

# Building a Standard Measurement Platform

Marcelo Bagnulo, Trevor Burbridge, Sam Crawford, Philip Eardley, Juergen Schoenwaelder, and Brian Trammell

## ABSTRACT

Network management is achieved through a large number of disparate solutions for different technologies and parts of the end-to-end network. Gaining an overall view, and especially predicting the impact on a service user, is difficult. Recently, a number of proprietary platforms have emerged to conduct end-to-end testing from user premises; however, these are limited in scale, interoperability, and the ability to compare like-for-like results. In this article we show that these platforms share similar architectures and can benefit from the standardization of key interfaces, test definitions, information model, and protocols. We take the SamKnows platform as a use case and propose an evolution from its current proprietary protocols to standardized protocols and tests. In particular, we propose to use extensions of the IETF's IPFIX and NETCONF/YANG in the platform. Standardization will allow measurement capabilities to be included on many more network elements and user devices, providing a much more comprehensive view of user experience and enabling problems and performance bottlenecks to be identified and addressed.

## INTRODUCTION

Network management has always been the Cinderella child: remembered only reluctantly, and after dealing with the more exciting business of adding new applications and link layer technologies. This has led to network management comprising a whole series of more or less independent solutions cobbled together to tackle what are seen as more or less independent problems. From a technical point of view, this approach becomes harder to sustain with the ever increasing scale, scope, and rate of change of the Internet: the tens of thousands of autonomously administered networks, the move from centralized client-server applications to fully distributed cloud computing, and the proliferation of super-fast access speeds. From a socio-economic perspective, the problem is even more urgent because the Internet is now vital

infrastructure; while the technical expert knows that the underlying technology is best effort, from a user's perspective the Internet is a critical utility — in fact, in a U.K. survey in 2011 of “things you couldn't live without,” the Internet was rated second, ahead of water, a cooker, and a mobile phone, and only behind sunshine!

In order to gain deeper insight about Internet performance, several large-scale measurement platforms (e.g., SamKnows1, RIPE Atlas2, and Netalyzr [1]) have been deployed over the last few years. They encompass several thousand probes distributed across the Internet. Most of these probes are located in homes and can perform active tests against measurement servers located in the core of the Internet, and they are used to assess the performance of Internet access from the perspective of an end user. Currently, these measurement platforms are mainly used for benchmarking Internet service providers (ISPs).

While these platforms have been proven useful, we believe they have not realized their full potential due to their proprietary and consequently closed nature. All the current platforms use proprietary protocols and proprietary metrics/tests, which has negative side effects, including:

- The platforms cannot interact with each other. Measurement platforms clearly benefit from the so-called network effect, meaning that a platform is more useful the more probes it has. Because of that, platforms would benefit from some form of interaction that would allow the pooling of the probes from multiple platforms. Different levels of interaction can be envisioned, ranging from enabling one platform to use the probes of another platform to supporting coordinated tests that are executed in parallel on different platforms. This is not possible today as deployed platforms are bespoke systems that cannot interact with other systems through standard interfaces.

- The deployment of platforms is expensive. Since each platform uses proprietary protocols, whenever a new party wants to deploy probes to perform its own measurements, it cannot buy off-the-shelf components that natively interact to

---

Marcelo Bagnulo is with Universidad Carlos III de Madrid.

Trevor Burbridge and Philip Eardley are with BT.

Sam Crawford is with SamKnows.

Juergen Schoenwaelder is with Jacobs University.

Brian Trammell is with ETH Zurich.

*The network is aimed at larger network operators, primarily those associated with research networks connected to Internet2 and/or GEANT, to support performance troubleshooting by their network engineering staffs.*

form a measurement platform. Instead, it should be possible to have components from multiple vendors.

• Because the tests and the metrics themselves are proprietary, it is not safe to compare results of metrics that are supposedly the same, since their definitions may differ. This also implies that the results obtained from different platforms cannot be aggregated to obtain a larger data set, severely reducing the usefulness of the obtained results.

The result is that the community has invested a large amount of resources deploying these platforms, but it is not possible to take full advantage of what such infrastructure could provide.

In this article, we propose to extend a set of standard protocols to build a large-scale measurement platform. As a result, it would be possible to build a measurement platform using standard off-the-shelf components and build federations of measurement platforms. The context of this work is the efforts being carried out in several standards development organizations (SDOs): in the Internet Engineering Task Force's (IETF's) LMAP and IPPM Working Groups; in the Broadband Forum's WT-304 activity; and in the IEEE's project P802.16.3.

In this article we follow an evolutionary approach. We examine the state of the art in measurement platforms, focusing on the SamKnows platform as a case study, and we propose an evolution of the platform toward open standards. We start by describing the current architecture of existing platforms and then describe how to evolve to a standard architecture, the information model, standard metrics, and standard protocols.

## STATE OF THE ART IN MEASUREMENT PLATFORMS

Here we examine three current active network measurement platforms, RIPE Atlas, perfSONAR, and SamKnows, as background for identifying the requirements for a standard measurement platform.

### RIPE ATLAS

Atlas (<http://atlas.ripe.net>) is an active Internet measurement network developed and run by the RIPE Network Coordination Centre (NCC). Small hardware probes, produced by RIPE, perform basic active measurements: ICMP ping, traceroute, and DNS, HTTP, and SSL certificate retrieval queries. It is aimed at engineers with some networking expertise, for whom its primary advantage is that it provides multiple vantage points from which these measurements can be run. Data collected from these probes are also available for research purposes, and provide a data source for RIPE Labs research studies. Users build up credit by hosting RIPE probes, and can use this credit to run their own measurements as well.

Currently, slightly fewer than 4000 probes are connected to the network at any given time; basic tests include ping and traceroute to a selection of DNS servers, as well as ping measure-

ments to the first two hops to the Internet via IPv4 and IPv6.

Control is distributed among a set of RIPE Atlas controllers, running proprietary control and reporting protocols designed by RIPE running over Secure Shell (SSH), making use of multiple SSH channels to logically separate different traffic types. Results are centralized at the controllers, and are available via a web interface as well as a RESTful application programming interface (API).

### PERFSONAR

perfSONAR [2] is a service-based platform for distributed passive and active measurements. It features a three-layer architecture, including a user interface layer and a services layer backed up by a measurement point layer. The controllers and collectors reside in the services layer.

The basic services provided include:

- A measurement point (MP) service, which performs measurements and is a service wrapper around an entity in the measurement point layer
- Measurement archive (MA) service, providing access to historical monitoring data and/or storing the results produced by an MP
- A transformation service (TS), which provides transformation of data (e.g., aggregation, correlation, filtering) provided by other services
- A lookup service (LS), which provides discovery of other services
- An authentication service (AS), which works with the lookup service to provide access control restrictions to measurement services
- A resource protector service (RP), which avoids overload of shared measurement resources

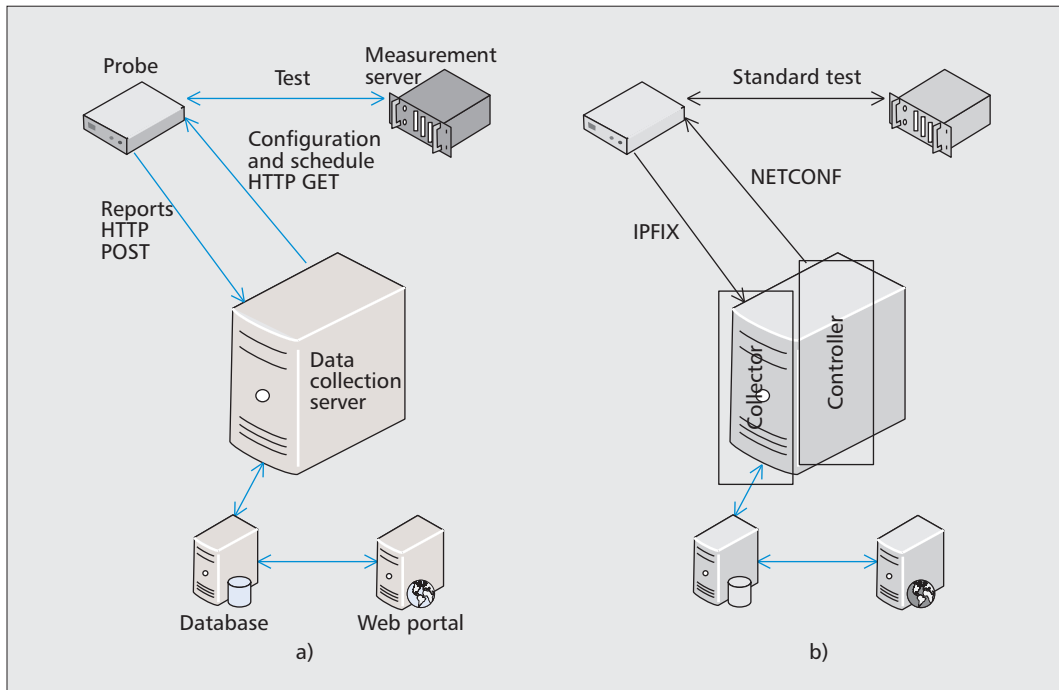
The measurements supported include ping and traceroute, as with RIPE Atlas, as well as one-way delay via One-Way Active Measurement Protocol (OWAMP) [2] and achievable bandwidth testing. perfSONAR can also integrate passive observation of network flows and Simple Network Management Protocol (SNMP) counters. Multiple user interfaces and data analysis tools run atop the perfSONAR service, which is itself concerned primarily with providing distributed access to measurement data and capabilities.

The protocols used between the different components of perfSONAR use NMWG3 messages, with the XML schema defined by the Open Grid Forum. These messages are carried in a Simple Object Access Protocol (SOAP) message carried over HTTP.

The network is aimed at larger network operators, primarily those associated with research networks connected to Internet2 and/or GEANT, to support performance troubleshooting by their network engineering staffs.

### SAMKNOWS

Currently, the SamKnows platform consists of three distinct components: measurement probes, measurement servers, and management infrastructure, as depicted in Fig. 1a.



**Figure 1.** Measurement platform architecture: a) the current SamKnows architecture; b) the proposed reference architecture overlaid on top of the SamKnows architecture.

*These applications on the measurement servers are relatively trivial; almost all of the measurement logic is performed on the client side (at the probe). Measurement servers can be deployed anywhere, and probes can be configured to test against them in any desired combination and frequency.*

The measurement probes are small Linux-based hardware devices. The SamKnows platform currently accounts for over 40,000 probes. These are deployed in volunteers' homes, inline with their existing home network (the probes can operate as Ethernet bridges or routers). The probes run measurements against measurement servers and real endpoints on the Internet when the end user is not actively using their connection. Cross-traffic is identified on both the wired and wireless interfaces. Probes pull their configuration from the management infrastructure upon startup and check for updates periodically thereafter. The testing configuration, frequency, and other associated parameters are all configured remotely.

Measurement servers are Linux servers running a set of custom server applications to support the tests discussed below. These applications on the measurement servers are relatively trivial; almost all of the measurement logic is performed on the client side (at the probe). Measurement servers can be deployed anywhere, and probes can be configured to test against them in any desired combination and frequency.

The probes support a wide range of measurements, including:

- UDP round-trip latency and packet loss (long-lived test, sampled periodically)
- UDP one-way jitter (singleton)
- DNS resolution time and failure rate
- Web browsing (measures transaction time to download a page)
- ICMP round-trip latency and packet loss
- Video streaming (measuring buffer under-runs using a fixed streaming rate)
- Latency and packet loss under load

- TCP throughput, downstream and upstream, using one or multiple concurrent TCP connections.

The management infrastructure consists of three key components.

**The Data Collection Server** — This handles all management interaction with the probes. All communications are conducted over HTTPS (TLS). The communications protocol is a simple but proprietary one. Requests for configuration updates use GET requests, with the current configuration version in the query string, and the response is the tuple <latest\_version, package\_url>. If the client determines that an upgrade is required, it fetches the configuration package (a tar.gz) from the package\_url. Measurement results are uploaded by the probes using the POST method, with a comma-delimited body containing the measurement results (one row per result). Each row contains at a minimum the tuple <probe\_id, metric\_id, timestamp\_utc>. Additional fields are determined by the metric type. Measurement results are written locally to disk on the data collection server and queued for import into the database.

**The Database** — This stores all probe metadata, recent raw measurement data, and summarized historical data. Raw measurement data is imported in bulk from the data collection server frequently (once per minute). This raw data is kept in the database for approximately three months, after which it is archived to flat files. During this time, summaries of the raw data are generated (at reduced resolution) for presentation to end users. Multiple database servers are typically deployed (some operating with master-

The multiple timing options allow for different measuring purposes, for example, calendar based timing allows to target a specific time while instantaneous tests allows to schedule a tests as soon as possible to troubleshoot an ongoing event.

master replication, others as read-only slaves) for redundancy and throughput.

**Web Reporting Portal and Web Services** — This presents the summarized results (stored in the database) to end users and client-side applications.

Note that only the probe initiates all communications. No results are collected or stored on the measurement servers; that is all handled by the probes.

## A STANDARD ARCHITECTURE FOR MEASUREMENT PLATFORMS

A reference architecture provides a common framework and helps identify the protocols that are needed between the different elements of the architecture. In order to hint how the SamKnows platform could evolve toward our reference architecture, Fig. 1b shows them overlaid. The reference architecture contains the following elements:

- *Measurement agents* perform network measurements. They are pieces of code that can be executed in specialized hardware (hardware probe, as in the case of SamKnows) or on a general-purpose device (e.g., a PC or mobile phone). Measurements may be active (agents generate test traffic), passive (agents observe user traffic), or some hybrid form of the two. A measurement agent can perform two distinct roles: either a *measurement client* (hereafter MC) or a *measurement server* (hereafter MS). They correspond to the SamKnows measurement probes and measurement servers, respectively. Note that the MC initiates a test, while the MS simply responds to the MC.

- A *controller* manages MCs by informing them which tests they should perform and when, and also where to report the measurement results and when. We refer to them as the test and report schedules. This is a fundamental component since it is in charge of scheduling measurement activities performed by the MCs.

- A *collector* accepts measurement results from the MCs once their tests are complete. The controller and collector functions are both performed by the SamKnows Data Collection Server.

We believe these are the main components for which standardization is critical, although a measurement platform may include other components, such as a results database, which receives results from collector(s) and processes and stores them; and data analysis tools, which use the data to isolate faults, present results (similar to the SamKnows Web portal), and interact with an operator's operations, administration, and management (OAM) systems.

Having identified the components of the reference architecture, we can easily identify the protocols involved:

- Protocols between the MCs and the MSs. These are the actual tests performed by the platforms and are covered later.
- The protocol between the MC and the controller. We call this the control protocol and cover it in a later section.

- The protocol between the MC and the collector. We call this the report protocol, described further on in the article.

An additional component that is relevant and useful to be standardized is an API to retrieve measurement results data from the platform. This would enable a number of applications. For example, in the case of an ISP that is using the measurement platform to monitor its network, this would allow it to export the measurement data into its OAM systems in a standard way. While we acknowledge that it is a key component, it is out of the scope of any standardization work at this point in time, which is why we do not cover this in any detail in this article.

## INFORMATION MODEL

Before defining the control and report protocols, it is sensible first to define the information model: an abstract, protocol-neutral definition of the data to be transferred. Later we present proposals for protocols and associated data models that implement the information model: NETCONF-YANG for the control protocol and IPFIX for the report protocol.

We believe this is a powerful approach that will prove useful. While in this article we argue for a standardized measurement platform, we also believe that some deployments will use other protocols due to environmental constraints. For example, some of the already deployed DSL access networks may decide to use a transport based on the Broadband Forum protocol TR069 [4]. Defining a protocol-independent information model allows these platforms to use different protocols while still exchanging the same information with the same control and reporting capabilities (e.g., the controller could specify the same calendar-based schedule of the same test with the same configuration parameters). Thus, a single information model ensures a very high level of interoperability between different control and reporting protocols.

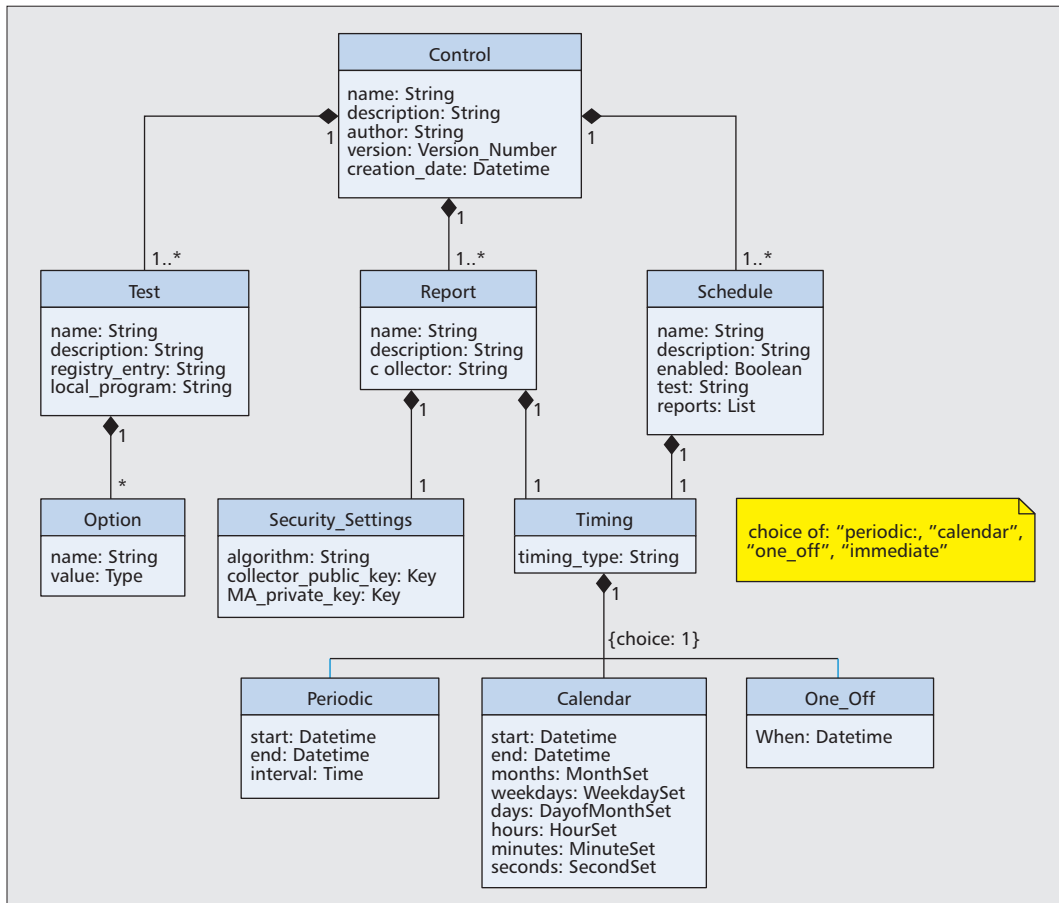
The information model encompasses the elements described next.

The control protocol carries information about the test and report schedules.

The test schedule defines which tests an MC has to perform and with what test parameters (including the MSs to test against). It also covers how to reschedule tests in case connectivity is temporarily lost (e.g., a device is turned off) or in situations where there is too much cross traffic to execute a test. Finally, it contains information about when the tests should be performed.

The report schedule defines when test results are reported (how often), where to (the collector's address), in what format, and what to do if reporting fails.

The structure for the information model of the test and report schedule is presented in UML in Fig. 2. In this proposal the schedule specifies which preconfigured tests to conduct and in which reports the results should be included. It allows for the test and report configuration to be done irregularly and the test schedule to be updated separately. Schedule timing options include periodic, calendar-based, and



In order to close that gap, we propose to complement the IPPM metric definition by specifying the open parameters that fundamentally affect the test (like the packet type) and leaving open only a few parameters that do not change the nature of the test (like the source or destination address).

Figure 2. Proposed information model for the control protocol, in UML

one-off scheduled or instantaneous tests. The multiple timing options allow for different measuring purposes; for example, calendar-based timing allows a specific time to be targeted (e.g. to perform measurement at night that are unlikely to disturb users or peak time traffic, to observe the highest load on the network), while instantaneous tests allows a test to be scheduled as soon as possible to troubleshoot an ongoing event.

The information model for the report protocol includes the MC's identifier; the time of the report; a description of the test (essentially an "echo" of the test schedule, which may be done by reference to a template); and the actual measurement results (which are highly test-specific).

## IPPM-BASED TESTS

The test protocols are executed between an MC and an MS. The SamKnows platform performs proprietary tests based on standard Internet protocols (e.g. TCP, UDP, ICMP, DNS, etc). For example, in order to measure UDP latency, the SamKnows probes send UDP packets, of a certain length from a set of ports and use a periodic schedule with a certain rate. When reporting the results, they exclude the outliers using a specific statistical method (e.g., providing the 95th percentile mean or interval). However, another platform measuring UDP latency would most likely make different choices for the above; thus, it would not be safe to compare the results obtained from the different platforms.

The IPPM Working Group at the IETF has defined a large set of metrics for delay, packet loss, jitter, and many others. A natural approach would be to use these metrics as the tests in a standard measurement platform. However, if we try to map the tests that are actually performed by the SamKnows platform to the defined IPPM metrics, we find that the IPPM metrics are not well defined enough to be useful as test descriptors. The problem is that the IPPM metrics leave too many degrees of freedom to the actual implementation. For example, all the IPPM metrics leave the packet type as an open parameter. This means that referencing a particular IPPM metric does not define whether the packets are TCP, UDP, ICMP, or something else. There seems to be a clear gap where further standards could help.

In order to close that gap, we propose to complement the IPPM metric definition by specifying the open parameters that fundamentally affect the test (like the packet type) and leaving open only a few parameters that do not change the nature of the test (like the source or destination address).

We can map some of the tests available in the SamKnows platform to existing IPPM metrics plus additional specification. As a few examples to illustrate this operation:

- SamKnows defines a UDP latency test. IPPM defines a round-trip delay metric in [2]. In order to bridge the gap between the two of them, we need to specify the packet

*IPFIX organizes data records into Messages. A Message is a sequence of Sets preceded by a Message Header which, among other things, includes an Observation Domain ID (identifying where the records in the Message were measured) and an Export Time.*

type (UDP packet, payload length, and content), scheduling types (periodic in this case), and output type (raw or 95th percentile mean). The source and destination ports and addresses are parameters as well as the time of execution.

- SamKnows defines an ICMP packet loss test. IPPM defines a Round Trip Loss metric in [5]. To use the IPPM specification we need to specify the packet type (ICMP echo request and reply messages), the scheduling and the output type, as above. The input parameters are the source and destination addresses as well as the time of execution are open parameters.

By extending the IPPM specifications we create standardized metrics, so that different implementations and platforms produce measurement results that are comparable. We have specified several other tests in [6].

Once there are well defined standard tests, they can be referred to in the control and report protocols so that an MC can unambiguously understand from the controller what test to perform (and later the collector knows what test the MC has done).

## A NETCONF/YANG-BASED CONTROL PROTOCOL

The control protocol is executed between the controller and an MC to configure test and report schedules, which was defined in a protocol-neutral fashion. In this section we propose to use the IETF standard protocol NETCONF and the IETF standard data modeling language YANG. This would be an alternative to the proprietary protocol currently used in the SamKnows platform.

During the last 10 years, the IETF has developed a generic protocol to support device configuration called NETCONF and an associated data modeling language called YANG [7]. The NETCONF protocol provides a remote procedure call mechanism running over a secure transport (SSH or TLS). On top of the generic RPC layer, a number of specific operations are defined to retrieve and edit a device's configuration (e.g., get-config, edit-config, copy-config, delete-config). In addition, there are standard operations to support coarse and fine-grained locking or to implement configuration change transactions over a number of devices. The configuration data manipulated by NETCONF is a structured document conceptually stored in a configuration datastore and serialized using XML. The structure and semantics of the configuration data manipulated by NETCONF is defined using the YANG data modeling language. In addition, YANG allows a data modeler to define:

- New operations that extend the core set of generic configuration management operations provided by NETCONF
- Notifications that can be emitted by a device when certain events occur

The NETCONF protocol was originally targeted at devices such as routers or switches in provider backbone networks and large enter-

prise networks. In these environments, configuration changes are usually pushed to the devices by a management application that can initiate NETCONF sessions as needed. In large-scale measurement platforms, however, the MCs are behind a network address translator and hence must establish the communication session and (periodically) pull their configuration from the controller. The usage of NETCONF therefore requires the provision of a "call home" mechanism allowing devices to initiate the establishment of a NETCONF session. While this is currently not supported by the standardized NETCONF transports, it seems relatively easy to add this feature to the NETCONF over TLS transport: the NETCONF client is configured with a schedule indicating when to establish a TCP connection to a NETCONF configuration server. The NETCONF server then acts as a TCP client establishing the TCP connection and as a TLS client establishing a TLS session. At this point, client and server roles are swapped, such that the MC takes the role of a NETCONF server and the controller takes the role of a NETCONF client, pushing any pending configuration changes to the device.

Consider the following example. Suppose a controller wants to request an MC with the IP address 192.0.2.1 (say) to perform a UDP latency test to a destination IP address 203.0.113.1, using source port 23677 and destination port 34567. The test is a singleton test performed at 08:00 UTC. The test is to be performed without cross-traffic, and the output type is raw. The use of NETCONF/YANG to send the information in this example is depicted in Fig. 3.

While NETCONF can, in principle, also be used to push measurement results to a collector, it seems that IPFIX is a much better fit for this task, as described below. The configuration parameters needed by an IPFIX exporter can easily be configured via NETCONF since there is already a standard IPFIX configuration data model [8], as described in the next section.

## AN IPFIX-BASED REPORT PROTOCOL

As mentioned earlier, the SamKnows platform uses HTTP with a proprietary protocol on top of it to convey test results data from the MCs to the Collector. In this section, we propose the use of the IETF standard IPFIX protocol for that purpose.

IPFIX [9] is a unidirectional, transport-independent export protocol for binary data records, with a focus on network measurement and operations applications. The structure of the data records is described in-band by Templates, which refer to Information Elements (IEs) from a common data model managed by the Internet Assigned Numbers Authority. The basic IEs cover most layer 3 and layer 4 measurement needs, and the information elements can be extended [10].

IPFIX organizes data records into messages. A message is a sequence of sets preceded by a message header, which, among other things, includes an observation domain ID (identifying where the records in the message were measured) and an export time (when the message was originally sent).

A set contains records preceded by a set header, which contains a set ID identifying the type of the records the set contains. Template sets, identified by a special set ID, contain templates, which are sequences of IE identifiers and lengths; these define the fields of the records they describe. A template's ID matches the set ID of the sets containing records described by the template. Since many records may be described by a single template, IPFIX's data representation is more efficient than those based on inline record structures (e.g., XML, JSON).

In IPFIX terminology [11], the MC encompasses both the metering process (MP) and the exporting process (EP), while the collector is the collecting process (CP). IPFIX is used between the EP/MC and the collector/CP.

We next explore how to use IPFIX to report measurement results by defining a Template.

Part of the information can be conveyed using the fields in the IPFIX header:

- Information about the MA: The MA identifier can be sent in the Observation Domain field of the IPFIX header.
- Information about the time of the report: The Export Time field can be used to convey this information.

The information describing the test is included in a template set that contains multiple IEs for each of the different pieces of information we need to convey. This includes:

- An identifier of the metric used for the test. In order to convey that we need to define a new IE, let us call it metricIdentifier.
- An identifier of the scheduling strategy used to perform the test. Again, this will be a new IE, called testSchedule.
- An identifier of the output format. A new IE, outputType, is needed.
- An identifier of the environment, notably, whether cross traffic was present during the execution of the test. A new IE is needed for this testEnvironment.
- The input parameters for the test. Most of these can be expressed using existing IEs, such as sourceIPv4Address, destinationIPv4Address, and so on.

The information describing the test results widely varies with each test, but can include the time each packet was sent and received, the number of sent and lost packets, and other information. Again, most of these can be expressed using existing IEs, and some new ones can be defined if needed for a particular test.

As an example, suppose an MC wants to report the result from the UDP latency test requested by the Controller in the previous section using IPFIX. The IPFIX report message for this test is depicted in Fig. 4.

## CONCLUSIONS

Network operators have many disparate technology and network management tools, but few that provide an overall assessment of what user experience might be like. Growing interest has led to a number of over-the-top measurement platforms. While these platforms all differ, we have shown that it would be possible to abstract common architectural components, and could share some

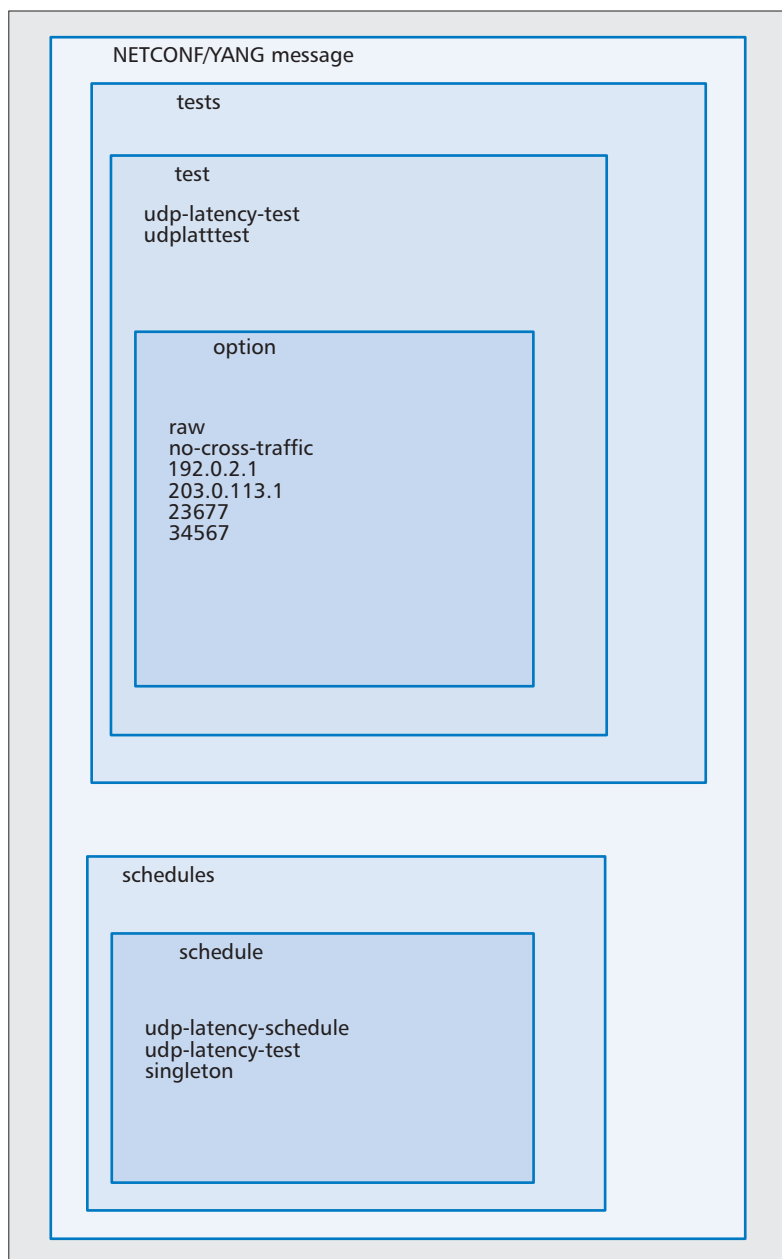


Figure 3. Example of a NETCONF/YANG control message for a UDP latency test.

common standard interfaces and data models that would enable a degree of interoperability.

Standardization discussions have already commenced in the IETF and Broadband Forum. In this article we have taken the SamKnows platform as a case study and have proposed an evolution of the platform toward standard protocols. In particular, we have shown how this can be done by using NETCONF and IPFIX for the control and report protocols, respectively. We have also presented two other platforms, RIPE Atlas and perfSONAR, to illustrate that several deployed platforms share a similar architecture, which hints that it would be feasible for the different platforms to adopt the proposed standard solution.

While traditional network management tools are applied across the breadth of the network, the emerging end-user premise tools are current-

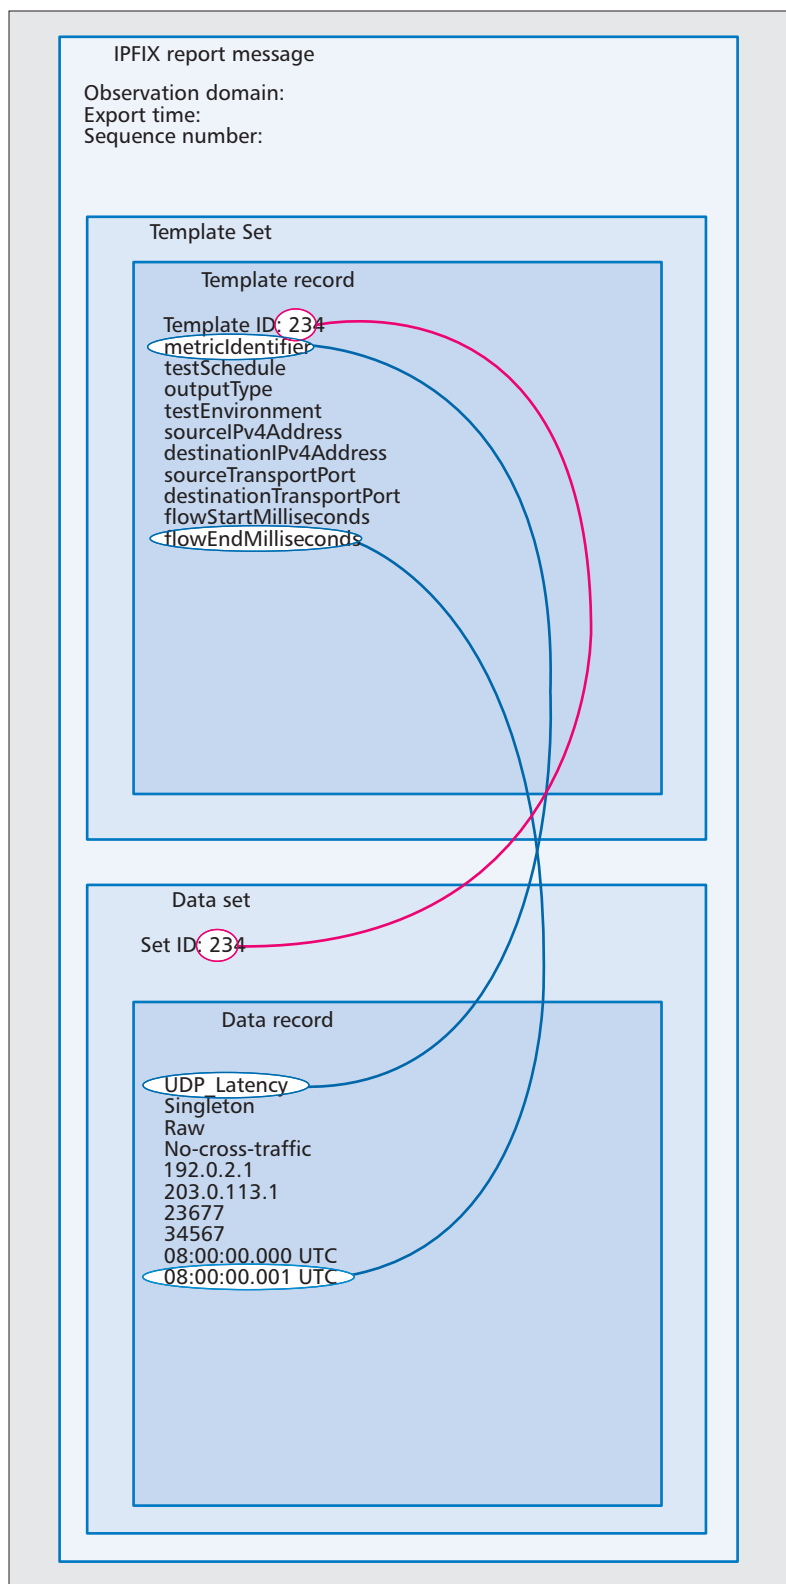


Figure 4. Example of an IPFIX report message for a UDP latency test.

ly limited to selective deployments on user devices and dedicated measurement boxes. Standardization will allow these approaches to break through into mainstream network and service management. We can imagine that network operators can embed measurement capabilities in a wide range of network and CPE devices

(e.g., home gateways) as well as internal network elements. All of these devices can be controlled by a single framework, and the measurement results can be collected to provide a comprehensive end-to-end view, as well as between known network locations. Tests can be operated continuously across all lines or on a randomly selected subset of lines for purposes such as capacity planning and network design. Problems can be investigated through adapting the test configurations, schedules, and selection of lines tested. This can help identify problems in the network, and with equipment or suppliers, and to isolate whether the issue is in the shared part of the network, a vendor hardware problem affecting many users, unique to a single user line, in the home network or an over-the-top service. Comparable data can also be shared between horizontal or vertically arranged network operators, with service providers, and with other parties such as regulators.

#### ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for helping us improve this article. We would also like to thank the Guest Editors, P. Calyam, C. Dovrolis, R. Kettimuthu, B. Tierney, J. Zurawski, and L. Jorgenson.

M. Bagnulo, P. Eardley, T. Burbridge, S. Crawford, and J. Schoenwaelder are partially funded by the EU FP7 Leone project. B. Trammell is partially funded by the EU FP7 mPlane project.

#### REFERENCES

- [1] C. Kreibich *et al.*, "Netalyzr: Illuminating The Edge Network" *Internet Measurement Conf. (IMC)*, Melbourne, Australia, 2010.
- [2] S. Shalunov *et al.*, "A One-Way Active Measurement Protocol (OWAMP)," IETF RFC 4656, Sept. 2006.
- [3] A. Hanemann *et al.*, "PerfSONAR: A Service Oriented Architecture for MultiDomain Network Monitoring," *Proc. 3rd Int'l. Conf. Service Oriented Computing*, Springer Verlag, LNCS 3826, Dec. 2005.
- [4] BBF TR-069, "CPE WAN Management Protocol v1.1," Issue 1, Amendment 2, Broadband Forum, Dec. 2007.
- [5] G. Almes, S. Kalidindi, and M. Zekauskas, "A Round-Trip Delay Metric for IPPM," IETF RFC 2681, Sept. 1999.
- [6] A. Morton, "Round-Trip Packet Loss Metrics," RFC 6673, Aug. 2012.
- [7] M. Bagnulo *et al.*, "A Registry for Commonly Used Metrics, Independent Registries," draft-bagnulo-ippm-new-registry-independent-00 (work in progress), Jan. 2013.
- [8] J. Schönwälder *et al.*, "Network Configuration Management Using NETCONF and YANG," *IEEE Commun. Mag.*, vol. 48, no. 9, Sept. 2010.
- [9] G. Muenz, B. Claise, and P. Aitken, "Configuration Data Model for the IP Flow Information Export (IPFIX) and Packet Sampling (PSAMP) Protocols," IETF RFC 6728, Oct. 2012.
- [10] B. Trammell and E. Boschi, "An Introduction to IP Flow Information Export (IPFIX)," *IEEE Commun. Mag.*, vol. 49, no. 4, Apr. 2011.
- [11] B. Trammell and B. Claise, "Guidelines for Authors and Reviewers of IPFIX Information Elements," draft-ietf-ipfix-ie-doctors-07 (work in progress), Oct. 2012.
- [12] G. Sadasivan *et al.*, "Architecture for IP Flow Information Export," IETF RFC 5470, Mar. 2009.

#### BIOGRAPHIES

PHILIP EARDLEY works in BT's Research & Innovation division on ways to measure and improve the quality of Broadband networks. He is an active member of IETF, where he contributes to the LMAP working group and co-chairs the Multipath TCP working group. He leads Leone, and previously initiated and led Trilogy, both EU FP7 collaborative research projects. He is co-author of the book *IP for 3G*.



---

JÜRGEN SCHÖNWÄLDER is the head of the Computer Networks and Distributed Systems (CNDS) research group at Jacobs University Bremen. His research interests are network management, distributed systems, wireless sensor networks, and network security. He is an active member of IETF, where he has edited more than 30 network management related specifications and standards. He is currently participating in the Leone and Flamingo EU FP7 projects.

BRIAN TRAMMELL ([trammell@tik.ee.ethz.ch](mailto:trammell@tik.ee.ethz.ch)) is a senior researcher in ETH Zurich's Communication Systems Group, where his primary research interests are Internet measurement, performance, and security. He is a member of the Internet Architecture Board, co-chair of the IETF's IP Performance Metrics Working Group, and an active contributor to measurement-related standards in the IETF.

TREVOR BURBRIDGE is a principal researcher in the Network Innovation Practice of BT Research, working on ways to both measure and improve the quality of broadband networks. Previous areas of research have included RFID and sensor networks, the privacy and security of pervasive computing, overlay routing and discovery networks, and application integration and

messaging. He is the co-author of several contributions to the IETF LMAP Working Group and an active member of the EU FP7 Leone project on large-scale network measurement.

SAM CRAWFORD is the founder and technical director of SamKnows Ltd, a London-based company providing measurement services to regulators and ISPs around the world. Since starting in 2009, the company's client list has grown to include the American, British, Brazilian, Singaporean, and Canadian telecom regulators, as well as the European Commission and many large ISPs. They currently have over 60,000 probes deployed across more than 40 countries. He is particularly interested in probe-based active measurements and how these can be evolved to better relate them to real-world QoE.

MARCELO BAGNULO ([marcelo@it.uc3m.es](mailto:marcelo@it.uc3m.es)) received an electrical engineering degree from the Universidad de la República Oriental del Uruguay and a Ph.D. degree in telecommunications from Universidad Carlos III de Madrid. He is an associate professor in the Telematics Department at Universidad Carlos III de Madrid. He is a co-author of over a dozen IETF RFCs and over 50 academic papers.