

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 14, 2015

M. Bagnulo
UC3M
T. Burbridge
BT
S. Crawford
SamKnows
J. Schoenwaelder
V. Bajpai
Jacobs University Bremen
September 10, 2014

Large MeAsurement Platform Protocol
draft-bagnulo-lmap-http-03

Abstract

This documents specifies the LMAP protocol based on HTTP for the Control and Report in Large Scale Measurement Platforms.

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 March 14, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

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

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

Table of Contents

1. Introduction	2
2. Overview	4
3. Naming Considerations	4
4. Information model	6
5. Transport protocol	7
5.1. Pre-configured information	7
5.2. Control Protocol	7
5.2.1. Retrieving Instructions	8
5.2.2. Handling communication failures	10
5.2.3. Pushing Information from the Controller to the MA	10
5.3. Report protocol	11
5.3.1. Handling communication failures	12
6. LMAP Data Model	13
6.1. Timing Information	13
6.2. Channels	15
6.3. Configuration	15
6.4. Instruction	16
6.5. Measurement Supression	16
6.6. Measurement Task Configurations	16
6.7. Measurement Schedules	17
6.8. Logging	18
6.9. Capability and Status	18
6.10. Reporting	19
7. Security considerations	22
8. IANA Considerations	24
9. Acknowledgments	24
10. References	24
10.1. Normative References	24
10.2. Informative References	25
Authors' Addresses	25

1. Introduction

A Large MeAsurement Platform (LMAP) is an infrastructure deployed in the Internet that enables performing measurements from a very large number of vantage points.

The main components of a LMAP are the following:

- o The Measurement Agents (MAs): these are the processes that perform the measurements. The measurements can be both active or passive measurements.

- o The Controller: this is the element that controls the MAs. In particular it provides configuration information and it instructs the MA to perform a set of measurements.
- o The Collector: this is the repository where the MAs send the results of the measurements that they have performed.

These and other terms used in this document are defined in [I-D.ietf-lmap-framework]. We only include the definition of the main elements in this document so it is self-contained and can be read without the need to consult other documents. The reader is referred to the terminology draft for further details.

In order for a LMAP to work, the following protocols are required:

- o Measurement protocols: These are the protocols used between the MA and the Measurement Peer in active measurements. These are the actual packets being used for the measurement operations.
- o Control Protocol. This is the protocol between the Controller and the MAs. This protocol is used to convey measurement Instruction(s) from the Controller to the MA as well as logging, failure and capabilities information from the MA to the Controller.
- o Report Protocol. This is the protocol between the MAs and the Collector. This protocol conveys information about the results of the measurements performed by the MA to the Collector.

Both the Control protocol and the Report protocol have essentially two parts: a transport and a data model. The data model represents the information about measurement instructions and logging/failure/capabilities (in the Control protocol) and the information about measurement results (in the Report protocol) that is being exchanged between the parties. The transport is the underlying protocol used to exchange that information. This document specifies the use of HTTP 1.1 [RFC7230] [RFC7231] [RFC7232] [RFC7233] [RFC7234] [RFC7235] as a transport for the Control and the Report protocol. This document also defines the data model for the Control and Report protocols. The data model described in this document follows the information model described in [I-D.ietf-lmap-information-model]. The Measurement protocols are out of the scope for this document.

At this stage, the goal of this document is to explore different options that can be envisioned to use the HTTP protocol to exchange LMAP information and to foster discussion about which one to use (if any). Because of that, the document contains several discussion paragraphs that explore different alternative approaches to perform the same function.

2. Overview

This section provides an overview of the architecture envisioned for a LMAP using HTTP as transport protocol. As we described in the previous section, a LMAP is formed by a large number of MAs, one or more Controllers and one or more Collectors. We assume that before the MAs are deployed, it is possible to pre-configure some information in them. Typically this includes information about the MA itself (like its identifier), security information (like some certificates) and information about the Controller(s) available in the measurement platform. Once that the MA is deployed it will retrieve additional configuration information from the pre configured Controller. After obtaining the configuration information, the MA is ready to receive Instructions from the Controller and initiate measurement tasks. The MA will perform the following operations:

- o It will obtain Instructions from one of the configured Controllers. These Instructions include information about the set of measurement tasks to be performed, a schedule for the execution of the measurements as well as a set of report channels. This information is downloaded by the MA from the Controller. The MA will periodically check whether there are new Instructions available from the Controller. This document specifies how the MA uses the HTTP protocol to retrieve information from the Controller.
- o The MA will execute measurement tasks either by passively listening to traffic or by actively sending and receiving measurement packets. How this is done is out of the scope of this document.
- o After one or more measurements have been performed, the MA reports the results to the Collector. The timing of these uploads is specified in the measurement Instruction i.e. each measurement specified in a measurement Instruction contains a report information, defining when the MA should report the results back to the Collector. This document specifies how the MA uses the HTTP protocol to upload the measurement results to the Collector.
- o In addition, the MA will periodically report back to the Controller information about its capabilities (like the number of interfaces it has, the corresponding IP addresses, the set of measurement methods it supports, etc) and also logging information (whether some of the requested measurement tasks failed and related information).

3. Naming Considerations

In this section we define how the different elements of the LMAP architecture are identified and named.

The Controller and the Collectors can be assumed to have both an IP address and a Fully Qualified Domain Name (FQDN). It is natural to use these as identifiers for these elements. In this document we will use FQDNs, but IP addresses can be used as well.

The MAs on the other hand, are likely to be executed in devices located in the end user premises and are likely to be located behind a NAT box. It is reasonable to assume they have neither a public IP address nor a FQDN. We propose then that the MAs are identified using an Universally Unique Identifier URN as defined in RFC 4122 [RFC4122]. In particular each MA has a version 4 UUID, which is randomly or pseudo randomly generated.

DISCUSSION:

MA ID Configuration: Some open issues related to this are: a) whether the MA ID is configured before or after the MA is deployed, b) if configured after deployment whether the MA ID is generated locally and posted or fetched from the Controller and c) whether this is within the scope of this (or other) specification if any. These issues seem also to be related to the nature of the MA platform (whether the MA is a software downloaded into a general purpose device or it is a special purpose hardware box). Consider the case that the MA is located in a special purpose hardware box, then having the MA ID pre configured before deployment requires a per device customization that is expensive. It would be more costly efficient to reuse an existent (hopefully) unique identifier available in the hardware (such as a MAC address) to serve as a one-time pre configured identifier to be used to fetch (or post a self generated) the MA ID from the Controller once the MA is deployed. The requirement for such one-time identifier is that they must be unique (which is not always true for the MACs). About the local generation of the MA ID (as opposed to fetch it from the Controller), the generation process performed in the MA MUST be idempotent, i.e. if the MA was factory-reset then the server would still see it with the same MA ID when it came back up. This is probably easier to achieve if it is generated in the Controller and then fetched by the MA. Finally, it is not clear at this stage if this needs to be specified in this document or in the information model document or left open to the implementers. Group identifiers. In some cases, like the case of measurements in mobile devices, it may be important because of privacy considerations for the MA not to have a unique identifier. It is possible then to assign "Group identifiers" to a set of devices that share relevant characteristics from the measurement perspective (e.g. devices from the same operator, with the same type of contract or other relevant feature). In this case, the MAs within the same group would retrieve common measurement

Instructions from the controller by presenting the same Group ID and would report results including the Group ID in the report. This would imply that it would not be possible for the platform to correlate specific measurement data with any given MA. The downside of this is that some MAs may be over-represented while other under-represented in the measurement data and it would not be possible to detect this case (for instance a given MA may have reported 20 results while another one only one). In order to deal with this issue, the MA behaviour must be programmed accordingly (e.g. the MA should not perform more than one measurement every given period of time). In addition, it should be noted that privacy is only achieved in a holistic way. This means that really anonymity of the MA is incompatible with strong authentication. In particular, if a measurement platform's goal is to keep MAs anonymous, it cannot require any form of strong authentication (other than weak group authentication e.g. a password shared by a group), which has security implications. In particular, the threat for report forgery (i.e. enabling an attacker to submit forged reports as discussed in the security considerations) increases.

There are additional naming considerations related to:

- o The measurements. In order to enable a Controller to properly convey a measurement schedule, it must be possible for the Controller to specify a measurement to be performed while providing the needed input parameters. While this is critical, it is out of the scope of this document. There is a proposed registry for metrics/measurements in [I-D.bagnulo-ippm-new-registry-independent])
- o The resources being exchanged, namely, the configuration information, the measurement Instructions and the reports. These are being discussed in the upcoming sections.

4. Information model

The information model for LMAP is described [I-D.ietf-lmap-information-model]. It contains basically two models one for the control information (i.e. the Instructions from the Controller to the MA) and a model for the Report information. We briefly describe their overall structure here.

The control information (or Instruction) has the following five elements:

- o The Set of Measurement Task Configurations: This element defines the measurements/test that the MA will perform without defining the schedule when they will be performed.

- o The Set of Report Channels: This element defines the set of collectors as well as the reporting schedules for the reports.
- o The Set of Measurement Schedules for Repeated Tasks: defines the schedules for the repeated measurements, by referencing the measurement tasks defined in the second element.
- o Suppression information

Summary of Report information model here.

Summary of Capability and Status information model here.

Summary of Logging information model here.

5. Transport protocol

5.1. Pre-configured information

As we mentioned earlier, the MAs contain pre-configured information before being deployed. The pre-configured information is the following:

- o The UUID for the MA. This should be pre-configured so that the Controller is aware of the MA and can feed configuration information and measurement Instructions to it.
- o Information about one or more Controllers. The MA MUST have enough information to create the URL for the Instruction resources. This includes the the FQDN of each of the Controller or the IP addresses of the Controller, as well as the well-known path prefix and its identifier.
- o The certificate for the Certification authority that is used in the platform to generate the certificates for the Controller and the Collector. See the Security considerations section below.
- o The security related information for the MA (it can be a certificate for the MA and the corresponding private key, or simply a key/password depending on the security method used, see the security considerations section below).

5.2. Control Protocol

The Control protocol is used by the MA to retrieve Instruction information from the Controller. In this section we describe how to use HTTP to transport Instructions. The Instruction information is structured as defined in the LMAP Information model [I-D.ietf-lmap-information-model] as described in the previous section. The MA uses the Control protocol to retrieve all the resources described above, namely, the Agent information, the Set of Measurement Task Configurations, the Set of Report Channels, the Set of Measurement Schedules for Repeated Tasks and the Set of

Measurement Schedules for Isolated Tasks. The main difference from the HTTP perspective is that the MA MUST have the URL for the Agent Information resource pre-configured as described in the previous section, while the URLs for all the other resources are contained in the Agent Information resource itself.

5.2.1. Retrieving Instructions

In order to retrieve the Instruction resources from the Controller the MA can use either the GET or the POST method using the corresponding URL.

5.2.1.1. Using the GET method

One way of using the GET method to retrieve configuration information is to explicitly name the configuration information resources and then apply the GET method. The MA retrieves its Instruction when it is first connected to the network and periodically after that. The frequency for the periodical retrieval is contained in the Agent Information (???).

The URL for the Agent Information resource is formed as the FQDN of the Controller, a well-known path prefix and the MA UUID. The well-known path prefix is /.well-known/lmap/ma-info. The URL for the remaining resources that compose the Instruction are contained in the Agent Information.

Agent Information retrieval: In order to retrieve the Agent information the MA uses the HTTP GET method follows:

```
GET /.well-known/lmap/ma-info/ < ma-iid> HTTP/1.1
Host: FQDN or IP of the Controller
Accept: application/json (as per [RFC7159])
```

The Agent Information should contain the Configuration Retrieval Schedule (i.e. how often the MA should retrieve configuration information) and also the Measurement Instruction Retrieval Schedule (i.e. how often the MA should retrieve the Measurement Instruction from the Controller). COMMENT: this is missing from the Data Model

The retrieval of the remaining resources of the Instruction using the GET method is analogous, only that the URL is extracted from the Agent Information file rather than constructed with pre-configured information.

The format for the response should be described here

Periodical Instruction retrieval: After having downloaded the initial Instruction information, the MA will periodically look for updated Instruction information. The frequency with which the MA polls for the new Instructions from the Controller is contained in the last Agent Information downloaded. In order to retrieve the Agent Information, the MA uses the GET method as follows:

```
GET /.well-known/lmap/ma-info/ma-iid/ HTTP/1.1
Host: FQDN or IP of the Controller
Accept: application/json (as per [RFC7159])
If-None-Match: the eTag of the last retrieved Agent Information
(an alternative option here is to use If-Modified-Since, not sure
which one is best)
```

For the other Instruction resources, the GET method is applied in the same way just that the URL used are the ones retrieved in the last Agent Information.

The format for the response should be described here

Alternatively, instead of explicitly naming the Instruction resources for each MA, it is possible to perform a query using the GET method as well. In this case, the MA could perform a GET for the following URI `http://controller.example.org/?ma=maid & q=ma-info` (similar queries can be constructed for the other Instruction resources). (I am not sure how to express in this case the condition that the MA wishes to retrieve the configuration if it is newer than the last one it downloaded.)

5.2.1.2. Using the POST method

An alternative to retrieve Instruction resources is to use the POST method to perform a query (similar to the query using GET). In this case there is no explicit naming of the Instruction information of each MA, but a general Instruction resource and the POST method is used to convey a query for the Instruction information of a particular MA. For the case of the Agent Information resource, this would look like as follows:

```
POST /.well-known/lmap/ma-info/ma-iid/ HTTP/1.1
Host: controller.example.com
Content-Type: application/lmap-maid+json
Accept: application/lmap-config+json
{
  "ma-id" : "550e8400-e29b-11d4-a716-446655440000",
}
```

The reply for this query would contain the actual configuration information as follows:

```
HTTP/1.1 200 OK
Content-Length: xxx
Content-Type: application/lmap-config+json
{
  // whatever config goes here
}
```

In this case, the URLs contained in the Agent information can be generic and not MA specific, since the MA will use the POST method including its own identifier when retrieving the Instruction resources.

The argument for this approach is that this is much more extensible since the POST can carry complex information and there is no need to "press" arguments into the strict hierarchy of URIs.

We need to describe how to use this to retrieve newer information in the periodic case.

5.2.2. Handling communication failures

The cases that the MA is unable to retrieve the Instructions are handled as follows:

- o The MA will use a timeout for the communication of TIMEOUT seconds. The value of TIMEOUT MUST be configurable via the aforementioned Configuration Information retrieval protocol. The default value for the TIMEOUT is 3 seconds. If after the timeout, the communication with the Controller has not been established, the MA will retry doing an exponential backoff and doing a round robin between the different Controllers it has available.
- o If a HTTP error message (5xx) is received from the Controller as a response to the GET request, the MA will retry doing an exponential back-off and doing a round robin between the different Controllers it has available. The 5xx error codes indicate that this Controller is currently incapable of performing the requested operation.

5.2.3. Pushing Information from the Controller to the MA

The previous sections described how the MA periodically polls the Controller to retrieve Instruction information. The frequency of the downloads is configurable. The question is whether this is enough or a mechanism for pushing Instruction information is needed. Such method would enable to contact the MA in any moment and take actions

like triggering a measurement right away or for instance to stop an ongoing measurement (e.g. because it is disturbing the network). The need for such a mechanism is likely to depend on the use case of the platform. Probably the ISP use case is more likely to require this feature than the regulator/benchmarking use case. It is probably useful then to provide this as an optional feature.

The main challenge in order to provide this feature is that the MAs are likely to be placed behind NATs, so it is not possible for the Controller to initiate a communication with the MA unless there is a binding in the NAT to forward the packets to the MA. There are several options that can be considered to enable this communication:

- o The MA can use one of the NAT control protocols, such as PCP or UPNP. If this approach is used, the MA will create a binding in the NAT opening a hole. After that, the MA should inform the Controller about which is the IP address and port available for communication. It would be possible to re-use existing protocols to forward this information. The problem with this is that the NAT may not support these protocols or they may not be activated. In any case, a solution should try to use them in the case they are available.
- o If it is not possible to use a NAT control protocol, then the MA can open a hole in the NAT by establishing a connection to the Controller and keeping it open. This allows the Controller to push information to the MA through that connection. One concern with this approach is that the MA is playing the role of the client and the Controller is playing the role of the server (the MA is initiating the TCP connection), but it would be the Controller who would use the PUT method towards the MA reversing the roles. An alternative approach is that the MA has a long running GET pending which is answered by the server if the measurement Instruction changes (or the server times out, in which case the MA restarts the long running GET. More discussion is needed about whether one of these options is acceptable or not. In addition, this would imply that the Controller should maintain as many open sessions as MAs it is managing, which imposes additional burden in the Controller. There are security considerations as well, but these are covered in the Security Considerations section below.

5.3. Report protocol

The MA after performing the measurements reports the results to a collector. There can be more than one collector within a LMAP framework. Each collector is identified by its FQDN or IP address which is retrieved as part of the Agent information from a pre-configured controller as previously discussed. The number of

Collectors that the MA uploads the results to as well as the schedule when it does so is defined in the measurement Instruction previously downloaded from the Controller. The MA themselves are identified by a UUID.

There are two options that can be considered for the MA to upload reports to the Collector either to use the PUT method or to use the POST method.

If the PUT method option is used, then the MA need to perform the PUT method using an explicit name for the report resource it is transferring to the Collector. The name of the resource is contained in the Agent Information previously retrieved by the MA

The other option is for the MA to use the POST method to upload the measurement reports to one or more Collectors. In this case,, the POST message body can contain the identifier of the MA and additional information describing the report in addition to the report itself.

One argument to consider is that PUT is idempotent. This means that if the network is bad at some point and the MA is not sure whether its request made it through, it can send it a second (or nth) time, and it is guaranteed that the request will have exactly the same effect as sending it for the first time. POST does not by itself guarantee this. This can be achieved by verifying the report data itself, and contrast it with data already stored in the Collector database.

5.3.1. Handling communication failures

The MA will use a timeout for the communication with the Collector of TIMEOUT seconds. The value of TIMEOUT MUST be configurable via the aforementioned Configuration Information retrieval protocol. The default value for the TIMEOUT is 3 seconds.

If the MA is uploading the report to several Collectors and it manages to establish the communication before TIMEOUT seconds with at least one of them, but not with one or more of the other Collectors, then the MA gives up after TIMEOUT seconds and it MAY issue an alarm. The definition of how to do that operation is out of the scope of this document.

If the MA is uploading the report to only one Collector, and it does not manages to establish a communication before TIMEOUT seconds, then it retry doing an exponential backoff and doing a round robin between the different Collectors it has available.

Similarly, if an HTTP error message (5xx) is received from the Collector as a response to the PUT request, the MA will retry doing an exponential backoff and doing a round robin between the different Collectors it has available. The 5xx error codes indicate that this Collector is currently incapable of performing the requested operation.

In order to support this, the information model must express the difference between a report sent to multiple collectors and multiple collectors used for fallback.

6. LMAP Data Model

This section will contain the data model in json.

6.1. Timing Information

An example immediate timing object with no defined randomness is shown below:

```
{
  "timings": [
    {
      "id": 1,
      "ma_periodic_end": 1410017611,
      "ma_periodic_interval": 3600000,
      "ma_periodic_start": 1410017613,
      "ma_randomness_spred": 0,
      "ma_timing_name": "hourly"
    },
    {
      "id": 3,
      "ma_periodic_end": 1410017611,
      "ma_periodic_interval": 86400,
      "ma_periodic_start": 1410017613,
      "ma_randomness_spred": 0,
      "ma_timing_name": "daily"
    },
    {
      "id": 2,
      "ma_periodic_end": 1410017611,
      "ma_periodic_interval": 3600000,
      "ma_periodic_start": 1410017613,
      "ma_randomness_spred": 0,
      "ma_timing_name": "hourly"
    }
  ]
}
```

```
    "id": 4,
    "ma_calendar_days_of_month": "",
    "ma_calendar_days_of_week": "tuesday, thursday, sunday",
    "ma_calendar_end": 1410017613,
    "ma_calendar_hours": "18",
    "ma_calendar_minutes": "04",
    "ma_calendar_months": "",
    "ma_calendar_seconds": "42",
    "ma_calendar_start": 1410017612,
    "ma_calendar_timezone_offset": 2,
    "ma_randomness_spred": 0,
    "ma_timing_name": "tuesday-thursday-sunday"
  },
  {
    "id": 5,
    "ma_calendar_days_of_month": "",
    "ma_calendar_days_of_week": "",
    "ma_calendar_end": 1410017619,
    "ma_calendar_hours": "0, 6 12 18",
    "ma_calendar_minutes": "0",
    "ma_calendar_months": "",
    "ma_calendar_seconds": "0",
    "ma_calendar_start": 1410017612,
    "ma_calendar_timezone_offset": 2,
    "ma_randomness_spred": 21600000,
    "ma_timing_name": "once-every-six-hours"
  },
  {
    "id": 6,
    "ma_one_off_time": 1410017613,
    "ma_randomness_spred": 0,
    "ma_timing_name": "immediate"
  },
  {
    "id": 7,
    "ma_one_off_time": 1410017613,
    "ma_randomness_spred": 0,
    "ma_timing_name": "immediate"
  },
  {
    "id": 8,
    "ma_randomness_spred": 12345,
    "ma_timing_name": "startup"
  }
]
}
```

6.2. Channels

An example channel object using the aforementioned timing object is shown below:

```
{
  "channels": [
    {
      "id": 1,
      "ma_channel_credentials": "MIIFeZCCAvsCAQEdQYJ...",
      "ma_channel_interface_name": "eth0",
      "ma_channel_name": "default-collector-channel",
      "ma_channel_target": "collector.example.org"
    },
    {
      "id": 2,
      "ma_channel_credentials": "MIIFeZCCAvsCAQEdQYJ...",
      "ma_channel_interface_name": "eth0",
      "ma_channel_name": "default-controller-channel",
      "ma_channel_target": "controller.example.org"
    }
  ]
}
```

6.3. Configuration

An example config object using the aforementioned channel objects is shown below:

```
{
  "config": [
    {
      "id": 1,
      "ma_agent_id": "550e8400-e29b-41d4-a716-446655440000",
      "ma_channel_name": "default-controller-channel",
      "ma_control_channel_fail_tresh": "10",
      "ma_credentials": "MIIFeZCCAvsCAQEdQYJ...",
      "ma_device_id": "01:23:45:67:89:ab",
      "ma_group_id": "550e8400-e29b-41d4-a716-446655440123",
      "ma_report_ma_id_flag": "1"
    }
  ]
}
```

6.4. Instruction

The instruction object is essentially a wrapper around suppression, schedule, task, channel objects.

6.5. Measurement Supression

An example supression object used by the aforementioned instruction object is shown below:

```
{
  "supression": [
    {
      "id": 1,
      "ma_supression_enabled": 0,
      "ma_supression_end": 0,
      "ma_supression_schedule_names": "icmp-latency-immediate",
      "ma_supression_start": 1410037509,
      "ma_supression_stop_ongoing_task": 0,
      "ma_supression_task_names": "iperf-server"
    }
  ]
}
```

6.6. Measurement Task Configurations

An example task object used by the aforementioned instruction object is shown below:

```
{
  "tasks": [
    {
      "id": 1,
      "ma_role": "",
      "ma_task_cycle_id": "1",
      "ma_task_name": "udp-latency-test",
      "ma_task_options": "",
      "ma_task_registry_entry": "urn:...",
      "ma_task_supress_default": "true"
    },
    {
      "id": 5,
      "ma_role": "",
      "ma_task_cycle_id": "1",
      "ma_task_name": "reporting-daily",
      "ma_task_options": "",
      "ma_task_registry_entry": "urn:...",
      "ma_task_supress_default": "true"
    }
  ]
}
```



```

    },
    {
      "id": 2,
      "ma_role": "",
      "ma_task_cycle_id": "1",
      "ma_task_name": "icmp-latency-test",
      "ma_task_options": "",
      "ma_task_registry_entry": "urn:...",
      "ma_task_suppress_default": "true"
    },
    {
      "id": 3,
      "ma_role": "",
      "ma_task_cycle_id": "1",
      "ma_task_name": "iperf-server",
      "ma_task_options": "{\\"name\\":\\"role\\",
        \\"value\\":\\"server\\"}",
      "ma_task_registry_entry": "server",
      "ma_task_suppress_default": "false"
    },
    {
      "id": 4,
      "ma_role": "",
      "ma_task_cycle_id": "1",
      "ma_task_name": "lmap-reporting-task",
      "ma_task_options": "",
      "ma_task_registry_entry": "lmap-reportd",
      "ma_task_suppress_default": "true"
    }
  ]
}

```

6.7. Measurement Schedules

An example schedule object used by the aforementioned instruction object is shown below:

```
{
  "schedules": [
    {
      "id": 1,
      "ma_sched_channel_interface_select": "0",
      "ma_sched_channel_names": "default-collector-channel",
      "ma_sched_task_downstream_config_names": "reporting-daily",
      "ma_sched_task_output_selection": "1",
      "ma_schedule_name": "reporting-immediate",
      "ma_schedule_task_name": "icmp-latency-test",
      "ma_timing_name": "immediate"
    }
  ]
}
```

6.8. Logging

An example log object is shown below:

```
{
  "logging": [
    {
      "id": 1,
      "ma_log_agent_id": "0e49b32b01falle4bcaf10ddb1bd23b5",
      "ma_log_code": "200",
      "ma_log_description": "OK",
      "ma_log_event_time": 1404313752
    }
  ]
}
```

6.9. Capability and Status

An example status object is shown below:

```
{
  "status": [
    {
      "id": 1,
      "ma_agent_id": "c54c284a01ee11e48dd310ddb1bd23b5",
      "ma_condition_code": "8081",
      "ma_condition_text": "Cond_Text",
      "ma_device_id": "urn:dev:mac:0024beffffe804ff1",
      "ma_firmware": "4560",
      "ma_hardware": "TL-MR3020",
      "ma_interface_dns_server": "8.8.8.8",
      "ma_interface_gateway": "192.168.1.1",
      "ma_interface_ip_address": "192.168.1.10",
      "ma_interface_name": "eth0",
      "ma_interface_speed": "100Mbps",
      "ma_interface_type": "100baseTX",
      "ma_last_config": "140423245",
      "ma_last_instruction": "140431312",
      "ma_last_measurement": "1404315031",
      "ma_last_report": "1404315053",
      "ma_link_layer_addr": "01:23:45:67:89:ab",
      "ma_task_name": "Report",
      "ma_task_registry": "urn:ietf:lmmap:report:http_report",
      "ma_task_role": "Role",
      "ma_version": "Busybox"
    }
  ]
}
```

6.10. Reporting

An example report object is shown below:

```
{
  "reporting": [
    {
      "id": 1,
      "ma_report_agent_id": "550e8400-e29b-41d4-a716-446655440000",
      "ma_report_date": 1404315031,
      "ma_report_group_id": "550e8400-e29b-41d4-a716-446655440123",
      "ma_report_result_conflict_task": "0",
      "ma_report_result_cross_traffic": 20,
      "ma_report_result_end_time": 1404315031,
      "ma_report_result_start_time": 1404315031,
      "ma_report_result_values": "result_values",
      "ma_report_task_column_labels": "\"start-time\"",
        "\"conflicting-tasks\"", "\"cross-traffic\"", "\"mean\"",
    }
  ]
}
```

```

        \ "min\ ", \ "max\ " ",
        "ma_role": " ",
        "ma_task_cycle_id": "1",
        "ma_task_name": "udp-latency-test",
        "ma_task_options": " ",
        "ma_task_registry_entry": "urn:...",
        "ma_task_supress_default": "true"
    },
    {
        "id": 2,
        "ma_report_agent_id": "550e8400-e29b-41d4-a716-446655440000",
        "ma_report_date": 1404315031,
        "ma_report_group_id": "550e8400-e29b-41d4-a716-446655440123",
        "ma_report_result_conflict_task": "0",
        "ma_report_result_cross_traffic": 20,
        "ma_report_result_end_time": 1404315031,
        "ma_report_result_start_time": 1404315031,
        "ma_report_result_values": "result_values",
        "ma_report_task_column_labels": "\ "start-time\ ",
        \ "conflicting-tasks\ ", \ "cross-traffic\ ",
        \ "mean\ ", \ "min\ ", \ "max\ " ",
        "ma_role": " ",
        "ma_task_cycle_id": "1",
        "ma_task_name": "icmp-latency-test",
        "ma_task_options": " ",
        "ma_task_registry_entry": "urn:...",
        "ma_task_supress_default": "true"
    },
    {
        "id": 3,
        "ma_report_agent_id": "550e8400-e29b-41d4-a716-446655440000",
        "ma_report_date": 1404315031,
        "ma_report_group_id": "550e8400-e29b-41d4-a716-446655440123",
        "ma_report_result_conflict_task": "0",
        "ma_report_result_cross_traffic": 20,
        "ma_report_result_end_time": 1404315031,
        "ma_report_result_start_time": 1404315031,
        "ma_report_result_values": "result_values",
        "ma_report_task_column_labels": "\ "start-time\ ",
        \ "conflicting-tasks\ ", \ "cross-traffic\ ",
        \ "mean\ ", \ "min\ ", \ "max\ " ",
        "ma_role": " ",
        "ma_task_cycle_id": "1",
        "ma_task_name": "iperf-server",
        "ma_task_options": "{\\\\"name\\\\" : \\\\"role\\\\" ,
        \\\\"value\\\\" : \\\\"server\\\\" }",
        "ma_task_registry_entry": "server",
        "ma_task_supress_default": "false"
    }

```

```
    },
    {
      "id": 4,
      "ma_report_agent_id": "550e8400-e29b-41d4-a716-446655440000",
      "ma_report_date": 1404315031,
      "ma_report_group_id": "550e8400-e29b-41d4-a716-446655440123",
      "ma_report_result_conflict_task": "0",
      "ma_report_result_cross_traffic": 20,
      "ma_report_result_end_time": 1404315031,
      "ma_report_result_start_time": 1404315031,
      "ma_report_result_values": "result_values",
      "ma_report_task_column_labels": "\"start-time\",
      \"conflicting-tasks\", \"cross-traffic\",
      \"mean\", \"min\", \"max\"",
      "ma_role": "",
      "ma_task_cycle_id": "1",
      "ma_task_name": "lmap-reporting-task",
      "ma_task_options": "",
      "ma_task_registry_entry": "lmap-reportd",
      "ma_task_suppress_default": "true"
    },
    {
      "id": 5,
      "ma_report_agent_id": "550e8400-e29b-41d4-a716-446655440000",
      "ma_report_date": 1404315031,
      "ma_report_group_id": "550e8400-e29b-41d4-a716-446655440123",
      "ma_report_result_conflict_task": "0",
      "ma_report_result_cross_traffic": 20,
      "ma_report_result_end_time": 1404315031,
      "ma_report_result_start_time": 1404315031,
      "ma_report_result_values": "result_values",
      "ma_report_task_column_labels": "\"start-time\",
      \"conflicting-tasks\", \"cross-traffic\",
      \"mean\", \"min\", \"max\"",
      "ma_role": "",
      "ma_task_cycle_id": "1",
      "ma_task_name": "reporting-daily",
      "ma_task_options": "",
      "ma_task_registry_entry": "urn:...",
      "ma_task_suppress_default": "true"
    }
  ]
}
```

7. Security considerations

Large Measurement Platforms may result in a security hazard if they are not properly secured. This is so because they encompass a large number of MAs that can be managed and coordinated easily to generate traffic and they can potentially be used for generating DDoS attacks or other forms of security threats.

From the perspective of the protocols described in this documents, we can identify the following threats:

- o Hijacking: Probably the worst threat is that an attacker takes over the control of one or more MAs. In this case the attacker would be able to instruct the MAs to generate traffic or to eavesdrop traffic in their location. It is then critical that the MA is able to strongly authenticate the Controller. An alternative way to achieve this attack is to alter the communication between the Controller and the MAs. In order to prevent this form of attack, integrity protection of the communication between the Controller and the MAs is required.
- o Polluting: Another type of attack is that an attacker is able to pollute the Collectors database by providing false results. In this case, the attacker would attempt to impersonate one or more MAs and upload fake results in the Collector. In order to prevent this, the authentication of the MAs with the Collector is needed. An alternative way to achieve this is for an attacker to alter the communication between the MA and the Collector. In order to prevent this form of attack, integrity protection of the communication between the MA and the Collector is needed.
- o Disclosure: Another threat is that an attacker may gather information about the MAs and their configuration and the Measurement schedules. In order to do that, it would connect to the Controller and download the information about one or more MAs. This can be prevented by using MA authentication with the Controller. An alternative mean to achieve this would be for the attacker to eavesdrop the communication between the MA and the Controller. In order to prevent this, confidentiality in the communication between the MA and the Controller is required. Similarly, an attacker may wish to obtain measurement result information by eavesdropping the communication between the MA and the Collector. In order to prevent this, confidentiality in the communication between the MA and the Collector is needed.

In order to address all the identified threats, the HTTPS protocol must be used for LMAP (i.e. using HTTP over TLS). HTTPS provides confidentiality, integrity protection and authentication, satisfying all the aforementioned needs. Ideally, mutual authentication should be used. In any case, server side authentication MUST be used. In

order to achieve that, both the Controller and the Collector MUST have certificates. The certificate of the CA used to issue the certificates for the Controller and the Collector MUST be pre configured in the MAs, so they can properly authenticate them. As mentioned earlier, ideally, mutual authentication should be used. However, this implies that certificates for the MAs are needed. Certificate management for a large number of MAs may be expensive and cumbersome. Moreover, the major threats identified are the ones related to hijacking of the MAs, which are prevented by authenticating the Controller. MAs authentication is needed to prevent Polluting and Disclosure threats, which are less severe. So, in this case, alternative (cheaper) methods for authenticating MAs can be considered. The simplest method would be to simply use the MA UUID as a token to retrieve information. Since the MA UUID is 128 bit long, it is hard to guess. It would be also possible to use a password and use the HTTP method for authentication. It is not obvious that managing passwords for a large number of MAs is easier than managing certificates though.

An additional security consideration is posed by the mechanism to push information from the Controller to the MAs. If this method is used, it would be possible its abuse by an attacker to control the MAs. This threat is prevented by the use of HTTPS. If HTTPS is used in the established connection between the MA and the Controller, the only effect that a packet generated by an external attacker to the MA or the Controller would be to reset the HTTPS connection, requiring the connection to be re-established.

It is required in this document that both the Controller and that the Collector are authenticated using digital certificates. The current specification allows for the MA to have information about the certificate of the Certification authority used for generating the Controller and Collector certificates while the actual certificates are exchanged in band using TLS. Another (more secure) option is to perform certificate pinning i.e. to configure in the MAs the actual certificates rather than the certification authority certificate. Another measure to increase the security would be to limit the domains that the FQDNs of the Controller and/or the Collector (e.g. only names in the exmample.org domain).

Large scale measurements can have privacy implications, especially in some scenarios like mobile devices performing measurements. In this memo we have considered using Group IDs to the MA in order to avoid the possibility for the platform to track each individual MA that is feeding results.

8. IANA Considerations

Registration of the well-known URL

9. Acknowledgments

We would like to thank Vlad Victor Ungureanu (Jacobs University Bremen) for providing us external support.

Marcelo Bagnulo, Trevor Burbridge, Sam Crawford, Juergen Schoenwaelder and Vaibhav Bajpai 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.

10. References

10.1. Normative References

- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", RFC 4122, July 2005.
- [RFC7159] Bray, T., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, March 2014.
- [RFC7230] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.
- [RFC7231] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, June 2014.
- [RFC7232] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests", RFC 7232, June 2014.
- [RFC7233] Fielding, R., Lafon, Y., and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Range Requests", RFC 7233, June 2014.
- [RFC7234] Fielding, R., Nottingham, M., and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Caching", RFC 7234, June 2014.
- [RFC7235] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Authentication", RFC 7235, June 2014.

[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-02 (work in progress), August 2014.

10.2. Informative References

[I-D.bagnulo-ippm-new-registry-independent]
Bagnulo, M., Burbridge, T., Crawford, S., Eardley, P., and
A. Morton, "A registry for commonly used metrics.
Independent registries", draft-bagnulo-ippm-new-registry-
independent-01 (work in progress), July 2013.

[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-08 (work in progress), August 2014.

Authors' Addresses

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
British Telecom
Adastral Park, Martlesham Heath
IPswitch
ENGLAND

Email: trevor.burbridge@bt.com

Sam Crawford
SamKnows

Email: sam@samknows.com

Juergen Schoenwaelder
Jacobs University Bremen
Campus Ring 1
28759 Bremen
Germany

Email: j.schoenwaelder@jacobs-university.de

Vaibhav Bajpai
Jacobs University Bremen
Campus Ring 1
28759 Bremen
Germany

Email: v.bajpai@jacobs-university.de

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 24, 2014

T. Burbridge
P. Eardley
British Telecom
M. Bagnulo
Universidad Carlos III de Madrid
J. Schoenwaelder
Jacobs University
October 21, 2013

Information Model for Large-Scale Measurement Platforms (LMAP)
draft-burbridge-lmap-information-model-01

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 MA 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 MA 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

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 April 19, 2014.

Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. LMAP Information Model	3
2.1. Information Structure	4
2.2. Pre-Configuration Information	5
2.3. Configuration Information	5
2.4. Instruction Information	6
2.5. Logging Information	9
2.6. Status Information	10
2.7. Reporting Information	11
2.8. Channels	12
2.9. Timing Information	13
2.9.1. Periodic Timing	14
2.9.2. Calendar Timing	14
2.9.3. One-Off Timing	15
2.9.4. Immediate Timing	15
2.9.5. Timing Randomness	15
3. IANA Considerations	15
4. Security Considerations	16
5. Acknowledgements	16
6. Informative References	16
Authors' Addresses	16

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. The main components of a large-scale measurement platform are the Measurement Agents (hereafter MAs), the Controllers and the Collectors.

The MAs are the elements actually performing the measurements. The MAs are controlled by one or more Controllers 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 MAs to perform a set of measurements at a certain point in time. The MAs execute the instructions from the Controller and once they have done so they report the results of the measurements to the Collector. The overall framework for a Large Measurement platform and the terminology used in this document is described in detail in [I-D.ietf-lmap-framework].

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

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

As defined in [RFC3444], the LMAP IM defines the concepts involved in a large-scale measurement platform at a high level of abstraction, independently 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 device (e.g. home gateway, smartphone, PC, router etc.).

2. LMAP Information Model

2.1. Information Structure

The information described herein relates to the information stored, received or transmitted by a Measurement Agent as described within the LMAP framework [I-D.ietf-lmap-framework]. As such, some subsets of this information model are applicable to the measurement Controller, Collector and systems that pre-configure the Measurement Agent. The information described in these models will be transmitted across the protocols and interfaces between the Measurement Agent and such systems according to a Data Model.

For clarity the information model is divided into six sections:

1. Pre-Configuration Information. Information pre-configured on the Measurement Agent prior to any communication with other components of the LMAP architecture, specifically detailing how to register with a Controller
2. Configuration Information. Information delivered to the MA on registration with a Controller or updated during a later communication, in particular detailing how to retrieve measurement and reporting instruction information from a Controller along with information specifically about the MA
3. Instruction Information. Information that is received by the MA from the Controller pertaining to the measurement and reporting configuration. This includes measurement configuration, report channel configuration, measurement schedules and measurement suppression information
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. 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 the Collector 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 reported by the MA as optional fields in the reporting communication

to the Collector.

2.2. 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.

This pre-configuration information needs to include an URL of the Controller where configuration information can be retrieved 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) as well as the timing for that communication. All this is expressed as the Configuration Channel. In addition to the Configuration Channel information, the MA's security information is configured which can be either a certificate and a private key or a password, depending on the security solution used.

Detail of the information model elements:

1. MA MAC: MAC Address
2. Configuration Channel: Channel
3. MA Certificate: Certificate (optional)
4. MA ID: random UUID (optional)
5. MA password: string (optional)

The detail of the Channel object is described later since it is common to several parts of the information model.

2.3. Configuration Information

During registration or at any later point at which the MA contacts the Controller, the choice of Controller and details for the timing of communication with the Controller can be changed. For example the pre-configured Controller may be replaced with a specific Controller that is more appropriate to the MA device type, location of characteristics of the network (e.g. access technology type or broadband product). The initial communication timing object may also be replaced with one more relevant to routine communications between the MA and the Controller.

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. 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. Where the Measurement Group ID is set an additional flag (the Report MA ID flag) is required to control whether the Measurement Agent ID is to be reported. This allows the MA to remain anonymous which may be particularly useful to prevent tracking of mobile MA devices.

The configuration information will also contain information about different communication channels that the MA will have with different elements of the infrastructure. Each channel specifies a URL, security information and timing information for the communication.

Detail of the additional information model elements:

1. Measurement Agent ID: UUID
2. Measurement Group ID (optional): String
3. Report MA ID flag (optional): Boolean
4. Instruction Channel: Channel (DISCUSSION: shouldn't we split this into 4 different channels i.e. the Measurement Task Configuration channel, the Report Channel channel, the Measurement Schedules channel and the Measurement Suppression channel?)
5. Status Channel: Channel
6. Logging Channel: Channel

2.4. Instruction Information

The Instruction information model has four sub-elements:

1. Measurement Task Configurations: Set
2. Report Channels: Set
3. Measurement Schedules: Set
4. Measurement Suppression: Object

Conceptually each Measurement Task Configuration defines the parameters of a Measurement 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 Measurement Schedule).

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 Measurement Schedule element), so it does not by itself instruct the MA to actually perform this measurement task.

The Measurement Task Configuration will include a local short name for reference by the Measurement Schedule, along with a registry entry [I-D.bagnulo-ippm-new-registry] that defines the Measurement Task. The MA itself will resolve the registry entry to a local executable program. In addition the Measurement Task is specialised through a set of configuration Options. The nature and number of these Options will depend upon the Measurement Task and will be defined in the Measurement Task Registry. In addition the Measurement Task Configuration may optionally also be given 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 is manually incremented when an Option change is implemented which could mean that two sets of results should not be directly compared.

A Report Channel defines how to report results to a single Collector. Several Report Channels can be defined to enable results to be split or duplicated across different report intervals or destinations. E.g. a single Collector may have three Report Channels, one reporting hourly, another reporting daily and a third on which to send immediate results for on-demand measurement tasks. The details of the Channel element is described later as it is common to several objects.

A Measurement Schedule contains the instruction from the Controller to the MA to execute a single or repeated series of Measurement Tasks. Each Measurement Schedule contains basically three elements: a reference to a list of Measurement Task Configuration, a reference to a set of one or more Report Channels, and a timing object for the schedule. The schedule basically states what measurement task to run, how to report the results, and when to run the measurement task. Multiple measurement tasks in the list will be executed in order with

minimal gaps. Note that the Controller can instruct the MA to report to several Collectors by specifying several Report Channels.

Example: a Measurement Schedule references a single Measurement Task Configuration for the UDP latency defined in the previous example. It references the Report Channel in the previous example to send results immediately as available to the specified Collector. The timing is specified to run the configured Measurement Task Configuration every hour at 23 minutes past the hour.

Measurement Suppression information is used to over-ride the Measurement Schedule and stop measurements from 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.

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 Measurement Tasks Configurations will be relatively static. The Measurement 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 Measurement Suppression command without removing the existing Measurement Schedules that would continue to apply after the Measurement 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 Measurement Task Configurations and Reporting Channels to help ensure consistency and reduce errors.

Definition of the information model elements:

1. Measurement Task Configurations: Set

1. Measurement Task Configuration: Object

1. Task Name (used for referral from the Measurement Schedules): String
2. Registry Entry: URN
3. Options: Set (optional)
 1. Interface name (reference by name to one of the Interfaces defined in the Status information): String

4. Measurement Cycle ID: String (optional)
 2. Report Channels: Set
 1. Report Channel: Channel
 3. Measurement Schedules: Set
 1. Measurement Schedule: Object
 1. Schedule Name: String
 2. Measurement Task Configuration Names (reference by Name to one of the measurement tasks defined in the Measurement Task Configuration set): List
 1. Task Name: String
 3. Report Channel Names (reference by Name to one of the measurement tasks defined in the Measurement Task Configuration set): Set
 1. Channel Name: String
 4. Measurement Timing: Timing
 4. Measurement Suppression: Object (optional)
 1. Start: datetime
 2. End: datetime
 3. Set of Measurement Task Configuration Names (optional - default all)
 1. Task Name: String
- 2.5. Logging Information
- The MA will report back success/failure and status information to the Controller. These messages will fall into a number of different categories:
1. Success/failure messages in response to information updates from the Controller. For example:
 - * "Report Channel 'hourly db' configured"

- * "Measurement Schedule does not conform to schema, Row 211"
- 2. Status updates from the operation of the MA. For example:
 - * "out of memory: cannot record result"
 - * "Collector 'collector.example.com' not responding"

Each log message will have the following Information model elements:

1. Log Time: datetime
2. Log Event: Object

2.6. Status Information

In addition to the information reported by the MA through the logging information, the MA will hold further status information that can be retrieved by a Controller. One category of additional information that has not been defined in earlier sections is the availability of Measurement Tasks on that MA.

MA Status information model elements:

1. MA ID: String
2. MA Device: String
3. MA hardware: String (optional)
4. MA firmware: String (optional)
5. MA software: String (optional)
6. MA Interfaces: set
 1. If name: String
 2. If type: String (one of eth, wlan, TBC)
 3. If speed: Integer (expressed in Mbps)
 4. Link Layer Address: String
 5. IP address: Set
 1. Protocol: String (one of v4, v6)

- 2. Address: String
- 6. Gateway: Set (optional)
 - 1. Protocol: String (one of v4, v6)
 - 2. Address: String
- 7. DNS server: Set (optional)
 - 1. Protocol: String (one of v4, v6)
 - 2. Address: String
- 7. Last Measurement: datetime
- 8. Last Report: datetime
- 9. Last Instruction: datetime
- 10. Last Configuration: datetime
- 11. Supported Measurements: Set
 - 1. Registry Entry: URN
 - 2. Version: String (optional)

2.7. Reporting Information

At a point in time specific by the Report Channel, the MA will communicate a set of measurement results to the Collector. These measurement results should be communicated within the context in which they were collected.

The report is structured hierarchically to avoid repetition of report, Measurement Agent 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). In addition optional further MA context information can be included at this point such as the line sync speed or ISP and product if known by the MA.

After the MA information the results are reported grouped into the different Measurement Tasks. Each Task starts with replicating the Measurement Task Configuration information before the result headers (titles for data columns) and the result data rows.

Information model elements:

1. Report Date: datetime
2. Measurement Agent ID: String (optional)
3. Measurement Group ID: String (optional)
4. MA Context: Set (optional)
 1. Context Item: Object
5. Measurement Task: Set
 1. Measurement Task Configuration: Object
 2. Result Headers: List
 1. Column Name: String
 3. Result Data: List
 1. Result Row: Object
 1. Measurement Time: datetime
 2. Cross-traffic: Integer (optional)
 3. Result Columns: List
 1. Column Data

2.8. Channels

A Channel defines a communication channel between the MA and other element of the measurement framework i.e. with the Collector to report results back, to Controller to retrieve Instructions or other information exchanged between the parties. Several Channels can be defined to enable results to be split or duplicated across different report intervals or destinations. E.g. a single Collector may have three Report Channels, one reporting hourly, another reporting daily and a third on which to send immediate results for on-demand measurement tasks.

Each Channel contains the details of the target (including location and security information such as the certificate), and the timing for the communication i.e. when to establish the communication. 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 of the target URL (which will be retrieved later on using a communication protocol such as SSL). The Channel can use the same timing information object as a Measurement Schedule and the Controller Communication Timing defined earlier. There are several options, such as immediately after the results are obtained or at a given interval or calendar based cycle). As with the Measurement task Configuration, each Channel is also given a local short name by which it can be referenced from a Measurement Schedule or other elements.

Example: A Channel using for reporting results may specify that results are to be sent to the URL (<https://collector.foo.org/report/>), using the appropriate digital certificate to establish a secure channel. The Channel specifies that the results are to be sent immediately as available and not batched.

Channel: Object

1. Channel Name (used for referral from other objects): String
2. Target: URL
3. Certificate: X.509 Certificate
4. Communication Timing: Timing

2.9. Timing Information

The Timing information object used throughout the information models can take one of four different forms:

1. Periodic. Specifies a start, end and interval time in milliseconds
2. Calendar: Specifies a calendar based pattern - e.g. 22 minutes past each hour of the day on weekdays
3. One Off: A single instance occurring at a specific time
4. Immediate: Should occur as soon as possible

Optionally each of the first three options may also specify a randomness that should be evaluated and applied separately to each indicated event.

2.9.1. Periodic Timing

Information model elements:

1. 1. Timing Name: String
2. 2. Start: datetime (optional)
3. 3. End: datetime (optional)
4. 4. Interval: Integer (in milliseconds)
5. 5. Randomness: Timing Randomness (optional)

2.9.2. Calendar Timing

Information model elements:

1. Timing Name: String
2. Start: datetime (optional)
3. End: datetime (optional)
4. Months: Set (optional - default [1-12])
 1. Month: Integer
5. Weekdays: Set (optional - default [Mon-Sun])
 1. Weekday: String (one off Mon, Tue, Wed, Thu, Fri, Sat Sun)
6. Days: Set (optional - default [1-31])
 1. Day: Integer
7. Hours: Set (optional - default [1-24])
 1. Hour: Integer
8. Minutes: Set (optional - default [1-60])
 1. Minute: Integer
9. Seconds: Set (optional - default [1-60])
 1. Second: Integer

10. Randomness: Timing Randomness (optional)

2.9.3. One-Off Timing

Information model elements:

1. Time: datetime
2. Randomness: Timing Randomness (optional)

2.9.4. Immediate Timing

The immediate timing object has no further information elements. The measurement or report is simply to be done as soon as possible.

2.9.5. Timing Randomness

The Timing randomness object specifies a random distribution that can be applied to any scheduled execution event such as a measurement or report. The intention is to be able to spread the load on the Controller, Collector and network in an automated manner for a large number of Measurement Agents. The randomness is expressed as a distribution (e.g. Poison, Normal, Uniform etc.) along with the spread over which the distribution should be applied. In addition optional upper and lower bounds can be applied to control extreme spread of timings.

Information model elements:

1. Distribution: String
2. Upper Cut: Integer (optional)
3. Lower Cut: Integer (optional)
4. Spread: Integer

3. IANA Considerations

This document makes no request of IANA .

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

4. 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. The second consideration is that no mandated information items 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 would, 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.

5. Acknowledgements

6. Informative References

- [I-D.bagnulo-ippm-new-registry]
Bagnulo, M., Burbridge, T., Crawford, S., Eardley, P., and A. Morton, "A registry for commonly used metrics", draft-bagnulo-ippm-new-registry-00 (work in progress), January 2013.
- [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-00 (work in progress), October 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3444] Pras, A. and J. Schoenwaelder, "On the Difference between Information Models and Data Models", RFC 3444, January 2003.

Authors' Addresses

Trevor Burbridge
British Telecom
Adastral Park, Martlesham Heath
Ipswich, IP5 3RE
UK

Phone:
Fax:
Email:
URI:

Philip Eardley
British Telecom
Adastral Park, Martlesham Heath
Ipswich, IP5 3RE
UK

Phone:
Fax:
Email:
URI:

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

Phone:
Fax:
Email:
URI:

Juergen Schoenwaelder
Jacobs University

Phone:
Fax:
Email:
URI:

IPPM Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: April 07, 2014

B. Claise
A. Akhter
Cisco Systems, Inc.
October 04, 2013

Performance Metrics Registry
draft-claise-ippm-perf-metric-registry-01.txt

Abstract

This document specifies an IANA registry for Performance Metrics, for both active monitoring and passive monitoring, along with the initial content. This document also gives a set of guidelines for Performance Metrics requesters and reviewers.

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 April 07, 2014.

Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Open Issues	2
2. Introduction	2
2.1. Terminology	4
3. Guidelines for Performance Metric Requesters and Reviewers .	5
3.1. Performance Metrics Template	5
3.2. Other Guidelines	6
4. Initial Set of Performance Metrics	6
4.1. Existing Performetrics Metrics, Based on the RFC6390 Template	6
4.2. Mapping Some IPPM Performance Metrics in the Registry . .	8
4.2.1. IPPM Performance Metric Mapping Experiment	9
4.2.2. Which IPPM Performance Metrics?	12
5. Performance Metrics in the IPFIX Registry	12
6. Security Considerations	13
7. IANA Considerations	13
8. Acknowledgments	13
9. References	13
9.1. Normative References	13
9.2. Informative References	14
Authors' Addresses	16

1. Open Issues

1. Check whether the "Initial Set of Performance Metrics" is up to date with the latest Performance Metrics published in XRBLOCK.
2. Do we want to organize the Performance Metrics list into different layers? IP, transport layer stats, application stats, etc?
3. "IPPM Performance Metric Mapping Experiment" for IPDV must be validated.
4. The community will have to agree on which Performance Metrics (along with the specific values of the measurements parameters) are operationally relevant
5. Define "Measurement Parameter"

2. Introduction

The IETF specifies and uses Performance Metrics of protocols and applications transported over its protocols. Performance metrics are such an important part of the operations of IETF protocols that [RFC6390] specifies guidelines for their development.

The definition and use of performance metrics in the IETF happens in various working groups (WG), most notably:

The "IP Performance Metrics" (IPPM) WG [IPPM] is the WG primarily focusing on Performance Metrics definition at the IETF.

The "Metric Blocks for use with RTP's Extended Report Framework" WG [XRBLOCK] 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" WG [BMWG] proposed some Performance Metrics part of the benchmarking methodology.

The "IP Flow Information eXport" (IPFIX) WG [IPFIX] Information elements related to performance metrics are currently proposed.

The "Performance Metrics for Other Layers" (PMOL) concluded WG [PMOL], defined some Performance Metrics related to Session Initiation Protocol (SIP) voice quality [RFC6035].

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

However, despite the abundance and importance of performance metrics, there are still some problems for the industry: first, how to discover which Performance Metrics have already specified, and second, how to avoid Performance Metrics redefinition. Only someone with a broad IETF knowledge would be able to find its way among all the different Performance Metrics specified in the different WGs. The way in which IETF manages namespaces is with IANA registries, and there is currently no Performance Metrics Registry in IANA.

This document specifies an IANA registry for Performance Metrics, along with the initial content, taken from the Performance Metrics already specified at the IETF. Firstly, from the Performance Metrics already specified by the RFC630 template (mentioned later on in the document), and secondly from the existing set of IPPM Performance Metrics. This second category requires a mapping to the RFC6390 template. This Performance Metric Registry is applicable to Performance Metrics issued from active monitoring, passive monitoring, or from the end point calculation. Therefore, it must be relevant to work developed in the following WGs: IPPM, LMAP, XRBLOCK, BMWG, and IPFIX. Finally, this document gives a set of guidelines for Performance Metrics requesters and reviewers.

Based on [RFC5226] Section 4.3, this document is processed as Best Current Practice (BCP) [RFC2026].

The IPPM Metrics Registry [RFC4148] was an attempt to create such a Performance Metrics registry. However, that registry was reclassified as obsolete with [RFC6248], "RFC 4148 and the IP Performance Metrics (IPPM) Registry of Metrics Are Obsolete", and consequently withdrawn.

A couple of interesting quotes from RFC 6248 might help understand the issues related to that registry.

1. "It is not believed to be feasible or even useful to register every possible combination of Type P, metric parameters, and Stream parameters using the current structure of the IPPM Metrics Registry."
2. "The registry structure has been found to be insufficiently detailed to uniquely identify IPPM metrics."
3. "Despite apparent efforts to find current or even future users, no one responded to the call for interest in the RFC 4148 registry during the second half of 2010."

2.1. Terminology

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

The terms Performance Metric and Performance Metrics Directorate are direct quotes from [RFC6390], and copied over in this document for the readers convenience.

Performance Metric: A Performance Metric is a quantitative measure of performance, specific to an IETF-specified protocol or specific to an application transported over an IETF-specified protocol. Examples of Performance Metrics are the FTP response time for a complete file download, the DNS response time to resolve the IP address, a database logging time, etc.

Performance Metrics Directorate: The Performance Metrics Directorate is a directorate that provides guidance for Performance Metrics development in the IETF. The Performance Metrics Directorate should be composed of experts in the performance community, potentially selected from the IP Performance Metrics (IPPM), Benchmarking Methodology (BMWG), and Performance Metrics for Other Layers (PMOL) WGs.

Performance Metrics Registry: The IANA registry containing the Performance Metrics. This registry is initially populated from this document.

Measurement Parameter: NOT SURE HOW TO DEFINE THIS

3. Guidelines for Performance Metric Requesters and Reviewers

3.1. Performance Metrics Template

"Guidelines for Considering New Performance Metric Development", [RFC6390] defines a framework and a process for developing Performance Metrics for protocols above and below the IP layer (such as IP-based applications that operate over reliable or datagram transport protocols). These metrics can be used to characterize traffic on live networks and services. As such, RFC 6390 does not define any Performance Metrics.

RFC 6390 scope covers guidelines for the Performance Metrics directorate members for considering new Performance Metrics and suggests how the Performance Metrics Directorate will interact with the rest of the IETF. Its mission is mentioned at [performance-metrics-directorate]. In practice, a weekly cron job discovers all the IETF drafts that refers to RFC 6390, or that contains the keyword "performance metric". Once discovered, the different drafts are assigned a Performance Metric Directorate reviewer. One of the primary task is to ensure that the RFC 6390 template is correctly applied, making sure that the Performance Metric semantic is correctly specified.

RFC 6390, specified in Section 5.4, proposes a template for Performance Metrics specifications:

Normative

- o Metric Name
- o Metric Description
- o Method of Measurement or Calculation

- o Units of Measurement
- o Measurement Point(s) with potential Measurement Domain
- o Measurement Timing

Informative

- o Implementation
- o Verification
- o Use and Applications
- o Reporting Model

The template specified in Section 5.4 of "Guidelines for Considering New Performance Metric Development", [RFC6390] MUST be used as a basis for the Performance Metrics Registry Definition.

3.2. Other Guidelines

RFC 6390 lacks a naming convention for Performance Metrics, but specifies that "Performance Metric names are RECOMMENDED to be unique within the set of metrics being defined for the protocol layer and context.". Imposing an unique Performance Metric name, while ideal, is not practicable in real live. Indeed, some metrics have already been specified, and the name clashes appeared already. Therefore, all Performance Metrics specified in the registry MUST have an unique performance metric Id. Regarding naming convention, the Performance Metric Names SHOULD be meaningful and easily searchable in the registry.

The group of experts (the reviewers) MUST check the requested Performance Metric for completeness, accuracy of the template description, and for correct naming according to [RFC6390].

Requests for Performance Metric that duplicate the functionality of existing Performance Metrics SHOULD be declined.

4. Initial Set of Performance Metrics

4.1. Existing Performance Metrics, Based on the RFC6390 Template

This section contains a list of Performance Metrics specified according to [RFC6390], either in RFCs, or IETF drafts currently in the RFC editor queue. This list should serve as initial content for the Performance Metrics Registry.

Performance Metric Id	Performance Metric	Reference
1	Threshold in RTP	[RFC6958], appendix A
2	Sum of Burst Durations in RTP	[RFC6958], appendix A
3	RTP Packets lost in bursts	[RFC6958], appendix A
4	Total RTP packets expected in bursts	[RFC6958], appendix A
5	Number of bursts in RTP	[RFC6958], appendix A
6	Sum of Squares of Burst Durations in RTP	[RFC6958], appendix A
7	Number of RTP packets discarded Metric	[RFC7002], appendix A
8	Threshold in RTP	[RFC7003], appendix A
9	RTP Packets discarded in bursts	[RFC7003], appendix A
10	Total RTP packets expected in bursts	[RFC7003], appendix A
11	RTP Burst Loss Rate	[RFC7004], appendix A
12	RTP Gap Loss Rate	[RFC7004], appendix A
13	RTP Burst Duration Mean	[RFC7004], appendix A
14	RTP Burst duration variance	[RFC7004], appendix A
15	RTP Burst Discard Rate	[RFC7004], appendix A
16	RTP Gap Discard Rate	[RFC7004], appendix A
17	Number of discarded frames in RTP	[RFC7004], appendix A
18	Number of duplicate frames in RTP	[RFC7004], appendix A
19	Number of full lost frames in RTP	[RFC7004], appendix A
20	Number of partial lost frames in RTP	[RFC7004], appendix A
21	De-jitter buffer nominal delay in RTP	[RFC7005], appendix A
22	De-jitter buffer maximum delay in RTP	[RFC7005], appendix A

23	De-jitter buffer high water mark in RTP	[RFC7005], appendix A
24	De-jitter buffer low water mark in RTP:	[RFC7005], appendix A

Table 1: List of Existing Performance Metrics Specified at the IETF

4.2. Mapping Some IPPM Performance Metrics in the Registry

The IPPM WG [IPPM] specified some Measurement Parameters (or measurement characteristics), for example Type-P [RFC2330], packet distribution, etc.

The IPPM WG also specified Performance Metrics. For example:

A One-way Delay Metric for IPPM [RFC2679]

A One-way Packet Loss Metric for IPPM [RFC2680]

A Round-trip Delay Metric for IPPM [RFC2681]

Those Performance Metrics are based on specific values for the Measurement Parameters. For example: the mean packet loss at IP layer, based on a periodic packet distribution, represented with percentile 95th.

The Performance Metrics Registry should contain the IPPM-specified Performance Metrics that are operationally relevant, as opposed to all Performance Metrics, resulting of all the potential combination of Measurement Parameters.

In a typical Large-Scale Measurement of Broadband Performance (LMAP) environment, some information can complement the test to be run:

Measurement Parameters configured part of the test definition

run-time parameters observed during the test

If a test definition requests the round-trip delay metric to a DNS server to be metered "now", the DNS server is a Measurement Parameter configured part of the test definition. Some run-time parameters observed during the test complement the test report: the IP address of the DNS server, the measurement start time, the measurement end time, the DSCP, the TTL, etc.

Those run-time parameters are not part of the Performance Metric definition, while the specific values for the Measurement Parameters are part of it.

4.2.1. IPPM Performance Metric Mapping Experiment

This section is an illustration on how the IP Packet Delay Variation (IPDV) Performance Metric [RFC3393] maps to the RFC 6390 template. Note that the delay variation is sometimes called "jitter", as mentioned in the section 1.1 of [RFC3393], and in section 1 of [RFC5481].

Normative Reference

Performance Metric Element ID

TBD1: The next available Performance Metric Element ID in the Performance Metric Registry.

Metric Name

Packet Delay Variation for UDP Packet with Periodic Distribution reported as 95th percentile

Metric Description

The difference between the one-way-delay of the selected packets, reported as the positive 95th percentile.

The default measurement parameters are

- o L, a packet length in bits, in case of active probing. L = 200 bits.

- o 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.

- o Inter packets time of 20 msec

o etc. (I have not reviewed all the parameters of [RFC3393])

If any of those measurement parameters is not the default value, its value must be stored with the performance metric value, as context information. THIS IS UP TO DISCUSSION.

Method of Measurement or Calculation

As documented in Section 4.1 of [RFC5481]: If we have packets in a stream consecutively numbered $i = 1, 2, 3, \dots$ falling within the test interval, then $IPDV(i) = D(i) - D(i-1)$ where $D(i)$ denotes the one-way delay of the i th packet of a stream.

One-way delays are the difference between timestamps applied at the ends of the path, or the receiver time minus the transmission time.

So $D(2) = R2 - T2$. With this timestamp notation, it can be shown that IPDV also represents the change in inter-packet spacing between transmission and reception:

$$IPDV(2) = D(2) - D(1) = (R2 - T2) - (R1 - T1) = (R2 - R1) - (T2 - T1)$$

Units of Measurement

As documented in Section 8.3 of [RFC5481]: With IPDV, it is interesting to report on a positive percentile, and an inter-quantile range is appropriate to reflect both positive and negative tails (e.g., 5% to 95%). If the IPDV distribution is symmetric around a mean of zero, then it is sufficient to report on the positive side of the distribution.

The unit of measurement is percentile 95th.

Measurement Point(s) with potential Measurement Domain

As documented in Section 4.1 of [RFC5481]: Both IPDV and PDV are derived from the one-way-delay metric. One-way delay requires knowledge of time at two points, e.g., the source

and destination of an IP network path in end-to-end measurement. Therefore, both IPDV and PDV can be categorized as 2-point metrics because they are derived from one-way delay. Specific methods of measurement may make assumptions or have a priori knowledge about one of the measurement points, but the metric definitions themselves are based on information collected at two measurement points.

Measurement Timing

As documented in Section 4.1 of [RFC5481]: The mean of all IPDV(i) for a stream is usually zero. However, a slow delay change over the life of the stream, or a frequency error between the measurement system clocks, can result in a non-zero mean.

See also <http://tools.ietf.org/html/rfc5481#section-6.3> for "clock stability and error" considerations.

See also <http://tools.ietf.org/html/rfc5481#section-8.5> for "clock Sync Options" considerations.

Informative Reference

Implementation

As documented in Section 4.1 of [RFC5481]: Note that IPDV can take on positive and negative values (and zero). One way to analyze the IPDV results is to concentrate on the positive excursions. However, this approach has limitations that are discussed in more detail below (see Section 5.3 of [RFC5481]).

Verification

Not Applicable

Use and Applications

See section 7 " Applicability of the Delay Variation Forms and Recommendations" of [RFC5481]:

Reporting Model

As mentioned previously: If any of those measurement parameters is not the default, its value must be stored with the performance metric value, as context information.

4.2.2. Which IPPM Performance Metrics?

Not all possible combinations of Measurement Parameters for all IPPM Performance Metrics will populate the Performance Metrics Registry. The criteria for selecting the Performance Metrics are (based on discussion with Brian Trammell):

- (1) interpretable by the user
- (2) implementable by the software designer
- (3) deployable by network operators, without major impact on the networks
- (4) accurate, for interoperability and deployment across vendors

Which IPPM Performance Metrics will be selected for the Performance Registry is out of the scope of this document, for now. What is envisioned is a RFC similar to "Basic Requirements for IPv6 Customer Edge Routers", [RFC6204], but for Performance Metrics: "Basic Performance Metrics Requirements for IP Packet SLA Monitoring with Active Probing", or something similar. This document would explain the list of Performance Metrics (from the Performance Metrics Registry, so with fixed Measurement Parameters), along with some proposed run time parameters, depending on the deployment scenario.

5. Performance Metrics in the IPFIX Registry

There are multiple proposals to add performance metrics Information Elements in the IPFIX IANA registry [iana-ipfix-assignments], to be used with the IPFIX protocol [RFC7011]. This is perfectly legal according the "Information Model for IPFIX" [RFC7012] and "Guidelines for Authors and Reviewers of IPFIX Information Elements" [RFC7013].

Simply adding some text in the Information Element Description field might be a solution if this description is compliant with the RFC6390 template definition. However, this is not an ideal solution. On the

top of having potentially long descriptions, this imposes a specific formatting for the description field of the performance metrics-related Information Elements, while none is imposed for the non performance metrics-related ones.

The preferred approach is for the Performance Metrics to be self-described in their own registry. When the Performance Metrics needs to be defined in the IPFIX IANA registry, the new Information Element can simply refer to the specific entry in the Performance Metrics registry.

6. Security Considerations

This draft doesn't introduce any security considerations. However, the definition of Performance Metrics may introduce some security concerns, and should be reviewed with security in mind.

7. IANA Considerations

This document refers to an initial set of Performance Metrics. The list of these Information Elements is given in the "Initial Set of Performance Metrics" Section. The Internet Assigned Numbers Authority (IANA) has created a new registry for Performance Metrics called "Performance Metrics", and filled it with the initial list in Section 4.

New assignments for Performance Metric will be administered by IANA through Expert Review [RFC5226], i.e., review by one of a group of experts appointed by the IESG upon recommendation of the Ops Area Directors. The experts will initially be drawn from the Working Group Chairs and document editors of the Performance Metrics directorate [performance-metrics-directorate].

8. Acknowledgments

Thanks to Carlos Pignataro for improving the text of version 00.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, March 2009.
- [RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, October 2011.
- [RFC6958] Clark, A., Zhang, S., Zhao, J., and Q. Wu, "RTP Control Protocol (RTCP) Extended Report (XR) Block for Burst/Gap Loss Metric Reporting", RFC 6958, May 2013.
- [RFC7002] Clark, A., Zorn, G., and Q. Wu, "RTP Control Protocol (RTCP) Extended Report (XR) Block for Discard Count Metric Reporting", RFC 7002, September 2013.
- [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.
- [RFC7004] Zorn, G., Schott, R., Wu, Q., and R. Huang, "RTP Control Protocol (RTCP) Extended Report (XR) Blocks for Summary Statistics Metrics Reporting", RFC 7004, September 2013.
- [RFC7005] Clark, A., Singh, V., and Q. Wu, "RTP Control Protocol (RTCP) Extended Report (XR) Block for De-Jitter Buffer Metric Reporting", RFC 7005, September 2013.

9.2. Informative References

- [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.
- [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.
- [RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", BCP 108, RFC 4148, August 2005.
- [RFC6035] Pendleton, A., Clark, A., Johnston, A., and H. Sinnreich, "Session Initiation Protocol Event Package for Voice Quality Reporting", RFC 6035, November 2010.
- [RFC6204] Singh, H., Beebe, W., Donley, C., Stark, B., and O. Troan, "Basic Requirements for IPv6 Customer Edge Routers", RFC 6204, April 2011.
- [RFC6248] Morton, A., "RFC 4148 and the IP Performance Metrics (IPPM) Registry of Metrics Are Obsolete", RFC 6248, April 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.
- [RFC7012] Claise, B. and B. Trammell, "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, September 2013.
- [RFC7013] Trammell, B. and B. Claise, "Guidelines for Authors and Reviewers of IP Flow Information Export (IPFIX) Information Elements", BCP 184, RFC 7013, September 2013.
- [iana-ipfix-assignments]
Internet Assigned Numbers Authority, ., "IP Flow Information Export Information Elements (<http://www.iana.org/assignments/ipfix/>)", .
- [performance-metrics-directorate]
IETF, ., "Performance Metrics Directorate (<http://www.ietf.org/iesg/directorate/performance-metrics.html>)", .

- [BMWG] IETF, ., "Benchmarking Methodology (BMWG) Working Group,
<http://datatracker.ietf.org/wg/bmwg/charter/>", .
- [IPFIX] IETF, ., "IP Flow Information eXport (IPFIX) Working
Group, <http://datatracker.ietf.org/wg/ipfix/charter/>", .
- [IPPM] IETF, ., "IP Performance Metrics (IPPM) Working Group,
<http://datatracker.ietf.org/wg/ippm/charter/>", .
- [PMOL] IETF, ., "IPerformance Metrics for Other Layers (PMOL)
Working Group,
<http://datatracker.ietf.org/wg/pmol/charter/>", .
- [XRBLOCK] IETF, ., "Metric Blocks for use with RTCP's Extended
Report Framework (XRBLOCK),
<http://datatracker.ietf.org/wg/xrblock/charter/>", .

Authors' Addresses

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

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

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

Email: aakhter@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.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
2. Outline of an LMAP-based measurement system	5
3. Terminology	9
4. Constraints	12
4.1. The measurement system is under the direction of a single organisation	13
4.2. Each MA may only have a single Controller at any point in time	13
5. Protocol Model	13
5.1. Bootstrapping process	14
5.2. Control Protocol	15
5.2.1. Configuration	15
5.2.2. Instruction	16
5.2.3. Capabilities, Failure and Logging Information	20
5.3. Operation of Measurement Tasks	22
5.3.1. Starting and Stopping Measurement Tasks	22
5.3.2. Overlapping Measurement Tasks	23
5.4. Report Protocol	24
5.4.1. Reporting of Subscriber's service parameters	25
5.5. Operation of LMAP over the underlying packet transfer mechanism	26
5.6. Items beyond the scope of the initial LMAP work	27
5.6.1. End-user-controlled measurement system	28
6. Deployment considerations	28
6.1. Controller and the measurement system	28
6.2. Measurement Agent	29
6.2.1. Measurement Agent on a networked device	30
6.2.2. Measurement Agent embedded in site gateway	30
6.2.3. Measurement Agent embedded behind site NAT /firewall	30
6.2.4. Multi-homed Measurement Agent	31
6.2.5. Measurement Agent embedded in ISP network	31

6.3.	Measurement Peer	32
6.4.	Deployment examples	32
7.	Security considerations	35
8.	Privacy considerations	37
8.1.	Categories of entities with information of interest . . .	38
8.2.	Examples of sensitive information	38
8.3.	Different privacy issues raised by different sorts of Measurement Methods	39
8.4.	Privacy analysis of the communication models	40
8.4.1.	MA Bootstrapping	40
8.4.2.	Controller <-> Measurement Agent	41
8.4.3.	Collector <-> Measurement Agent	42
8.4.4.	Measurement Peer <-> Measurement Agent	42
8.4.5.	Measurement Agent	44
8.4.6.	Storage and reporting of Measurement Results	45
8.5.	Threats	45
8.5.1.	Surveillance	45
8.5.2.	Stored data compromise	45
8.5.3.	Correlation and identification	46
8.5.4.	Secondary use and disclosure	46
8.6.	Mitigations	47
8.6.1.	Data minimisation	47
8.6.2.	Anonymity	48
8.6.3.	Pseudonymity	49
8.6.4.	Other mitigations	49
9.	IANA considerations	50
10.	Acknowledgments	50
11.	History	51
11.1.	From -00 to -01	51
11.2.	From -01 to -02	51
11.3.	From -02 to -03	52
11.4.	From -03 to -04	52
11.5.	From -04 to -05	53
11.6.	From -05 to -06	54
11.7.	From -06 to -07	54
11.8.	From -07 to -08	54
11.9.	From -08 to -09	55
11.10.	From -09 to -10	55
11.11.	From -10 to -11	55
11.12.	From -11 to -12	55
11.13.	From -12 to -13	55
11.14.	From -13 to -14	55
12.	Informative References	55
	Authors' Addresses	57

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:

- o Operators: to help plan their network and identify faults
- o 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:

- o 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.
- o 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).
- o 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.
- o 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 $xx.05 + \text{Unif}[0,180]$ 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 + $\text{Unif}[0,180]$ 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:

- o a Bootstrapping process before the MA can take part in the other three phases.
- o a Control Protocol, which delivers Instruction Messages from a Controller to a MA (amongst other things).

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

- o (optional): indicated by round brackets
- o [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):

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

- o 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

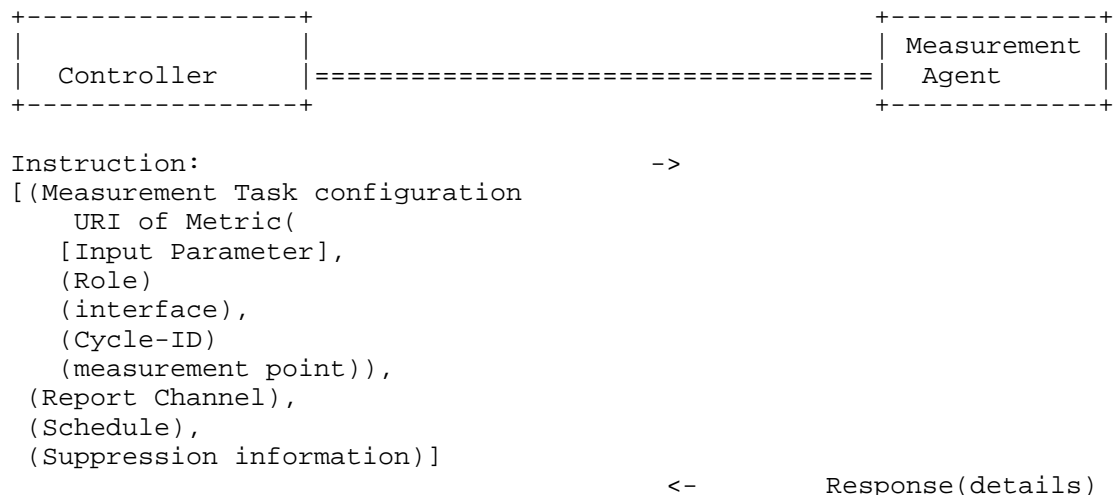
Configuration allows the Controller to update the MA about some or all of the information that it obtained during the bootstrapping process: the MA-ID, the (optional) Group-ID and the Control Channel. The Measurement System might use Configuration for several reasons. For example, the bootstrapping process 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):

- o 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.
- o 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
 - o 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
 - o 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:


- o 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.
- o 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.
- o 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.

- o 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.
- o 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".
- o a start time, at which suppression begins. If absent, then Suppression begins immediately.
- o an end time, at which suppression ends. If absent, then Suppression continues until the MA receives an un-Suppress message.
- o 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:

- o the Measurement Method (roles) that the MA supports
- o the measurement protocol types and roles that the MA supports
- o the interfaces that the MA has
- o the version of the MA
- o the version of the hardware, firmware or software of the device with the MA
- o its Instruction (this could be useful if the Controller thinks something has gone wrong, and wants to check what Instruction the MA is using)
- o 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:


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

- o the MA failed to record the Measurement Results because it (unexpectedly) is out of spare memory
- o a Report failed to deliver Measurement Results because the Collector (unexpectedly) is not responding
- o 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:

- o the last time the MA ran a Measurement Task
- o the last time the MA sent a Measurement Report
- o the last time the MA received an Instruction Message
- o 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.



```

(Instruction:
  [(Request Capabilities),
   (Request Failure Information),
   (Request Logging Information),
   (Request Instruction)])
                                ->
                                <-      (Capabilities),
                                           (Failure Information),
                                           (Logging Information),
                                           (Instruction)

```


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:

- o checking that there is traffic of interest;
- o 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:

- o the MA checking that there is no cross-traffic. In other words, a check that the end-user isn't already sending traffic;
- o 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;

- o sending traffic that probes the path to check it isn't overloaded;
- o 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:

- o For 'upload' tests, the MA not sending traffic
- o 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.



```

                                <-    Report:
                                      [MA-ID &/or Group-ID],
                                      [Measurement Result],
                                [details of Measurement Task],
                                      [Cycle-ID]
ACK                                ->
  
```

The Report contains:

- o the MA-ID or a Group-ID (to anonymise results)
- o 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.
- o 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.
- o the Cycle-ID, if one was included in the Instruction.
- o perhaps the Subscriber's service parameters (see Section 5.4.1).
- o 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:

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

- o a 'push' protocol (that is, from the Controller to the MA)
- o a multicast protocol (from the Controller to a group of MAs)
- o 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.
- o 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:

- o 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.

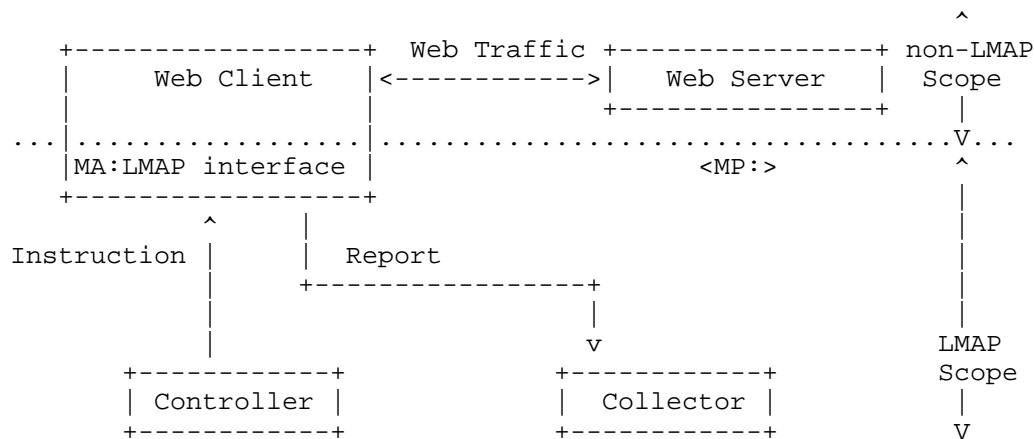
A device may participate in some Measurement Methods as a Measurement Agent and in others as a Measurement Peer.

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.

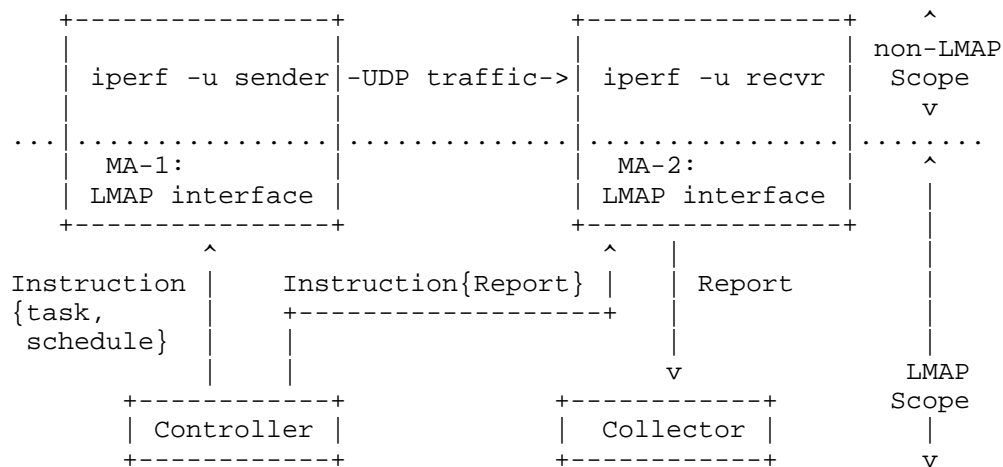
A further example is the case of a traceroute like measurement. In this case, for each packet sent, the router where the TTL expires is performing the MP function. So for a given Measurement Task, there is one MA involved and several MPs, one per hop.

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:

- o limit the number of reporting connections from a single MA (simultaneous, and connections per unit time).
- o limit the transmission rate from a single MA.
- o limit the memory/storage consumed by a single MA's reports.
- o efficiently reject reporting connections from unknown sources.
- o 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.

- o 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.
- o 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.
- o 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.
- o 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:

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

Examples of Internet Service Provider sensitive information:

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

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

- o the MA's identifier (MA-ID)
- o the address that identifies the Control Channel, such as the Controller's FQDN
- o 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.

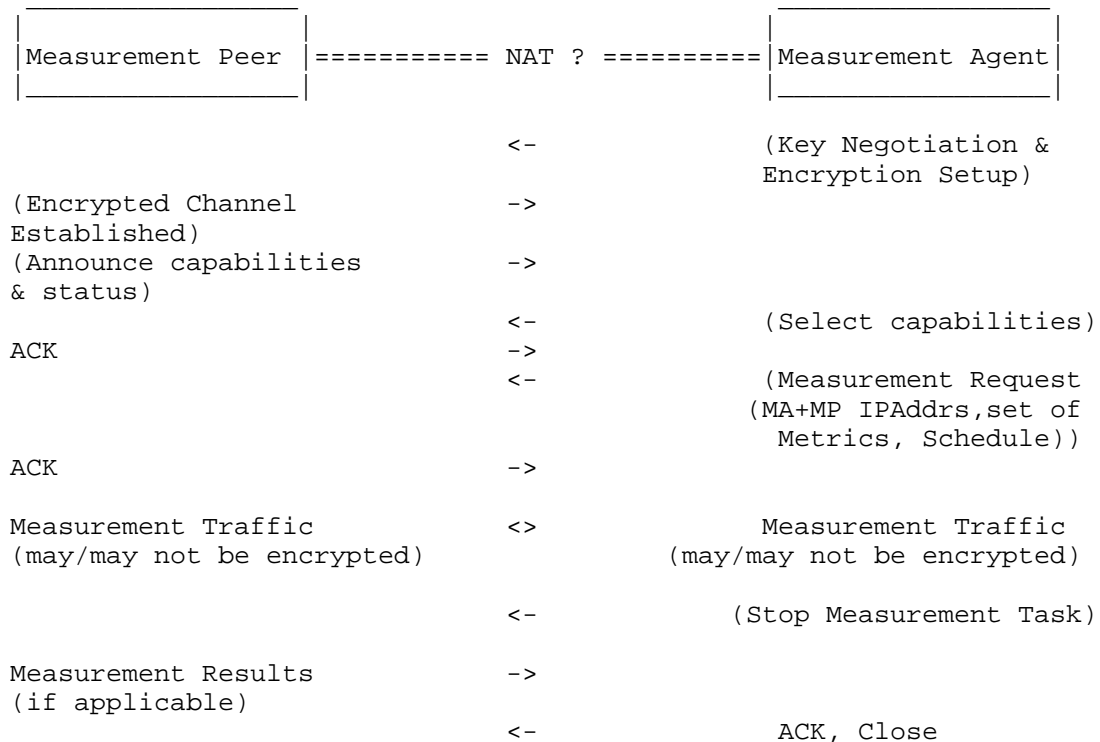
The Measurement Results are the additional sensitive information included in the Collector-MA exchange. Organisations collecting LMAP measurements have the responsibility for data control. Thus, the Results and other information communicated in the Collector protocol must be secured.

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 "mpl90".

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 Burbidge 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

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

11.2. From -01 to -02

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

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

11.3. From -02 to -03

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

11.4. From -03 to -04

- o 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:

- o tweaked definitions, especially of Measurement Agent and Measurement Peer
- o Instruction - left to each implementation & deployment of LMAP to decide on the granularity at which an Instruction Message works
- o words added about overlapping Measurement Tasks (Measurement System can handle any way they choose; Report should mention if the Task overlapped with another)
- o 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
- o new concept of Data Transfer Task (and therefore adjustment of the Channel concept)
- o 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
- o various other smaller improvements, arising from the WGLC
- o 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

- o 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)
- o 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)
- o 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).
- o the Measurement Method, specified as a URI to a registry entry - rather than a URN

- o 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)
- o 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).
- o removed mention of Data Transfer Tasks. This abstraction is left to the information model i-d
- o added Deployment sub-section about Measurement Agent embedded in ISP Network
- o various other smaller improvements, arising from the 2nd WGLC

11.6. From -05 to -06

- o clarified terminology around Measurement Methods and Tasks. Since within a Method there may be several different roles (requester and responder, for instance)
- o 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.
- o The previous bullet also means there is no need to make a distinction between active and passive Measurement Tasks, so this distinction is removed.
- o removed Configuration Protocol - Configuration is part of the Instruction and so uses the Control Protocol.

11.7. From -06 to -07

- o Clarifications and nits

11.8. From -07 to -08

- o 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.
- o added mention of "measurement point designations" in Measurement Task configuration and Report Protocol.

11.9. From -08 to -09

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

11.10. From -09 to -10

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

11.11. From -10 to -11

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

11.12. From -11 to -12

- o Fixing nits from IETF Last call and authors.

11.13. From -12 to -13

- o IESG changes.

11.14. From -13 to -14

- o 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/sdcp-s-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

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.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1	Introduction	3
2	Use Cases	3
2.1	Internet Service Provider (ISP) Use Case	3
2.2	Regulator Use Case	4
3	Details of ISP Use Case	5
3.1	Understanding the quality experienced by customers	5
3.2	Understanding the impact and operation of new devices and technology	6
3.3	Design and planning	6
3.4	Monitoring Service Level Agreements	7
3.5	Identifying, isolating and fixing network problems	7
4	Details of Regulator Use Case	8
4.1	Providing transparent performance information	8
4.2	Measuring broadband deployment	9
4.3	Monitoring traffic management practices	9
6	Conclusions	11
7	Security Considerations	13
8	IANA Considerations	14
	Contributors	14
	Informative References	14
	Authors' Addresses	17

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:

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

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

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

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

- o Understanding the impact and operation of new devices and technology. As a new product is deployed, or a new technology introduced into the network, it is essential that its operation and its impact 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:

- o Check whether a customer can be upgraded to a new broadband option
- o Understand the impact of IPv6 before it is made available to customers. Questions such as these could be assessed: will v6 packets get through? what will the latency be to major websites? what transition mechanisms will be most appropriate?
- o Check whether a new capability can be signaled using TCP options (how often it will be blocked by a middlebox? - along the lines of the experiments described in [Extend TCP]);
- o Investigate a 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 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:

- o Accurate - the measurement results must be correct and not influenced by errors or side effects. The results should be reproducible and consistent over time.
- o Comparable - common metrics should be used across different ISPs and service offerings, and over time, so that measurement results can be compared.

- o Meaningful - the metrics used for measurements need to reflect what end users value about their broadband Internet access service.
- o Reliable - the number and distribution of measurement agents, and the statistical processing of the raw measurement data, 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/AMIfv9451jiJXzG-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: March 30, 2014

M. Bagnulo
UC3M
A. Morton
AT&T Labs
P. Eardley
BT
September 30, 2013

A Registry for Performance Metrics
draft-mornulo-ippm-registry-00

Abstract

This memo investigates a scheme to organize registry entries, especially those defined in RFCs prepared in the IP Performance Metrics (IPPM) Working Group of the IETF, and applicable to all 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 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. Specifically, this memo proposes a way to organize registry entries into columns that are well-defined, permitting consistent development of entries over time. Also, this fosters development of registry entries based on existing reference RFCs for performance metrics, and requires expert review for every entry before IANA action.

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 March 22, 2014.

Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Background and Motivation	4
2. Scope	5
3. Registry Columns and Sub-Columns	5
3.1. Metric ID	6
3.2. Metric Name	6
3.3. Metric Description	7
3.4. Method of Measurement	7
3.4.1. Reference Method	7
3.4.2. Fixed Parameters	7
3.4.3. Schedule	8
3.4.4. Output Type	8
3.4.5. Run-time Parameters	9
3.5. Metric Units	10
3.6. Measurement Point	10
3.7. Timing	10
3.8. Other	10
4. Example of allocation	10
4.1. UDP latency metric	10
5. Security Considerations	13
6. IANA Considerations	13
7. Acknowledgements	13
8. References	13
8.1. Normative References	13
8.2. Informative References	14
Authors' Addresses	14

1. Introduction

This memo investigates a scheme to organize registry entries, 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 registry is based on the template defined in [RFC6390] expanded with further details to fully cover the needs of a registry.

The authors of [draft-bagnulo-ippm-new-registry] and [draft-bagnulo-ippm-new-registry-independent] made important contributions to this memo in the registry column structure, and the problem of registry development in general. We also acknowledge input from the authors of [draft-claise-ippm-perf-metric-registry], especially the value of an Element ID and the need for naming conventions.

1.1. Background and Motivation

The motivation for having such registry is to allow a controller to request a measurement agent to execute a measurement using a specific metric. Such 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 for each. The idea is that the entries in the registry represent different measurement tests, whilst the run-time parameters set things like source and destination

addresses that don't change the fundamental nature of the test. The downside of this approach is that it could result in an explosion in the 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 this document defines that the registry only includes commonly used metrics that are well defined; hence we require both reference specification required AND expert review policies for the assignment of values in the registry.

There are a couple of side benefits of having such registry. 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 is properly defined.

2. Scope

Specifically, this memo proposes a way to organize registry entries into columns that are well-defined, permitting consistent development of entries over time. Also, this fosters development of registry entries based on existing reference RFCs for performance metrics, and requires expert review for every entry before IANA action.

In this memo, we attempt a combinatoric registry, where all factors that can be reasonably specified ARE specified, and changing even one factor would require a new registry entry (row). It is believed that this exercise can also be instructive for a registry based on independent factors, [draft-bagnulo-ippm-new-registry-independent] but that topic is beyond the scope of this effort.

Entries in the registry must reference an existing RFC or other recognized standard, and are subject to expert review. The expert review must make sure that the proposed metric is operationally useful. This means that the metric has proven to be useful in operational/real scenarios.

3. Registry Columns and Sub-Columns

This section describes the columns and sub-columns proposed for the registry. Below, columns are described at the 3.x heading level, and sub-columns are at the 3.x.y heading level. The Figure below illustrates this organization.

Taken as a whole, each entries (row) in the registry gives a

registered instance of a metric with sufficient specificity to promote comparable results across independent implementations. In other words, a **complete description** of a Metric Instance. Some instances may not require entries in all sub-columns, but this is preferred to more general organization because each sub-column serves as a check-list item and helps to avoid omissions during registration and expert review. The columns are extracted directly from [RFC6390] while the sub-columns provide additional detailed required for each column.

ID	Name	Description	Method	Units	Measurement Points	Timing	Other

Figure 1: Registry columns

We describe the content of each of the columns next, some of them contain sub-columns.

3.1. Metric ID

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

3.2. Metric Name

The current guidance from Section 13 of [RFC2330], where Type-P is a feature of all IPPM metric names, is:

"... we introduce the generic notion of a "packet of type P", where in some contexts P will be explicitly defined (i.e., exactly what type of packet we mean), partially defined (e.g., "with a payload of B octets"), or left generic. Thus we may talk about generic IP-type-P-connectivity or more specific IP-port-HTTP-connectivity. Some metrics and methodologies may be fruitfully defined using generic type P definitions which are then made specific when performing actual measurements. Whenever a metric's value depends on the type of the packets involved in the metric, the metric's name will include either a specific type or a phrase such as "type-P". ..."

Registry entries are a context where Type-P must be defined.

IPPM Metric names have also included the typically included the stream type, to distinguish between singleton and sample metrics (see [RFC2330] for the definition of these terms).

Based on this, the metric name is composed in the following way:

P_type-Descriptive_name-Schedule-Output-Other, where:

- o P-type is a text describing the P-type selected
- o Descriptive_name describes the nature of the metric
- o The schedule describes the so-called stream type
- o The output describes the expected output of the metric, in particular, the type of statistic which is outputted, if it is one.
- o Other, describes other consideration that affects the nature of the metric, for example the presence or absence of cross traffic

3.3. Metric Description

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.

3.4. Method of Measurement

This column is composed by the following sub-columns:

Method				
Reference Method	Fixed Parameters	Schedule	Output Type	Run-time Param

3.4.1. Reference Method

This sub-column provides references to relevant sections of the specifications or 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.

3.4.2. Fixed Parameters

In the case that the metric is defined as a more specific instance of a broader metric defined in the specification pointed in the "Reference method" column, this is done by defining as fixed some of the open parameters defined in the broader metric. If this should be the case of the specific entry in the registry, this sub-column specifies the values of these parameters in the Registry.

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

3.4.3. Schedule

Principally, two different schedules 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.

Some metrics, such as those intended for passive monitoring or RTCP and RTCP-XR metrics, will not specify an entry for this column.

Each entry for this sub-column contains the following information:

- o Value: The name of the packet stream scheduling discipline
- o Schedule 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.
- o Reference: the specification where the stream is defined

+-----+-----+-----+		
	Schedule	
+-----+	+-----+	+-----+
Value	Schedule Parameters	Reference
+-----+	+-----+	+-----+

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.

3.4.4. Output Type

For some entries, 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:

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

Output type		
Value	Data format	Reference

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. 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").

3.4.5. Run-time Parameters

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 actual 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 measurement-specific information.

3.5. Metric Units

The measured results of a metric must be expressed using some standard dimension or units of measure. This column provides the units (and if possible, the data format, whose specification will simplify both measurement implementation and collection/storage tasks, see the Output Type column below).

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

3.6. Measurement Point

Measurement Point(s) with potential Measurement Domain: A pointer to the specification that defines whether the metric is specific to a given measurement point or measurement domain. A canonical reference path is defined in [I-D.ietf-ippm-lmap-path].

3.7. Timing

A pointer to the specification where the acceptable range of timing intervals or sampling intervals are defined, if any.

3.8. Other

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 of allocation

In this section we provide a few example of allocations.

4.1. UDP latency metric

The registry entry for for the Xth percentile mean of the UDP latency using a Poisson stream of packets would look like this:

ID: 344 (for example, typically assigned by IANA)

Name: UDP-Latency-Poisson-Xth_percentile_mean.

Description: This metric is a specific instance of the Round trip metric defined in RFC2681 and it measures the Xth percentile mean of the UDP latency of a Poisson stream of packets.

Method:

Reference Method: The methodology for this metric is defined as Type-P-Round-trip-Delay-Poisson-Stream in RFC 2681.

Fixed Parameters:

P-Type:

IPv4 header values:

DSCP: set to 0

TTL set to 255

Protocol: Set to 17 (UDP)

UDP header values: Checksum: the checksum must be calculated

Payload

Sequence number: 8-byte integer

Timestamp: 8 byte integer. Expressed as 64-bit NTP timestamp as per section 6 of RFC 5905

No padding

Timeout: 3 seconds

Schedule:

Value: Poisson

Schedule Parameters:

lambda: the parameter defining the Poisson distribution. Lambda is the mean number of distinct measurements per second in the sample.

T0: time to begin a test

Tf: time to end a test

T0 and Tf are both in seconds and use the date (yyyy-mm-dd) and NTP 64 bit timestamp. T0 includes any control handshaking before the test stream or singleton. Tf is the time the last test data is sent. As a result, we have that the time when test devices may close the test socket is Tf + Waiting Time (the time to wait before declaring a packet lost is fixed for each metric) and the Total duration of the test: $Tf - T0 + \text{Waiting Time}$

Reference: The Poisson scheduling is defined in section 11.1.1.1 of RFC 2330

Output Type

Value: Xth percentile mean

Data format:

Reference:

Run-time Param

Source IP Address

Destination IP Address

Source UDP port

Destination UDP port

Initial time T0

end time Tf

Rate lambda

X

Units: milliseconds

Measurement Points: The metric is not specific to any particular measurement point.

Timing: between microseconds and seconds

Other

5. Security Considerations

This registry has no known implications on Internet Security.

6. IANA Considerations

Metrics previously defined in IETF were registered in the IANA IPPM METRICS REGISTRY, however this process was discontinued when the registry structure was found to be inadequate, and the registry was declared Obsolete [RFC6248].

The form of metric registration will be finalized in the future, and no IANA Action is requested at this time.

7. Acknowledgements

The author thanks Brian Trammell for suggesting the term "Run-time Parameters", which led to the distinction between run-time and fixed parameters implemented in this memo.

8. References

8.1. Normative References

- [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.

8.2. Informative References

- [I-D.ietf-ippm-lmap-path] Bagnulo, M., Burbridge, T., Crawford, S., Eardley, P., and A. Morton, "A Reference Path and Measurement Points for LMAP", draft-ietf-ippm-lmap-path-00 (work in progress), July 2013.
- [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.
- [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.

Authors' Addresses

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>

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/>

Philip Eardley
British Telecom
Adastral Park, Martlesham Heath
Ipswich
ENGLAND

Email: philip.eardley@bt.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

M. Bagnulo
UC3M
A. Morton
AT&T Labs
P. Eardley
BT
October 21, 2013

A(nother) Registry for Performance Metrics
draft-mornulo-ippm-registry-columns-01

Abstract

This memo investigates a scheme to organize registry entries, especially those defined in RFCs prepared in the IP Performance Metrics (IPPM) Working Group of the IETF, and applicable to all 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 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. Specifically, this memo proposes a way to organize registry entries into columns that are well-defined, permitting consistent development of entries over time. Also, this fosters development of registry entries based on existing reference RFCs for performance metrics, and requires expert review for every entry before IANA action.

This version contains an example registry entry for a passive endpoint metric based on RFC7003, an example active metric entry based on RFC3393 and RFC5481, and an example pure passive metric based on RFC5472. Also, this version *continues* to allow blank entries in columns which have no applicability to a specific metric, or class of metrics. This is preferred to more general registry organization because each column serves as a check-list item and helps to avoid omissions during registration and expert review.

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 April 24, 2014.

Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	5
1.1. Background and Motivation	5
2. Scope	7
3. Registry Categories and Columns	7
3.1. Registry Indexes	8
3.1.1. Element ID	8
3.1.2. Metric Name	8
3.2. Metric Definition	9
3.2.1. Reference Definition	9
3.2.2. Fixed Parameters	9
3.2.3. Metric Units	9
3.3. Method of Measurement	10
3.3.1. Reference Method	10
3.3.2. Stream Type and Stream Parameters	10
3.3.3. Output Type and Data Format	10
3.3.4. Run-time Parameters and Data Format	11
3.4. Comments and Remarks	11
4. Example RTCP-XR Registry Entry	12
4.1. Registry Indexes	12
4.1.1. Element ID	12
4.1.2. Metric Name	12
4.2. Metric Definition	12
4.2.1. Reference Definition	12
4.2.2. Fixed Parameters	12
4.2.3. Metric Units	13
4.3. Method of Measurement	13
4.3.1. Reference Method	13
4.3.2. Stream Type and Stream Parameters	14
4.3.3. Output Type and Data Format	14
4.3.4. Run-time Parameters and Data Format	14
4.4. Comments and Remarks	16
5. Example IPPM Active Registry Entry	16
5.1. Registry Indexes	16
5.1.1. Element ID	16
5.1.2. Metric Name	16
5.2. Metric Definition	16
5.2.1. Reference Definition	16
5.2.2. Fixed Parameters	17
5.2.3. Metric Units	17
5.3. Method of Measurement	17
5.3.1. Reference Method	17
5.3.2. Stream Type and Stream Parameters	17
5.3.3. Output Type and Data Format	18
5.3.4. Run-time Parameters and Data Format	18
5.4. Comments and Remarks	19
6. Example IPFIX RTT Pair Matching Registry Entry	19

6.1. Registry Indexes	19
6.1.1. Element ID	19
6.1.2. Metric Name	19
6.2. Metric Definition	19
6.2.1. Reference Definition	20
6.2.2. Fixed Parameters	20
6.2.3. Metric Units	20
6.3. Method of Measurement	20
6.3.1. Reference Method	20
6.3.2. Stream Type and Stream Parameters	21
6.3.3. Output Type and Data Format	21
6.3.4. Run-time Parameters and Data Format	21
6.4. Comments and Remarks	21
7. Security Considerations	21
8. IANA Considerations	22
9. Acknowledgements	22
10. References	22
10.1. Normative References	22
10.2. Informative References	23
Authors' Addresses	23

1. Introduction

This memo investigates a scheme to organize registry entries, 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.

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 examine the aspects of metric specifications which need to be registered, we will see that none of the existing metric templates fully satisfies the needs of a registry.

The authors of [draft-bagnulo-ippm-new-registry] and [draft-bagnulo-ippm-new-registry-independent] made important contributions to this memo in the registry column structure, and the problem of registry development in general. We also acknowledge input from the authors of [draft-claise-ippm-perf-metric-registry], especially the value of an Element ID and the need for naming conventions.

1.1. Background and Motivation

The motivation for having such registry is to allow a controller to request a measurement agent to execute a measurement using a specific metric. Such 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 for each. The idea is that the entries in the registry represent different measurement tests, whilst the run-time parameters set things like source and destination addresses that don't change the fundamental nature of the test. The downside of this approach is that it could result in an explosion in the 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 this document defines that the registry only includes commonly used metrics that are well defined; hence we require both reference specification required AND expert review policies for the assignment of values in the registry.

There are several side benefits of having such a registry. 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 is properly defined.

The registry forms part 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 Instructions that will be communicated to measurement agents, whose task is then to execute the Instruction 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 the 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

Specifically, this memo proposes a way to organize registry entries into columns that are well-defined, permitting consistent development of entries over time. Also, this fosters development of registry entries based on existing reference RFCs for performance metrics, and requires expert review for every entry before IANA action.

In this memo, we attempt a combinatoric registry, where all factors that can be reasonably specified ARE specified, and changing even one factor would require a new registry entry (row). It is believed that this exercise can also be instructive for a registry based on independent factors, [draft-bagnulo-ippm-new-registry-independent] but that topic is beyond the scope of this effort.

Entries in the registry must reference an existing RFC or other recognized standard, and are subject to expert review. The expert review must make sure that the proposed metric is operationally useful. This means that the metric has proven to be useful in operational/real scenarios.

3. Registry Categories and Columns

This section briefly describes the categories and columns proposed for the registry, as this is likely to be a topic for discussion and revision. 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.

Taken as a whole, the entries in the columns give a registered instance of a metric with sufficient specificity to promote comparable results across independent implementations. In other words, a *complete description* of a Metric Instance. Some instances may not require entries in all columns, but this is preferred to more general organization because each column serves as a check-list item and helps to avoid omissions during registration and expert review.

Registry Categories and Columns, shown as

Category			

Column		Column	
Registry Indexes			

Element ID		Metric Name	
Metric Definition			

Reference Definition		Fixed Parameters	Metric Units
Method of Measurement			

Reference Method		Stream Type and Param	Output Type Run-time Param
Comments and Remarks			

3.1. Registry Indexes

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

3.1.1. Element ID

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

3.1.2. Metric Name

A metric naming convention is TBD.

The current guidance from Section 13 of [RFC2330], where Type-P is a feature of all IPPM metric names, is:

"... we introduce the generic notion of a "packet of type P", where in some contexts P will be explicitly defined (i.e., exactly what type of packet we mean), partially defined (e.g., "with a payload of B octets"), or left generic. Thus we may talk about generic IP-type-P-connectivity or more specific IP-port-HTTP-connectivity. Some metrics and methodologies may be fruitfully defined using generic type P definitions which are then made specific when performing

actual measurements. Whenever a metric's value depends on the type of the packets involved in the metric, the metric's name will include either a specific type or a phrase such as "type-P". ..."

Registry entries are a context where Type-P must be defined.

IPPM Metric names have also included the typically included the stream type, to distinguish between singleton and sample metrics (see [RFC2330] for the definition of these terms).

3.2. 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.

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

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.

3.2.3. Metric Units

The measured results of a metric must be expressed using some standard dimension or units of measure. This column provides the units (and if possible, the data format, whose specification will simplify both measurement implementation and collection/storage tasks, see the Output Type column below).

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. 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.

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.

Some metrics, such as those intended for passive monitoring or RTCP and RTCP-XR metrics, will not specify an entry for this column.

Each entry for this column contains the following information:

- o Value: The name of the packet stream scheduling discipline
- o 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.
- o 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:

- o Value: The name of the output type
- o Data Format: provided to simplify the communication with collection systems and implementation of measurement devices.
- o 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. 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. 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 measurement-specific information.

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 RTCP-XR Registry Entry

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

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.

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. Section 3.2 of [RFC7003] provides the reference information for this category.

4.2.1. 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 0xFFFFF MUST be reported.

4.2.2. 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.

4.2.3. Metric Units

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

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. 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].

4.3.1. 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 RTP packet as this metrics block. If not, this metrics block MUST be discarded.

4.3.2. Stream Type and Stream Parameters

Since RTP-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.

4.3.3. Output Type and Data Format

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

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

4.3.4. 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 10.47.16.5

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 224.2.17.12/127

t=2873397496 2873404696

a=recvonly

m=audio 49170 RTP/AVP 0
```

```
m=video 51372 RTP/AVP 99
a=rtpmap:99 h263-1998/90000
```

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 IPPM Active Registry Entry

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

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.

One possibility based on IPPM's framework is:

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

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

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

5.2.2. Fixed Parameters

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

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

5.2.3. 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.

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.

5.3.1. Reference Method

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

5.3.2. 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.

5.3.3. 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.

5.3.4. Run-time Parameters and Data Format

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. In related registry entries, some of the parameters below may be designated as Fixed Parameters instead.

- o `Src`, the IP address of a host (32-bit value for IPv4, 128-bit value for IPv6)
- o `Dst`, the IP address of a host (32-bit value for IPv4, 128-bit value for IPv6)
- o `T`, a time (start of test interval, 128-bit NTP Date Format, see section 6 of [RFC5905])
- o `Tf`, a time (end of test interval, 128-bit NTP Date Format, see section 6 of [RFC5905])
- o `Tl`, 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]).

- o 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]).
- o I(i), I(i+1), $i \geq 0$, 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, $I(0) = T_0$ and assuming that n is the largest index, $I(n) = T_f$ (pairs of 64-bit NTP Timestamp Format, see section 6 of [RFC5905]).

5.4. 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 IPDV.

6. Example IPFIX RTT Pair Matching Registry Entry

This section gives an example registry entry for the passive metric described in section 2.5.2.1 of [RFC5472], for Round-Trip Time (RTT) Measurements with Packet Pair Matching (Single-Point).

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.

6.2. 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 2.5.2.1 of [RFC5472] provides the reference information for this category.

6.2.1. Reference Definition

Observations of both directions are required to correlate request/response packet pairs.

Pair matching techniques are described in [Brow00].

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.

Protocol (Pair Type): TCP (SYN/SYN_ACK)

Note: other possibilities are DNS, ICMP, SNMP or TCP (DATA/ACK), discussed in [Brow00].

6.2.3. Metric Units

The measured results are expressed in microseconds, which follows the format of Information Elements per observed packet, see section 8.4.3 of[RFC5477] titled "observationTimeMicroseconds".

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 the TCP(SYN/SYN_ACK) RTT metric, the guidance on methods of measurement is in slides 12 and 15 of [Brow00].

Recognition of request response pairs is a REQUIRED function, as is the correlation of data from both directions of transmission, see section 2.5.2.1 of [RFC5472].

The method requires the collection of the following Information Elements per packet:

- o Packet arrival time: observationTimeMicroseconds, see section 8.4.3 of[RFC5477]
- o TCP header: ipPayloadPacketSection, see section 8.5.2 of[RFC5477]

6.3.2. Stream Type and Stream Parameters

Since IPFIX passive Measurements are conducted on live/production network traffic, the measurement methods rely on user-generated packet flows. Such flows are not described in this column.

6.3.3. Output Type and Data Format

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

- o Value: RTT in microseconds
- o Data Format: (There may be some precedent to follow here, but otherwise use 64-bit NTP Timestamp Format, see section 6 of [RFC5905]).
- o Reference: Section 2.5.2.1 of [RFC5472]

6.3.4. 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 list of Run-time parameters is not specified for purely passive metrics, as there are infinite possibilities.

A likely Run-time parameter is the Destination host, which may be given as a Fully-Qualified Domain Name as done in [Brow00], or an IP address of the host (32-bit value for IPv4, 128-bit value for IPv6).

6.4. Comments and Remarks

Additional (Informational) details for this entry, from [Brow00]:

Can't get RTT for every packet, only those which are ACKed.

Overlapping packets (resent) are counted as lost, but not queued. This means the first copy of resent packets are used for RTTs, giving a high RTT estimate.

7. Security Considerations

This registry has no known implications on Internet Security.

8. IANA Considerations

Metrics previously defined in IETF were registered in the IANA IPPM METRICS REGISTRY, however this process was discontinued when the registry structure was found to be inadequate, and the registry was declared Obsolete [RFC6248].

The form of metric registration will be finalized in the future, and no IANA Action is requested at this time.

9. Acknowledgements

The author thanks Brian Trammell for suggesting the term "Run-time Parameters", which led to the distinction between run-time and fixed parameters implemented in this memo.

10. References

10.1. Normative References

- [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.
- [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

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>

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/>

Philip Eardley
British Telecom
Adastral Park, Martlesham Heath
Ipswich
ENGLAND

Email: philip.eardley@bt.com

