Network Working Group
Internet-Draft
Intended status: Informational
Expires: October 31, 2015

P. Eardley
BT
A. Morton
AT&T Labs
M. Bagnulo
UC3M
T. Burbridge
BT
P. Aitken
Brocade
A. Akhter
Consultant

April 29, 2015

A framework for Large-Scale Measurement of Broadband Performance (LMAP)
draft-ietf-lmap-framework-14

Abstract

Measuring broadband service on a large scale requires a description of the logical architecture and standardisation of the key protocols that coordinate interactions between the components. The document presents an overall framework for large-scale measurements. It also defines terminology for LMAP (Large-Scale Measurement of Broadband Performance).

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 October 31, 2015.

Table of Contents

1. Introduction

There is a desire to be able to coordinate the execution of broadband measurements and the collection of measurement results across a large scale set of Measurement Agents (MAs). These MAs could be software based agents on PCs, embedded agents in consumer devices (such as TVs or gaming consoles), embedded in service provider controlled devices such as set-top boxes and home gateways, or simply dedicated probes. MAs may also be embedded on a device that is part of an ISP's network, such as a DSLAM (Digital Subscriber Line Access Multiplexer), router, Carrier Grade NAT (Network Address Translator) or ISP Gateway. It is expected that a measurement system could easily encompass a few hundred thousand or even millions of such MAs. Such a scale presents unique problems in coordination, execution and measurement result collection. Several use cases have been proposed for large-scale measurements including:

- Operators: to help plan their network and identify faults
- Regulators: to benchmark several network operators and support public policy development

Further details of the use cases can be found in [I-D.ietf-lmap-use-cases]. The LMAP framework should be useful for these, as well as other use cases, such as to help end users run diagnostic checks like a network speed test.

The LMAP Framework has three basic elements: Measurement Agents, Controllers and Collectors.

Measurement Agents (MAs) initiate the actual measurements, which are called Measurement Tasks in the LMAP terminology. In principle, there are no restrictions on the type of device in which the MA function resides.

The Controller instructs one or more MAs and communicates the set of Measurement Tasks an MA should perform and when. For example it may instruct a MA at a home gateway: "Measure the 'UDP latency' with www.example.org; repeat every hour at xx.05". The Controller also manages a MA by instructing it how to report the Measurement Results, for example: "Report results once a day in a batch at 4am". We refer to these as the Measurement Schedule and Report Schedule.

The Collector accepts Reports from the MAs with the Results from their Measurement Tasks. Therefore the MA is a device that gets Instructions from the Controller, initiates the Measurement Tasks, and reports to the Collector. The communications between these three

LMAP functions are structured according to a Control Protocol and a Report Protocol.

The desirable features for a large-scale Measurement Systems we are designing for are:

- Standardised - in terms of the Measurement Tasks that they perform, the components, the data models and protocols for transferring information between the components. Amongst other things, standardisation enables meaningful comparisons of measurements made of the same metric at different times and places, and provides the operator of a Measurement System with criteria for evaluation of the different solutions that can be used for various purposes including buying decisions (such as buying the various components from different vendors). Today's systems are proprietary in some or all of these aspects.

- Large-scale - [I-D.ietf-lmap-use-cases] envisages Measurement Agents in every home gateway and edge device such as set-top boxes and tablet computers, and located throughout the Internet as well [RFC7398]. It is expected that a Measurement System could easily encompass a few hundred thousand or even millions of Measurement Agents. Existing systems have up to a few thousand MAs (without judging how much further they could scale).

- Diversity - a Measurement System should handle Measurement Agents from different vendors, that are in wired and wireless networks, can execute different sorts of Measurement Task, are on devices with IPv4 or IPv6 addresses, and so on.

- Privacy Respecting - the protocols and procedures should respect the sensitive information of all those involved in measurements.

2. Outline of an LMAP-based measurement system

In this section we provide an overview of the whole Measurement System. New LMAP-specific terms are capitalised; Section 3 provides a terminology section with a compilation of all the LMAP terms and their definition. Section 4 onwards considers the LMAP components in more detail.

Other LMAP specifications will define an information model, the associated data models, and select/extend one or more protocols for the secure communication: firstly, a Control Protocol, from a Controller to instruct Measurement Agents what performance metrics to measure, when to measure them, how/when to report the measurement results to a Collector; secondly, a Report Protocol, for a Measurement Agent to report the results to the Collector.

The Figure below shows the main components of a Measurement System, and the interactions of those components. Some of the components are outside the scope of initial LMAP work.

The MA performs Measurement Tasks. One possibility is that the MA is observes existing traffic. Another possibility is for the MA to generate (or receive) traffic specially created for the purpose and measure some metric associated with its transfer. The Figure includes both possibilities (in practice, it may be more usual for a MA to do one) whilst Section 6.4 shows some examples of possible arrangements of the components.

The MAs are pieces of code that can be executed in specialised hardware (hardware probe) or on a general-purpose device (like a PC or mobile phone). A device with a Measurement Agent may have multiple physical interfaces (Wi-Fi, Ethernet, DSL (Digital Subscriber Line); and non-physical interfaces such as PPPoE (Point-to-Point Protocol over Ethernet) or IPsec) and the Measurement Tasks may specify any one of these.

The Controller manages a MA through use of the Control Protocol, which transfers the Instruction to the MA. This describes the Measurement Tasks the MA should perform and when. For example the Controller may instruct a MA at a home gateway: "Count the number of TCP SYN packets observed in a 1 minute interval; repeat every hour at xx.05 + Unif[0,180] seconds". The Measurement Schedule determines when the Measurement Tasks are executed. The Controller also manages a MA by instructing it how to report the Measurement Results, for example: "Report results once a day in a batch at 4am + Unif[0,180] seconds; if the end user is active then delay the report 5 minutes". The Report Schedule determines when the Reports are uploaded to the Collector. The Measurement Schedule and Report Schedule can define one-off (non-recurring) actions ("Do measurement now", "Report as soon as possible"), as well as recurring ones.

The Collector accepts a Report from a MA with the Measurement Results from its Measurement Tasks. It then provides the Results to a repository (see below).

A Measurement Method defines how to measure a Metric of interest. It is very useful to standardise Measurement Methods, so that it is meaningful to compare measurements of the same Metric made at different times and places. It is also useful to define a registry for commonly-used Metrics [I-D.ietf-ippm-metric-registry] so that a Metric with its associated Measurement Method can be referred to simply by its identifier in the registry. The registry will hopefully be referenced by other standards organisations. The

Measurement Methods may be defined by the IETF, locally, or by some other standards body.

Broadly speaking there are two types of Measurement Method. In both types a Measurement Agent measures a particular Observed Traffic Flow. It may involve a single MA simply observing existing traffic - for example, the Measurement Agent could count bytes or calculate the average loss for a particular flow. On the other hand, a Measurement Method may involve multiple network entities, which perform different roles. For example, a "ping" Measurement Method, to measure the round trip delay , would consist of an MA sending an ICMP (Internet Control Message Protocol) ECHO request to a responder in the Internet. In LMAP terms, the responder is termed a Measurement Peer (MP), meaning that it helps the MA but is not managed by the Controller. Other Measurement Methods involve a second MA, with the Controller instructing the MAs in a coordinated manner. Traffic generated specifically as part of the Measurement Method is termed Measurement Traffic; in the ping example, it is the ICMP ECHO Requests and Replies. The protocols used for the Measurement Traffic are out of the scope of initial LMAP work, and fall within the scope of other IETF WGs such as IPPM (IP Performance Metrics).

A Measurement Task is the action performed by a particular MA at a particular time, as the specific instance of its role in a Measurement Method. LMAP is mainly concerned with Measurement Tasks, for instance in terms of its Information Model and Protocols.

For Measurement Results to be truly comparable, as might be required by a regulator, not only do the same Measurement Methods need to be used to assess Metrics, but also the set of Measurement Tasks should follow a similar Measurement Schedule and be of similar number. The details of such a characterisation plan are beyond the scope of work in IETF although certainly facilitated by IETF’s work.

Both control and report messages are transferred over a secure Channel. A Control Channel is between the Controller and a MA; the Control Protocol delivers Instruction Messages to the MA and Capabilities, Failure and Logging Information in the reverse direction. A Report Channel is between a MA and Collector, and the Report Protocol delivers Reports to the Collector.

Finally we introduce several components that are outside the scope of initial LMAP work and will be provided through existing protocols or applications. They affect how the Measurement System uses the Measurement Results and how it decides what set of Measurement Tasks to perform. As shown in the Figure, these components are: the bootstrapper, Subscriber parameter database, data analysis tools, and Results repository.

The MA needs to be bootstrapped with initial details about its Controller, including authentication credentials. The LMAP work considers the bootstrap process, since it affects the Information Model. However, LMAP does not define a bootstrap protocol, since it is likely to be technology specific and could be defined by the Broadband Forum, CableLabs or IEEE depending on the device. Possible protocols are SNMP (Simple Network Management Protocol), NETCONF (Network Configuration Protocol) or (for Home Gateways) CPE WAN Management Protocol (CWMP) from the Auto Configuration Server (ACS) (as specified in TR-069 [TR-069]).

A Subscriber parameter database contains information about the line, such as the customer's broadband contract (perhaps 2, 40 or 80Mb/s), the line technology (DSL or fibre), the time zone where the MA is located, and the type of home gateway and MA. These parameters are already gathered and stored by existing operations systems. They may affect the choice of what Measurement Tasks to run and how to interpret the Measurement Results. For example, a download test suitable for a line with an 80Mb/s contract may overwhelm a 2Mb/s line.

A Results repository records all Measurement Results in an equivalent form, for example an SQL (Structured Query Language) database, so that they can easily be accessed by the data analysis tools.

The data analysis tools receive the results from the Collector or via the Results repository. They might visualise the data or identify which component or link is likely to be the cause of a fault or degradation. This information could help the Controller decide what follow-up Measurement Task to perform in order to diagnose a fault. The data analysis tools also need to understand the Subscriber's service information, for example the broadband contract.

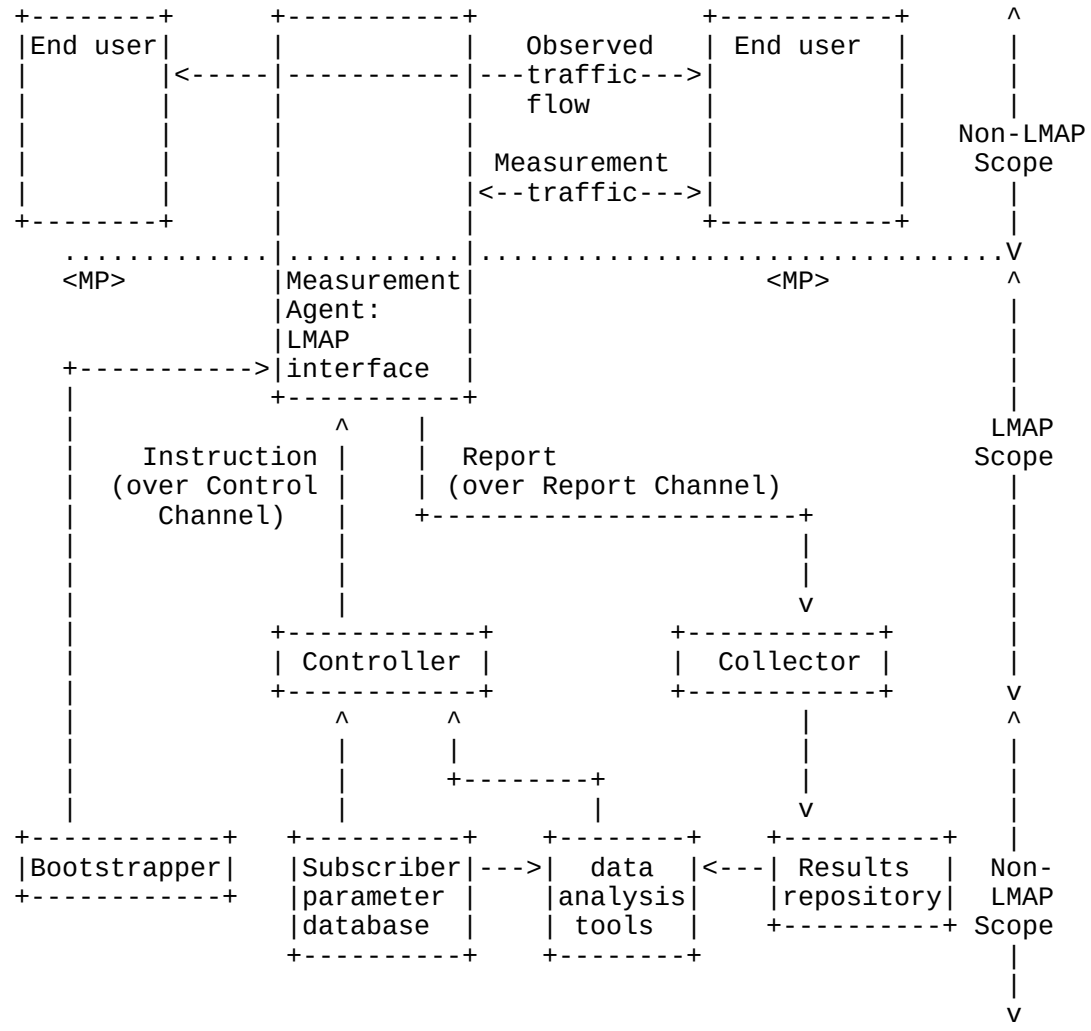

Schematic of main elements of an LMAP-based Measurement System (showing the elements in and out of the scope of initial LMAP work)

3. Terminology

This section defines terminology for LMAP. Please note that defined terms are capitalized.

Bootstrap: A process that integrates a Measurement Agent into a Measurement System.

Capabilities: Information about the performance measurement capabilities of the MA, in particular the Measurement Method roles and measurement protocol roles that it can perform, and the device hosting the MA, for example its interface type and speed, but not dynamic information.

Channel: A bi-directional logical connection that is defined by a specific Controller and MA, or Collector and MA, plus associated security.

Collector: A function that receives a Report from a Measurement Agent.

Configuration: A process for informing the MA about its MA-ID, (optional) Group-ID and Control Channel.

Controller: A function that provides a Measurement Agent with its Instruction.

Control Channel: A Channel between a Controller and a MA over which Instruction Messages and Capabilities, Failure and Logging Information are sent.

Control Protocol: The protocol delivering Instruction(s) from a Controller to a Measurement Agent. It also delivers Capabilities, Failure and Logging Information from the Measurement Agent to the Controller. It can also be used to update the MA's Configuration. It runs over the Control Channel.

Cycle-ID: A tag that is sent by the Controller in an Instruction and echoed by the MA in its Report. The same Cycle-ID is used by several MAs that use the same Measurement Method for a Metric with the same Input Parameters. Hence the Cycle-ID allows the Collector to easily identify Measurement Results that should be comparable.

Data Model: The implementation of an Information Model in a particular data modelling language [RFC3444].

Environmental Constraint: A parameter that is measured as part of the Measurement Task, its value determining whether the rest of the Measurement Task proceeds.

Failure Information: Information about the MA's failure to action or execute an Instruction, whether concerning Measurement Tasks or Reporting.

Group-ID: An identifier of a group of MAs.

Information Model: The protocol-neutral definition of the semantics of the Instructions, the Report, the status of the different elements of the Measurement System as well of the events in the system [RFC3444].

Input Parameter: A parameter whose value is left open by the Metric and its Measurement Method and is set to a specific value in a Measurement Task. Altering the value of an Input Parameter does not change the fundamental nature of the Measurement Task.

Instruction: The description of Measurement Tasks for a MA to perform and the details of the Report for it to send. It is the collective description of the Measurement Task configurations, the configuration of the Measurement Schedules, the configuration of the Report Channel(s), the configuration of Report Schedule(s), and the details of any suppression.

Instruction Message: The message that carries an Instruction from a Controller to a Measurement Agent.

Logging Information: Information about the operation of the Measurement Agent, which may be useful for debugging.

Measurement Agent (MA): The function that receives Instruction Messages from a Controller and operates the Instruction by executing Measurement Tasks (using protocols outside the initial LMAP work scope and perhaps in concert with one or more other Measurement Agents or Measurement Peers) and (if part of the Instruction) by reporting Measurement Results to a Collector or Collectors.

Measurement Agent Identifier (MA-ID): a UUID [RFC4122] that identifies a particular MA and is configured as part of the Bootstrapping process.

Measurement Method: The process for assessing the value of a Metric; the process of measuring some performance or reliability parameter associated with the transfer of traffic.

Measurement Peer (MP): The function that assists a Measurement Agent with Measurement Tasks and does not have an interface to the Controller or Collector.

Measurement Result: The output of a single Measurement Task (the value obtained for the parameter of interest or Metric).

Measurement Schedule: The schedule for performing Measurement Tasks.

Measurement System: The set of LMAP-defined and related components that are operated by a single organisation, for the purpose of measuring performance aspects of the network.

Measurement Task: The action performed by a particular Measurement Agent that consists of the single assessment of a Metric through operation of a Measurement Method role at a particular time, with all of the role's Input Parameters set to specific values.

Measurement Traffic: the packet(s) generated by some types of Measurement Method that involve measuring some parameter associated with the transfer of the packet(s).

Metric: The quantity related to the performance and reliability of the network that we'd like to know the value of.

Observed Traffic Flow: In RFC 7011, a Traffic Flow (or Flow) is defined as a set of packets or frames passing an Observation Point in the network during a certain time interval. All packets belonging to a particular Flow have a set of common properties, such as packet header fields, characteristics, and treatments. A Flow measured by the LMAP system is termed an Observed Traffic Flow. Its properties are summarized and tabulated in Measurement Results (as opposed to raw capture and export).

Report: The set of Measurement Results and other associated information (as defined by the Instruction). The Report is sent by a Measurement Agent to a Collector.

Report Channel: A Channel between a Collector and a MA over which Report messages are sent.

Report Protocol: The protocol delivering Report(s) from a Measurement Agent to a Collector. It runs over the Report Channel.

Report Schedule: the schedule for sending Reports to a Collector.

Subscriber: An entity (associated with one or more users) that is engaged in a subscription with a service provider.

Suppression: the temporary cessation of Measurement Tasks.

4. Constraints

The LMAP framework makes some important assumptions, which constrain the scope of the initial LMAP work.

4.1. The measurement system is under the direction of a single organisation

In the LMAP framework, the Measurement System is under the direction of a single organisation that is responsible for any impact that its measurements have on a user's quality of experience and privacy. Clear responsibility is critical given that a misbehaving large-scale Measurement System could potentially harm user experience, user privacy and network security.

However, the components of an LMAP Measurement System can be deployed in administrative domains that are not owned by the measuring organisation. Thus, the system of functions deployed by a single organisation constitutes a single LMAP domain which may span ownership or other administrative boundaries.

4.2. Each MA may only have a single Controller at any point in time

A MA is instructed by one Controller and is in one Measurement System. The constraint avoids different Controllers giving a MA conflicting instructions and so means that the MA does not have to manage contention between multiple Measurement (or Report) Schedules. This simplifies the design of MAs (critical for a large-scale infrastructure) and allows a Measurement Schedule to be tested on specific types of MA before deployment to ensure that the end user experience is not impacted (due to CPU, memory or broadband-product constraints). However, a Measurement System may have several Controllers.

5. Protocol Model

A protocol model [RFC4101] presents an architectural model for how the protocol operates and needs to answer three basic questions:

1. What problem is the protocol trying to address?

2. What messages are being transmitted and what do they mean?

3. What are the important, but unobvious, features of the protocol?

An LMAP system goes through the following phases:

- a Bootstrapping process before the MA can take part in the other three phases.

- a Control Protocol, which delivers Instruction Messages from a Controller to a MA (amongst other things).

- the actual Measurement Tasks, which measure some performance or reliability parameter(s) associated with the transfer of packets.

- a Report Protocol, which delivers Reports containing the Measurement Results from a MA to a Collector.

The diagrams show the various LMAP messages and uses the following convention:

- (optional): indicated by round brackets

- [potentially repeated]: indicated by square brackets

The protocol model is closely related to the Information Model [I-D.ietf-lmap-information-model], which is the abstract definition of the information carried by the protocol. (If there is any difference between this document and the Information Model, the latter is definitive, since it is on the standards track.) The purpose of both is to provide a protocol and device independent view, which can be implemented via specific protocols. LMAP defines a specific Control Protocol and Report Protocol, but others could be defined by other standards bodies or be proprietary. However it is important that they all implement the same Information Model and protocol model, in order to ease the definition, operation and interoperability of large-scale Measurement Systems.

5.1. Bootstrapping process

The primary purpose of bootstrapping is to enable a MA to be integrated into a Measurement System. The MA retrieves information about itself (like its identity in the Measurement System) and about the Controller, the Controller learns information about the MA, and they learn about security information to communicate (such as certificates and credentials).

Whilst this memo considers the bootstrapping process, it is beyond the scope of initial LMAP work to define a bootstrap mechanism, as it depends on the type of device and access.

As a result of the bootstrapping process the MA learns information with the following aims ([I-D.ietf-lmap-information-model] defines the consequent list of information elements):

- its identifier, either its MA-ID or a device identifier such as one of its MAC or both.

- (optionally) a Group-ID. A Group-ID would be shared by several MAs and could be useful for privacy reasons. For instance,

reporting the Group-ID and not the MA-ID could hinder tracking of a mobile device

- the Control Channel, which is defined by:
  - the address which identifies the Control Channel, such as the Controller's FQDN (Fully Qualified Domain Name) [RFC1035])
  - security information (for example to enable the MA to decrypt the Instruction Message and encrypt messages sent to the Controller)

The details of the bootstrapping process are device /access specific. For example, the information could be in the firmware, manually configured or transferred via a protocol like TR-069 [TR-069]. There may be a multi-stage process where the MA contacts a 'hard-coded' address, which replies with the bootstrapping information.

The MA must learn its MA-ID before getting an Instruction, either during Bootstrapping or via Configuration (Section 5.2.1).

5.2. Control Protocol

The primary purpose of the Control Protocol is to allow the Controller to configure a Measurement Agent with an Instruction about what Measurement Tasks to do, when to do them, and how to report the Measurement Results (Section 5.2.2). The Measurement Agent then acts on the Instruction autonomously. The Control Protocol also enables the MA to inform the Controller about its Capabilities and any Failure and Logging Information (Section 5.2.2). Finally, the Control Protocol allows the Controller to update the MA's Configuration.

5.2.1. Configuration

Configuration allows the Controller to update the MA about some or all of the information that it obtained during the bootstrapping process: the MA-ID, the (optional) Group-ID and the Control Channel. The Measurement System might use Configuration for several reasons. For example, the bootstrapping process could 'hard code' the MA with details of an initial Controller, and then the initial Controller could configure the MA with details about the Controller that sends Instruction Messages. (Note that a MA only has one Control Channel, and so is associated with only one Controller, at any moment.)

Note that an implementation may choose to combine Configuration information and an Instruction Message into a single message.

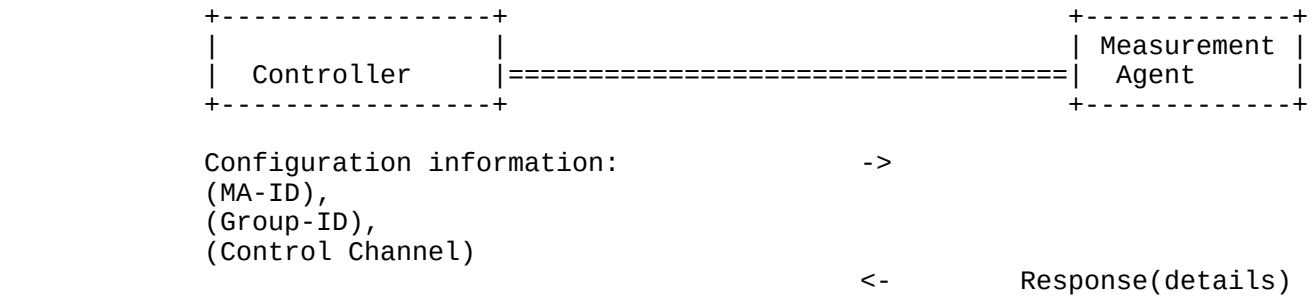

5.2.2. Instruction

The Instruction is the description of the Measurement Tasks for a Measurement Agent to do and the details of the Measurement Reports for it to send. In order to update the Instruction the Controller uses the Control Protocol to send an Instruction Message over the Control Channel.

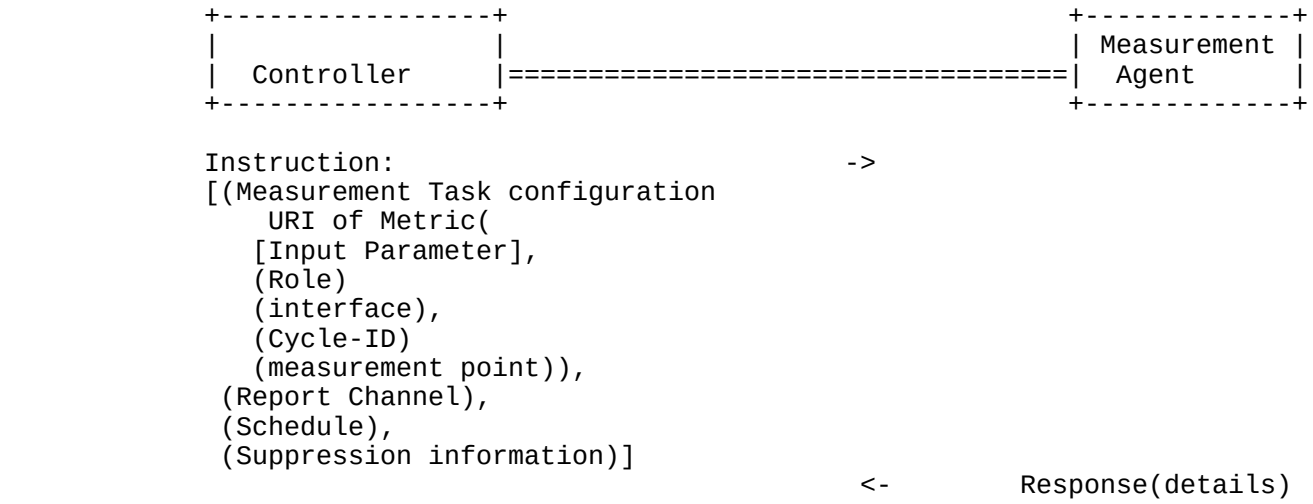

The Instruction defines information with the following aims ([I-D.ietf-lmap-information-model] defines the consequent list of information elements):

- the Measurement Task configurations, each of which needs:
  - the Metric, specified as a URI to a registry entry; it includes the specification of a Measurement Method. The registry could

be defined by a standards organisation or locally by the operator of the Measurement System. Note that, at the time of writing, the IETF works on such a registry specification [I-D.ietf-ippm-metric-registry].

* the Measurement Method role. For some Measurement Methods, different parties play different roles; for example (see Section 6.4) an iperf sender and receiver. Each Metric and its associated Measurement Method will describe all measurement roles involved in the process.

* a boolean flag (suppress or do-not-suppress) indicating if such a Measurement Task is impacted by a Suppression message (see Section 5.2.2.1). Thus, the flag is an Input Parameter.

* any Input Parameters that need to be set for the Metric and the Measurement Method. For example, the address of a Measurement Peer (or other Measurement Agent) that may be involved in a Measurement Task , or traffic filters associated with the Observed Traffic Flow.

* if the device with the MA has multiple interfaces, then the interface to use (if not defined, then the default interface is used).

* optionally, a Cycle-ID.

* optionally, the measurement point designation [RFC7398] of the MA and, if applicable, of the MP or other MA. This can be useful for reporting.

- configuration of the Schedules, each of which needs:

  * the timing of when the Measurement Tasks are to be performed, or the Measurement Reports are to be sent. Possible types of timing are periodic, calendar-based periodic, one-off immediate and one-off at a future time

- configuration of the Report Channel(s), each of which needs:

  * the address of the Collector, for instance its URL

  * security for this Report Channel, for example the X.509 certificate

- Suppression information, if any (see Section 5.2.1.1)

A single Instruction Message may contain some or all of the above parts. The finest level of granularity possible in an Instruction Message is determined by the implementation and operation of the Control Protocol. For example, a single Instruction Message may add or update an individual Measurement Schedule - or it may only update the complete set of Measurement Schedules; a single Instruction Message may update both Measurement Schedules and Measurement Task configurations - or only one at a time; and so on. However, Suppression information always replaces (rather than adds to) any previous Suppression information.

The MA informs the Controller that it has successfully understood the Instruction Message, or that it cannot action the Instruction - for example, if it doesn't include a parameter that is mandatory for the requested Metric and Measurement Method, or it is missing details of the target Collector.

The Instruction Message instructs the MA; the Control Protocol does not allow the MA to negotiate, as this would add complexity to the MA, Controller and Control Protocol for little benefit.

5.2.2.1. Suppression

The Instruction may include Suppression information. The main motivation for Suppression is to enable the Measurement System to eliminate Measurement Traffic, because there is some unexpected network issue for example. There may be other circumstances when Suppression is useful, for example to eliminate inessential Reporting traffic (even if there is no Measurement Traffic).

The Suppression information may include any of the following optional fields:

- a set of Measurement Tasks to suppress; the others are not suppressed. For example, this could be useful if a particular Measurement Task is overloading a Measurement Peer with Measurement Traffic.

- a set of Measurement Schedules to suppress; the others are not suppressed. For example, suppose the Measurement System has defined two Schedules, one with the most critical Measurement Tasks and the other with less critical ones that create a lot of Measurement Traffic, then it may only want to suppress the second.

- a set of Reporting Schedules to suppress; the others are not suppressed. This can be particularly useful in the case of a Measurement Method that doesn't generate Measurement Traffic; it

may need to continue observing traffic flows but temporarily suppress Reports due to the network footprint of the Reports.

- if all the previous fields are included then the MA suppresses the union - in other words, it suppresses the set of Measurement Tasks, the set of Measurement Schedules, and the set of Reporting Schedules.

- if the Suppression information includes neither a set of Measurement Tasks nor a set of Measurement Schedules, then the MA does not begin new Measurement Tasks that have the boolean flag set to "suppress"; however, the MA does begin new Measurement Tasks that have the flag set to "do-not-suppress".

- a start time, at which suppression begins. If absent, then Suppression begins immediately.

- an end time, at which suppression ends. If absent, then Suppression continues until the MA receives an un-Suppress message.

- a demand that the MA immediately ends on-going Measurement Task(s) that are tagged for suppression. (Most likely it is appropriate to delete the associated partial Measurement Result(s).) This could be useful in the case of a network emergency so that the operator can eliminate all inessential traffic as rapidly as possible. If absent, the MA completes on-going Measurement Tasks.

An un-Suppress message instructs the MA no longer to suppress, meaning that the MA once again begins new Measurement Tasks, according to its Measurement Schedule.

Note that Suppression is not intended to permanently stop a Measurement Task (instead, the Controller should send a new Measurement Schedule), nor to permanently disable a MA (instead, some kind of management action is suggested).

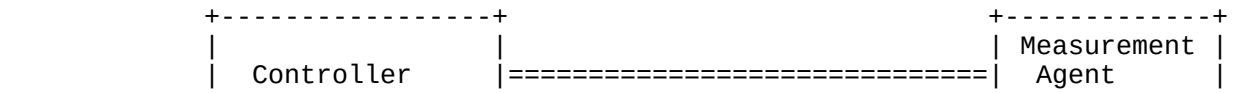


```
Suppress:
[(Measurement Task),                  ->
 (Measurement Schedule),
 [start time],
 [end time],
 [on-going suppressed?]]

Un-suppress                           ->
```

5.2.3. Capabilities, Failure and Logging Information

The Control Protocol also enables the MA to inform the Controller about various information, such as its Capabilities and any Failures. It is also possible to use a device-specific mechanism which is beyond the scope of the initial LMAP work.

Capabilities are information about the MA that the Controller needs to know in order to correctly instruct the MA, such as:

- the Measurement Method (roles) that the MA supports
- the measurement protocol types and roles that the MA supports
- the interfaces that the MA has
- the version of the MA
- the version of the hardware, firmware or software of the device with the MA
- its Instruction (this could be useful if the Controller thinks something has gone wrong, and wants to check what Instruction the MA is using)
- but not dynamic information like the currently unused CPU, memory or battery life of the device with the MA.

Failure Information concerns why the MA has been unable to execute a Measurement Task or deliver a Report, for example:

- the Measurement Task failed to run properly because the MA (unexpectedly) has no spare CPU cycles

- the MA failed to record the Measurement Results because it (unexpectedly) is out of spare memory
- a Report failed to deliver Measurement Results because the Collector (unexpectedly) is not responding
- but not if a Measurement Task correctly doesn't start. For example, the first step of some Measurement Methods is for the MA to check there is no cross-traffic.

Logging Information concerns how the MA is operating and may help debugging, for example:

- the last time the MA ran a Measurement Task
- the last time the MA sent a Measurement Report
- the last time the MA received an Instruction Message
- whether the MA is currently Suppressing Measurement Tasks

Capabilities, Failure and Logging Information are sent by the MA, either in response to a request from the Controller (for example, if the Controller forgets what the MA can do or otherwise wants to resynchronize what it knows about the MA), or on its own initiative (for example when the MA first communicates with a Controller or if it becomes capable of a new Measurement Method). Another example of the latter case is if the device with the MA re-boots, then the MA should notify its Controller in case its Instruction needs to be updated; to avoid a "mass calling event" after a widespread power restoration affecting many MAs, it is sensible for an MA to pause for a random delay, perhaps in the range of one minute or so.

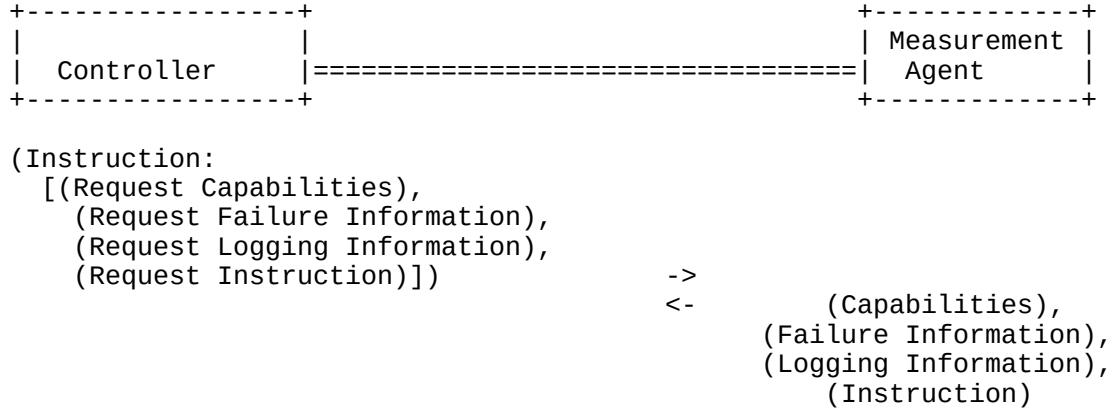

5.3. Operation of Measurement Tasks

This LMAP framework is neutral to what the actual Measurement Task is. It does not define Metrics and Measurement Methods, these are defined elsewhere.

The MA carries out the Measurement Tasks as instructed, unless it gets an updated Instruction. The MA acts autonomously, in terms of operation of the Measurement Tasks and reporting of the Results; it doesn't do a 'safety check' with the Controller to ask whether it should still continue with the requested Measurement Tasks.

The MA may operate Measurement Tasks sequentially or in parallel (see Section 5.3.2).

5.3.1. Starting and Stopping Measurement Tasks

This LMAP framework does not define a generic start and stop process, since the correct approach depends on the particular Measurement Task; the details are defined as part of each Measurement Method. This section provides some general hints. The MA does not inform the Controller about Measurement Tasks starting and stopping.

Before beginning a Measurement Task the MA may want to run a pre-check. (The pre-check could be defined as a separate, preceding Task or as the first part of a larger Task.)

For Measurement Tasks that observe existing traffic, action could include:

- checking that there is traffic of interest;
- checking that the device with the MA has enough resources to execute the Measurement Task reliably. Note that the designer of the Measurement System should ensure that the device's capabilities are normally sufficient to comfortably operate the Measurement Tasks.

For Measurement Tasks that generate Measurement Traffic, a pre-check could include:

- the MA checking that there is no cross-traffic. In other words, a check that the end-user isn't already sending traffic;
- the MA checking with the Measurement Peer (or other Measurement Agent) involved in the Measurement Task that it can handle a new Measurement Task. For example, the Measurement Peer may already be handling many Measurement Tasks with other MAs;

- sending traffic that probes the path to check it isn't overloaded;
- checking that the device with the MA has enough resources to execute the Measurement Task reliably.

It is possible that similar checks continue during the Measurement Task, especially one that is long-running and/or creates a lot of Measurement Traffic, and might lead to it being abandoned whilst in-progress. A Measurement Task could also be abandoned in response to a "suppress" message (see Section 5.2.1). Action could include:

- For 'upload' tests, the MA not sending traffic
- For 'download' tests, the MA closing the TCP connection or sending a TWAMP (Two-Way Active Measurement Protocol) Stop control message [RFC5357].

The Controller may want a MA to run the same Measurement Task indefinitely (for example, "run the 'upload speed' Measurement Task once an hour until further notice"). To avoid the MA generating traffic forever after a Controller has permanently failed (or communications with the Controller have failed), the MA can be configured with a time limit; if the MA doesn't hear from the Controller for this length of time, then it stops operating Measurement Tasks.

5.3.2. Overlapping Measurement Tasks

It is possible that a MA starts a new Measurement Task before another Measurement Task has completed. This may be intentional (the way that the Measurement System has designed the Measurement Schedules), but it could also be unintentional - for instance, if a Measurement Task has a 'wait for X' step which pauses for an unexpectedly long time. This document makes no assumptions about the impact of one Measurement Task on another.

The operator of the Measurement System can handle (or not) overlapping Measurement Tasks in any way they choose - it is a policy or implementation issue and not the concern of LMAP. Some possible approaches are: to configure the MA not to begin the second Measurement Task; to start the second Measurement Task as usual; for the action to be an Input Parameter of the Measurement Task; and so on.

It may be important to include in the Measurement Report the fact that the Measurement Task overlapped with another.

5.4. Report Protocol

The primary purpose of the Report Protocol is to allow a Measurement Agent to report its Measurement Results to a Collector, along with the context in which they were obtained.

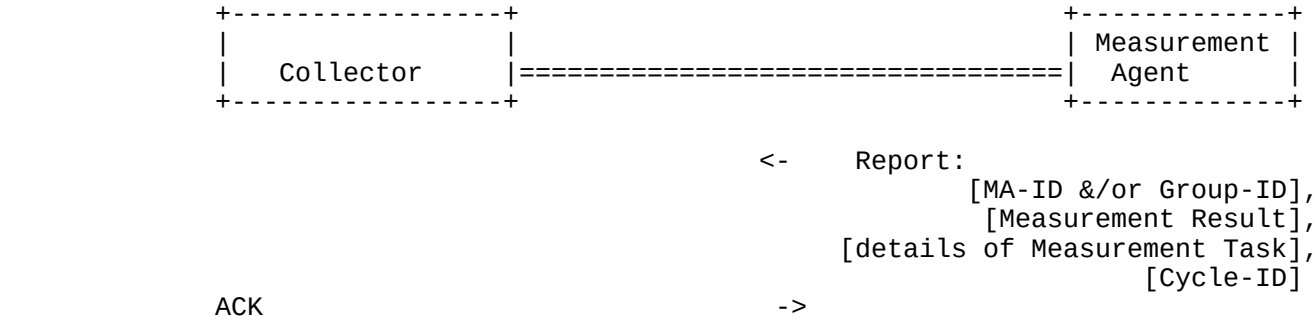

The Report contains:

- the MA-ID or a Group-ID (to anonymise results)
- the actual Measurement Results, including the time they were measured. In general the time is simply the MA's best estimate and there is no guarantee on the accuracy or granularity of the information. It is possible that some specific analysis of a particular Measurement Method's Results will impose timing requirements.
- the details of the Measurement Task (to avoid the Collector having to ask the Controller for this information later). For example, the interface used for the measurements.
- the Cycle-ID, if one was included in the Instruction.
- perhaps the Subscriber's service parameters (see Section 5.4.1).
- the measurement point designation of the MA and, if applicable, the MP or other MA, if the information was included in the Instruction. This numbering system is defined in [RFC7398] and allows a Measurement Report to describe abstractly the path measured (for example, "from a MA at a home gateway to a MA at a DSLAM"). Also, the MA can anonymise results by including measurement point designations instead of IP addresses (Section 8.6.2).

The MA sends Reports as defined by the Instruction. It is possible that the Instruction tells the MA to report the same Results to more than one Collector, or to report a different subset of Results to different Collectors. It is also possible that a Measurement Task may create two (or more) Measurement Results, which could be reported differently (for example, one Result could be reported periodically, whilst the second Result could be an alarm that is created as soon as the measured value of the Metric crosses a threshold and that is reported immediately).

Optionally, a Report is not sent when there are no Measurement Results.

In the initial LMAP Information Model and Report Protocol, for simplicity we assume that all Measurement Results are reported as-is, but allow extensibility so that a Measurement System (or perhaps a second phase of LMAP) could allow a MA to:

- label, or perhaps not include, Measurement Results impacted by, for instance, cross-traffic or a Measurement Peer (or other Measurement Agent) being busy
- label Measurement Results obtained by a Measurement Task that overlapped with another
- not report the Measurement Results if the MA believes that they are invalid
- detail when Suppression started and ended

As discussed in Section 6.1, data analysis of the results should carefully consider potential bias from any Measurement Results that are not reported, or from Measurement Results that are reported but may be invalid.

5.4.1. Reporting of Subscriber’s service parameters

The Subscriber’s service parameters are information about his/her broadband contract, line rate and so on. Such information is likely to be needed to help analyse the Measurement Results, for example to help decide whether the measured download speed is reasonable.

The information could be transferred directly from the Subscriber parameter database to the data analysis tools. If the subscriber’s service parameters are available to the MAs, they could be reported with the Measurement Results in the Report Protocol. How (and if) the MA knows such information is likely to depend on the device type.

The MA could either include the information in a Measurement Report or separately.

5.5. Operation of LMAP over the underlying packet transfer mechanism

The above sections have described LMAP's protocol model. Other specifications will define the actual Control and Report Protocols, possibly operating over an existing protocol, such as REST-style HTTP(S). It is also possible that a different choice is made for the Control and Report Protocols, for example NETCONF-YANG [RFC6241] and IPFIX (Internet Protocol Flow Information Export) [RFC7011] respectively.

From an LMAP perspective, the Controller needs to know that the MA has received the Instruction Message, or at least that it needs to be re-sent as it may have failed to be delivered. Similarly the MA needs to know about the delivery of Capabilities and Failure information to the Controller and Reports to the Collector. How this is done depends on the design of the Control and Report Protocols and the underlying packet transfer mechanism.

For the Control Protocol, the underlying packet transfer mechanism could be:

- a 'push' protocol (that is, from the Controller to the MA)
- a multicast protocol (from the Controller to a group of MAs)
- a 'pull' protocol. The MA periodically checks with Controller if the Instruction has changed and pulls a new Instruction if necessary. A pull protocol seems attractive for a MA behind a NAT or firewall (as is typical for a MA on an end-user's device), so that it can initiate the communications. It also seems attractive for a MA on a mobile device, where the Controller might not know how to reach the MA. A pull mechanism is likely to require the MA to be configured with how frequently it should check in with the Controller, and perhaps what it should do if the Controller is unreachable after a certain number of attempts.
- a hybrid protocol. In addition to a pull protocol, the Controller can also push an alert to the MA that it should immediately pull a new Instruction.

For the Report Protocol, the underlying packet transfer mechanism could be:

- a 'push' protocol (that is, from the MA to the Collector)

- o perhaps supplemented by the ability for the Collector to 'pull' Measurement Results from a MA.

5.6. Items beyond the scope of the initial LMAP work

There are several potential interactions between LMAP elements that are beyond the scope of the initial LMAP work:

1. It does not define a coordination process between MAs. Whilst a Measurement System may define coordinated Measurement Schedules across its various MAs, there is no direct coordination between MAs.

2. It does not define interactions between the Collector and Controller. It is quite likely that there will be such interactions, optionally intermediated by the data analysis tools. For example, if there is an "interesting" Measurement Result then the Measurement System may want to trigger extra Measurement Tasks that explore the potential cause in more detail; or if the Collector unexpectedly does not hear from a MA, then the Measurement System may want to trigger the Controller to send a fresh Instruction Message to the MA.

3. It does not define coordination between different Measurement Systems. For example, it does not define the interaction of a MA in one Measurement System with a Controller or Collector in a different Measurement System. Whilst it is likely that the Control and Report Protocols could be re-used or adapted for this scenario, any form of coordination between different organisations involves difficult commercial and technical issues and so, given the novelty of large-scale measurement efforts, any form of inter-organisation coordination is outside the scope of the initial LMAP work. Note that a single MA is instructed by a single Controller and is only in one Measurement System.

   * An interesting scenario is where a home contains two independent MAs, for example one controlled by a regulator and one controlled by an ISP. Then the Measurement Traffic of one MA is treated by the other MA just like any other end-user traffic.

4. It does not consider how to prevent a malicious party "gaming the system". For example, where a regulator is running a Measurement System in order to benchmark operators, a malicious operator could try to identify the broadband lines that the regulator was measuring and prioritise that traffic. It is assumed this is a policy issue and would be dealt with through a code of conduct for instance.

5. It does not define how to analyse Measurement Results, including how to interpret missing Results.

6. It does not specifically define a end-user-controlled Measurement System, see sub-section 5.6.1.

5.6.1. End-user-controlled measurement system

This framework concentrates on the cases where an ISP or a regulator runs the Measurement System. However, we expect that LMAP functionality will also be used in the context of an end-user-controlled Measurement System. There are at least two ways this could happen (they have various pros and cons):

1. an end-user could somehow request the ISP- (or regulator-) run Measurement System to test his/her line. The ISP (or regulator) Controller would then send an Instruction to the MA in the usual LMAP way.

2. an end-user could deploy their own Measurement System, with their own MA, Controller and Collector. For example, the user could implement all three functions onto the same end-user-owned end device, perhaps by downloading the functions from the ISP or regulator. Then the LMAP Control and Report Protocols do not need to be used, but using LMAP's Information Model would still be beneficial. A Measurement Peer (or other MA involved in a Measurement Task) could be in the home gateway or outside the home network; in the latter case the Measurement Peer is highly likely to be run by a different organisation, which raises extra privacy considerations.

In both cases there will be some way for the end-user to initiate the Measurement Task(s). The mechanism is outside the scope of the initial LMAP work, but could include the user clicking a button on a GUI or sending a text message. Presumably the user will also be able to see the Measurement Results, perhaps summarised on a webpage. It is suggested that these interfaces conform to the LMAP guidance on privacy in Section 8.

6. Deployment considerations

6.1. Controller and the measurement system

The Controller should understand both the MA's LMAP Capabilities (for instance what Metrics and Measurement Methods it can perform) and about the MA's other capabilities like processing power and memory. This allows the Controller to make sure that the Measurement Schedule

of Measurement Tasks and the Reporting Schedule are sensible for each MA that it instructs.

An Instruction is likely to include several Measurement Tasks. Typically these run at different times, but it is also possible for them to run at the same time. Some Tasks may be compatible, in that they do not affect each other's Results, whilst with others great care would need to be taken. Some Tasks may be complementary. For example, one Task may be followed by a traceroute Task to the same destination address, in order to learn the network path that was measured.

The Controller should ensure that the Measurement Tasks do not have an adverse effect on the end user. Tasks, especially those that generate a substantial amount of Measurement Traffic, will often include a pre-check that the user isn't already sending traffic (Section 5.3). Another consideration is whether Measurement Traffic will impact a Subscriber's bill or traffic cap.

A Measurement System may have multiple Controllers (but note the overriding principle that a single MA is instructed by a single Controller at any point in time (Section 4.2)). For example, there could be different Controllers for different types of MA (home gateways, tablets) or locations (Ipswich, Edinburgh, Paris), for load balancing or to cope with failure of one Controller.

The measurement system also needs to consider carefully how to interpret missing Results. The correct interpretation depends on why the Results are missing (perhaps related to measurement suppression or delayed Report submission), and potentially on the specifics of the Measurement Task and Measurement Schedule. For example, the set of packets represented by a Flow may be empty; that is, an Observed Traffic Flow may represent zero or more packets. The Flow would still be reported according to schedule.

6.2. Measurement Agent

The MA should be cautious about resuming Measurement Tasks if it re-boots or has been off-line for some time, as its Instruction may be stale. In the former case it also needs to ensure that its clock has re-set correctly, so that it interprets the Schedule correctly.

If the MA runs out of storage space for Measurement Results or can't contact the Controller, then the appropriate action is specific to the device and Measurement System.

The Measurement Agent could take a number of forms: a dedicated probe, software on a PC, embedded into an appliance, or even embedded

into a gateway. A single site (home, branch office etc.) that is participating in a measurement could make use of one or multiple Measurement Agents or Measurement Peers in a single measurement.

The Measurement Agent could be deployed in a variety of locations. Not all deployment locations are available to every kind of Measurement Agent. There are also a variety of limitations and trade-offs depending on the final placement. The next sections outline some of the locations a Measurement Agent may be deployed. This is not an exhaustive list and combinations may also apply.

6.2.1. Measurement Agent on a networked device

A MA may be embedded on a device that is directly connected to the network, such as a MA on a smartphone. Other examples include a MA downloaded and installed on a subscriber's laptop computer or tablet when the network service is provided on wired or other wireless radio technologies, such as Wi-Fi.

6.2.2. Measurement Agent embedded in site gateway

A Measurement Agent embedded with the site gateway, for example a home router or the edge router of a branch office in a managed service environment, is one of better places the Measurement Agent could be deployed. All site-to-ISP traffic would traverse through the gateway. So, Measurement Methods that measure user traffic could easily be performed. Similarly, due to this user traffic visibility, a Measurement Method that generates Measurement Traffic could ensure it does not compete with user traffic. Generally NAT and firewall services are built into the gateway, allowing the Measurement Agent the option to offer its Controller-facing management interface outside of the NAT/firewall. This placement of the management interface allows the Controller to unilaterally contact the Measurement Agent for instructions. However, a Measurement Agent on a site gateway (whether end-user service-provider owned) will generally not be directly available for over the top providers, the regulator, end users or enterprises.

6.2.3. Measurement Agent embedded behind site NAT /firewall

The Measurement Agent could also be embedded behind a NAT, a firewall, or both. In this case the Controller may not be able to unilaterally contact the Measurement Agent unless either static port forwarding or firewall pin holing is configured. Configuring port forwarding could use protocols such as PCP [RFC6887], TR-069 [TR-069] or UPnP [UPnP]. To open a pin hole in the firewall, the Measurement Agent could send keepalives towards the Controller (and perhaps use these also as a network reachability test).

6.2.4. Multi-homed Measurement Agent

If the device with the Measurement Agent is single homed then there is no confusion about what interface to measure. Similarly, if the MA is at the gateway and the gateway only has a single WAN-side and a single LAN-side interface, there is little confusion - for Measurement Methods that generate Measurement Traffic, the location of the other MA or Measurement Peer determines whether the WAN or LAN is measured.

However, the device with the Measurement Agent may be multi-homed. For example, a home or campus may be connected to multiple broadband ISPs, such as a wired and wireless broadband provider, perhaps for redundancy or load- sharing. It may also be helpful to think of dual stack IPv4 and IPv6 broadband devices as multi-homed. More generally, Section 3.2 of [RFC7368] describes dual-stack and multi-homing topologies that might be encountered in a home network, [RFC6419] provides the current practices of multi-interfaces hosts, and the Multiple Interfaces (mif) working group covers cases where hosts are either directly attached to multiple networks (physical or virtual) or indirectly (multiple default routers, etc.). In these cases, there needs to be clarity on which network connectivity option is being measured.

One possibility is to have a Measurement Agent per interface. Then the Controller’s choice of MA determines which interface is measured. However, if a MA can measure any of the interfaces, then the Controller defines in the Instruction which interface the MA should use for a Measurement Task; if the choice of interface is not defined then the MA uses the default one. Explicit definition is preferred if the Measurement System wants to measure the performance of a particular network, whereas using the default is better if the Measurement System wants to include the impact of the MA’s interface selection algorithm. In any case, the Measurement Result should include the network that was measured.

6.2.5. Measurement Agent embedded in ISP network

A MA may be embedded on a device that is part of an ISP’s network, such as a router or switch. Usually the network devices with an embedded MA will be strategically located, such as a Carrier Grade NAT or ISP Gateway. [RFC7398] gives many examples where a MA might be located within a network to provide an intermediate measurement point on the end-to-end path. Other examples include a network device whose primary role is to host MA functions and the necessary measurement protocol.

6.3. Measurement Peer

A Measurement Peer participates in some Measurement Methods. It may have specific functionality to enable it to participate in a particular Measurement Method. On the other hand, other Measurement Methods may require no special functionality. For example if the Measurement Agent sends a ping to example.com then the server at example.com plays the role of a Measurement Peer; or if the MA monitors existing traffic, then the existing end points are Measurement Peers.

A device may participate in some Measurement Methods as a Measurement Agent and in others as a Measurement Peer.

Measurement Schedules should account for limited resources in a Measurement Peer when instructing a MA to execute measurements with a Measurement Peer. In some measurement protocols, such as [RFC4656] and [RFC5357], the Measurement Peer can reject a measurement session or refuse a control connection prior to setting-up a measurement session and so protect itself from resource exhaustion. This is a valuable capability because the MP may be used by more than one organisation.

6.4. Deployment examples

In this section we describe some deployment scenarios that are feasible within the LMAP framework defined in this document.

A very simple example of a Measurement Peer (MP) is a web server that the MA is downloading a web page from (such as www.example.com) in order to perform a speed test. The web server is a MP and from its perspective, the MA is just another client; the MP doesn’t have a specific function for assisting measurements. This is described in the figure below.

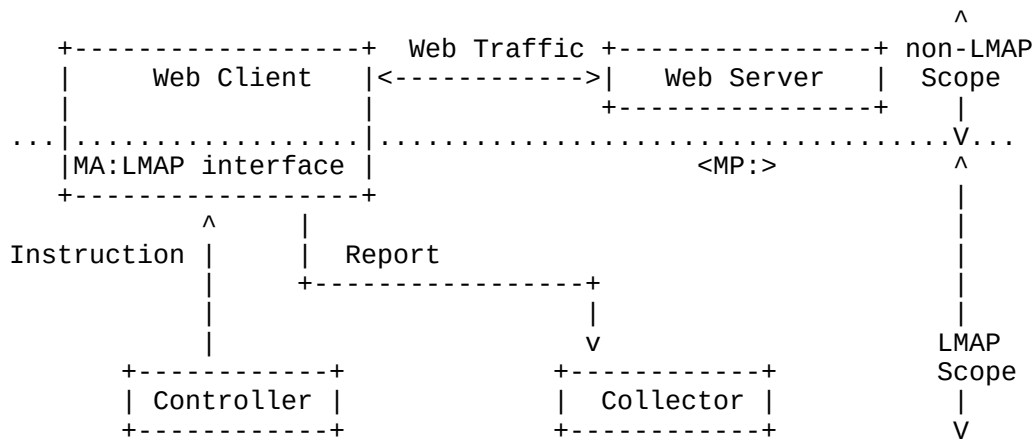

Schematic of LMAP-based Measurement System,
with Web server as Measurement Peer

Another case that is slightly different than this would be the one of a TWAMP-responder. This is also a MP, with a helper function, the TWAMP server, which is specially deployed to assist the MAs that perform TWAMP tests. Another example is with a ping server, as described in Section 2.

A further example is the case of a traceroute like measurement. In this case, for each packet sent, the router where the TTL expires is performing the MP function. So for a given Measurement Task, there is one MA involved and several MPs, one per hop.

In the figure below we depict the case of an OWAMP (One-Way Active Measurement Protocol) responder acting as an MP. In this case, the helper function in addition reports results back to the MA. So it has both a data plane and control interface with the MA.

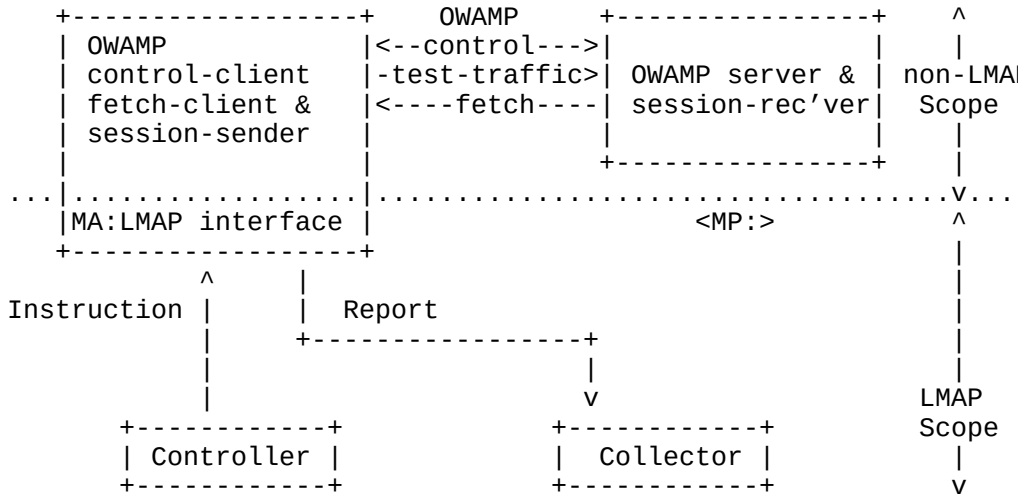

Schematic of LMAP-based Measurement System,
with OWAMP server as Measurement Peer

However, it is also possible to use two Measurement Agents when performing one way Measurement Tasks, as described in the figure below. Both MAs are instructed by the Controller: MA-1 to send the traffic and MA-2 to measure the received traffic and send Reports to the Collector. Note that the Measurement Task at MA-2 can listen for traffic from MA-1 and respond multiple times without having to be rescheduled.

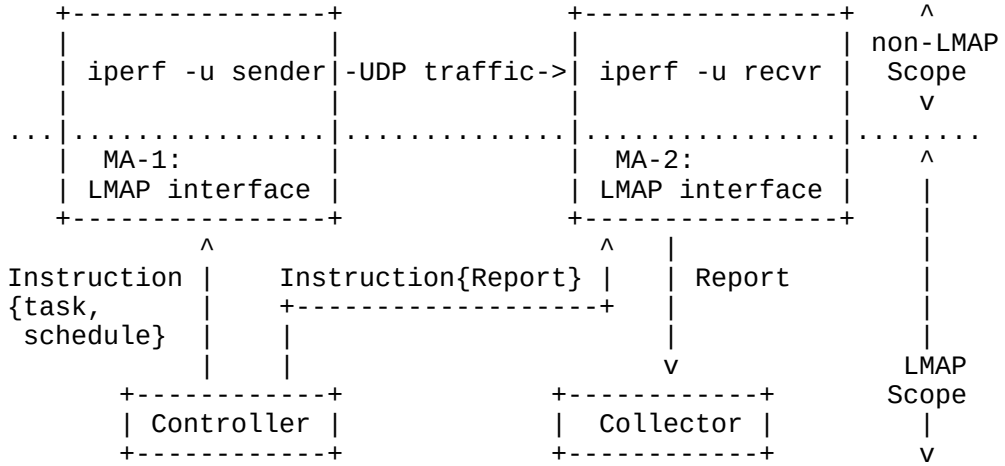

Schematic of LMAP-based Measurement System, with two
Measurement Agents cooperating to measure UDP traffic

Next, we consider Measurement Methods that meter the Observed Traffic Flow. Traffic generated in one point in the network flowing towards a given destination and the traffic is observed in some point along the path. One way to implement this is that the endpoints generating and receiving the traffic are not instructed by the Controller; hence they are MPs. The MA is located along the path with a monitor function that measures the traffic. The MA is instructed by the Controller to monitor that particular traffic and to send the Report to the Collector. It is depicted in the figure below.

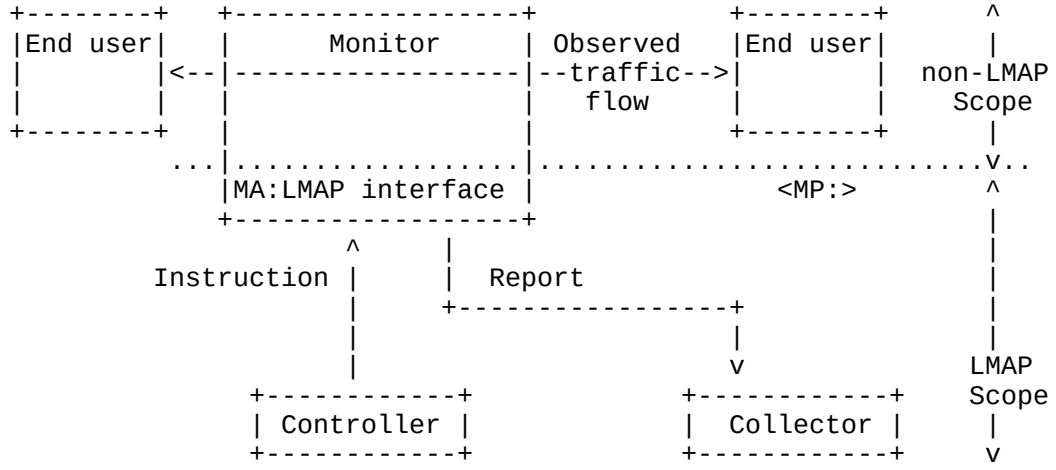

Schematic of LMAP-based Measurement System,
with a Measurement Agent monitoring traffic

7. Security considerations

The security of the LMAP framework should protect the interests of the measurement operator(s), the network user(s) and other actors who could be impacted by a compromised measurement deployment. The Measurement System must secure the various components of the system from unauthorised access or corruption. Much of the general advice contained in section 6 of [RFC4656] is applicable here.

The process to upgrade the firmware in an MA is outside the scope of the initial LMAP work, just as is the protocol to bootstrap the MAs. However, systems which provide remote upgrade must secure authorised access and integrity of the process.

We assume that each Measurement Agent (MA) will receive its Instructions from a single organisation, which operates the Controller. These Instructions must be authenticated (to ensure that they come from the trusted Controller), checked for integrity (to

ensure no-one has tampered with them) and not vulnerable to replay attacks. If a malicious party can gain control of the MA they can use it to launch DoS attacks at targets, create a platform for pervasive monitoring [RFC7258], reduce the end user's quality of experience and corrupt the Measurement Results that are reported to the Collector. By altering the Measurement Tasks and/or the address that Results are reported to, they can also compromise the confidentiality of the network user and the MA environment (such as information about the location of devices or their traffic). The Instruction Messages also need to be encrypted to maintain confidentiality, as the information might be useful to an attacker.

Reporting by the MA must be encrypted to maintain confidentiality, so that only the authorised Collector can decrypt the results, to prevent the leakage of confidential or private information. Reporting must also be authenticated (to ensure that it comes from a trusted MA and that the MA reports to a genuine Collector) and not vulnerable to tampering (which can be ensured through integrity and replay checks). It must not be possible to fool a MA into injecting falsified data and the results must also be held and processed securely after collection and analysis. See section 8.5.2 below for additional considerations on stored data compromise, and section 8.6 on potential mitigations for compromise.

Since Collectors will be contacted repeatedly by MAs using the Collection Protocol to convey their recent results, a successful attack to exhaust the communication resources would prevent a critical operation: reporting. Therefore, all LMAP Collectors should implement technical mechanisms to:

- limit the number of reporting connections from a single MA (simultaneous, and connections per unit time).
- limit the transmission rate from a single MA.
- limit the memory/storage consumed by a single MA's reports.
- efficiently reject reporting connections from unknown sources.
- separate resources if multiple authentication strengths are used, where the resources should be separated according to each class of strength.

A corrupted MA could report falsified information to the Collector. Whether this can be effectively mitigated depends on the platform on which the MA is deployed, but where the MA is deployed on a customer-controlled device then the reported data is to some degree inherently untrustworthy. Further, a sophisticated party could distort some

Measurement Methods, perhaps by dropping or delaying packets for example. This suggests that the network operator should be cautious about relying on Measurement Results for action such as refunding fees if a service level agreement is not met.

As part of the protocol design, it will be decided how LMAP operates over the underlying protocol (Section 5.5). The choice raises various security issues, such as how to operate through a NAT and how to protect the Controller and Collector from denial of service attacks.

The security mechanisms described above may not be strictly necessary if the network's design ensures the LMAP components and their communications are already secured, for example potentially if they are all part of an ISP's dedicated management network.

Finally, there are three other issues related to security: privacy (considered in Section 8 below), availability and 'gaming the system'. While the loss of some MAs may not be considered critical, the unavailability of the Collector could mean that valuable business data or data critical to a regulatory process is lost. Similarly, the unavailability of a Controller could mean that the MAs do not operate a correct Measurement Schedule.

A malicious party could "game the system". For example, where a regulator is running a Measurement System in order to benchmark operators, an operator could try to identify the broadband lines that the regulator was measuring and prioritise that traffic. Normally, this potential issue is handled by a code of conduct. It is outside the scope of the initial LMAP work to consider the issue.

8. Privacy considerations

The LMAP work considers privacy as a core requirement and will ensure that by default the Control and Report Protocols operate in a privacy-sensitive manner and that privacy features are well-defined.

This section provides a set of privacy considerations for LMAP. This section benefits greatly from the timely publication of [RFC6973]. Privacy and security (Section 7) are related. In some jurisdictions privacy is called data protection.

We begin with a set of assumptions related to protecting the sensitive information of individuals and organisations participating in LMAP-orchestrated measurement and data collection.

8.1. Categories of entities with information of interest

LMAP protocols need to protect the sensitive information of the following entities, including individuals and organisations who participate in measurement and collection of results.

- Individual Internet users: Persons who utilise Internet access services for communications tasks, according to the terms of service of a service agreement. Such persons may be a service Subscriber, or have been given permission by the Subscriber to use the service.

- Internet service providers: Organisations who offer Internet access service subscriptions, and thus have access to sensitive information of individuals who choose to use the service. These organisations desire to protect their Subscribers and their own sensitive information which may be stored in the process of performing Measurement Tasks and collecting Results.

- Regulators: Public authorities responsible for exercising supervision of the electronic communications sector, and which may have access to sensitive information of individuals who participate in a measurement campaign. Similarly, regulators desire to protect the participants and their own sensitive information.

- Other LMAP system operators: Organisations who operate Measurement Systems or participate in measurements in some way.

Although privacy is a protection extended to individuals, we discuss data protection by ISPs and other LMAP system operators in this section. These organisations have sensitive information involved in the LMAP system, and many of the same dangers and mitigations are applicable. Further, the ISPs store information on their Subscribers beyond that used in the LMAP system (for instance billing information), and there should be a benefit in considering all the needs and potential solutions coherently.

8.2. Examples of sensitive information

This section gives examples of sensitive information which may be measured or stored in a Measurement System, and which is to be kept private by default in the LMAP core protocols.

Examples of Subscriber or authorised Internet user sensitive information:

- Sub-IP layer addresses and names (MAC address, base station ID, SSID)
- IP address in use
- Personal Identification (real name)
- Location (street address, city)
- Subscribed service parameters
- Contents of traffic (activity, DNS queries, destinations, equipment types, account info for other services, etc.)
- Status as a study volunteer and Schedule of Measurement Tasks

Examples of Internet Service Provider sensitive information:

- Measurement device identification (equipment ID and IP address)
- Measurement Instructions (choice of measurements)
- Measurement Results (some may be shared, others may be private)
- Measurement Schedule (exact times)
- Network topology (locations, connectivity, redundancy)
- Subscriber billing information, and any of the above Subscriber information known to the provider.
- Authentication credentials (such as certificates)

Other organisations will have some combination of the lists above. The LMAP system would not typically expose all of the information above, but could expose a combination of items which could be correlated with other pieces collected by an attacker (as discussed in the section on Threats below).

8.3. Different privacy issues raised by different sorts of Measurement Methods

Measurement Methods raise different privacy issues depending on whether they measure traffic created specifically for that purpose, or whether they measure user traffic.

Measurement Tasks conducted on user traffic store sensitive information, however briefly this storage may be. We note that some

authorities make a distinction on time of storage, and information that is kept only temporarily to perform a communications function is not subject to regulation (for example, active queue management, deep packet inspection). Such Measurement Tasks could reveal all the websites a Subscriber visits and the applications and/or services they use. This issue is not specific to LMAP. For instance, IPFIX has discussed similar issues (see section 11.8 of [RFC7011]), but mitigations described in the sections below were considered beyond their scope.

Other types of Measurement Task are conducted on traffic which is created specifically for the purpose. Even if a user host generates Measurement Traffic, there is limited sensitive information about the Subscriber present and stored in the Measurement System:

- IP address in use (and possibly sub-IP addresses and names)

- Status as a study volunteer and Schedule of Measurement Tasks

On the other hand, for a service provider the sensitive information like Measurement Results is the same for all Measurement Tasks.

From the Subscriber perspective, both types of Measurement Task potentially expose the description of Internet access service and specific service parameters, such as subscribed rate and type of access.

8.4. Privacy analysis of the communication models

This section examines each of the protocol exchanges described at a high level in Section 5 and some example Measurement Tasks, and identifies specific sensitive information which must be secured during communication for each case. With the protocol-related sensitive information identified, we can better consider the threats described in the following section.

From the privacy perspective, all entities participating in LMAP protocols can be considered "observers" according to the definition in [RFC6973]. Their stored information potentially poses a threat to privacy, especially if one or more of these functional entities has been compromised. Likewise, all devices on the paths used for control, reporting, and measurement are also observers.

8.4.1. MA Bootstrapping

Section 5.1 provides the communication model for the Bootstrapping process.

Although the specification of mechanisms for Bootstrapping the MA are beyond the initial LMAP work scope, designers should recognize that the Bootstrapping process is extremely powerful and could cause an MA to join a new or different LMAP system with a different Controller and Collector, or simply install new Metrics with associated Measurement Methods (for example to record DNS queries). A Bootstrap attack could result in a breach of the LMAP system with significant sensitive information exposure depending on the capabilities of the MA, so sufficient security protections are warranted.

The Bootstrapping process provides sensitive information about the LMAP system and the organisation that operates it, such as

- the MA's identifier (MA-ID)
- the address that identifies the Control Channel, such as the Controller's FQDN
- Security information for the Control Channel

During the Bootstrap process for an MA located at a single subscriber's service demarcation point, the MA receives a MA-ID which is a persistent pseudonym for the Subscriber. Thus, the MA-ID is considered sensitive information because it could provide the link between Subscriber identification and Measurements Results.

Also, the Bootstrap process could assign a Group-ID to the MA. The specific definition of information represented in a Group-ID is to be determined, but several examples are envisaged including use as a pseudonym for a set of Subscribers, a class of service, an access technology, or other important categories. Assignment of a Group-ID enables anonymisation sets to be formed on the basis of service type/grade/rates. Thus, the mapping between Group-ID and MA-ID is considered sensitive information.

8.4.2. Controller <-> Measurement Agent

The high-level communication model for interactions between the LMAP Controller and Measurement Agent is illustrated in Section 5.2. The primary purpose of this exchange is to authenticate and task a Measurement Agent with Measurement Instructions, which the Measurement Agent then acts on autonomously.

Primarily IP addresses and pseudonyms (MA-ID, Group-ID) are exchanged with a capability request, then measurement-related information of interest such as the parameters, schedule, metrics, and IP addresses of measurement devices. Thus, the measurement Instruction contains sensitive information which must be secured. For example, the fact

that an ISP is running additional measurements beyond the set reported externally is sensitive information, as are the additional Measurements Tasks themselves. The Measurement Schedule is also sensitive, because an attacker intending to bias the results without being detected can use this information to great advantage.

An organisation operating the Controller having no service relationship with a user who hosts the Measurement Agent *could* gain real-name mapping to a public IP address through user participation in an LMAP system (this applies to the Measurement Collection protocol, as well).

8.4.3. Collector <-> Measurement Agent

The high-level communication model for interactions between the Measurement Agent and Collector is illustrated in Section 5.4. The primary purpose of this exchange is to authenticate and collect Measurement Results from a MA, which the MA has measured autonomously and stored.

The Measurement Results are the additional sensitive information included in the Collector-MA exchange. Organisations collecting LMAP measurements have the responsibility for data control. Thus, the Results and other information communicated in the Collector protocol must be secured.

8.4.4. Measurement Peer <-> Measurement Agent

A Measurement Method involving Measurement Traffic raises potential privacy issues, although the specification of the mechanisms is beyond the scope of the initial LMAP work. The high-level communications model below illustrates the various exchanges to execute such a Measurement Method and store the Results.

We note the potential for additional observers in the figures below by indicating the possible presence of a NAT, which has additional significance to the protocols and direction of initiation.

The various messages are optional, depending on the nature of the Measurement Method. It may involve sending Measurement Traffic from the Measurement Peer to MA, MA to Measurement Peer, or both. Similarly, a second (or more) MAs may be involved. (Note: For simplicity, the Figure and description don’t show the non-LMAP functionality that is associated with the transfer of the Measurement Traffic and is located at the devices with the MA and MP.)

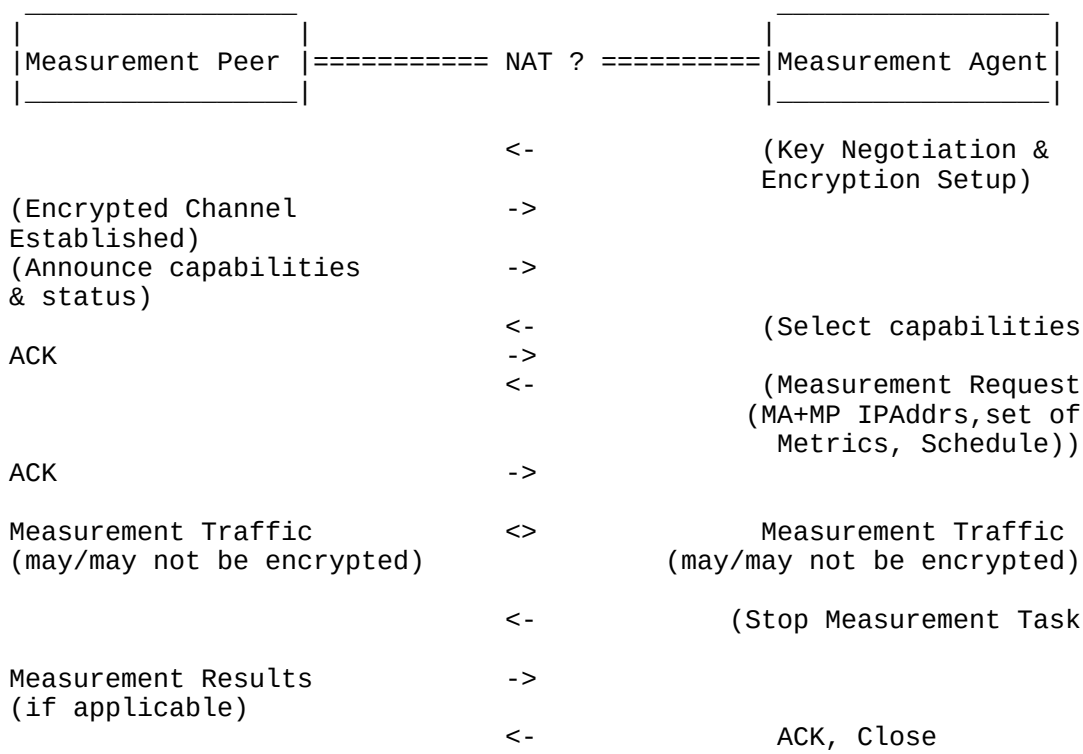

This exchange primarily exposes the IP addresses of measurement devices and the inference of measurement participation from such traffic. There may be sensitive information on key points in a service provider's network included. There may also be access to measurement-related information of interest such as the Metrics, Schedule, and intermediate results carried in the Measurement Traffic (usually a set of timestamps).

The Measurement Peer may be able to use traffic analysis (perhaps combined with traffic injection) to obtain interesting insights about the Subscriber. As a simple example, if the Measurement Task includes a pre-check that the end-user isn't already sending traffic, the Measurement Peer may be able to deduce when the Subscriber is away on holiday, for example.

If the Measurement Traffic is unencrypted, as found in many systems today, then both timing and limited results are open to on-path observers.

8.4.5. Measurement Agent

Some Measurement Methods only involve a single Measurement Agent observing existing traffic. They raise potential privacy issues, although the specification of the mechanisms is beyond the scope of the initial LMAP work.

The high-level communications model below illustrates the collection of user information of interest with the Measurement Agent performing the monitoring and storage of the Results. This particular exchange is for measurement of DNS Response Time, which most frequently uses UDP transport. (Note: For simplicity, the Figure and description don't show the non-LMAP functionality that is associated with the transfer of the Measurement Traffic and is located at the devices with the MA.)

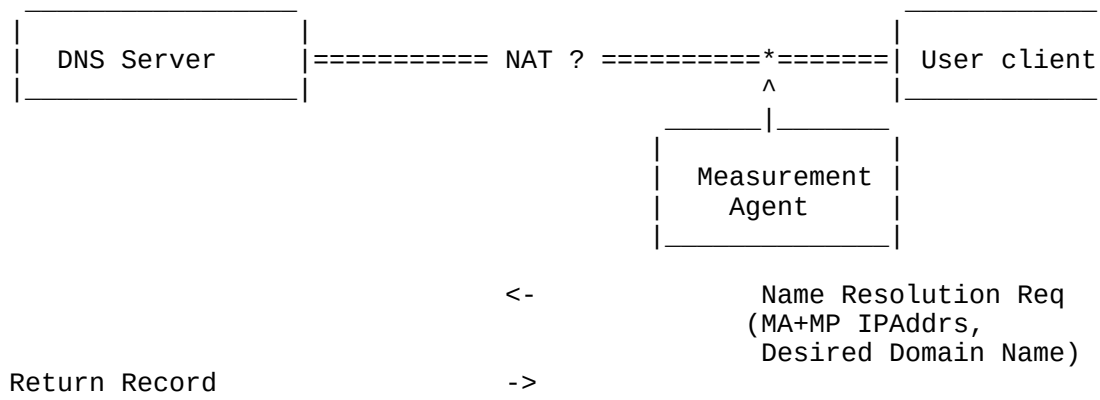

In this particular example, the MA monitors DNS messages in order to measure that DNS response time. The Measurement Agent may be embedded in the user host, or it may be located in another device capable of observing user traffic. The MA learns the IP addresses of measurement devices and the intent to communicate with or access the services of a particular domain name, and perhaps also information on key points in a service provider's network, such as the address of one of its DNS servers.

In principle, any of the user sensitive information of interest (listed above) can be collected and stored in the monitoring scenario and so must be secured.

It would also be possible for a Measurement Agent to source the DNS query itself. But then there are few privacy concerns.

8.4.6. Storage and reporting of Measurement Results

Although the mechanisms for communicating results (beyond the initial Collector) are beyond the initial LMAP work scope, there are potential privacy issues related to a single organisation's storage and reporting of Measurement Results. Both storage and reporting functions can help to preserve privacy by implementing the mitigations described below.

8.5. Threats

This section indicates how each of the threats described in [RFC6973] apply to the LMAP entities and their communication and storage of "information of interest". Denial of Service (DOS) and other attacks described in the Security section represent threats as well, and these attacks are more effective when sensitive information protections have been compromised.

8.5.1. Surveillance

Section 5.1.1 of [RFC6973] describes Surveillance as the "observation or monitoring of and individual's communications or activities." Hence all Measurement Methods that measure user traffic are a form of surveillance, with inherent risks.

Measurement Methods which avoid periods of user transmission indirectly produce a record of times when a subscriber or authorised user has used their network access service.

Measurement Methods may also utilise and store a Subscriber's currently assigned IP address when conducting measurements that are relevant to a specific Subscriber. Since the Measurement Results are time-stamped, they could provide a record of IP address assignments over time.

Either of the above pieces of information could be useful in correlation and identification, described below.

8.5.2. Stored data compromise

Section 5.1.2 of [RFC6973] describes Stored Data Compromise as resulting from inadequate measures to secure stored data from unauthorised or inappropriate access. For LMAP systems this includes deleting or modifying collected measurement records, as well as data theft.

The primary LMAP entity subject to compromise is the repository, which stores the Measurement Results; extensive security and privacy

threat mitigations are warranted. The Collector and MA also store sensitive information temporarily, and need protection. The communications between the local storage of the Collector and the repository is beyond the scope of the initial LMAP work, though this communications channel will certainly need protection as well as the mass storage itself.

The LMAP Controller may have direct access to storage of Subscriber information (location, billing, service parameters, etc.) and other information which the controlling organisation considers private, and again needs protection.

Note that there is tension between the desire to store all raw results in the LMAP Collector (for reproducibility and custom analysis), and the need to protect the privacy of measurement participants. Many of the compromise mitigations described in section 8.6 below are most efficient when deployed at the MA, therefore minimising the risks with stored results.

8.5.3. Correlation and identification

Sections 5.2.1 and 5.2.2 of [RFC6973] describe Correlation as combining various pieces of information to obtain desired characteristics of an individual, and Identification as using this combination to infer identity.

The main risk is that the LMAP system could unwittingly provide a key piece of the correlation chain, starting with an unknown Subscriber's IP address and another piece of information. For example, a Subscriber utilised Internet access from 2000 to 2310 UTC, because the Measurement Tasks were deferred, or sent a name resolution for www.example.com at 2300 UTC.

If a user's access with another system already gave away sensitive info, correlation is clearly easier and can result in re-identification, even when an LMAP conserves sensitive information to great extent.

8.5.4. Secondary use and disclosure

Sections 5.2.3 and 5.2.4 of [RFC6973] describes Secondary Use as unauthorised utilisation of an individual's information for a purpose the individual did not intend, and Disclosure is when such information is revealed causing other's notions of the individual to change, or confidentiality to be violated.

Measurement Methods that measure user traffic are a form of Secondary Use, and the Subscribers' permission should be obtained beforehand.

It may be necessary to obtain the measured ISP's permission to conduct measurements, for example when required by the terms and conditions of the service agreement, and notification is considered good measurement practice.

For Measurement Methods that measure Measurement Traffic the Measurement Results provide some limited information about the Subscriber or ISP and could result in Secondary Uses. For example, the use of the Results in unauthorised marketing campaigns would qualify as Secondary Use. Secondary use may break national laws and regulations, and may violate individual's expectations or desires.

8.6. Mitigations

This section examines the mitigations listed in section 6 of [RFC6973] and their applicability to LMAP systems. Note that each section in [RFC6973] identifies the threat categories that each technique mitigates.

8.6.1. Data minimisation

Section 6.1 of [RFC6973] encourages collecting and storing the minimal information needed to perform a task.

LMAP results can be useful for general reporting about performance and for specific troubleshooting. They need different levels of information detail, as explained in the paragraphs below.

For general results, the results can be aggregated into large categories (the month of March, all subscribers West of the Mississippi River). In this case, all individual identifications (including IP address of the MA) can be excluded, and only relevant results are provided. However, this implies a filtering process to reduce the information fields, because greater detail was needed to conduct the Measurement Tasks in the first place.

For troubleshooting, so that a network operator or end user can identify a performance issue or failure, potentially all the network information (IP addresses, equipment IDs, location), Measurement Schedule, service configuration, Measurement Results, and other information may assist in the process. This includes the information needed to conduct the Measurements Tasks, and represents a need where the maximum relevant information is desirable, therefore the greatest protections should be applied. This level of detail is greater than needed for general performance monitoring.

As regards Measurement Methods that measure user traffic, we note that a user may give temporary permission (to enable detailed

troubleshooting), but withhold permission for them in general. Here the greatest breadth of sensitive information is potentially exposed, and the maximum privacy protection must be provided. The Collector may perform pre-storage minimisation and other mitigations (below) to help preserve privacy.

For MAs with access to the sensitive information of users (e.g., within a home or a personal host/handset), it is desirable for the results collection to minimise the data reported, but also to balance this desire with the needs of troubleshooting when a service subscription exists between the user and organisation operating the measurements.

8.6.2. Anonymity

Section 6.1.1 of [RFC6973] describes a way in which anonymity is achieved: "there must exist a set of individuals that appear to have the same attributes as the individual", defined as an "anonymity set".

Experimental methods for anonymisation of user identifiable data (and so particularly applicable to Measurement Methods that measure user traffic) have been identified in [RFC6235]. However, the findings of several of the same authors is that "there is increasing evidence that anonymisation applied to network trace or flow data on its own is insufficient for many data protection applications as in [Bur10]." Essentially, the details of such Measurement Methods can only be accessed by closed organisations, and unknown injection attacks are always less expensive than the protections from them. However, some forms of summary may protect the user’s sensitive information sufficiently well, and so each Metric must be evaluated in the light of privacy.

The techniques in [RFC6235] could be applied more successfully in Measurement Methods that generate Measurement Traffic, where there are protections from injection attack. The successful attack would require breaking the integrity protection of the LMAP Reporting Protocol and injecting Measurement Results (known fingerprint, see section 3.2 of [RFC6973]) for inclusion with the shared and anonymised results, then fingerprinting those records to ascertain the anonymisation process.

Beside anonymisation of measured Results for a specific user or provider, the value of sensitive information can be further diluted by summarising the results over many individuals or areas served by the provider. There is an opportunity enabled by forming anonymity sets [RFC6973] based on the reference path measurement points in [RFC7398]. 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 anonymisation applies to the Internet Service Provider, where their Internet gateway would be referred to as "mp190".

Another anonymisation technique is for the MA to include its Group-ID instead of its MA-ID in its Measurement Reports, with several MAs sharing the same Group-ID.

8.6.3. Pseudonymity

Section 6.1.2 of [RFC6973] indicates that pseudonyms, or nicknames, are a possible mitigation to revealing one's true identity, since there is no requirement to use real names in almost all protocols.

A pseudonym for a measurement device's IP address could be an LMAP-unique equipment ID. However, this would likely be a permanent handle for the device, and long-term use weakens a pseudonym's power to obscure identity.

8.6.4. Other mitigations

Data can be de-personalised by blurring it, for example by adding synthetic data, data-swapping, or perturbing the values in ways that can be reversed or corrected.

Sections 6.2 and 6.3 of [RFC6973] describe User Participation and Security, respectively.

Where LMAP measurements involve devices on the Subscriber's premises or Subscriber-owned equipment, it is essential to secure the Subscriber's permission with regard to the specific information that will be collected. The informed consent of the Subscriber (and, if different, the end user) may be needed, including the specific purpose of the measurements. The approval process could involve showing the Subscriber their measured information and results before instituting periodic collection, or before all instances of collection, with the option to cancel collection temporarily or permanently.

It should also be clear who is legally responsible for data protection (privacy); in some jurisdictions this role is called the 'data controller'. It is always good practice to limit the time of personal information storage.

Although the details of verification would be impenetrable to most subscribers, the MA could be architected as an "app" with open source-code, pre-download and embedded terms of use and agreement on

measurements, and protection from code modifications usually provided by the app-stores. Further, the app itself could provide data reduction and temporary storage mitigations as appropriate and certified through code review.

LMAP protocols, devices, and the information they store clearly need to be secure from unauthorised access. This is the hand-off between privacy and security considerations (Section 7). The Data Controller has the (legal) responsibility to maintain data protections described in the Subscriber's agreement and agreements with other organisations.

Finally, it is recommended that each entity in section 8.1, (individuals, ISPs, Regulators, others) assess the risks of LMAP data collection by conducting audits of their data protection methods.

9. IANA considerations

There are no IANA considerations in this memo.

10. Acknowledgments

This document originated as a merger of three individual drafts: draft-eardley-lmap-terminology-02, draft-akhter-lmap-framework-00, and draft-eardley-lmap-framework-02.

Thanks to Juergen Schoenwaelder for his detailed review of the terminology. Thanks to Charles Cook for a very detailed review of -02. Thanks to Barbara Stark and Ken Ko for many helpful comments about later versions.

Thanks to numerous people for much discussion, directly and on the LMAP list (apologies to those unintentionally omitted): Alan Clark, Alissa Cooper, Andrea Soppera, Barbara Stark, Benoit Claise, Brian Trammell, Charles Cook, Dan Romascanu, Dave Thorne, Frode Soerensen, Greg Mirsky, Guangqing Deng, Jason Weil, Jean-Francois Tremblay, Jerome Benoit, Joachim Fabini, Juergen Schoenwaelder, Jukka Manner, Ken Ko, Lingli Deng, Mach Chen, Matt Mathis, Marc Ibrahim, Michael Bugenhagen, Michael Faath, Nalini Elkins, Radia Perlman, Rolf Winter, Sam Crawford, Sharam Hakimi, Steve Miller, Ted Lemon, Timothy Carey, Vaibhav Bajpai, Vero Zheng, William Lupton.

Philip Eardley, Trevor Burbridge and Marcelo Bagnulo work in part on the Leone research project, which receives funding from the European Union Seventh Framework Programme [FP7/2007-2013] under grant agreement number 317647.

11. History

First WG version, copy of draft-folks-lmap-framework-00.

11.1. From -00 to -01

- new sub-section of possible use of Group-IDs for privacy
- tweak to definition of Control protocol
- fix typo in figure in S5.4

11.2. From -01 to -02

- change to INFORMATIONAL track (previous version had typo'd Standards track)
- new definitions for Capabilities Information and Failure Information
- clarify that diagrams show LMAP-level information flows. Underlying protocol could do other interactions, eg to get through NAT or for Collector to pull a Report
- add hint that after a re-boot should pause random time before re-register (to avoid mass calling event)
- delete the open issue "what happens if a Controller fails" (normal methods can handle)
- add some extra words about multiple Tasks in one Schedule
- clarify that new Schedule replaces (rather than adds to) and old one. Similarly for new configuration of Measurement Tasks or Report Channels.
- clarify suppression is temporary stop; send a new Schedule to permanently stop Tasks
- alter suppression so it is ACKed
- add un-suppress message
- expand the text on error reporting, to mention Reporting failures (as well as failures to action or execute Measurement Task & Schedule)
- add some text about how to have Tasks running indefinitely

- add that optionally a Report is not sent when there are no Measurement Results
- add that a Measurement Task may create more than one Measurement Result
- clarify /amend /expand that Reports include the "raw" Measurement Results - any pre-processing is left for lmap2.0
- add some cautionary words about what if the Collector unexpectedly doesn’t hear from a MA
- add some extra words about the potential impact of Measurement Tasks
- clarified various aspects of the privacy section
- updated references
- minor tweaks

11.3. From -02 to -03

- alignment with the Information Model [burbridge-lmap-information-model] as this is agreed as a WG document
- One-off and periodic Measurement Schedules are kept separate, so that they can be updated independently
- Measurement Suppression in a separate sub-section. Can now optionally include particular Measurement Tasks &/or Schedules to suppress, and start/stop time
- for clarity, concept of Channel split into Control, Report and MA-to-Controller Channels
- numerous editorial changes, mainly arising from a very detailed review by Charles Cook
-

11.4. From -03 to -04

- updates following the WG Last Call, with the proposed consensus on the various issues as detailed in http://tools.ietf.org/agenda/89/slides/slides-89-lmap-2.pdf. In particular:

- tweaked definitions, especially of Measurement Agent and Measurement Peer
- Instruction - left to each implementation & deployment of LMAP to decide on the granularity at which an Instruction Message works
- words added about overlapping Measurement Tasks (Measurement System can handle any way they choose; Report should mention if the Task overlapped with another)
- Suppression: no defined impact on Passive Measurement Task; extra option to suppress on-going Active Measurement Tasks; suppression doesn't go to Measurement Peer, since they don't understand Instructions
- new concept of Data Transfer Task (and therefore adjustment of the Channel concept)
- enhancement of Results with Subscriber's service parameters - could be useful, don't define how but can be included in Report to various other sections
- various other smaller improvements, arising from the WGLC
- Appendix added with examples of Measurement Agents and Peers in various deployment scenarios. To help clarify what these terms mean.

11.5. From -04 to -05

- clarified various scoping comments by using the phrase "scope of initial LMAP work" (avoiding "scope of LMAP WG" since this may change in the future)
- added a Configuration Protocol - allows the Controller to update the MA about information that it obtained during the bootstrapping process (for consistency with Information Model)
- Removed over-detailed information about the relationship between the different items in Instruction, as this seems more appropriate for the information model. Clarified that the lists given are about the aims and not a list of information elements (these will be defined in draft-ietf-information-model).
- the Measurement Method, specified as a URI to a registry entry - rather than a URN

- MA configured with time limit after which, if it hasn't heard from Controller, then it stops running Measurement Tasks (rather than this being part of a Schedule)
- clarified there is no distinction between how capabilities, failure and logging information are transferred (all can be when requested by Controller or by MA on its own initiative).
- removed mention of Data Transfer Tasks. This abstraction is left to the information model i-d
- added Deployment sub-section about Measurement Agent embedded in ISP Network
- various other smaller improvements, arising from the 2nd WGLC

11.6. From -05 to -06

- clarified terminlogy around Measurement Methods and Tasks. Since within a Method there may be several different roles (requester and responder, for instance)
- Suppression: there is now the concept of a flag (boolean) which indicates whether a Task is by default gets suppressed or not. The optional suppression message (with list of specific tasks /schedules to suppress) over-rides this flag.
- The previous bullet also means there is no need to make a distinction between active and passive Measurement Tasks, so this distinction is removed.
- removed Configuration Protocol - Configuration is part of the Instruction and so uses the Control Protocol.

11.7. From -06 to -07

- Clarifications and nits

11.8. From -07 to -08

- Clarifications resulting from WG 3rd LC, as discussed in https://tools.ietf.org/agenda/90/slides/slides-90-lmap-0.pdf, plus comments made in the IETF-90 meeting.
- added mention of "measurement point designations" in Measurement Task configuration and Report Protocol.

11.9. From -08 to -09

- Clarifications and changes from the AD review (Benoit Claise) and security directorate review (Radia Perlman).

11.10. From -09 to -10

- More changes from the AD review (Benoit Claise).

11.11. From -10 to -11

- More changes from the AD review (Benoit Claise).

11.12. From -11 to -12

- Fixing nits from IETF Last call and authors.

11.13. From -12 to -13

- IESG changes.

11.14. From -13 to -14

- Fixing Figure 1.

12. Informative References

[Bur10] Burkhart, M., Schatzmann, D., Trammell, B., and E. Boschi, "The Role of Network Trace anonymisation Under Attack", January 2010.

[TR-069] TR-069, , "CPE WAN Management Protocol", http://www.broadband-forum.org/technical/trlist.php, November 2013.

[UPnP] ISO/IEC 29341-x, , "UPnP Device Architecture and UPnP Device Control Protocols specifications", http://upnp.org/sdcps-and-certification/standards/, 2011.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC4101] Rescorla, E. and IAB, "Writing Protocol Models", RFC 4101, June 2005.

[RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", RFC 4122, July 2005.

[RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.

[RFC7011] Claise, B., Trammell, B., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, September 2013.

[RFC7368] Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil, "IPv6 Home Networking Architecture Principles", RFC 7368, October 2014.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, May 2014.

[I-D.ietf-lmap-use-cases] Linsner, M., Eardley, P., Burbridge, T., and F. Sorensen, "Large-Scale Broadband Measurement Use Cases", draft-ietf-lmap-use-cases-06 (work in progress), February 2015.

[I-D.ietf-ippm-metric-registry] Bagnulo, M., Claise, B., Eardley, P., Morton, A., and A. Akhter, "Registry for Performance Metrics", draft-ietf-ippm-metric-registry-02 (work in progress), February 2015.

[RFC6419] Wasserman, M. and P. Seite, "Current Practices for Multiple-Interface Hosts", RFC 6419, November 2011.

[RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.

[I-D.ietf-lmap-information-model] Burbridge, T., Eardley, P., Bagnulo, M., and J. Schoenwaelder, "Information Model for Large-Scale Measurement Platforms (LMAP)", draft-ietf-lmap-information-model-05 (work in progress), April 2015.

[RFC6235] Boschi, E. and B. Trammell, "IP Flow Anonymization Support", RFC 6235, May 2011.

[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.

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

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

[RFC3444] Pras, A. and J. Schoenwaelder, "On the Difference between Information Models and Data Models", RFC 3444, January 2003.

[RFC7398] Bagnulo, M., Burbridge, T., Crawford, S., Eardley, P., and A. Morton, "A Reference Path and Measurement Points for Large-Scale Measurement of Broadband Performance", RFC 7398, February 2015.

Authors' Addresses

Philip Eardley
BT
Adastral Park, Martlesham Heath
Ipswich
ENGLAND

Email: philip.eardley@bt.com

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

Email: acmorton@att.com

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
ENGLAND

Email: trevor.burbridge@bt.com

Paul Aitken
Brocade
Edinburgh, Scotland
UK

Email: paitken@brocade.com

Aamer Akhter
Consultant
118 Timber Hitch
Cary, NC
USA

Email: aakhter@gmail.com