

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 09, 2014

M. Bagnulo
UC3M
T. Burbridge
BT
S. Crawford
SamKnows
P. Eardley
BT
A. Morton
AT&T Labs
July 08, 2013

A Reference Path and Measurement Points for LMAP
draft-ietf-ippm-lmap-path-00

Abstract

This document defines a reference path for Large-scale Measurement of Broadband Access Performance (LMAP) and measurement points for commonly used performance metrics.

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 09, 2014.

Copyright Notice

Copyright (c) 2013 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction 2

2. Purpose and Scope 3

3. Terms and Definitions 3

 3.1. Reference Path 3

 3.2. Subscriber 3

 3.3. Dedicated Component (Links or Nodes) 4

 3.4. Shared Component (Links or Nodes) 4

 3.5. Resource Transition Point 4

4. Reference Path 4

5. Measurement Points 6

6. Translation Between Ref. Path and Tech. X 7

7. Security considerations 9

8. IANA Considerations 9

9. Acknowledgements 9

10. References 9

 10.1. Normative References 9

 10.2. Informative References 10

Authors' Addresses 10

1. Introduction

This document defines a reference path for Large-scale Measurement of Broadband Access Performance (LMAP). The series of IP Performance Metrics (IPPM) RFCs have developed terms that are generally useful for path description (section 5 of [RFC2330]). There are a limited number of additional terms needing definition here, and they will be defined in this memo.

The reference path is usually needed when attempting to communicate precisely about the components that comprise the path, often in terms of their number (hops) and geographic location. This memo takes the path definition further, by establishing a set of measurement points along the path and ascribing a unique designation to each point. This topic has been previously developed in section 5.1 of [RFC3432], and as part of the updated framework for composition and aggregation, section 4 of [RFC5835] (which may also figure in the LMAP work effort). Section 4.1 of [RFC5835] defines the term "measurement point".

Measurement points and the paths they cover are often described in general terms, like "end-to-end", "user-to-user", or "access". These terms are insufficient for scientific method: What is an end? Where is a user located? Is the home network included?

The motivation for this memo is to provide an unambiguous framework to describe measurement coverage, or scope of the reference path. This is an essential part of the metadata to describe measurement results. Measurements conducted over different path scopes are not a valid basis for performance comparisons.

2. Purpose and Scope

The scope of this memo is to define a reference path for LMAP activities with sufficient level of detail to determine the location of different measurement points without ambiguity.

The bridge between the reference path and specific network technologies (with differing underlying architectures) is within the scope of this effort. Both wired and wireless technologies are in-scope.

The purpose is to create an efficient way to describe the location of the measurement point(s) used to conduct a particular measurement so that the measurement result will adequately described in this regard. This should serve many measurement uses, including diagnostic (where the same metric may be measured over many different path scopes) and comparative (where the same metric may be measured on different network infrastructures).

3. Terms and Definitions

This section defines key terms and concepts for the purposes of this memo.

3.1. Reference Path

A reference path is a serial combination of routers, switches, links, radios, and processing elements that comprise all the network elements traversed by each packet between the source and destination hosts. The reference path is intended to be equally applicable to all networking technologies, therefore the components are generically defined, but their functions should have a clear counterpart or be obviously omitted in any network technology.

3.2. Subscriber

An entity possessing one or more hosts participating in an Internet access service.

3.3. Dedicated Component (Links or Nodes)

All resources of a Dedicated component (typically a link or node on the Reference Path) are allocated to serving the traffic of an individual Subscriber. Resources include transmission time-slots, queue space, processing for encapsulation and address/port translation, and others. A Dedicated component can affect the performance of the Reference Path, or the performance of any sub-path where the component is involved.

3.4. Shared Component (Links or Nodes)

A component on the Reference Path is designated a Shared component when the traffic associated with multiple Subscribers is served by common resources.

3.5. Resource Transition Point

A point between Dedicated and Shared components on a Reference Path that may be a point of significance, and is identified as a transition between two types of resources.

4. Reference Path

This section defines a reference path for Internet Access.

```
Subsc. -- Private -- Private -- Access -- Intra IP -- GRA -- Transit
device   Net #1   Net #2   Demarc.   Access   GW   GRA GW
```

```
... Transit -- GRA -- Service -- Private -- Private -- Destination
   GRA GW   GW   Demarc.   Net #n   Net #n+1   Host
```

GRA = Globally Routable Address, GW = Gateway

The following are descriptions of reference path components that may not be clear from their name alone.

- o Subsc. (Subscriber) device - This is a host that normally originates and terminates communications conducted over the IP packet transfer service.

- o Private Net #x - This is a network of devices owned and operated by the Internet Access Service Subscriber. In some configurations, one or more private networks and the device that provides the Access Service Demarcation point are collapsed in a single device (and ownership may shift to the service provider), and this should be noted as part of the path description.
- o Access (Service) Demarcation point - this varies by technology but is usually defined as the Ethernet interface on a residential gateway or modem where the scope of access packet transfer service begins and ends. In the case of a WiFi Service, this would be an Air Interface within the intended service boundary (e.g., walls of the coffee shop). The Demarcation point may be within an integrated endpoint using an Air Interface (e.g., LTE UE). Ownership may not affect the demarcation point; a Subscriber may own all equipment on their premises, but it is likely that the service provider will certify such equipment for connection to their access network, or a third-party will certify standards compliance.
- o Intra IP Access - This is the first point in the access architecture beyond the Access Service Demarc. where a globally routable IP address is exposed and used for routing. In architectures that use tunneling, this point may be equivalent to the GRA GW. This point could also collapse to the device providing the Access Service Demarc., in principle. Only one Intra IP Access point is shown, but they can be identified in any access or transit network.
- o GRA GW - the point of interconnection between the access administrative domain and the rest of the Internet, where routing will depend on the GRAs in the IP header.
- o Transit GRA GW - Networks that intervene between the Subscriber's Access network and the Destination Host's network are designated "transit" and involve two GRA GW.

Use of multiple IP address families in the measurement path must be noted, as the conversions between IPv4 and IPv6 certainly influence the visibility of a GRA for each family.

In the case that a private address space is used throughout an access architecture, then the Access Service Demarc. and the Intra IP Access points must use the same address space and be separated by the shared and dedicated access link infrastructure, such that a test between these points produces a useful assessment of access performance.

5. Measurement Points

A key aspect of measurement points, beyond the definition in section 4.1 of [RFC5835], is that the innermost IP header and higher layer information must be accessible through some means. This is essential to measure IP metrics. There may be tunnels and/or other layers which encapsulate the innermost IP header, even adding another IP header of their own.

In general, measurement points cannot always be located exactly where desired. However, the definition in [RFC5835] and the discussion in section 5.1 of [RFC3432] indicate that allowances can be made: for example, deterministic errors that can be quantified are ideal.

The Figure below illustrates the assignment of measurement points to selected components of the reference path.

```

Subsc. -- Private -- Private -- Access -- Intra IP -- GRA -- Transit
device   Net #1   Net #2   Demarc.   Access   GW   GRA GW
mp000                                mp100   mp150   mp190   mp200

... Transit -- GRA -- Service -- Private -- Private -- Destination
  GRA GW   GW   Demarc.   Net #n   Net #n+1   Host
  mpX90   mp890 mp800                                mp900

```

GRA = Globally Routable Address, GW = Gateway

The numbering for measurement points (mpNNN) allows for considerable local use of unallocated numbers.

Notes:

- o Some use the terminology "on-net" and "off-net" when referring to Internet Service Provider (ISP) measurement coverage. With respect to the reference path, tests between mp100 and mp190 are "on-net".
- o Widely deployed broadband access measurements have used pass-through devices[SK] (at the subscriber's location) directly connected to the service demarcation point: this would be located at mp100.
- o The networking technology used at all measurement points must be indicated, especially the interface standard and configured speed.

- o If it can be shown that a link connecting to a measurement point has reliably deterministic or negligible performance, then the remote end of the connecting link is an equivalent point for some methods of measurement (To Be Specified Elsewhere). In any case, the presence of such a link must be reported.
- o Many access network architectures have a traffic aggregation point (e.g., CMTS or DSLAM) between mp100 and mp150. We designate this point mp120, but it won't currently fit in the figure.
- o A Carrier Grade NAT (CGN) deployed in the Subscriber's access network would be positioned between mp100 and mp190, and the egress side of the CGN will typically be designated mp150.
- o In the case that a private address space is used in an access architecture, then mp100 may need to use the same address space as its remote measurement point counterpart, so that a test between these points produces a useful assessment of network performance. Tests between mp000 and mp100 could use private address space, and when the egress side of a CGN is at mp150, then the private address side of the CGN could be designated mp149 for tests with mp100.
- o Measurement points at Transit GRA GWs are numbered mpX00 and mpX90, where X is the lowest positive integer not already used in the path.

6. Translation Between Ref. Path and Tech. X

This section and those that follow are intended to provide a more exact mapping between particular network technologies and the reference path.

We provide an example for 3G Cellular access below.

Subscriber	--	Private	--	Access Srv	-----	GRA	---	Transit	...
device		Net #1		Demarc.		GW		GRA GW	
mp000				mp100		mp190		mp200	




GRA = Globally Routable Address, GW = Gateway, UE = User Equipment,
 RAN = Radio Access Network, GGSN = Gateway GPRS Support Node.

We next provide a few examples of DSL access. Consider first the case where:

- o The Customer Premises Equipment (CPE) is a NAT device that is configured with a public IP address.
- o The CPE is a home router that has also an incorporated a WiFi access point and this is the only networking device in the home network, all endpoints attach directly to the CPE though the WiFi access.

We believe this is a fairly common configuration in some parts of the world and fairly simple as well.

This case would map into the defined reference measurement points as follows:




GRA = Globally Routable Address, GW = Gateway

Consider next the case where:

- o The Customer Premises Equipment (CPE) is a NAT device that is configured with a private IP address.
- o There is a Carrier Grade NAT (CGN) located deep into the Access ISP network.
- o The CPE is a home router that has also an incorporated a WiFi access point and this is the only networking device in the home network, all endpoints attach directly to the CPE though the WiFi access.

We believe is becoming a fairly common configuration in some parts of the world.

This case would map into the defined reference measurement points as follows:



```

|--UE--|-----CPE/NAT-----|-----|-CGN-|-----|
      |---Access Network--|

```

GRA = Globally Routable Address, GW = Gateway

7. Security considerations

Specification of a Reference Path and identification of measurement points on the path represent agreements among interested parties, and they present no threat to the readers of this memo or to the Internet itself.

8. IANA Considerations

TBD

9. Acknowledgements

Thanks to Matt Mathis for review and comments.

10. References

10.1. Normative 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.
- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.
- [RFC6673] Morton, A., "Round-Trip Packet Loss Metrics", RFC 6673, August 2012.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.
- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999.

- [RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.
- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, March 2009.
- [RFC5835] Morton, A. and S. Van den Berghe, "Framework for Metric Composition", RFC 5835, April 2010.

10.2. Informative References

- [RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", BCP 108, RFC 4148, August 2005.
- [RFC6248] Morton, A., "RFC 4148 and the IP Performance Metrics (IPPM) Registry of Metrics Are Obsolete", RFC 6248, April 2011.
- [SK] Crawford, Sam., "Test Methodology White Paper", SamKnows Whitebox Briefing Note
<http://www.samknows.com/broadband/index.php>, July 2011.

Authors' Addresses

Marcelo Bagnulo
Universidad Carlos III de Madrid
Av. Universidad 30
Leganes, Madrid 28911
SPAIN

Phone: 34 91 6249500
Email: marcelo@it.uc3m.es
URI: <http://www.it.uc3m.es>

Trevor Burbridge
British Telecom
Adastral Park, Martlesham Heath
IPswitch
ENGLAND

Email: trevor.burbridge@bt.com

Sam Crawford
SamKnows

Email: sam@samknows.com

Phil Eardley
British Telecom
Adastral Park, Martlesham Heath
IPswitch
ENGLAND

Email: philip.eardley@bt.com

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown, NJ
USA

Email: acmorton@att.com