Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 06, 2014

P. Fan
L. Huang
China Mobile
M. Chen
Huawei Technologies
N. Kumar
Cisco Systems
July 05, 2013

IP Packet Loss Rate Measurement Testing and Problem Statement
draft-fan-opsawg-packet-loss-01

Abstract

This document describes common methods for measuring packet loss rate and their effectiveness. Issues encountered when using the methods and necessary considerations are also discussed and recommended.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on January 06, 2014.

Copyright Notice

Table of Contents

1. Introduction

IP packet loss rate is one of the important metrics that are frequently used to measure IP performance of a data path or link. A general framework of IP performance metrics is provided in [RFC2330], including fundamental concepts definition and issues related to defining sound metrics and methodologies. [RFC2680] and [RFC6673] further define metrics for one-way and round-trip packet loss.

In practical network operation, a number of methods are used by network engineers to calculate packet loss rate, and one of the common ways is to use ping. By checking ping statistics, people expect to get the idea of traffic transmission condition on the link. This document gives an overview of the frequently used methods for

measuring IP packet loss rate, and describes a test on packet loss rate measurement with multiple methods using routers from different vendors. Issues that should be taken into consideration during the measurement using different methods are discussed. Causes analysis and processing mechanisms of routers are also covered. It is expected that an operable measurement scheme with consistent testing results and equal treatment of network components can be reached.

2. Methods for Packet Loss Rate Measurement

This section describes common methods for measuring packet loss rate.

2.1. Active Approach

2.1.1. Ping

Ping (ICMP echo request/reply) is a useful tool to examine the connectivity and performance of a path between two nodes in the network. The source node generates echo request packets with configured size, interval, count and other settings, and the destination node sends back an echo reply packet once it receives a request. Then we count the packets sent out and received and get the round-trip packet loss rate on the link between source and destination. This approach is clear and convenient, and is frequently used by engineers when packet loss rate is needed.

In practical network operation, the ping testing can be initiated manually and directly on the node by engineers, for example through the command line interface (CLI) of a router, or activated indirectly by instructions, for example through SNMP messages sent from network management system.

No matter through CLI or SNMP, ping testing can be conducted directly on the endpoint devices of the link to be tested, or other nodes as long as the request/reply packets pass through the link. Those nodes are often referred to as probes, which can be a router or a PC server, directly connected or indirectly reachable to the endpoints. Usually the probes and paths to the endpoints are not supposed to be congested to avoid affecting the ping testing result.

2.1.2. OWAMP and TWAMP

The One-way Active Measurement Protocol (OWAMP, [RFC4656]) and Two-Way Active Measurement Protocol (TWAMP, [RFC5357]) are defined by the IP Performance Metrics (IPPM) working group. They provide a method and protocol for measuring delay and packet loss of IP flows, and are designed for wide scale deployment in the network to provide ubiquitous performance data. Both OWAMP and TWAMP use control

protocol and test protocol. The control protocol is used to negotiate test session between test endpoints, start and stop the test, and fetch the test result for OWAMP. The test protocol runs over UDP and conducts the test.

OWAMP can be used to perform one-way packet loss measurement, and requires synchronized time defined by GPS. The test results are collected at the receiving endpoints and returned using the control protocol. TWAMP is more simplified, and used for two-way packet loss measurement. The opposite endpoint is regarded as a reflector, and the test results are collected at the sender.

2.1.3. Proprietary Tools

There are some other proprietary performance measurement tools incorporating embedded and external probes. The probes generate and inject extra packets into the network to mimic the service flows that are intended to be tested. The performance of the target service flows can be evaluated by measuring the performance of the injected packets. Compared with Ping, these proprietary tools normally support more services, which include not only ICMP, but TCP, UDP, HTTP, etc.

The embedded proprietary tools have been widely implemented by routers to provide automatic detection of IP performance. Examples of this kind of tools include RPM (Juniper), IPSLA (Cisco), NQA (Huawei/H3C), SAA (ALU), etc. By necessary configurations on the router, the embedded tools support multi-service testing of multiple queues on an interface. Packet loss rate can be measured with ICMP ping function of the tool. Routers send out ICMP packets automatically according to the configured parameters, so the embedded tool is working in a similar way as ping method described above.

2.2. Passive Approach

2.2.1. Interface Statistics Report

Forwarding devices maintain statistics report of every interface. The report shows the detailed status of the interface as well as traffic information, including inbound and outbound speed and packet count. For a typical router, traffic statistics show number of packets transmitted and discarded by an interface, and even on the basis of QoS queue, so the entire packet loss rate of a link or packet loss rates regarding different queues can be calculated. Traffic data on the report can be displayed through CLI or obtained using SNMP which allows automatic packet loss sampling.

2.2.2. Coloring Based Performance Measurement

The concept of coloring based performance measurement is introduced in [I-D.tempia-opsawg-p3m], and [I-D.chen-coloring-based-ipfpm-framework] defines a framework for coloring based IP Flow Performance Measurement (IPFPM). By periodically setting/changing one or more bits of the IP header of the packets that belong to an IP flow to "color" the packets into different colors, the IP flow is split into different consecutive blocks. Packets in the same block have the same color and packets in consecutive blocks have different colors. This method gives a way to a measurement node to count and calculate, without inserting any extra auxiliary OAM packets, packet loss based on each color block. Since the measurement is based on the real traffic data, the measurement results will reflect the real performance of the tested flow.

3. Test on Packet Loss Rate Measurement

This section describes test result on packet loss rate measurement using different methods. Test equipment covers routers from several vendors. Results show the diverse outcome of the methods used, and the diverse responding mechanism of routers.

3.1. Basic Test Information

The basic topology of testing can be depicted as follows.

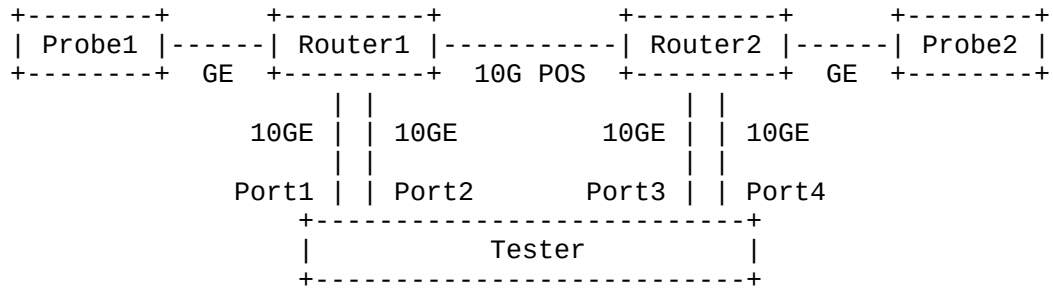

Figure 1: Basic topology for packet loss rate test

Two routers are connected by a 10G POS link, and each router is connected to the tester by two 10GE links. The tester generates unidirectional/bidirectional traffic between port 1 and port 3, and between port 2 and port 4, with frame length of 400 bytes. The total volume of traffic injected into a router by the tester is more than 10G, leading to congestion when the traffic passes through the 10G POS link between the two routers. Routers and probes generate ping packets for testing, with frame length of 400 and DSCP field of 0.

We tested routers from 3 vendors, indicated as A, B, and C in the following parts of discussion. The tester generated different levels of congestion, and we tested packet loss rates on the 10G POS interconnection link on those congestion levels with CLI, SNMP, and interface statistics report.

3.2. Ping with CLI vs. SNMP

Some routing boxes by default treat ping packets generated with CLI and SNMP in different ways. The following is a test on this issue.

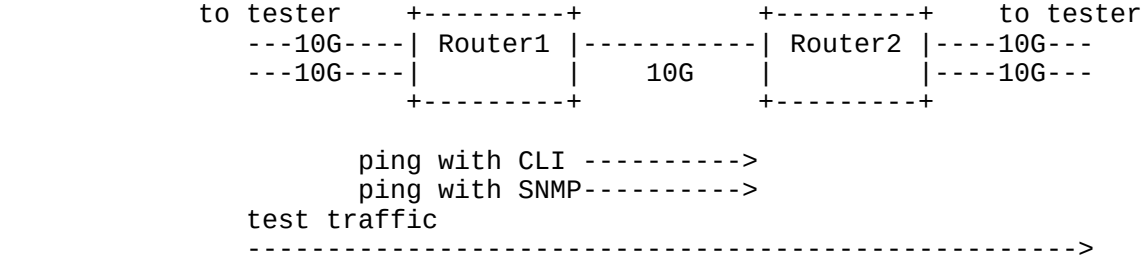

The tester generates test traffic at 20 Gbps, and sends the traffic into a router of vendor A. The traffic goes through the 10G interconnecting link and past the router of vendor B on the other end. We use ping with CLI and SNMP on router A to test packet loss rate on the interconnecting link. The DSCP fields of test traffic and ping packets are all left to be 0..

By default, router A forwards the test traffic with the basic priority, like BE class. The ping packets with CLI are also treated as of best effort class, but ping packets with SNMP are given a higher priority, some class like network control. So the two kinds of ping are actually testing packet loss of streams in different classes. The test result verifies the issue. Ping with SNMP shows no packet loss, and ping with CLI shows a packet loss rate of around 50%.

The forwarding class of ICMP packets can be configured on router A. In the following tests we put all traffic in the same basic class.

3.3. Ping Behaviors of Routers

We considered the following test cases (TCs) when investigating packet loss rate with ping on the link between two different routers.

TC 1: Router sends ICMP echo request packets with SNMP instruction to the peering router.

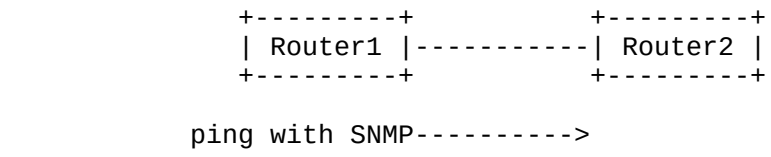

TC 2: Router sends ICMP echo request packets with CLI to the peering router.

TC 3: Router sends ICMP echo request packets with SNMP instruction to the probe behind the peering router.

TC 4: Router sends ICMP echo request packets with CLI to the probe behind the peering router.

TC 5: Probe behind router sends ICMP echo request packets to the probe behind the peering router.

The link between the two routers is injected bidirectional or unidirectional test traffic to cause congestion. The packet loss rate of test traffic is calculated with the Rx and Tx rate on the tester. We use router A, B and C in pairs and get the ICMP packet loss rate in each test case. The comparison of the packet loss rate of ICMP and test traffic shows diverse behaviors of ping process on routers. The following tables show the test results

Pkt loss rate of test traffic		ICMP pkt loss rate (echo req drct: A->B)					ICMP pkt loss rate (echo req drct: B->A)				
A->B	B->A	TC1	TC2	TC3	TC4	TC5	TC1	TC2	TC3	TC4	TC5
48.60%	48.60%	54%	56%	80%	76%	73%	54%	54%	58%	58%	77%
28%	28%	27%	30%	61%	58%	47%	32%	32%	27%	21%	53%
7.60%	7.60%	9%	12%	15%	18%	21%	13%	15%	11%	11%	21%
48.60%	No traffic	54%	56%	57%	54%	54%	62%	56%	54%	48%	56%
28%	No traffic	31%	33%	32%	33%	33%	36%	34%	34%	35%	35%
7.60%	No traffic	14%	13%	12%	9%	14%	14%	13%	11%	12%	14%
No traffic	48.60%	1%	0%	54%	50%	47%	1%	1%	0%	1%	50%
No traffic	28%	0%	0%	26%	31%	28%	0%	0%	0%	0%	28%
No traffic	7.60%	0%	0%	10%	9%	9%	0%	0%	0%	0%	8%

Table 1: Test result when interconnecting router A and router B

Pkt loss rate of test traffic		ICMP pkt loss rate (echo req drct: A->B)					ICMP pkt loss rate (echo req drct: C->A)				
A->C	C->A	TC1	TC2	TC3	TC4	TC5	TC1	TC2	TC3	TC4	TC5
48.70%	44.70%	58%	54%	57%	58%	53%	57%	55%	48%	57%	56%
28%	22.40%	38%	31%	37%	33%	35%	30%	33%	33%	37%	35%
7.70%	7.30%	14%	13%	13%	13%	12%	16%	13%	15%	16%	14%
48.80%	No traffic	50%	54%	51%	53%	55%	54%	56%	55%	59%	57%
28%	No traffic	27%	29%	32%	32%	33%	35%	30%	35%	33%	33%
7.60%	No traffic	11%	10%	15%	15%	13%	11%	11%	15%	15%	13%
No traffic	44.50%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
No traffic	22.60%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
No traffic	7.74%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%

Table 2: Test result when interconnecting router A and router C

Pkt loss rate of	ICMP pkt loss rate	ICMP pkt loss rate

test traffic		(echo req drct: C->B)			(echo req drct: B->C)		
C->B	B->C	TC1	TC2	TC5	TC1	TC2	TC5
48.76%	44.69%	1%	0%	54%	0%	1%	50%
28.04%	22.29%	0%	0%	40%	0%	1%	29%
7.62%	7.62%	0%	0%	11%	0%	0%	8%
48.69%	No traffic	0%	0%	0%	0%	0%	0%
28.03%	No traffic	0%	0%	0%	0%	0%	0%
7.62%	No traffic	0%	0%	0%	0%	0%	0%
No traffic	44.50%	1%	0%	51%	0%	1%	51%
No traffic	22.29%	0%	0%	29%	0%	0%	29%
No traffic	7.74%	0%	0%	9%	0%	0%	10%

Table 3: Test result when interconnecting router C and router B

The behaviors of the three vendors' routers are summarized here, and we leave the discussion on reasons for the behaviors to the next section.

Router A: Ping by router A with SNMP, CLI and by the probe behind router A lead to similar usable results. However, all the methods encounter larger errors when the test traffic is less congested.

Router B: Ping by router B with SNMP and CLI will not report correctly the packet loss rate of test traffic. Ping by the probe behind router B gives usable result of packet loss rate, but also with certain errors.

Router C: Ping by router C with SNMP, CLI and by the probe behind router C will not report correctly the packet loss rate of test traffic.

We can further highlight the outcomes when testing the packet loss rate on the interconnection link between each pair of routers.

Router A - router B: If one wants to get relatively accurate value of packet loss rate in all congestion scenarios, he is advised to use ping between probes (test case 5), or have A generate ping to the probe behind B.

Router A - router C: All the test methods will only reflect the outbound packet loss rate of A.

Router B - router C: Packet loss rate is difficult to measure with this combination- only using ping between probes (test case 5) can reflect the outbound packet loss rate of B.

3.4. Statistics Report of Routers

We also checked the interface statistics reports given with CLI on the 3 routers, and we confirmed that the outbound packet loss rate of an interface obtained from the statistics report was in accordance with the actual packet loss rate of test traffic. The following table shows the test result.

Router	Outbound pkt loss rate of test traffic	Outbound pkt loss rate shown on statistics report
A	48.52%	48.52%
B	48.52%	48.52%
C	44.60%	44.60%

Table 4: Test result when referring to the statistics report on routers

4. Measurement Issues

This section describes issues encountered when measuring the packet loss rate of a link using different testing methods.

4.1. Issues with Ping

Routers from every vendor have their unique processing procedure when sending and receiving ICMP packets, thus resulting in diverse ping packet loss rates, as described in the section above. Errors exist using the ping method, and in some cases ping no longer reflects the actual packet loss rate correctly. Relevant issues that have to be taken into account include:

Forwarding class: When sending ping packets locally, routers are likely to put the packets into a certain QoS queue/class although the DSCP field of ICMP packets is kept zero. QoS queue of ping may be different than that of the traffic to be measured, and even ping packets sent by CLI commands and SNMP are in different queues by default. Usually forwarding class can be adjusted by CLI or SNMP commands.

Inner priority: For some routers, although ping traffic and service traffic will not be treated differently by QoS, packets sent out by the router itself, for example ping packets, are put into an inner high priority while other forwarding service traffic into low priority. These kinds of inner priority are valid within the interior of routers and do not rewrite the packets. One of the

purposes of using the priorities is to get the protocol packets (ping included) processed in prior. These priorities are set by vendor and may not be able to adjust, so in this case ping will not give the correct packet loss rate as ping packets are not processed and discarded together with service traffic.

Ingress line card: If the ping testing is conducted on a probe which is connected or IP reachable to the router, then the ping packets will be treated by the router as forwarding traffic, eliminating the queue and priority issues. However, the location of interfaces through which ingress traffic is received matters when using some types of routers. In this case, the router employs a polling schedule which allows traffic from different line cards or modules to get forwarding chance. For a card with small volume of traffic, the chance will be little but not none. So if ping packets come through a card different from the high-volume service traffic, the packets would probably get enough forwarding resources as ping traffic itself requires little bandwidth. As a result, ping will suffer little from congestion and shows disaccord in packet loss rate.

Internal rate limitation: Routers normally have rate limitation towards CPU, which is considered a kind of protection to the control plane of routers. So if a packet is sent to CPU for processing rather than line card ASIC (e.g. in many routers, an ICMP echo reply packet received in response to an earlier echo request packet sent by the router will be sent to the CPU), it might be influenced by the rate limiter. Typical rate limitation of ICMP packets would be 1000 pps, though the value is highly dependent on vendor implementation and can be configured. In practical deployment, if there is a large number of ICMP packets sending to a router, the ping test packets may be dropped, causing test errors. This problem did not arise in our test in section 3 as the ICMP traffic is rather small.

4.2. Issues with OWAMP and TWAMP

OWAMP and TWAMP fall into the category of active measurement, so the general issues of active measurement apply to them. When using the two methods, one is advised to make sure that the measurement traffic will have the same drop probability as non-measurement traffic. However, it is usually difficult to guarantee this, as too many factors effect the behavior of traffic.

4.3. Issues with Proprietary Tools

Since the proprietary tools are implemented by vendors independently, interoperability is one of the major issues when using the tools,

especially for one-way measurement. Besides, these tools also share the common issues of active measurement. The accuracy of results depends on the rate, numbers and interval of the injected packets. It also needs to guarantee that the injected packets follows the same path as the tested packets, otherwise the results cannot reflect the real performance.

Although these tools provide automatic testing method, the basic principle is still to ping from the router itself. So it is believed toolset method will experience the same issues about class and priority as local ping from router does. However, we did not test diagnosis toolsets, and the discussion is left to be further continued.

4.4. Issues with Interface Statistics Report

Interface statistic is the most direct and accurate way to get performance of an interface. Packet loss rate calculated from traffic statistics is in accordance with the expected value. By referring to statistics collected from the endpoint routers, bidirectional packet loss rate can easily be obtained.

However, this approach requires access to routers, while in some scenarios it is difficult to do that. For example, if we would like to know the inbound packet loss rate of the interconnection link to another service operator, we may have to rely on statistics provided by the peering router. Normally, this information is not easily shared by interworking operators.

4.5. Issues with Coloring Based Performance Measurement

The challenge for coloring based performance measurement is that there are not so many bits in the IP header that can be used for IP packet coloring. Operators have to carefully think of the color bits selection to make sure that the setting and changing of the color bits will not affect the normal packet forwarding and process.

5. Considerations and Recommendations

We summarize the above analysis here and come to the following considerations:

a. The ping method to measure packet loss rate is easy to be influenced by the diverse processing mechanism of ICMP packet within routers. If this method is to be used on a router, one is advised to make sure that the ICMP packets experience the same forwarding and discarding courses as the service traffic (of which the packet loss rate is to be measured) does, otherwise the

measurement will not make sense. When measuring with ping, the following points are also worth reminding:

* Packet loss rate given by measurement with ping is a value related to a certain forwarding class in which the ICMP packets are forwarded. So it is not a scientific way to say what the packet loss rate is on a link if traffic is transmitted in more than one class on the link.

* Measurement with ping is enough if one only wants to get a general, qualitative picture of packet loss. But if one is to measure precisely and quantitatively, possible errors (sometimes very large errors) should be taken into account.

* If configured in the right way on router, ping with CLI and SNMP lead to similar results.

b. It is more likely to get good results if a probe is used to perform ping measurement (though not 100% guaranteed), but following issues also need to be considered.

* If the probe is directly connected to a router, then a router port is occupied. This will be a problem for routers with limited or expensive port resources, as the probing traffic is usually extremely small.

* If the probe is more than one hop away from a router, load of the path to the router is supposed to be under the congestion level.

c. Interface statistics report gives us the most accurate value of pack loss rate, and the value is irrelevant to router platforms. From the report we can find numbers of packets being received, transmitted, and discarded in different classes within a period of time, thus we get packet loss rate. Actually this is indeed how packet loss rate is defined.

* Referring to report requires access to routers, which may be easier if routers are within a single administrative area. However it gets annoying if more routers are evolved, for instance measurement on a long path with a number of routers.

* Router interface report only gives the outbound packet loss rate. If we want to see if traffic in the other direction is congested, we'll have to check the upstream routers in that direction. This will be difficult on certain links, say, interconnection link to another provider.

6. Security Considerations

TBD.

7. IANA Considerations

This memo includes no request to IANA.

8. Acknowledgements

The authors would like to thank Brian Trammell for the kind comments.

9. References

9.1. Normative References

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

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

[RFC4656] Shalunov, S. and B. Teitelbaum, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.

[RFC5357] Hedayat, K. and R. Krzanowski, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.

[RFC6673] Morton, A., "Round-Trip Packet Loss Metrics", RFC 6673, August 2012.

[RFC792] Postel, J., "Internet Control Message Protocol", RFC 792, September 1981.

9.2. Informative References

[I-D.chen-coloring-based-ipfpm-framework]
Chen, M., Liu, H., and Y. Yin, "Coloring based IP Flow Performance Measurement Framework", draft-chen-coloring-based-ipfpm-framework-01 (Work in Progress), February 2013.

[I-D.tempia-opsawg-p3m]
Capello, A., Cociglio, M., Castaldelli, L., and A. Bonda, "A packet based method for passive performance monitoring", draft-tempia-opsawg-p3m-03 (Work in Progress), February 2013.

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

Authors' Addresses

Peng Fan
China Mobile
32 Xuanwumen West Street, Xicheng District
Beijing 100053
P.R. China

Email: fanpeng@chinamobile.com

Lu Huang
China Mobile
32 Xuanwumen West Street, Xicheng District
Beijing 100053
P.R. China

Email: huanglu@chinamobile.com

Mach(Guoyi) Chen
Huawei Technologies

Email: mach.chen@huawei.com

Nagendra Kumar
Cisco Systems

Email: naikumar@cisco.com