

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 16, 2014

P. Eardley
T. Burbridge
BT
A. Morton
AT&T Labs
July 15, 2013

A framework for large-scale measurements
draft-eardley-lmap-framework-02

Abstract

Measuring broadband service on a large scale requires standardisation of the logical architecture and a description of the key protocols that coordinate interactions between the components. The document presents an overall framework for large-scale measurements and discusses which elements could be standardised in the IETF. It is intended to assist the work of the LMAP working group.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 16, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

Internet-Draft

LMAP Framework

July 2013

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Outline of framework	4
3.	Constraints	6
3.1.	Measurement system is under the direction of a single organisation	6
3.2.	Each MA may only have a single Controller at any point in time	7
3.3.	A Measurement Agent acts autonomously	7
4.	Work items for LMAP WG	8
4.1.	Information Model	9
4.2.	Control Protocol	10
4.3.	Report Protocol	10
5.	Related work required but out of scope of LMAP	10
5.1.	Standard measurement tests	10
5.2.	Characterisation plan	11
5.3.	Other elements	11
6.	IANA Considerations	12
7.	Security Considerations	12
8.	Acknowledgements	13
9.	Changes	13
9.1.	from -00 to -01	13
10.	Informative References	14
	Authors' Addresses	14

1. Introduction

The Large-Scale Measurement of Broadband Performance (LMAP) working group standardizes the LMAP measurement system for performance measurements of broadband access devices such as home and enterprise edge routers, personal computers, mobile devices, set top box, whether wired or wireless. Measuring portions of the Internet on a large scale is essential for accurate characterizations of performance over time and geography.

[use-cases] discusses several use cases have been proposed for large-scale measurements:

- o Operators: to help plan their network and identify faults
- o End Users: to run diagnostic checks, such as a network speed test
- o Regulators: to benchmark several network operators and support public policy development

The LMAP framework should be useful for all these.

The goal is to have the measurements (made using the same metrics and mechanisms) for a large number of points on the Internet, and to have the results collected and stored in the same form.

There are existing measurement systems. However, they typically lack some of the desirable features for a large-scale measurement system:

- o Standardised - in terms of the tests that they perform, the components, the data models and protocols for transferring information between the components. For example so that it is meaningful to compare measurements made of the same metric at different times and places. For example so that the operator of a measurement system can buy the various components from different vendors. Today's systems are proprietary in some or all of these

aspects.

- o Extensible - it should be easy to add or modify tests, for example an improved test methodology or to measure a performance metric not previously considered important (e.g., bufferbloat).
- o Large-scale - [[use-cases](#)] envisages Measurement Agents in every home gateway and edge device such as set-top-boxes and tablet computers. Existing systems have up to a few thousand Measurement Agents (without judging how much further they could scale).

- o Diversity - a measurement system should handle different types of Measurement Agent - for example Measurement Agents may come from different vendors, be in wired and wireless networks and be on devices with IPv4 or IPv6 addresses.

[2.](#) Outline of framework

The LMAP framework for large-scale measurements has four elements:

- o Measurement Agent (MA)
- o Measurement Peer
- o Controller
- o Collector

In addition there are some components that are outside LMAP but useful within the context of a large scale measurement system:

- o Initialiser
- o Subscriber Parameter Database
- o Results Database
- o Data Analysis Tools

- o Operator's OAM (Operations Administration and Management)

a large-scale measurement system essentially has three sets of communications:

- o several measurement protocols between a Measurement Agent and a Measurement Peer
- o a Control Protocol between a Controller and a MA
- o a Report Protocol between a MA and a Collector.

A Measurement Agent and a Measurement Peer jointly perform an active measurement test, by generating test traffic and measuring some metric associated with its transfer over the path from one to the other; for example the time taken to transfer a 'test file'. A MA may also conduct passive testing through the observation of traffic (i.e. without the involvement of a Measurement Peer); for example an end user's mix of applications.

The MA interacts with the Controller and Collector, and a Measurement Peer only takes part in active tests (and does not interact with the Controller and Collector).

The MA functions are implemented either in specialised hardware or as code on general purpose devices like a PC, tablet or smartphone. The Measurement Peer may be an LMAP device or a normal, non-LMAP device (for example if the MA measures the time for a DNS response or a webpage download from www.example.com).

The Controller manages a MA by instructing it which tests it should perform and when. For example it may instruct a MA at a home gateway: "Run the 'download speed test' with the test server at the end user's first IP point in the network; if the end user is active then delay the test and re-try 1 minute later, with up to 3 re-tries; repeat every hour at $xx.05 + \text{Unif}[0,180]$ seconds". The Controller also manages a MA by instructing it how to report the test results, for example: "Report results once a day in a batch at 4am + $\text{Unif}[0,180]$ seconds; if the end user is active then delay the report 5 minutes". As well as regular tests, a Controller can initiate a one-off test ("Do test now", "Report as soon as possible"). These are called the Test and Report Schedule.

The Collector accepts a Report from a MA with the results from its tests. It may also do some processing on the results – for instance to eliminate outliers, as they can severely impact the aggregated results.

Therefore the MA is a LMAP-specific device that initiates the test, gets instructions from the Controller and reports to the Collector.

It is possible that communications between two Collectors, two Controllers and a Controller and Collector may be useful in some use cases, perhaps to help a measurement system scale. Work on such a protocol is out of scope of LMAP (?)

The Initialiser, Subscriber Parameter Database, Results Database, Data Analysis Tools and OAM are out of scope of LMAP. They may be provided through existing protocols or applications and are likely to be part of a complete large-scale measurement system. See [Section 5](#) for further discussion.

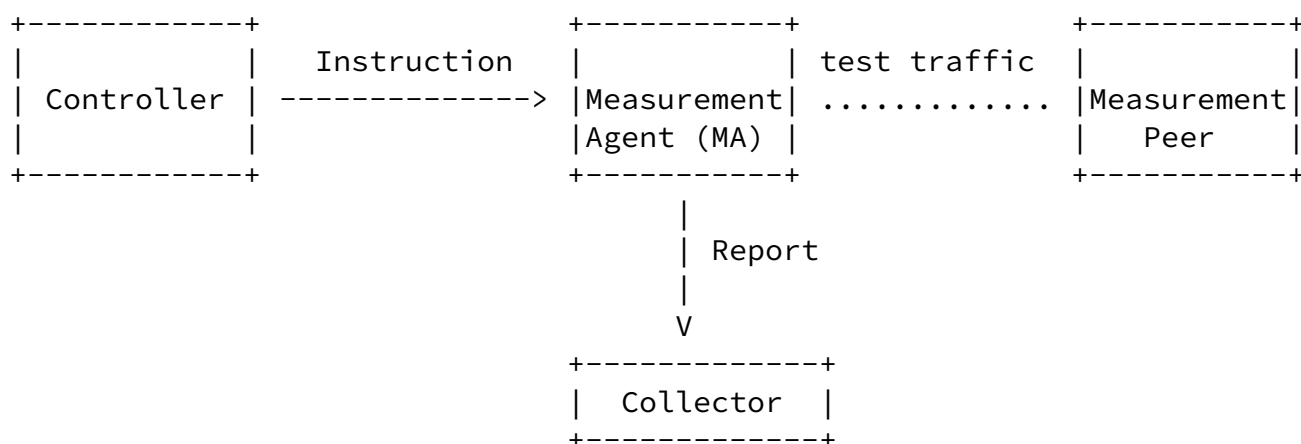


Figure 1: Schematic of main elements of LMAP framework

[3.](#) Constraints

[3.1.](#) Measurement system is under the direction of a single organisation

In the LMAP framework (as defined in the WG's charter) the measurement system is under the direction of a single organisation that is responsible both for the data and the quality of experience delivered to its users. 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.

Note that different LMAP measurement systems may overlap, in the sense that the active measurement packets of one measurement system appear along with normal user traffic to another measurement system. For instance, imagine an operator with an MA on the home gateway and an end user with an MA on their laptop. Rather than making separate measurements, an organisation might share its measurement data, or a suitably anonymised version of it, with another organisation. However, any form of coordination between different organisation 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 LMAP.

[3.2.](#) Each MA may only have a single Controller at any point in time

The constraint avoids different Controllers giving a MA conflicting instructions and so means that the MA does not have to manage contention between multiple Test (or Report) Schedules. This simplifies the design of MAs (critical for a large-scale infrastructure) and allows a Test Schedule to be tested on specific types of MA before deployment to ensure that the home user experience

is not impacted (due to CPU, memory or broadband-product constraints).

An operator may have several Controllers, perhaps with a Controller for different types of MA (home gateways, tablets) or location (Ipswich, Edinburgh).

To avoid problems with NAT and firewalls, it is likely that the MA 'pulls' the configuration from its Controller, as identified by the Initialiser.

- o Open issue: Should there be negotiation between a Controller and its MA, or should the Controller simply instruct the MA by sending its Test and Report Schedules?
 - * The argument for negotiation is that occasionally the MA may be updated with enhanced versions of existing tests. It is easier for the Controller to learn the MAs capabilities directly from the MA than from a management system. It avoids any mis-synchronisation.
 - * The argument against negotiation is that it makes the Controller-MA protocol more complicated, increases the MAs resource requirements and increases the complexity of the Controller when it decides how to schedule tests across numerous MAs or when it deploys a new Test Schedule to potentially millions of MAs.
- o Open issue: what happens if a Controller fails, how is the MA is homed onto a new one?

[3.3.](#) A Measurement Agent acts autonomously

Once the MA gets its Test and Report Schedules from its Controller then it acts autonomously, in terms of operation of the tests and reporting of the result.

Firstly, this means that the MA initiates Measurement Tasks. For the typical case where the MA is on a home gateway or edge device, this means that the MA initiates a 'download speed test' by asking a

Measurement Peer to send the file. The main rationale is that, for a

test that should be performed when there is no user traffic on the link, the MA knows whether the end user is active and therefore whether to start the test or delay it. Having the Schedule on the MA also avoids it having to check frequently with the Controller. Further, if the MA is behind a NAT then the Measurement Peer naturally learns its public-facing IP address.

Secondly, it is useful for the MA and perhaps the Measurement Peer to make some 'admission control' checks at the initiation of the Measurement Task to ensure that desired test conditions are present. The exchange of initialization packets between the MA and Measurement Peer ensures basic connectivity between them. Also, the MA may delay Measurement Task may if the associated subscriber is active, or the Measurement Peer may reject a testing request if it is overloaded. It has also been suggested that, in extremis, the Controller may want the ability to send a Measurement Suppression message to an MA, which causes the Measurement Tasks to be temporarily stopped.

Last, it is easier to secure the reporting process, for example with a unique certificate for each MA-Collector pair, so that the Collector is confident the results really do originate from the MA. All measurement results are sent from the MA.

4. Work items for LMAP WG

This Section considers the work that the LMAP working group needs to tackle. [Section 5](#) considers other work that needs doing that would be beyond the scope of the LMAP WG.

The main work items are:

- o Information Model, the abstract definition of the information carried from the Controller to the MA and the information carried from the MA to the Collector.
- o Control protocol and the associated data model: The definition of how instructions are delivered from a Controller to a MA; this includes a Data Model consistent with the Information Model plus a transport protocol.
- o Report protocol and the associated data model: The definition of how the Report is delivered from a MA to a Collector; this includes a Data Model consistent with the Information Model plus a transport protocol.

[4.1.](#) Information Model

The Information Model provides a protocol and device independent view of the information carried from the Controller to the MA and the information carried from the MA to the Collector. It can be implemented via a Control Protocol and Report Protocol, as defined by the LMAP WG. It is also possible that other Control and Report Protocols could be defined by other standards bodies or proprietary, however it is important that they all implement the same Information Model, in order to ease the definition, operation and interoperability of large-scale measurement systems.

The Information Model also includes information that is pre-configured on the MA in order that it can start communicating with a Controller.

An initial proposal for the Information Model is in [\[information-model\]](#).

The Information Model is divided into two main parts, each of which may be broken down into sub-parts:

- o information about the Instruction: Information that is received by the MA from the Controller pertaining to the measurement and reporting configuration. This includes:
 - * what measurements to do: the Measurement Task could be defined by reference to a registry entry, along with any parameters that need to be set (such as the address of the Measurement Peer) and any Environmental Constraint (such as, delay the test if the end user is active)
 - * when to do them: the Measurement Schedule details the timings of regular tests, one-off tests, and if regularly tests should be temporarily suppressed
 - * how to report the Measurement Results: via Reporting Channel(s), each of which defines a target Collector and Report Schedule
- o information about the Report: Information transmitted from the MA to the Collector including measurement results and the context in which they were conducted. This includes:
 - * the MAs identifier, or perhaps a Group-ID to anonymise results

- * the actual Measurement Results, including the time they were measured

- * the details of the Measurement Task (to avoid the Collector having to ask the Controller for this information later)

It is important to consider how to divide the Information Model into (sub-)parts, so that each (sub-)part can be updated independently at different times and regularities, as discussed in [[information-model](#)]

[4.2.](#) Control Protocol

The Control protocol and its associated data model define how instructions are delivered from a Controller to a MA; this includes a Data Model consistent with the Information Model plus a transport protocol. This may be a simple instruction - response protocol, and LMAP will specify how it operates over an existing protocol (to be selected, perhaps REST-style HTTP(s) or NETCONF).

[4.3.](#) Report Protocol

The Report protocol and the associated data model: The definition of how the Report is delivered from a MA to a Collector; this includes a Data Model consistent with the Information Model plus a transport protocol (to be selected, perhaps REST-style HTTP(s) or IPFIX).

[5.](#) Related work required but out of scope of LMAP

This section considers the items that need to be agreed between deployers of large-scale measurement systems, but that are out of scope of the LMAP WG ([Section 4](#) considers items within its scope).

[5.1.](#) Standard measurement tests

Standardised methods are needed for each metric that is measured. A registry for commonly-used metrics [[registry](#)] is also required, so that a test can be defined simply by its identifier in the registry. The methods and registry would hopefully also be referenced by other standards organisations.

- o Such activities are in scope of the IPPM working group (possibly

re-chartered) and not LMAP.

A new (or revised) test may need to be uploaded to MAs. How this is done is out of scope of the IETF; it could be as a firmware upgrade for a home hub, or new app for a PC, etc and may be device-specific.

[5.2.](#) Characterisation plan

Each organisation operating an LMAP system and collecting measurements for comparison purposes needs to conduct the same measurements according to the same sampling plan (ie size and schedule) and make the results available in the same format. The scope of comparison determines the set of organisations needing to agree on the common characterisation plan; for example those falling within the same regulatory environment in a particular country or region. Such agreements are certainly facilitated by IETF's work, but the details of the plan are beyond the scope of work in IETF.

[5.3.](#) Other elements

Other elements may be useful within the context of a large scale measurement system and worthy of standardisation, but are outside the scope of the LMAP WG: Initialiser, Subscriber Parameter Database, Results Database, Data Analysis Tools and operator's OAM.

An Initialiser configures a MA with details about its Controller, including authentication credentials. A bootstrap protocol is likely to be technology specific and so for different types of device could be defined by the Broadband Forum, DOCSIS or IEEE. Possible protocols are SNMP, NETCONF or (for Home Gateways) CPE WAN Management Protocol (CWMP) from the Auto Configuration Server (ACS) (as specified in TR-069).

A Subscriber Parameter Database contains information about the line, for example the customer's broadband contract (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 are all factors which 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. Another example is if the Controller wants to run a one-off test to diagnose a fault, then it should understand what problem the customer is experiencing and what tests have already been run. The subscribers' service parameters are already gathered and stored by existing operations systems.

A Results Database records all measurements in an equivalent form, for example an SQL database [[schulzrinne](#)], so that they can be easily accessed by the Data Analysis Tools whilst the LMAP system implementor can choose local solutions for each component. The Data Analysis Tools also need to understand subscriber service information, for example the broadband contract.

The Data Analysis Tools receive the results from the Collector or via

the Results Database. They might visualise the data or identify which component or link is likely to be the cause of a fault or degradation.

The operator's OAM (Operations, Administration and Management) uses the results from the tools.

[6.](#) IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

[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.

We assume that each Measurement Agent will receive test configuration, scheduling and reporting instructions from a single organisation (operator of 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 be prevented from replay. If a malicious party can gain control of the Measurement Agent they can use the MA capabilities to launch DoS attacks at targets, reduce the network user experience and corrupt the measurement results that are reported to the Collector. By altering the tests that are operated and/or the Collector address 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 reporting of the MA must also be secured to maintain confidentiality. The results must be encrypted such that only the authorised Collector can decrypt the results to prevent the leakage of confidential or private information. In addition it must be authenticated that the results have come from the expected MA and that they have not been tampered with. It must not be possible to spoof an MA to inject falsified data into the measurement platform or to corrupt the results of a real MA.

Availability should also be considered. 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 continue to operate an incorrect test schedule or fail to initiate.

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. This potential issue is currently handled by a code of conduct. It is outside the scope of the LMAP WG to consider the issue.

Concerning privacy and data protection, the role of the LMAP framework should be to ensure that only authorised data is collected and that this data is returned securely to the framework operator. Data should be stored securely and onward sharing of data to other parties should be controlled according to local data protection regulations. Depending upon the ownership/placement of the MA, local data protection laws, the tests being operated and existing user

agreements, it is possible that additional consent may need to be secured from parties such as the home broadband user. Having the measurement system under the direction of a single organisation clarifies the responsibility for data protection.

The next versions of [[lmap-yang](#)] and [[lmap-ipfix](#)] will also include further consideration of security.

[8.](#) Acknowledgements

Thanks to numerous people for much discussion, directly and on the LMAP list. This document tries to capture the current conclusions.

Philip Eardley and Trevor Burbridge 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.

[9.](#) Changes

[9.1.](#) from -00 to -01

aligned with terminology in [draft-eardley-lmap-terminology](#)

introduced aspects mentioned in the LMAP WG charter

introduced aspects from the Information model in [draft-burbridge-lmap-information-model](#)

Eardley, et al.

Expires January 16, 2014

[Page 13]

Internet-Draft

LMAP Framework

July 2013

[10.](#) Informative References

[RFC6241] "Network Configuration Protocol (NETCONF)",
<<http://tools.ietf.org/html/rfc6241>>.

[information-model]

Burbridge, T., Eardley, P., Bagnulo, M., and J. Schoenwaelder, "Information Model for Large-Scale Measurement Platforms (LMAP)", <<http://tools.ietf.org/html/draft-burbridge-lmap-information-model>>.

[lmap-ipfix]
"An LMAP application for IPFIX",
<<http://tools.ietf.org/html/draft-bagnulo-lmap-ipfix>>.

[lmap-netconf]
"Considerations on using NETCONF with LMAP Measurement Agents",
<<http://tools.ietf.org/html/draft-schoenw-lmap-netconf>>.

[lmap-yang]
"A YANG Data Model for LMAP Measurement Agents",
<<http://tools.ietf.org/html/draft-schoenw-lmap-yang>>.

[registry]
"A registry for commonly used metrics. Independent registries", <<http://tools.ietf.org/html/draft-bagnulo-ippm-new-registry-independent>>.

[schulzrinne]
"Large-Scale Measurement of Broadband Performance: Use Cases, Architecture and Protocol Requirements", <<http://tools.ietf.org/html/draft-schulzrinne-lmap-requirements>>.

[use-cases]
"Large-Scale Broadband Measurement Use Cases",
<<http://tools.ietf.org/html/draft-linsner-lmap-use-cases>>.

[yang-api]
"YANG-API Protocol", <<http://tools.ietf.org/html/rfc6241>>.

Adastral Park, Martlesham Heath
Ipswich
ENGLAND

Email: philip.eardley@bt.com

Trevor Burbridge
British Telecom
Adastral Park, Martlesham Heath
Ipswich
ENGLAND

Email: trevor.burbridge@bt.com

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

Email: acmorton@att.com