            IPv4, IPv6, and IPv4-IPv6 Coexistence:
        Updates for the IP Performance Metrics (IPPM) Framework

Abstract

   This memo updates the IP Performance Metrics (IPPM) framework defined
   by RFC 2330 with new considerations for measurement methodology and
   testing. It updates the definition of standard-formed packets to
   include IPv6 packets, deprecates the definition of minimal IP packet,
   and augments distinguishing aspects, referred to as Type-P, for test
   packets in RFC 2330. This memo identifies that IPv4-IPv6 coexistence
   can challenge measurements within the scope of the IPPM framework.
   Example use cases include, but are not limited to, IPv4-IPv6
   translation, NAT, and protocol encapsulation. IPv6 header
   compression and use of IPv6 over Low-Power Wireless Area Networks
   (6LoWPAN) are considered and excluded from the standard-formed packet
   evaluation.

Table of Contents

1. Introduction

   The IETF IP Performance Metrics (IPPM) working group first created a
   framework for metric development in [RFC2330]. This framework has
   stood the test of time and enabled development of many fundamental
   metrics. It has been updated in the area of metric composition
   [RFC5835] and in several areas related to active stream measurement
   of modern networks with reactive properties [RFC7312].

   The IPPM framework [RFC2330] recognized (in Section 13) that many
   aspects of an IP packet can influence its processing during transfer
   across the network.

   In Section 15 of [RFC2330], the notion of a "standard-formed" packet
   is defined. However, the definition was never expanded to include
   IPv6, even though the authors of [RFC2330] explicitly identified the
   need for this update in Section 15: "the version field is 4 (later,
   we will expand this to include 6)".

   In particular, IPv6 Extension Headers and protocols that use IPv6
   header compression are growing in use. This memo seeks to provide
   the needed updates to the original definition in [RFC2330].

2. Requirements Language

   The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
   "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and
   "OPTIONAL" in this document are to be interpreted as described in
   BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all
   capitals, as shown here.

3. Scope

   The purpose of this memo is to expand the coverage of IPPM to include
   IPv6, highlight additional aspects of test packets, and make them
   part of the IPPM framework.

   The scope is to update key sections of [RFC2330], adding
   considerations that will aid the development of new measurement
   methodologies intended for today's IP networks. Specifically, this
   memo expands the Type-P examples in Section 13 of [RFC2330] and
   expands the definition (in Section 15 of [RFC2330]) of a standard-
   formed packet to include IPv6 header aspects and other features.

   Other topics in [RFC2330] that might be updated or augmented are
   deferred to future work. This includes the topics of passive and
   various forms of hybrid active/passive measurements.

4. Packets of Type-P

   A fundamental property of many Internet metrics is that the measured
   value of the metric depends on characteristics of the IP packet(s)
   used to make the measurement. Potential influencing factors include
   IP header fields and their values, as well as higher-layer protocol
   headers and their values. Consider an IP-connectivity metric: one
   obtains different results depending on whether one is interested in,
   for example, connectivity for packets destined for well-known TCP
   ports or unreserved UDP ports, those with invalid IPv4 checksums, or
   those with TTL or Hop Limit of 16. In some circumstances, these
   distinctions will result in special treatment of packets in
   intermediate nodes and end systems -- for example, if Diffserv
   [RFC2474], Explicit Congestion Notification (ECN) [RFC3168], Router
   Alert [RFC6398], Hop-by-Hop extensions [RFC7045], or Flow Labels
   [RFC6437] are used, or in the presence of firewalls or RSVP
   reservations.

   Because of this distinction, we introduce the generic notion of a
   "packet of Type-P", where in some contexts P will be explicitly
   defined (i.e., exactly what type of packet we mean), partially
   defined (e.g., "with a payload of B octets"), or left generic. Thus,
   we may talk about generic IP-Type-P-connectivity or more specific
   IP-port-HTTP-connectivity. Some metrics and methodologies may be
   fruitfully defined using generic Type-P definitions, which are then
   made specific when performing actual measurements.

   Whenever a metric's value depends on the type of the packets
   involved, the metric's name will include either a specific type or a
   phrase such as "Type-P". Thus, we will not define an
   "IP-connectivity" metric but instead an "IP-Type-P-connectivity"
   metric and/or perhaps an "IP-port-HTTP-connectivity" metric. This
   naming convention serves as an important reminder that one must be
   conscious of the exact type of traffic being measured.

   If the information constituting Type-P at the Source is found to have
   changed at the Destination (or at a measurement point between the
   Source and Destination, as in [RFC5644]), then the modified values
   MUST be noted and reported with the results. Some modifications
   occur according to the conditions encountered in transit (such as
   congestion notification) or due to the requirements of segments of
   the Source-to-Destination path. For example, the packet length will
   change if IP headers are converted to the alternate version/address
   family or optional Extension Headers are added or removed. Even
   header fields like TTL/Hop Limit that typically change in transit may
   be relevant to specific tests. For example, Neighbor Discovery
   Protocol (NDP) [RFC4861] packets are transmitted with the Hop Limit
   value set to 255, and the validity test specifies that the Hop Limit

MUST have a value of 255 at the receiver, too. So, while other tests
may intentionally exclude the TTL/Hop Limit value from their Type-P
definition, for this particular test, the correct Hop Limit value is
of high relevance and MUST be part of the Type-P definition.

Local policies in intermediate nodes based on examination of IPv6
Extension Headers may affect measurement repeatability. If
intermediate nodes follow the recommendations of [RFC7045],
repeatability may be improved to some degree.

A closely related note: It would be very useful to know if a given
Internet component (like a host, link, or path) treats equally a
class C of different types of packets. If so, then any one of those
types of packets can be used for subsequent measurement of the
component. This suggests we should devise a metric or suite of
metrics that attempt to determine class C (a designation that has no
relationship to address assignments, of course).

Load-balancing over parallel paths is one particular example where
such a class C would be more complex to determine in IPPM
measurements. Load balancers and routers often use flow identifiers,
computed as hashes (of specific parts) of the packet header, for
deciding among the available parallel paths a packet will traverse.
Packets with identical hashes are assigned to the same flow and
forwarded to the same resource in the load balancer's (or router's)
pool. The presence of a load balancer on the measurement path, as
well as the specific headers and fields that are used for the
forwarding decision, are not known when measuring the path as a black
box. Potential assessment scenarios include the measurement of one
of the parallel paths, and the measurement of all available parallel
paths that the load balancer can use. Therefore, knowledge of a load
balancer's flow definition (alternatively, its class-C-specific
treatment in terms of header fields in scope of hash operations) is a
prerequisite for repeatable measurements. A path may have more than
one stage of load-balancing, adding to class C definition complexity.

5. Standard-Formed Packets

   Unless otherwise stated, all metric definitions that concern IP
   packets include an implicit assumption that the packet is standard-
   formed. A packet is standard-formed if it meets all of the following
   REQUIRED criteria:

   + It includes a valid IP header. See below for version-specific
      criteria.

   + It is not an IP fragment.

+   The Source and Destination addresses correspond to the intended
    Source and Destination, including Multicast Destination addresses.

+   If a transport header is present, it contains a valid checksum and
    other valid fields.

For an IPv4 packet (as specified in [RFC791] and the RFCs that update
it) to be standard-formed, the following additional criteria are
REQUIRED:

o   The version field is 4.

o   The Internet Header Length (IHL) value is >= 5; the checksum is
    correct.

o   Its total length as given in the IPv4 header corresponds to the
    size of the IPv4 header plus the size of the payload.

o   Either the packet possesses sufficient TTL to travel from the
    Source to the Destination if the TTL is decremented by one at each
    hop or it possesses the maximum TTL of 255.

o   It does not contain IP options unless explicitly noted.

For an IPv6 packet (as specified in [RFC8200] and any future updates)
to be standard-formed, the following criteria are REQUIRED:

o   The version field is 6.

o   Its total length corresponds to the size of the IPv6 header (40
    octets) plus the length of the payload as given in the IPv6
    header.

o   The payload length value for this packet (including Extension
    Headers) conforms to the IPv6 specifications.

o   Either the packet possesses sufficient Hop Limit to travel from
    the Source to the Destination if the Hop Limit is decremented by
    one at each hop or it possesses the maximum Hop Limit of 255.

o   Either the packet does not contain IP Extension Headers or it
    contains the correct number and type of headers as specified in
    the packet and the headers appear in the standard-conforming order
    (Next Header).

o   All parameters used in the header and Extension Headers are found
    in the "Internet Protocol Version 6 (IPv6) Parameters" registry
    specified in [IANA-6P].

Two mechanisms require some discussion in the context of standard-
formed packets, namely IPv6 over Low-Power Wireless Area Networks
(6LowPAN) [RFC4944] and Robust Header Compression (ROHC) [RFC3095].
6LowPAN, as defined in [RFC4944] and updated by [RFC6282] with header
compression and [RFC6775] with neighbor discovery optimizations,
proposes solutions for using IPv6 in resource-constrained
environments. An adaptation layer enables the transfer of IPv6
packets over networks having an MTU smaller than the minimum IPv6
MTU. Fragmentation and reassembly of IPv6 packets, as well as the
resulting state that would be stored in intermediate nodes, poses
substantial challenges to measurements. Likewise, ROHC operates
statefully in compressing headers on subpaths, storing state in
intermediate hosts. The modification of measurement packets' Type-P
by ROHC and 6LowPAN requires substantial work, as do requirements
with respect to the concept of standard-formed packets for these two
protocols. For these reasons, we consider ROHC and 6LowPAN packets
to be out of the scope of the standard-formed packet evaluation.

The topic of IPv6 Extension Headers brings current controversies into
focus, as noted by [RFC6564] and [RFC7045]. However, measurement use
cases in the context of the IPPM framework, such as in situ OAM
[IOAM-DATA] in enterprise environments, can benefit from inspection,
modification, addition, or deletion of IPv6 extension headers in
hosts along the measurement path.

[RFC8250] endorses the use of the IPv6 Destination Option for
measurement purposes, consistent with other relevant and approved
IETF specifications.

The following additional considerations apply when IPv6 Extension
Headers are present:

o Extension Header inspection: Some intermediate nodes may inspect
   Extension Headers or the entire IPv6 packet while in transit. In
   exceptional cases, they may drop the packet or route via a
   suboptimal path, and measurements may be unreliable or
   unrepeatable. The packet (if it arrives) may be standard-formed,
   with a corresponding Type-P.

o Extension Header modification: In Hop-by-Hop headers, some TLV-
   encoded options may be permitted to change at intermediate nodes
   while in transit. The resulting packet may be standard-formed,
   with a corresponding Type-P.

o Extension Header insertion or deletion: Although such behavior is
   not endorsed by current standards, it is possible that Extension
   Headers could be added to, or removed from, the header chain. The
   resulting packet may be standard-formed, with a corresponding
   Type-P. This point simply encourages measurement system designers
   to be prepared for the unexpected and notify users when such
   events occur. There are issues with Extension Header insertion
   and deletion, of course, such as exceeding the path MTU due to
   insertion, etc.

o A change in packet length (from the corresponding packet observed
   at the Source) or header modification is a significant factor in
   Internet measurement and REQUIRES a new Type-P to be reported with
   the test results.

It is further REQUIRED that if a packet is described as having a
"length of B octets", then 0 <= B <= 65535; and if B is the payload
length in octets, then B <= (65535-IP header size in octets,
including any Extension Headers). The jumbograms defined in
[RFC2675] are not covered by the above length analysis, but if the
IPv6 Jumbogram Payload Hop-by-Hop Option Header is present, then a
packet with corresponding length MUST be considered standard-formed.
In practice, the path MTU will restrict the length of standard-formed
packets that can successfully traverse the path. Path MTU Discovery
for IP version 6 (PMTUD, [RFC8201]) or Packetization Layer Path MTU
Discovery (PLPMTUD, [RFC4821]) is recommended to prevent
fragmentation.

So, for example, one might imagine defining an IP-connectivity metric
as "IP-Type-P-connectivity for standard-formed packets with the IP
Diffserv field set to 0", or, more succinctly,
"IP-Type-P-connectivity with the IP Diffserv field set to 0", since
standard-formed is already implied by convention. Changing the
contents of a field, such as the Diffserv Code Point, ECN bits, or
Flow Label may have a profound effect on packet handling during
transit, but does not affect a packet's status as standard-formed.
Likewise, the addition, modification, or deletion of extension
headers may change the handling of packets in transit hosts.

[RFC2330] defines the "minimal IP packet from A to B" as a particular
type of standard-formed packet often useful to consider. When
defining IP metrics, no packet smaller or simpler than this can be
transmitted over a correctly operating IP network. However, the
concept of the minimal IP packet has not been employed (since typical
active measurement systems employ a transport layer and a payload),
and its practical use is limited. Therefore, this memo deprecates
the concept of the "minimal IP packet from A to B".

6. NAT, IPv4-IPv6 Transition, and Compression Techniques

   This memo adds the key considerations for utilizing IPv6 in two
   critical conventions of the IPPM framework, namely packets of Type-P
   and standard-formed packets. The need for coexistence of IPv4 and
   IPv6 has originated transitioning standards like the framework for
   IPv4/IPv6 translation in [RFC6144] or the IP/ICMP translation
   algorithms in [RFC7915] and [RFC7757].

   The definition and execution of measurements within the context of
   the IPPM framework is challenged whenever such translation mechanisms
   are present along the measurement path. In use cases like IPv4-IPv6
   translation, NAT, protocol encapsulation, or IPv6 header compression
   may result in modification of the measurement packet's Type-P along
   the path. All these changes MUST be reported. Example consequences
   include, but are not limited to:

   o Modification or addition of headers or header field values in
      intermediate nodes. IPv4-IPv6 transitioning or IPv6 header
      compression mechanisms may result in changes of the measurement
      packets' Type-P, too. Consequently, hosts along the measurement
      path may treat packets differently because of the Type-P
      modification. Measurements at observation points along the path
      may also need extra context to uniquely identify a packet.

   o Network Address Translators (NAT) on the path can have an
      unpredictable impact on latency measurement (in terms of the
      amount of additional time added) and possibly other types of
      measurements. It is not usually possible to control this impact
      as testers may not have any control of the underlying network or
      middleboxes. There is a possibility that stateful NAT will lead
      to unstable performance for a flow with specific Type-P, since
      state needs to be created for the first packet of a flow and state
      may be lost later if the NAT runs out of resources. However, this
      scenario does not invalidate the Type-P for testing; for example,
      the purpose of a test might be exactly to quantify the NAT's
      impact on delay variation. The presence of NAT may mean that the
      measured performance of Type-P will change between the source and
      the destination. This can cause an issue when attempting to
      correlate measurements conducted on segments of the path that
      include or exclude the NAT. Thus, it is a factor to be aware of
      when conducting measurements.

   o Variable delay due to internal state. One side effect of changes
      due to IPv4-IPv6 transitioning mechanisms is the variable delay
      that intermediate nodes experience for header modifications.
      Similar to NAT, the allocation of internal state and establishment
      of context within intermediate nodes may cause variable delays,
      depending on the measurement stream pattern and position of a
      packet within the stream. For example, the first packet in a
      stream will typically trigger allocation of internal state in an
      intermediate IPv4-IPv6 transition host. Subsequent packets can
      benefit from lower processing delay due to the existing internal
      state. However, large interpacket delays in the measurement
      stream may result in the intermediate host deleting the associated
      state and needing to re-establish it on arrival of another stream
      packet. It is worth noting that this variable delay due to
      internal state allocation in intermediate nodes can be an explicit
      use case for measurements.

   o Variable delay due to packet length. IPv4-IPv6 transitioning or
      header compression mechanisms modify the length of measurement
      packets. The modification of the packet size may or may not
      change how the measurement path treats the packets.

7. Security Considerations

   The security considerations that apply to any active measurement of
   live paths are relevant here as well. See [RFC4656] and [RFC5357].

   When considering the privacy of those involved in measurement or
   those whose traffic is measured, the sensitive information available
   to potential observers is greatly reduced when using active
   techniques that are within this scope of work. Passive observations
   of user traffic for measurement purposes raise many privacy issues.
   We refer the reader to the privacy considerations described in the
   Large Scale Measurement of Broadband Performance (LMAP) framework
   [RFC7594], which covers active and passive techniques.

8. IANA Considerations

   This document has no IANA actions.

9. References

9.1. Normative References

   [RFC791]   Postel, J., "Internet Protocol", STD 5, RFC 791,
               DOI 10.17487/RFC0791, September 1981,
               <https://www.rfc-editor.org/info/rfc791>.

   [RFC2119]   Bradner, S., "Key words for use in RFCs to Indicate
               Requirement Levels", BCP 14, RFC 2119,
               DOI 10.17487/RFC2119, March 1997,
               <https://www.rfc-editor.org/info/rfc2119>.

   [RFC2330]   Paxson, V., Almes, G., Mahdavi, J., and M. Mathis,
               "Framework for IP Performance Metrics", RFC 2330,
               DOI 10.17487/RFC2330, May 1998,
               <https://www.rfc-editor.org/info/rfc2330>.

   [RFC2474]   Nichols, K., Blake, S., Baker, F., and D. Black,
               "Definition of the Differentiated Services Field (DS
               Field) in the IPv4 and IPv6 Headers", RFC 2474,
               DOI 10.17487/RFC2474, December 1998,
               <https://www.rfc-editor.org/info/rfc2474>.

   [RFC2675]   Borman, D., Deering, S., and R. Hinden, "IPv6 Jumbograms",
               RFC 2675, DOI 10.17487/RFC2675, August 1999,
               <https://www.rfc-editor.org/info/rfc2675>.

   [RFC3095]   Bormann, C., Burmeister, C., Degermark, M., Fukushima, H.,
               Hannu, H., Jonsson, L-E., Hakenberg, R., Koren, T., Le,
               K., Liu, Z., Martensson, A., Miyazaki, A., Svanbro, K.,
               Wiebke, T., Yoshimura, T., and H. Zheng, "RObust Header
               Compression (ROHC): Framework and four profiles: RTP, UDP,
               ESP, and uncompressed", RFC 3095, DOI 10.17487/RFC3095,
               July 2001, <https://www.rfc-editor.org/info/rfc3095>.

   [RFC3168]   Ramakrishnan, K., Floyd, S., and D. Black, "The Addition
               of Explicit Congestion Notification (ECN) to IP",
               RFC 3168, DOI 10.17487/RFC3168, September 2001,
               <https://www.rfc-editor.org/info/rfc3168>.

   [RFC4656]   Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and
               M. Zekauskas, "A One-way Active Measurement Protocol
               (OWAMP)", RFC 4656, DOI 10.17487/RFC4656, September 2006,
               <https://www.rfc-editor.org/info/rfc4656>.

   [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU
              Discovery", RFC 4821, DOI 10.17487/RFC4821, March 2007,
              <https://www.rfc-editor.org/info/rfc4821>.

   [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman,
              "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861,
              DOI 10.17487/RFC4861, September 2007,
              <https://www.rfc-editor.org/info/rfc4861>.

   [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler,
              "Transmission of IPv6 Packets over IEEE 802.15.4
              Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007,
              <https://www.rfc-editor.org/info/rfc4944>.

   [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and
              J. Babiarz, "A Two-Way Active Measurement Protocol
              (TWAMP)", RFC 5357, DOI 10.17487/RFC5357, October 2008,
              <https://www.rfc-editor.org/info/rfc5357>.

   [RFC5644] Stephan, E., Liang, L., and A. Morton, "IP Performance
              Metrics (IPPM): Spatial and Multicast", RFC 5644,
              DOI 10.17487/RFC5644, October 2009,
              <https://www.rfc-editor.org/info/rfc5644>.

   [RFC5835] Morton, A., Ed. and S. Van den Berghe, Ed., "Framework for
              Metric Composition", RFC 5835, DOI 10.17487/RFC5835, April
              2010, <https://www.rfc-editor.org/info/rfc5835>.

   [RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for
              IPv4/IPv6 Translation", RFC 6144, DOI 10.17487/RFC6144,
              April 2011, <https://www.rfc-editor.org/info/rfc6144>.

   [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6
              Datagrams over IEEE 802.15.4-Based Networks", RFC 6282,
              DOI 10.17487/RFC6282, September 2011,
              <https://www.rfc-editor.org/info/rfc6282>.

   [RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and
              Usage", BCP 168, RFC 6398, DOI 10.17487/RFC6398, October
              2011, <https://www.rfc-editor.org/info/rfc6398>.

   [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme,
              "IPv6 Flow Label Specification", RFC 6437,
              DOI 10.17487/RFC6437, November 2011,
              <https://www.rfc-editor.org/info/rfc6437>.

   [RFC6564] Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J., and
              M. Bhatia, "A Uniform Format for IPv6 Extension Headers",
              RFC 6564, DOI 10.17487/RFC6564, April 2012,
              <https://www.rfc-editor.org/info/rfc6564>.

   [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and
              C. Bormann, "Neighbor Discovery Optimization for IPv6 over
              Low-Power Wireless Personal Area Networks (6LoWPANs)",
              RFC 6775, DOI 10.17487/RFC6775, November 2012,
              <https://www.rfc-editor.org/info/rfc6775>.

   [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing
              of IPv6 Extension Headers", RFC 7045,
              DOI 10.17487/RFC7045, December 2013,
              <https://www.rfc-editor.org/info/rfc7045>.

   [RFC7312] Fabini, J. and A. Morton, "Advanced Stream and Sampling
              Framework for IP Performance Metrics (IPPM)", RFC 7312,
              DOI 10.17487/RFC7312, August 2014,
              <https://www.rfc-editor.org/info/rfc7312>.

   [RFC7757] Anderson, T. and A. Leiva Popper, "Explicit Address
              Mappings for Stateless IP/ICMP Translation", RFC 7757,
              DOI 10.17487/RFC7757, February 2016,
              <https://www.rfc-editor.org/info/rfc7757>.

   [RFC7915] Bao, C., Li, X., Baker, F., Anderson, T., and F. Gont,
              "IP/ICMP Translation Algorithm", RFC 7915,
              DOI 10.17487/RFC7915, June 2016,
              <https://www.rfc-editor.org/info/rfc7915>.

   [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC
              2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174,
              May 2017, <https://www.rfc-editor.org/info/rfc8174>.

   [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6
              (IPv6) Specification", STD 86, RFC 8200,
              DOI 10.17487/RFC8200, July 2017,
              <https://www.rfc-editor.org/info/rfc8200>.

   [RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed.,
              "Path MTU Discovery for IP version 6", STD 87, RFC 8201,
              DOI 10.17487/RFC8201, July 2017,
              <https://www.rfc-editor.org/info/rfc8201>.

   [RFC8250] Elkins, N., Hamilton, R., and M. Ackermann, "IPv6
              Performance and Diagnostic Metrics (PDM) Destination
              Option", RFC 8250, DOI 10.17487/RFC8250, September 2017,
              <https://www.rfc-editor.org/info/rfc8250>.

9.2. Informative References

   [IANA-6P] IANA, "Internet Protocol Version 6 (IPv6) Parameters",
              <https://www.iana.org/assignments/ipv6-parameters>.

   [IOAM-DATA]
              Brockners, F., Bhandari, S., Pignataro, C., Gredler, H.,
              Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov,
              P., Chang, R., daniel.bernier@bell.ca, d., and J. Lemon,
              "Data Fields for In-situ OAM", Work in Progress,
              draft-ietf-ippm-ioam-data-03, June 2018.

   [RFC7594] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T.,
              Aitken, P., and A. Akhter, "A Framework for Large-Scale
              Measurement of Broadband Performance (LMAP)", RFC 7594,
              DOI 10.17487/RFC7594, September 2015,
              <https://www.rfc-editor.org/info/rfc7594>.

Acknowledgements

Authors' Addresses

   Al Morton
   AT&T Labs
   200 Laurel Avenue South
   Middletown, NJ 07748
   United States of America
   Phone: +1 732 420 1571
   Fax: +1 732 368 1192
   Email: acm@researh.att.com

   Joachim Fabini
   TU Wien
   Gusshausstrasse 25/E389
   Vienna 1040
   Austria
   Phone: +43 1 58801 38813
   Fax: +43 1 58801 38898
   Email: Joachim.Fabini@tuwien.ac.at
   URI: http://www.tc.tuwien.ac.at/about-us/staff/joachim-fabini/

   Nalini Elkins
   Inside Products, Inc.
   Carmel Valley, CA 93924
   United States of America
   Email: nalini.elkins@insidethestack.com

   Michael S. Ackermann
   Blue Cross Blue Shield of Michigan
   Email: mackermann@bcbsm.com

   Vinayak Hegde
   Consultant
   Brahma Sun City, Wadgaon-Sheri
   Pune, Maharashtra 411014
   India
   Phone: +91 9449834401
   Email: vinayakh@gmail.com
   URI: http://www.vinayakhegde.com