, the extended METRIC Object can be descrined as:

```
       0                   1                   2                   3
       0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |    Flags  |C|0|  Type = 12    | Stat Function |      0x2      |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |       24      |      sec      |            3600               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                              5                                |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                             0.2                               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                            99.9                               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                             20                                |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                             25                                |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The second example takes the example of statistical distribution in [I-D.ietf-ippm-pam], where the path delay metric is statistically defined in the form of:

- not to exceed 30 ms for any packet;
- to not exceed 25 ms for 99.999% of packets;
    - to not exceed 20 ms for 99% of packets

Assuming similar VIR, SVIR, availability period and time interval duration. In these conditions, he extended METRIC Object can be descrined as:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Flags    |C|1|  Type = 12    |   Histogram   |      0x3      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       24      |      sec      |             3600              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               5                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              0.2                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              99                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              20                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            99.999                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              25                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              30                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Once the PCE processes these PRECISION METRICT Objects, the PCE will calculate the VIR and SVIR of the different path alternatives and check them against the requested VIR and SVIR. How the PCE calculate the VIR and SVIR is out of scope of this document.

4.2.2. Example on the usage of PRECISION METRIC Object

The example considers a PCC sending a path computation request to the PCE, including a PRECISION METRIC object detailing path delay described in terms of SLO, and a METRIC object indicating that the path loss must not exceed the value of M. The two objects are inserted in the PCReq message as follows:

o First PRECISION METRIC object coded as in the previous examples, depending on the applicable SLO.

o Second METRIC object with B=1, T=14, metric-value=M

In case the PRECISION METRIC contains flag C = 1, as per [RFC5440], in case there is a path satisfying the set of constraints and there is no policy that prevents the return of the computed metric, then the PCE inserts in its response one PRECISION METRIC object with T=12 and the corresponding SLO description for that path (i.e., all the fields contained in the definition of the PRECISION METRIC Object). Additionally, the PCE MAY insert a second METRIC object with B=1, T=14, metric-value=computed path loss.

5. Related work

In the case of deterministic networking, other documents like [I-D.xiong-pce-detnet-bounded-latency] and [I-D.zhang-pce-enhanced-detnet] propose extensions to PCE adapted to deterministic service capabilities. As part of those capabilities specific metrics are considered. Such metrics could be considered as SLOs that can be handled as PAM. This document presents a generic form of using precision availability metrics in PCEP messages, and then permitting its applicability to broader networking scenarios. Thus, this extension could be used instead of ad-hoc extensions in [I-D.xiong-pce-detnet-bounded-latency] and [I-D.zhang-pce-enhanced-detnet].

6. Security and operational considerations

Same security and operational considerations as described in [RFC5440] apply also in this document.

Other security considerations will be addressed in future versions of the document.

7. IANA Considerations

This document defines a new object class for the PCEP. IANA is requested to allocate the following codepoint in the PCEP "Objects" registry.

Value	Description	Reference
TBD1	PRECISION METRIC object	This document

Additional IANA considerations required by this extension will be documented in future document versions.

8. Informative References

[I-D.ietf-ippm-pam]
Mirsky, G., Halpern, J. M., Min, X., Clemm, A., Strassner, J., and J. François, "Precision Availability Metrics for Services Governed by Service Level Objectives (SLOs)", Work in Progress, Internet-Draft, draft-ietf-ippm-pam-09, 1 December 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-ippm-pam-09>.

[I-D.ietf-teas-ietf-network-slices]
Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "A Framework for

Network Slices in Networks Built from IETF Technologies", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slices-25, 14 September 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slices-25>.

[I-D.xiong-pce-detnet-bounded-latency]
Xiong, Q., Liu, P., and R. Gandhi, "PCEP Extension for DetNet Bounded Latency", Work in Progress, Internet-Draft, draft-xiong-pce-detnet-bounded-latency-03, 8 June 2023, <https://datatracker.ietf.org/doc/html/draft-xiong-pce-detnet-bounded-latency-03>.

[I-D.zhang-pce-enhanced-detnet]
Zhang, L., Geng, X., and T. Zhou, "PCEP for Enhanced DetNet", Work in Progress, Internet-Draft, draft-zhang-pce-enhanced-detnet-03, 9 July 2023, <https://datatracker.ietf.org/doc/html/draft-zhang-pce-enhanced-detnet-03>.

[IEEE.754.2019]
"754-2019 - IEEE Standard for Floating-Point Arithmetic", 22 July 2019, <https://ieeexplore.ieee.org/document/8766229>.

[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>.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <https://www.rfc-editor.org/info/rfc4655>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <https://www.rfc-editor.org/info/rfc5440>.

[RFC5878] Brown, M. and R. Housley, "Transport Layer Security (TLS) Authorization Extensions", RFC 5878, DOI 10.17487/RFC5878, May 2010, <https://www.rfc-editor.org/info/rfc5878>.

[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>.

Acknowledgements

This work has been partially funded by the European Commission Horizon Europe SNS JU PREDICT-6G project (GA 101095890), and the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union NextGenerationEU UNICO 5G I+D "Towards a smart and efficient telecom infrastructure meeting current and future industry needs" (TIMING) project (TSI-063000-2021-145, -148, -149).

Authors' Addresses

Luis M. Contreras
Telefonica
Ronda de la Comunicacion, s/n
28050 Madrid
Spain
Email: luismiguel.contrerasmurillo@telefonica.com
URI: http://lmcontreras.com

Fernando Agraz
Universitat Politecnica de Catalunya
08034 Barcelona
Spain
Email: fernando.agraz@upc.edu

Salvatore Spadaro
Universitat Politecnica de Catalunya
08034 Barcelona
Spain
Email: salvatore.spadaro@upc.edu