

Internet Engineering Task Force  
Internet-Draft  
Intended status: Informational  
Expires: January 7, 2010

R. Geib, Ed.  
Deutsche Telekom  
R. Fardid  
Covad Communications  
July 6, 2009

IPPM standard compliance testing  
draft-geib-ippm-metrictest-00

## Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 7, 2010.

## Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

## Abstract

This document specifies tests to determine if multiple, independent, and interoperable implementations of a metrics specification document

Internet-Draft

IPPM standard compliance testing

July 2009

are at hand so that the metrics specification can be advanced to an Internet standard. Results of different IPPM implementations can be compared if they measure under the same underlying network conditions. Results are compared using state of the art statistical methods.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">1.1.</a>	Requirements Language . . . . .	<a href="#">4</a>
<a href="#">2.</a>	Basic idea . . . . .	<a href="#">4</a>
<a href="#">3.</a>	Verification of equivalence by statistic measurements . . . . .	<a href="#">5</a>
<a href="#">4.</a>	Recommended Metric Verification Measurement Process . . . . .	<a href="#">12</a>
<a href="#">5.</a>	Acknowledgements . . . . .	<a href="#">14</a>
<a href="#">6.</a>	Contributors . . . . .	<a href="#">14</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">15</a>
<a href="#">8.</a>	Security Considerations . . . . .	<a href="#">15</a>
<a href="#">9.</a>	References . . . . .	<a href="#">15</a>
<a href="#">9.1.</a>	Normative References . . . . .	<a href="#">15</a>
<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">15</a>
	Authors' Addresses . . . . .	<a href="#">16</a>

Internet-Draft

IPPM standard compliance testing

July 2009

## 1. Introduction

Draft bradner-metrictest [[bradner-metrictest](#)] states:

The Internet Standards Process [RFC2026](#) [[RFC2026](#)] requires that for a IETF specification to advance beyond the Proposed Standard level, at least two genetically unrelated implementations must be shown to interoperate correctly with all features and options. There are two distinct reasons for this requirement.

In the case of a protocol specification, the notion of "interoperability" is reasonably intuitive - the implementations must successfully "talk to each other", while exercising all features and options.

In the case of a specification for a performance metric, network latency for example, exactly what constitutes "interoperation" is less obvious. The IESG didn't yet decide how to judge "metric specification interoperability" in the context of the IETF Standards Process and this new draft suggests a methodology which (hopefully) is suitable for IPPM metrics. General applicability of the methods proposed in the following should however not be excluded.

A metric specification describes a method of testing and a way to report the results of this testing. One example of such a metric would be a way to test and report the latency that data packets would incur while being sent from one network location to another.

Since implementations of testing metrics are by their nature stand-alone and do not interact with each other, the level of the interoperability called for in the IETF standards process cannot be simply determined by seeing that the implementations interact properly. Instead, verifying equivalence by proofing that different implementations verifiably give statistically equivalent results Verifiable equivalence may take the place of interoperability.

This document defines the process of verifying equivalence by using a specified test set up to create the required separate data sets (which may be seen as samples taken from the same underlying distribution) and then apply state of the art statistical methods to verify equivalence of the results. To illustrate application of the process defined her, validating compliance with [RFC2679](#) [[RFC2679](#)] is picked as an example. While test set ups may vary with the metrics to be validated, the statistical methods will not. Documents defining test setups to validate other metrics should be created by the IPPM WG, once the process proposed here has been agreed upon.

## [1.1.](#) Requirements Language

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 [RFC 2119](#) [[RFC2119](#)].

## [2.](#) Basic idea

Two different IPPM implementations are expected to measure statistically equivalent results, if they both measure a metric under the same networking conditions. Formulating the measurement in statistical terms: separate samples are collected (by separate metric implementations) from the same underlying statistical process (the same network conditions). The "statistical hypothesis" to be tested is the expectation, that both samples expose statistically equivalent properties. This requires careful test design:

- o The error induced by the sample size must be small enough to minimize its influence on the test result. This may have to be respected, especially if two implementations measure with different average probing rates.
- o If time series are compared, the implementation with the lowest probing frequency determines the smallest temporal interval for which results can be compared.
- o Every comparison must be repeated several times based on different measurement data to avoid random indications of compatibility (or

the lack of it).

- o The measurement test set up must be self-consistent to the largest possible extent. This means, network conditions, paths and IPPM metric implementations SHOULD be identical for the compared implementations to the largest possible degree to minimize the influence of the test and measurement set up on the result. This includes e.g. aspects of the stability and non-ambiguity of routes taken by the measurement packets. See [RFC 2330](#) for a discussion on self-consistency [RFC 2330](#) [[RFC2330](#)].

State of the art statistical methods are proposed for a comparison of measurement results in the hope that user friendly tools required to perform the necessary statistical analysis are easily accessible.  
[editor: this sentence may be reworded or deleted, if the expectation doesn't hold].

Let's assume a one way delay measurement comparison between system A, probing with a frequency of 2 probes per second and system B probing

at a rate of 2 probes every 3 minutes. To ensure reasonable confidence in results, sample metrics are calculated from at least 5 singletons per compared time interval. This means, sample delay values are calculated for each system for identical 6 minute intervals for the whole test duration. Per 6 minute interval, the sample metric is calculated from 720 singletons for system A and from 6 singletons for system B). Note, that if outliers are not filtered, moving averages are an option for an evaluation too. The minimum move of an averaging interval is three minutes in our example.

The test set up for the delay measurement is chosen to minimize errors by locating one system of each implementation at the same end of two separate sites, between which delay is measured for the metric test. Both measurement sites are connected by one IPSEC tunnel, so that all measurement packets cross the Internet with the same IP addresses. Both measurement systems measure simultaneously and the local links are dimensioned to avoid congestion caused by the probing traffic itself.

The measured delay values are reported with a resolution above the measurement error and above the synchronisation error. This is done to avoid comparing these errors between two different metric

implementations instead of comparing the IPPM metric implementation itself.

The overall duration of the test is chosen so that more than 1000 six minute measurement intervals are collected. The amount of data collected allows separate comparisons for e.g. 200 consecutive 6 minute intervals. intervals, during which routes were instable, are discarded prior to evaluation.

### 3. Verification of equivalence by statistic measurements

Following the definition of statistical precision [[Precision](#)], a measurement process can be characterised by two properties:

- o Accuracy, which is the degree of conformity of a measured quantity to its actual (true) value.
- o Precision, also called reproducibility or repeatability, the degree to which repeated measurements show the same or similar results.

Figure 1 further clarifies the difference between accuracy and precision of a measurement.

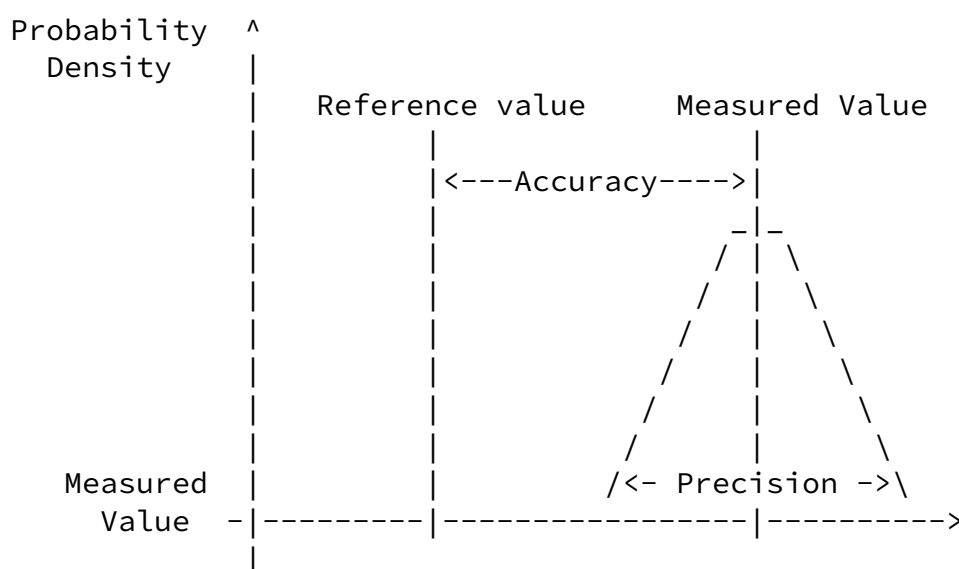


Figure 1

The Framework for IP Performance Metrics ([RFC 2330](#), [[RFC2330](#)]) expects that a "methodology for a metric should have the property that it is repeatable: if the methodology is used multiple times under identical conditions, it should result in consistent measurements." This means, an IPPM implementation is expected to measure a metric with high precision.

Further, [RFC2330](#) expects that a "a methodology for a given metric exhibits continuity if, for small variations in conditions, it results in small variations in the resulting measurements. Slightly more precisely, for every positive epsilon, there exists a positive delta, such that if two sets of conditions are within delta of each other, then the resulting measurements will be within epsilon of each other." A small variation in conditions in the context of a metric comparison can be seen as two implementations measuring the same metric along the same path.

Two guidelines for an IPPM conformant metric implementation can be taken from these principles:

- o A single IPPM conformant implementation MUST under otherwise identical network conditions produce highly precise results for repeated measurements of the same metric.
- o Two different implementations measuring the same IPPM metric MUST produce results with a rather limited difference if measuring under to the largest extent possible identical network conditions.

In a metric test, both conditions must hold, meaning that repeated tests of two implementations MUST produce precise results for all repetition intervals.

A suitable statistical test and a level of confidence to define whether differences are rather limited and whether a measurement is highly precise are specified below.

[RFC 2330](#) prefers the "empirical distribution function" EDF to describe collections of measurements. [RFC 2330](#) uses the EDF to test goodness of fit of an IPPM flow's inter packet spacing to a Poisson process. To do that, [RFC 2330](#) uses the Anderson-Darling test with a 5% significance. [RFC 2330](#) further determines, that "unless otherwise stated, IPPM goodness-of-fit tests are done using 5% significance."

The principles suggested by [RFC 2330](#) are applied to compare the implementation of IPPM metrics as follows:

- o The empirical distribution function of the singletons or samples resulting from the measurement of a particular metric is forming the basis of a comparison of two IPPM implementations. Note that a parametric description of this distribution is not required.
- o The hypothesis to be validated by an IPPM metric test is that two implementations of an IPPM metric draw probes from the same underlying distribution. The hypothesis is true, if samples of two tested metric implementations follow the same distribution by a significance of 95%. Note that the distribution function from which the probes are drawn itself is irrelevant.
- o The samples taken by two implementations to be tested are compared by an Anderson-Darling k sample test. The Anderson-Darling k sample test is the generalization of the classical Anderson-Darling goodness of fit test, and it is used to test the hypothesis that k independent samples belong to the same population without specifying their common distribution function. [Editor: I couldn't find a complete documentation of that test on the web by a fast search, but a reference to a publication is there and code seems to be available too. Other tests which are documented in Wikipedia for that purpose are Kolmogorov-Smirnov and Chi-Square. it is proposed to make Anderson Darling k sample obligatory/a MUST if code can be appended to this draft. If not, Anderson Darling k sample is recommended and Kolmogorov-Smirnov or Chi Square are optional].

Getting back to the chosen example delay measurement, the captured delays may have been captured singletons ranging from an absolute minimum Delay  $D_{min}$  to values  $D_{min} + 5$  ms. To compare distributions,

the set of singletons of a chosen evaluation interval (e.g. the data



of one of the five 1800 minute capture sequences, see above) is sorted for the frequency of singletons per  $D_{min} + N * 0.5 \text{ ms}$  ( $n = 1, 2, \dots$ ). After that, a comparison of the two probe sets with any of the mentioned tests may be applied.

While constructing the example, some additional rules to calculate and compare samples have been respected. The following two rules are of importance for the IPPM metric tests:

- o To compare different probes of a common underlying distribution in terms of metrics characterising a communication network requires to respect the temporal nature for which the assumption of common underlying distribution may hold. Any singletons or samples to be compared MUST be captured within the same time interval.
- o Whenever sample metrics, samples of singletons or rates are used to characterise measured metrics of a time-interval, at least 5 events of a relevant metric MUST be present to ensure a minimum confidence into the reported value (see Wikipedia on confidence [Rule of thumb]). Note that this criterion is to be respected e.g. when comparing packet loss metrics. Any packet loss measurement interval to be compared with the results of another implementation needs to contain at least five lost packets to have a minimum confidence that these losses didn't happen randomly.
- o The minimum number of singletons or samples to be compared by an Anderson-Darling test is 100 per tested metric implementation. Note that the Anderson-Darling test detects small differences in distributions fairly well and will fail for high number of compared results ([RFC2330](#) mentions an example with 8192 measurements to guarantee a failure of an Anderson-Darling test).

Comparing "Accuracy" of IPPM implementations based on averages and variations may require prior checks for the absence of long range dependency within the compared measurements. Large outliers as typically occurring in the case of long range dependency, can have a serious impact on mean values. The median or percentiles may be more robust measures on which to compare the accuracy of different IPPM implementations. An idea may be to consider data up to a certain percentile, calculate the mean for data up to this percentile and then compare the means of the two implementations. This could be repeated for different percentiles. If long range dependencies impact is limited to large outliers, the method may work for lower percentiles. Whether this makes sense must be confirmed by a statistician, so this attempt requires further study.

IPPM metrics are captured by time series. Time series can be checked

for correlation. There are two expectations on statistical time series properties which should be met by separate measurements probing the same underlying network performance distribution:

- o The Autocorrelation indicates, whether there are any repeating patterns within a time series. For the purpose of this document, it does not matter whether there is autocorrelation in a measurement. It is however expected, that two measurements expose the same autocorrelation on identical "lag" intervals. If calculable, the autocorrelation lies within an interval [-1;1], (see Wikipedia on autocorrelation [[Autocorrelation](#)]).
- o The correlation coefficient "indicates the strength of a linear relationship between two random variables." The two random variables in the case of this document are the measurement time series of the IPPM implementations to be compared. The expectation is, that both are strongly correlated and the resulting correlation coefficient is close to 1, (see Wikipedia on correlation [[Correlation](#)]).

A metric test can derive additional statistics from time series analysis. Further, formulation of a test hypothesis is possible for autocorrelation and the correlation coefficient. It is however not clear, whether an appropriate statistical test to validate the hypothesis by 95% significance exists. Applicability of time series analysis for a metric test requires further input from statisticians.

In the absence of any metric test on time series, any test result SHOULD provide the autocorrelation of the compared metrics time series by lags from 1 to 10. In addition, the value of the correlation coefficient SHOULD be provided. Autocorrelation and Correlation coefficient are expected to be rather close to the value 1.

As mentioned earlier, the time series analysis requires application of identical time intervals to allow a comparison. In our delay example, single sample delay metric values are calculated for 9 minute intervals. If 200 consecutive sample delay metrics with the same start and end interval are available for each implementation, autocorrelation can be calculated for different  $n * 9$  minute lags. The autocorrelation calculated for the time series of each implementation should be very close to the autocorrelation of the other implementation for the same time lag. Further, the correlation coefficient for both time series should be close to 1.

The way to prove that two IPPM metric measurements provide compatible

results then could be performed stepwise:

- o First prove that the two compared implementations have the same precision by comparing statistics of the distribution of singletons (or samples) of a metric by comparing the EDF of the samples captured by the two implementations.
- o Second indicate that two compared implementations produce strongly correlated time series of which each one individually has the same autocorrelation as the other one.

Clock synchronization effects require special attention. Accuracy of one-way active delay measurements for any metrics implementation depends on clock synchronization between the source and destination of tests. Ideally, one-way active delay measurement ([RFC 2679](#), [\[RFC2679\]](#)) test endpoints either have direct access to independent GPS or CDMA-based time sources or indirect access to nearby NTP primary (stratum 1) time sources, equipped with GPS receivers. Access to these time sources may not be available at all test locations associated with different Internet paths, for a variety of reasons out of scope of this document.

When secondary (stratum 2 and above) time sources are used with NTP running across the same network, whose metrics are subject to comparative implementation tests, network impairments can affect clock synchronization, distort sample one-way values and their interval statistics. It is RECOMMENDED to discard sample one-way delay values for any implementation, when one of the following reliability conditions is met:

- o Delay is measured and is finite in one direction, but not the other.
- o Absolute value of the difference between the sum of one-way measurements in both directions and round-trip measurement is greater than X% of the latter value.

Examination of the second condition requires RTT measurement for reference, e.g., based on TWAMP ([RFC5357](#), [RFC 5357](#) [\[RFC5357\]](#)), in conjunction with one-way delay measurement.

Specification of X% to strike a balance between identification of unreliable one-way delay samples and misidentification of reliable samples under a wide range of Internet path RTTs probably requires further study.

An IPPM compliant metric implementation whose measurement requires synchronized clocks is however expected to provide precise measurement results. Any IPPM metric implementation MUST be of a precision of 1 ms (+/- 500 us) with a confidence of 95% if the metric is captured

along an Internet path which is stable and not congested during a measurement duration of an hour or more. [Editor: this latter definition may avoid NTP (stratum 2 or worse) synchronized IPPM implementations from becoming IPPM compliant. However internal PC clock synched implementations can't be rejected that way. Ideas on criteria to deal with the latter are welcome. May drift be one, as GPS synched implementations shouldn't have one or the same on origin and destination, respectively].

Metric tests should be executed under conditions which are identical to the largest possible or necessary extent. As "identical network conditions" are fundamental to the methodology proposed by this document, more input and a thorough discussion is needed to define these. Some thoughts are:

- o In a laboratory environment, NTP synchronisation may have a less serious impact. In a real network, improper synchronisation will be harder to conceal.
- o OWD measurements are of highest precision with well synchronized measurement systems measuring delays along a stable not congested path. Care must be taken to avoid comparing noise and the measurement error respectively instead of the delay.
- o Packet loss, delay variation and packet reordering require a sufficient number of these events to allow for a metric test with the desired confidence. While one could wait for congestion or execute the test across known bottlenecks, this may incur some effort. A question is, whether to test these metrics under laboratory conditions. To generalise this question: can laboratory metric tests be tolerated for metrics whose precision doesn't depend on synchronized clocks?

- o Packet loss and delay variation probably allow for a relaxed definition of "identical test conditions", as it may be sufficient for test packets to share the congested interface or paths to test for these metrics.
- o In a laboratory environment, "stationary" networking conditions can be produced without having to care about parallel resources, applied by carriers to increase capacity. In a commercial network, hashing functions (on addresses and ports) determine which set of resources all the packets in a flow will traverse. Testing in the lab may not remove the parallel resources, but it can provide some time stability that's never assured in live network testing.

- o Applicability of tunnels to avoid the impact of unknown parallel resources applied by networks traversed by measurement packets during a test should be investigated.
- o To determine if some aspects of the metric specifications are clear and unambiguous, some specific conditions in the lab may be simulated to determine if implementations measure them as expected. This it should be tested whether all implementors read the spec the same way. Further, reducing some sources of variation right at the start, will make the job of statistical comparison simpler.
- o Getting access to operator information like load and packet loss counters of a network which was used during a metric test is improbable. But testing across a real network still is desirable for a metric test.

#### 4. Recommended Metric Verification Measurement Process

The proposal made by the authors of [bradner-metrictest](#) [[bradner-metrictest](#)] is picked up and slightly enhanced:

"In order to meet their obligations under the IETF Standards Process the IESG must be convinced that each metric specification advanced to

Draft Standard or Internet Standard status is clearly written, that there are the required multiple verifiably equivalent implementations, and that all options have been implemented.

"In the context of this memo, metrics are designed to measure some characteristic of a data network. An aim of any metric definition should be that it should be specified in a way that can reliably measure the specific characteristic in a repeatable way."

Each metric, statistic or option of those to be validated must be compared against a reference measurement or another implementation by at least 5 different basic data sets, each on with sufficient size to reach the specified level of confidence.

"In the same way, sequentially running different implementations of software that perform the tests described in the metric document on a stable network, or simultaneously on a network that may or may not be stable should produce essentially the same results."

Following these assumptions any recommendation for the advancement of a metric specification needs to be accompanied by an implementation report, as is the case with all requests for the advancement of IETF specifications. The implementation report needs to include a

specific plan to test the specific metrics in the RFC in lab or real-world networks and reports of the tests performed with two or more implementations of the software. The test plan should cover key parts of the specification, specify the accuracy required for each measured metric and thus define the meaning of "statistically equivalent" for the specific metrics being tested. Ideally, the test plan would co-evolve with the development of the metric, since that's when people have the most context in their thinking regarding the different subtleties that can arise.

In particular, the implementation report MUST as a minimum document:

- o The metric compared and the RFC specifying it, including the chosen options (like e.g. the implemented selection function in the case of IPDV).
- o A complete specification of the measurement stream (mean rate, statistical distribution of packets, packet size (or mean packet

size and their distribution), DSCP and any other measurement stream property which could result in deviating results. Deviations in results can be caused also if chosen IP addresses and ports of different implementations can result in different layer 2 or layer 3 paths due to operation of Equal Cost Multi-Path routing in an operational network

- o The duration of each measurement to be used for a metric validation, the number of measurement points collected for each metric during each measurement interval (i.e. the probe size) and the level of confidence derived from this probe size for each measurement interval
- o The result of the statistical tests performed for each metric validation.
- o The measurement configuration and set up
- o A parameterization of laboratory conditions and applied traffic and network conditions allowing reproduction of these laboratory conditions for readers of the implementation report.

"All of the tests for each set MUST be run in the same direction between the same two points on the same network. The tests SHOULD be run simultaneously unless the network is stable enough to ensure that the path the data takes through the network will not change between tests."

It is RECOMMENDED to avoid effects falsifying results of real data networks, if validation measurements are taken over them. Obviously,

the conditions met there can't be reproduced. As the measurement equipment compared is designed to reliably quantify real network performance, validating metrics under real network conditions is desirable of course.

Data networks may forward packets differently in the case of:

- o Different packet sizes chosen for different metric implementations. A proposed countermeasure is selecting the same packet size when validating results of two samples or a sample against an original distribution.

- o Selection of differing IP addresses and ports used by different metric implementations during metric validation tests. If ECMP is applied on IP or MPLS level, different paths can result (note that it may be impossible to detect an MPLS ECMP path from an IP endpoint). A proposed counter measure is to connect the measurement equipment to be compared by a NAT device, or establishing a single tunnel to transport all measurement traffic. The aim is to have the same IP addresses and port for all measurement packets or to avoid ECMP by a layer 2 tunnel.
- o Different IP options.
- o Different DSCP.

The test design may have to be adapted for the purpose of the measurement. Creation of delay and delay variation probes is simple and straightforward, also if the measurement runs across a real data network. Collecting a large number of packet loss samples on a real data network while being sure that operational conditions are stable may not be feasible. Further discussion on test designs to verify specific metrics may indeed be required.

## 5. Acknowledgements

Gerhard Hasslinger commented a first version of this document, suggested statistical tests and the evaluation of time series information. Henk Uijterwaal pushed this work and Mike Hamilton reviewed the document before publication.

## 6. Contributors

Scott Bradner, Vern Paxson and Allison Mankin drafted `bradner-metrictest` [[bradner-metrictest](#)], and major parts of it are quoted in this document. Al Morton and Scott Bradner commented this draft

before publication.

## 7. IANA Considerations



This memo includes no request to IANA.

## 8. Security Considerations

This draft does not raise any specific security issues.

## 9. References

### 9.1. Normative References

- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", [RFC 2330](#), May 1998.
- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", [RFC 2679](#), September 1999.

### 9.2. Informative References

- [Autocorrelation]  
N., N., "Autocorrelation", December 2008.
- [Correlation]  
N., N., "Correlation", June 2009.
- [Precision]  
N., N., "Accuracy and precision", June 2009.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", [RFC 5357](#), October 2008.
- [Rule of thumb]  
N., N., "Confidence interval", October 2008.

[bradner-metrictest]

Bradner, S., Mankin, A., and V. Paxson, "Advancement of metrics specifications on the IETF Standards Track", draft -bradner-metricstest-03, (work in progress), July 2007.

#### Authors' Addresses

Ruediger Geib (editor)  
Deutsche Telekom  
Heinrich Hertz Str. 3-7  
Darmstadt, 64295  
Germany

Phone: +49 6151 628 2747  
Email: Ruediger.Geib@telekom.de

Reza Fardid  
Covad Communications  
2510 Zanker Road  
San Jose, CA 95131  
USA

Phone: +1 408 434-2042  
Email: RFardid@covad.com

