Call Identifier: SMART 2012/0046

Full title: Study on European Internet Traffic: Monitoring Tools and Analysis

Deliverable D5: Second Experts' Workshop Report

Date of preparation: Version 1.0, May 2015

List of participant organisations:

Role	Participant name	Participant short name	Country
Tenderer	Hebrew University of Jerusalem	HUJI	Israel
Supporting partner	Alcatel Lucent Bell N.V.	ALB	Belgium
Supporting partner	Martel GmbH	MAR	Switzerland
Supporting partner	Universidad Autónoma de Madrid	UAM	Spain
Supporting partner	Universite Pierre et Marie Curie – Paris 6	UPMC	France

Editor of this deliverable: Professor Scott Kirkpatrick
Organisation: Hebrew University of Jerusalem
e-mail: kirk@cs.huji.ac.il

TABLE OF CONTENTS

1. SECOND EXPERT WORKSHOP: OVERVIEW AND HIGHLIGHTS

1.1 Overview

This document reports on the 2nd Experts' Workshop organized within the EC-funded "Study on European Internet Traffic: Monitoring Tools and Analysis".

The Internet is an important critical infrastructure, but efforts to monitor this complex system have been diverse and uncoordinated. This study analyses existing Internet monitoring tools and methodologies and provides concrete recommendations about the needs and the next steps that Europe should take in this area.

The outcomes of this study will be:

1. An up-to-date, and as-complete-as-possible cartography of existing monitoring tools and methodologies.
2. A gap analysis of the needs for new methods and tools, taking into account how the Internet is evolving today and considering future Internet design and policy directions. The gap analysis can point both to possible new tools and methods as well as innovative ways to use current tools and methods.
3. A proof-of-concept showcase for the tangible ways some of those tools and methods can be used with real data.
4. A set of recommendations on how to close the gaps that have been identified, and suggestions for mechanisms that could support useful Internet monitoring for stakeholders in Europe.

We are gathering information for the study through two workshops as well as a questionnaire. We will present our results on April 22, 2015, as a satellite workshop preceding the Seventh International Traffic Monitoring and Analysis workshop (TMA2015) in Barcelona, Spain. The final workshop will be jointly announced by the EC and by TMA2015.

Here we present the information gathered at - and the conclusions of - our second expert workshop, held on May 19 and 20, 2014, in Brussels. In this meeting we reviewed with our Expert Council members the draft version of our results in the measurement study to date, indicated the areas in which we are focusing our gap analysis, and discussed potential use case prototypes that we will develop for our final public workshop in 2015. We also heard from and discussed several areas that were not completely addressed in our first workshop, held in October 2013. After listing the attendees and describing the agenda of the workshop, we provide in this deliverable a summary of the input obtained from our two days of discussion. Material presented as slides at the meeting are included in two appendices to this document, available for downloading on the study workshop website, at http://internet-monitoring-study.eu .

1.2 Objectives of this workshop

In our first workshop we presented the initial version of our survey questionnaire aimed at research and testing infrastructure owners and operators, and obtained feedback. We also received offers to help in targeting the appropriate recipients for these questions, and incorporated these suggestions and contacts into our process.

In this second workshop we reviewed the almost completed survey with the experts present and heard several presentations of the overall view and major issues facing network owners and regulators.

The workshop then focused on the gap analysis, the - as yet - unmet needs seen in our discussions and by our experts. This required a reality-based status from the survey and agreed-upon objectives, which we discussed. The discussion focused on the likely challenges that will be addressed, covering “telescopes,” or large scale behaviour, and “microscopes,” which expose detailed traffic flow information. On the large scale, identifying growth rates and trends, and rapid identification of distributed anomalies (which might be attacks) are likely key questions. On the finer scale, characterizing flows, in terms of user intentions and needs, was discussed. We plan to present a vision incorporating both types of observation into an Internet Observatory. Finally, this part of the study can shed some light on how the FIRE (Future Internet Research and Experimentation) efforts on federation might work to integrate measurement systems in both the research experimentation and commercial domains.

Section 2 of this report details the process followed to get maximum value from our wide-ranging inputs. In subsection 2.1, we provide the instructions sent to each participant and the agenda we followed (at least approximately). In subsection 2.2, we list the experts who participated and others consulted who gave us valuable advice. This was an interesting group. We brought together measurement experts and a senior technical advisor to a major regulator, the US FCC. We heard the practices of several major European telecom operators and of an independent commercial measurement group. Finally, we discussed measurements from the edge of the Internet, by active means, in greater detail than was possible in our first meeting.

Section 3 outlines the content of the meeting. We adopt an outline-like style for this preliminary report, basing the report on detailed notes which the SMART measurement team drafted at the conclusion of the meeting, while our impressions were freshest. Several of the participants brought detailed presentations of their tools or of their practices, in addition to the discussion points that were requested in their invitations.

In reviewing potentially useful outcomes of the study during our initiation meeting, we identified three of particular value. First, identifying new sources of information beyond the classic active and passive measurement tools. Second, finding, or creating, new opportunities for data sharing to make it possible to manage Internet performance on a wider scale, while managing privacy concerns and exposures. Finally, as we survey the tools and analytical techniques in current use, we anticipate identifying opportunities to employ modern “big data” centered machine learning tools on Internet traffic data, to see activities previously hidden in the noise, or to identify causes and development paths for complex behaviours. Each of these topics arose in the course of our first experts’ workshop, and are evident as requirements that were discussed in this second workshop. We provide a Section 4 summarizing these requirements and their sources.

2. WORKSHOP ARRANGEMENTS

2.1 Agenda

The following Agenda was distributed before the meeting:

SMART Internet Monitoring Study

Expert Council Workshop 2

Brussels, 19-20 May 2014

Website: http://internet-monitoring-study.eu/index.php/workshops/workshop-1

Venue: Meeting Room 0/54, Offices of the European Commission,

Avenue de Beaulieu 33 (BU33), 1160 Brussels, Belgium

Monday, 19 May, 2014

09:00: Coffee and introductions

09:45: Kickoff: FIRE unit and workshop objectives

- Georgios Tselentis, European Commission
- Scott Kirkpatrick, HUJI

10:00: Survey of Tools and Methodologies

Presenting our Deliverable: Jorge Lopez deVergara-Mendez, Universidad Autonoma de Madrid

Questions -- Is this the state of the art? What is missing? Which other questions should we be asking? Who requires network measurement data? How do they use it?

Views and Discussion – the big picture:

- Henning Schulzrinne, FCC CTO and Columbia University Professor
- Maria Teresa Herrera Zamorana, Telefonica
- Bert Wijnen, RIPE NCC

Scribe: Martin Potts

13:00 – 14:30: Lunch

14:30 - 17:30: Gap Analysis

Presenter: Eunsook Eunah Kim, Martel

Views and discussion: Unmet Needs

- Peter Heinzmann, CNLab, Switzerland

- Marcin Pilarski, Orange Poland and Warsaw Technical University
- Boris Banjanin, MG-Soft Corp

Specific questions: What additional data would regulators, consumers, and carriers use if they could have it? In what ways is the landscape changing (new technologies, new players, new regulatory initiatives, etc.)? Who else should we be speaking with? How do the needs of mobile and wireless stretch our measurement and monitoring capabilities?

Scribe: Jerker Wilander

19:30: Hosted dinner for the group, Aux Armes de Bruxelles

Tuesday, 20 May

09:00 – 12:00: Next steps for the Infrastructure Community; Use Cases and Requirements

Views and discussion -- new modes and methods:

- Yuval Shavitt, Tel Aviv University
- Arjuna Sathiaseelan, Cambridge University
- Walter deDonato, University de Napoli
- Antonio Skarmeta Gomez, Universidad de Murcia

Scribe: Jorge Lopez de Vergara Mendez, UAM

12:00 – 13:30: Lunch

14:00 – 15:00: Working meeting

Next steps to finalize our survey document and draft an initial Gap Analysis, collect and organize this workshop's many inputs.

Scribe: Scott Kirkpatrick

3. WORKSHOP INPUTS

3.1 FIRE unit and workshop objectives

The declared meeting objectives were to review with our Expert Council members the draft version of our results in the Internet Measurement Study to date, indicate the areas in which we are focusing our gap analysis, and propose a set of Use Case prototypes that we will develop for our final public workshop in 2015. We also need to hear from, and discuss, several areas that were not completely addressed in our first workshop, held in October 2013.

Georgios Tselentis, European Commission, explained the overall objectives as being to identify:

- Missing tools and methodologies for large-scale monitoring and measurement
- Potential users for the methodologies and measurements (possibly combinations of them)
- Recommendations (that could be used in future EC Workprogrammes and towards regulation and policies)

He reminded that the European stakeholders are, by definition, fragmented.

Scott Kirkpatrick, HUJI, presented the project timeline:

- 1st Workshop: October 2013
- 2nd Workshop: May 2014
- 3rd Workshop: April 2015

He explained that the key outputs are:

- 3 post-Workshop reports
- Interim study report, focusing on Task 1: Survey of the state of the art
- Interim study focusing on Task 2: Gap analysis
- Technical reports focusing on Task 3: Use Cases
- Final report, due April 2015

He informed the audience that we are approximately half-way through the study and, whilst we are still gathering information, the gap analysis has already started.

3.2 Session 1: Survey of Tools and Methodologies

Presenter: Jorge Lopez de Vergara-Mendez, Universidad Autonoma de Madrid

Jorge presented the questionnaire that has been distributed (revised after the 1st Workshop) and the results received so far. He added that the questionnaire is also in the EU Survey platform. The questions include:

- The type of operator (public, private, access under agreement)
- The users of the network measurement data (own use, end user, researchers, other companies, administrators)

- Number of nodes (10s to 1000s)
- Scope (regional, country, continent, world)
- Data format (standard (open) or proprietary (closed, commercial))

It was discussed if the effort of abiding by standards was worthwhile. Some comments made were:

- Using a standard data format (eg. xml) should require little effort
- Common measurement definitions should be used (eg. for latency, sampling rates, RTT definition, ports used, does the measurement use ICMP, UDP, TCP, ….)
- For example, in the FCC's MBA (Measuring Broadband America) Verizon AT&T should at least measure the same thing and it should be possible to compare the values (measurement timescales, are the long-tail measurements considered, …) and to justify any claims made by the companies
- The use of a standardised format for the results might also be useful when collecting historical data (though measurement software will inevitably be updated in the meantime and, therefore, results may not be comparable)
- Parameters should, at least, be well-documented

Jorge explained that we only have 19 responses so far, but these represent a good spread of countries and show that the questions can be answered by organisations that run a large measurement infrastructure.

As well as questionnaires, Jorge reported that the project had also made interviews (8 done so far). Some of the comments collected from the interviews so far are:

- Lack of measurement interoperability and the fact that data cannot be combined easily
- Lack of ability to transfer data between European regulators
- Mobile networks are even harder to monitor, both in L3 & L4 and in QoE
- Network virtualisation complicates the matter further
- There is no compatibility between tools from different vendors
- It is challenging to monitor high speeds (>100Gps)
- There are no tools for monitoring QoE
- Regulators and End-users do not have the skills or the tools to monitor net neutrality (Scott suggested that the tool "Glasnost" can do this)
- End users do not know how to test their SLA

Jorge raised a few questions for the audience:

- Do the questions in the questionnaire represent the state of the art?
- What is missing?
- Which other questions should we ask?
- Who requires network measurement data?
- How do they use it?

3.2.1 Henning Schulzrinne, FCC CTO and Columbia University

Henning described the scope of the American market, which he divided into 2 main areas:

- Edge to BIAS (Broadband Internet Access Service), comprising:
  - Transit providers & CDN providers
  - IXP (border of the access to the customer)
- BIAS to customer, consisting of:
  - Many (1,000s) of very small providers – serving 20% of US customers
  - A few large providers – serving 80% of customers:
  - Wired: 6 cable operators, 6 traditional telcos,
  - Wireless: 4 wireless providers (some of whom are also wireline providers)

The Access network technologies are: Coax, DSL, HFC and FTTH.

Coax / HFC: As customers are offered increasingly higher bandwidth, there is less sharing (typically 30 customers per coax "cell"), so there is little congestion and less scope for operators to apply traffic preference/manipulation (eg. queueing).

Traditional DSL: 3-4Mbps (characterised by long lines to the nearest exchange and little fibre). The MBA survey shows that there can be congestion on the backhaul ("middle mile") since, due to the fact that there is little competition, operators have not upgraded their networks sufficiently as demand has increased.

*High-speed DSL (*20-40Mbps) in non-rural areas.
20% of homes have fibre available (though not all subscribe). There is some overlap with cable offering in densely populated areas.

The expansion of **fibre** build-outs is in:

- Rural areas (with government infrastructure support funding).
- Rural Electric Collectives (these companies have poles (*rights of way*), trucks, staff, existing customer relationships. The organisations are also not-for-profit, so can recover costs over a longer period. They do not install Internet over power line, but may deploy fibre to a nearby pole, then provide wireless access to homes (various licensed or non-licensed frequency bands). This can be a competitive option for customers in outlying locations with poor DSL capabilities.

Satellite can offer 12Mbps downstream, but the longer latency is a problem for some services, and usually a restrictive capacity cap per month is applied. Nevertheless, there are 1 million satellite customers (mainly in areas where there is no other coverage, or only low-grade DSL).

Henning informed that the amount of fibre in the US Access Network has increased dramatically and has been pushed deeper into the network. P2P (BitTorrent) services are generally "controlled" by capacity limit caps (typically 300-500GB per month).

With dedicated DSL Access lines, or shared Cable TV connections (eg. 10Gbps per 30 households), there is normally little congestion. Where congestion does occur, it is usually at:

1. *Middle mile* (DSL networks), where the local ISP has to rely on (eg.) AT&T for middle

mile connectivity and it is expensive, and

2. *Peering connections* (congestion here is artificial). If there is no competition, then providers don't want to upgrade the capacity. They want to be paid by the transit providers (or Netflix) for carrying their traffic with a high QoS.

Peter Heinzmann explained that the situation is not the same in Europe. In Switzerland, for example, congestion is seen in the cable network. Yuval Shavitt added that the congestion experienced in Israel was on the links into/out of the country.

Other US facts:

- About 80% of the Internet traffic is video (most is domestic).
- Complaints about the Quality of Experience are almost exclusively about content which originates and terminates within the country.
- It is being discussed whether to continue to charge on a flat rate basis, or per GB.
- There is not much competition in either Access or Transit.
- The Access provider is not necessarily the same as the Transit provider. However, cable companies now also offer long-distance Transit and content providers connect directly to Access providers. Relatively cheap Internet Exchange points are emerging (eg. Comcast Equinix), which are neutral places where dozens of carriers connect). All Netflix content is stored very close to the regional Access providers. Content Providers who are too small to run their own CDN use Akamai, Level 3, etc.

3.2.1.1 FCC Open Internet NPRM (May2014)

By "Open Internet" is meant:

- Transparency (disclosure of network management practices)
- No blocking
- No unreasonable discrimination

However, recent court decisions upheld Verizon's challenges to the latter two items, on the grounds that the Internet was defined in the 1986 telecommunications act as an *information service* not a *communication service*. This definition was done deliberately by the politicians at the time, in order to make the Internet more competitive.

Communications services have an obligation to be neutral (they may charge by weight or speed or bulk, but cannot be unreasonably discriminating among their customers).

However, if ISPs are allowed to prioritise video traffic from one organisation (for a fee), then it is likely that other traffic will suffer. It is also not clear what business arrangements would be allowed.

The FCC in May 2014 has issued a Notice of Proposed Rule Making (NPRM), which seeks comments on the 1996 Telecommunications Act Section 706 or Title II (general telecoms regulation), Sections 201 and 202, which deal with the aspects of ISPs being *just and reasonable* and *non-discriminatory*. The public release which accompanies the NPRM asks the question "should prioritisation be banned outright", though the term *prioritisation* is not defined (could be blocking, throttling, adding delay, …)

Comments are invited by September 10th, 2014.

Henning commented that the legal framework may be simpler in other countries. Business arrangements for IP services (including your own IP services) in Europe (and Asia) are made according to transparent SLAs. This does not prevent that specialised services can be given paid prioritisation, but this has to be done in a non-discriminatory way. Nevertheless, moving classical services to IP services generally requires some sort of QoS mechanism in order to protect low-bandwidth traffic which is sensitive to packet loss or delay from high-speed video. This is a global problem.

3.2.2 Maria Teresa Herrera Zamorana, Telefonica

Telefonica collects data from 250,000 *elements* every 15 mins (where elements include everything with limited resources, such as: nodes, links, radio cells). The data is analysed to determine how close they are operating to their limits. When measurement thresholds are reached, new equipment is ordered, or the network is re-designed / re-dimensioned.

She added that observations show that there are less fluctuations in measurement results at points where traffic is aggregated; variations are worse closer to the customer. She explained that the goal is to make the network engineering process *predictable*, rather than making a *best guess* and then measuring and reacting. Congestion experienced by users is also often due to too few content servers being installed to serve all the download demand, rather than too few network resources.

Yuval Shavitt commented that it has been shown, in Israel, that caches can save 50% of traffic.

With reference to a book by Warren Weaver, she considered that networking problems fall into one of three categories:]

- Simple problems (can be solved with analytical formula)
- Disorganised complexity problems (can be solved using statistics)
- Organised complexity problems (do not have solutions)

Voice traffic belongs to the first category: each customer has the same demand (homogeneous) for a 4 Khz channel, the traffic follows a Poisson law, no memory is needed, the system is linear.

For data, we have to apply heuristics (time-series analysis in the aggregated levels). It is possible to predict traffic patterns, but it needs much more data (especially in the radio access, also since the number of customers in a cell can change). Latency, bandwidth, jitter, delay, packet loss have an impact on QoE.

Maria Teresa also referred to Cooper's Law, which defines how wireless capacity has grown 1 million fold in the last 50 years. However, the following figures show that this growth has been mainly through an increase in the number of cells:

- Efficiency (has increased 20 fold)
- Spectrum (has increased 25 fold)
- Number of cells (has increased 2,000 fold)

She informed that Sony mobile handsets collect data to measure QoS and that operators also install the tool "Carrier IQ", so that they can measure details of phone coverage and performance (weak signal coverage, location, battery running out). Carrier IQ is installed on EVERY new phone, but - to save the battery - only takes measurements when a call is being made. However, potentially one phone is making a call from every cell at any particular time.

Another way to predict future traffic patterns is to analyse the published penetration of apps (eg. Whatsapp) and extrapolate their usage characteristics. She expected that the growth in IoT applications would have an impact on Traffic Engineering.

3.2.3 Bert Wijnen, RIPE NCC

Bert explained that the goal of **RIPE Atlas** tests is to continuously monitor the "overall health" of the Internet and to detect problems. RIPE Atlas probe measurements can detect cable cuts, or problems in individual ISP networks (IP prefix hijacks). Being under the control of its members, RIPE Atlas measurements are made according to whatever the RIPE community wants.

Standard Atlas tests are: Ping, Traceroute, DNS, SSL. RIPE Atlas probes do not measure bandwidth (Sam Knows probes do). By making Ping tests, it is possible to compare network performance in Europe with (for example) Africa.

He said that there are 5'000 Atlas probes installed, whereas Sam Knows has 40'000. He added that Sam Knows uses reference servers, whereas Atlas measures actual traffic, so can see real user experience. Atlas probes are making Traceroutes all the time (they now have a 2-year history). In a separate effort, RIPE NCC has collected routing table information from all over the world (they have a 15-year history).

Apart from RIPE NCC's own measurements, the probes are also used by users to define their own tests. For example, a user can ping his/her own Webserver from 1000 probes to test its reachability and response time (the test takes about 5 mins.). Making such measurements costs a user *credits* and hosting a probe gives a user *credits*. Anchor probes (50) around the world are located near the core of the network and hosting an Anchor probe gets many more *credits*. If a proposed measurement is interesting for the RIPE community, then (if RIPE can publish the measurements) they will give the necessary *credits*. *Credits* can also be transferred between users. Users have a form-filling GUI to help them set up a measurement (Traceroute Ping, target (Domain name/IP address), start time, end time, maximum bandwidth consumed by the test, …). More details can be found at: https://atlas.ripe.net/doc/udm and https://atlas.ripe.net/doc/measurement-creation/api

The RIPE Atlas team ultimately decides what tests are offered through the GUI.

The results are collected in a Hadoop cluster and the user has another GUI to download the data by day, week, at the end of the test, etc. Users can indicate if they want to keep the data from their measurements private or made public.

Bert also explained about the **RIPEstat** database. This provides information about the ownership of every IP address and AS number. With this information and RIPE Atlas measurements, ISPs can compare the performance of their prefix with others in the same country.

Some RIPE success stories are:

- IXP: Measuring the effect of installing L-root in Belgrade/SOX
- DNS: Looking for most popular instances of .FR anycast servers
- Events: Measuring Internet outage in Sudan

3.3 Session 2: Gap Analysis

3.3.1 Eunah Kim, Martel

Standardization work on all fronts is ongoing. Relevant work on QoS and QoE are found in most standardization organisations. ITU and ETSI perform similar work on QoS and QoE. 3GPP and IEEE orient work towards mobile QoS and QoE. Other efforts are found in the Broadband Forum (BBF) and the Grid Forum.

There is no strong agreement on the definition of QoS and QoE.

Discussion in the audience put a rather low value to the standardization efforts. "Standards are very good since there are so many to choose from". It takes too much time, with too many separate efforts that fail to provide a common definition. There should rather be work on measurement tools especially for the consumer that bypasses both standardization and regulation bodies.

There is a need to divide by application type. KPIs will be required for different applications.
Ofcom UK is attempting to define easily understood (for a consumer) labels to characterize usage requirements for different applications. The Leone project is working in this area.

3.3.2 Peter Heinzmann, CNLab

A gap that needs to be filled is a simple measurement that is easily understood. This might be difficult and it might be different for different access technologies. Average performance is not always a good measure.

Cable operators (FCC opinion) sometimes suffer from problems on segments that are shared by many users. In the US, the cable industry has a very low user confidence (even lower than the airline industry).

Reliability and performance figures and measurements are missing.

The problem of advertised speed and throttling is not as simple as first viewed. Shaping of speed is going to create further problems since TCP will adjust. The throttling will lead to delay increases. Buffer Bloat problems are also related to this.

Measurements from a home depend on the equipment at home. There will be a need to combine "probe" and user measurements.

We will need measurements on the backbone.

The stakeholders are:

- Operators/carriers
- Users
- Regulators
- Consumer organisations
- Media (cnlab clients)
- Companies – public and non for profit

Content distribution networks are important.

Connect Test Europe is a yearly test for the operators. This sometimes leads to a push on operators to make investments to ensure that any bad performance in one year is eliminated by the next test in the following year.

Bottlenecks in network performance are found in the home, the access network, servers, backbones, international content providers. This might vary according to the applications used. The BEREC application classification structure and requirements on networking performance parameters shows difficulties (see slides, BEREC report and SMAR-report).

The access network is often a bottleneck, but is not the only one. Bottlenecks are found all over; at home, 100Mbit equipment, Cable and 2/4 pin Ethernet cables. This makes fair measurements difficult.

Mobile is the most difficult, since it includes all of the problems found in fixed networks, as well as radio contention and signal strength.

3.3.3 Marcin Pilarski, Orange and Warsaw University of Technology

Marcin presented the results of the “Sniffer” experiment from the OneLab2 project.

This experiment observes and tracks the growth of various Storage Networks (Grids, Clouds, Content Delivery Networks, Information Centric Networks) by traffic intercept and pattern discovery, content servers discovery, and content server clustering.

He explained that the problems faced are that tracerouting is time-consuming and that, therefore, using caching is a better option. The geo-tagging database has deficiencies, for example, low accuracy for many addresses which leads to inaccurate clustering and poor visualization. Also, only URLs from Poland are intercepted, which means that the content coverage is limited.

He said that analysing measurements is difficult because ISPs don’t publish their structure, but it is known that Google aims to provide users with the best possible result by using their IP address.

One possibility to improve QoE is to place more content deep in the infrastructure where only the local caches are seen. For example, movies which are accessed three times in the last 48 hours are cached. Working with transparent caching and traffic engineering routing to support caching was discussed.

An issue is pricing and “what’s in it for me” in peering and hosting.

Finally, he said that Orange Poland wants to extend the RIPE infrastructure.

3.3.4 Boris Banjanin, MG-SOFT (partner of Leone)

Boris said that MG-Soft provides a “Net Inspector” network management tool which performs client/server fault and performance management, monitoring and management of system status and resources, intelligent alarm management system, distributed polling with load balancing, includes network discovery, IP SLA and NetFlow monitoring and flexible integration with other sources. Net Inspector is able to manage large volumes and they have many customers.

He added that Net Inspector has a 3-level threshold structure. They are currently working on improving the anomaly-detection, by using better statistics for automating the calculation of the threshold values used to determine if an alarm should be escalated and the ability to take into account the effects of busy hour traffic. It is well known that Internet traffic is not level across a day or week, for example.

The Leone network is relatively small (ca. 200 SamKnows probes), but SamKnows has given them larger datasets from BT's network to feed into the tool to test the ability of the algorithms to detect the root cause of problems.

Henning Schulzrinne commented that data is available from 8'000 SamKnows nodes in the US from 12 ISPs.

3.4 Session 3: Next steps for the Infrastructure Community; Use Cases and Requirements

3.4.1 Yuval Shavitt - DIMES

Yuval presented the DIMES measurement tool, which started in 2003, when measurements were done with instrumentation boxes, almost all of which were placed in universities. Having physical boxes meant, at that time, that it was difficult to deploy them in large numbers. Dimes changed from instrumentation boxes to software agents, which enabled large scale deployment and in the commercial Internet.

DIMES continued to develop thanks to the FP7 projects MOMENT and OneLab2 and certainly had an impact on other projects for large scale network measurements, such as Ono.

DIMES could not extend further once the EC funding ended, but nevertheless, data continues to be collected from 500 active agents every day, in 200 ASes. Approximately 2-3 million measurements are made daily. Users can get DIMES data for free, since the **measurement data is open.** Users usually ask for a day or a week of traceroutes or topologies. They define the conditions for the measurements, and then choose the groups of agents and targets, in a "click-and-select" manner. One of the latest developments was an interface to the measurement agents, so that they can be controlled to permit individualized experimentation, as is done in RIPE's and BISmark's hardware probes. It is, however, difficult for DIMES to provide any support to process the raw data and control the measurements, without funding.

Henning Schulzrinne commented that it was also not easy to get funding in the US at the moment. NSF funding is low, and not so much is given to computer networking; only really new topics get funded. Also, the US funding mechanism makes it difficult to have sustainable projects. In addition, projects that only produce data are not doing research, so would not be funded by the NSF.

It was questioned if was easier to use RIPE Atlas? Certainly there are more probes in RIPE Atlas (21'000) and with a web form it is also quite easy to set up a measurement experiment. On the other hand, users have to pay (credit system) to make experiments (unless the experiment is considered useful for the RIPE community).

Henning Schulzrinne asked how DIMES maintains its user population. He said that the FCC makes Press Releases to attract more people to host probes. Yuval said that they occasionally publish papers which generate a spike in the number of users.

The DIMES measurements have been used to build a map of discovered PoPs, using link delay times and network motifs. Approximately 4'400 PoPs have been discovered, with over 50'000 IP addresses within each one.

DIMES has also developed applications for mobile phones, which measure the performance of cellular networks and discover the routing infrastructure. This is being deployed by Telefonica. They can also send packet trains to do bandwidth measurements.

3.4.2 Arjuna Sathiaseelan - Public Access WiFi Services (PAWS)

Arjuna's presentation of Public Access WiFi Services (PAWS) was about the digital divide and how to enable Internet access in non-developed countries, even in Europe.

There is a strong correlation between affordability and the level and quality of broadband penetration.

The PAWS project, funded by the RCUK is led by Cambridge University, and includes partners from Nottingham University, BT, SamKnows, BISmark and Nottingham City Council. The aim of the project is to provide a free "Less than Best Effort (LBE)" access with no QoS. 20 users ("sharers") share their access line (2Mbps downstream and 512Kbps upstream) with other users. The sharer's Access Point is connected to a gateway via a VPN to avoid snooping from other sharers. The PAWS partners will monitor the traffic and see what people really need.

An area in Nottingham with low Internet access was chosen for the trial. 1'067 households have been selected in total: 23% from BT, 23% from Sky, 21% from Virgin and 33% of the APs had an unknown ISP. 61% of them are "FON enabled". The difference between FON and PAWS is that FON customers pay for the service.

20 PAWS routers have been installed in 20 sharers' homes. By sharing their connection, the sharers are given compensation. 8 of these routers are being used by 15 citizens, with one deployed in a public space. The 12 remaining routers are currently used as measurement points. There have been some recruitment problems.

Henning Schulzrinne said that in the US, there is a so-called "Internet Essentials" scheme: If you are below a certain income, you can get a home Internet access with a cable company for $10.

The basis for the measurements are BISmark metrics:

- Availability: checked with a UDP probe ("heartbeat") every minute.
- Throughput with netperf (3 parallel TCP connections) every 6 hours.
- Last mile latency (ping to first non-NATed IP, every 10 mins.).
- End to end RTT, to different servers, every 10 mins.
- Loss, with a Distributed Internet Traffic Generator (D-ITG).

It was observed that many routers were switched off most of the time. One sharer could not pay the bill, and switched to being a (free) user of PAWS.

Measurements showed that the available peak and off-peak throughput bandwidth changed appreciably.

It was concluded that sharing a 2Mbps downstream link in a fibre access network with a total capacity of (10Mbps upstream and 100Mbps downstream) is possible at both peak and off-peak times. However, sharing a 512Kbps in the upstream direction need QoS control functionality on the home APs.

The latency in the *last mile* was about 10ms (median) for fibre, but some reached more than 40ms. Packet loss was negligible (less than 0.7%).

Regarding the usage, the main source of traffic was advertisements (60% of the traffic for both PAWS users and sharers), probably because of free apps.
Henning Schulzrinne commented that the amount of traffic due to advertisements is not consistent with other traffic studies, and suggested that the reason for such a high level in this study could be the demographic group of people in the trial.

Lessons learnt:

- Urban areas (even deprived ones) have good network infrastructure.
- Fibre networks are ideal for network sharing. ADSL lines needs better QoS support as they have less overall capacity.
- Advertisements (at least in this example) are a main driver of traffic.

Issues:

- How to measure unused capacity? (to determine if - and when - sharing would be feasible).
- Usage caps and measurements. How do they work together?

The project will be making more deployments.

3.4.3 Antonio Gómez Skarmeta, Universidad de Murcia, Spain

Antonio said that there are many definitions and views regarding the Internet of Things, M2M, and O2O as well as many experiences in different environments. There are several issues in the deployment and measurement of IoT scenarios. The key points are that the devices are constrained in terms of power and there are many legacy devices that are non-IP.

He explained that the FP7 IoT6 project has defined an ecosystem of IoT, in the context of IPv6. It describes how sensor ecosystems extend beyond the local sensor network to the Internet Wide Area (including, for example, aggregators and virtual machines), making solutions much more complex to apply.

There are different layers, from the sensors themselves to the gateways, platform and applications. Scalability is important, since measurements may be made using a huge number of sensors.

In order to evaluate IoT solutions under real world situations, experimentation environments will be needed that allow the technical evaluation of IoT solutions under realistic conditions, the assessment of the social acceptance of new IoT solutions and the quantification of service usability and performance with end users in the loop.

Some approaches can be seen in the Plugtests from ETSI TS 103 104 (April 2013).

The ProbeIT project is also producing a taxonomy of parameters that have to be measured in IoT deployments.

Challenges:

- Scalability.
- The heterogeneity of devices (including legacy devices) and device technologies means that interoperability has to be supported at different layers.
- It is difficult to get statistically-significant.measurements when repeating the same experiment across testbeds, or even on a single testbed.
- IoT testbeds should be federated with other testbeds, for end to end experimentation.
- Concurrency of access to the infrastructure: horizontalization of infrastructure and pervasiveness of the IoT environment.
- Data-layer solutions experimentation, not just communication layer experimentation. Sensors are also mobile.
- Realism of the environment for the experimentation. Realistic mobility.
- Human users have to be included into the experimentation loop. Designing mechanisms that support adequate user involvement during experimentation.
- Mobility, and handling such mobility.
- Privacy and security.
- Making deployments in the real-world.

Related activities:

- ANA4IoT experiments and measurements in OpenLab: Future Internet architecture for the IoT. They have made several measurements of typical QoS parameters using either publicly available measurement tools, or ones developed themselves.
- Universidad de Murcia (UMU) has been using the SmartSantander FIRE facility. The number of sensors there is not huge, but is much larger than the NITOS FIRE testbed. They are using about 1’000 real sensors and UMU is simulating another 10’000.

Gaps:

- A way is needed to correlate information from the different sensors (similar to what was presented by MG-Soft, but for sensor networks).
- Tools are needed to automatically classify the type of sensors that are working in the network (if they are legacy, what interfaces are provided, etc.)

3.4.4 Walter de Donato, University de Napoli, Italy

Walter stated that large scale measurement platforms are necessary for studying residential Internet access networks. Among the several approaches and platforms are: BISmark (gateway based) and Hobbit (application based). These are two complementary platforms in one architecture that share common features of:

- Support for pre-existing (ie. already well-tested) measurement tools.
- Implicit management of measurements targets' resources. Avoid overloading targets (supports unmanaged existing services).
- Automatic remote upgrade/configuration. User intervention is required only when strictly necessary.
- Accessible real-time reports. Transparent results.

The specific feature differences between BISmark and Hobbit are:

BISmark:

- Customized OpenWRT
- Remote access to router console for troubleshooting
- Captive portal-based one-time device registration
- Monitoring of gateway health
- Cross traffic-aware measurements
- Opt-in passive measurements
- Mutual exclusive scheduling algorithm for heavy measurements

Hobbit:

- Multiplatform client based on Qt libraries and bash/awk ports
- Identification of connection ISP and service plan details
- Possibility to temporarily suspend the measurements
- Flexible measurements (when, which tool, how to run it)
- Controlled overlap scheduling algorithm for " heavy measurements"
- User aware of current activities

BISmark deployments are currently worldwide with 417 vantage points (141 active), 417 users and measured access networks, in 176 cities.

BISmark is being used in many studies: broadband performance, wireless performance, web performance, network usage, PAWS, etc.

Hobbit is exclusively in Italy with 489 vantage points, (57 active, 70% Windows, 12% MacOS, 18% Linux), 380 users, 1'665 networks, and in 341 cities.

Hobbit measurements are grouped geographically. 90% of the users use DSL. There is an Internet Performance Map for Italy available on their website: http://hobbit.comics.unina.it.

A management server configures the measurement experiments, collects all the data from the vantage points and performs troubleshooting.

Hobbit is used, for example, to determine to what extent ISPs offer the advertised performance, and if users chose the right service plan. Usually, the same performance could be obtained with a cheaper service plan (especially for high-end service plans).

BISmark measurement and tools are:

- Netperf
    - Upstream throughput (multiple TCP flows)

    - Downstream throughput (multiple TCP flows)
- Ping
    - Round-Trip Latency (ICMP)
- Distributed Internet Traffic Generator (D-ITG)
    - Round-trip Jitter
    - Round-Trip Packet Loss
- Shaperprobe
    - Upstream/downstream capacity
    - Upstream/downstream shape rate
- nslookup
    - DNS dely
    - DNS failure rate
- paris-traceroute
    - forward/reverse IP level path
- Netperf + ICMP
    - Round-trip latency under load

Hobbit measurement tools are:

- D-ITG for all:
    - Upstream, downstream throughput (both TCP and UDP), Round-trip latency (UDP),Round-trip jitter, Round-trip packet loss, and BitTorrent upstream and downstream. BitTorrent is emulated by generating traffic to well-known ports.

Lessons learnt:

- Gateway and application-based approaches have complementary approaches. They might cooperate to get more insights on performance.
- Encouraging end users to participate is challenging, while losing them is very easy; users blame the probe for any problem they experience.
- Form factor matters. Users often trust commodity hardware over custom hardware. For instance, a common router vs. a development platform.
- The duration of measurements is important for some metrics. Long-term throughput might be very different from short-term throughput. → This has affected SamKnows too.
- Proper scheduling of measurements is needed to enable scalability while managing overlap among measurements.
- Using fine granularity when storing results is a good practice.

Gaps:

- There is no standard way to query the ADSL modem about the sync speed they have achieved.
- Large scale and dense deployment of vantage points is needed to obtain more accurate insights on performance by geographical location and ISP.
- Cooperation among available platforms would improve performance analysis effectiveness.

- Access to technology-specific layer 2 parameters (e.g. DSL negotiated bitrate, signal attenuation, SNR, interleaving/fastpath, etc.) would be useful for tuning measurement tools and better interpretation of results.
- Layer 2 technology detection techniques, for enabling technology-aware measurement techniques. This information is not easy to obtain. SNMP is not available or open. ETSI TR-069 could help, but this is enabled only on the operator's side.

4. ISSUES AND REQUIREMENTS EMERGING FROM OUR WORK SO FAR

From our work so far, we can already conclude that there is a general consensus that more information is needed about performance and capacity, both deep in the Internet and delivered to the end users. The crowd-sourcing opportunities available with current large deployments of active measurement probes contain information that is not being fully used. Maria Teresa Herrera noted that this information constitutes “disorganized complexity”, so statistical methods must be used to extract it. However, networking experts have traditionally preferred to exactly identify sources of problems, so routing information may be required alongside performance measurements to more precisely identify the sources of problems.

An idea for an “Edge Use Case” tool that was proposed by Henning Schulzrinne would be one that integrated simple measurements (useful for regulators, ISPs and end users) into the home router (like what SamKnows and BT are purportedly doing in the BT Home Hub).

The minimum set of edge tests that would be sufficient to discover poorly-performing networks would be:

- Ping,
- Upstream /downstream speedtest,
- Delay to access to content (i) in the same ISP domain and (ii) from outside the ISP.

As noted in the first paragraph, it could be advantageous to add routing information to supplement the test results from “(ii)”, above.

He also expressed his support for the LMAP standardisation activities, which he saw as complementary.

However, standards approaches such as LMAP will not settle issues or result in interoperable measurement tools that work with different manufacturers’ component. It is not clear how this can be achieved (through the IPPM WG in IETF?), or if ad hoc approaches are more productive.

In Europe, there is reluctance by regulators to take transparency very far, yet CNLab’s results appear (in Switzerland) to have had a positive effect. The weaker performers in one year (after first complaining about the inaccuracy of the tests) have improved their results by upgrading their service by the time the tests are repeated and published again one year later.

Another open issue is the “business model” for user-defined testing. Can a stable sustainable system, comprising networks of active probes, be achieved? Is it RIPE’s “credits” model, or a system which is completely free to the user who performs tests? If it is completely free, where is the incentive to deploy and maintain it?

There is also a trade-off between running tests, such as the as the FCC’s mobile test package or OOKLA’s SpeedTest, sufficiently frequently (e.g. by running them automatically in the background) to give warning of problems, or an accurate overall network characterization vs the likely degradation of the user’s bandwidth. Similarly, can file transfer measurements be arranged so that they do not overload target sites, thereby compromising the measurements? Perhaps with wider adoption of such measurement tools, the workload required from any single observation point can be decreased, which will also decrease any interference with user workload.

The Internet of Things does not seem to have a measurement infrastructure to monitor it yet.

How will this develop, and what new approaches to scaling of solutions will be required? How can the public have better access to the IoT when it is present in their city?

Several technical limitations of present measurement tools have been identified. First is the lack of Layer 2 tools that can serve as “telescopes”, identifying problems outside the immediate domain. Second is the difficulty of measuring the ISP performance from the home without complicating the measurements by the details of the networking, wired or wireless, within the home. Moving all active measurements to run in firmware on the DSL-, or cable-, or fibre-interface modem seems the natural approach. However, although the Unix platforms found in such modems are reasonably standardized, the systems are not instrumented in standard ways and much useful information, such as the line speed at which the modem has synchronized, is not always available.

Our user surveys and interviews also provided a number of issues and requirements, some of which have been reinforced by the presentation and discussions in the 1st and 2nd Workshops. These include:

- There is a lack of measurement interoperability and data from different measurements cannot be combined easily.
- There is no compatibility between tools from different vendors.
- There is a lack of ability to transfer data between European regulators.
- Mobile networks are even harder to monitor, both in L3 & L4 and in QoE.
- Network virtualisation complicates the matter further.
- It is challenging to monitor high speeds (>100Gps).
- There are no tools for monitoring QoE.
- Regulators and end-users do not have the skills or the tools to monitor net neutrality.
- End users do not know how to test their SLA.

In summary, therefore, we have identified three areas in which there is promise for significant progress in the near future, and are targeting our use cases to expose some of the opportunities in these areas.

- Measurement interoperability within the scope of the equipment controlled by a single operator is a tough but achievable goal, and can be approached with an eye to meeting the still-developing standards of the LMAP group. Good visualization tools will be a key to broader use and sharing of such tools. Because the high backbone speeds and rapid response required to deal with problems are fundamental challenges, this use case is being developed in partnership with a major operator.
- The surprising success of the Swiss edge measurement efforts in publicizing comparative results and seeing an overall improvement in response should be publicized and extended to other countries. The SamKnows deployments already facilitated in collaboration with European regulators are a still tentative step in this direction across Europe more broadly.

  More work is needed to explore the information available in the variability of these data – are the average metrics capturing QoS to the whole population, or are there unhappy customers in the tails of the distributions? And can this measurement capability be extended to become a standard feature of all edge modems, whether cable, fibre, satellite, or twisted pair?

- Can active measurements from the edge of the Internet become telescopes, spotting possible issues deeper within the network as well as monitoring the quality of delivery over the "last mile"? And is the scale-up needed to encompass the Internet's edge at the mobile device or, in the future, the Internet of (every)Things even possible? The greater number of observers potentially available with an almost cost-free software observing client and a base in the smart phone community suggests a use case which will look at this potential. In particular, truly large scale measurement deployments offer the possibility at looking at short-lived problems or persistent problems with a characteristic diurnal temporal fingerprint from many different but simultaneous access paths. This is the most future-oriented of our use cases under current study.