

Network Working Group  
Internet Draft

G. Almes, Advanced Network & Services  
W. Cervený, Advanced Network & Services  
P. Krishnaswamy, BellCore  
J. Mahdavi, Pittsburgh Supercomputer Center  
M. Mathis, Pittsburgh Supercomputer Center  
V. Paxson, Lawrence Berkeley Labs

Expiration Date: January 1997

July 1996

Framework for IP Provider Metrics  
<[draft-almes-ippm-framework-00.txt](#)>

## 1. Status of this Memo

This document is an Internet Draft. 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'`.

To learn the current status of any Internet Draft, please check the ``[id-abstracts.txt](#)'' listing contained in the Internet Drafts shadow directories on [ftp.is.co.za](#) (Africa), [nic.nordu.net](#) (Europe), [munnari.oz.au](#) (Pacific Rim), [ds.internic.net](#) (US East Coast), or [ftp.isi.edu](#) (US West Coast).

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

## 2. Introduction

The purpose of this memo is to define a general framework for particular metrics to be developed by the IP Provider Metrics (IPPM) effort within the Benchmarking Methodology Working Group (BMWG) of the Operational Requirements Area (OR) of the IETF.

We begin by laying out several criteria for the metrics that we adopt. These criteria are designed to promote an IPPM effort that will maximise an accurate common understanding by Internet users and Internet providers of the performance and reliability both of end-to-end paths through the Internet and of specific 'IP clouds' that

ID

Framework for IP Provider Metrics

July 1996

comprise portions of those paths.

We next define some Internet vocabulary that will allow us to speak clearly about Internet components such as routers, paths, and clouds.

We next define three fundamental concepts, metrics, measurement methodology, and uncertainties/errors, that will allow us to speak clearly about specific metrics. Given these concepts, we proceed to discuss how they relate to the analytical framework shared by many aspects of the Internet engineering discipline. We then introduce the notion of empirically defined metrics, and continue to discuss two forms of composition.

In some sections of the memo, we will surround some commentary text with the brackets {Comment: ... }. We stress that this commentary is only commentary, and is not itself part of the framework document or a proposal of particular metrics. In some cases this commentary will discuss some of the properties of metrics that might be envisioned, but the reader should assume that any such discussion is intended only to shed light on points made in the framework document, and not to suggest any specific metrics.

### 3. Criteria for IP Provider Metrics

The overarching goal of the IP Provider Metrics effort is to achieve a situation in which users and providers of Internet transport service have an accurate common understanding of the performance and reliability of the Internet component 'clouds' that they user/provide.

To achieve this, performance and reliability metrics for paths through the Internet must be developed. In several meetings of the BMWG criteria for these metrics have been specified:

- + The metrics must be concrete and well-defined,
- + A methodology for a metric should have the property that it is repeatable: if the methodology is used multiple times under identical conditions, the same measurements should result in the same measurements.
- + The metrics must exhibit no bias for IP clouds implemented with identical technology,
- + The metrics must exhibit understood and fair bias for IP clouds

implemented with non-identical technology,

ID Framework for IP Provider Metrics July 1996

- + The metrics must be useful to users and providers in understanding the performance they experience or provide,
- + The metrics must avoid inducing artificial performance goals.

#### [4. Terminology for Paths and Clouds](#)

The following list defines terms that need to be precise in the development of path metrics. We proceed from low-level notions of host, router, and link, then proceed to define the notions of path and notions of IP cloud and exchange that allow us to segment a path into relevant pieces.

host A computer capable of communicating using the Internet protocols; includes "routers".

link A single link-level connection between two (or more) hosts; includes leased lines, ethernet, frame relay clouds, etc.

router

A host which facilitates network-level communication between hosts by forwarding IP packets.

path A sequence of the form  $\langle h_0, l_1, h_1, \dots, l_n, h_n \rangle$ , where  $n \geq 0$ , each  $h_i$  is a host, each  $l_i$  is a link between  $h_{i-1}$  and  $h_i$ , each  $h_1 \dots h_{n-1}$  is a router. In an appropriate operational configuration, the links and routers in the path facilitate network-layer communication of packets from  $h_0$  to  $h_n$ . Note that path is a unidirectional concept.

subpath

Given a path, a subpath is any subsequence of the given path which is itself a path. (Thus, the first and last element of a subpath is a host.)

## cloud

An undirected (possibly cyclic) graph whose vertices are routers and whose edges are links that connect pairs of routers. Formally, ethernet, frame relay clouds, and other links that connect more than two routers are modelled as fully-connected meshes of graph edges. Note that to connect to a cloud means to connect to a router of the cloud over a link; this link is not itself part of the cloud.

## exchange

A special case of a link, an exchange directly connects either a host to a cloud and/or one cloud to another cloud.

Almes et al.

[Page 3]

---

ID

Framework for IP Provider Metrics

July 1996

## cloud subpath

A subpath of a given path, all of whose hosts are routers of a given cloud.

## path digest

A sequence of the form  $\langle h_0, e_1, C_1, \dots, e_n, h_n \rangle$ , where  $n \geq 0$ ,  $h_0$  and  $h_n$  are hosts, each  $e_1 \dots e_n$  is an exchange, and each  $C_1 \dots C_{n-1}$  is a cloud subpath.

## [5. Three Fundamental Concepts](#)

### [5.1. Metrics](#)

In the operational Internet, there are several quantities related to the performance and reliability of the Internet that we'd like to know the value of. When such a quantity is carefully specified, we term the quantity a metric. We anticipate that there will be separate RFCs for each metric (or for each closely related group of metrics).

In some cases, there might be no obvious means to effectively measure the metric; this is allowed, and even understood to be very useful in some cases. It is required, however, that the specification of the metric be as clear as possible about what quantity is being specified. Thus, difficulty in practical measurement is sometimes allowed, but ambiguity in meaning is not.

Each metric will be defined in terms of standard units of measurement. The international metric system will be used, with the following points specifically noted:

- + When a unit is expressed in simple meters (for distance/length) or seconds (for duration), appropriate related units based on thousands or thousandths of acceptable units are acceptable. Thus, distances expressed in kilometers (Km), durations expressed in milliseconds (msec), or microseconds (usec) are allowed, but not centimeters (because the prefix is not in terms of thousands or thousandths).
- + When a unit is expressed in a combination of units, appropriate related units based on thousands or thousandths of acceptable units are acceptable, but all such thousands/thousandths must be grouped at the beginning. Thus, kilo-meters per second (Km/sec) is allowed, but meters per millisecond is not.

- + The unit of information is the bit.
- + When metric prefixes are used with bits or with combinations including bits, those prefixes will have their metric meaning (related to decimal 1000), and not the meaning conventional with computer storage (related to decimal 1024). In any RFC that defines a metric whose units include bits, this convention will be followed and will be repeated to ensure clarity for the reader.
- + When a time is given, it will be taken in UTC.

Note that these points apply to the specifications for metrics and not, for example, to packet formats where octets will likely be used in preference/addition to bits.

Finally, we note that some metrics may be defined purely in terms of other metrics; such metrics are call 'derived metrics'.

## [5.2.](#) Measurement Methodology

For a given set of well-defined metrics, a number of distinct measurement methodologies may exist. A partial list includes:

- + Direct measurement of a performance metric using injected test traffic. Example: measurement of the round-trip delay of an IP

- packet of a given size over a given route at a given time.
- + Projection of a metric from lower-level measurements. Example: given accurate measurements of propagation delay and bandwidth for each step along a path, projection of the complete delay for the path for an IP packet of a given size.
- + Estimation of a constituent metric from a set of more aggregated measurements. Example: given accurate measurements of delay for a given one-hop path for IP packets of different sizes, estimation of propagation delay for the link of that one-hop path.
- + Estimation of a given metric at one time from a set of related metrics at other times. Example: given an accurate measurement of flow capacity at a past time, together with a set of accurate delay measurements for that past time and the current time, and given a model of flow dynamics, estimate the flow capacity that would be observed at the current time.

This list is by no means exhaustive. The purpose is to point out the variety of measurement techniques.

When a given metric is specified, a given measurement approach might be noted and discussed. That approach, however, is not formally part of the specification.

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.

Backing off a little from the word 'identical' in the previous paragraph, we could more accurately use the word 'continuity' to describe a property of a given methodology: 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. At this point, this should be taken as a heuristic driving our intuition about one kind of robustness property rather than as a precise notion.

A metric that has at least one methodology that exhibits continuity is said itself to exhibit continuity.

Note that some metrics, such as hop-count along a path, are integer-valued and therefore cannot exhibit continuity in quite the sense given above.

Note further that, in practice, it may not be practical to know (or be able to quantify) the conditions relevant to a measurement at a given time. For example, since the instantaneous load (in packets to be served) at a given router in a high-speed wide-area network can vary widely over relatively brief periods and will be very hard for an external observer to quantify, various statistics of a given metric may be more repeatable, or may better exhibit continuity. In that case those particular statistics should be specified when the metric is specified.

Finally, some measurement methodologies may be 'conservative' in the sense that a measurement that may themselves modify the value of the performance metric they attempt to measure. {Comment: for example, in a wide-area high-speed network under modest load, a test using several small 'ping' packets to measure delay would likely not interfere (much) with the delay properties of that network as observed by others. The corresponding statement about tests using a large flow to measure flow capacity would likely fail.}

### [5.3](#). Measurements, Uncertainties, and Errors

Even the very best measurement methodologies for the very most well behaved metrics will exhibit errors. Those who develop such measurement methodologies, however, should strive to:

- + minimise their uncertainties/errors,
  - + understand and document the sources of uncertainty/error, and
  - + quantify the amounts of uncertainty/error.
- by doing so, the measurement community will work together to improve our ability to understand the performance and reliability of the Internet.

For example, when developing a method for measuring delay, understand

how any errors in your clocks introduce errors into your delay measurement, and quantify this effect as well as you can. In some cases, this will result in a requirement that a clock be at least up to a certain quality if it is to be used to make a certain measurement.

As a second example, consider the timing error due to measurement overheads within computer making the measurement, as opposed to delays due to the Internet component being measured. The former is a measurement error, while the latter reflects the metric of interest. Note that one technique that can help avoid this overhead is the use of a packet filter/sniffer, running on a separate computer that records network packets and timestamps them accurately. The resulting trace can then be analysed to assess the test traffic, minimising the effect of measurement host delays, or at least allowing those delays to be accounted for.

Finally, we note that derived metrics (defined above) or metrics that exhibit spatial or temporal composition (defined below) offer occasion for the analysis of measurement uncertainty of related measurements to be themselves related.

## 6. Metrics and the Analytical Framework

As the Internet has evolved from the early packet-switching studies of the 1960s, the Internet engineering community has evolved a common analytical framework of concepts. This analytical framework, or A-frame, used by designers and implementers of protocols, by those involved in measurement, and by those who study computer network performance using the tools of simulation and analysis, has great advantage to our work. A major objective here is to generate network characterisations that are consistent in both analytical and practical settings, since this will maximise the chances that non-empirical network study can be better correlated with, and used to further our understanding of, real network behavior.

Whenever possible, therefore, we would like to develop and leverage the A-frame. Thus, whenever a metric to be specified is understood to be closely related to concepts (such as the Internet components

defined above) within the A-frame, we will attempt to specify the



metric in the A-frame's terms. In such a specification we will develop the A-frame by precisely defining the concepts needed for the metric, then leverage the A-frame by defining the metric in terms of those concepts.

Such a metric will be called an 'analytically specified metric' or, more simply an analytical metric.

{Comment: Examples of such analytical metrics might include:

propagation time of a link

The time, in seconds, required by a single bit to travel from the output port on one Internet host across a single link to another Internet host.

bandwidth of a link for packets of size k

The capacity, in bits/second, where only those bits of the IP packet are counted, for a packet of size k bytes.

route

The path, as defined in [Section 4](#), from A to B at a given time.

hop count of a route

The value 'n' of the route path.  
}

Note that we make no a priori list of just what A-frame concepts will emerge in these specifications, but we do encourage their use and urge that they be carefully specified so that, as our set of metrics develops, so will a specified set of A-frame concepts technically consistent with each other and consonant with the common understanding of those concepts within the general Internet community.

These A-frame concepts will be intended to abstract from actual Internet components in such a way that:

- + the essential function of the component is retained,
- + properties of the component relevant to the metrics we aim to create are retained,
- + a subset of these component properties are defined as analytical metrics, and
- + those properties of actual Internet components not relevant to defining the metrics we aim to create are dropped.

{Comment: for example, when considering a router in the context of packet forwarding, we might model the router as a component that receives packets on an input link, queues them on a FIFO packet queue

---

of finite size, employs tail-drop when the packet queue is full, and forwards them on an output link. The transmission speed (in bits/second) of the input and output links, the latency in the router (in seconds), and the maximum size of the packet queue (in bits) are relevant analytical metrics.}

In some cases, such analytical metrics used in relation to a router will be very closely related to specific metrics of the performance of Internet paths. For example, an obvious formula  $(L + P/B)$  involving the latency in the router ( $L$ ), the packet size (in bits) ( $P$ ), and the transmission speed of the output link ( $B$ ) might closely approximate the increase in packet delay due to the insertion of a given router along a path.

We stress, however, that well-chosen and well-specified A-frame concepts and their analytical metrics will support more general metric creation effort in less obvious ways.

{Comment: for example, when considering the flow capacity of a path, it may be of real value to be able to model each of the routers along the path as packet forwarders as above. Techniques for estimating the flow capacity of a path might use the maximum packet queue size as a parameter in decidedly non-obvious ways. For example, as the maximum queue size increases, so will the ability of the router to continuously move traffic along an output link despite fluctuations in traffic from an input link. Estimating this increase, however, remains a research topic.}

The key role of these concepts is to abstract the properties of the Internet components relevant to given metrics. Judgement is required to avoid making assumptions that bias the modeling and metric effort toward one kind of design.

{Comment: for example, routers might not use tail-drop, even though tail-drop might be easier to model analytically.}

Note that, when we specify A-frame concepts and analytical metrics, we will inevitably make simplifying assumptions. Further, as noted above, judgement is required in making these assumptions in order to make them best suit our purposes.

Finally, note that different elements of the A-frame might well make different simplifying assumptions. For example, the abstraction of a router used to further the definition of delay might treat the router's packet queue as a single FIFO queue, but the abstraction of a router used to further the definition of the handling of an RSVP-

enabled packet might treat the router's packet queue to support bounded delay -- a contradictory assumption. This is not to say that

we make contradictory assumptions at the same time, but that two different parts of our work might refine the simpler base concept in two divergent ways for different purposes.

## [7.](#) Empirically Specified Metrics

There are useful performance and reliability metrics that do not fit so neatly into the A-frame, usually because the A-frame lacks the complexity or power for dealing with them. For example, "the best flow capacity achievable along a path using an [RFC-1122](#)-compliant TCP" would be good to be able to measure, but we have no analytical framework of sufficient complexity to allow us to cast that flow capacity as an analytical metric.

These notions can still be well specified by instead describing a reference methodology for measuring them.

Such a metric will be called an 'empirically specified metric', or more simply, an empirical metric.

Such empirical metrics should have three properties:

- + we should have a clear definition for each in terms of real-world Internet components,
- + we should have at least one effective means to measure them, and
- + to the extent possible, we should have an (necessarily incomplete) understanding of the metric in terms of the A-frame so that we can use our measurements to reason about the performance and reliability of A-frame components and of aggregations of A-frame components.

## [8.](#) Two Forms of Composition

### [8.1.](#) Spatial Composition of Metrics

In some cases, it may be realistic and useful to define metrics in such a fashion that they exhibit spatial composition.

By spatial composition, we mean a characteristic of some path metrics, in which the metric as applied to a (complete) path can also be defined for various subpaths (cf. definition above), and in which the appropriate A-frame concepts for the metric suggest useful relationships between the metric applied to these various subpaths (including the complete path, the various cloud subpaths of a given path digest, and even single routers along the path). The effectiveness of spatial composition depends:

- + on the usefulness in analysis of these relationships as applied to the relevant A-frame components, and
- + on the practical use of the corresponding relationships as applied to metrics and to measurement methodologies.

{Comment: for example, consider some metric for delay of a 100-byte packet across a path P, and consider further a path digest  $\langle h_0, e_1, C_1, \dots, e_n, h_n \rangle$  of P. The definition of such a metric might include a conjecture that the delay across P is very nearly the sum of the corresponding metric across the exchanges ( $e_i$ ) and clouds ( $C_i$ ) of the given path digest. The definition would further include a note on how a corresponding relation applies to relevant A-frame components, both for the path P and for the exchanges and clouds of the path digest.}

When the definition of a metric includes a conjecture that the metric across the path is related to the metric across the subpaths of the path, that conjecture constitutes a claim that the metric exhibits spatial composition. The definition should then include:

- + the specific conjecture applied to the metric,
- + a justification of the practical utility of the composition in terms of making accurate measurements of the metric on the path, and
- + a justification of the usefulness of the composition in terms of making analysis of the path using A-frame concepts more effective.

## [8.2.](#) Temporal Composition of Formal Models and Empirical Metrics

In some cases, it may be realistic and useful to define metrics in such a fashion that they exhibit temporal composition.

By temporal composition, we mean a characteristic of some path metrics, in which the metric as applied to a path at a given time  $T$  is also defined for various times  $t_0 < t_1 < \dots < t_n < T$ , and in which the appropriate A-frame concepts for the metric suggests useful relationships between the metric applied at times  $t_0, \dots, t_n$  and the metric applied at time  $T$ . The effectiveness of temporal composition depends:

- + on the usefulness in analysis of these relationships as applied to the relevant A-frame components, and
- + on the practical use of the corresponding relationships as applied to metrics and to measurement methodologies.

{Comment: for example, consider some metric for the expected flow capacity across a path  $P$  during the five-minute period surrounding the time  $T$ , and suppose further that we have the corresponding values for each of the four previous five-minute periods  $t_0, t_1, t_2$ , and  $t_3$ .

The definition of such a metric might include a conjecture that the flow capacity at time  $T$  can be estimated from a certain kind of extrapolation from the values of  $t_0, \dots, t_3$ . The definition would further include a note on how a corresponding relation applies to relevant A-frame components.

Note: any (spatial or temporal) compositions involving flow capacity are likely to be subtle, and temporal compositions are generally more subtle than spatial compositions, so the reader should understand that the foregoing example is intentionally naive.}

When the definition of a metric includes a conjecture that the metric across the path at a given time  $T$  is related to the metric across the path for a set of other times, that conjecture constitutes a claim that the metric exhibits temporal composition. The definition should then include:

- + the specific conjecture applied to the metric,
- + a justification of the practical utility of the composition in terms of making accurate measurements of the metric on the path, and
- + a justification of the usefulness of the composition in terms of making analysis of the path using A-frame concepts more effective.

This memo raises no security issues.

## 10. References

[1] Vern Paxson, <ftp://ftp.ee.lbl.gov/papers/metrics-framework-INET96.ps.Z>

## 11. Authors' Addresses

Guy Almes <almes@advanced.org>  
Advanced Network & Services, Inc.  
200 Business Park Drive  
Armonk, NY 10504  
USA  
Phone: +1 914/273-7863

Bill Cervený <cervený@advanced.org>  
Advanced Network & Services, Inc.  
200 Business Park Drive

Almes et al.

[Page 12]

---

ID	Framework for IP Provider Metrics	July 1996
----	-----------------------------------	-----------

Armonk, NY 10504  
USA

Padma Krishnaswamy <kri@bellcore.com>  
Bell Communications Research  
445 South Street  
Morristown, NJ 07960  
USA

Jamshid Mahdavi <mahdavi@psc.edu>  
Pittsburgh Supercomputing Center  
4400 5th Avenue  
Pittsburgh, PA 15213  
USA

Matt Mathis <mathis@psc.edu>  
Pittsburgh Supercomputing Center  
4400 5th Avenue

Pittsburgh, PA 15213  
USA

Vern Paxson <vern@ee.lbl.gov>  
MS 50B/2239  
Lawrence Berkeley National Laboratory  
University of California  
Berkeley, CA 94720  
USA  
Phone: +1 510/486-7504