IPPM Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 4, 2014

A. Akhter
B. Claise
Cisco Systems, Inc.
March 3, 2014

Passive Performance Metrics Sub-Registry
draft-akhter-ippm-registry-passive-01.txt

Abstract

This memo defines the Passive Performance Metrics sub-registry of the Performance Metric Registry. This sub-registry will contain Passive Performance Metrics, especially those defined in RFCs prepared in the IP Performance Metrics (IPPM) Working Group of the IETF, and possibly applicable to other IETF metrics.

IPPM Passive metric registration is meant to allow wider adoption of common metrics in an inter-operable way. There are challenges with metric interoperability and adoption (to name a few) due to flexible input parameters, confusion between many similar metrics, and varying output formats.

This memo proposes a way to organize registry entries into columns that are well-defined, permitting consistent development of entries over time (a column may be marked NA if it is not applicable for that metric). The design is intended to foster development of registry entries based on existing reference RFCs, whilst each column serves as a check-list item to avoid omissions during the registration process. Every entry in the registry, before IANA action, requires Expert review as defined by concurrent IETF work in progress "Registry for Performance Metrics" (draft-manyfolks-ippm-metric-registry).

The document contains example entries for the Passive Performance Metrics sub-registry: a registry entry for a passive metric based on octetTotalCount as defined in RFC5102 and a protocol specific passive metric based on RTP packets lost as defined in RFC3550. The examples are for Informational purposes and do not create any entry in the IANA registry.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The IETF has been specifying and continues to specify Performance Metrics. While IP Performance Metrics (IPPM) is the working group (WG) primarily focusing on Performance Metrics definition at the IETF, other working groups, have also specified Performance Metrics:

> The "Metric Blocks for use with RTCP's Extended Report Framework" [XRBLOCK] WG recently specified many Performance Metrics related to "RTP Control Protocol Extended Reports (RTCP XR)" [RFC3611], which establishes a framework to allow new information to be conveyed in RTCP, supplementing the original report blocks defined

in "RTP: A Transport Protocol for Real-Time Applications", [RFC3550].

The Benchmarking Methodology" [BMWG] WG proposed some Performance Metrics as part of the benchmarking methodology.

The IP Flow Information eXport WG (IPFIX) [IPFIX] has existing and proposed Information Elements related to performance metrics.

The Performance Metrics for Other Layers (PMOL) [PMOL], a concluded working group, defined some Performance Metrics related to Session Initiation Protocol (SIP) voice quality [RFC6035], as well as guidelines for defining performance metrics [RFC6390]

It is expected that more and more Performance Metrics will be defined in the future, not only IP based metrics, but also protocol-specific ones and application-specific ones.

However, there is currently no Performance Metrics registry in IANA. "Registry for Performance Metrics" [I-D.manyfolks-ippm-metric-registry] defines a common registry for metrics. The registry proposes the creation of two sub-registries, one for active metrics and another for passive measurements.

There is a sister document for the active metric sub-registry in "Active Performance Metric Sub-Registry" [I-D.mornuley-ippm-registry-active].

This document defines the Passive Performance Measurements Sub-Registry of the Performance Metric Registry. This sub-registry will contain passive performance metrics that meet the criteria set by the IETF and review of the Performance Metric Experts. It is expected that the majority of the metrics will have been defined elsewhere within the IETF working groups such as IPPM, BMWG, IPFIX, etc.

This sub-registry is part of the Performance Metric Registry [I-D.manyfolks-ippm-metric-registry] which specifies that all sub-registries must contain at least the following common fields: the identifier, the name, the status, the requester, the revision, the revision date, the description for each entry, and the reference specifications used as the foundation for the Registered Performance Metric (see [I-D.manyfolks-ippm-metric-registry]). In addition to these common fields the passive metrics sub-registry has additional fields that provide the necessary background for interoperability and adoption.

2. Background and Motivation:

(from draft-mornuley-ippm-registry-active):

One clear motivation for having such a registry is to allow a controller to request a measurement agent to perform a measurement using a specific metric (see [I-D.ietf-lmap-framework]). Such a request can be performed using any control protocol that refers to the value assigned to the specific metric in the registry. Similarly, the measurement agent can report the results of the measurement and by referring to the metric value it can unequivocally identify the metric that the results correspond to.

There are several side benefits of having a registry with well-chosen entries. First, the registry could serve as an inventory of useful and used metrics that are normally supported by different implementations of measurement agents. Second, the results of the metrics would be comparable even if they are performed by different implementations and in different networks, as the metric and method is unambiguously defined.

3. Scope

[I-D.manyfolks-ippm-metric-registry] defines the overall structure for a Performance Metric Registry and provides guidance for defining a sub registry.

This document defines the Passive Performance Metrics Sub-registry; passive metrics are those where the measurements are based the observation of on user traffic. Specifically, this traffic has not been generated for the purpose of measurement.

A row in the registry corresponds to one Registered Performance Metric, with entries in the various columns specifying the metric. Section 4 defines the additional columns for a Registered Passive Performance Metric.

As discussed in [I-D.manyfolks-ippm-metric-registry], each entry (row) must be tightly defined; the definition must leave open only a few parameters that do not change the fundamental nature of the measurement (such as source and destination addresses), and so promotes comparable results across independent implementations. Also, each registered entry must be based on existing reference RFCs (or other standards) for performance metrics, and must be operationally useful and have significant industry interest. This is ensured by expert review for every entry before IANA action.

4. Passive Registry Categories and Columns

This section defines the categories and columns of the registry. Below, categories are described at the 4.x heading level, and columns are at the 4.x.y heading level. There are three categories, divided into common information (from [I-D.manyfolks-ippm-metric-registry]), metric definition and an open Comments section.

4.1. Common Registry Indexes and Information

This category has multiple indexes to each registry entry. It is defined in [I-D.manyfolks-ippm-metric-registry]:

4.1.1. Identifier

Defined in [I-D.manyfolks-ippm-metric-registry]. Definition text to be copied once source is stable.

4.1.2. Name

Defined in [I-D.manyfolks-ippm-metric-registry], same comment as above.

4.1.3. Status

Defined in [I-D.manyfolks-ippm-metric-registry], same comment as above.

4.1.4. Requester

Defined in [I-D.manyfolks-ippm-metric-registry], same comment as above.

4.1.5. Revision

Defined in [I-D.manyfolks-ippm-metric-registry], same comment as above.

4.1.6. Revision Date

Defined in [I-D.manyfolks-ippm-metric-registry], same comment as above.

4.1.7. Description

Defined in [I-D.manyfolks-ippm-metric-registry], same comment as the above.

4.1.8. Reference Specification(s)

Defined in [I-D.manyfolks-ippm-metric-registry], same comment as the above.

4.2. Metric Definition

This category includes columns to prompt all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters, which are left open in the origin definition but have a particular value defined by the performance metric.

4.2.1. Reference Definition

This entry provides references to relevant sections of the RFC(s) defining the metric, as well as any supplemental information needed to ensure an unambiguous definition for implementations.

4.2.2. Fixed Parameters

Fixed Parameters are input factors whose value must be specified in the Registry. The measurement system uses these values.

Where referenced metrics supply a list of Parameters as part of their descriptive template, a sub-set of the Parameters will be designated as Fixed Parameters. For example, for RTP packet loss calculation relies on the validation of a packet as RTP which is a multi-packet validation controlled by MIN_SEQUENTIAL as defined by [RFC3550]. Varying MIN_SEQUENTIAL values can alter the loss report and this value could be set as a fixed parameter.

A Parameter which is Fixed for one Registry entry may be designated as a Run-time Parameter for another Registry entry.

4.3. Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous method for implementations.

4.3.1. Reference Implementation

This entry provides references to relevant sections of the RFC(s) describing the method of measurement, as well as any supplemental information needed to ensure unambiguous interpretation for implementations referring to the RFC text.

Specifically, this section should include pointers to pseudocode or actual code that could be used for an unambigious implementation.

4.3.2. Traffic Filter Criteria

The filter specifies the traffic constraints that the passive measurement method used is valid (or invalid) for. This includes valid packet sampling ranges, width of valid traffic matches (eg. all traffic on interface, UDP packets packets in a flow (eg. same RTP session).

It is possible that the measurement method may not have a specific limitation. However, this specific registry entry with it’s combination of fixed parameters implies restrictions. These restrictions would be listed in this field.

4.3.3. Measurement Timing

Measurement timing defines the behavior of the measurement method with respect to timing.

Is the measurement continuous?

If the measurement is sampled, what is the format of sampling? (eg random packet, random time, etc.)

How long is the measurement interval?

4.3.4. Output Type(s) and Data Format

For entries which involve a stream and many singleton measurements, a statistic may be specified in this column to summarize the results to a single value. If the complete set of measured singletons is output, this will be specified here.

Some metrics embed one specific statistic in the reference metric definition, while others allow several output types or statistics.

Each entry in the output type column contains the following information:

- Value: The name of the output type
- Data Format: provided to simplify the communication with collection systems and implementation of measurement devices.
- Reference: the specification where the output type is defined

The output type defines the type of result that the metric produces. It can be the raw result(s) or it can be some form of statistic. The specification of the output type must define the format of the output. In some systems, format specifications will simplify both measurement implementation and collection/storage tasks. Note that if two different statistics are required from a single measurement (for example, both "Xth percentile mean" and "Raw"), then a new output type must be defined ("Xth percentile mean AND Raw").

4.3.5. Metric Units

The measured results must be expressed using some standard dimension or units of measure. This column provides the units.

When a sample of singletons (see [RFC2330] for definitions of these terms) is collected, this entry will specify the units for each measured value.

4.3.6. Run-time Parameters and Data Format

Run-Time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete. However, the values of these parameters is not specified in the Registry, rather these parameters are listed as an aid to the measurement system implementor or user (they must be left as variables, and supplied on execution).

Where metrics supply a list of Parameters as part of their descriptive template, a sub-set of the Parameters will be designated as Run-Time Parameters.

The Data Format of each Run-time Parameter SHALL be specified in this column, to simplify the control and implementation of measurement devices.

Examples of Run-time Parameters include IP addresses, measurement point designations, start times and end times for measurement, and other information essential to the method of measurement.

4.4. Comments and Remarks

Besides providing additional details which do not appear in other categories, this open Category (single column) allows for unforeseen issues to be addressed by simply updating this Informational entry.

5. Example Generalized Passive Octet Count Entry

tbd

This section is Informational.

This section gives an example registry entry for a generalized the passive metric octetDeltaCount described in [RFC5102].

5.1. Registry Indexes

This category includes multiple indexes to the registry entries, the element ID and metric name.

5.1.1. Element Identifier

An integer having enough digits to uniquely identify each entry in the Registry.

TBD by IANA.

5.1.2. Metric Name

A metric naming convention is TBD.

One possibility based on the proposal in [I-D.manyfolks-ippm-metric-registry]:

Pas_IP-Octet-Delta-General

5.1.3. Status

Current

5.1.4. Requester

TBD

5.1.5. Revision

0

5.1.6. Revision Date

TBD

5.1.7. Metric Description

A delta count of the number of octets observed.

5.1.8. Reference Specification(s)

octetDeltaCount described in section 5.10.1 of [RFC5102]

5.2. Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

5.2.1. Reference Definition

octetDeltaCount described in section 5.10.1 of [RFC5102]

5.2.2. Fixed Parameters

As this is the generalised version of the IP delta count metric, there are no fixed parameters.

5.3. Method of Measurement

5.3.1. Reference Implementation

For <metric>.

<section reference>

5.3.2. Traffic Filter Criteria

This measurement only covers IP packets and the IP payload (including the IP header) of these packets. Non-IP packets (BPDUs, ISIS) will not be accounted. Layer 2 overhead (Ethernet headers, MPLS, QinQ, etc.) will also not be represented in the measurement.

5.3.3. Measurement Timing

This is a continous measurement of the IP octets seen in the traffic selection scope (run-time parameter).

The measurement interval is a run time parameter.

There is no sampling.

5.3.4. Output Type(s) and Data Format

It is possible that multiple observation intervals are reported in a single report. In such a case concatination of the interval reports (deltaOctetCount, start-time, end-time) is allowed.

The delta octet count metric reports a observation start time and end time.

- Value: observation-start-time and observation-end-time
- Data Format: 64-bit NTP Time-stamp Format
- Reference: section 6 of [RFC5905]

5.3.5. Metric Units

The measured results are expressed in octets with a data format of unsigned64 as described in [RFC5102]

5.3.6. Run-time Parameters and Data Format

Run-time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

- samplingTimeInterval, length of time a single report covers. unsigned32 microseconds [RFC5477]
- observationInterface, ifindex of interface to monitor. -1 represents all interfaces. -2 representings WAN facing and -3 represnets LAN facing. unsigned32.
- observation direction, unsigned8 where 0 represents incoming traffic on interface, 1 outgoing and 2 represents both incoming and outgoing.

5.4. Comments and Remarks

Additional (Informational) details for this entry

6. Example 5min Passive Egress Octet Count Entry on WAN Interface

tbd

This section is Informational.

This section gives an example registry entry for accounting of outgoing WAN IP traffic the passive metric in terms of octetDeltaCount, as described in [RFC5102].

6.1. Registry Indexes

This category includes multiple indexes to the registry entries, the element ID and metric name.

6.1.1. Element Identifier

An integer having enough digits to uniquely identify each entry in the Registry.

TBD by IANA.

6.1.2. Metric Name

A metric naming convention is TBD.

One possibility based on the proposal in [I-D.manyfolks-ippm-metric-registry]:

Pas_IP-Octet-Delta-WAN-egress

6.1.3. Status

Current

6.1.4. Requester

TBD

6.1.5. Revision

0

6.1.6. Revision Date

TBD

6.1.7. Metric Description

A delta count of the number of octets observed outgoing on WAN interface.

6.1.8. Reference Specification(s)

octetDeltaCount described in section 5.10.1 of [RFC5102]

6.2. Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

6.2.1. Reference Definition

octetDeltaCount described in section 5.10.1 of [RFC5102]

6.2.2. Fixed Parameters

As this is a specific version of Pas_IP-Octet-Delta-General that performs metering of all outgoing WAN traffic.

- samplingTimeInterval= 300000000, length of time a single report covers. unsigned32 microseconds [RFC5477]
- observationInterface= -2, ifindex of interface to monitor. -1 represents all interfaces. -2 representings WAN facing and -3 represnets LAN facing. unsigned32.
- observation direction= 1, unsigned8 where 0 represents incoming traffic on interface, 1 outgoing and 2 represents both incoming and outgoing.

6.3. Method of Measurement

6.3.1. Reference Implementation

For <metric>.

<section reference>

6.3.2. Traffic Filter Criteria

This measurement only covers IP packets observed in the WAN outgoing direction. The bytes counted are the IP payload (including the IP header) of these packets. Non-IP packets (BPDUs, ISIS) will not be accounted. Layer 2 overhead (Ethernet headers, MPLS, QinQ, etc.) will also not be represented in the measurement.

6.3.3. Measurement Timing

This is a continous measurement of the IP octets seen in the traffic selection scope (run-time parameter), each of a 5 minute duration.

There is no sampling.

6.3.4. Output Type(s) and Data Format

It is possible that multiple observation intervals are reported in a single report. In such a case concatination of the interval reports (deltaOctetCount, start-time, end-time) is allowed.

The delta octet count metric reports a observation start time and end time.

- Value: observation-start-time and observation-end-time
- Data Format: 64-bit NTP Time-stamp Format
- Reference: section 6 of [RFC5905]

6.3.5. Metric Units

The measured results are expressed in octets with a data format of unsigned64 as described in [RFC5102]

6.3.6. Run-time Parameters and Data Format

There are no run-time parameters for this registry entry.

6.4. Comments and Remarks

Additional (Informational) details for this entry

7. Example Passive RTP Lost Packet Count

tbd

8. Example BLANK Registry Entry

This section is Informational. (?)

This section gives an example registry entry for the <type of metric and specification reference> .

8.1. Registry Indexes

This category includes multiple indexes to the registry entries, the element ID and metric name.

8.1.1. Element Identifier

An integer having enough digits to uniquely identify each entry in the Registry.

8.1.2. Metric Name

A metric naming convention is TBD.

8.1.3. Status

Current

8.1.4. Requester

TBD

8.1.5. Revision

0

8.1.6. Revision Date

TBD

8.1.7. Metric Description

A metric Description is TBD.

8.1.8. Reference Specification(s)

Section YY, RFCXXXX

8.2. Metric Definition

8.2.1. Reference Definition

< possible section reference>

8.2.2. Fixed Parameters

Fixed Parameters are input factors that must be determined and embedded in the measurement system for use when needed. The values of these parameters is specified in the Registry.

<list fixed parameters>

8.3. Method of Measurement

8.3.1. Reference Implementation

For <metric>.

<section reference>

8.3.2. Traffic Filter Criteria

<list filter criteria limitations and allowances >

8.3.3. Measurement Timing

< list timing requirements and limitations >

8.3.4. Output Type(s) and Data Format

The output types define the type of result that the metric produces.

- Value:
- Data Format: (There may be some precedent to follow here, but otherwise use 64-bit NTP Time-stamp Format, see section 6 of [RFC5905]).
- Reference: <section reference>

8.3.5. Metric Units

The measured results are expressed in <units>,

<section reference>.

8.3.6. Run-time Parameters and Data Format

Run-time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

<list of run-time parameters>

<reference(s)>.

8.4. Comments and Remarks

Additional (Informational) details for this entry

9. Security Considerations

This registry has no known implications on Internet Security.

10. IANA Considerations

IANA is requested to create The Passive Performance Metric Sub-registry within the Performance Metric Registry defined in [I-D.manyfolks-ippm-metric-registry]. The Sub-registry will contain the following categories and (bullet) columns, (as defined in section 3 above):

Common Registry Indexes and Info

- Identifier
- Name
- Status
- Requester
- Revision
- Revision Date
- Description
- Reference Specification(s)

Metric Definition

- Reference Definition
- Fixed Parameters

Method of Measurement

- Reference Implementation

- o Traffic Filter Criteria
- o Measurement Timing
- o Output Type(s) and Data format
- o Metric Units
- o Run-time Parameters

Comments and Remarks

11. Acknowledgements

The authors thank the prior work of Al Morton, Marcelo Bagnulo and Phil Eardley in "draft-mornuley-ippm-registry-active" which was used both as a template for this document and source of text.

12. References

12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

12.2. Informative References

[BMWG] IETF, , "Benchmarking Methodology (BMWG) Working Group, http://datatracker.ietf.org/wg/bmwg/charter/", .

[I-D.ietf-lmap-framework] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A framework for large-scale measurement platforms (LMAP)", draft-ietf-lmap-framework-03 (work in progress), January 2014.

[I-D.manyfolks-ippm-metric-registry] Bagnulo, M., Claise, B., Eardley, P., and A. Morton, "Registry for Performance Metrics", draft-manyfolks-ippm-metric-registry-00 (work in progress), February 2014.

[I-D.mornuley-ippm-registry-active] Morton, A., Bagnulo, M., and P. Eardley, "Active Performance Metric Sub-Registry", draft-mornuley-ippm-registry-active-00 (work in progress), February 2014.

[IPFIX] IETF, , "IP Flow Information eXport (IPFIX) Working Group, http://datatracker.ietf.org/wg/ipfix/charter/", .

[PMOL] IETF, , "IP Performance Metrics for Other Layers (PMOL) Working Group, http://datatracker.ietf.org/wg/pmol/charter/", .

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.

[RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.

[RFC3611] Friedman, T., Caceres, R., and A. Clark, "RTP Control Protocol Extended Reports (RTCP XR)", RFC 3611, November 2003.

[RFC5102] Quittek, J., Bryant, S., Claise, B., Aitken, P., and J. Meyer, "Information Model for IP Flow Information Export", RFC 5102, January 2008.

[RFC5477] Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", RFC 5477, March 2009.

[RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.

[RFC6035] Pendleton, A., Clark, A., Johnston, A., and H. Sinnreich, "Session Initiation Protocol Event Package for Voice Quality Reporting", RFC 6035, November 2010.

[RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, October 2011.

[XRBLOCK] IETF, , "Metric Blocks for use with RTCP’s Extended Report Framework (XRBLOCK), http://datatracker.ietf.org/wg/xrblock/charter/", .

Authors’ Addresses

Aamer Akhter
Cisco Systems, Inc.
7025 Kit Creek Road
RTP, NC 27709
USA

Email: aakhter@cisco.com

Benoit Claise
Cisco Systems, Inc.
De Kleetlaan 6a b1
1831 Diegem
Belgium

Phone: +32 2 704 5622
Email: bclaise@cisco.com

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.

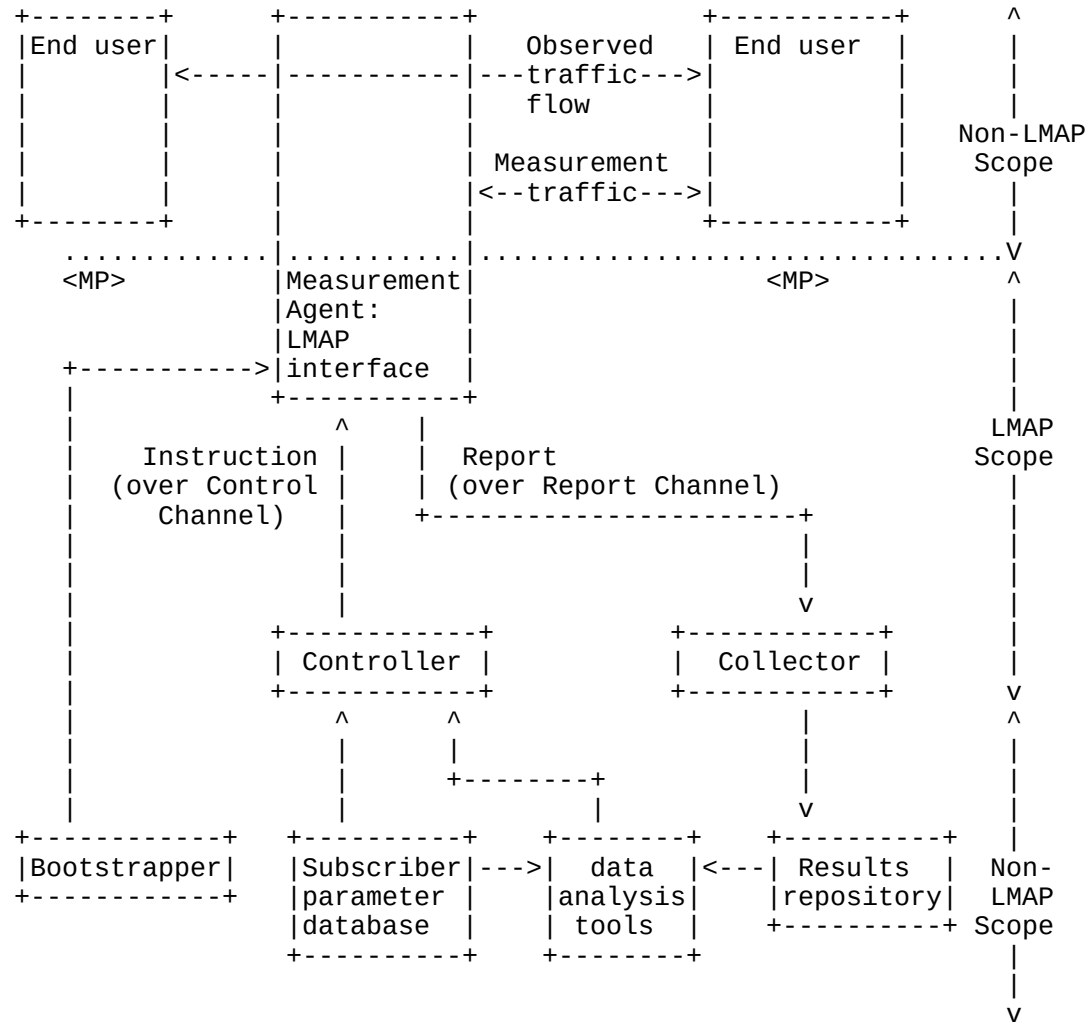

Schematic of main elements of an LMAP-based Measurement System (showing the elements in and out of the scope of initial LMAP work)

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.

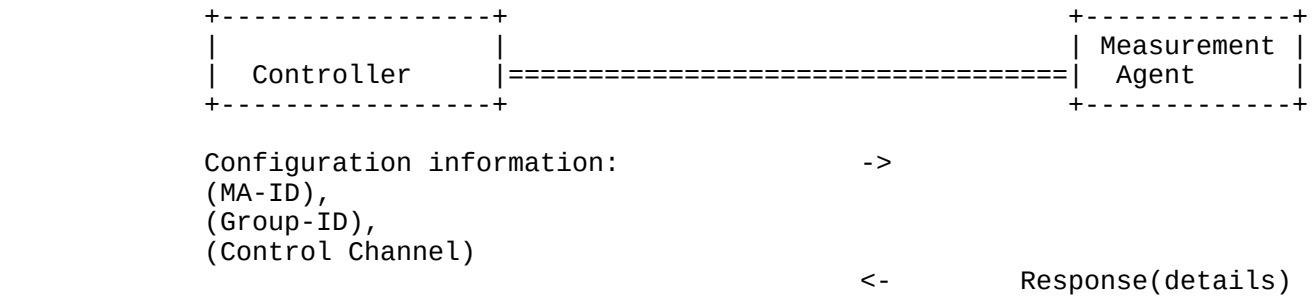

5.2.2. Instruction

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.

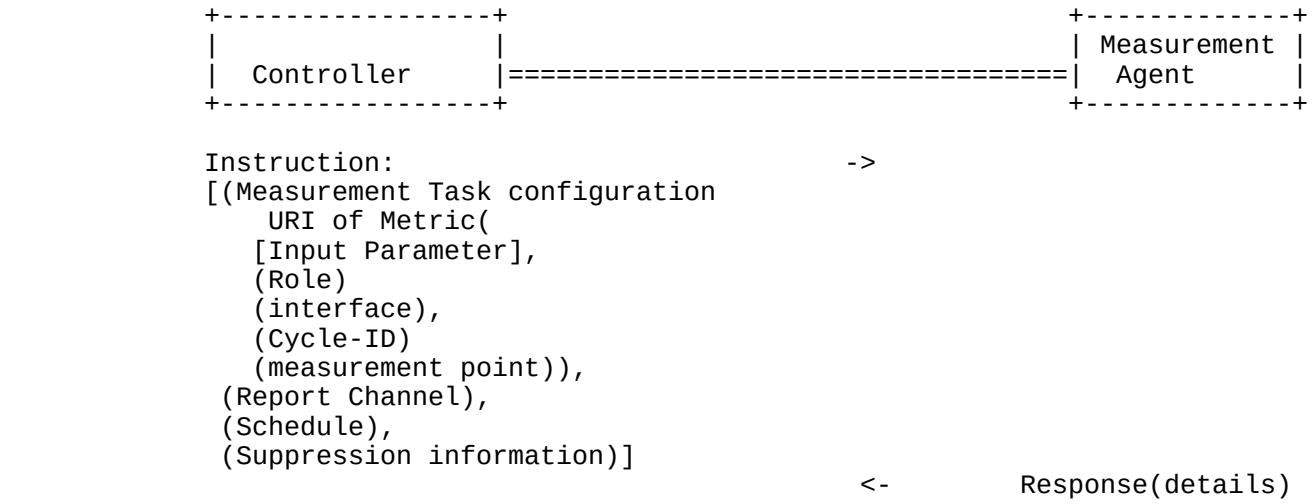

The Instruction defines information with the following aims ([I-D.ietf-lmap-information-model] defines the consequent list of information elements):

- 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).

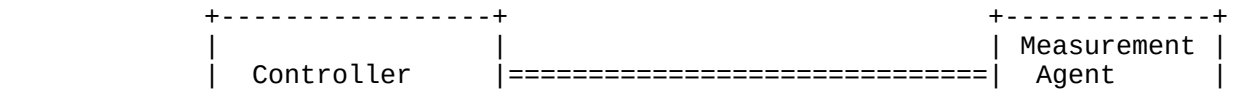


```
Suppress:
[(Measurement Task),                  ->
 (Measurement Schedule),
 [start time],
 [end time],
 [on-going suppressed?]]

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.

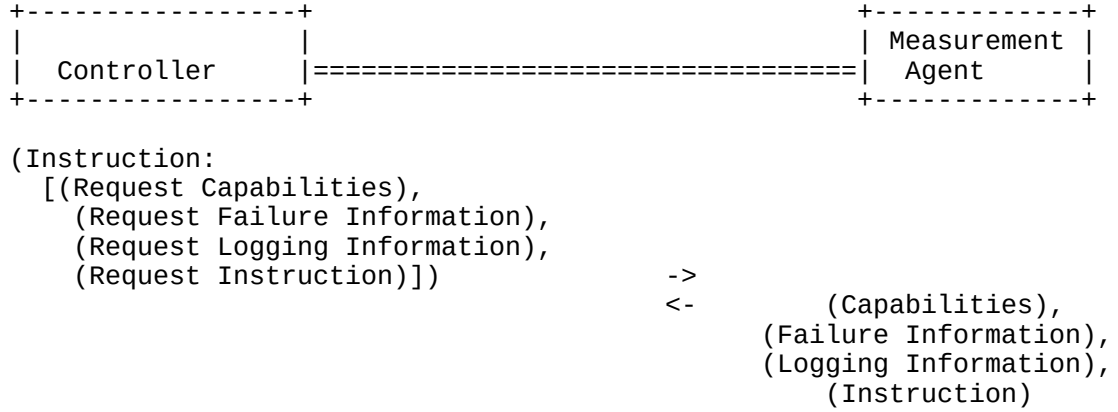

5.3. Operation of Measurement Tasks

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.

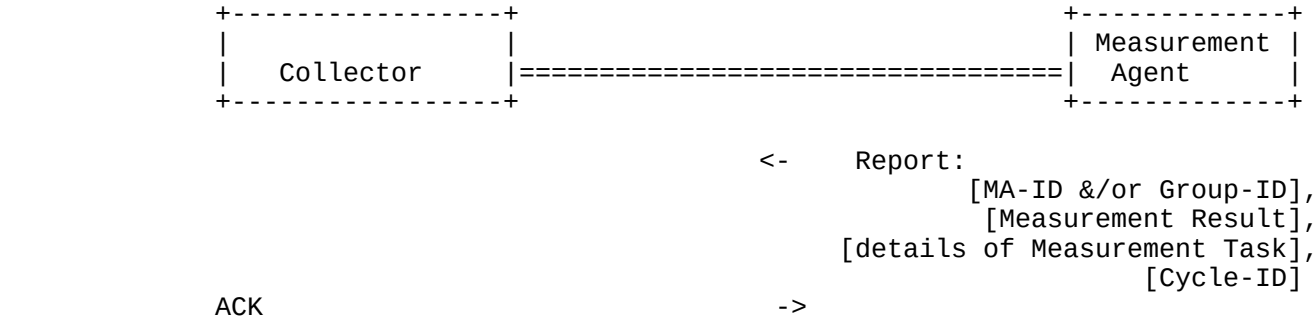

The Report contains:

- 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.

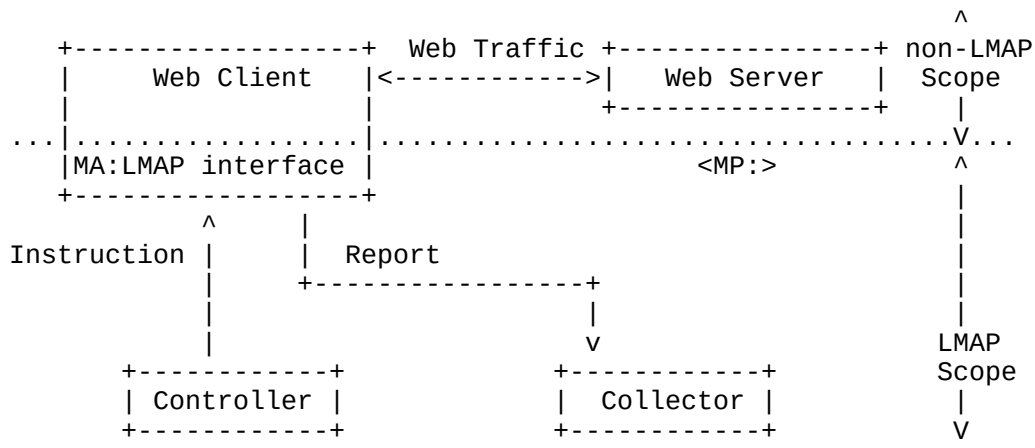

Schematic of LMAP-based Measurement System,
with Web server as Measurement Peer

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.

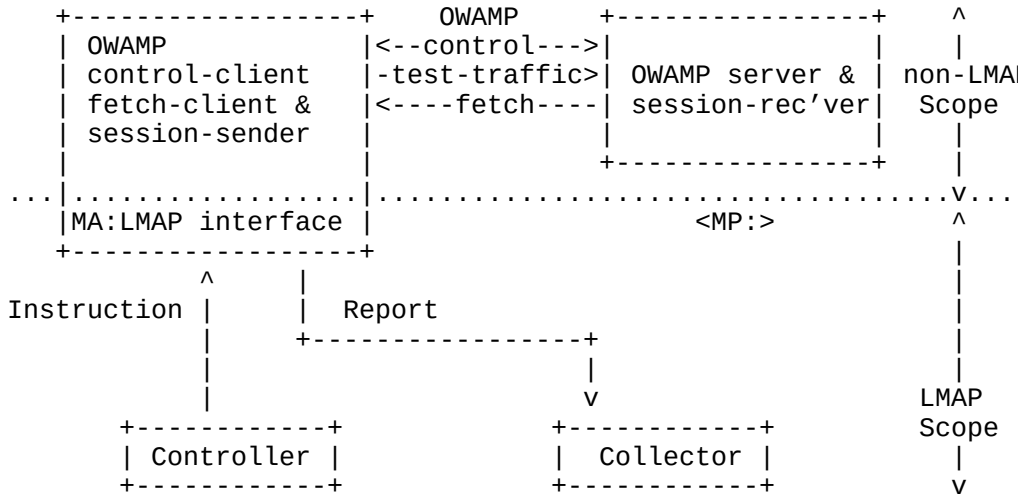

Schematic of LMAP-based Measurement System,
with OWAMP server as Measurement Peer

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.

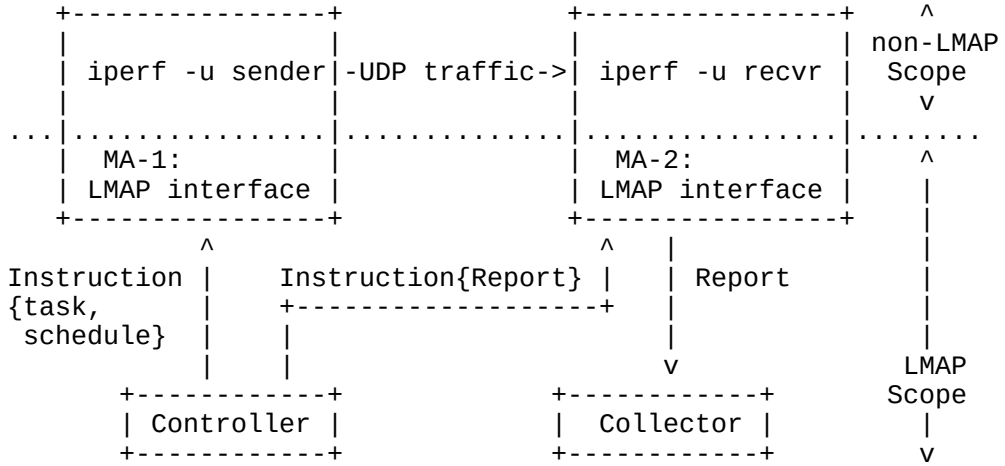

Schematic of LMAP-based Measurement System, with two
Measurement Agents cooperating to measure UDP traffic

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.

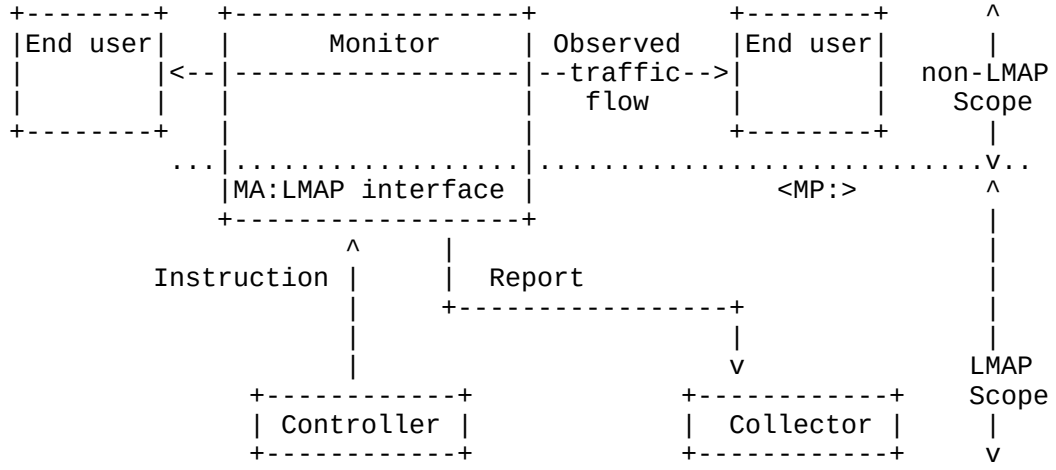

Schematic of LMAP-based Measurement System,
with a Measurement Agent monitoring traffic

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.)

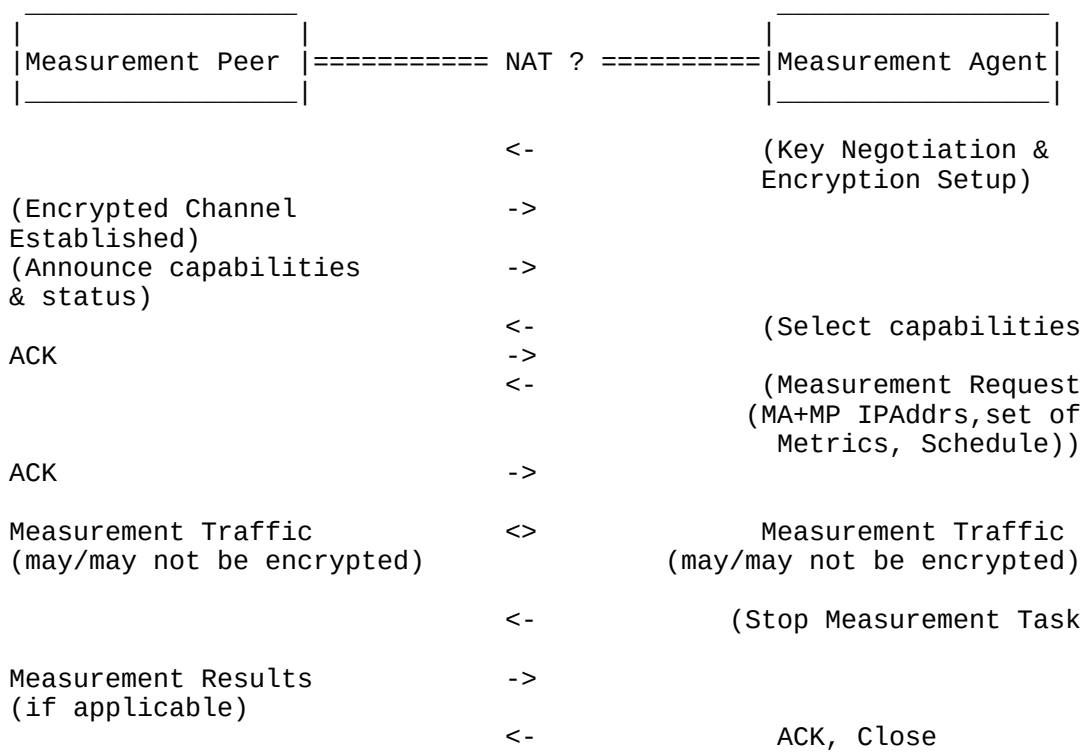

This exchange primarily exposes the IP addresses of measurement devices and the inference of measurement participation from such traffic. There may be sensitive information on key points in a service provider's network included. There may also be access to measurement-related information of interest 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.)

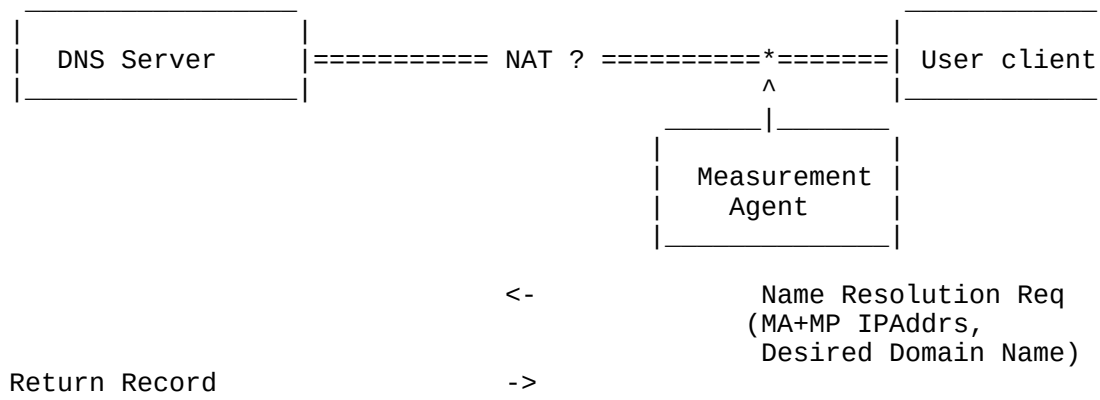

In this particular example, the MA monitors DNS messages in order to measure that DNS response time. The Measurement Agent may be embedded in the user host, or it may be located in another device capable of observing user traffic. The MA learns the IP addresses of measurement devices and the intent to 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

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 23, 2017

T. Burbridge
P. Eardley
BT
M. Bagnulo
Universidad Carlos III de Madrid
J. Schoenwaelder
Jacobs University Bremen
April 21, 2017

Information Model for Large-Scale Measurement Platforms (LMAP)
draft-ietf-lmap-information-model-18

Abstract

This Information Model applies to the Measurement Agent within a Large-Scale Measurement Platform. As such it outlines the information that is (pre-)configured on the Measurement Agent or exists in communications with a Controller or Collector within an LMAP framework. The purpose of such an Information Model is to provide a protocol and device independent view of the Measurement Agent that can be implemented via one or more Control and Report protocols.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of This Memo

Table of Contents

1. Introduction

A large-scale measurement platform is a collection of components that work in a coordinated fashion to perform measurements from a large number of vantage points. A typical use case is the execution of broadband measurements [RFC7536]. The main components of a large-scale measurement platform are the Measurement Agents (hereafter MAs), the Controller(s) and the Collector(s).

The MAs are the elements actually performing the measurements. The MAs are controlled by exactly one Controller at a time and the

Collectors gather the results generated by the MAs. In a nutshell, the normal operation of a large-scale measurement platform starts with the Controller instructing a set of one or more MAs to perform a set of one or more Measurement Tasks at a certain point in time. The MAs execute the instructions from a Controller, and once they have done so, they report the results of the measurements to one or more Collectors. The overall framework for a large-scale measurement platform as used in this document is described in detail in [RFC7594].

A large-scale measurement platform involves basically three types of protocols, namely, a Control protocol (or protocols) between a Controller and the MAs, a Report protocol (or protocols) between the MAs and the Collector(s) and several measurement protocols between the MAs and Measurement Peers (MPs), used to actually perform the measurements. In addition some information is required to be configured on the MA prior to any communication with a Controller.

This document defines the information model for both Control and the Report protocols along with pre-configuration information that is required on the MA before communicating with the Controller, broadly named as the LMAP Information Model. The measurement protocols are out of the scope of this document.

As defined in [RFC3444], the LMAP Information Model defines the concepts involved in a large-scale measurement platform at a high level of abstraction, independent of any specific implementation or actual protocol used to exchange the information. It is expected that the proposed information model can be used with different protocols in different measurement platform architectures and across different types of MA devices (e.g., home gateway, smartphone, PC, router). A YANG data model implementing the information model can be found in [I-D.ietf-lmap-yang].

The definition of an Information Model serves a number of purposes:

1. To guide the standardisation of one or more Control and Report protocols and data models

2. To enable high-level inter-operability between different Control and Report protocols by facilitating translation between their respective data models such that a Controller could instruct sub-populations of MAs using different protocols

3. To form agreement of what information needs to be held by an MA and passed over the Control and Report interfaces and support the functionality described in the LMAP framework

4. To enable existing protocols and data models to be assessed for their suitability as part of a large-scale measurement system

2. Notation

This document uses a programming language-like notation to define the properties of the objects of the information model. An optional property is enclosed by square brackets, [ ], and a list property is indicated by two numbers in angle brackets, <m..n>, where m indicates the minimal number of values, and n is the maximum. The symbol * for n means no upper bound.

The object definitions use a couple of base types that are defined as follows:

int	A type representing signed or unsigned integer numbers. This information model does not define a precision nor does it make a distinction between signed and unsigned number ranges. This type is also used to represent enumerations.
boolean	A type representing a boolean value.
string	A type representing a human-readable string consisting of a (possibly restricted) subset of Unicode and ISO/IEC 10646 [ISO.10646] characters.
datetime	A type representing a date and time using the Gregorian calendar. The datetime format MUST conform to RFC 3339 [RFC3339].
uuid	A type representing Universally Unique IDentifier (UUID(as defined in RFC 4122 [RFC4122]. The UUID values are expected to be unique within an installation of a large-scale measurement system.
uri	A type representing a Uniform Resource Identifier as defined in STD 66 [RFC3986].
ip-address	A type representing an IP address. This type supports both IPv4 and IPv6 addresses.
counter	A non-negative integer that monotonically increases. Counters may have discontinuities and they are not expected to persist across restarts.
credentials	An opaque type representing credentials needed by a cryptographic mechanism to secure communication. Data

models must expand this opaque type as needed and required by the security protocols utilized.

data An opaque type representing data obtained from measurements.

Names of objects are generally assumed to be unique within an implementation.

3. LMAP Information Model

The information described herein relates to the information stored, received or transmitted by a Measurement Agent as described within the LMAP framework [RFC7594]. As such, some subsets of this information model are applicable to the measurement Controller, Collector and any device management system that pre-configures the Measurement Agent. The information described in these models will be transmitted by protocols using interfaces between the Measurement Agent and such systems according to a Data Model.

The information model is divided into six aspects. Firstly the grouping of information facilitates reader understanding. Secondly, the particular groupings chosen are expected to map to different protocols or different transmissions within those protocols.

1. Pre-Configuration Information. Information pre-configured on the Measurement Agent prior to any communication with other components of the LMAP architecture (i.e., the Controller, Collector and Measurement Peers), specifically detailing how to communicate with a Controller and whether the device is enabled to participate as an MA.

2. Configuration Information. Update of the pre-configuration information during the registration of the MA or subsequent communication with the Controller, along with the configuration of further parameters about the MA (rather than the Measurement Tasks it should perform) that were not mandatory for the initial communication between the MA and a Controller.

3. Instruction Information. Information that is received by the MA from the Controller pertaining to the Measurement Tasks that should be executed. This includes the task execution Schedules (other than the Controller communication Schedule supplied as (pre)configuration information) and related information such as the Task Configuration, communication Channels to Collectors and schedule Event and Timing information. It also includes Task Suppression information that is used to over-ride normal Task execution.

4. Logging Information. Information transmitted from the MA to the Controller detailing the results of any configuration operations along with error and status information from the operation of the MA.

5. Capability and Status Information. Information on the general status and capabilities of the MA. For example, the set of measurements that are supported on the device.

6. Reporting Information. Information transmitted from the MA to one or more Collectors including measurement results and the context in which they were conducted.

In addition the MA may hold further information not described herein, and which may be optionally transferred to or from other systems including the Controller and Collector. One example of information in this category is subscriber or line information that may be extracted by a task and reported by the MA in the reporting communication to a Collector.

It should also be noted that the MA may be in communication with other management systems which may be responsible for configuring and retrieving information from the MA device. Such systems, where available, can perform an important role in transferring the pre-configuration information to the MA or enabling/disabling the measurement functionality of the MA.

The granularity of data transmitted in each operation of the Control and Report Protocols is not dictated by the Information Model. For example, the Instruction object may be delivered in a single operation. Alternatively, Schedules and Task Configurations may be separated or even each Schedule/Task Configuration may be delivered individually. Similarly the Information Model does not dictate whether data is read, write, or read/write. For example, some Control Protocols may have the ability to read back Configuration and Instruction information which have been previously set on the MA. Lastly, while some protocols may simply overwrite information (for example refreshing the entire Instruction Information), other protocols may have the ability to update or delete selected items of information.

The information modeled by the six aspects of the information model is supported by a number of common information objects. These objects are also described later in this document and comprise of:

a. Schedules. A set of Schedules tells the MA to execute Actions. An Action of a Schedule leads to the execution of a Task. Without a Schedule no Task (including measurements or reporting

or communicating with the Controller) is ever executed. Schedules are used within the Instruction to specify what tasks should be performed, when, and how to direct their results. A Schedule is also used within the pre-Configuration and Configuration information in order to execute the Task or Tasks required to communicate with the Controller. A specific Schedule can only be active once. Attempts to start a Schedule while the same Schedule is still running will fail.

b. Channels. A set of Channel objects are used to communicate with a number of endpoints (i.e., the Controller and Collectors). Each Channel object contains the information required for the communication with a single endpoint such as the target location and security details.

c. Task Configurations. A set of Task Configurations is used to configure the Tasks that are run by the MA. This includes the registry entries for the Task and any configuration parameters, represented as Task Options. Task Configurations are referenced from a Schedule in order to specify what Tasks the MA should execute.

d. Events. A set of Event objects that can be referenced from the Schedules. Each Schedule always references exactly one Event object that determines when the schedule is executed. An Event object specifies either a singleton or series of events that indicate when Tasks should be executed. A commonly used kind of Event objects are Timing objects. For Event objects specifying a series of events, it is generally a good idea to configure an end time and to refresh the end time as needed to ensure that MAs that loose connectivity to their controller do not continue executing Schedules forever.

Figure 1 illustrates the structure in which these common information objects are referenced. The references are achieved by each object (Task Configuration, Event) being given a short textual name that is used by other objects. The objects shown in parenthesis are part of the internal object structure of a Schedule. Channels are not shown in the diagram since they are only used as an option by selected Task Configurations but are similarly referenced using a short text name.

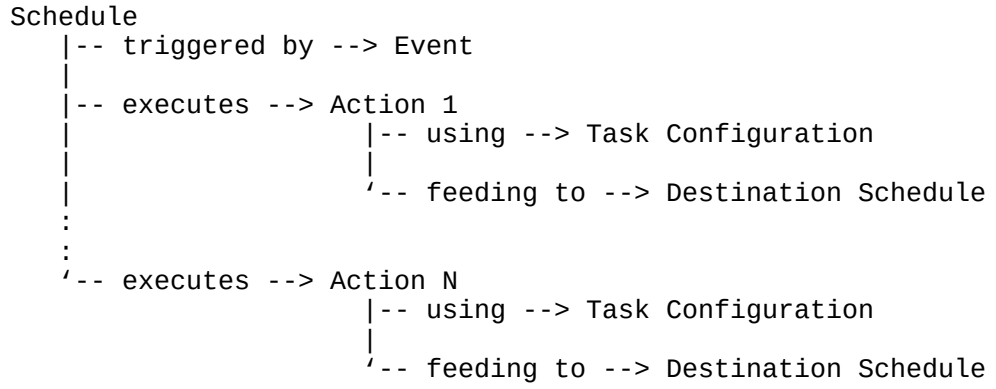

Figure 1: Relationship between Schedules, Events, Actions, Task Configurations, and Destination Schedules

The primary function of an MA is to execute Schedules. A Schedule, which is triggered by an Event, executes a number of Actions. An Action refers to a Configured Task and it may feed results to a Destination Schedule. Both, Actions and Configured Tasks can provide parameters, represented as Action Options and Task Options.

Tasks can implement a variety of different functions. While in terms of the Information Model, all Tasks have the same structure, it can help conceptually to think of different Task categories:

1. Measurement Tasks measure some aspect of network performance or traffic. They may also capture contextual information from the MA device or network interfaces such as the device type or interface speed.

2. Data Transfer Tasks support the communication with a Controller and Collectors:

   A. Reporting Tasks report the results of Measurement Tasks to Collectors

   B. Control Task(s) implement the Control Protocol and communicate with the Controller.

3. Data Analysis Tasks can exist to analyse data from other Measurement Tasks locally on the MA

4. Data Management Tasks may exist to clean-up, filter or compress data on the MA such as Measurement Task results

Figure 1 indicates that Actions can produce data that is fed into Destination Schedules. This can by used by Actions implementing Measurement Tasks to feed measurement results to a Schedule that triggers Actions implementing Reporting Tasks. Data fed to a Destination Schedule is consumed by the first Action of the Destination Schedule if the Destination Schedule is using sequential or pipelined execution mode and it is consumed by all Actions of the Destination Schedule if the Destination Schedule is using parallel execution mode.

3.1. Pre-Configuration Information

This information is the minimal information that needs to be pre-configured to the MA in order for it to successfully communicate with a Controller during the registration process. Some of the Pre-Configuration Information elements are repeated in the Configuration Information in order to allow an LMAP Controller to update these items. The pre-configuration information also contains some elements that are not under the control of the LMAP framework (such as the device identifier and device security credentials).

This Pre-Configuration Information needs to include a URL of the initial Controller from where configuration information can be communicated along with the security information required for the communication including the certificate of the Controller (or the certificate of the Certification Authority which was used to issue the certificate for the Controller). All this is expressed as a Channel. While multiple Channels may be provided in the Pre-Configuration Information they must all be associated with a single Controller (e.g., over different interfaces or network protocols).

Where the MA pulls information from the Controller, the Pre-Configuration Information also needs to contain the timing of the communication with the Controller as well as the nature of the communication itself (such as the protocol and data to be transferred). The timing is represented as an Event that invokes a Schedule that executes the Task(s) responsible for communication with the Controller. It is this Task (or Tasks) that implement the Control protocol between the MA and the Controller and utilises the Channel information. The Task(s) may take additional parameters, as defined by a Task Configuration.

Even where information is pushed to the MA from the Controller (rather than pulled by the MA), a Schedule still needs to be supplied. In this case the Schedule will simply execute a Controller listener Task when the MA is started. A Channel is still required for the MA to establish secure communication with the Controller.

It can be seen that these Channels, Schedules and Task Configurations for the initial MA-Controller communication are no different in terms of the Information Model to any other Channel, Schedule or Task Configuration that might execute a Measurement Task or report the measurement results (as described later).

The MA may be pre-configured with an MA ID, or may use a Device ID in the first Controller contact before it is assigned an MA ID. The Device ID may be a MAC address or some other device identifier expressed as a URI. If the MA ID is not provided at this stage, then it must be provided by the Controller during Configuration.

3.1.1. Definition of ma-preconfig-obj

```
object {
    [uuid                ma-preconfig-agent-id;]
     ma-task-obj         ma-preconfig-control-tasks<1..*>;
     ma-channel-obj      ma-preconfig-control-channels<1..*>;
     ma-schedule-obj     ma-preconfig-control-schedules<1..*>;
    [uri                 ma-preconfig-device-id;]
     credentials         ma-preconfig-credentials;
} ma-preconfig-obj;
```

The ma-preconfig-obj describes information that needs to be available to the MA in order to bootstrap communication with a Controller. The ma-preconfig-obj consists of the following elements:

ma-preconfig-agent-id: An optional uuid uniquely identifying the measurement agent.

ma-preconfig-control-tasks: An unordered set of task objects.

ma-preconfig-control-channels: An unordered set of channel objects.

ma-preconfig-control-schedules: An unordered set of scheduling objects.

ma-preconfig-device-id: An optional identifier for the device.

ma-preconfig-credentials: The security credentials used by the measurement agent.

3.2. Configuration Information

During registration or at any later point at which the MA contacts the Controller (or vice-versa), the choice of Controller, details for the timing of communication with the Controller or parameters for the

communication Task(s) can be changed (as captured by the Channels, Schedules and Task Configurations objects). For example the pre-configured Controller (specified as a Channel or Channels) may be over-ridden with a specific Controller that is more appropriate to the MA device type, location or characteristics of the network (e.g., access technology type or broadband product). The initial communication Schedule may be over-ridden with one more relevant to routine communications between the MA and the Controller.

While some Control protocols may only use a single Schedule, other protocols may use several Schedules (and related data transfer Tasks) to update the Configuration Information, transfer the Instruction Information, transfer Capability and Status Information and send other information to the Controller such as log or error notifications. Multiple Channels may be used to communicate with the same Controller over multiple interfaces (e.g., to send logging information over a different network).

In addition the MA will be given further items of information that relate specifically to the MA rather than the measurements it is to conduct or how to report results. The assignment of an ID to the MA is mandatory. If the MA Agent ID was not optionally provided during the pre-configuration then one must be provided by the Controller during Configuration. Optionally a Group ID may also be given which identifies a group of interest to which that MA belongs. For example the group could represent an ISP, broadband product, technology, market classification, geographic region, or a combination of multiple such characteristics. Additional flags control whether the MA ID or the Group ID are included in Reports. The reporting of a Group ID without the MA ID may allow the MA to remain anonymous, which may be particularly useful to prevent tracking of mobile MA devices.

Optionally an MA can also be configured to stop executing any Instruction Schedule if the Controller is unreachable. This can be used as a fail-safe to stop Measurement and other Tasks being conducted when there is doubt that the Instruction Information is still valid. This is simply represented as a time window in seconds since the last communication with the Controller after which an Event is generated that can trigger the suspension of Instruction Schedules. The appropriate value of the time window will depend on the specified communication Schedule with the Controller and the duration for which the system is willing to tolerate continued operation with potentially stale Instruction Information.

While Pre-Configuration Information is persistent upon device reset or power cycle, the persistency of the Configuration Information may be device dependent. Some devices may revert back to their pre-

configuration state upon reboot or factory reset, while other devices may store all Configuration and Instruction information in persistent storage. A Controller can check whether an MA has the latest Configuration and Instruction information by examining the Capability and Status information for the MA.

3.2.1. Definition of ma-config-obj

```
object {
    uuid                ma-config-agent-id;
    ma-task-obj         ma-config-control-tasks<1..*>;
    ma-channel-obj      ma-config-control-channels<1..*>;
    ma-schedule-obj     ma-config-control-schedules<1..*>;
    credentials         ma-config-credentials;
   [string              ma-config-group-id;]
   [string              ma-config-measurement-point;]
   [boolean             ma-config-report-agent-id;]
   [boolean             ma-config-report-group-id;]
   [boolean             ma-config-report-measurement-point;]
   [int                 ma-config-controller-timeout;]
} ma-config-obj;
```

The ma-config-obj consists of the following elements:

ma-config-agent-id: A uuid uniquely identifying the measurement agent.

ma-config-control-tasks: An unordered set of task objects.

ma-config-control-channels: An unordered set of channel objects.

ma-config-control-schedules: An unordered set of scheduling objects.

ma-config-credentials: The security credentials used by the measurement agent.

ma-config-group-id: An optional identifier of the group of measurement agents this measurement agent belongs to.

ma-config-measurement-point: An optional identifier for the measurement point indicating where the measurement agent is located on a path (see [RFC7398] for further details).

ma-config-report-agent-id:	An optional flag indicating whether the agent identifier (ma-config-agent-id) is included in reports. The default value is true.
ma-config-report-group-id:	An optional flag indicating whether the group identifier (ma-config-group-id) is included in reports. The default value is false.
ma-config-report-measurement-point:	An optional flag indicating whether the measurement point (ma-config-measurement-point) should be included in reports. The default value is false.
ma-config-controller-timeout:	A timer is started after each successful contact with a controller. When the timer reaches the controller-timeout (measured in seconds), an event is raised indicating that connectivity to the controller has been lost (see ma-controller-lost-obj).

3.3. Instruction Information

The Instruction information model has four sub-elements:

1. Instruction Task Configurations

2. Report Channels

3. Instruction Schedules

4. Suppression

The Instruction supports the execution of all Tasks on the MA except those that deal with communication with the Controller (specified in (pre-)configuration information). The Tasks are configured in Instruction Task Configurations and included by reference in the Actions of Instruction Schedules that specify when to execute them. The results can be communicated to other Schedules or a Task may implement a Reporting Protocol and communicate results over Report Channels. Suppression is used to temporarily stop the execution of

new Tasks as specified by the Instruction Schedules (and optionally to stop ongoing Tasks).

A Task Configuration is used to configure the mandatory and optional parameters of a Task. It also serves to instruct the MA about the Task including the ability to resolve the Task to an executable and specifying the schema for the Task parameters.

A Report Channel defines how to communicate with a single remote system specified by a URL. A Report Channel is used to send results to a single Collector but is no different in terms of the Information Model to the Control Channel used to transfer information between the MA and the Controller. Several Report Channels can be defined to enable results to be split or duplicated across different destinations. A single Channel can be used by multiple (reporting) Task Configurations to transfer data to the same Collector. A single Reporting Task Configuration can also be included in multiple Schedules. E.g., a single Collector may receive data at three different cycle rates, one Schedule reporting hourly, another reporting daily and a third specifying that results should be sent immediately for on-demand measurement tasks. Alternatively multiple Report Channels can be used to send Measurement Task results to different Collectors. The details of the Channel element is described later as it is common to several objects.

Instruction Schedules specify which Actions to execute according to a given triggering Event. An Action extends a Configured Task with additional specific parameters. An Event can trigger the execution of a single Action or it can trigger a repeated series of Actions. The Schedule also specifies how to link Tasks output data to other Schedules.

Measurement Suppression information is used to over-ride the Instruction Schedule and temporarily stop measurements or other Tasks from running on the MA for a defined or indefinite period. While conceptually measurements can be stopped by simply removing them from the Measurement Schedule, splitting out separate information on Measurement Suppression allows this information to be updated on the MA on a different timing cycle or protocol implementation to the Measurement Schedule. It is also considered that it will be easier for a human operator to implement a temporary explicit suppression rather than having to move to a reduced Schedule and then roll-back at a later time.

It should be noted that control schedules and tasks cannot be suppressed as evidenced by the lack of suppression information in the Configuration. The control schedule must only reference tasks listed as control tasks (i.e., within the Configuration information).

A single Suppression object is able to enable/disable a set of Instruction Tasks that are tagged for suppression. This enables fine grained control on which Tasks are suppressed. Suppression of both matching Actions and Measurement Schedules is supported. Support for disabling specific Actions allows malfunctioning or mis-configured Tasks or Actions that have an impact on a particular part of the network infrastructure (e.g., a particular Measurement Peer) to be targeted. Support for disabling specific Schedules allows for particularly heavy cycles or sets of less essential Measurement Tasks to be suppressed quickly and effectively. Note that Suppression has no effect on either Controller Tasks or Controller Schedules.

Suppression stops new Tasks from executing. In addition, the Suppression information also supports an additional Boolean that is used to select whether on-going tasks are also to be terminated.

Unsuppression is achieved through either overwriting the Measurement Suppression information (e.g., changing 'enabled' to False) or through the use of an End time such that the Measurement Suppression will no longer be in effect beyond this time.

The goal when defining these four different elements is to allow each part of the information model to change without affecting the other three elements. For example it is envisaged that the Report Channels and the set of Task Configurations will be relatively static. The Instruction Schedule, on the other hand, is likely to be more dynamic, as the measurement panel and test frequency are changed for various business goals. Another example is that measurements can be suppressed with a Suppression command without removing the existing Instruction Schedules that would continue to apply after the Suppression expires or is removed. In terms of the Controller-MA communication this can reduce the data overhead. It also encourages the re-use of the same standard Task Configurations and Reporting Channels to help ensure consistency and reduce errors.

3.3.1. Definition of ma-instruction-obj

```
object {
    ma-task-obj         ma-instruction-tasks<0..*>;
    ma-channel-obj      ma-instruction-channels<0..*>;
    ma-schedule-obj     ma-instruction-schedules<0..*>;
   [ma-suppression-obj  ma-instruction-suppressions<0..*>;]
} ma-instruction-obj;
```

An ma-instruction-obj consists of the following elements:

ma-instruction-tasks: A possibly empty unordered set of task objects.

ma-instruction-channels: A possibly empty unordered set of channel objects.

ma-instruction-schedules: A possibly empty unordered set of schedule objects.

ma-instruction-suppressions: An optional possibly empty unordered set of suppression objects.

3.3.2. Definition of ma-suppression-obj

```
object {
    string              ma-suppression-name;
   [ma-event-obj        ma-suppression-start;]
   [ma-event-obj        ma-suppression-end;]
   [string              ma-suppression-match<0..*>;]
   [boolean             ma-suppression-stop-running;]
} ma-suppression-obj;
```

The ma-suppression-obj controls the suppression of schedules or actions and consists of the following elements:

ma-suppression-name: A name uniquely identifying a suppression.

ma-suppression-start: The optional event indicating when suppression starts. If not present, the suppression starts immediately, i.e., as if the value would be 'immediate'.

ma-suppression-end: The optional event indicating when suppression ends. If not present, the suppression does not have a defined end, i.e., the suppression remains for an indefinite period of time.

ma-suppression-match: An optional and possibly empty unordered set of match patterns. The suppression will apply to all schedules (and their actions) that have a matching value in their ma-schedule-suppression-tags and all actions that have a matching value in their ma-action-suppression-tags. Pattern matching is done using glob style pattern (see below).

ma-suppression-stop-running: An optional boolean indicating whether suppression will stop any running matching schedules or actions. The default value for this boolean is false.

Glob style pattern matching is following POSIX.2 fnmatch() [POSIX.2] without special treatment of file paths:

*	matches a sequence of characters
?	matches a single character
[seq]	matches any character in seq
[!seq]	matches any character not in seq

A backslash followed by a character matches the following character. In particular:

*	matches *
\?	matches ?
\\	matches \

A sequence seq may be a sequence of characters (e.g., [abc] or a range of characters (e.g., [a-c]).

3.4. Logging Information

The MA may report on the success or failure of Configuration or Instruction communications from the Controller. In addition further operational logs may be produced during the operation of the MA and updates to capabilities may also be reported. Reporting this information is achieved in exactly the same manner as scheduling any other Task. We make no distinction between a Measurement Task conducting an active or passive network measurement and one which solely retrieves static or dynamic information from the MA such as capabilities or logging information. One or more logging tasks can be programmed or configured to capture subsets of the Logging Information. These logging tasks are then executed by Schedules which also specify that the resultant data is to be transferred over the Controller Channels.

The type of Logging Information will fall into three different categories:

1. Success/failure/warning messages in response to information updates from the Controller. Failure messages could be produced due to some inability to receive or parse the Controller communication, or if the MA is not able to act as instructed. For example:

* "Measurement Schedules updated OK"

* "Unable to parse JSON"

* "Missing mandatory element: Measurement Timing"

* "'Start' does not conform to schema - expected datetime"

* "Date specified is in the past"

* "'Hour' must be in the range 1..24"

* "Schedule A refers to non-existent Measurement Task Configuration"

* "Measurement Task Configuration X registry entry Y not found"

* "Updated Measurement Task Configurations do not include M used by Measurement Schedule N"

2. Operational updates from the MA. For example:

* "Out of memory: cannot record result"

* "Collector 'collector.example.com' not responding"

* "Unexpected restart"

* "Suppression timeout"

* "Failed to execute Measurement Task Configuration H"

3. Status updates from the MA. For example:

* "Device interface added: eth3"

* "Supported measurements updated"

* "New IP address on eth0: xxx.xxx.xxx.xxx"

This Information Model document does not detail the precise format of logging information since it is to a large extent protocol and MA specific. However, some common information can be identified.

3.4.1. Definition of ma-log-obj

```
object {
    uuid                ma-log-agent-id;
    datetime            ma-log-event-time;
    int                 ma-log-code;
    string              ma-log-description;
} ma-log-obj;
```

The ma-log-obj models the generic aspects of a logging object and consists of the following elements:

ma-log-agent-id: A uuid uniquely identifying the measurement agent.

ma-log-event-time: The date and time of the event reported in the logging object.

ma-log-code: A machine readable code describing the event.

ma-log-description: A human readable description of the event.

3.5. Capability and Status Information

The MA will hold Capability Information that can be retrieved by a Controller. Capabilities include the device interface details available to Measurement Tasks as well as the set of Measurement Tasks/Roles (specified by registry entries) that are actually installed or available on the MA. Status information includes the times that operations were last performed such as contacting the Controller or producing Reports.

3.5.1. Definition of ma-capability-obj

```
object {
    string                  ma-capability-hardware;
    string                  ma-capability-firmware;
    string                  ma-capability-version;
   [string                  ma-capability-tags<0..*>;]
   [ma-capability-task-obj  ma-capability-tasks<0..*>;]
} ma-capability-obj;
```

The ma-capability-obj provides information about the capabilities of the measurement agent and consists of the following elements:

ma-capability-hardware: A description of the hardware of the device the measurement agent is running on.

ma-capability-firmware: A description of the firmware of the device the measurement agent is running on.

ma-capability-version: The version of the measurement agent.

ma-capability-tags: An optional unordered set of tags that provide additional information about the capabilities of the measurement agent.

ma-capability-tasks: An optional unordered set of capability objects for each supported task.

3.5.2. Definition of ma-capability-task-obj

```
object {
    string              ma-capability-task-name;
    ma-registry-obj     ma-capability-task-functions<0..*>;
    string              ma-capability-task-version;
} ma-capability-task-obj;
```

The ma-capability-task-obj provides information about the capability of a task and consists of the following elements:

ma-capability-task-name: A name uniquely identifying a task.

ma-capability-task-functions: A possibly empty unordered set of registry entries identifying functions this task implements.

ma-capability-task-version: The version of the measurement task.

3.5.3. Definition of ma-status-obj

```
object {
    uuid                      ma-status-agent-id;
   [uri                       ma-status-device-id;]
    datetime                  ma-status-last-started;
    ma-status-interface-obj   ma-status-interfaces<0..*>;
   [ma-status-schedule-obj    ma-status-schedules<0..*>;]
   [ma-status-suppression-obj ma-status-suppressions<0..*>;]
} ma-status-obj;
```

The ma-status-obj provides status information about the measurement agent and consists of the following elements:

ma-status-agent-id: A uuid uniquely identifying the measurement agent.

ma-status-device-id: A URI identifying the device.

ma-status-last-started: The date and time the measurement agent last started.

ma-status-interfaces: An unordered set of network interfaces available on the device.

ma-status-schedules: An optional unordered set of status objects for each schedule.

ma-status-suppressions: An optional unordered set of status objects for each suppression.

3.5.4. Definition of ma-status-schedule-obj

```
object {
    string                  ma-status-schedule-name;
    string                  ma-status-schedule-state;
    int                     ma-status-schedule-storage;
    counter                 ma-status-schedule-invocations;
    counter                 ma-status-schedule-suppressions;
    counter                 ma-status-schedule-overlaps;
    counter                 ma-status-schedule-failures;
    datetime                ma-status-schedule-last-invocation;
   [ma-status-action-obj    ma-status-schedule-actions<0..*>;]
} ma-status-schedule-obj;
```

The ma-status-schedule-obj provides status information about the status of a schedule and consists of the following elements:

ma-status-schedule-name: The name of the schedule this status object refers to.

ma-status-schedule-state: The state of the schedule. The value 'enabled' indicates that the schedule is currently enabled. The value 'suppressed' indicates that the schedule is currently suppressed. The value 'disabled' indicates that the schedule is currently disabled. The value 'running' indicates that the schedule is currently running.

ma-status-schedule-storage: The amount of secondary storage (e.g., allocated in a file

	system) holding temporary data allocated to the schedule in bytes. This object reports the amount of allocated physical storage and not the storage used by logical data records. Data models should use a 64-bit integer type.
ma-status-schedule-invocations	Number of invocations of this schedule. This counter does not include suppressed invocations or invocations that were prevented due to an overlap with a previous invocation of this schedule.
ma-status-schedule-suppressions	Number of suppressed executions of this schedule.
ma-status-schedule-overlaps	Number of executions prevented due to overlaps with a previous invocation of this schedule.
ma-status-schedule-failures	Number of failed executions of this schedule. A failed execution is an execution where at least one action failed.
ma-status-schedule-last-invocation:	The date and time of the last invocation of this schedule.
ma-status-schedule-actions:	An optional ordered list of status objects for each action of the schedule.

3.5.5. Definition of ma-status-action-obj

```
object {
    string              ma-status-action-name;
    string              ma-status-action-state;
    int                 ma-status-action-storage;
    counter             ma-status-action-invocations;
    counter             ma-status-action-suppressions;
    counter             ma-status-action-overlaps;
    counter             ma-status-action-failures;
    datetime            ma-status-action-last-invocation;
    datetime            ma-status-action-last-completion;
    int                 ma-status-action-last-status;
    string              ma-status-action-last-message;
    datetime            ma-status-action-last-failed-completion;
    int                 ma-status-action-last-failed-status;
    string              ma-status-action-last-failed-message;
} ma-status-action-obj;
```

The ma-status-action-obj provides status information about an action of a schedule and consists of the following elements:

ma-status-action-name: The name of the action of a schedule this status object refers to.

ma-status-action-state: The state of the action. The value 'enabled' indicates that the action is currently enabled. The value 'suppressed' indicates that the action is currently suppressed. The value 'disabled' indicates that the action is currently disabled. The value 'running' indicates that the action is currently running.

ma-status-action-storage: The amount of secondary storage (e.g., allocated in a file system) holding temporary data allocated to the action in bytes. This object reports the amount of allocated physical storage and not the storage used by logical data records. Data models should use a 64-bit integer type.

ma-status-action-invocations	Number of invocations of this action. This counter does not include suppressed invocations or invocations that were prevented due to an overlap with a previous invocation of this action.
ma-status-action-suppressions	Number of suppressed executions of this action.
ma-status-action-overlaps	Number of executions prevented due to overlaps with a previous invocation of this action.
ma-status-action-failures	Number of failed executions of this action.
ma-status-action-last-invocation:	The date and time of the last invocation of this action.
ma-status-action-last-completion:	The date and time of the last completion of this action.
ma-status-action-last-status:	The status code returned by the last execution of this action.
ma-status-action-last-message:	The status message produced by the last execution of this action.
ma-status-action-last-failed-completion:	The date and time of the last failed completion of this action.
ma-status-action-last-failed-status:	The status code returned by the last failed execution of this action.
ma-status-action-last-failed-message:	The status message produced by the last failed execution of this action.

3.5.6. Definition of ma-status-suppression-obj

```
object {
    string                ma-status-suppression-name;
    string                ma-status-suppression-state;
} ma-status-suppression-obj;
```

The ma-status-suppression-obj provides status information about that status of a suppression and consists of the following elements:

ma-status-suppression-name: The name of the suppression this status object refers to.

ma-status-suppression-state: The state of the suppression. The value 'enabled' indicates that the suppression is currently enabled. The value 'active indicates that the suppression is currently active. The value 'disabled' indicates that the suppression is currently disabled.

3.5.7. Definition of ma-status-interface-obj

```
object {
    string              ma-status-interface-name;
    string              ma-status-interface-type;
   [int                 ma-status-interface-speed;]
   [string              ma-status-interface-link-layer-address;]
   [ip-address          ma-status-interface-ip-addresses<0..*>;]
   [ip-address          ma-status-interface-gateways<0..*>;]
   [ip-address          ma-status-interface-dns-servers<0..*>;]
} ma-status-interface-obj;
```

The ma-status-interface-obj provides status information about network interfaces and consists of the following elements:

ma-status-interface-name: A name uniquely identifying a network interface.

ma-status-interface-type: The type of the network interface.

ma-status-interface-speed: An optional indication of the speed of the interface (measured in bits-per-second).

ma-status-interface-link-layer-address:	An optional link-layer address of the interface.
ma-status-interface-ip-addresses:	An optional ordered list of IP addresses assigned to the interface.
ma-status-interface-gateways:	An optional ordered list of gateways assigned to the interface.
ma-status-interface-dns-servers:	An optional ordered list of DNS servers assigned to the interface.

3.6. Reporting Information

At a point in time specified by a Schedule, the MA will execute tasks that communicate a set of measurement results to the Collector. These Reporting Tasks will be configured to transmit task results over a specified Report Channel to a Collector.

It should be noted that the output from Tasks does not need to be sent to communication Channels. It can alternatively, or additionally, be sent to other Tasks on the MA. This facilitates using a first Measurement Task to control the operation of a later Measurement Task (such as first probing available line speed and then adjusting the operation of a video testing measurement) and also to allow local processing of data to output alarms (e.g., when performance drops from earlier levels). Of course, subsequent Tasks also include Tasks that implement the reporting protocol(s) and transfer data to one or more Collector(s).

The Report generated by a Reporting Task is structured hierarchically to avoid repetition of report header and Measurement Task Configuration information. The report starts with the timestamp of the report generation on the MA and details about the MA including the optional Measurement Agent ID and Group ID (controlled by the Configuration Information).

Much of the report Information is optional and will depend on the implementation of the Reporting Task and any parameters defined in the Task Configuration for the Reporting Task. For example some Reporting Tasks may choose not to include the Measurement Task Configuration or Action parameters, while others may do so dependent on the Controller setting a configurable parameter in the Task Configuration.

It is possible for a Reporting Task to send just the Report header (datetime and optional agent ID and/or Group ID) if no measurement data is available. Whether to send such empty reports again is dependent on the implementation of the Reporting Task and potential Task Configuration parameter.

The handling of measurement data on the MA before generating a Report and transfer from the MA to the Collector is dependent on the implementation of the device, MA and/or scheduled Tasks and not defined by the LMAP standards. Such decisions may include limits to the measurement data storage and what to do when such available storage becomes depleted. It is generally suggested that implementations running out of storage stop executing new measurement tasks and retain old measurement data.

No context information, such as line speed or broadband product are included within the report header information as this data is reported by individual tasks at the time they execute. Either a Measurement Task can report contextual parameters that are relevant to that particular measurement, or specific tasks can be used to gather a set of contextual and environmental data at certain times independent of the reporting schedule.

After the report header information the results are reported grouped according to different Measurement Task Configurations. Each Task section optionally starts with replicating the Measurement Task Configuration information before the result headers (titles for data columns) and the result data rows. The Options reported are those used for the scheduled execution of the Measurement Task and therefore include the Options specified in the Task Configuration as well as additional Options specified in the Action. The Action Options are appended to the Task Configuration Options in exactly the same order as they were provided to the Task during execution.

The result row data includes a time for the start of the measurement and optionally an end time where the duration also needs to be considered in the data analysis.

Some Measurement Tasks may optionally include an indication of the cross-traffic although the definition of cross-traffic is left up to each individual Measurement Task. Some Measurement Tasks may also output other environmental measures in addition to cross-traffic such as CPU utlilisation or interface speed.

Whereas the Configuration and Instruction information represent information transmitted via the Control Protocol, the Report represents the information that is transmitted via the Report Protocol. It is constructed at the time of sending a report and

represents the inherent structure of the information that is sent to the Collector.

3.6.1. Definition of ma-report-obj

```
object {
    datetime              ma-report-date;
   [uuid                  ma-report-agent-id;]
   [string                ma-report-group-id;]
   [string                ma-report-measurement-point;]
   [ma-report-result-obj  ma-report-results<0..*>;]
} ma-report-obj;
```

The ma-report-obj provides the meta-data of a single report and consists of the following elements:

ma-report-date: The date and time when the report was sent to a collector.

ma-report-agent-id: An optional uuid uniquely identifying the measurement agent.

ma-report-group-id: An optional identifier of the group of measurement agents this measurement agent belongs to.

ma-report-measurement-point: An optional identifier for the measurement point indicating where the measurement agent is located on a path (see [RFC7398] for further details).

ma-report-results: An optional and possibly empty unordered set of result objects.

3.6.2. Definition of ma-report-result-obj

```
object {
    string                  ma-report-result-schedule-name;
    string                  ma-report-result-action-name;
    string                  ma-report-result-task-name;
   [ma-option-obj           ma-report-result-options<0..*>;]
   [string                  ma-report-result-tags<0..*>;]
    datetime                ma-report-result-event-time;
    datetime                ma-report-result-start-time;
   [datetime                ma-report-result-end-time;]
   [string                  ma-report-result-cycle-number;]
    int                     ma-report-result-status;
   [ma-report-conflict-obj  ma-report-result-conflicts<0..*>;]
   [ma-report-table-obj     ma-report-result-tables<0..*>;]
} ma-report-result-obj;
```

The ma-report-result-obj provides the meta-data of a result report of a single executed action. It consists of the following elements:

ma-report-result-schedule-name: The name of the schedule that produced the result.

ma-report-result-action-name: The name of the action in the schedule that produced the result.

ma-report-result-task-name: The name of the task that produced the result.

ma-report-result-options: An optional ordered joined list of options provided by the task object and the action object when the action was started.

ma-report-result-tags: An optional unordered set of tags. This is the joined set of tags provided by the task object and the action object and schedule object when the action was started.

ma-report-result-event-time: The date and time of the event that triggered the schedule of the action that produced the reported result values. The date and time does not include any added randomization.

ma-report-result-start-time: The date and time of the start of the action that produced the reported result values.

ma-report-result-end-time: An optional date and time indicating when the action finished.

ma-report-result-cycle-number: An optional cycle number derived from ma-report-result-event-time. It is the time closest to ma-report-result-event-time that is a multiple of the ma-event-cycle-interval of the event that triggered the execution of the schedule. The value is only present in an ma-report-result-obj if the event that triggered the execution of the schedule has a defined ma-event-cycle-interval. The cycle number is represented in the format YYYYMMDD.HHMMSS where YYYY represents the year, MM the month (1..12), DD the day of the months (01..31), HH the hour (00..23), MM the minute (00..59), and SS the second (00..59). The cycle number is using Coordinated Universal Time (UTC).

ma-report-result-status: The status code returned by the execution of the action.

ma-report-result-conflicts: A possibly empty set of conflict actions that might have impacted the measurement results being reported.

ma-report-result-tables: An optional and possibly empty unordered set of result tables.

3.6.3. Definition of ma-report-conflict-obj

```
object {
    string  ma-report-conflict-schedule-name;
    string  ma-report-conflict-action-name;
    string  ma-report-conflict-task-name;
} ma-report-conflict-obj;
```

The ma-report-conflict-obj provides the information about conflicting action that might have impacted the measurement results. It consists of the following elements:

ma-report-result-schedule-name: The name of the schedule that may have impacted the result.

ma-report-result-action-name: The name of the action in the schedule that may have impacted the result.

ma-report-result-task-name: The name of the task that may have impacted the result.

3.6.4. Definition of ma-report-table-obj

```
object {
   [ma-registry-obj     ma-report-table-functions<0..*>;]
   [string]             ma-report-table-column-labels<0..*>;]
   [ma-report-row-obj   ma-report-table-rows<0..*>;]
} ma-report-table-obj;
```

The ma-report-table-obj represents a result table and consists of the following elements:

ma-report-table-functions: An optional and possibly empty unordered set of registry entries identifying the functions for which results that are reported.

ma-report-table-column-labels: An optional and possibly empty ordered list of column labels.

ma-report-table-rows: A possibly empty ordered list of result rows.

3.6.5. Definition of ma-report-row-obj

```
object {
    data               ma-report-row-values<0..*>;
} ma-report-row-obj;
```

The ma-report-row-obj represents a result row and consists of the following elements:

ma-report-row-values: A possibly empty ordered list of result values. When present, it contains an ordered list of values that align to the set of column labels for the report.

3.7. Common Objects: Schedules

A Schedule specifies the execution of a single or repeated series of Actions. An Action extends a Configured Task with additional specific parameters. Each Schedule contains basically two elements:

an ordered list of Actions to be executed and an Event object triggering the execution of the Schedule. The Schedule states what Actions to run (with what configuration) and when to run the Actions. A Schedule may optionally have an Event that stops the execution of the Schedule or a maximum duration after which a schedule is stopped.

Multiple Actions contained as an ordered list of a single Measurement Schedule will be executed according to the execution mode of the Schedule. In sequential mode, Actions will be executed sequentially and in parallel mode, all Actions will be executed concurrently. In pipelined mode, data produced by one Action is passed to the subsequent Action. Actions contained in different Schedules execute in parallel with such conflicts being reported in the Reporting Information where necessary. If two or more Schedules have the same start time, then the two will execute in parallel. There is no mechanism to prioritise one schedule over another or to mutex scheduled tasks.

As well as specifying which Actions to execute, the Schedule also specifies how to link the data outputs from each Action to other Schedules. Specifying this within the Schedule allows the highest level of flexibility since it is even possible to send the output from different executions of the same Task Configuration to different destinations. A single Task producing multiple different outputs is expected to properly tag the different result. An Action receiving the output can then filter the results based on the tag if necessary. For example, a Measurement Task might report routine results to a data Reporting Task in a Schedule that communicates hourly via the Broadband PPP interface, but also outputs emergency conditions via an alarm Reporting Task in a different Schedule communicating immediately over a GPRS channel. Note that task-to-task data transfer is always specified in association with the scheduled execution of the sending task - there is no need for a corresponding input specification for the receiving task. While it is likely that an MA implementation will use a queue mechanism between the Schedules or Actions, this Information Model does not mandate or define a queue. The Information Model, however, reports the storage allocated to Schedules and Actions so that storage usage can be monitored. Furthermore, it is recommended that MA implementations by default retain old data and stop the execution of new measurement tasks if the MA runs out of storage capacity.

When specifying the task to execute within the Schedule, i.e., creating an Action, it is possible to add to the Action option parameters. This allows the Task Configuration to determine the common characteristics of a Task, while selected parameters (e.g., the test target URL) are defined within as option parameters of the Action in the schedule. A single Tasks Configuration can even be

used multiple times in the same schedule with different additional parameters. This allows for efficiency in creating and transferring the Instruction. Note that the semantics of what happens if an option is defined multiple times (either in the Task Configuration, Action or in both) is not standardised and will depend upon the Task. For example, some tasks may legitimately take multiple values for a single parameter.

Where Options are specified in both the Action and the Task Configuration, the Action Options are appended to those specified in the Task Configuration.

Example: An Action of a Schedule references a single Measurement Task Configuration for measuring UDP latency. It specifies that results are to be sent to a Schedule with a Reporting Action. This Reporting Task of the Reporting Action is executed by a separate Schedule that specifies that it should run hourly at 5 minutes past the hour. When run this Reporting Action takes the data generated by the UDP latency Measurement Task as well as any other data to be included in the hourly report and transfers it to the Collector over the Report Channel specified within its own Schedule.

Schedules and Actions may optionally also be given tags that are included in result reports sent to a Collector. In addition, schedules can be given suppression tags that may be used to select Schedules and Actions for suppression.

3.7.1. Definition of ma-schedule-obj

```
  object {
      string              ma-schedule-name;
      ma-event-obj        ma-schedule-start;
     [ma-event-obj        ma-schedule-end;]
     [int                 ma-schedule-duration;]
      ma-action-obj       ma-schedule-actions<0..*>;
      string              ma-schedule-execution-mode;
     [string              ma-schedule-tags<0..*>;]
     [string              ma-schedule-suppression-tags<0..*>;]
  } ma-schedule-obj;
```

The ma-schedule-obj is the main scheduling object. It consists of the following elements:

ma-schedule-name: A name uniquely identifying a scheduling object.

ma-schedule-start:	An event object indicating when the schedule starts.
ma-schedule-end:	An optional event object controlling the forceful termination of scheduled actions. When the event occurs, all actions of the schedule will be forced to terminate gracefully.
ma-schedule-duration:	An optional duration in seconds for the schedule. All actions of the schedule will be forced to terminate gracefully after the duration number of seconds past the start of the schedule.
ma-schedule-actions:	A possibly empty ordered list of actions to invoke when the schedule starts.
ma-schedule-execution-mode:	Indicates whether the actions should be executed sequentially, in parallel, or in a pipelined mode (where data produced by one action is passed to the subsequent action). The default execution mode is pipelined.
ma-schedule-tags:	An optional unordered set of tags that are reported together with the measurement results to a collector.
ma-schedule-suppression-tags:	An optional unordered set of suppression tags that are used to select schedules to be suppressed.

3.7.2. Definition of ma-action-obj

```
  object {
      string              ma-action-name;
      string              ma-action-config-task-name;
     [ma-option-obj       ma-action-task-options<0..*>;]
     [string              ma-action-destinations<0..*>;]
     [string              ma-action-tags<0..*>;]
     [string              ma-action-suppression-tags<0..*>;]
   } ma-action-obj;
```

The ma-action-obj models a task together with its schedule specific task options and destination schedules. It consists of the following elements:

ma-action-name:	A name uniquely identifying an action of a scheduling object.
ma-action-config-task-name:	A name identifying the configured task to be invoked by the action.
ma-action-task-options:	An optional and possibly empty ordered list of options (name-value pairs) that are passed to the task by appending them to the options configured for the task object.
ma-action-destinations:	An optional and possibly empty unordered set of names of destination schedules that consume output produced by this action.
ma-action-tags:	An optional unordered set of tags that are reported together with the measurement results to a collector.
ma-action-suppression-tags:	An optional unordered set of suppression tags that are used to select actions to be suppressed.

3.8. Common Objects: Channels

A Channel defines a bi-directional communication mechanism between the MA and a Controller or Collector. Multiple Channels can be defined to enable results to be split or duplicated across different Collectors.

Each Channel contains the details of the remote endpoint (including location and security credential information such as a certificate). The timing of when to communicate over a Channel is specified by the Schedule which executes the corresponding Control or Reporting Task. The certificate can be the digital certificate associated to the FQDN in the URL or it can be the certificate of the Certification Authority that was used to issue the certificate for the FQDN (Fully Qualified Domain Name) of the target URL (which will be retrieved later on using a communication protocol such as TLS). In order to establish a secure channel, the MA will use its own security credentials (in the Configuration Information) and the given credentials for the individual Channel end-point.

As with the Task Configurations, each Channel is also given a text name by which it can be referenced as a Task Option.

Although the same in terms of information, Channels used for communication with the Controller are referred to as Control Channels whereas Channels to Collectors are referred to as Report Channels. Hence Control Channels will be referenced from Control Tasks executed by a Control Schedule, whereas Report Channels will be referenced from within Reporting Tasks executed by an Instruction Schedule.

Multiple interfaces are also supported. For example the Reporting Task could be configured to send some results over GPRS. This is especially useful when such results indicate the loss of connectivity on a different network interface.

Example: A Channel used for reporting results may specify that results are to be sent to the URL (https://collector.example.org/report/), using the appropriate digital certificate to establish a secure channel.

3.8.1. Definition of ma-channel-obj

```
object {
    string              ma-channel-name;
    url                 ma-channel-target;
    credentials         ma-channel-credentials;
    [string             ma-channel-interface-name;]
} ma-channel-obj;
```

The ma-channel-obj consists of the following elements:

ma-channel-name: A unique name identifying the channel object.

ma-channel-target: A URL identifying the target channel endpoint.

ma-channel-credentials: The security credentials needed to establish a secure channel.

ma-channel-interface-name: An optional name of the network interface to be used. If not present, the IP protocol stack will select a suitable interface.

3.9. Common Objects: Task Configurations

Conceptually each Task Configuration defines the parameters of a Task that the Measurement Agent (MA) may perform at some point in time. It does not by itself actually instruct the MA to perform them at any particular time (this is done by a Schedule). Tasks can be

Measurement Tasks (i.e., those Tasks actually performing some type of passive or active measurement) or any other scheduled activity performed by the MA such as transferring information to or from the Controller and Collectors. Other examples of Tasks may include data manipulation or processing Tasks conducted on the MA.

A Measurement Task Configuration is the same in information terms to any other Task Configuration. Both measurement and non-measurement Tasks may have registry entries to enable the MA to uniquely identify the Task it should execute and retrieve the schema for any parameters that may be passed to the Task. Registry entries are specified as a URI and can therefore be used to identify the Task within a namespace or point to a web or local file location for the Task information. As mentioned previously, these URIs may be used to identify the Measurement Task in a public namespace [I-D.ietf-ippm-metric-registry].

Example: A Measurement Task Configuration may configure a single Measurement Task for measuring UDP latency. The Measurement Task Configuration could define the destination port and address for the measurement as well as the duration, internal packet timing strategy and other parameters (for example a stream for one hour and sending one packet every 500 ms). It may also define the output type and possible parameters (for example the output type can be the 95th percentile mean) where the measurement task accepts such parameters. It does not define when the task starts (this is defined by the Schedule element), so it does not by itself instruct the MA to actually perform this Measurement Task.

The Task Configuration will include a local short name for reference by a Schedule. Task Configurations may also refer to registry entries as described above. In addition the Task can be configured through a set of configuration Options. The nature and number of these Options will depend upon the Task. These options are expressed as name-value pairs although the ’value’ may be a structured object instead of a simple string or numeric value. The implementation of these name-value pairs will vary between data models.

An Option that must be present for Reporting Tasks is the Channel reference specifying how to communicate with a Collector. This is included in the task options and will have a value that matches a channel name that has been defined in the Instruction. Similarly Control Tasks will have a similar option with the value set to a specified Control Channel.

A Reporting Task might also have a flag parameter, defined as an Option, to indicate whether to send a report without measurement results if there is no measurement result data pending to be

transferred to the Collector. In addition many tasks will also take as a parameter which interface to operate over.

In addition the Task Configuration may optionally also be given tags that can carry a Measurement Cycle ID. The purpose of this ID is to easily identify a set of measurement results that have been produced by Measurement Tasks with comparable Options. This ID could be manually incremented or otherwise changed when an Option change is implemented which could mean that two sets of results should not be directly compared.

3.9.1. Definition of ma-task-obj

```
object {
    string              ma-task-name;
    ma-registry-obj     ma-task-functions<0..*>;
   [ma-option-obj       ma-task-options<0..*>;]
   [string              ma-task-tags<0..*>;]
} ma-task-obj;
```

The ma-task-obj defines a configured task that can be invoked as part of an action. A configured task can be referenced by its name and it contains a possibly empty set of URIs to link to registry entries. Options allow the configuration of task parameters (in the form of name-value pairs). The ma-task-obj consists of the following elements:

ma-task-name: A name uniquely identifying a configured task object.

ma-task-functions: A possibly empty unordered set of registry entries identifying the functions of the configured task.

ma-task-options: An optional and possibly empty ordered list of options (name-value pairs) that are passed to the configured task.

ma-task-tags: An optional unordered set of tags that are reported together with the measurement results to a collector.

3.9.2. Definition of ma-option-obj

```
object {
    string              ma-option-name;
   [object              ma-option-value;]
} ma-option-obj;
```

The ma-option-obj models a name-value pair and consists of the following elements:

ma-option-name: The name of the option.

ma-option-value: The optional value of the option.

The ma-option-obj is used to define Task Configuration Options. Task Configuration Options are generally task specific. For tasks associated with an entry in a registry, the registry may define well-known option names (e.g., the so-called parameters in the IPPM metric registry [I-D.ietf-ippm-metric-registry]). Control and Reporting Tasks need to know the Channel they are going to use. The common option name for specifying the channel is "channel" where the option's value refers to the name of an ma-channel-obj.

3.10. Common Objects: Registry Information

Tasks and actions can be associated with entries in a registry. A registry object refers to an entry in a registry (identified by a URI) and it may define a set of roles.

3.10.1. Definition of ma-registry-obj

```
object {
    uri                ma-registry-uri;
   [string             ma-registry-role<0..*>;]
} ma-registry-obj;
```

The ma-registry-obj refers to an entry of a registry and it defines the associated role(s). The ma-registry-obj consists of the following elements:

ma-registry-uri: A URI identifying an entry in a registry.

ma-registry-role: An optional and possibly empty unordered set of roles for the identified registry entry.

3.11. Common Objects: Event Information

The Event information object used throughout the information models can initially take one of several different forms. Additional forms may be defined later in order to bind the execution of schedules to additional events. The initially defined Event forms are:

1. Periodic Timing: Emits multiple events periodically according to an interval time defined in seconds

2. Calendar Timing: Emits multiple events according to a calendar based pattern, e.g., 22 minutes past each hour of the day on weekdays

3. One Off Timing: Emits one event at a specific date and time

4. Immediate: Emits one event as soon as possible

5. Startup: Emits an event whenever the MA is started (e.g., at device startup)

6. Controller Lost: Emits an event when connectivity to the controller has been lost

7. Controller Connected: Emits an event when connectivity to the controller has been (re-)established

Optionally each of the Event options may also specify a randomness that should be evaluated and applied separately to each indicated event. This randomness parameter defines a uniform interval in seconds over which the start of the task is delayed from the starting times specified by the event object.

Both the Periodic and Calendar timing objects allow for a series of Actions to be executed. While both have an optional end time, it is best practice to always configure an end time and refresh the information periodically to ensure that lost MAs do not continue their tasks forever.

Startup events are only created on device startup, not when a new Instruction is transferred to the MA. If scheduled task execution is desired both on the transfer of the Instruction and on device restart then both the Immediate and Startup timing needs to be used in conjunction.

The datetime format used for all elements in the information model MUST conform to RFC 3339 [RFC3339].

3.11.1. Definition of ma-event-obj

```
object {
   string               ma-event-name;
   union {
       ma-periodic-obj             ma-event-periodic;
       ma-calendar-obj             ma-event-calendar;
       ma-one-off-obj              ma-event-one-off;
       ma-immediate-obj            ma-event-immediate;
       ma-startup-obj              ma-event-startup;
       ma-controller-lost-obj      ma-event-controller-lost;
       ma-controller-connected-obj ma-event-controller-connected;
   }
   [int                 ma-event-random-spread;]
   [int                 ma-event-cycle-interval;]
} ma-event-obj;
```

The ma-event-obj is the main event object. Event objects are identified by a name. A generic event object itself contains a more specific event object. The set of specific event objects should be extensible. The initial set of specific event objects is further described below. The ma-event-obj also includes an optional uniform random spread that can be used to randomize the start times of schedules triggered by an event. The ma-event-obj consists of the following elements:

ma-event-name: The name uniquely identifies an event object. Schedules refer to event objects by this name.

ma-event-periodic: The ma-event-periodic is present for periodic timing objects.

ma-event-calendar: The ma-event-calendar is present for calendar timing objects.

ma-event-one-off: The ma-event-one-off is present for one-off timing objects.

ma-event-immediate: The ma-event-immediate is present for immediate event objects.

ma-event-startup: The ma-event-startup is present for startup event objects.

ma-event-controller-lost: The ma-event-controller-lost is present for connectivity to controller lost event objects.

ma-event-controller-connected: The ma-event-controller-connected is present for connectivity to a controller established event objects.

ma-event-random-spread: The optional ma-event-random-spread adds a random delay defined in seconds to the event object. No random delay is added if ma-event-random-spread does not exist.

ma-event-cycle-interval: The optional ma-event-cycle-interval defines the duration of the time interval in seconds that is used to calculate cycle numbers. No cycle number is calculated if ma-event-cycle-interval does not exist.

3.11.2. Definition of ma-periodic-obj

```
object {
   [datetime           ma-periodic-start;]
   [datetime           ma-periodic-end;]
    int                ma-periodic-interval;
} ma-periodic-obj;
```

The ma-periodic-obj timing object has an optional start and an optional end time plus a periodic interval. Schedules using an ma-periodic-obj are started periodically between the start and end time. The ma-periodic-obj consists of the following elements:

ma-periodic-start: The optional date and time at which Schedules using this object are first started. If not present it defaults to immediate.

ma-periodic-end: The optional date and time at which Schedules using this object are last started. If not present it defaults to indefinite.

ma-periodic-interval: The interval defines the time in seconds between two consecutive starts of tasks.

3.11.3. Definition of ma-calendar-obj

Calendar Timing supports the routine execution of Schedules at specific times and/or on specific dates. It can support more flexible timing than Periodic Timing since the execution of Schedules

does not have to be uniformly spaced. For example a Calendar Timing could support the execution of a Measurement Task every hour between 6pm and midnight on weekdays only.

Calendar Timing is also required to perform measurements at meaningful times in relation to network usage (e.g., at peak times). If the optional timezone offset is not supplied then local system time is assumed. This is essential in some use cases to ensure consistent peak-time measurements as well as supporting MA devices that may be in an unknown timezone or roam between different timezones (but know their own timezone information such as through the mobile network).

The calendar elements within the Calendar Timing do not have defaults in order to avoid accidental high-frequency execution of Tasks. If all possible values for an element are desired then the wildcard * is used.

```
  object {
     [datetime            ma-calendar-start;]
     [datetime            ma-calendar-end;]
     [string              ma-calendar-months<0..*>;]
     [string              ma-calendar-days-of-week<0..*>;]
     [string              ma-calendar-days-of-month<0..*>;]
     [string              ma-calendar-hours<0..*>;]
     [string              ma-calendar-minutes<0..*>;]
     [string              ma-calendar-seconds<0..*>;]
     [int                 ma-calendar-timezone-offset;]
  } ma-calendar-obj;
```

ma-calendar-start: The optional date and time at which Schedules using this object are first started. If not present it defaults to immediate.

ma-calendar-end: The optional date and time at which Schedules using this object are last started. If not present it defaults to indefinite.

ma-calendar-months: The optional set of months (1-12) on which tasks scheduled using this object are started. The wildcard * means all months. If not present, it defaults to no months.

ma-calendar-days-of-week: The optional set of days of a week ("Mon", "Tue", "Wed", "Thu", "Fri",

	"Sat", "Sun") on which tasks scheduled using this object are started. The wildcard * means all days of the week. If not present, it defaults to no days.
ma-calendar-days-of-month:	The optional set of days of a months (1-31) on which tasks scheduled using this object are started. The wildcard * means all days of a months. If not present, it defaults to no days.
ma-calendar-hours:	The optional set of hours (0-23) on which tasks scheduled using this object are started. The wildcard * means all hours of a day. If not present, it defaults to no hours.
ma-calendar-minutes:	The optional set of minutes (0-59) on which tasks scheduled using this object are started. The wildcard * means all minutes of an hour. If not present, it defaults to no hours.
ma-calendar-seconds:	The optional set of seconds (0-59) on which tasks scheduled using this object are started. The wildcard * means all seconds of an hour. If not present, it defaults to no seconds.
ma-calendar-timezone-offset:	The optional timezone offest in hours. If not present, it defaults to the system’s local timezone.

If a day of the month is specified that does not exist in the month (e.g., 29th of Feburary) then those values are ignored.

3.11.4. Definition of ma-one-off-obj

```
object {
    datetime            ma-one-off-time;
} ma-one-off-obj;
```

The ma-one-off-obj timing object specifies a fixed point in time. Schedules using an ma-one-off-obj are started once at the specified date and time. The ma-one-off-obj consists of the following elements:

ma-one-off-time: The date and time at which Schedules using this object are started.

3.11.5. Definition of ma-immediate-obj

```
object {
                        // empty
} ma-immediate-obj;
```

The ma-immediate-obj event object has no further information elements. Schedules using an ma-immediate-obj are started as soon as possible.

3.11.6. Definition of ma-startup-obj

```
object {
                        // empty
} ma-startup-obj;
```

The ma-startup-obj event object has no further information elements. Schedules or suppressions using an ma-startup-obj are started at MA initialization time.

3.11.7. Definition of ma-controller-lost-obj

```
object {
                        // empty
} ma-controller-lost-obj;
```

The ma-controller-lost-obj event object has no further information elements. The ma-controller-lost-obj indicates that connectivity to the controller has been lost. This is determined by a timer started after each successful contact with a controller. When the timer reaches the controller-timeout (measured in seconds), an ma-controller-lost-obj event is generated. This event may be used to start a suppression.

3.11.8. Definition of ma-controller-connected-obj

```
object {
                        // empty
} ma-controller-connected-obj;
```

The ma-controller-connected-obj event object has no further information elements. The ma-controller-connected-obj indicates that connectivity to the controller has been established again after it was lost. This event may be used to end a suppression.

4. Example Execution

The example execution has two event sources E1 and E2 and three schedules S1, S2, and S3. The schedule S3 is started by events of event source E2 while the schedules S1 and S2 are both started by events of the event source E1. The schedules S1 and S2 have two actions each and schedule S3 has a single action. The event source E2 has no randomization while the event source E1 has the randomization r.

Figure 2 shows a possible timeline of an execution. The time T is progressing downwards. The dotted vertial line indicates progress of time while a dotted horizontal line indicates which schedule are triggered by an event. Tilded lines indicate data flowing from an action to another schedule. Actions within a schedule are named A1, A2, etc.

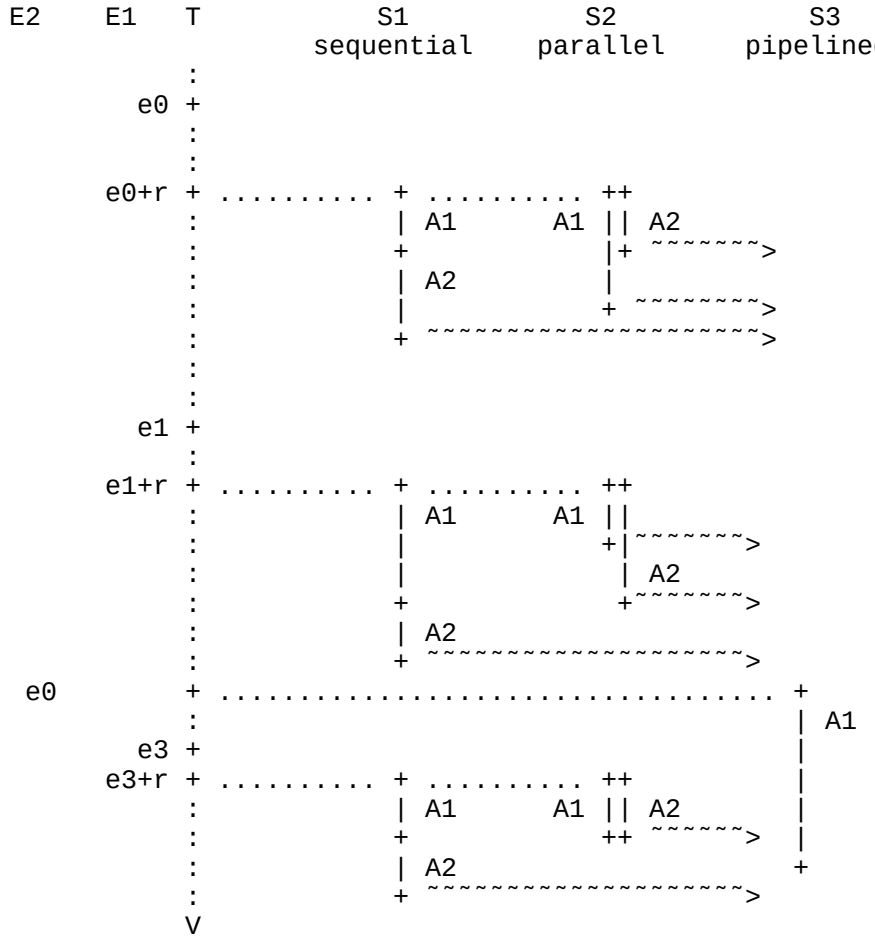

Figure 2: Example Execution

Note that implementations must handle possible concurrency issues. In the example execution, action A1 of schedule S3 is consuming the data that has been forwarded to schedule S3 while additional data is arriving from action A2 of schedule S2.

5. IANA Considerations

This document makes no request of IANA.

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

6. Security Considerations

This Information Model deals with information about the control and reporting of the Measurement Agent. There are broadly two security considerations for such an Information Model. Firstly the Information Model has to be sufficient to establish secure communication channels to the Controller and Collector such that other information can be sent and received securely. Additionally, any mechanisms that the Network Operator or other device administrator employs to pre-configure the MA must also be secure to protect unauthorized parties from modifying pre-configuration information. These mechanisms are important to ensure that the MA cannot be hijacked, for example to participate in a distributed denial of service attack.

The second consideration is that no mandated information items should pose a risk to confidentiality or privacy given such secure communication channels. For this latter reason items such as the MA context and MA ID are left optional and can be excluded from some deployments. This may, for example, allow the MA to remain anonymous and for information about location or other context that might be used to identify or track the MA to be omitted or blurred. Implementations and deployments should also be careful about exposing device-ids when this is not strictly needed.

An implementation of this Information Model should support all the security and privacy requirements associated with the LMAP Framework [RFC7594]. In addition, users of this Information Model are advised to choose identifiers for Group IDs, tags or names of information model objects (e.g., configured tasks, schedules or actions) that do not reveal any sensitive information to people authorized to process measurement results but who are not authorized to know details about the Measurement Agents that were used to perform the measurement.

7. Acknowledgements

Several people contributed to this specification by reviewing early versions and actively participating in the LMAP working group (apologies to those unintentionally omitted): Vaibhav Bajpai, Michael Bugenhagen, Timothy Carey, Alissa Cooper, Kenneth Ko, Al Morton, Dan Romascanu, Henning Schulzrinne, Andrea Soppera, Barbara Stark, and Jason Weil.

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

Juergen Schoenwaelder was partly funded by Flamingo, a Network of Excellence project (ICT-318488) supported by the European Commission under its Seventh Framework Programme.

8. References

8.1. Normative References

[ISO.10646] International Organization for Standardization, "Information Technology - Universal Multiple-Octet Coded Character Set (UCS)", ISO Standard 10646:2014, 2014.

[POSIX.2] The IEEE and The Open Group, "The Open Group Base Specifications Issue 7", IEEE Standard 1003.1-2008, 2016.

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

[RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, DOI 10.17487/RFC3339, July 2002, <http://www.rfc-editor.org/info/rfc3339>.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <http://www.rfc-editor.org/info/rfc3986>.

[RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", RFC 4122, DOI 10.17487/RFC4122, July 2005, <http://www.rfc-editor.org/info/rfc4122>.

8.2. Informative References

[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-10 (work in progress), November 2016.

[I-D.ietf-lmap-yang] Schoenwaelder, J. and V. Bajpai, "A YANG Data Model for LMAP Measurement Agents", draft-ietf-lmap-yang-10 (work in progress), January 2017.

[RFC3444] Pras, A. and J. Schoenwaelder, "On the Difference between Information Models and Data Models", RFC 3444, DOI 10 .17487/RFC3444, January 2003, <http://www.rfc-editor.org/info/rfc3444>.

[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, DOI 10.17487/RFC7398, February 2015, <http://www.rfc-editor.org/info/rfc7398>.

[RFC7536] Linsner, M., Eardley, P., Burbridge, T., and F. Sorensen, "Large-Scale Broadband Measurement Use Cases", RFC 7536, DOI 10.17487/RFC7536, May 2015, <http://www.rfc-editor.org/info/rfc7536>.

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

Appendix A. Change History

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

A.1. Non-editorial changes since -17

- The information model is subdivided into aspects and not sections.
- Changes to address the GEN-ART review comments.

A.2. Non-editorial changes since -16

- Addressing Alissa Cooper’s review comments.

A.3. Non-editorial changes since -15

- The reference to the framework is now informational.

A.4. Non-editorial changes since -14

- Clarified that the cycle number is in UTC.

A.5. Non-editorial changes since -13

- Removed the ma-config-device-id from the ma-config-obj.
- Added ma-config-report-group-id and clarified how two flags ma-config-report-agent-id and ma-config-report-group-id work.

A.6. Non-editorial changes since -12

- Renamed the ma-metrics-registry-obj to ma-registry-obj since tasks may refer to different registries (not just a metrics registry).
- Clarifications and bug fixes.

A.7. Non-editorial changes since -11

- Clarifications and bug fixes.

A.8. Non-editorial changes since -10

- Rewrote the text concerning the well-known "channel" option name.
- Added ma-report-result-event-time, ma-report-result-cycle-number, and ma-event-cycle-interval.
- Added ma-capability-tags.
- Added a new section showing an example execution.
- Several clarifications and bug fixes.

A.9. Non-editorial changes since -09

- Added ma-status-schedule-storage and ma-status-action-storage.
- Removed suppress-by-default.
- Moved ma-report-result-metrics of the ma-report-result-obj to ma-report-table-metrics of the ma-report-table-obj so that the relationship between metrics and result tables is clear.
- Added ma-report-conflict-obj.
- Added ma-report-result-status to ma-report-result-obj.
- Several clarifications and bug fixes.

A.10. Non-editorial changes since -08

- Refactored the ma-report-task-obj into the ma-report-result-obj.
- Introduced the ma-report-table-obj so that a result can contain multiple tables.
- Report schedule, action, and task name as part of the ma-report-result-obj.
- Report conflicts per ma-report-result-obj and not per ma-report-row-obj.
- Report the start/end time as part of the ma-report-result-obj.

A.11. Non-editorial changes since -07

- Added ma-schedule-end and ma-schedule-duration.
- Changed the granularity of scheduler timings to seconds.
- Added ma-status-suppression-obj to report the status of suppressions as done in the YANG data model.
- Added counters to schedule and action status objects to match the counters in the YANG data model.
- Using tags to pass information such as a measurement cycle identifier to the collector.
- Using suppression tags and glob-style matching to select schedules and actions to be suppressed.

A.12. Non-editorial changes since -06

- The default execution mode is pipelined (LI12)
- Added text to define which action consumes data in sequential, pipelines, and parallel execution mode (LI11)
- Added ma-config-measurement-point, ma-report-measurement-point, and ma-config-report-measurement-point to configure and report the measurement point (LI10)
- Turned ma-suppression-obj into a list that uses a start event and a stop event to define the start and end of suppression; this unifies the handling of suppression and loss of controller connectivity (LI09)

- Added ma-controller-lost-obj and ma-controller-ok-obj event objects (LI09)
- Added ma-status-schedule-obj to report the status of a schedule and refactored ma-task-status-obj into ma-status-action-obj to report the status of an action (LI07, LI08)
- Introduced a common ma-metric-registry-obj that identifies a metric and a set of associated roles and added this object to expose metric capabilities and to support the configuration of metrics and to report the metrics used (LI06)
- Introduced ma-capability-obj and ma-capability-task-obj to expose the capabilities of a measurement agent (LI05)
- Use ’ordered list’ or ’unordered set’ instead of list, collection, etc. (LI02)
- Clarification that Actions are part of a Schedule (LI03)
- Deleted terms that are not strictly needed (LI04)

A.13. Non-editorial changes since -05

- A task can now reference multiply registry entries.
- Consistent usage of the term Action and Task.
- Schedules are triggered by Events instead of Timings; Timings are just one of many possible event sources.
- Actions feed into other Schedules (instead of Actions within other Schedules).
- Removed the notion of multiple task outputs.
- Support for sequential, parallel, and pipelined execution of Actions.

Authors’ Addresses

Trevor Burbridge
BT
Adastral Park, Martlesham Heath
Ipswich IP5 3RE
United Kingdom

Email: trevor.burbridge@bt.com

Philip Eardley
BT
Adastral Park, Martlesham Heath
Ipswich IP5 3RE
United Kingdom

Email: philip.eardley@bt.com

Marcelo Bagnulo
Universidad Carlos III de Madrid
Av. Universidad 30
Leganes, Madrid 28911
Spain

Email: marcelo@it.uc3m.es

Juergen Schoenwaelder
Jacobs University Bremen
Campus Ring 1
Bremen 28759
Germany

Email: j.schoenwaelder@jacobs-university.de

INTERNET-DRAFT
Intended Status: Informational
Expires: August 15, 2015

Marc Linsner
Cisco Systems
Philip Eardley
Trevor Burbridge
BT
Frode Sorensen
Nkom

February 11, 2015

Large-Scale Broadband Measurement Use Cases
draft-ietf-lmap-use-cases-06

Abstract

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

Table of Contents

1 Introduction

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

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

2 Use Cases

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

2.1 Internet Service Provider (ISP) Use Case

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

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

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

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

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

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

2.2 Regulator Use Case

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

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

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

A regulator’s role in the development and enforcement of broadband Internet access service policies requires that the measurement approaches meet a high level of verifiability, accuracy and provider-independence to support valid and meaningful comparisons of Internet access service performance. Standards can help regulators’ shared needs for scalable, cost-effective, scientifically robust solutions to the measurement and collection of broadband Internet access

service performance information.

3 Details of ISP Use Case

3.1 Understanding the quality experienced by customers

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

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

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

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

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

and by sales and marketing.

3.2 Understanding the impact and operation of new devices and technology

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

- Check whether a customer can be upgraded to a new broadband option

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

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

- Investigate a quality of service mechanism (e.g. checking whether Diffserv markings are respected on some path); and so on.

3.3 Design and planning

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

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

With information gained from measurement results, capacity planning and network design can be more effective. Such planning typically uses simulations to emulate the measured performance of the current network and understand the likely impact of new capacity and potential changes to the topology. Simulations, informed by data from a limited panel of probes, can help quantify the advantage that a new technology brings and support the business case for larger roll-out.

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

potential impact if some new application becomes popular.

3.4 Monitoring Service Level Agreements

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

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

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

3.5 Identifying, isolating and fixing network problems

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

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

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

A more extensive deployment of the measurement capability to every broadband line would enable an operator to identify issues unique to

a single customer. Overall, large-scale measurements can help an operator help an operator fix the fault more rapidly and/or allow the affected customers to be informed what's happening. More accurate information enables the operator to reassure customers and take more rapid and effective action to cure the problem.

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

4 Details of Regulator Use Case

4.1 Providing transparent performance information

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

The published information needs to be:

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

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

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

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

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

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

4.2 Measuring broadband deployment

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

4.3 Monitoring traffic management practices

A regulator may want to monitor traffic management practices or compare the performance of Internet access service with specialized services offered in parallel to but separate from Internet access service (for example IPTV). A regulator could monitor for departures from application agnosticism such as blocking or throttling of traffic from specific applications, or preferential treatment of specific applications. A measurement system could send, or passively monitor, application-specific traffic and then measure

in detail the transfer of the different packets. Whilst it is relatively easy to measure port blocking, it is a research topic how to detect other types of differentiated treatment. The paper, "Glasnost: Enabling End Users to Detect Traffic Differentiation" [M-Labs NSDI 2010] and follow-on tool "Glasnost" [Glasnost] is an example of work in this area.

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

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

5 Implementation Options

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

One type of probe is a special hardware device that is connected directly to the home gateway. The devices are deployed to a carefully selected panel of end users and they perform measurements according to a defined schedule. The schedule can run throughout the day, to allow continuous assessment of the network. Careful design ensures that measurements do not detrimentally impact the home user experience or corrupt the results by testing when the user is also using the broadband line. The system is therefore tightly controlled by the operator of the measurement system. One advantage of this approach is that it is possible to get reliable benchmarks for the performance of a network with only a few devices. One disadvantage is that it would be expensive to deploy hardware devices on a mass scale sufficient to understand the performance of the network at the granularity of a single broadband user.

Another type of probe involves implementing the measurement

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

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

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

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

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

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

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

6 Conclusions

Large-scale measurements of broadband performance are useful for both network operators and regulators. Network operators would like to use measurements to help them better understand the quality experienced by their customers, identify problems in the network and design

network improvements. Regulators would like to use measurements to help promote competition between network operators, stimulate the growth of broadband access and monitor 'net neutrality'. There are other use cases that are not the focus of the initial LMAP charter (although it is expected that the mechanisms developed would be readily applied), for example end users would like to use measurements to help identify problems in their home network and to monitor the performance of their broadband provider.

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

A measurement capability is needed across a wide number of heterogeneous environments. Tests may be needed in the home network, in the ISP's network or beyond; they may be measuring a fixed or wireless network; they may measure just the access network or across several networks; at least some of which are not operated by the measurement provider.

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

Regardless of the tests being operated, there needs to be a way to demand or schedule the tests. Most use cases need a regular schedule of measurements, but sometimes ad hoc testing is needed, for example for troubleshooting. It needs to be ensured that measurements do not affect the user experience and are not affected by user traffic (unless desired). In addition there needs to be a common way to collect the results. Standardization of this control and reporting functionality allows the operator of a measurement system to buy the various components from different vendors.

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

diagnostics and network planning.

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

7 Security Considerations

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

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

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

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

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

* from the regulator’s perspective, the concerns include distortion of the measurement tests or alteration of the measurement results. Also, a malicious network operator could try to identify the broadband lines that the regulator was

measuring and prioritise that traffic ("game the system").

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

4. a measurement system that does not indicate who is responsible for the collection and processing of personal data and who is responsible for fulfilling the rights of users. The responsible party (often termed the "data controller") should, as good practice, consider issues such as defining:- the purpose for which the data is collected and used; how the data is stored, accessed, and processed; how long it is retained for; and how the end user can view, update, and even delete their personal data. If anonymized personal data is shared with a third party, the data controller should consider the possibility that the third party can de-anonymize it by combining it with other information.

These security and privacy issues will need to be considered carefully by any measurement system. In the context of LMAP, the [framework] considers them further along with some potential mitigations. Other LMAP documents will specify protocol(s) that enable the measurement system to instruct a probe about what measurements to make and that enable the probe to report the measurement results. Those documents will need to discuss solutions to the security and privacy issues. However, the protocol documents will not consider the actual usage of the measurement information; many use cases can be envisaged and, earlier in this document, we have described some likely ones for the network operator and regulator.

8 IANA Considerations

None

Contributors

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

James Miller jamesmilleresquire@gmail.com

Rachel Huang rachel.huang@huawei.com

Informative References

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

[Extend TCP] Michio Honda, Yoshifumi Nishida, Costin Raiciu, Adam Greenhalgh, Mark Handley and Hideyuki Tokuda. "Is it Still Possible to Extend TCP?" Proc. ACM Internet Measurement Conference (IMC), November 2011, Berlin, Germany. http://www.ietf.org/proceedings/82/slides/IRTF-1.pdf

[framework] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., Akhter, A. "A framework for large-scale measurement platforms (LMAP)", http://datatracker.ietf.org/doc/draft-ietf-lmap-framework/

[RFC6973] Cooper, A., Tschofenig, H.z., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, July 2013.

[RFC7258] Farrell, S., Tschofenig, H., "PPervasive Monitoring Is an Attack", RFC 7258, May 2014.

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

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

[M-Labs NSDI 2010] M-Lab, "Glasnost: Enabling End Users to Detect Traffic Differentiation", http://www.measurementlab.net/download/AMIfv945ljiJXzG-fgUrZSTu2hs1xRl5Oh-rpGQMWL305BNQh-BSq5oBoYU4a7zqXOvrztpJhK9gwk5unOe-fOzj4X-vOQz_HRrnYU-aFd0rv332RDReRfOYkJuagysstN3GZ__lQHTS8_UHJTWkrwyqIUjffVeDxQ/

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

[P.800] ITU-T, "SERIES P: TELEPHONE TRANSMISSION QUALITY Methods for objective and subjective assessment of quality", https://www.itu.int/rec/dologin_pub.asp?lang=e&id=T-REC-P.800-199608-I!!PDF-E&type=items

[MOS] Wikipedia, "Mean Opinion Score", http://en.wikipedia.org/wiki/Mean_opinion_score

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

Authors' Addresses

Marc Linsner
Cisco Systems, Inc.
Marco Island, FL
USA

EMail: mlinsner@cisco.com

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

Email: philip.eardley@bt.com

Trevor Burbridge
BT
B54 Room 77, Adastral Park, Martlesham
Ipswich, IP5 3RE
UK

Email: trevor.burbridge@bt.com

Frode Sorensen
Norwegian Communications Authority (Nkom)
Lillesand
Norway

Email: frode.sorensen@nkom.no

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 17, 2014

A. Morton
AT&T Labs
M. Bagnulo
UC3M
P. Eardley
BT

February 13, 2014

Active Performance Metric Sub-Registry
draft-mornuley-ippm-registry-active-00

Abstract

This memo defines the Active Performance Metrics sub-registry of the Performance Metric Registry. This sub-registry will contain Active Performance Metrics, especially those defined in RFCs prepared in the IP Performance Metrics (IPPM) Working Group of the IETF, and possibly applicable to other IETF metrics. Three aspects make IPPM metric registration difficult: (1) Use of the Type-P notion to allow users to specify their own packet types. (2) Use of flexible input variables, called Parameters in IPPM definitions, some of which determine the quantity measured and others of which should not be specified until execution of the measurement. (3) Allowing flexibility in choice of statistics to summarize the results on a stream of measurement packets.

This memo proposes a way to organize registry entries into columns that are well-defined, permitting consistent development of entries over time (a column may marked NA if it is not applicable for that metric). The design is intended to foster development of registry entries based on existing reference RFCs, whilst each column serves as a check-list item to avoid omissions during the registration process. Every entry in the registry, before IANA action, requires Expert review as defined by concurrent IETF work in progress "Registry for Performance Metrics" (draft-manyfolks-ippm-metric-registry).

The document contains two examples: a registry entry for an active Performance Metric entry based on RFC3393 and RFC5481, and a registry entry for an end-point Performance Metric based on RFC 7003. The examples are for Informational purposes and do not create any entry in the IANA registry.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

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 August 17, 2014.

Copyright Notice

Table of Contents

1. Introduction

[ISSUES

1. REAL-TIME OR INPUT PARAMETER [CONSISTENT WITH REGISTRY I-D] closed - just Parameter

2. CHANGED STREAM PARAMETER TO STREAM INPUT PARAMETER I didn't find any instances of this change - closed

3. I PREFER KEEPING THE CATEGORY-COLUMN HIERARCHY - ok we keep it

4. RATHER THAN BLANK COLUMNS, SHOULD WE HAVE 'NOT APPLICABLE' [MAYBE EVEN IANA REGISTERED??] sounds good to Al, used NA.

5. THE EXAMPLES ARE INFORMATIONAL NOT STANDARDS TRACK yes of course - -Closed.

-=-

Note: Efforts to synchronize terminology with [I-D.manyfolks-ippm-metric-registry] will likely be incomplete until both drafts are stable.

This memo defines the Active Performance Metrics sub-registry of the Performance Metric Registry. This sub-registry will contain Active Performance Metrics, especially those defined in RFCs prepared in the IP Performance Metrics (IPPM) Working Group of the IETF, according to their framework [RFC2330]. Three aspects make IPPM metric registration difficult: (1) Use of the Type-P notion to allow users to specify their own packet types. (2) Use of Flexible input variables, called Parameters in IPPM definitions, some which determine the quantity measured and others which should not be specified until execution of the measurement. (3) Allowing flexibility in choice of statistics to summarize the results on a stream of measurement packets. This memo uses terms and definitions from the IPPM literature, primarily [RFC2330], and the reader is assumed familiar with them or may refer questions there as necessary.

This sub-registry is part of the Performance Metric Registry [I-D.manyfolks-ippm-metric-registry] which specifies that all sub-registries must contain at least the following fields: the identifier, the name, the status, the requester, the revision, the revision date, the description for each entry, and the reference specifications used as the foundation for the Registered Performance Metric (see [I-D.manyfolks-ippm-metric-registry]).

Although there are several standard templates for organizing specifications of performance metrics (see [RFC2679] for an example of the traditional IPPM template, based to large extent on the Benchmarking Methodology Working Group’s traditional template in [RFC1242], and see [RFC6390] for a similar template), none of these templates was intended to become the basis for the columns of an IETF-wide registry of metrics. As we examined the aspects of metric specifications which need to be registered, it was clear that none of the existing metric templates fully satisfies the particular needs of a registry.

1.1. Background and Motivation

One clear motivation for having such a registry is to allow a controller to request a measurement agent to execute a measurement using a specific metric (see [I-D.ietf-lmap-framework]). Such a request can be performed using any control protocol that refers to the value assigned to the specific metric in the registry. Similarly, the measurement agent can report the results of the measurement and by referring to the metric value it can unequivocally identify the metric that the results correspond to.

There was a previous attempt to define a metric registry RFC 4148 [RFC4148]. However, it was obsoleted by RFC 6248 [RFC6248] because it was "found to be insufficiently detailed to uniquely identify IPPM

metrics... [there was too much] variability possible when characterizing a metric exactly" which led to the RFC4148 registry having "very few users, if any".

Our approach learns from this by tightly defining each entry in the registry with only a few parameters open, if any. The idea is that entries in the registry represent different measurement methods. Each may require run-time parameters to set factors like source and destination addresses, which do not change the fundamental nature of the measurement and can be set just before measurement execution. The downside of this approach is that it could result in a large number of entries in the registry. We believe that less is more in this context - it is better to have a reduced set of useful metrics rather than a large set of metrics with questionable usefulness. Therefore it is required for all registries within the Performance Metric Registry (see [I-D.manyfolks-ippm-metric-registry]) that the registry only includes commonly used metrics that are well defined; hence we require expert review policies for the approval and assignment of entries in this sub-registry.

There are several side benefits of having a registry with well-chosen entries. First, the registry could serve as an inventory of useful and used metrics that are normally supported by different implementations of measurement agents. Second, the results of the metrics would be comparable even if they are performed by different implementations and in different networks, as the metric and method is unambiguously defined.

The registry constitutes a key component of a 'Characterization Plan'. It describes various factors that need to be set by the party controlling the measurements, for example: specific values for the parameters associated with the selected registry entry (for instance, source and destination addresses); and how often the measurement is made. The Characterization Plan determines the individual Measurement Tasks which Measurement Agents will be instructed to do and which they then execute autonomously.

Measurement Instructions might look something like: "Dear measurement agent: Please start test DNS(example.com) and RTT(server.com,150) every day at 2000 GMT. Run the DNS test 5 times and the RTT test 50 times. Do that when the network is idle. Generate both raw results and 99th percentile mean. Send measurement results to collector.com in IPFIX format". The Characterization Plan depends on the requirements of the controlling party. For instance the broadband consumer might want a one-off measurement made immediately to one specific server; a regulator might want the same measurement made once a day until further notice to the 'top 10' servers; whilst an operator might want a varying series of tests (some of which will be

beyond those defined in an IETF registry) as determined from time to time by their operational support system. While the registries defined in this document help to define the Characterization Plan, its full specification falls outside the scope of this document, and other IETF work as currently chartered.

2. Scope

[I-D.manyfolks-ippm-metric-registry] defines the overall structure for a Performance Metric Registry and provides guidance for defining a sub registry.

This document defines the Active Performance Metrics Sub-registry; active metrics are those where the packets measured have been specially generated for the purpose.

A row in the registry corresponds to one Registered Performance Metric, with entries in the various columns specifying the metric. Section 3 defines the columns for a Registered Active Performance Metric.

As discussed in [I-D.manyfolks-ippm-metric-registry], each entry (row) must be tightly defined; the definition must leave open only a few parameters that do not change the fundamental nature of the measurement (such as source and destination addresses), and so promotes comparable results across independent implementations. Also, each registered entry must be based on existing reference RFCs (or other standards) for performance metrics, and must be operationally useful and have significant industry interest. This is ensured by expert review for every entry before IANA action.

3. Registry Categories and Columns

This section defines the categories and columns of the registry. Below, categories are described at the 3.x heading level, and columns are at the 3.x.y heading level. The Figure below illustrates this organization. An entry (row) therefore gives a complete description of a Registered Metric.

Each column serves as a check-list item and helps to avoid omissions during registration and expert review. In some cases an entry (row) may have some columns without specific entries, marked Not Applicable (NA).

```
    Registry Categories and Columns, shown as
                                                    Category
                                                    ------------------
                                                    Column |  Column |
   Common Registry Indexes and Information
   ------------------------------------------------------------------------
   ID | Name | Status |Request | Rev| Rev.Date | Description | Ref Spec|


   Metric Definition
   ----------------------------------------
   Reference Definition | Fixed Parameters |


   Method of Measurement
   ------------------------------------------------------------------
   Reference Method | Stream Type    | Output | Output | Run-time |
                    | and Parameters | Type   | Units  | Param

   Comments and Remarks
   --------------------
```

3.1. Common Registry Indexes and Information

This category has multiple indexes to each registry entry. It is defined in [I-D.manyfolks-ippm-metric-registry]:

3.1.1. Identifier

Defined in [I-D.manyfolks-ippm-metric-registry]. In order to have the document self contained, we could copy the definition from [I-D.manyfolks-ippm-metric-registry] here, but i guess we should do that once the definition in [I-D.manyfolks-ippm-metric-registry] is stable.

3.1.2. Name

Defined in [I-D.manyfolks-ippm-metric-registry], same comment than above.

3.1.3. Status

Defined in [I-D.manyfolks-ippm-metric-registry], same comment than above.

3.1.4. Requester

Defined in [I-D.manyfolks-ippm-metric-registry], same comment than above.

3.1.5. Revision

Defined in [I-D.manyfolks-ippm-metric-registry], same comment than above.

3.1.6. Revision Date

Defined in [I-D.manyfolks-ippm-metric-registry], same comment than above.

3.1.7. Description

Defined in [I-D.manyfolks-ippm-metric-registry], same comment as the previous.

3.1.8. Reference Specification(s)

Defined in [I-D.manyfolks-ippm-metric-registry], same comment as the previous.

3.2. Metric Definition

This category includes columns to prompt all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters, which are left open in the RFC but have a particular value defined by the performance metric.

3.2.1. Reference Definition

This entry provides references to relevant sections of the RFC(s) defining the metric, as well as any supplemental information needed to ensure an unambiguous definition for implementations.

3.2.2. Fixed Parameters

Fixed Parameters are input factors whose value must be specified in the Registry. The measurement system uses these values.

Where referenced metrics supply a list of Parameters as part of their descriptive template, a sub-set of the Parameters will be designated as Fixed Parameters. For example, Fixed Parameters determine most or

all of the IPPM Framework convention "packets of Type-P" as described in [RFC2330], such as transport protocol, payload length, TTL, etc.

A Parameter which is Fixed for one Registry entry may be designated as a Run-time Parameter for another Registry entry.

3.3. Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous method for implementations.

3.3.1. Reference Method

This entry provides references to relevant sections of the RFC(s) describing the method of measurement, as well as any supplemental information needed to ensure unambiguous interpretation for implementations referring to the RFC text.

3.3.2. Stream Type and Stream Parameters

Principally, two different streams are used in IPPM metrics, Poisson distributed as described in [RFC2330] and Periodic as described in [RFC3432]. Both Poisson and Periodic have their own unique parameters, and the relevant set of values is specified in this column.

Each entry for this column contains the following information:

- Value: The name of the packet stream scheduling discipline
- Stream Parameters: The values and formats of input factors for each type of stream. For example, the average packet rate and distribution truncation value for streams with Poisson-distributed inter-packet sending times.
- Reference: the specification where the stream is defined

The simplest example of stream specification is Singleton scheduling, where a single atomic measurement is conducted. Each atomic measurement could consist of sending a single packet (such as a DNS request) or sending several packets (for example, to request a webpage). Other streams support a series of atomic measurements in a "sample", with a schedule defining the timing between each transmitted packet and subsequent measurement.

3.3.3. Output Type and Data Format

For entries which involve a stream and many singleton measurements, a statistic may be specified in this column to summarize the results to a single value. If the complete set of measured singletons is output, this will be specified here.

Some metrics embed one specific statistic in the reference metric definition, while others allow several output types or statistics.

Each entry in the output type column contains the following information:

- Value: The name of the output type
- Data Format: provided to simplify the communication with collection systems and implementation of measurement devices.
- Reference: the specification where the output type is defined

The output type defines the type of result that the metric produces. It can be the raw results or it can be some form of statistic. The specification of the output type must define the format of the output. In some systems, format specifications will simplify both measurement implementation and collection/storage tasks. Note that if two different statistics are required from a single measurement (for example, both "Xth percentile mean" and "Raw"), then a new output type must be defined ("Xth percentile mean AND Raw").

3.3.4. Metric Units

The measured results must be expressed using some standard dimension or units of measure. This column provides the units.

When a sample of singletons (see [RFC2330] for definitions of these terms) is collected, this entry will specify the units for each measured value.

3.3.5. Run-time Parameters and Data Format

Run-Time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete. However, the values of these parameters is not specified in the Registry, rather these parameters are listed as an aid to the measurement system implementor or user (they must be left as variables, and supplied on execution).

Where metrics supply a list of Parameters as part of their descriptive template, a sub-set of the Parameters will be designated as Run-Time Parameters.

The Data Format of each Run-time Parameter SHALL be specified in this column, to simplify the control and implementation of measurement devices.

Examples of Run-time Parameters include IP addresses, measurement point designations, start times and end times for measurement, and other information essential to the method of measurement.

3.4. Comments and Remarks

Besides providing additional details which do not appear in other categories, this open Category (single column) allows for unforeseen issues to be addressed by simply updating this Informational entry.

4. Example IPPM Active Registry Entry

This section is Informational.

This section gives an example registry entry for the active metric described in [RFC3393], on Packet Delay Variation.

4.1. Registry Indexes

This category includes multiple indexes to the registry entries, the element ID and metric name.

4.1.1. Element ID

An integer having enough digits to uniquely identify each entry in the Registry.

4.1.2. Metric Name

A metric naming convention is TBD.

One possibility based on IPPM’s framework is:

Act_IP-UDP-One-way-pdv-95th-percentile-Poisson

4.1.3. Metric Description

An assessment of packet delay variation with respect to the minimum delay observed on the stream.

4.1.4. Other Info Columns not provided in Example

4.2. Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

4.2.1. Reference Definition

See sections 2.4 and 3.4 of [RFC3393]. Singleton delay differences measured are referred to by the variable name "ddT".

4.2.2. Fixed Parameters

Since the metric's reference supplies a list of Parameters as part of its descriptive template, a sub-set of the Parameters have been designated as designated as Fixed Parameters for this entry.

- F, a selection function defining unambiguously the packets from the stream selected for the metric. See section 4.2 of [RFC5481] for the PDV form.

- L, a packet length in bits. L = 200 bits.

- Tmax, a maximum waiting time for packets to arrive at Dst, set sufficiently long to disambiguate packets with long delays from packets that are discarded (lost). Tmax = 3 seconds.

- Type-P, as defined in [RFC2330], which includes any field that may affect a packet's treatment as it traverses the network. The packets are IP/UDP, with DSCP = 0 (BE).

4.3. Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

4.3.1. Reference Method

See section 2.6 and 3.6 of [RFC3393] for singleton elements.

4.3.2. Stream Type and Stream Parameters

Poisson distributed as described in [RFC2330], with the following Parameters.

- o lambda, a rate in reciprocal seconds (for Poisson Streams). lambda = 1 packet per second

- o Upper limit on Poisson distribution (values above this limit will be clipped and set to the limit value). Upper limit = 30 seconds.

4.3.3. Output Type and Data Format

See section 4.3 of [RFC3393] for details on the percentile statistic.

The percentile = 95.

Data format is a 32-bit unsigned floating point value.

Individual results (singletons) should be represented by the following triple

- o T1 and T2, times as described below in the Run-time parameters section.

- o ddT as defined in section 2.4 of [RFC3393]

if needed. The result format for ddT is *similar to* the short format in [RFC5905] (32 bits) and is as follows: the first 16 bits represent the *signed* integer number of seconds; the next 16 bits represent the fractional part of a second.

4.3.4. Metric Units

See section 3.3 of [RFC3393] for singleton elements.

[RFC2330] recommends that when a time is given, it will be expressed in UTC.

The timestamp format (for T, Tf, etc.) is the same as in [RFC5905] (64 bits) and is as follows: the first 32 bits represent the unsigned integer number of seconds elapsed since 0h on 1 January 1900; the next 32 bits represent the fractional part of a second that has elapsed since then.

4.3.5. Run-time Parameters and Data Format

Since the metric’s reference supplies a list of Parameters as part of its descriptive template, a sub-set of the Parameters have been designated as Run-Time Parameters for this entry. In related registry entries, some of the parameters below may be designated as Fixed Parameters instead.

- Src, the IP address of a host (32-bit value for IPv4, 128-bit value for IPv6)

- Dst, the IP address of a host (32-bit value for IPv4, 128-bit value for IPv6)

- T, a time (start of test interval, 128-bit NTP Date Format, see section 6 of [RFC5905])

- Tf, a time (end of test interval, 128-bit NTP Date Format, see section 6 of [RFC5905])

- T1, the wire time of the first packet in a pair, measured at MP(Src) as it leaves for Dst (64-bit NTP Timestamp Format, see section 6 of [RFC5905]).

- T2, the wire time of the second packet in a pair, measured at MP(Src) as it leaves for Dst (64-bit NTP Timestamp Format, see section 6 of [RFC5905]).

- I(i),I(i+1), i >=0, pairs of times which mark the beginning and ending of the intervals in which the packet stream from which the measurement is taken occurs. Here, I(0) = T0 and assuming that n is the largest index, I(n) = Tf (pairs of 64-bit NTP Timestamp Format, see section 6 of [RFC5905]).

4.4. Comments and Remarks

Lost packets represent a challenge for delay variation metrics. See section 4.1 of [RFC3393] and the delay variation applicability statement[RFC5481] for extensive analysis and comparison of PDV and an alternate metric, IPDV.

5. Example RTCP-XR Registry Entry

This section is Informational.

This section gives an example registry entry for the end-point metric described in RFC 7003 [RFC7003], for RTCP-XR Burst/Gap Discard Metric reporting.

5.1. Registry Indexes

This category includes multiple indexes to the registry entries, the element ID and metric name.

5.1.1. Element ID

An integer having enough digits to uniquely identify each entry in the Registry.

5.1.2. Metric Name

A metric naming convention is TBD.

5.1.3. Metric Description

TBD.

5.1.4. Other Info Columns not provided in Example

5.2. Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters. Section 3.2 of [RFC7003] provides the reference information for this category.

5.2.1. Reference Definition

Packets Discarded in Bursts:

The total number of packets discarded during discard bursts. The measured value is unsigned value. If the measured value exceeds 0xFFFFFD, the value 0xFFFFFE MUST be reported to indicate an over-range measurement. If the measurement is unavailable, the value 0xFFFFFF MUST be reported.

5.2.2. Fixed Parameters

Fixed Parameters are input factors that must be determined and embedded in the measurement system for use when needed. The values of these parameters is specified in the Registry.

Threshold: 8 bits, set to value = 3 packets.

The Threshold is equivalent to Gmin in [RFC3611], i.e., the number of successive packets that must not be discarded prior to and following a discard packet in order for this discarded packet to be regarded as part of a gap. Note that the Threshold is set in accordance with the Gmin calculation defined in Section 4.7.2 of [RFC3611].

Interval Metric flag: 2 bits, set to value 11=Cumulative Duration

This field is used to indicate whether the burst/gap discard metrics are Sampled, Interval, or Cumulative metrics [RFC6792]:

I=10: Interval Duration - the reported value applies to the most recent measurement interval duration between successive metrics reports.

I=11: Cumulative Duration - the reported value applies to the accumulation period characteristic of cumulative measurements.

Senders MUST NOT use the values I=00 or I=01.

5.3. Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations. For the Burst/Gap Discard Metric, it appears that the only guidance on methods of measurement is in Section 3.0 of [RFC7003] and its supporting references. Relevant information is repeated below, although there appears to be no section titled "Method of Measurement" in [RFC7003].

5.3.1. Reference Method

Metrics in this block report on burst/gap discard in the stream arriving at the RTP system. Measurements of these metrics are made at the receiving end of the RTP stream. Instances of this metrics block use the synchronization source (SSRC) to refer to the separate auxiliary Measurement Information Block [RFC6776], which describes measurement periods in use (see [RFC6776], Section 4.2).

This metrics block relies on the measurement period in the Measurement Information Block indicating the span of the report. Senders MUST send this block in the same compound RTCP packet as the Measurement Information Block. Receivers MUST verify that the measurement period is received in the same compound RTCP packet as this metrics block. If not, this metrics block MUST be discarded.

5.3.2. Stream Type and Stream Parameters

Since RTCP-XR Measurements are conducted on live RTP traffic, the complete description of the stream is contained in SDP messages that proceed the establishment of a compatible stream between two or more communicating hosts. See Run-time Parameters, below.

5.3.3. Output Type and Data Format

The output type defines the type of result that the metric produces.

- Value: Packets Discarded in Bursts
- Data Format: 24 bits
- Reference: Section 3.2 of [RFC7003]

5.3.4. Metric Units

The measured results are apparently expressed in packets, although there is no section of [RFC7003] titled "Metric Units".

5.3.5. Run-time Parameters and Data Format

Run-Time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete. However, the values of these parameters is not specified in the Registry, rather these parameters are listed as an aid to the measurement system implementor or user (they must be left as variables, and supplied on execution).

The Data Format of each Run-time Parameter SHALL be specified in this column, to simplify the control and implementation of measurement devices.

SSRC of Source: 32 bits As defined in Section 4.1 of [RFC3611].

SDP Parameters: As defined in [RFC4566]

Session description v= (protocol version number, currently only 0)

o= (originator and session identifier : username, id, version number, network address)

s= (session name : mandatory with at least one UTF-8-encoded character)

i=* (session title or short information) u=* (URI of description)

e=* (zero or more email address with optional name of contacts)

p=* (zero or more phone number with optional name of contacts)

c=* (connection information--not required if included in all media)

b=* (zero or more bandwidth information lines) One or more Time descriptions ("t=" and "r=" lines; see below)

z=* (time zone adjustments)

k=* (encryption key)

a=* (zero or more session attribute lines)

Zero or more Media descriptions (each one starting by an "m=" line; see below)

m= (media name and transport address)

i=* (media title or information field)

c=* (connection information -- optional if included at session level)

b=* (zero or more bandwidth information lines)

k=* (encryption key)

a=* (zero or more media attribute lines -- overriding the Session attribute lines)

An example Run-time SDP description follows:

v=0

o=jdoe 2890844526 2890842807 IN IP4 192.0.2.5

s=SDP Seminar i=A Seminar on the session description protocol

u=http://www.example.com/seminars/sdp.pdf e=j.doe@example.com (Jane Doe)

c=IN IP4 233.252.0.12/127

t=2873397496 2873404696

a=recvonly

m=audio 49170 RTP/AVP 0

m=video 51372 RTP/AVP 99

a=rtpmap:99 h263-1998/90000

5.4. Comments and Remarks

TBD.

6. Example BLANK Registry Entry

This section is Informational. (?)

This section gives an example registry entry for the <type of metric and specification reference> .

6.1. Registry Indexes

This category includes multiple indexes to the registry entries, the element ID and metric name.

6.1.1. Element ID

An integer having enough digits to uniquely identify each entry in the Registry.

6.1.2. Metric Name

A metric naming convention is TBD.

6.1.3. Metric Description

A metric Description is TBD.

6.1.4. Other Info Columns not provided in Example

6.2. Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the RFC reference and values of input factors, called fixed parameters.

<possible section reference>.

6.2.1. Reference Definition

6.2.2. Fixed Parameters

Fixed Parameters are input factors that must be determined and embedded in the measurement system for use when needed. The values of these parameters is specified in the Registry.

<list fixed parameters>

6.3. Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous methods for implementations.

6.3.1. Reference Method

For <metric>.

<section reference>

6.3.2. Stream Type and Stream Parameters

<list of stream parameters>.

<references>

6.3.3. Output Type and Data Format

The output type defines the type of result that the metric produces.

- Value:
- Data Format: (There may be some precedent to follow here, but otherwise use 64-bit NTP Timestamp Format, see section 6 of [RFC5905]).
- Reference: <section reference>

6.3.4. Metric Units

The measured results are expressed in <units>,

<section reference>.

6.3.5. Run-time Parameters and Data Format

Run-time Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

<list of run-time parameters>

<reference(s)>.

6.4. Comments and Remarks

Additional (Informational) details for this entry

7. Security Considerations

This registry has no known implications on Internet Security.

8. IANA Considerations

IANA is requested to create The Active Performance Metric Sub-registry within the Performance Metric Registry defined in [I-D.manyfolks-ippm-metric-registry]. The Sub-registry will contain the following categories and (bullet) columns, (as defined in section 3 above):

Common Registry Indexes and Info

- Identifier
- Name
- Status
- Requester
- Revision
- Revision Date
- Description
- Reference Specification(s)

Metric Definition

- Reference Definition
- Fixed Parameters

Method of Measurement

- Reference Method
- Stream Type and Parameters
- Output type and Data format

- o Metric Units
- o Run-time Parameters

Comments and Remarks

9. Acknowledgements

The authors thank Brian Trammell for suggesting the term "Run-time Parameters", which led to the distinction between run-time and fixed parameters implemented in this memo, and the IPFIX metric with Flow Key as an example.

10. References

10.1. Normative References

[I-D.manyfolks-ippm-metric-registry] Bagnulo, M., Claise, B., Eardley, P., and A. Morton, "Registry for Performance Metrics", Internet Draft (work in progress) draft-manyfolks-ippm-metric-registry, 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.

[RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999.

[RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.

[RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.

[RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.

[RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", RFC 3432, November 2002.

[RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", RFC 4737, November 2006.

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

[RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.

10.2. Informative References

[Brow00] Brownlee, N., "Packet Matching for NeTraMet Distributions", March 2000.

[I-D.ietf-lmap-framework] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A framework for large-scale measurement platforms (LMAP)", draft-ietf-lmap-framework-03 (work in progress), January 2014.

[RFC1242] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.

[RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", BCP 108, RFC 4148, August 2005.

[RFC5472] Zseby, T., Boschi, E., Brownlee, N., and B. Claise, "IP Flow Information Export (IPFIX) Applicability", RFC 5472, March 2009.

[RFC5477] Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", RFC 5477, March 2009.

[RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, March 2009.

[RFC6248] Morton, A., "RFC 4148 and the IP Performance Metrics (IPPM) Registry of Metrics Are Obsolete", RFC 6248, April 2011.

[RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, October 2011.

[RFC7003] Clark, A., Huang, R., and Q. Wu, "RTP Control Protocol (RTCP) Extended Report (XR) Block for Burst/Gap Discard Metric Reporting", RFC 7003, September 2013.

Authors' Addresses

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown,, NJ 07748
USA

Phone: +1 732 420 1571
Fax: +1 732 368 1192
Email: acmorton@att.com
URI: http://home.comcast.net/~acmacm/

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

Philip Eardley
British Telecom
Adastral Park, Martlesham Heath
Ipswich
ENGLAND

Email: philip.eardley@bt.com