

Network Working Group
Request for Comments: 2498
Category: Experimental

J. Mahdavi
Pittsburgh Supercomputing Center
V. Paxson
Lawrence Berkeley National Laboratory
January 1999

IPPM Metrics for Measuring Connectivity

Status of this Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

1. Introduction

Connectivity is the basic stuff from which the Internet is made. Therefore, metrics determining whether pairs of hosts (IP addresses) can reach each other must form the base of a measurement suite. We define several such metrics, some of which serve mainly as building blocks for the others.

This memo defines a series of metrics for connectivity between a pair of Internet hosts. It builds on notions introduced and discussed in RFC 2330, the IPPM framework document. The reader is assumed to be familiar with that document.

The structure of the memo is as follows:

- + An analytic metric, called Type-P-Instantaneous-Unidirectional-Connectivity, will be introduced to define one-way connectivity at one moment in time.
- + Using this metric, another analytic metric, called Type-P-Instantaneous-Bidirectional-Connectivity, will be introduced to define two-way connectivity at one moment in time.
- + Using these metrics, corresponding one- and two-way analytic metrics are defined for connectivity over an interval of time.

- + Using these metrics, an analytic metric, called Type-P1-P2-Interval-Temporal-Connectivity, will be introduced to define a useful notion of two-way connectivity between two hosts over an interval of time.
- + Methodologies are then presented and discussed for estimating Type-P1-P2-Interval-Temporal-Connectivity in a variety of settings.

Careful definition of Type-P1-P2-Interval-Temporal-Connectivity and the discussion of the metric and the methodologies for estimating it are the two chief contributions of the memo.

2. Instantaneous One-way Connectivity

2.1. Metric Name:

Type-P-Instantaneous-Unidirectional-Connectivity

2.2. Metric Parameters:

- + Src, the IP address of a host
- + Dst, the IP address of a host
- + T, a time

2.3. Metric Units:

Boolean.

2.4. Definition:

Src has *Type-P-Instantaneous-Unidirectional-Connectivity* to Dst at time T if a type-P packet transmitted from Src to Dst at time T will arrive at Dst.

2.5. Discussion:

For most applications (e.g., any TCP connection) bidirectional connectivity is considerably more germane than unidirectional connectivity, although unidirectional connectivity can be of interest for some security applications (e.g., testing whether a firewall correctly filters out a "ping of death"). Most applications also require connectivity over an interval, while this metric is instantaneous, though, again, for some security applications instantaneous connectivity remains of interest. Finally, one might not have instantaneous connectivity due to a transient event such as a full queue at a router, even if at nearby instants in time one does have connectivity. These points are addressed below, with this metric serving as a building block.

Note also that we have not explicitly defined *when* the packet arrives at Dst. The TTL field in IP packets is meant to limit IP packet lifetimes to 255 seconds (RFC 791). In practice the TTL field can be strictly a hop count (RFC 1812), with most Internet hops being much shorter than one second. This means that most packets will have nowhere near the 255 second lifetime. In principle, however, it is also possible that packets might survive longer than 255 seconds. Consideration of packet lifetimes must be taken into account in attempts to measure the value of this metric.

Finally, one might assume that unidirectional connectivity is difficult to measure in the absence of connectivity in the reverse direction. Consider, however, the possibility that a process on Dst's host notes when it receives packets from Src and reports this fact either using an external channel, or later in time when Dst does have connectivity to Src. Such a methodology could reliably measure the unidirectional connectivity defined in this metric.

3. Instantaneous Two-way Connectivity

3.1. Metric Name:

Type-P-Instantaneous-Bidirectional-Connectivity

3.2. Metric Parameters:

- + A1, the IP address of a host
- + A2, the IP address of a host
- + T, a time

3.3. Metric Units:

Boolean.

3.4. Definition:

Addresses A1 and A2 have *Type-P-Instantaneous-Bidirectional-Connectivity* at time T if address A1 has Type-P-Instantaneous-Unidirectional-Connectivity to address A2 and address A2 has Type-P-Instantaneous-Unidirectional-Connectivity to address A1.

3.5. Discussion:

An alternative definition would be that A1 and A2 are fully connected if at time T address A1 has instantaneous connectivity to address A2, and at time T+dT address A2 has instantaneous connectivity to A1, where T+dT is when the packet sent from A1 arrives at A2. This definition is more useful for measurement, because the measurement

can use a reply from A2 to A1 in order to assess full connectivity. It is a more complex definition, however, because it breaks the symmetry between A1 and A2, and requires a notion of quantifying how long a particular packet from A1 takes to reach A2. We postpone discussion of this distinction until the development of interval-connectivity metrics below.

4. One-way Connectivity

4.1. Metric Name:

Type-P-Interval-Unidirectional-Connectivity

4.2. Metric Parameters:

- + Src, the IP address of a host
 - + Dst, the IP address of a host
 - + T, a time
 - + dT, a duration
- {Comment: Thus, the closed interval [T, T+dT] denotes a time interval.}

4.3. Metric Units:

Boolean.

4.4. Definition:

Address Src has *Type-P-Interval-Unidirectional-Connectivity* to address Dst during the interval [T, T+dT] if for some T' within [T, T+dT] it has Type-P-instantaneous-connectivity to Dst.

5. Two-way Connectivity

5.1. Metric Name:

Type-P-Interval-Bidirectional-Connectivity

5.2. Metric Parameters:

- + A1, the IP address of a host
 - + A2, the IP address of a host
 - + T, a time
 - + dT, a duration
- {Comment: Thus, the closed interval [T, T+dT] denotes a time interval.}

5.3. Metric Units:

Boolean.

5.4. Definition:

Addresses A1 and A2 have **Type-P-Interval-Bidirectional-Connectivity** between them during the interval [T, T+dT] if address A1 has *Type-P-Interval-Unidirectional-Connectivity* to address A2 during the interval and address A2 has *Type-P-Interval-Unidirectional-Connectivity* to address A1 during the interval.

5.5. Discussion:

This metric is not quite what's needed for defining "generally useful" connectivity - that requires the notion that a packet sent from A1 to A2 can elicit a response from A2 that will reach A1. With this definition, it could be that A1 and A2 have full-connectivity but only, for example, at time T1 early enough in the interval [T, T+dT] that A1 and A2 cannot reply to packets sent by the other. This deficiency motivates the next metric.

6. Two-way Temporal Connectivity

6.1. Metric Name:

Type-P1-P2-Interval-Temporal-Connectivity

6.2. Metric Parameters:

- + Src, the IP address of a host
 - + Dst, the IP address of a host
 - + T, a time
 - + dT, a duration
- {Comment: Thus, the closed interval [T, T+dT] denotes a time interval.}

6.3. Metric Units:

Boolean.

6.4. Definition:

Address Src has **Type-P1-P2-Interval-Temporal-Connectivity** to address Dst during the interval [T, T+dT] if there exist times T1 and T2, and time intervals dT1 and dT2, such that:

- + T1, T1+dT1, T2, T2+dT2 are all in [T, T+dT].
- + T1+dT1 <= T2.
- + At time T1, Src has Type-P1 instantaneous connectivity to Dst.
- + At time T2, Dst has Type-P2 instantaneous connectivity to Src.
- + dT1 is the time taken for a Type-P1 packet sent by Src at time T1 to arrive at Dst.
- + dT2 is the time taken for a Type-P2 packet sent by Dst at time T2 to arrive at Src.

6.5. Discussion:

This metric defines "generally useful" connectivity -- Src can send a packet to Dst that elicits a response. Because many applications utilize different types of packets for forward and reverse traffic, it is possible (and likely) that the desired responses to a Type-P1 packet will be of a different type Type-P2. Therefore, in this metric we allow for different types of packets in the forward and reverse directions.

6.6. Methodologies:

Here we sketch a class of methodologies for estimating Type-P1-P2-Interval-Temporal-Connectivity. It is a class rather than a single methodology because the particulars will depend on the types P1 and P2.

6.6.1. Inputs:

- + Types P1 and P2, addresses A1 and A2, interval [T, T+dT].
 - + N, the number of packets to send as probes for determining connectivity.
 - + W, the "waiting time", which bounds for how long it is useful to wait for a reply to a packet.
- Required: $W \leq 255$, $dT > W$.

6.6.2. Recommended values:

dT = 60 seconds.
W = 10 seconds.
N = 20 packets.

6.6.3. Algorithm:

- + Compute N *sending-times* that are randomly, uniformly distributed over $[T, T+dT-W]$.
- + At each sending time, transmit from A1 a well-formed packet of type P1 to A2.
- + Inspect incoming network traffic to A1 to determine if a successful reply is received. The particulars of doing so are dependent on types P1 & P2, discussed below. If any successful reply is received, the value of the measurement is "true". At this point, the measurement can terminate.
- + If no successful replies are received by time $T+dT$, the value of the measurement is "false".

6.6.4. Discussion:

The algorithm is inexact because it does not (and cannot) probe temporal connectivity at every instant in time between $[T, T+dT]$. The value of N trades off measurement precision against network measurement load. The state-of-the-art in Internet research does not yet offer solid guidance for picking N . The values given above are just guidelines.

6.6.5. Specific methodology for TCP:

A TCP-port-N1-port-N2 methodology sends TCP SYN packets with source port N1 and dest port N2 at address A2. Network traffic incoming to A1 is interpreted as follows:

- + A SYN-ack packet from A2 to A1 with the proper acknowledgement fields and ports indicates temporal connectivity. The measurement terminates immediately with a value of "true". {Comment: if, as a side effect of the methodology, a full TCP connection has been established between A1 and A2 -- that is, if A1's TCP stack acknowledges A2's SYN-ack packet, completing the three-way handshake -- then the connection now established between A1 and A2 is best torn down using the usual FIN handshake, and not using a RST packet, because RST packets are not reliably delivered. If the three-way handshake is not completed, however, which will occur if the measurement tool on A1 synthesizes its own initial SYN packet rather than going through A1's TCP stack, then A1's TCP stack will automatically terminate the connection in a reliable fashion as A2 continues transmitting the SYN-ack in an attempt to establish the connection. Finally, we note that using A1's TCP stack to conduct the measurement complicates the methodology in that the stack may retransmit the initial SYN packet, altering the number of probe packets sent.}

- + A RST packet from A2 to A1 with the proper ports indicates temporal connectivity between the addresses (and a *lack* of service connectivity for TCP-port-N1-port-N2 - something that probably should be addressed with another metric).
- + An ICMP port-unreachable from A2 to A1 indicates temporal connectivity between the addresses (and again a *lack* of service connectivity for TCP-port-N1-port-N2). {Comment: TCP implementations generally do not need to send ICMP port-unreachable messages because a separate mechanism is available (sending a RST). However, RFC 1122 states that a TCP receiving an ICMP port-unreachable MUST treat it the same as the equivalent transport-level mechanism (for TCP, a RST).}
- + An ICMP host-unreachable or network-unreachable to A1 (not necessarily from A2) with an enclosed IP header matching that sent from A1 to A2 *suggests* a lack of temporal connectivity. If by time T+dT no evidence of temporal connectivity has been gathered, then the receipt of the ICMP can be used as additional information to the measurement value of "false".

{Comment: Similar methodologies are needed for ICMP Echo, UDP, etc.}

7. Acknowledgments

The comments of Guy Almes, Martin Horneffer, Jeff Sedayao, and Sean Shapira are appreciated.

8. Security Considerations

As noted in RFC 2330, active measurement techniques, such as those defined in this document, can be abused for denial-of-service attacks disguised as legitimate measurement activity. Furthermore, testing for connectivity can be used to probe firewalls and other security mechanisms for weak spots.

9. References

- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.
- [RFC1122] Braden, R., Editor, "Requirements for Internet Hosts -- Communication Layers", STD, 3, RFC 1122, October 1989.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J. and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.
- [RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

10. Authors' Addresses

Jamshid Mahdavi
Pittsburgh Supercomputing Center
4400 5th Avenue
Pittsburgh, PA 15213
USA

EMail: mahdavi@psc.edu

Vern Paxson
MS 50A-3111
Lawrence Berkeley National Laboratory
University of California
Berkeley, CA 94720
USA

Phone: +1 510/486-7504
EMail: vern@ee.lbl.gov

11. Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

