

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 12, 2009

H. Uijterwaal
RIPE NCC
December 9, 2008

A One-Way Packet Duplication Metric
draft-ietf-ippm-duplicate-06.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of 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 June 12, 2009.

Abstract

When a packet is sent from one host to the other, one normally expects that exactly one copy of the packet that was sent arrives at the destination. It is, however, possible that a packet is either lost or that multiple copies arrive.

In earlier work a metric for packet loss has been defined. This metric quantifies the case where a packet that is sent, does not arrive at its destination within a reasonable time. In this memo, a metric for another case is defined: a packet is sent, but multiple copies arrive. The document also discusses streams and methods to summarize the results of streams.

Internet-Draft

Packet Duplication Metric

December 2008

Table of Contents

1.	Introduction	4
1.1.	Requirements notation	4
1.2.	Motivation	4
2.	A Singleton Definition for one-way packet arrival count	5
2.1.	Metric Name	5
2.2.	Metrics Parameters	5
2.3.	Metric Units	5
2.4.	Definition	5
2.5.	Discussion	6
2.6.	Methodology	7
2.7.	Errors and uncertainties	7
2.8.	Reporting the metric	7
3.	A Singleton Definition for one-way packet duplication	7
3.1.	Metric Name	7
3.2.	Metrics Parameters	7
3.3.	Metric Units	7
3.4.	Definition	7
3.5.	Discussion	8
4.	Definition for samples for one-way Packet Duplication	8
4.1.	Poisson Streams	8
4.1.1.	Metric Name	8
4.1.2.	Metric Parameters	8
4.1.3.	Metric Units	8
4.1.4.	Definition	8
4.1.5.	Methodology	9
4.1.6.	Errors and uncertainties	9
4.1.7.	Reporting the metric	9
4.2.	Periodic Streams	9
4.2.1.	Metric Name	9
4.2.2.	Metric Parameters	9
4.2.3.	Metric Units	9
4.2.4.	Definition	9
4.2.5.	Methodology	9
4.2.6.	Errors and uncertainties	10
4.2.7.	Reporting the metric	10
5.	Some statistics definitions for one-way Duplication	10
5.1.	Type-P-one-way-packet-duplication-fraction	10
5.2.	Type-P-one-way-replicated-packet-rate	10
5.3.	Examples	11
6.	Security Considerations	11
7.	IANA Considerations	12

[8.](#) Acknowledgements [12](#)
[9.](#) References [13](#)
 [9.1.](#) Normative References [13](#)
 [9.2.](#) Informative References [13](#)
Author's Address [13](#)

Intellectual Property and Copyright Statements [14](#)

1. Introduction

This document defines a metric for one-way packet duplication across Internet paths. It builds on the IPPM Framework document [[RFC2330](#)]; the reader is assumed to be familiar with that document.

This document follows the same structure as the document for one-way Packet Loss [[RFC2680](#)]; the reader is assumed to be familiar with that document as well.

The structure of this memo is as follows:

- o First, a singleton metric, called Type-P-one-way-packet-arrival-count, is introduced to measure the number of arriving packets for each packet sent.
- o Then, a singleton metric, called Type-P-one-way-packet-duplication, is defined to describe a single instance of packet duplication.
- o Next, this singleton metric is used to define samples, Type-P-one-way-Packet-Duplication-Poisson-Stream and Type-P-one-way-Packet-Duplication-Periodic-Stream. These are introduced to measure duplication in a series of packets sent with either Poisson-distributed [[RFC2680](#)] or periodic [[RFC3432](#)] intervals between the packets.
- o Finally, statistics that summarize the properties of these samples are introduced.

1.1. Requirements notation

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](#)].

Although [RFC 2119](#) was written with protocols in mind, the key words are used in this document for similar reasons. They are used to ensure the results of measurements from two different implementations are comparable, and to note instances when an implementation could perturb the network.

[1.2.](#) Motivation

When a packet is sent from one host to the other, one normally expects that exactly one copy of the packet that was sent arrives at the destination. It is, however, possible that a packet is either lost or that multiple copies arrive.

In earlier work a metric for packet loss has been defined [[RFC2680](#)]. This metric distinguishes between cases where the packet arrives and where the packet does not arrive within a reasonable time. In this

memo, a metric for a third outcome is defined: a single packet is sent but multiple copies arrive.

As this document describes a case similar to the one discussed in [[RFC2680](#)], all considerations from that document on timing and accuracy apply.

[2.](#) A Singleton Definition for one-way packet arrival count

[2.1.](#) Metric Name

Type-P-one-way-packet-arrival-count

[2.2.](#) Metrics Parameters

- o Src, the IP address of a host
- o Dst, the IP address of a host
- o T, the wire time of a packet at the source
- o T₀, the maximum waiting time for a packet to arrive at the destination.

[2.3.](#) Metric Units

An integer number

[2.4.](#) Definition

Two packets are considered identical if and only if:

- o Both contain identical information fields. The recipient thus could take either packet and use the data in an application. The other packet does not contain any additional information.
- o Both packets appear to have been sent by one and the same host, to one and the same destination. Host are identified by their IP addresses.
- o Both contain valid, but not necessarily identical IP header fields.

The recipient thus could take either packet and use it in an application, the other copy does not contain any additional information.

The value of a Type-P-one-way-packet-arrival-count is a positive integer number indicating the number of (uncorrupted and identical) copies received by dst in the interval $[T, T+T_0]$ for a packet sent by src at time T.

If a packet is sent, but it is lost or does not arrive in the interval $[T, T+T_0]$, then the metric is undefined. Applications MAY report an "impossible" value (for example, -1) to indicate this condition instead of undefined.

If a packet is fragmented during transport and if, for whatever reason, re-assembly does not occur, then the packet will be deemed lost. It is thus not included in the Type-P-one-way-packet-arrival-count.

[2.5.](#) Discussion

This metric counts the number of packets arriving for each packet sent. The time-out value T_0 SHOULD be set to a value when the application could potentially still use the packet and would not

discard it automatically.

The IP headers do not necessarily have to be identical. This can happen, for example, if two packets take a different route resulting in a different TTL.

If this metrics is used in parallel with the Packet Loss Metric [[RFC2680](#)], the value of T0 should be the same for both cases in order to keep the results comparable.

The metric only counts packets that are not corrupted during transmission and may have been resent automatically by lower layers or intermediate devices. Packets that were corrupted during transmission but nevertheless still arrived at dst are not counted.

Clocks do have to be synchronized between src and dst such that it is possible to uniquely and accurately determine the interval [T, T+T0] at both sides.

If this metric is used in an active measurement system, the system MUST NOT send multiple packets with identical information fields, in order to avoid that all packets will be declared duplicates. This metric can be used inside a passive measurement system as well, using packets generated by another source. However, if the source can send two identical packets within the interval [T, T+T0], this will be incorrectly labelled as a duplicate, resulting in a false positive. It is up to the implementor to estimate if this scenario is likely to happen and the rate of false positives is acceptable.

In IPv4, "information fields" refers to all data following the IPv4 header. The equivalent of this in IPv6 is all information following the IPv6 header and the hop-by-hop options header.

[2.6.](#) Methodology

The basic technique to measure this metrics follows the methodology described in [[RFC2680](#)], [section 2.6](#), with one exception.

[RFC2680] does not specify that the receiving host should be able to receive multiple copies of a single packet, as it only needs one copy to determine the metrics. Implementations for this metric should

obviously be capable to receive multiple copies.

[2.7.](#) Errors and uncertainties

Refer to [section 2.7 of \[RFC2680\]](#)

[2.8.](#) Reporting the metric

Refer to [section 2.8 of \[RFC2680\]](#)

[3.](#) A Singleton Definition for one-way packet duplication

[3.1.](#) Metric Name

Type-P-one-way-packet-duplication

[3.2.](#) Metrics Parameters

- o Src, the IP address of a host
- o Dst, the IP address of a host
- o T, the wire time of a packet at the source
- o T0, the maximum waiting time for a packet to arrive at the destination.

[3.3.](#) Metric Units

An integer number.

[3.4.](#) Definition

The value of a Type-P-one-way-packet-duplication is a positive integer number indicating the number of (uncorrupted and identical) additional copies of an individual packet received by dst in the interval [T, T+T0] sent by src at time T.

If a packet is sent and only one copy arrives in the interval [T, T+T0], then the metric is 0. If no copy arrives in this interval, then the metric is undefined. Applications MAY report an "impossible" value (for example, -1) to indicate this condition.

[3.5.](#) Discussion

This metric is equal to

Type-P-one-way-packet-arrival-count - 1

This metric is expected to be used for applications that need to know duplication for an individual packet. All considerations regarding methodology, errors and reporting from the previous section apply.

[4. Definition for samples for one-way Packet Duplication](#)

[4.1. Poisson Streams](#)

[4.1.1. Metric Name](#)

Type-P-one-way-Packet-Duplication-Poisson-Stream

[4.1.2. Metric Parameters](#)

- o Src, the IP address of a host
- o Dst, the IP address of a host
- o Ts, a time
- o Tf, a time. Ts and Tf specify the time interval when packets can be sent for this stream.
- o T0, the maximum waiting time for a packet to arrive at the destination.
- o lambda, a rate in reciprocal seconds

[4.1.3. Metric Units](#)

A sequence of pairs; the elements of each pair are:

- o T, a time
- o Type-P-one-way-packet-arrival-count for the packet sent at T.

[4.1.4. Definition](#)

Given Ts, Tf and lambda, we compute a pseudo-random Poisson process beginning at or before Ts, with average rate lambda and ending at or after Tf. Those time values greater than or equal to Ts, and less than or equal to Tf are then selected. At each of the times in this process, we obtain the value of Type-P-one-way-packet-arrival-count. The value of the sample is the sequence made up of the resulting {time, duplication} pairs. If there are no such pairs, the sequence is of length zero and the sample is said to be empty.

[4.1.5.](#) Methodology

Refer to [\[RFC2680\], section 3.6.](#)

[4.1.6.](#) Errors and uncertainties

Refer to [\[RFC2680\], section 3.7.](#)

[4.1.7.](#) Reporting the metric

Refer to [\[RFC2680\], section 3.8.](#)

[4.2.](#) Periodic Streams

[4.2.1.](#) Metric Name

Type-P-one-way-Packet-Duplication-Periodic-Stream

[4.2.2.](#) Metric Parameters

- o Src, the IP address of a host
- o Dst, the IP address of a host
- o Ts, a time
- o Tf, a time. Ts and Tf specify the time interval when packets can be sent for this stream.
- o T0, the maximum waiting time for a packet to arrive at the destination.
- o lambda, a rate in reciprocal seconds

[4.2.3.](#) Metric Units

A sequence of pairs; the elements of each pair are:

- o T, a time
- o Type-P-one-way-packet-arrival-count for the packet sent at T.

[4.2.4.](#) Definition

At time Ts, we start sending packets with a constant rate lambda, until time Tf. For each packet sent, we obtain the value of Type-P-one-way-packet-arrival-count. The value of the sample is the sequence made up of the resulting {time, duplication} pairs. If there are no such pairs, the sequence is of length zero and the sample is said to be empty.

[4.2.5.](#) Methodology

Refer to [\[RFC3432\], section 4.5.](#)

[4.2.6.](#) Errors and uncertainties

Refer to [\[RFC3432\], section 4.6.](#)

[4.2.7.](#) Reporting the metric

Refer to [\[RFC3432\], section 4.7.](#)

[5.](#) Some statistics definitions for one-way Duplication

Note: the statistics described in this section can be used for both Type-P-one-way-Packet-Duplication-Poisson-Stream and Type-P-one-way-Packet-Duplication-Periodic-Stream. The application SHOULD report which sample was used as input.

[5.1.](#) Type-P-one-way-packet-duplication-fraction

This statistic gives the fraction of additional packets that arrived in a stream.

Given a Type-P-one-way-Packet-Duplication-Poisson-Stream, one first removes all values of Type-P-one-way-Packet-Duplication which are undefined. For the remaining pairs in the stream, one calculates: (Sum Type-P-one-Way-packet-arrival-count/Number of pairs left) - 1 (In other words, (number of packets received)/(number of packets sent and not lost).)

The number can be expressed as a percentage.

Note: this statistic is the equivalent of the Y.1540 IPDR [\[Y1540\]](#)

[5.2.](#) Type-P-one-way-replicated-packet-rate

This statistic gives the fraction of packets that was duplicated (one or more times) in a stream.

Given a Type-P-one-way-Packet-Duplication-Poisson-Stream, one first removes all values of Type-P-one-way-packet-arrival-count which are

undefined. For the remaining pairs in the stream, one counts the number of pairs with Type-P-one-Way-packet-arrival-count greater than 1. Then one calculates the fraction of packets that meet this criterion as a fraction of the total. (In other words: (number of duplicated packets)/(number of packets sent and not lost).).

The number can be expressed as a percentage.

Note: this statistic is the equivalent of the Y.1540 RIPR [[Y1540](#)]

[5.3.](#) Examples

Consider a stream of 4 packets, sent as:

(1, 2, 3, 4)

and arriving as:

- o Case 1: (1, 2, 3, 4)
- o Case 2: (1, 1, 2, 2, 3, 3, 4, 4)
- o Case 3: (1, 1, 1, 2, 2, 2, 3, 3, 3, 4, 4, 4)
- o Case 4: (1, 1, 1, 2, 3, 3, 3, 4)

Case 1: No packets are duplicated in a stream and both the Type-P-one-way-packet-duplication-fraction and the type-P-one-way-packet-replicated-packet-rate are 0.

Case 2: Every packet is duplicated once and the Type-P-one-way-packet-duplication-fraction is 100%. The type-P-one-way-replicated-packet-rate is 100% too.

Case 3: Every packet is duplicated twice, so the Type-P-one-way-packet-duplication-fraction is 200%. The type-P-one-way-replicated-packet-rate is still 100%.

Case 4: Half the packets are duplicated twice and the other half are not duplicated. The Type-P-one-way-packet-duplication-fraction is again 100% and this number does not show the difference with case 2. However, the type-P-one-way-packet-replicated-packet-rate is 50% in this case and 100% in case 2.

However, the type-P-one-way-packet-duplication-rate will not show the difference between case 2 and 3. For this, one has to look at the

Type-P-one-way-packet-duplication-fraction.

Finally, note that the order in which the packets arrived, do not affect the results. For example, these variations of case 2:

o Case 2a: (1, 1, 2, 2, 3, 3, 4, 4)

o Case 2b: (1, 2, 3, 4, 1, 2, 3, 4)

o Case 2c: (1, 2, 3, 4, 4, 3, 2, 1)

(as well as any other permutation) all yield the same results for Type-P-one-way-packet-duplication-fraction and the type-P-one-way-replicated-packet-rate.

6. Security Considerations

Conducting Internet measurements raises both security and privacy concerns. This memo does not specify an implementation of the

Uijterwaal

Expires June 12, 2009

[Page 11]

Internet-Draft

Packet Duplication Metric

December 2008

metrics, so it does not directly affect the security of the Internet nor of applications which run on the Internet. However, implementations of these metrics must be mindful of security and privacy concerns.

There are two types of security concerns: potential harm caused by the measurements, and potential harm to the measurements. The measurements could cause harm because they are active, and inject packets into the network. The measurement parameters MUST be carefully selected so that the measurements inject trivial amounts of additional traffic into the networks they measure. If they inject "too much" traffic, they can skew the results of the measurement, and in extreme cases cause congestion and denial of service.

The measurements themselves could be harmed by routers giving measurement traffic a different priority than "normal" traffic, or by an attacker injecting artificial measurement traffic. If routers can recognize measurement traffic and treat it separately, the measurements will not reflect actual user traffic. If an attacker injects artificial traffic that is accepted as legitimate, the loss rate will be artificially lowered. Therefore, the measurement methodologies SHOULD include appropriate techniques to reduce the probability measurement traffic can be distinguished from "normal" traffic. Authentication techniques, such as digital signatures, may be used where appropriate to guard against injected traffic attacks.

The privacy concerns of network measurement are limited by the active measurements described in this memo. Unlike passive measurements, there can be no release of existing user data.

7. IANA Considerations

IANA is asked to add this metrics to the IANA IP Performance Metrics (IPPM) Metrics Registry, see [[RFC4148](#)]. This section can be removed after this has been done and upon publication as a RFC.

8. Acknowledgements

The idea to write this draft came up in a meeting with Al Morton, Stanislav Shalunov, Emile Stephan and the author, on the IPPM reporting draft.

This document relies heavily on [[RFC2680](#)] and the author likes to thank the authors of that document for writing it.

Finally, thanks are due to Lars Eggert, Al Morton, Martin Swamy and

Matt Zekauskas for their comments.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", [RFC 2680](#), September 1999.

9.2. Informative References

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", [RFC 2330](#), May 1998.

- [RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", [RFC 3432](#), November 2002.
- [RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", [BCP 108](#), [RFC 4148](#), August 2005.
- [Y1540] "Y.1540 ITU-T Recommendation Y.1540 (2007), Internet protocol data communication service IP packet transfer and availability performance parameters.", 2007.

Author's Address

Henk Uijterwaal
RIPE NCC
Singel 258
1016 AB Amsterdam
The Netherlands

Phone: +31 20 535 4444
Email: henk@ripe.net

Uijterwaal

Expires June 12, 2009

[Page 13]

Internet-Draft

Packet Duplication Metric

December 2008

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.