Internet Engineering Task Force (IETF)
Request for Comments: 7398
Category: Informational
ISSN: 2070-1721

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

February 2015

A Reference Path and Measurement Points for Large-Scale Measurement of Broadband Performance

Abstract

This document defines a reference path for Large-scale Measurement of Broadband Access Performance (LMAP) and measurement points for commonly used performance metrics. Other similar measurement projects may also be able to use the extensions described here for measurement point location. The purpose is to create an efficient way to describe the location of the measurement point(s) used to conduct a particular measurement.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7398.

Table of Contents

1. Introduction

This document defines a reference path for Large-scale Measurement of Broadband Access Performance (LMAP) or similar measurement projects. The series of IP Performance Metrics (IPPM) RFCs have developed terms that are generally useful for path description (see Section 5 of [RFC2330]). There are a limited number of additional terms defined in this memo.

The reference path (see Section 3.1 and Figure 1 of [Y.1541], including the accompanying discussion) is usually needed when attempting to communicate precisely about the components that comprise the path, and is often expressed 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 in Section 4 of [RFC5835]. Section 4.1 of [RFC5835] defines the term "measurement point".

Measurement points and the paths they inhabit are often described in general terms, like "end-to-end", "user-to-user", or "access". These terms alone are insufficient for the scientific method, since we need to clarify issues such as: What is an end? Where is a user located? Is the home network included?

As an illustrative example, consider a measurement agent in an LMAP system. When it reports its measurement results, rather than detailing its IP address and that of its measurement peer, it may prefer to describe the measured path segment abstractly (perhaps for privacy reasons), e.g., ’from a measurement agent at a home gateway to a measurement peer at a DSLAM.’ This memo provides the definition for such abstract ’measurement points’ and, therefore, the portion of ’reference path’ between them.

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. We note that additional measurement context information may be necessary to support a valid comparison of results.

1.1. Requirements Language

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 RFC 2119 [RFC2119].

2. Purpose and Scope

The scope of this memo is to define a reference path for LMAP activities with a sufficient level of detail to determine the location of different measurement points along a path without ambiguity. These conventions are likely to be useful in other measurement projects and to describe the applicable measurement scope for some metrics.

The connection between the reference path and specific network technologies (with differing underlying architectures) is within the scope of this memo, and examples are provided. 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 be adequately described in terms of scope or coverage. This should serve many measurement uses, including:

- diagnostic, where the same metric would be measured on different sub-paths bounded by measurement points (see Section 4.10 of [RFC5835]), for example, to isolate the sub-path contributing the majority of impairment levels observed on a path.

- comparison, where the same metric may be measured on equivalent portions of different network infrastructures, for example, to compare the performance of wired and wireless home network technologies.

3. Terms and Definitions

This section defines key terms and concepts for this memo.

3.1. Reference Path

A reference path is a serial combination of hosts, routers, switches, links, radios, and processing elements that comprise all the network elements traversed by each packet in a flow between the source and destination hosts. A reference path also indicates the various boundaries present, such as administrative boundaries. A reference path is intended to be equally applicable to all IP and link-layer

networking technologies. Therefore, the components are generically defined, but their functions should have a clear counterpart or be obviously omitted in any network architecture.

3.2. Subscriber

A subscriber is an entity (associated with one or more users) that is engaged in a subscription with a service provider. The subscriber is allowed to subscribe and unsubscribe to services and to register a user or a list of users authorized to enjoy these services. [Q.1741]

Both the subscriber and service provider are allowed to set the limits relative to the use that associated users make of subscribed services.

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

This is 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.

3.6. Service Demarcation Point

This is the point where a service managed by the service provider begins (or ends) and varies by technology. For example, this point is usually defined as the Ethernet interface on a residential gateway or modem where the scope of a 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., Long-Term Evolution User Equipment (LTE UE)). Ownership does not necessarily 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 network or that a third-party will certify standards compliance.

3.7. Managed and Unmanaged Sub-paths

Service providers are responsible for the portion of the path they manage. However, most paths involve a sub-path that is beyond the management of the subscriber's service provider. This means that private networks, wireless networks using unlicensed frequencies, and the networks of other service providers are designated as unmanaged sub-paths. The service demarcation point always divides managed and unmanaged sub-paths.

4. Reference Path

This section defines a reference path for Internet communication.

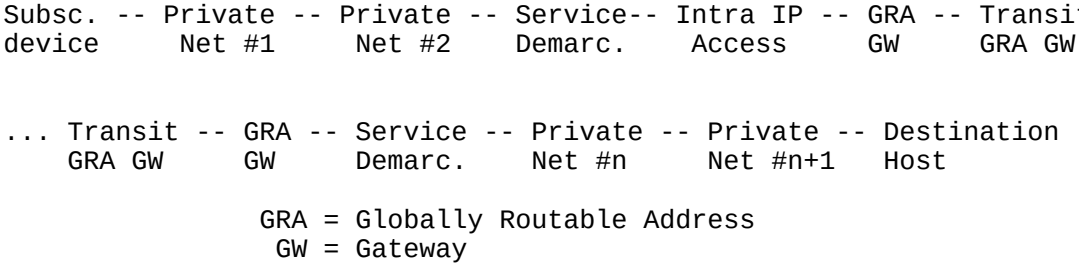

Figure 1: Reference Path

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

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

- Private Net #x - This is a network of devices owned and operated by the Internet service subscriber. In some configurations, one or more private networks and the device that provides the service demarcation point are collapsed in a single device (ownership may shift to the service provider); this should be noted as part of the path description.

- Intra IP Access - This is the first point in the access architecture, beyond the service demarcation, where a globally routable IP address is exposed and used for routing. In architectures that use tunneling, this point may be equivalent to the Globally Routable Address Gateway (GRA GW). This point could also collapse to the device providing the service demarcation, in principle. Only one Intra IP Access point is shown, but they can be identified in any access network.

- GRA GW - This is the point of interconnection between a service provider's administrative domain and the rest of the Internet, where routing will depend on the GRAs in the IP header.

- Transit GRA GW - If one or more networks intervene between the service provider's access networks of the subscriber and of the destination host, then such networks are designated "transit" and are bounded by two transit GRA GWs.

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 Intra IP Access points must use the same address space as the service demarcation point, and the Intra IP Access points must be selected such that a test between these points produces a useful assessment of access performance (e.g., includes both shared and dedicated access link infrastructure).

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 that 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, it is nearly ideal when there are deterministic errors that can be quantified between desired and actual measurement points.

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

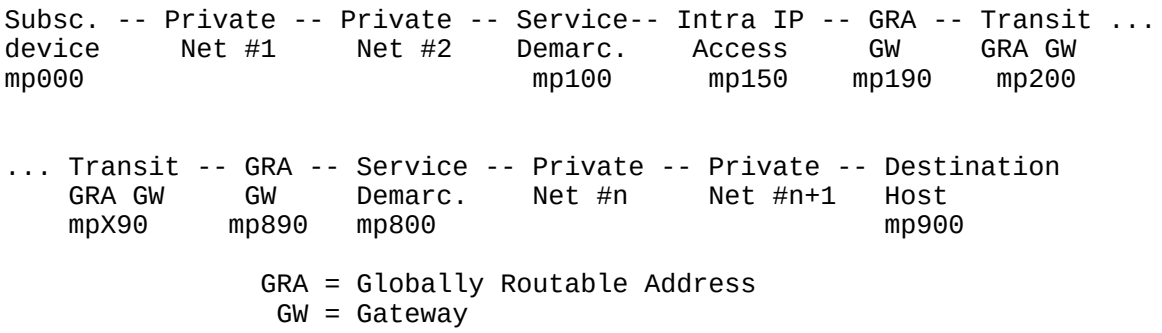

Figure 2: Reference Path with Measurement Point Designations

Each measurement point on a specific reference path MUST be assigned a unique number. To facilitate interpretation of the results, the measuring organization (and whoever it shares results with) MUST have an unambiguous understanding of what path or point was measured. In order to achieve this, a set of numbering recommendations follow.

When communicating the results of measurements, the measuring organization SHOULD supply a diagram similar to Figure 2 (with the technology-specific information in examples that follow) and MUST supply it when additional measurement point numbers have been defined and used (with sufficient detail to identify measurement locations in the path).

Ideally, the consumer of measurement results would know the location of a measurement point on the reference path from the measurement point number alone; the recommendations below provide a way to accomplish this goal. Although the initial numbering may be fully compliant with this system, changing circumstances could, over time, lead to gaps in network numbers or non-monotonic measurement point number assignments along the path. Such circumstances could include growth, consolidation, re-arrangement, and change of ownership of the network. These are examples of reasonable causes for numbering deviations that must be identified on the reference path diagram, as required above.

While the numbering of a measurement point is in the context of a particular path, for simplicity, the measuring organization SHOULD use the same numbering for a device (playing the same role) on all the measurement paths through it. Similarly, whilst the measurement point numbering is in the context of a particular measuring organization, organizations with similar technologies and

architectures are encouraged to coordinate on local numbering and diagrams.

The measurement point numbering system, mpXnn, has two independent parts:

1. The X in mpXnn indicates the network number. The network with the subscriber's device is network 0. The network of a different organization (administrative or ownership domains) SHOULD be assigned a different number. Each successive network number SHOULD be one greater than the previous network's number. Two circumstances make it necessary to designate X=9 in the destination host's network and X=8 for the service provider network at the destination:

   A. The number of transit networks is unknown.

   B. The number of transit networks varies over time.

2. The nn in mpXnn indicates the measurement point and is locally assigned by network X. The following conventions are suggested:

   A. 00 SHOULD be used for a measurement point at the subscriber's device and at the service demarcation point or GW nearest to the subscriber's device for transit networks.

   B. 90 SHOULD be used for a measurement point at the GW of a network (opposite from the subscriber's device or service demarcation).

   C. In most networks, measurement point numbers SHOULD monotonically increase from the point nearest the subscriber's device to the opposite network boundary on the path (but see item D for an exception).

   D. When a destination host is part of the path, 00 SHOULD be used for a measurement point at the destination host and at the destination's service demarcation point. Measurement point numbers SHOULD monotonically increase from the point nearest the destination's host to the opposite network boundary on the path ONLY in these networks. This directional numbering reversal allows consistent 00 designation for end hosts and service demarcation.

   E. 50 MAY be used for an intermediate measurement point of significance, such as a Network Address Translator (NAT).

F. 20 MAY be used for a traffic aggregation point, such as a Digital Subscriber Line Access Multiplexer (DSLAM) within a network.

G. Any other measurement points SHOULD be assigned unused integers between 01 and 99. The assignment SHOULD be stable for at least the duration of a particular measurement study and SHOULD avoid numbers that have been assigned to other locations within network X (unless the assignment is considered sufficiently stale). Subnetworks or domains within a network are useful locations for measurement points.

When supplying a diagram of the reference path and measurement points, the operator of the measurement system MUST indicate the reference path, the numbers (mpXnn) of the measurement points, and the technology-specific definition of any measurement point other than X00 and X90 with sufficient detail to clearly define its location (similar to the technology-specific examples in Section 6 of this document).

If the number of intermediate networks (between the source and destination) is not known or is unstable, then this SHOULD be indicated on the diagram, and results from measurement points within those networks need to be treated with caution.

Notes:

- The terminology "on-net" and "off-net" is sometimes used when referring to the subscriber’s Internet Service Provider (ISP) measurement coverage. With respect to the reference path, tests between mp100 and mp190 are "on-net".

- Widely deployed broadband Internet 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.

- The networking technology must be indicated for the measurement points used, especially the interface standard and configured speed (because the measurement connectivity itself can be a limiting factor for the results).

- If it can be shown that a link connecting to a measurement point has reliably deterministic performance or negligible impairments, then the remote end of the connecting link is an equivalent point for some methods of measurement (although those methods should describe this possibility in detail, it is not in scope to provide such methods here). In any case, the presence of a link and claimed equivalent measurement point must be reported.

- Some access network architectures may have an additional traffic aggregation device between mp100 and mp150. Use of a measurement point at this location would require a local number and diagram.

- A Carrier Grade NAT (CGN) deployed in the service provider's access network would be positioned between mp100 and mp190, and the egress side of the CGN may be designated mp150. mp150 is generally an intermediate measurement point in the same address space as mp190.

- In the case that private address space is used in an access architecture, mp100 may need to use the same address space as its "on-net" measurement point counterpart so that a test between these points produces a useful assessment of network performance. Tests between mp000 and mp100 could use a different private address space, and when the globally routable side of a CGN is at mp150, the private address side of the CGN could be designated mp149 for tests with mp100.

- Measurement points at transit GRA GWs are numbered mpX00 and mpX90, where X is the lowest positive integer not already used in the path. The GW of the first transit network is shown with point mp200 and the last transit network GW with mpX90.

6. Examples of Reference Paths with Various Technologies

This section and those that follow are intended to provide example mappings between particular network technologies and the reference path.

We provide an example for 3G cellular access below.

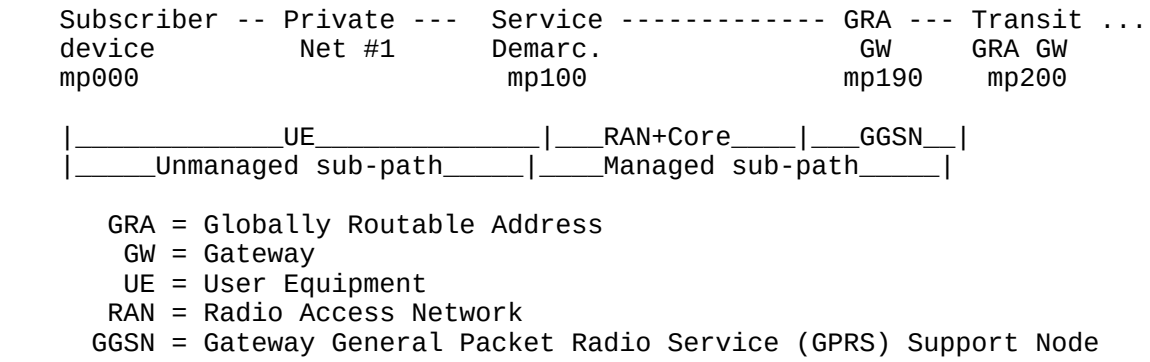

Figure 3: Example of Reference Path with 3G Cellular Access

Next, we provide an example of DSL access. Consider the case where:

- The Customer Premises Equipment (CPE) has a NAT device that is configured with a public IP address.

- The CPE consists of a wired residential GW and modem internally connected (via Private Net #2) to an embedded home router and WiFi access point (Private Net #1). All subscriber devices (UE) attach to the CPE through the WiFi access. mp100 is on the modem side of Private Net #2.

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

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

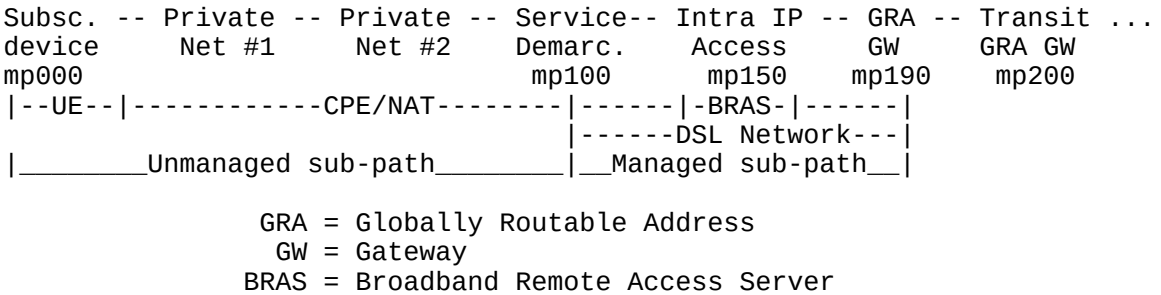

Figure 4: Example of Reference Path with DSL Access

Consider another access network case where:

- The CPE is a NAT device that is configured with a private IP address.
- There is a CGN located deep in the access ISP network.
- 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 through the WiFi access.

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

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

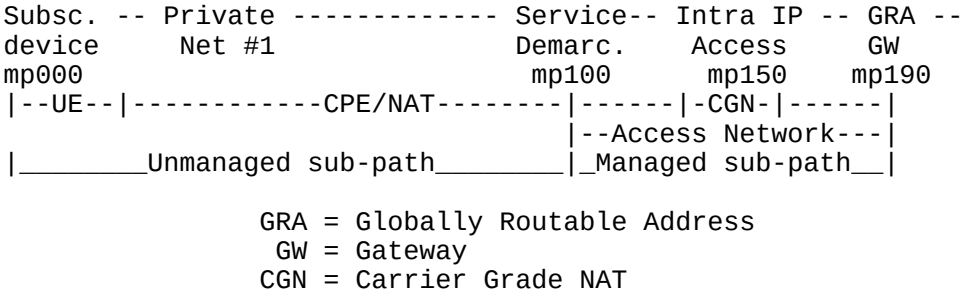

Figure 5: Example of Reference Path with CGN

7. Example of Reference Path with Resource Transition

This section gives an example of shared and dedicated portions with the reference path. This example shows two resource transition points.

Consider the case where:

- The CPE consists of a wired residential GW and modem (Private Net #2) connected to a WiFi access point (Private Net #1). The subscriber device (UE) attaches to the CPE through the WiFi access.
- The WiFi subnetwork (Private Net #1) shares unlicensed radio channel resources with other WiFi access networks (and potentially other sources of interference); thus, this is a shared portion of the path.

- The wired subnetwork (Private Net #2) and a portion of the service provider's network are dedicated resources (for a single subscriber); thus, there is a resource transition point between Private Net #1 and Private Net #2.

- Subscriber traffic shares common resources with other subscribers upon reaching the CGN; thus, there is a resource transition point and further network components are designated as shared resources.

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

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

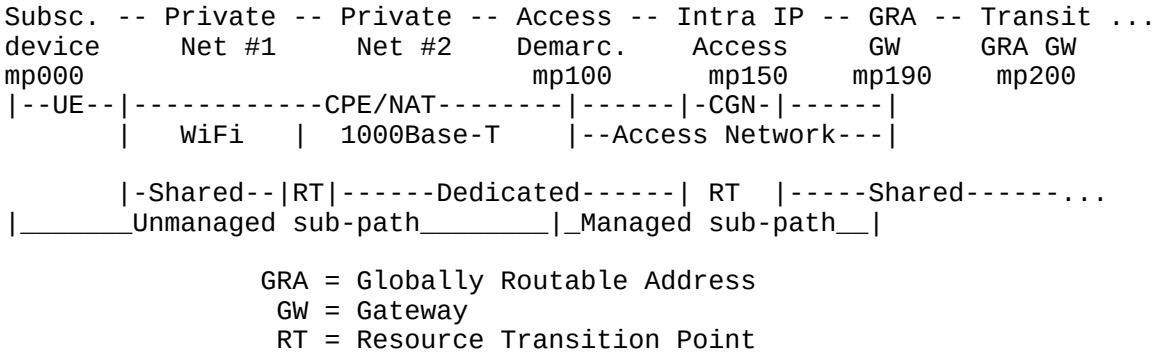

Figure 6: Example of Reference Path with Two Reference Transition Points

8. Security Considerations

Specification of a reference path and identification of measurement points on the path represent agreements among interested parties. They present no threat to the implementors of this memo, or to the Internet resulting from implementation of the guidelines provided here.

Attacks at end hosts or identified measurement points are possible. However, there is no requirement to include IP addresses of hosts or other network devices in a reference path with measurement points that is compliant with this memo. As a result, the path diagrams with measurement point designation numbers do not aid such attacks.

Most network operators' diagrams of reference paths will bear a close resemblance to similar diagrams in relevant standards or other publicly available documents. However, when an operator must include

atypical network details in their diagram, e.g., to explain why a longer latency measurement is expected, then the diagram reveals some topological details and should be marked as confidential and shared with others under a specific agreement.

When considering privacy of those involved in measurement or those whose traffic is measured, there may be sensitive information communicated to recipients of the network diagrams illustrating paths and measurement points described above. We refer the reader to the privacy considerations described in the Large Scale Measurement of Broadband Performance (LMAP) Framework [LMAP-FRAMEWORK], which covers active and passive measurement techniques and supporting material on measurement context. For example, the value of sensitive information can be further diluted by summarizing measurement results over many individuals or areas served by the provider. There is an opportunity enabled by forming anonymity sets described in [RFC6973] based on the reference path and measurement points in this memo. For example, all measurements from the subscriber device can be identified as "mp000", instead of using the IP address or other device information. The same anonymization applies to the Internet service provider, where their Internet gateway would be referred to as "mp190".

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, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998, <http://www.rfc-editor.org/info/rfc2330>.

[RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, November 2002, <http://www.rfc-editor.org/info/rfc3432>.

[RFC5835] Morton, A. and S. Van den Berghe, "Framework for Metric Composition", RFC 5835, April 2010, <http://www.rfc-editor.org/info/rfc5835>.

9.2. Informative References

[LMAP-FRAMEWORK]
Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A framework for large-scale measurement platforms (LMAP)", Work in Progress, draft-ietf-lmap-framework-10, January 2015.

[Q.1741] International Telecommunications Union, "IMT-2000 references to Release 9 of GSM-evolved UMTS core network", ITU-T Recommendation Q.1741.7, November 2011, <http://www.itu.int/rec/T-REC-Q.1741.7/en>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, July 2013, <http://www.rfc-editor.org/info/rfc6973>.

[SK] Crawford, S., "Test Methodology White Paper", SamKnows Whitebox Briefing Note , July 2011, <http://www.samknows.com/broadband/index.php>.

[Y.1541] International Telecommunications Union, "Network performance objectives for IP-based services", ITU-T Recommendation Y.1541, November 2011, <http://www.itu.int/rec/T-REC-Y.1541/en>.

Acknowledgments

Thanks to Matt Mathis, Charles Cook, Dan Romascanu, Lingli Deng, and Spencer Dawkins for review and comments.

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
BT
Adastral Park, Martlesham Heath
Ipswich
United Kingdom

EMail: trevor.burbridge@bt.com

Sam Crawford
SamKnows

EMail: sam@samknows.com

Philip Eardley
BT
Adastral Park, Martlesham Heath
Ipswich
United Kingdom

EMail: philip.eardley@bt.com

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

EMail: acmorton@att.com