Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 29, 2013

P. Eardley
T. Burbridge
BT
A. Morton
AT&T

February 25, 2013

A framework for large-scale measurements
draft-eardley-lmap-framework-01

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 discussions about the potential creation of the LMAP working group.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

[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

The LMAP framework should be useful for all these.

The key requirement is for large scale. A measurement system might have at least ~100k Measurement Agents.

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. 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 (eg 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.

This section forms the problem statement / aim of the proposed LMAP working group, in the context of the constraints and scope given below.

2. Outline of framework

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

- o Measurement Agent (MA)
- 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)

Two Measurement Agents (MAs) jointly perform an active measurement test, by generating test traffic and measuring some metric associated with its transfer over the path from one MA to the other; for example the time taken to transfer a ’test file’. Some tests may involve more than two MAs; for example, to measure ’latency under load’. A single MA may also conduct passive testing through the observation of traffic (ie without the involvement of a second MA); for example an end user’s mix of applications.

The MA functions are implemented either in specialised hardware or as code on general purpose devices like a PC, tablet or smartphone.

- o Comment: It may be useful to distinguish two types of MA - a ’complete MA’ that interacts with the Controller and Collector, and a ’remote MA’ that only takes part in active tests (and does not interact with the Controller and Collector). This is for further study.
- o Comment: A ’complete MA’ may want to run a test against a normal, non-LMAP device, for example 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 + 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 + 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.

- o Comment: 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. It is for further study whether such communications should be in or out of scope of LMAP.

- o Comment: 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.

An Initialiser configures a MA with details about its Controller, including authentication credentials. 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 Test Schedule. 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.

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 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.

```
+------------+                +-----------+               +-----------+
|            | Test & Report  |           | test traffic  | (remote)  |
| Controller | ---------------> |Measurement| ............. |Measurement|
|            | Schedule       |Agent (MA) |               |Agent (MA) |
+------------+                +-----------+               +-----------+
                                    |
                                    | Report
                                    |
                                    V
                              +-------------+
                              |  Collector  |
                              +-------------+
```

Figure 1: Schematic of main elements of LMAP framework

3. Constraints

3.1. Measurement system is under the direction of a single organisation

Explanation: In the LMAP framework 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 privacy and network security.

Given the novelty of large-scale measurement efforts, the expectation is that inter-organization coordination is an out-of-band consideration. There could be scenarios where measurement data, or a suitably anonymised version of it, is shared between organisations, via their Data Analysis Tools for instance. Consideration is outside the scope of LMAP.

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

Explanation: 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, the MA 'pulls' the configuration from its Controller, as identified by the Initialiser.

- 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 with versions of existing tests. It is easier for the Controller to learn the MA's 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 MA's 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.
- 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 tests. 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 remote MA 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 remote MA naturally learns its public-facing IP address.

Secondly, 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 required in LMAP

This Section considers the work that the prospective LMAP working group would tackle. Section 5 considers other work that needs doing that would be beyond the scope of the LMAP WG.

4.1. Defining the Test and Report Schedules

The Test and Report Schedules contain the instructions sent by the Controller to the MA. The Schedules could be combined into a single Schedule, this is for further study.

The Test Schedule would include things like:

- Which tests to operate and (if applicable) to which remote MAs
- The testing schedule
- Any test or environmental parameters
- How to react to the presence of user or other test traffic (if not inherent in the test design)

The Report Schedule would include things like:

- How often to report results (e.g. time or volume of data)
- Where to report results
- What to do if reporting fails

Defining Test and Report Schedule(s) is in scope of LMAP:

- o Information model: the abstract definition of the Schedule

- o Data model: which instantiates the information model in a particular language. It could be done using an existing IETF data modeling language, for example YANG as sketched in [lmap-yang].

- o Protocol: how the Test and Report Schedule(s) are delivered from the Controller to the MA. Possibilities would include NETCONF [RFC6241] as discussed in [lmap-netconf] or a RESTful interface [yang-api].

4.2. Defining the Report

The Report contain the measurement results sent by the Controller to the MA. The Report includes things like:

- o The results of the test (typically at least all the singleton measurements, including the time they were measured)

- o The MA's identifier

- o The test and its parameters (essentially an 'echo' of the Test Schedule with the parameters actually used, which avoids the Collector having to ask the Controller).

Defining the Report is in scope of LMAP:

- o Information model: the abstract definition of the Report

- o Data model: which instantiates the information model in a particular language. It could be done using an existing IETF data modeling language, for example IPFIX, as sketched in [lmap-ipfix] or

- o Protocol: how the Report is delivered from the MA to the Collector. Possibilities would include IPFIX, as sketched in [lmap-ipfix] or a RESTful interface.

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 prospective 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

The other elements discussed in Section 2 may also benefit from standardisation: Initialiser, Subscriber Parameter Database, Results Database, Data Analysis Tools and operator's OAM.

The Initialiser-MA protocol is likely to be technology specific and so for different types of device could be defined by the Broadband Forum, DOCSIS or IEEE. The Data Analysis Tools and operator's OAM are also beyond the scope of the IETF. For the Subscriber Parameter Database and Results Database, it is possible that there could be work to define a data model - it is suggested that this is for later study and should be out of the initial scope of IETF work.

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.

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. Informative References

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

[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>.

Authors' Addresses

Philip Eardley
BT

Trevor Burbridge
BT

Al Morton
AT&T