                  Large-Scale Broadband Measurement Use Cases

Abstract

   Measuring broadband performance on a large scale is important for
   network diagnostics by providers and users, as well as for public
   policy. Understanding the various scenarios and users of measuring
   broadband performance is essential to development of the Large-scale
   Measurement of Broadband Performance (LMAP) framework, information
   model, and protocol. This document details two use cases that can
   assist in developing that framework. The details of the measurement
   metrics themselves are beyond the scope of this document.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

   This document describes two use cases for the Large-scale Measurement
   of Broadband Performance (LMAP). The use cases contained in this
   document are (1) the Internet Service Provider Use Case and (2) the
   Regulator Use Case. In the first, a network operator wants to
   understand the performance of the network and the quality experienced
   by customers, while in the second, a regulator wants to provide
   information on the performance of the ISPs in their jurisdiction.
   There are other use cases that are not the focus of the initial LMAP
   work (for example, end users would like to use measurements to help
   identify problems in their home network and to monitor the
   performance of their broadband provider); it is expected that the
   same mechanisms are applicable.

   Large-scale measurements raise several security concerns, including
   privacy issues. These are summarized in Section 7 and considered in
   further detail in [Framework].

2. Use Cases

   From the LMAP perspective, there is no difference between fixed
   service and mobile (cellular) service used for Internet access.
   Hence, like measurements will take place on both fixed and mobile
   networks. Fixed services include technologies like Digital
   Subscriber Line (DSL), Cable, and Carrier Ethernet. Mobile services
   include all those advertised as 2G, 3G, 4G, and Long Term Evolution
   (LTE). A metric defined to measure end-to-end services will execute
   similarly on all access technologies. Other metrics may be access
   technology specific. The LMAP architecture covers both IPv4 and IPv6
   networks.

2.1. Internet Service Provider (ISP) Use Case

   A network operator needs to understand the performance of their
   networks, the performance of the suppliers (downstream and upstream
   networks), the performance of Internet access services, and the
   impact that such performance has on the experience of their
   customers. Largely, the processes that ISPs operate (which are based
   on network measurement) include:

   o Identifying, isolating, and fixing problems, which may be in the
      network, with the service provider, or in the end-user equipment.
      Such problems may be common to a point in the network topology
      (e.g., a single exchange), common to a vendor or equipment type
      (e.g., line card or home gateway), or unique to a single user line
      (e.g., copper access). Part of this process may also be helping

        users understand whether the problem exists in their home network
        or with a third-party application service instead of with their
        broadband (BB) product.

     o Design and planning. Through monitoring the end-user experience,
        the ISP can design and plan their network to ensure specified
        levels of user experience. Services may be moved closer to end
        users, services upgraded, the impact of QoS assessed, or more
        capacity deployed at certain locations. Service Level Agreements
        (SLAs) may be defined at network or product boundaries.

     o Understanding the quality experienced by customers. The network
        operator would like to gain better insight into the end-to-end
        performance experienced by its customers. "End-to-end" could, for
        instance, incorporate home and enterprise networks, and the impact
        of peering, caching, and Content Delivery Networks (CDNs).

     o Understanding the impact and operation of new devices and
        technology. As a new product is deployed, or a new technology
        introduced into the network, it is essential that its operation
        and its impact are measured. This also helps to quantify the
        advantage that the new technology is bringing and support the
        business case for larger roll-out.

2.2. Regulator Use Case

   A regulator may want to evaluate the performance of the Internet
   access services offered by operators.

   While each jurisdiction responds to distinct consumer, industry, and
   regulatory concerns, much commonality exists in the need to produce
   datasets that can be used to compare multiple Internet access service
   providers, diverse technical solutions, geographic and regional
   distributions, and marketed and provisioned levels and combinations
   of broadband Internet access services.

   Regulators may want to publish performance measures of different ISPs
   as background information for end users. They may also want to track
   the growth of high-speed broadband deployment, or to monitor the
   traffic management practices of Internet providers.

   A regulator's role in the development and enforcement of broadband
   Internet access service policies requires that the measurement
   approaches meet a high level of verifiability, accuracy, and
   provider-independence to support valid and meaningful comparisons of
   Internet access service performance. Standards can help regulators'

shared needs for scalable, cost-effective, scientifically robust
solutions to the measurement and collection of broadband Internet
access service performance information.

3. Details of ISP Use Case

3.1. Understanding the Quality Experienced by Customers

Operators want to understand the quality of experience (QoE) of their
broadband customers. The understanding can be gained through a
"panel", i.e., measurement probes deployed to several customers. A
probe is a device or piece of software that makes measurements and
reports the results, under the control of the measurement system.
Implementation options are discussed in Section 5. The panel needs
to include a representative sample of the operator's technologies and
broadband speeds. For instance, it might encompass speeds ranging
from below 8 Mbps to over 100 Mbps. The operator would like the
end-to-end view of the service, rather than just the access portion.
This involves relating the pure network parameters to something like
a 'mean opinion score' [MOS], which will be service dependent (for
instance, web-browsing QoE is largely determined by latency above a
few Mbps).

An operator will also want compound metrics such as "reliability",
which might involve packet loss, DNS failures, retraining of the
line, video streaming under-runs, etc.

The operator really wants to understand the end-to-end service
experience. However, the home network (Ethernet, Wi-Fi, powerline)
is highly variable and outside its control. To date, operators (and
regulators) have instead measured performance from the home gateway.
However, mobile operators clearly must include the wireless link in
the measurement.

Active measurements are the most obvious approach, i.e., special
measurement traffic is sent by -- and to -- the probe. In order not
to degrade the service of the customer, the measurement data should
only be sent when the user is silent, and it shouldn't reduce the
customer's data allowance. The other approach is passive
measurements on the customer's ordinary traffic; the advantage is
that it measures what the customer actually does, but it creates
extra variability (different traffic mixes give different results)
and, in particular, it raises privacy concerns. [RFC6973] discusses
privacy considerations for Internet protocols in general, while
[Framework] discusses them specifically for large-scale measurement
systems.

From an operator's viewpoint, understanding customer experience
enables it to offer better services. Also, simple metrics can be
more easily understood by senior managers who make investment
decisions and by sales and marketing.

3.2. Understanding the Impact and Operation of New Devices and
      Technology

Another type of measurement is to test new capabilities before they
are rolled out. For example, the operator may want to:

o Check whether a customer can be upgraded to a new broadband
   option.

o Understand the impact of IPv6 before it is made available to
   customers. Questions such as these could be assessed: Will v6
   packets get through? What will the latency be to major websites?
   What transition mechanisms will be most appropriate?

o Check whether a new capability can be signaled using TCP options
   (how often it will be blocked by a middlebox -- along the lines of
   the experiments described in [Extend-TCP]).

o Investigate a QoS mechanism (e.g., checking whether Diffserv
   markings are respected on some path).

3.3. Design and Planning

Operators can use large-scale measurements to help with their network
planning -- proactive activities to improve the network.

For example, by probing from several different vantage points the
operator can see that a particular group of customers has performance
below that expected during peak hours, which should help with
capacity planning. Naturally, operators already have tools to help
with this -- a network element reports its individual utilization
(and perhaps other parameters). However, making measurements across
a path rather than at a point may make it easier to understand the
network. There may also be parameters like bufferbloat that aren't
currently reported by equipment and/or that are intrinsically path
metrics.

With information gained from measurement results, capacity planning
and network design can be more effective. Such planning typically
uses simulations to emulate the measured performance of the current
network and understand the likely impact of new capacity and
potential changes to the topology. Simulations, informed by data

   from a limited panel of probes, can help quantify the advantage that
   a new technology brings and support the business case for larger
   roll-out.

   It may also be possible to use probes to run stress tests for risk
   analysis. For example, an operator could run a carefully controlled
   and limited experiment in which probing is used to assess the
   potential impact if some new application becomes popular.

3.4. Monitoring Service Level Agreements

   Another example is that the operator may want to monitor performance
   where there is a Service Level Agreement (SLA). This could be with
   its own customers; in particular, enterprises may have an SLA. The
   operator can proactively spot when the service is degrading near the
   point of the SLA limit and get information that will enable more
   informed conversations with the customer at contract renewal.

   An operator may also want to monitor the performance of its
   suppliers, to check whether they meet their SLA or to compare two
   suppliers if it is dual-sourcing. This could include its transit
   operator, CDNs, peering, video source, or local network provider for
   a global operator in countries where it doesn't have its own network.
   A virtual operator may monitor the whole underlying network.

   Through a better understanding of its own network and its suppliers,
   the operator should be able to focus investment more effectively --
   in the right place at the right time with the right technology.

3.5. Identifying, Isolating, and Fixing Network Problems

   Operators can use large-scale measurements to help identify a fault
   more rapidly and decide how to solve it.

   Operators already have Test and Diagnostic tools, where a network
   element reports some problem or failure to a management system.
   However, many issues are not caused by a point failure but something
   wider and so will trigger too many alarms, while other issues will
   cause degradation rather than failure and so not trigger any alarm.
   Large-scale measurements can help provide a more nuanced view that
   helps network management to identify and fix problems more rapidly
   and accurately. The network management tools may use simulations to
   emulate the network and so help identify a fault and assess possible
   solutions.

An operator can obtain useful information without measuring the
performance on every broadband line. By measuring a subset, the
operator can identify problems that affect a group of customers. For
example, the issue could be at a shared point in the network topology
(such as an exchange), or common to a vendor, or equipment type; for
instance, [IETF85-Plenary] describes a case where a particular home
gateway upgrade had caused a (mistaken!) drop in line rate.

A more extensive deployment of the measurement capability to every
broadband line would enable an operator to identify issues unique to
a single customer. Overall, large-scale measurements can help an
operator fix the fault more rapidly and/or allow the affected
customers to be informed of what's happening. More accurate
information enables the operator to reassure customers and take more
rapid and effective action to cure the problem.

Often, customers experience poor broadband due to problems in the
home network -- the ISP's network is fine. For example, they may
have moved too far away from their wireless access point.
Anecdotally, a large fraction of customer calls about fixed BB
problems are due to in-home wireless issues. These issues are
expensive and frustrating for an operator, as they are extremely hard
to diagnose and solve. The operator would like to narrow down
whether the problem is in the home (a problem with the home network,
edge device, or home gateway), in the operator's network, or with an
application service. The operator would like two capabilities:
firstly, self-help tools that customers use to improve their own
service or understand its performance better -- for example, to
reposition their devices for better Wi-Fi coverage; and secondly,
on-demand tests that the operator can run instantly, so that the call
center person answering the phone (or e-chat) could trigger a test
and get the result while the customer is still in an online session.

4. Details of Regulator Use Case

4.1. Providing Transparent Performance Information

Some regulators publish information about the quality of the various
Internet access services provided in their national market. Quality
information about service offers could include speed, delay, and
jitter. Such information can be published to facilitate end users'
choice of service provider and offer. Regulators may check the
accuracy of the marketing claims of Internet service providers and
may also encourage ISPs to all use the same metrics in their service
level contracts. The goal of these transparency mechanisms is to
promote competition for end users and potentially also help content,
application, service, and device providers develop their Internet
offerings.

The published information needs to be:

o Accurate - the measurement results must be correct and not
   influenced by errors or side effects. The results should be
   reproducible and consistent over time.

o Comparable - common metrics should be used across different ISPs
   and service offerings, and over time, so that measurement results
   can be compared.

o Meaningful - the metrics used for measurements need to reflect
   what end users value about their broadband Internet access
   service.

o Reliable - the number and distribution of measurement agents, and
   the statistical processing of the raw measurement data, need to be
   appropriate.

In practical terms, the regulators may measure network performance
from users towards multiple content and application providers,
including dedicated test measurement servers. Measurement probes are
distributed to a 'panel' of selected end users. The panel covers all
the operators and packages in the market, spread over urban,
suburban, and rural areas, and often includes both fixed and mobile
Internet access. Periodic tests running on the probes can, for
example, measure actual speed at peak and off-peak hours, but can
also measure other detailed quality metrics like delay and jitter.
Collected data goes afterwards through statistical analysis, deriving
estimates for the whole population. Summary information, such as a
service quality index, is published regularly, perhaps alongside more
detailed information.

The regulator can also facilitate end users to monitor the
performance of their own broadband Internet access service. They
might use this information to check that the performance meets that
specified in their contract or to understand whether their current
subscription is the most appropriate.

4.2. Measuring Broadband Deployment

Regulators may also want to monitor the improvement over time of
actual broadband Internet access performance in a specific country or
a region. The motivation is often to evaluate the effect of the
stimulated growth over time, when government has set a strategic goal
for high-speed broadband deployment, whether in absolute terms or
benchmarked against other countries. An example of such an
initiative is [DAE]. The actual measurements can be made in the same
way as described in Section 4.1.

4.3. Monitoring Traffic Management Practices

   A regulator may want to monitor traffic management practices or
   compare the performance of Internet access service with specialized
   services offered in parallel to, but separate from, Internet access
   service (for example, IPTV). A regulator could monitor for
   departures from application agnosticism such as blocking or
   throttling of traffic from specific applications, or preferential
   treatment of specific applications. A measurement system could send,
   or passively monitor, application-specific traffic and then measure
   in detail the transfer of the different packets. While it is
   relatively easy to measure port blocking, how to detect other types
   of differentiated treatment is a research topic in itself. The
   "Glasnost: Enabling End Users to Detect Traffic Differentiation"
   paper [M-Labs_NSDI-2010] and follow-on tool "Glasnost" [Glasnost]
   provide an example of work in this area.

   A regulator could also monitor the performance of the broadband
   service over time, to try and detect if the specialized service is
   provided at the expense of the Internet access service. Comparison
   between ISPs or between different countries may also be relevant for
   this kind of evaluation.

   The motivation for a regulator monitoring such traffic management
   practices is that regulatory approaches related to net neutrality and
   the open Internet have been introduced in some jurisdictions.
   Examples of such efforts are the Internet policy as outlined by the
   Body of European Regulators for Electronic Communications guidelines
   for quality of service [BEREC-Guidelines] and the US FCC's
   "Preserving the Open Internet" Report and Order [FCC-R&O]. Although
   legal challenges can change the status of policy, the take-away for
   LMAP purposes is that policy-makers are looking for measurement
   solutions to assist them in discovering biased treatment of traffic
   flows. The exact definitions and requirements vary from one
   jurisdiction to another.

5. Implementation Options

   There are several ways of implementing a measurement system. The
   choice may be influenced by the details of the particular use case
   and what the most important criteria are for the regulator, ISP, or
   third party operating the measurement system.

   One type of probe is a special hardware device that is connected
   directly to the home gateway. The devices are deployed to a
   carefully selected panel of end users, and they perform measurements
   according to a defined schedule. The schedule can run throughout the
   day, to allow continuous assessment of the network. Careful design

ensures that measurements do not detrimentally impact the home user
experience or corrupt the results by testing when the user is also
using the broadband line. The system is therefore tightly controlled
by the operator of the measurement system. One advantage of this
approach is that it is possible to get reliable benchmarks for the
performance of a network with only a few devices. One disadvantage
is that it would be expensive to deploy hardware devices on a mass
scale sufficient to understand the performance of the network at the
granularity of a single broadband user.

Another type of probe involves implementing the measurement
capability as a webpage or an "app" that end users are encouraged to
download onto their mobile phone or computing device. Measurements
are triggered by the end user; for example, the user interface may
have a button to "test my broadband now." One advantage of this
approach is that the performance is measured to the end user, rather
than to the home gateway, and so includes the home network. Another
difference is that the system is much more loosely controlled, as the
panel of end users and the schedule of tests are determined by the
end users themselves rather than the measurement system. While this
approach makes it easier to make measurements on a large scale, it is
harder to get comparable benchmarks, as the measurements are affected
by the home network; also, the population is self-selecting and so
potentially biased towards those who think they have a problem. This
could be alleviated by encouraging widespread downloading of the app
and careful post-processing of the results to reduce biases.

There are several other possibilities. For example, as a variant on
the first approach, the measurement capability could be implemented
as software embedded in the home gateway, which would make it more
viable to have the capability on every user line. As a variant on
the second approach, the end user could initiate measurements in
response to a request from the measurement system.

The operator of the measurement system should be careful to ensure
that measurements do not detrimentally impact users. Potential
issues include the following:

* Measurement traffic generated on a particular user's line may
   impact that end user's quality of experience. The danger is
   greater for measurements that generate a lot of traffic over a
   lengthy period.

* The measurement traffic may impact that particular user's bill or
   traffic cap.

   * The measurement traffic from several end users may, in
      combination, congest a shared link.

   * The traffic associated with the control and reporting of
      measurements may overload the network. The danger is greater
      where the traffic associated with many end users is synchronized.

6. Conclusions

   Large-scale measurements of broadband performance are useful for both
   network operators and regulators. Network operators would like to
   use measurements to help them better understand the quality
   experienced by their customers, identify problems in the network, and
   design network improvements. Regulators would like to use
   measurements to help promote competition between network operators,
   stimulate the growth of broadband access, and monitor 'net
   neutrality'. There are other use cases that are not the focus of the
   initial LMAP charter (although it is expected that the mechanisms
   developed would be readily applied); for example, end users would
   like to use measurements to help identify problems in their home
   network and to monitor the performance of their broadband provider.

   From consideration of the various use cases, several common themes
   emerge, while there are also some detailed differences. These
   characteristics guide the development of LMAP's framework,
   information model, and protocol.

   A measurement capability is needed across a wide number of
   heterogeneous environments. Tests may be needed in the home network,
   in the ISP's network, or beyond; they may be measuring a fixed or
   wireless network; they may measure just the access network or across
   several networks.

   There is a role for both standardized and non-standardized
   measurements. For example, a regulator would like to publish
   standardized performance metrics for all network operators, while an
   ISP may need their own tests to understand some feature special to
   their network. Most use cases need active measurements, which create
   and measure specific test traffic, but some need passive measurements
   of the end user's traffic.

   Regardless of the tests being operated, there needs to be a way to
   demand or schedule the tests. Most use cases need a regular schedule
   of measurements, but sometimes ad hoc testing is needed -- for
   example, for troubleshooting. It needs to be ensured that
   measurements do not affect the user experience and are not affected
   by user traffic (unless desired). In addition, there needs to be a

common way to collect the results. Standardization of this control
and reporting functionality allows the operator of a measurement
system to buy the various components from different vendors.

After the measurement results are collected, they need to be
understood and analyzed. Often, it is sufficient to measure only a
small subset of end users, but per-line fault diagnosis requires the
ability to test every individual line. Analysis requires accurate
definition and understanding of where the test points are, as well as
contextual information about the topology, line, product, and the
subscriber's contract. The actual analysis of results is beyond the
scope of LMAP, as is the key challenge of how to integrate the
measurement system into a network operator's existing tools for
diagnostics and network planning.

Finally, the test data, along with any associated network, product,
or subscriber contract data, is commercial or private information and
needs to be protected.

7. Security Considerations

Large-scale measurements raise several potential security, privacy
(data protection) [RFC6973], and business sensitivity issues:

1. A malicious party may try to gain control of probes to launch DoS
   (Denial of Service) attacks at a target. A DoS attack could be
   targeted at a particular end user or set of end users, a certain
   network, or a specific service provider.

2. A malicious party may try to gain control of probes to create a
   platform for pervasive monitoring [RFC7258] or for more targeted
   monitoring. [RFC7258] summarizes the threats as follows: "An
   attack may change the content of the communication, record the
   content or external characteristics of the communication, or
   through correlation with other communication events, reveal
   information the parties did not intend to be revealed." For
   example, a malicious party could distribute to the probes a new
   measurement test that recorded (and later reported) information of
   maleficent interest. Similar concerns also arise if the
   measurement results are intercepted or corrupted.

   * From the end user's perspective, the concerns include a
     malicious party monitoring the traffic they send and receive,
     who they communicate with, the websites they visit, and such
     information about their behavior as when they are at home and
     the location of their devices. Some of the concerns may be
     greater when the probe is on the end user's device rather than
     on their home gateway.

      * From the network operator's perspective, the concerns include
         the leakage of commercially sensitive information about the
         design and operation of their network, their customers, and
         suppliers. Some threats are indirect; for example, the
         attacker could reconnoiter potential weaknesses, such as open
         ports and paths through the network, which enabled it to launch
         an attack later.

      * From the regulator's perspective, the concerns include
         distortion of the measurement tests or alteration of the
         measurement results. Also, a malicious network operator could
         try to identify the broadband lines that the regulator was
         measuring and prioritize that traffic ("game the system").

   3. Another potential issue is a measurement system that does not
      obtain the end user's informed consent, fails to specify a
      specific purpose in the consent, or uses the collected information
      for secondary uses beyond those specified.

   4. Another potential issue is a measurement system that does not
      indicate who is responsible for the collection and processing of
      personal data and who is responsible for fulfilling the rights of
      users. The responsible party (often termed the "data controller")
      should, as good practice, consider such issues as defining:

      o the purpose for which the data is collected and used,

      o how the data is stored, accessed, and processed,

      o how long the data is retained, and

      o how the end user can view, update, and even delete their
         personal data.

      If anonymized personal data is shared with a third party, the data
      controller should consider the possibility that the third party
      can de-anonymize it by combining it with other information.

   These security and privacy issues will need to be considered
   carefully by any measurement system. In the context of LMAP,
   [Framework] considers them further, along with some potential
   mitigations. Other LMAP documents will specify one or more protocols
   that enable the measurement system to instruct a probe about what
   measurements to make and that enable the probe to report the
   measurement results. Those documents will need to discuss solutions
   to the security and privacy issues. However, the protocol documents

will not consider the actual usage of the measurement information.
Many use cases can be envisaged, and earlier in this document we
described some likely ones for the network operator and regulator.

8. Informative References

   [IETF85-Plenary]
            Crawford, S., "Large-Scale Active Measurement of Broadband
            Networks", 'example' from slide 18, November 2012,
            <http://www.ietf.org/proceedings/85/slides/
            slides-85-iesg-opsandtech-7.pdf>.

   [Extend-TCP]
            Honda, M., Nishida, Y., Raiciu, C., Greenhalgh, A.,
            Handley, M., and H. Tokuda, "Is it Still Possible to
            Extend TCP?", Proceedings of IETF 82, November 2011,
            <http://www.ietf.org/proceedings/82/slides/IRTF-1.pdf>.

   [Framework]
            Eardley, P., Morton, A., Bagnulo, M., Burbridge, T.,
            Aitken, P., and A. Akhter, "A framework for Large-Scale
            Measurement of Broadband Performance (LMAP)", Work in
            Progress, draft-ietf-lmap-framework-14, April 2015.

   [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>.

   [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an
            Attack", BCP 188, RFC 7258, May 2014,
            <http://www.rfc-editor.org/info/rfc7258>.

   [FCC-R&O] United States Federal Communications Commission,
            "Preserving the Open Internet; Broadband Industries
            Practices: Report and Order", FCC 10-201, December 2010,
            <http://hraunfoss.fcc.gov/edocs_public/attachmatch/
            FCC-10-201A1.pdf>.

   [BEREC-Guidelines]
            Body of European Regulators for Electronic Communications,
            "BEREC Guidelines for quality of service in the scope of
            net neutrality", <http://berec.europa.eu/eng/
            document_register/subject_matter/berec/download/0/
            1101-berec-guidelines-for-quality-of-service-_0.pdf>.

   [M-Labs_NSDI-2010]
              M-Lab, "Glasnost: Enabling End Users to Detect Traffic
              Differentiation", <http://www.measurementlab.net/
              download/AMIfv945ljiJXzG-fgUrZSTu2hs1xRl5Oh-
              rpGQMWL305BNQh-BSq5oBoYU4a7zqXOvrztpJhK9gwk5unOe-
              fOzj4X-vOQz_HRrnYU-aFd0rv332RDReRfOYkJuagysstN3GZ__lQHTS8_
              UHJTWkrwyqIUjffVeDxQ/>.

   [Glasnost] M-Lab tool "Glasnost", <http://mlab-live.appspot.com/
              tools/glasnost>.

   [MOS]   Wikipedia, "Mean Opinion Score", January 2015,
              <http://en.wikipedia.org/w/index.php?
              title=Mean_opinion_score&oldid=644494161>.

   [DAE]   Digital Agenda for Europe, COM(2010)245 final,
              "Communication from the Commission to the European
              Parliament, the Council, the European Economic and Social
              Committee and the Committee of the Regions",
              <http://eur-lex.europa.eu/legal-content/EN/TXT/
              PDF/?uri=CELEX:52010DC0245&from=EN>.

Contributors

   The information in this document is partially derived from text
   written by the following contributors:

      James Miller   jamesmilleresquire@gmail.com

      Rachel Huang   rachel.huang@huawei.com

Authors' Addresses

   Marc Linsner
   Cisco Systems, Inc.
   Marco Island, FL
   United States

   EMail: mlinsner@cisco.com

   Philip Eardley
   BT
   B54 Room 77, Adastral Park, Martlesham
   Ipswich, IP5 3RE
   United Kingdom

   EMail: philip.eardley@bt.com

   Trevor Burbridge
   BT
   B54 Room 70, Adastral Park, Martlesham
   Ipswich, IP5 3RE
   United Kingdom

   EMail: trevor.burbridge@bt.com

   Frode Sorensen
   Norwegian Communications Authority (Nkom)
   Lillesand
   Norway

   EMail: frode.sorensen@nkom.no