05-11-2015 updated 11-04-2016

M7.1 Wireless Crowd Source Performance Measurement and Verification

Milestone M7.1

Contractual Date:	31-10-2015
Actual Date:	05-11-2015 updated 11-04-2016
Grant Agreement No.:	691567
Activity:	7/SA3
Task Item:	Task 3
Lead Partner:	SWITCH
Document Code:	GN4-1-15-201EE
Authors:	Anna Wilson (HEAnet), James Healy (DCU), Vasileios Kokkinos (CTI), Kostas Stamos (CTI), Arne Oslebo (UNINETT), Jiri Melnikov (CESNET), Kurt Baumann (SWITCH), Brian Mortensen (NORDUnet)

© GEANT Limited on behalf of the GN4 Phase 1 project.

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 691567 (GN4-1).

Abstract

This document proposes an architecture design using JavaScript and NetTest/Boomerang, on essential, highly frequented web sources. This kind of performance measurement and verification is known as "wireless crowd sourced performance measurement and verification".

Table of Contents

Table of Figures

Executive Summary

Demand for a wireless performance measurement and verification (PM&V) system has caused an urgent need for research on wireless campus networks, technologies, and standards (e.g. IEEE802.11x or IETF-LMAP) that are most suited to promote a better understanding of performance measurement and verification procedures (manual or automated).

A brief overview of GÉANT's ongoing activities shows a traditional approach to performance measurement and verification, looking primarily for infrastructure information, by implementing hardware (HW) probes on the wireless campus network and answering the question "WHAT have we done?", but not giving any statements of HOW something was done (see Section 2.1.5). Therefore, the concept of performance verification becomes more and more important for two reasons:

- Elaboration of problem solving procedures to ensure an appropriate response when faced with performance issues.
- Strategically, by introducing short-, medium- and long-term improvements to the wireless network topology.

In response to end-user feedback, this document proposes an architecture design using JavaScript and NetTest/Boomerang, on essential, highly frequented web sources. This kind of performance measurement and verification is called "wireless crowd sourced performance measurement and verification". As the term implies, the accuracy of our measurements largely depends on the number of "complete" data sets. From a statistical point of view, the large amount of measured data (end-user activities) allows us to analyse the performance behaviour on a wireless network in a campus/conference environment.

Section 2 shows where the PM&V is embedded into the GÉANT Network and the missing elements are summarised in a conclusion, which defines our focus of a wireless, crowd-sourced PM&V.

In Section 3, a generic PM&V with its prerequisites/requirements and architecture elements, will be defined, and its architectural building blocks described.

In Section 4, the generic PM&V architecture is transferred into demonstrations, which shows real use cases with aspects of lessons learned.

Finally, Section 5 looks at recommendation for future steps. Three categories are in focus here: measurement verification, service automation, GUI development, and app deployments for mobile devices/smartphones.

The original version of this document detailed the status of GN4-1-SA3 T3 WiFiMon activities as at Q2. This updated version, published in the last month of the project, contains new sections on the WiFiMon testbed (Section 3.6) and HEAnet workshop (Section 4.5 and Appendix E).

1 Introduction

Crowd sourcing can be used to introduce new or more developed skill sets, or a larger work force to achieve some specific goals, such as the performance behaviour of an NREN's wireless campus network. The use of crowd-source-based network performance measurement services and technologies is set to continually increase among the NRENs in the near future. This creates a need to understand the behaviour of network performance issues, their localisation and verification. The outcome of these investigations will support the NRENs' NOC, help network administrators to optimise and improve their wireless network topology/footprint, and user satisfaction.

The mobile user access network – supporting smartphones, tablets or other portable devices, such as laptops that require internet access via WLAN, Wi-Fi campus network, or cellular network – connects users to their applications. Today, eduroam [EDUROAM] and wireless networks specifically designed for campus use allow academic users and users from the private sector to traverse their daily workload on laptops and on powerful smartphone devices [MOBILEWORLD]. Increased mobility is a huge growth of the user-population. A nearly tenfold increase of data traffic over 2014, means that by 2019, 24.3 exabyte per month is expected [MOBILEDATA], which implies the need to enable QoS, or OLAs for network performance measurement, verification procedures and visualisation of measured data.

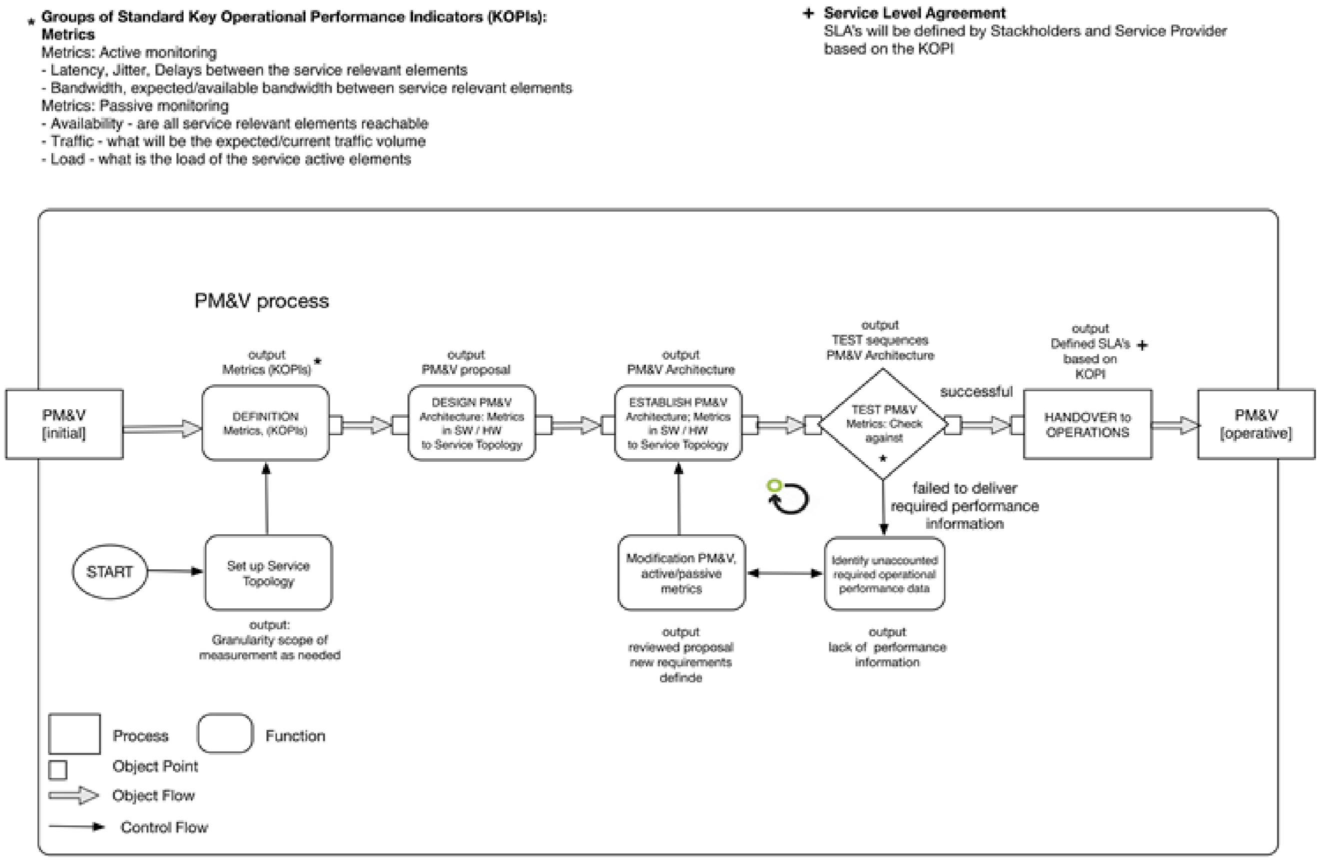
* Groups of Standard Key Operational Performance Indicators (KOPIs):
Metrics
Metrics: Active monitoring
- Latency, Jitter, Delays between the service relevant elements
- Bandwidth, expected/available bandwidth between service relevant elements
Metrics: Passive monitoring
- Availability - are all service relevant elements reachable
- Traffic - what will be the expected/current traffic volume
- Load - what is the load of the service active elements
+ Service Level Agreement
SLA's will be defined by Stackholders and Service Provider based on the KOPI
PM&V process
PM&V [initial]
output Metrics (KOPIs) *
DEFINITION Metrics, (KOPIs)
output PM&V proposal
DESIGN PM&V Architecture: Metrics in SW / HW to Service Topology
output PM&V Architecture
ESTABLISH PM&V Architecture: Metrics in SW / HW to Service Topology
output TEST sequences PM&V Architecture
TEST PM&V Metrics: Check against *
successful
output Defined SLA's + based on KOPI
HANDOVER to OPERATIONS
PM&V [operative]
failed to deliver required performance information
START
Set up Service Topology
output: Granularity scope of measurement as needed
Modification PM&V, active/passive metrics
output reviewed proposal new requirements definde
Identify unaccounted required operational performance data
output lack of performance information
Process
Function
Object Point
Object Flow
Control Flow

Figure 1.1: Generic PM&V process

From discussions with NRENs, universities and conference/workshop organisers, a Performance Monitoring & Verification (PM&V) process was defined, as depicted in Figure 1.1. The process above can be manually introduced for collecting and analysing measurement data or incorporated into automated procedures implemented on the building blocks of the PM&V architecture, see Section 3:

Four (4) main functions of PM&V are defined as follows (detailed description see Appendix A):

1. **Definition Metrics** - INPUT: Granularity of measurement scope / OUTPUT: Key Operational Performance Indicators (KOPIs).
2. **DESIGN PM&V Architecture** - INPUT: Defined Metrics in SW, HW / OUTPUT: PM&V proposal Metrics in SW/HW to the service topology.
3. **ESTABLISH PM&V Architecture** - INPUT: PM&V proposal / OUTPUT: PM&V Architecture, metrics implemented in SW/HW to the service topology.
4. **TEST PM&V metrics** - INPUT: PM&V architecture / OUTPUT: Test sequences on metrics, and validation to the OLAs.

The output of this process allows the introduction of an operative performance measurement and verification service, in GÉANT's case, e.g. Wireless PM&V as a Service (WPM&VaaS).

This document proposes a pragmatic solution to wireless crowd source performance measurement and verification. Its main purpose is to present an outline of an initial WPM&V architecture and design. The background section will also focus on standardisation, and will provide a service angle that can be integrated (modularly) into a generic PM&V framework/architecture model and facilitate a higher QoS on network service delivery on the campus, at multiple locations or on network domains.

The WPM&V architecture is actually a prototype, with the deployment efforts built from scratch. The motivation and requirements for the WPM&V deployments were discussed and collected as part of conference demonstrations and implementations at universities.

A brief overview of other related projects is also presented in Section 2. The building blocks of a WPM&V architecture are describe in Section 3. Finally, demonstrations are detailed in Section 4, and a conclusion with recommendations for future work is provided in Section 5.

The original version of this document detailed the status of GN4-1-SA3 T3 WiFiMon activities as at Q2, and was published on 5 November 2015. This updated version, published in the last month of the GN4-1 project, on 11 April 2016, contains the following new sections:

- Section 3.6 Testbed.
- Section 4.5 HEAnet Workshop and the related Appendix E HEAnet Conference: Comments on the Best Sessions.

Additional material about the PM&V process, IETF-LMAP metrics, terminology and the GUI can be found in Appendices A, B, C and D.

2 Background

New wireless technology utilises several powerful and complex technologies. These technologies most recently appear in the form of relevant WLAN standards, IEEE 802.11a/b/g/n, and in eduroam [EDUROAM]. Thus, measuring the quality (performance) of the wireless network can be particularly challenging. There is no one tool that is able to cover all aspects of performance verification and monitoring at once. So the goal continues to be to draw on developments that measure how 'healthy' the network is and to determine how individual end-users (clients) are faring with WiFi at a given place on the network, at a given time.

With eduroam, a new perspective has arisen that provides network extension capacities and enlarges the network footprint across the entire GÉANT network. This offering poses new questions to measurement methods, to essential key operational performance indicators, as well as on how to measure and collect performance measurement data on the verification of performance issues on the NRENs' wireless campus networks.

As a result, the term 'verification' is receiving more attention and a network quality statement is needed; statistics and knowledge of performance issues are no longer enough. It's also important to verify issues, to localise them and have, for example, a 'weather map' (visualisation), that shows the lack of performance on the campus network when introducing performance issue solving processes, respective of improvements made to the network topology.

Today three sources can report important 'nuggets' of information to performance measurement and verification [WIRELESSPERF]:

- *End-user-Device*: End-user-devices can be different, from the Operating System (OS) as well as from the manufacturer (vendors). So the network distinguishes between laptops and smartphone devices when connected to the wireless networks. The devices can deliver significant performance measurements, such as data rate, and signal strength as the WLAN adapter sees them.

- *Access Points (AP) / Controller:* The management interface can show a view of the real-time connection to the network, the service set identifier, as well as details on the connection and amount of traffic. What is not seen is traffic levels and the performance behaviour of applications.

- *Network Management Systems (NMS):* The functionality of the NMS today is to have visibility of aggregate and individual client's use of applications and how clean (or free of interference) the spectrum is into the network and measured at the APs; in real-time and with a history. This allows basic troubleshooting on clients, the authN/Z processes, successful or failed.

Using these three sources of reporting can only determine the operational status of the wireless network (e.g. up/down). Statistics are reported, but a quality statement about the performance behaviour on the network is lacking: delays (e.g. of critical services like DNS), jitter (e.g. web-sources like You-Tube videos running properly) or bandwidth from a client (user-device) perspective on the WLAN.

The missing reporting element is *the performance measurement and verification of users' behaviour on the wireless campus network, called end-user feedback*. This is important, as it will provide a complementary view to the objective measurements by hardware probes, infrastructure information and end-user feedback.

2.1 Wireless PM&V into GÉANT (R&E)

In order to arrange GÉANT's work on wireless PM&V, a number of activities from GN3plus [CAMPUSBP], and on standardisation, the IETF LMAP group activities were highlighted [IETFLMAP]. Thus we are focusing of four major activities in performance monitoring and verification:

- UNINETT; a pilot project– WLAN Monitoring and Controlling of Probes
- CESNET; WLAN Service, Monitoring and Quality testing
- GÉANT– eduroam.org ; A WLAN Probe on RIPE Atlas for large-scale. simple campus use and a Raspberry PI based 'expert' toolkit for NROs/advanced diagnostics
- UNINETT – IETF LMAP group.

2.1.1 UNINETT – WLAN Monitoring and Controlling the Probes

As part of a pilot project for improving the monitoring of WiFi quality, UNINETT has deployed several miniature WiFi monitoring probes in different campuses of Universities and Colleges in Norway [UNUINETT].

2.1.1.1 *Motivation*

Every year UNINETT arranges a workshop where customers are invited to discuss the latest news in network management, monitoring and measurements. This workshop is also an arena for customers to provide feedback to UNINETT about any issues they have related to these topics. In 2014, UNINETT received a great deal of feedback from customers who ***wanted better monitoring of WiFi quality***. Several customers said they often had complaints from professors or students about WiFi problems in a lecture hall or office. However, when the IT department went to check on the issue, everything was often working fine. Some have bought commercial tools for debugging WiFi networks, such as Fluke AirCheck [FLUKEAIRCHECK]. Unfortunately, WiFi problems are often transient, to find the problem monitoring is required over a long period of time. Commercial equipment is usually too expensive to leave unattended in a lecture hall, so people wanted cheap measurement hardware that they could install permanently or semi-permanently all over campus.

Based on the customer feedback, UNINETT decided to start a pilot project developing a WiFi monitoring probe based on Raspberry Pi computers [RASPBERRYPI]. So far probes have been installed in ten different campuses, and the feedback has been very positive.

2.1.1.2 *Hardware and Software*

The probes use a Raspberry PI Mod B+ [RASPBERRYPIB+]. This is a credit-card sized computer with enough CPU power and memory to perform all the measurements needed. It does not have a built in wireless network card, so a WiFi USB adapter was used. Rasbian was selected as the operating system, which is a Linux distribution for Raspberry PI based on Debian. With this approach, all of the wireless network metrics are able collected by existing open source tools or basic Linux utilities that already form part of Raspbian.

2.1.1.3 *Measurement and Metrics*

The probes collect more than 40 different metrics at configurable intervals. These range from low-level metrics about the WiFi infrastructure, to higher-level measurement about DNS and HTTP requests. Both IPv4 and IPv6 is measured. A full list of the available metrics is available in Appendix B.

2.1.1.4 *Controlling the Probes*

All measurements originate from a single script and are executed sequentially. It is not possible to run some measurements more often than others, and all probes belonging to the same domain run the same scripts. The probes are only connected to the WiFi network while it is actively running the measurements. As soon as it is finished measuring, it disconnects from the network until the next measurement session starts.

2.1.1.5 *Lessons Learned and Future Work*

From UNINETT’s own experience and from the feedback of pilot participants, it looks like the WiFi monitoring probes can be a very useful tool for measuring the quality of the WiFi network on campus. The pilot was pragmatic and focused on **what** was possible to measure instead of **how** things were done.

The control of the probes is very primitive. While it is easy to add or remove measurements from probes belonging to the same domain, it is not possible to control individual probes.

Further, there is no global control of when measurements are performed. When a probe is turned on, it immediately starts a measurement session before it sleeps, until it is time for the next measurements. If a large number of probes are turned on at the same time, it can quickly overload the infrastructure if all the probes carry out throughput tests at the same time.

For a production system, it would be better to use a proper protocol for both controlling the probes and for collecting the results. The IETF LMAP protocol is designed for this and looks like a good choice, see Section 2.1.4.

2.1.2 CESNET - WLAN Service and Monitoring

The CESNET WLAN is located within the *Czech Technical University* campus, where most places are covered by eduroam and organization's private WLAN services. The network is mostly based on CISCO technology. All CESNET's users are authenticated by RADIUS servers, regardless of the account type. To be able to fulfil the requirements for roaming, the network controllers of faculties and organizations in the campus are interconnected and monitored.

2.1.2.1 *Interconnection*

Almost every installed Access Point (AP) on the campus is connected to a CISCO controller. The controller is typically a switch with extended IOS or specialised operating system *AirOS* [AIROS]. Currently, there is a single controller per organisation. These controllers are interconnected by *CAPWAP* tunnels, which enables them to share the information about the connected users and access features such as user roaming between different organisations and different network spaces with a single eduroam account. A user can walk around the campus without any need of getting a new IP address or rerouting the data flow. *CAPWAP* tunnels encapsulate and encrypt the user data and send it to the original network controller.

2.1.2.2 *Service monitoring*

The two types of monitoring services are used. For a general view, the Nagios monitoring system [NAGIOS] is used with classic Open Source utilities such as a ping or iperf [IPERF]. To obtain a detailed summary of all devices, the proprietary CISCO tools are used. Every CISCO controller is connected to the monitoring system *CISCO* Prime [PRIME]. As new APs equipped with two radio modules were recently installed, it is possible to have a complete overview of the network infrastructure and radio stations transmitting in the 2.4GHz or 5GHz radio bands. A monitoring system enables the view and detection of possible anomalies in the WLAN connection and even finds its source (such as Bluetooth devices, microwaves, etc.). With more APs connected, there is also possibility to track the individual users on the map with accuracy of meter units.

2.1.2.3 *Load balancing*

With so many connected users, there is a need for load balancing of clients between different APs. Since there is no standard way to inform the client which AP it should connect, some reliable solutions were elaborated. A thin client scans the band for available WLANs and connects to one with a known SSID. Due to the number of APs broadcasting the same SSID, a client must choose one. This can be done on the basis of signal-strength, which AP was found first, or just randomly. The CESNET and campus network utilises this behaviour to offload the busiest APs by disconnecting the clients and letting them reconnect to close AP broadcasting the same SSID. Unfortunately, this technique is not always successful, as selecting the right AP depends on the client's geographical location.

2.1.2.4 *Quality testing*

There are currently no regular network quality (bandwidth, latency or jitter) measurements. Tests are used when users report connection problems or AP is not responding to pings/controller in time.

In these cases, the iperf software is used to find out the problems of the current AP. When iperf is not an option, a speed-test service is used [SPEEDTEST]. The current version of speedtest is based on OOKLA proprietary flash technology [OOKLA]. CESNET is now developing the replacement, specialized software similar to iperf, but based on modern web technologies. The idea is to develop a C++ websocket server and client-side JavaScript library. The software should be able to perform tests for a durations ranging from seconds to tens of minutes. Parameters such as average and maximum bandwidth, latency and jitter will also be measured.

2.1.3 GÉANT – eduroam.org – A WLAN Probe on RIPE Atlas

Within the GÉANT network, eduroam extends the network capabilities and enlarges the footprint accessible to the user. Roaming over multiple domains requires WiFi measurements, reliability and consistency of the network, to guarantee high end-user satisfaction based on QoS. Thus SA5 T3 is evaluating a number of technologies [EDUROAM]. However, the scope is not purely on the wireless performance aspects but also includes AAI measurements.

The eduroam community has pursued the following initiatives:

- In eduroam pulse, 2009-RIP AARNet tested SheevaPlug. SheevaPlug does not scale and there was also not a community.

- In Active Monitoring eduroam Node (AMeN), 2009–2014, an Alix System that deployed 14 sites was evaluated at SURFnet. The deployment of this activity is still running.

- The activity Srce Raspberry Pi (“Frankenprobe”), 2012 to present, concludes scanning for SSIDs and Signal Quality. The approach was not favoured as a solution.

- UNINETT Tradlos Probe, 2014 to present, is a project on Raspberry Pi (see Section 2.1.1). The focus of this project is on trouble shooting and continuous quality measurements.

- Janet/Loughborough Uni Probe, 2012–2014 (still running) focused on WPA supplicants and a script factory. Used probes are TP-Link MR3020, and are similar to Atlas probes v3. The numbers of probes are growing, from 20 up to 200.

- Use of RIPE Atlas probes, and extension of the probes with eduroam-specific metrics.

Following a review of several of these options by the eduroam team, and the developers involved with the above initiatives and other National Research Organisations (NROs), the eduroam task in GN3plus and GN4-1 has chosen to focus on two options to cover two use cases: RIPE Atlas and Frankenprobe.

2.1.3.1 *RIPE Atlas*

In GN3plus, eduroam carried out a trial of 50 Atlas probes. From this trial and following discussions with RIPE, it was determined that with some enhancement, the RIPE Atlas probes can deliver the following eduroam-specific measurements, as well as a range of general functionality such as ping and traceroute.

- Variants of eduroam SSID present.
- Different access points offering eduroam SSID.
- Try to connect using credentials.
- Report/check certificate.

RIPE NCC has the logistics in place to support wide-scale delivery, and is willing to collaborate on the features roadmap if GÉANT/eduroam can help to increase the community uptake. As 50% of the networks connected to eduroam already have some Atlas probe presence, this indicates good trust in this system for the community to adopt a widespread deployment. In GN3plus, the GÉANT project sponsored the RIPE Atlas initiative in 2014 and distributed several hundred probes. In GN4-1, focus is on the roadmap features and integration with eduroam supporting services.

2.1.3.2 *CARnet Frankenprobe*

The expert toolkit provided by the Raspberry Pi-based probe allows for advanced diagnostics when more basic systems either do not detect a problem or do not provide enough information to resolve more complex problems. It is intended to have a smaller deployment base. That means the “Frankenprobe” is foreseen for diagnostics, expert probes only.

2.1.4 UNINETT – IETF LMAP

The charter of the IETF LMAP working group states that they will standardise a measurement system of broadband access devices such as home and enterprise edge routers, personal computers, mobile devices or set top boxes. They will cover both wireless and wired measurements.

The working group is chartered to specify an information model, the associated data models and select/extend one or more protocols for secure communication. The first version of LMAP has a requirement that the measurement system is under the control of one single organisation. This requirement is there to make the initial work easier. Active and passive measurements are both supported and privacy issues are an important part of the core requirements.

The IETF document Large-Scale Broadband Measurement Use Cases [RFC7536] describes several use cases for LMAP. It divides the uses cases into two main categories, one for ISPs and one for regulators. The ISP use cases focus on how ISPs can measure their own network to document that the performance is good and to detect when problems occurs. Regulators can use LMAP to compare measurements from different ISPs so that they can get an objective view of the performance.

2.1.4.1 *Framework and Building blocks*

Figure 2.1 shows the main components of LMAP:

- Measurement Agent (MA) - performs measurement tasks.
- Measurement Peer (MP) - assists the MA with the measurement tasks.
- Controller - manages the MA.
- Collector - receives measurement results from the MA.

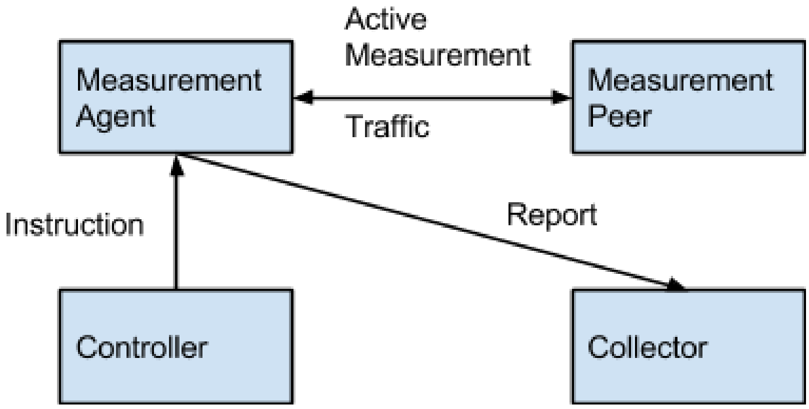

Figure 2.1: Framework - Building Blocks LMAP Architecture

2.1.4.2 *Measurement Agent*

The MA performs measurement tasks that the controller has instructed it to carry out. Some tasks may include an MP that assists the MA with the measurement task. One example of this is a throughput test using iperf, where the MP will be running an iperf server. The MA sends the results of the measurement tasks to a collector.

2.1.4.3 *Controller*

The Controller controls the MA through the control protocol. Using this protocol the Controller can tell an MA to carry out specific measurement tasks, when to do them and when to report the results.

One Controller can control multiple MAs, but to simplify the design, one MA can only be controlled by one single controller.

2.1.4.4 *LMAP tasks*

In LMAP, everything is defined as tasks. You can have a task that tells the MA to contact the controller to download new configuration, you can have measurement tasks that do the actual measurements and you can have reporting tasks that sends measurement results to a collector. Schedules are used for executing tasks at specific times. It is possible to pipe the results from one task to the input of another task. This can, for example, be used to aggregate results before sending them to the collector.

2.1.4.5 *Protocol*

The main work of the LMAP working group is to define the Control and Report protocol. It has now been decided to write a YANG module for LMAP and use RESTCONF as a protocol. RESTCONF is a REST based protocol for accessing information described by YANG modules.

While RESTCONF will be used as both a control and reporting protocol, there is some talk to also support IPFIX as a reporting protocol for measurements that generate a lot of results. A good example of this is flow-based measurements.

2.1.4.6 *Registry for Performance Metrics*

The LMAP working group will not define the measurement tasks. The idea is to have a repository of available tasks that the controller can point to when instructing the MA. The work on this registry is done in the Internet Protocol Performance Metrics (IPPM) working group [IPPM]. The registry will be populated with metrics from existing RFCs from many different working groups.

2.1.4.7 *Collaborative LMAP*

A private draft has been published describing several use cases for a collaborative LMAP, where multiple autonomous measurement systems can collaborate to solve end-to-end performance problems [LMAP]. So far, this document has not caused a great deal of discussion, since the LMAP charter explicitly states that the LMAP working group will not look at multi-domain monitoring. It does, however, show that there is an interest in multi-domain monitoring, and there is a good chance that the LMAP working group will start looking at this after the current work is finished.

2.1.5 Conclusion

In all activities described above, one important end-user requirement was improved monitoring of the WiFi (WLAN) quality.

WiFi monitoring probes are very useful for measuring the quality of the WiFi network on the campus (see Section 2.1.1 and 2.1.3); often the question of WHAT (e.g. bandwidth, latency and jitter) is possible to measure is answered and not on HOW things are done. Further there is not only a need for performance measurement also on AAI measurement described in form of use cases in Section 2.1.3.

Controlling of the probes can be very primitive; to control individual probes is difficult, and often not possible. Further, a global control of the measurements is not given. So no regular measurements will be performed. Performance measurements are mostly on best effort, initiated on user reports or investigated by the support team if the APs are not responding on ping/iperf scans. In Standardization, the IETF LMAP activities mentioned in Section 2.1.4.7 would allow that multiple autonomous measurement systems could collaborate to solve e2e performance measurement. This would mean regular measurements could be taken over multiple domains, which would present a good opportunity for ongoing activity, as the LMAP working group will looking for this type of reporting .

New deployments replacing specialised software (e.g. iperf) are based on web technologies, with JavaScript deployment on the client site. However, for a production-ready system, a proper protocol, described in Section 2.1.4, for both controlling the probes and collecting the results would fit, as focussed attention should be paid to a hybrid approach with less impact to the network traffic(see also Section 2.1.3).

The approaches described in this document focus on performance measurement/monitoring and their verification from an infrastructure point of view, so answering the WHAT will be measured. What is missing is the user feedback, the HOW things will be done. This will be discussed in Section 3, and will provide a basis for a WPM&V service.

3 Wireless PM&V Architecture

The Wireless PM&V (WPM&V) architecture is a greenfield approach, and comes up to a prototype network setup that provides a WPM&V service to GÉANT and to the NRENs NOC. Thus we act on the assumptions on the hypothesis underlying all our work as followed:

"…*It is possible to gather data from multiple sources, including browser-based measurements, in addition to traditional monitoring, and extract meaningful information on the performance of a WiFi network from that data…***"**

Every action taken in the development stage of this architecture serves to confirm or disprove some aspect of this hypothesis, as does the architecture described in Section 3.2 below.

To break this down, it means:

- That it is possible to run non-invasive performance tests within the environment of a browser or a mobile app.
- That a sufficient number of these tests are adequate to give an accurate picture of overall performance of the network in a given location.
- An adequate number of tests can be performed by embedding JavaScripts in a website or app that the end-users in a particular location use from day to day (e.g. university or conference website).
- Existing Layer 2/Layer 3 mapping data maintained by eduroam networks can be used to separate out the performance tests by access point.
- Once data on performance by access point is collected, we believe that the data can be analysed and visualised in order to provide useful troubleshooting and capacity planning information to the network administrator.

In Section 3.1 we will discuss the requirements and pre-requisites (generic functionality). This includes the data that needs to be gathered, the equipment that may be used on a given campus, and a walk through of the process, from WiFi client connection to report generation.

Once that is established, Section 3.2 will layout the building blocks that provide this information, and show how they may be assembled in order to execute the process described in Section 3.1.

Section 3.3 describes the functional and non-functional requirements for both single domain examples and how this could be expanded into a multi-domain approach.

Every campus is different, so Section 3.4 will describe the various models that we believe can be used, as circumstances require.

Finally, while we do not gather additional user data beyond what is already gathered for the operation of the eduroam service, we are using some of this data in a new way, so the privacy implications must be considered carefully, discussed in Section 3.5.

3.1 Requirements and Prerequisites

This section describes the objective, the prerequisites, and the functionality of the Wireless Monitoring applications, and their enhancements. A "walk through", a minimum data-set and its extension complete the requirements for defining the architecture building blocks in Section 3.2.

3.1.1 Objectives

What's needed is a simple toolset that allows us to have an insight into the Wireless performance of various locations across a campus, multiple campuses or sites. The following sections provide an outline of the functionality of the tool and the requirements that a network administrator/support team would need, to have some useful visibility of wireless network performance. It is envisaged that the SA3 team would establish this functionality, get the tool working and build additional functionality in an iterative fashion. The 'look and feel' is, for now, a secondary consideration, however, it is something that we should undertake at a later stage. The following prerequisites will help us to establish the data sets we require and the query statements, menu options and output formats for a GUI-based monitoring application. The project will need to agree the distilled data set to meet the functionality required. Additionally, the project will need to provide a 'walk through' from the client associating with an Access Point to the network admin interface to the monitoring system, through to the visualisation of the query outputs in the GUI.

3.1.2 Prerequisites

As prerequisites, we expect:

- Wireless network with eduroam identity provider and service provider capability
- Syslog server logs
- Radius server logs
- DHCP logs

3.1.3 Functionality of WiFiMon Application (Monitoring System)

The functionality of the wireless monitoring applications can be described as followed:

- Process for discovery of APs and the addition of a ’Discovery’ button in the GUI to activate this.
- The above process should provide output in the form of a list of APs that reside in the GUI’s database.
- Ability to group APs (from a list of those found) by ticking a checkbox for each AP and create a location name (group) and a ‘Group’ create button.

This will result in many selected and defined groups of AP (locations). A group of APs can also correspond to a physical location (room, auditorium, library, etc.). These groups should be displayed in the GUI as a collection of APs, labelled with the user-defined location name and descriptions such as:

- Query to extract top five best performing locations and display on a dashboard
- Query to extract bottom five performing locations and display on a dashboard
- Dashboards to be updated in real time

The definition of performance needs to be clarified, as well as the aggregated performance measurement data of all clients connecting to all APs in a selected location. In other words, the aggregated performance in Lecture Room A, the aggregated performance in Lecture Room B, the aggregated performance on the top floor of the Library, etc. While it may be very useful to track individual mobile devices (clients), perhaps it is best to focus on overall performance, rather than individual user or client issues. This may be investigated in future projects, however, the structure of the database should not preclude the further enhancement for this functionality, such as:

- Ability to select a location (based on the AP grouping as described) and run a report with a graphical representation of historical performance over a given time, e.g. Monday to Sunday.

- Ability to set an acceptable baseline (agreed) performance or threshold aggregated value to be indicated on the same graph, indicating the delta – actual performance vs. agreed aggregated threshold for what constitutes an acceptable level of performance.

So how is a globally acceptable threshold of acceptable performance defined? Is there a min download / upload speed aggregated across all clients at a specific location, such as a lecture hall with 12 APs, or a small classroom with one AP. This point requires further investigation, such as:

- Ability to generate (create) a report/graph of performance at a location in the past versus live performance. In other words, what was the performance like last week and what is it like at present? Is there an issue now that was not there last week? Pre-defined report query – click and run. The output to the GUI may be a graph that overlays the performance from a past week with the performance of a current week – or something similar.

3.1.4 Enhanced Functionality of WiFiMon Application

Enhancement of this application should focus on the search function:

- Ability to search the GUI for a username, device MAC-address, or device IP address, and to be able to track the performance of a device over time on a graph overlaid with the location information. End-users' privacy must be considered, see Section 3.5.

3.1.5 Walk Through

A walk through of an application can include, but is not limited to:

- A client device associates with a wireless access point, authenticates (or not) and receives an IP address from a DHCP server.
  - A DHCP log entry is created.
  - A syslog entry is created.
  - A radius log (authentication log) is created, allowing the match of a client MAC-address with a client IP address, an access point identifier, and date stamps.

This data is generated through controller/AP configuration to data collectors.

- This data may be parsed by scripts and auto populates database tables in the monitoring system database.
- The Web browser on a client initiates a connection to a frequently accessed web page, such as e-learning portal page, etc.
- The execution of JavaScript on the client runs a series of network performance tests.
- The execution of JavaScript is limited for each device, to avoid over-burdening the client browser.
- Performance test values for the client IP address and timestamp are recorded in a database table in the monitoring system database.
- The query language, hidden from network admin with GUI functionality, provides configuration options that may be tailored for a specific site, e.g. grouping APs to represent location (described in further developments).
- Pre-defined report generation and ad-hoc report capability are available to network admins through the GUI interface on the monitoring system (described in further developments).

3.1.6 Data Set – Performance of WiFi Location, Client and User Tracking

A data set can be listed, but not limited on:

- Time stamps for performance test values
- Time stamps for client IP address
- Access point identifier

- Client MAC-address
- Client IP address
- Client username

3.2 Architecture Building Blocks

The description of the technical architecture includes the building blocks with detailed description of the components, depicted in Figure 3.1, and references to specific technologies that can be utilized for the deployment. Terms and terminology of the architecture are explained in Appendix C.

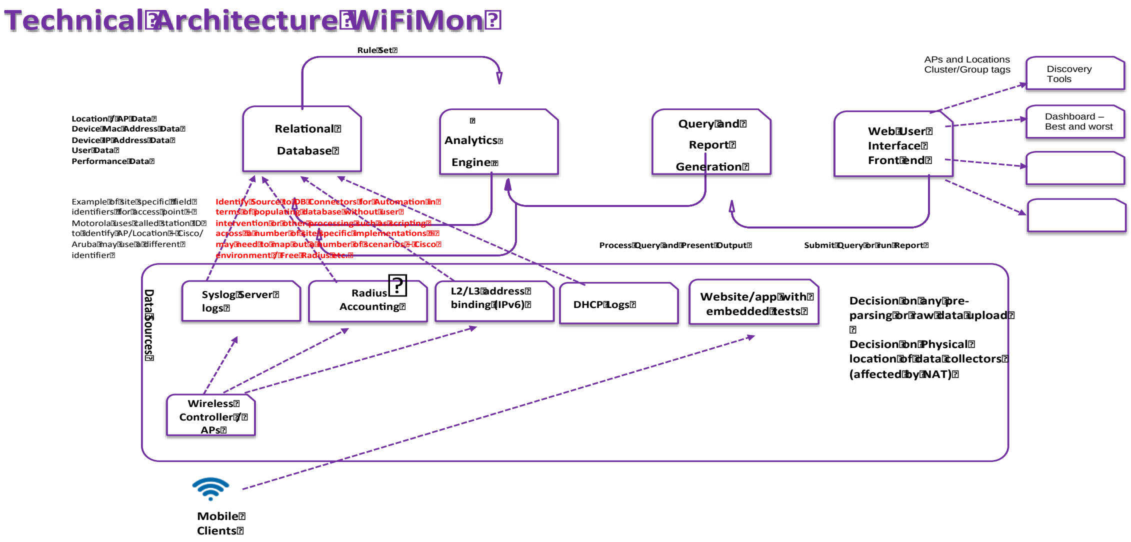
Technical Architecture WiFiMon
Rule Set
Location / AP Data
Device Mac Address Data
Device IP Address Data
User Data
Performance Data
Relational Database
Analytics Engine
Query and Report Generation
Web User Interface Front end
APs and Locations Cluster/Group tags
Discovery Tools
Dashboard – Best and worst
Example of site specific field identifiers for access point – Motorola uses called station ID to identify AP/Location – Cisco/ Aruba may use a different identifier
Identify Source to DB Connectors for Automation in terms of populating database without user intervention or other processing such as scripting across a number of site specific implementations – may need to map out a number of scenarios – Cisco environment / Free Radius etc.
Process Query and Present Output
Submit Query or run Report
Data Sources
Syslog Server logs
Radius Accounting
L2/L3 address binding (IPv6)
DHCP Logs
Website/app with embedded tests
Decision on any pre-parsing or raw data upload
Decision on Physical location of data collectors (affected by NAT)
Wireless Controller / APs
Mobile Clients

Figure 3.1: Technical Architecture - Building Blocks

As a consequence of changing requirements, we have to assume that during the project implementation, part of the proposed technologies may be changed.

3.2.1 Data Sources

The data source layer is responsible for: (i) Generating information (through websites with embedded test procedures, namely JavaScript code embedded in the website page that enables users to run process tests without intervention.) regarding the performance of the wireless network as experienced by the users / mobile clients, and (ii) exporting the data from data source collectors such as Syslog, DHCP logs, and Radius accounting to the Relational Database (RDB). That means in detail:

- Open Source tools, such as Boomerang and NetTest are embedded into websites in order to extract network-related information within a Web browser [BOOMERANG], [NETTEST]. Such information includes performance data, such as throughput on downloads and uploads of images with various sizes [THROUGHPUT], and round-trip time (RTT) through ping as experienced by the end-user [RTT]. The procedure takes place without user intervention and without overloading the wireless network. In fact, Boomerang is based on JavaScript that means to measure the performance of a website from the end user's point of view. NetTest provides a way for gathering network-related information within a Web browser. It uses network calls available in JavaScript and Flash [JAVASCRIPT], [FLASH].

- Wireless controllers and Access Points (APs) in a smaller implementation may provide the minimum amount of information, at the most basic level (see list below) for performance verification and monitoring on a wireless campus network. Some implementations may have a proprietary data source collector. Therefore, initial focus will be on Open Source implementations that can export the data-to-data source collectors, such as Syslog, DHCP logs, and Radius accounting. In some site-specific implementations the collectors are embedded in the controller/APs, while in other implementations, they may be set up as external collectors that receive the data from the controller/APs.

Thus the data sources will vary, depending on the type of data to be collected. In some cases, the performance data, along with a single source of device/AP, will meet all requirements as a minimum in order to allow the correlation of the collected data. Indicative data from collectors, such as that held by wireless controllers, access points, as well as performance data is listed below.

- Wireless controllers and AP information

- MAC-address of client device
- IP address of client device
- AP Mac-address
    - AP IP address
    - AP location identifier (e.g. NAS-identifier in sample Motorola data)
    - Connect info (speed of connection, not throughput; 802.11b,g,n.a as reported by APs)
    - Username associated with device (Radius username)
    - Correlated timestamp UTC
    - Client device name
    - DHCP logs
- Performance data [PERFORMANCE]:
    - Download throughput
    - Upload throughput
    - Round-trip time
    - User Agent / Web Browser Identifier
    - Geolocation information (accuracy is important and difficult to achieve without user intervention)

3.2.2 Relational Database

The raw data, (the data collected from the data source layer, that have not been subjected to processing or any other manipulation) such as Syslog Server logs, DHCP logs, Radius Accounting, performance data, etc.) will be automatically collected in the Relational Database (RDB). This means that the raw data generated from the data sources block (described in Section 3.3.1), will be automatically imported in the RDB, e.g. through scripts or some database connectors without manually import/correlation of the raw data by the end-user/network administrator. The RDB will use SQL as the language for querying and maintenance [SQL]. Depending on the required scalability, the database may be centrally installed, so that raw data from all APs will be sent to one database per site/campus, or across a number of site implementations. In every case, the database will be based on Open Source technologies, such as PostgreSQL or MySQL [POSTGRESQL], [MYSQL].

3.2.3 Analytics Engine

The Analytic Engine (AE) is the architecture block responsible for examining/analysing the RDB's raw data and preparation of reference data (namely, the data objects relevant to transactions, consisting of: sets of values, statuses or classification schema, such as raw data in transactions to visualization) [REFERENCE]. Thus the main functionality is to sort the raw data collected, analyse it and provide visualisations using tools that provide the greatest insight on the wireless network performance. Based on initial research, one candidate for the development of these tools is the Java-based framework Spring XD [SPRINGXD]. Spring XD is an Open Source project, a unified, distributed, and extensible service for data ingestion, real time analytics, batch processing, and data export.

3.2.4 Query and Report Generation

The Query and Report Generator (QaRG) is used to get specific information from the RDB and the AE. Its main purpose is (a) to search for usable information from these two architecture blocks and (b) to post this information in the form of reports or visualisation options to the Web-user Interface (Web-UI), the front end.

3.2.5 Web User Interface Front End

The available data from the RDB and the AE is accessible through a network admin Web-UI which allows data querying. Network administrators are the end-users of this Web-UI, which allows investigation of the collected performance reference data, and in turn, status checks of the wireless network. This block of the architecture is also responsible for projecting the collected data and allowing real-time visualisation options, such as:

- Collected data for a specific time period.
- Collected data for a specific AP.
- Min-max-mean values of download/upload/latency measurements.
- Non-normal measurements.
- Top five performing locations (see Section 3.1).

- Bottom five performing locations (see Section 3.1).

Additional functionalities for this architecture block includes, but is not limited to:

- The ability to discover a specific AP.
- The ability to select from a list of APs and group them together at a single location. One example is grouping a number of APs in the same room and to put a tag on this. From this, a more accurate picture of the network quality within that room over time may be built.
- The ability to type in a name of a user based on username, a MAC address of a device, or an IP address of a device, and to track the performance of a device over time on a graph overlaid with the location information.
- The ability to graph the results over time in a number of different forms: scatter plot, line graph, etc.
- The ability to generate/create a report/graph of performance at a location in the past versus current performance.

3.3 Single, Multi-Domain (Non-) Functional Requirements

We separate the requirements between the tester, the JavaScript that runs inside the user's browser and the back end, the combination of modules above that form the data gathering and analysis functions.

3.3.1 Tester Requirements – Functional

The functional requirements are listed below:

- The tester must perform download speed and latency calculations from inside a browser environment, or embedded within an app that is not dedicated to the purpose.
- The tester must, on completion of the tests, report the results to a backend server.
- The tester must, selectively, be able to be configured to run over IPv4 only, or IPv6 only, or both.

3.3.2 Tester Requirements – Non Functional

The non-functional requirements (those that judge the operation of a system, rather than specific behaviours) are listed below:

- Although the user should be aware of the testing, it should not be disruptive. The tester must operate transparently to the user, and without the user's intervention.
- The tester must not repeat tests from the same device, on the same network, within a specified cool-down time.

- If the network is performing poorly, the tester should not overload the network unnecessarily.

3.3.3 Back-end Requirements – Functional

The back-end requirements are listed below:

- The back-end must reliably accept test results from the tester.
- The back-end must be able to be selectively configured to receive results over IPv4 only, IPv6 only, or both.
- The back-end must store test results along with the datestamp, IP address and user-agent of the browser.
- The back-end must be able to correlate an IP address plus datestamp pair with the wireless access point serving that IP address at that time.
- The back-end must be able to accept data from a variety of sources: (to be expanded as the project proceeds):
- ○ [vendor name] RADIUS accounting logs.
- ○ [vendor name] DHCP server logs.
- ○ [vendor name] switch IPv6 layer 2/layer 3 mapping table.
- The back-end must provide a real time view of test results, separated by access point.
- The back-end must provide a long-term report view of test results.

3.3.4 Back-end Requirements – Non Functional

The non-functional requirements are listed below:

- The back-end must support federated single sign on.
- The back-end must not retain any more data than is required to perform the performance analysis (user-identifying data must be carefully dealt with, see Section 3.5.)

The above requirements cover the single domain case, where each system operates separately in a single institution. There is a case to be made for following this with a full, multi-domain service. This has two advantages: (1) some of the infrastructure may be able to be shared, and (2) it may be possible to create a GÉANT-wide view of actual user performance, a weather map.

3.3.5 Multi-Domain Back-end Requirements – Functional

For a multi-domain backend requirements are listed below:

- The results receiver must be able to operate in a NAT environment, and report back to a separate, central analysis back end.
- The analysis backend must be able to provide multiple views of the data based on the permissions of the user viewing.

- The analysis backend must provide a detailed view of a campus to authorised staff from that campus.
- The analysis backend must be able to provide a GÉANT-wide performance overview that does not reveal user-identifying information (see Section 3.5).

3.4 Implementation Models

Our first test site, TNC2015, was constructed with minimum infrastructure, almost entirely using resources that were already present for the operation of the site's network. In this section, we start with this model as a basic installation and describe additions that may be needed in certain environments.

The TNC2015 (details see Section 4.2) test site used:

- An off-the-shelf testing library installed on the TNC website.
- An off-site virtual machine to receive the test results.
- RADIUS accounting logs to provide the IP address->AP mapping.
- Rudimentary scripts for offline processing and visualisation of the data.

These are the basic building blocks of the implementation, as reflected in the architecture section (see 3.2) above. Every change to the model is an enhancement or addition to these building blocks.

3.4.1 Performance Tests

Browser tests form the basis of our methodology. However, mobile usage is becoming dominant in network environments, and browser usage on mobile is far lower than on the desktop or laptop. While a dedicated testing app is not likely to get the quantity of results needed for an accurate overview, the performance tests could be included in an app that is frequently used on the network (e.g. an institution's own timetable app, or a conference's scheduling app – see Section 5.3)

3.4.2 Test Result Gathering

The TNC model is inappropriate for campuses where devices are not assigned a globally unique IP address; this is because, when the results are reported back to the external VM, the (public) IP address of the report does not correspond to the (private, NATed) IP address of the user's device.

The two options here are to:

- Perform the tests over IPv6, and report the results over IPv6.
- Situate the results receiver on the local network so that it receives results with addresses in the same scope as the RADIUS logs (i.e. logging private addresses.)

3.4.3 IP–>AP Mapping

Some networks may not have this information in the RADIUS accounting logs. Also, we have yet to fully verify the accuracy of using RADIUS accounting logs for this information. Other models for gathering this data are:

- Combining DHCP logs, access point logs, and RADIUS accounting logs.
- (IPv4) Polling ARP information from campus switches.
- (IPv6) Polling neighbour discovery information from campus switches.

3.4.4 Processing and Visualisation

In a single domain environment, the models to follow are:

- Immediate troubleshooting information for a room: “How does performance in this room compare to the same time last week?”
- Immediate troubleshooting information for a device: “How does the network performance of the device with this MAC address compare to the average of devices in the same visited locations?”
- Identifying hotspots: “Where are the worst performing areas and times?”
- Capacity planning: “Which areas are showing early signs of inadequate capacity?”

If we gather certain data on a multi domain basis and process it centrally, the model can go further to answer these questions:

- Immediate troubleshooting and diagnosis: “What is the impact of the current trouble?”
- Longer term views: “What was the impact of a specific backbone change at a given time?”

3.5 Privacy

Wireless, crowd-sourced performance verification and measurement is dependent on valid data-sets collected from a huge number of end-users working on NRENs’ campus wireless networks. Regarding the data sets generated from these collectors (see Section 3.2.1), geographical location information, performance data and end-user information is available, which will be analysed to establish network performance expertise, and for advice to academic ICTs, to improve their wireless network design or topology.

When trying to understand the end-user’s point of view, a dialog about privacy, collecting end-user data and transparency must be initiated. In July 2015, a discussion at the eduPERT monthly call came up with a number of questions, from transparency, tracking and end-user behaviour on campus, to the authorised entities having access to the RADIUS accounting, and the authN/Z log files [EDUPERT].

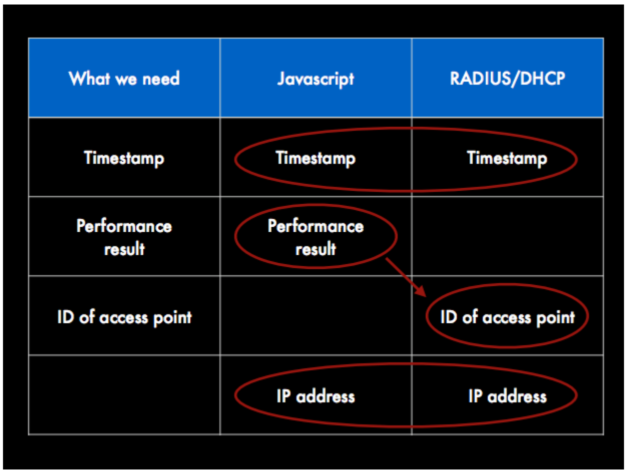

What we need	Javascript	RADIUS/DHCP
Timestamp	Timestamp	Timestamp
Performance result	Performance result	
ID of access point		ID of access point
	IP address	IP address

Figure 3.2: Performance relevant values

The project team is aware of end-user's concerns of the collection and analysing of "highly" sensitive data from the end-user's behaviour on the campus network. The fact that the eduroam habitat is used may bring forward or simplify the process of regulating data gathering. This area needs, investigation, especially the transparency of end-user data when AE analyse this. Referring to the eduroam service policy (IdP and SP) [EDUROAMSP], it is important to reinforce collectors on "what we need", "what JavaScript delivers" and "what is relevant". The RADIUS/DHCP logs for measuring performance and verification, as depicted in Figure 3.2.

Taking into account the eduroam authN/Z process, and referring to Section 3.1.6, we can argue that

- The value of the User-Name attribute in the request ('outer EAP-identity') and tunnelled EAP types (inner EAP identity) corresponds with the domain in the DNS (arbitrary@realm), and the client user name from the RADIUS-logs.

- The value of the Calling-Station-ID attribute in authentication requests is comparable with the client-MAC address of the user-device received from the RADIUS log.

However, the results of the authentication process (IEEE802.1x, EAP and RADIUS), provides the AP-ID, the domain of the client-user, the client-user name, the client-IP (DHCP logs), the client's MAC address and also allows, according the MAC address, verification of device type, e.g. Apple, Windows, Android, etc. used.

The fact that end-user values are collected and visible (e.g. by the network provider's administrator) through the eduroam concept does not offer a "carte blanche" for evaluating the whole spectrum of available parameters. Furthermore, the existence of an AE, which allows personal profile building so that the end-user can be tracked causes concern and needs special attention by the Project.

Transparency concerning the end-user is mandatory. Examples of realising this include:

- To inform the end-user through pop-ups, approving performance tests with 'accept' or 'deny' functions.

- To have links or pop-ups that explain the process of data collection.
- In case measurement tests are performed without user intervention and notification, ensure that sensitive/user-related data will be analysed with caution.

At the time of writing this milestone, initial discussions on privacy (transparency) may be initiated, but prototype level is too early to have a 'serious' dialog, as this is still an experimental stage. An extended policy discussion within the eduPERT community would make sense following the demonstration of a PoC during deployment phase, to provide a productive-ready service for the NRENs, GÉANT community.

3.6 Testbed

In order to check the WiFiMon functionality we have put in place a testbed [TESTBED] (in Dublin), configured as shown in Figure 3.3 below. Our intention was to focus first on the FreeRadius servers for each of the wireless vendors – Cisco and Aruba; for the future we envisage a number of scenarios for an eduroam configuration – Aruba with Ubuntu, Aruba with Windows Server, Cisco with Ubuntu, Cisco with Windows Server, Cisco with Windows Server and Cisco ACS.

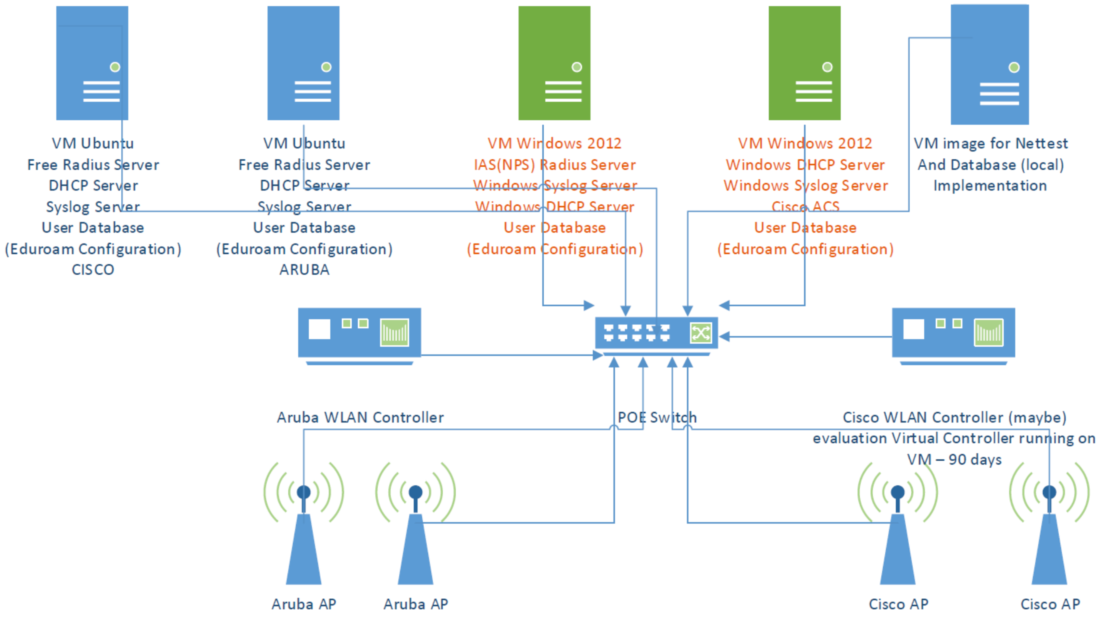

Figure 3.3: Testbed configuration

The testbed installation focuses on a case with the following specific requirements:

- FreeRadius server installed.
- FreeRadius configured to store the records in a remote SQL database.

- FreeRadius records should contain the information of the IP assigned to authorised users.

Some initial tests took place on 31 March 2016, confirming the functionality of the WiFiMon agent for measurements and correlation with the FreeRadius logs. Screenshots from these tests can be found in Figure 3.4 to Figure 3.9 below.

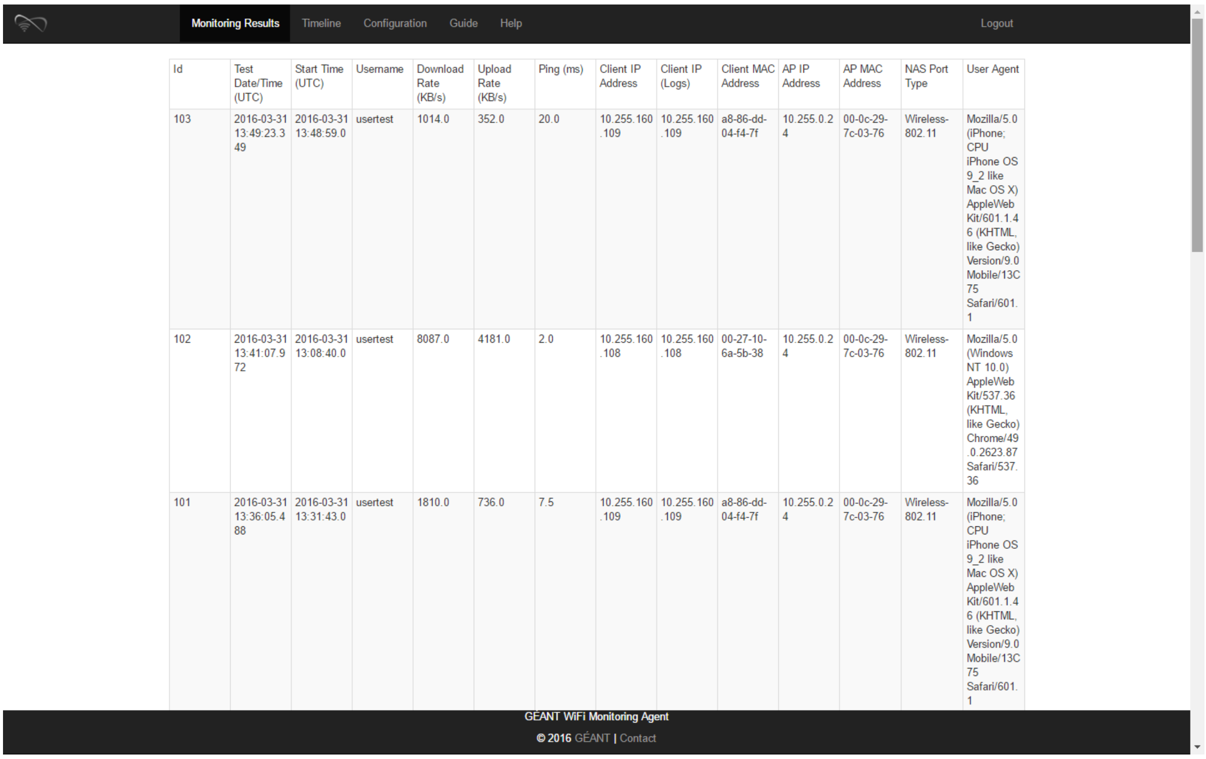

Id	Test Date/Time (UTC)	Start Time (UTC)	Username	Download Rate (KB/s)	Upload Rate (KB/s)	Ping (ms)	Client IP Address	Client IP (Logs)	Client MAC Address	AP IP Address	AP MAC Address	NAS Port Type	User Agent
103	2016-03-31 13:49:23.349	2016-03-31 13:48:59.0	usertest	1014.0	352.0	20.0	10.255.160.109	10.255.160.109	a8-86-dd-04-f4-7f	10.255.0.24	00-0c-29-7c-03-76	Wireless-802.11	Mozilla/5.0 (iPhone; CPU iPhone OS 9_2 like Mac OS X) AppleWebKit/601.1.46 (KHTML, like Gecko) Version/9.0 Mobile/13C75 Safari/601.1
102	2016-03-31 13:41:07.972	2016-03-31 13:08:40.0	usertest	8087.0	4181.0	2.0	10.255.160.108	10.255.160.108	00-27-10-6a-5b-38	10.255.0.24	00-0c-29-7c-03-76	Wireless-802.11	Mozilla/5.0 (Windows NT 10.0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/49.0.2623.87 Safari/537.36
101	2016-03-31 13:36:05.488	2016-03-31 13:31:43.0	usertest	1810.0	736.0	7.5	10.255.160.109	10.255.160.109	a8-86-dd-04-f4-7f	10.255.0.24	00-0c-29-7c-03-76	Wireless-802.11	Mozilla/5.0 (iPhone; CPU iPhone OS 9_2 like Mac OS X) AppleWebKit/601.1.46 (KHTML, like Gecko) Version/9.0 Mobile/13C75 Safari/601.1

GEANT WiFi Monitoring Agent

© 2016 GÉANT | Contact

Figure 3.4: Screenshot from WiFiMon UI: Monitoring results tab

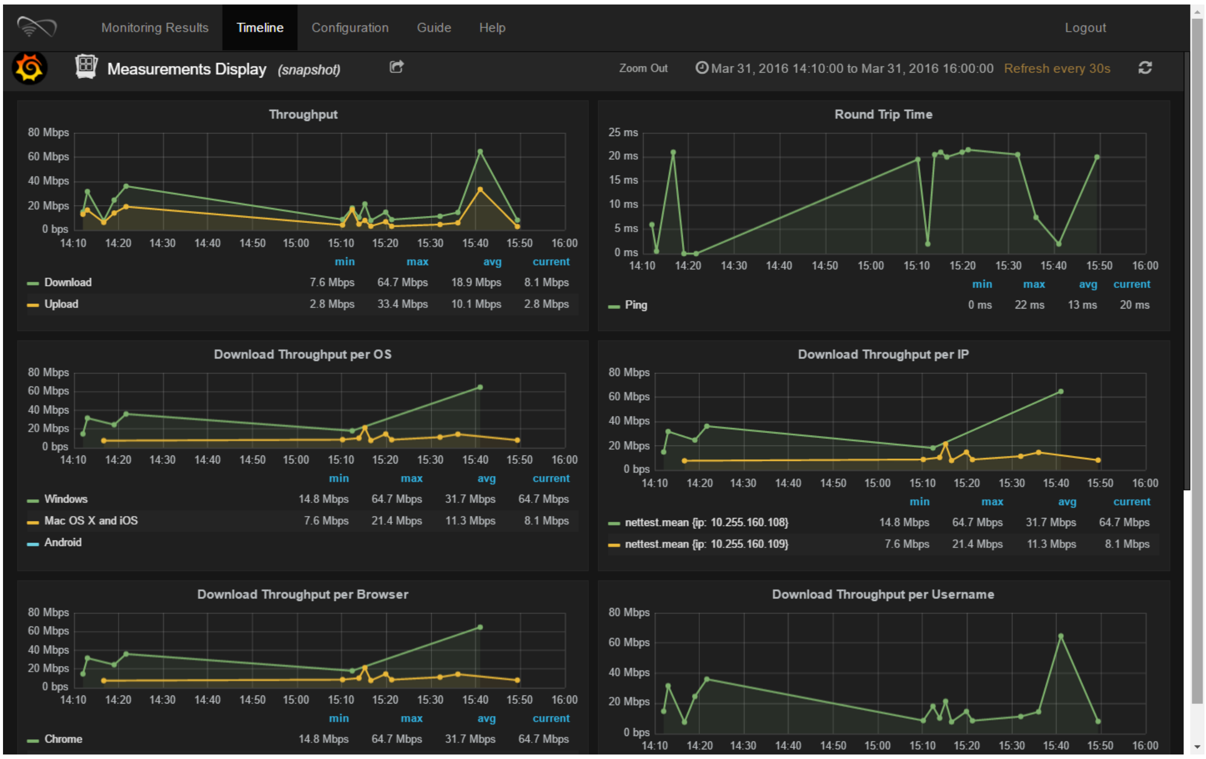

Throughput

	min	max	avg	current
Download	7.6 Mbps	64.7 Mbps	18.9 Mbps	8.1 Mbps
Upload	2.8 Mbps	33.4 Mbps	10.1 Mbps	2.8 Mbps

Round Trip Time

	min	max	avg	current
Ping	0 ms	22 ms	13 ms	20 ms

Download Throughput per OS

	min	max	avg	current
Windows	14.8 Mbps	64.7 Mbps	31.7 Mbps	64.7 Mbps
Mac OS X and iOS	7.6 Mbps	21.4 Mbps	11.3 Mbps	8.1 Mbps
Android				

Download Throughput per IP

	min	max	avg	current
nettest.mean {ip: 10.255.160.108}	14.8 Mbps	64.7 Mbps	31.7 Mbps	64.7 Mbps
nettest.mean {ip: 10.255.160.109}	7.6 Mbps	21.4 Mbps	11.3 Mbps	8.1 Mbps

Download Throughput per Browser

	min	max	avg	current
Chrome	14.8 Mbps	64.7 Mbps	31.7 Mbps	64.7 Mbps

Figure 3.5: Screenshot from WiFiMon UI: Timeline tab (1/2)

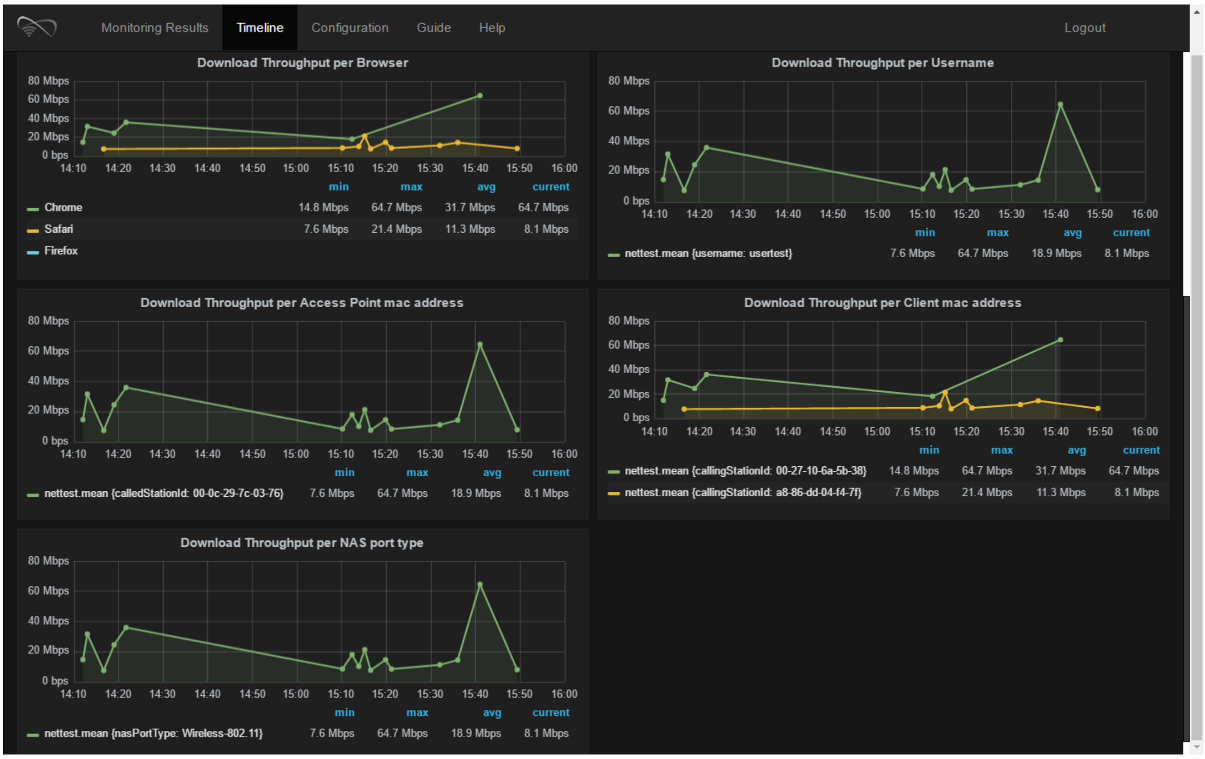
Monitoring Results
Timeline
Configuration
Guide
Help
Logout
Download Throughput per Browser
Download Throughput per Username
Download Throughput per Access Point mac address
Download Throughput per Client mac address
Download Throughput per NAS port type
80 Mbps
60 Mbps
40 Mbps
20 Mbps
0 bps
14:10
14:20
14:30
14:40
14:50
15:00
15:10
15:20
15:30
15:40
15:50
16:00

Download Throughput per Browser

	min	max	avg	current
Chrome	14.8 Mbps	64.7 Mbps	31.7 Mbps	64.7 Mbps
Safari	7.6 Mbps	21.4 Mbps	11.3 Mbps	8.1 Mbps
Firefox				

Download Throughput per Username

	min	max	avg	current
nettest.mean {username: usertest}	7.6 Mbps	64.7 Mbps	18.9 Mbps	8.1 Mbps

Download Throughput per Access Point mac address

	min	max	avg	current
nettest.mean {calledStationId: 00-0c-29-7c-03-76}	7.6 Mbps	64.7 Mbps	18.9 Mbps	8.1 Mbps

Download Throughput per Client mac address

	min	max	avg	current
nettest.mean {callingStationId: 00-27-10-6a-5b-38}	14.8 Mbps	64.7 Mbps	31.7 Mbps	64.7 Mbps
nettest.mean {callingStationId: a8-86-dd-04-f4-7f}	7.6 Mbps	21.4 Mbps	11.3 Mbps	8.1 Mbps

Download Throughput per NAS port type

	min	max	avg	current
nettest.mean {nasPortType: Wireless-802.11}	7.6 Mbps	64.7 Mbps	18.9 Mbps	8.1 Mbps

Figure 3.6: Screenshot from WiFiMon UI: Timeline tab (2/2)

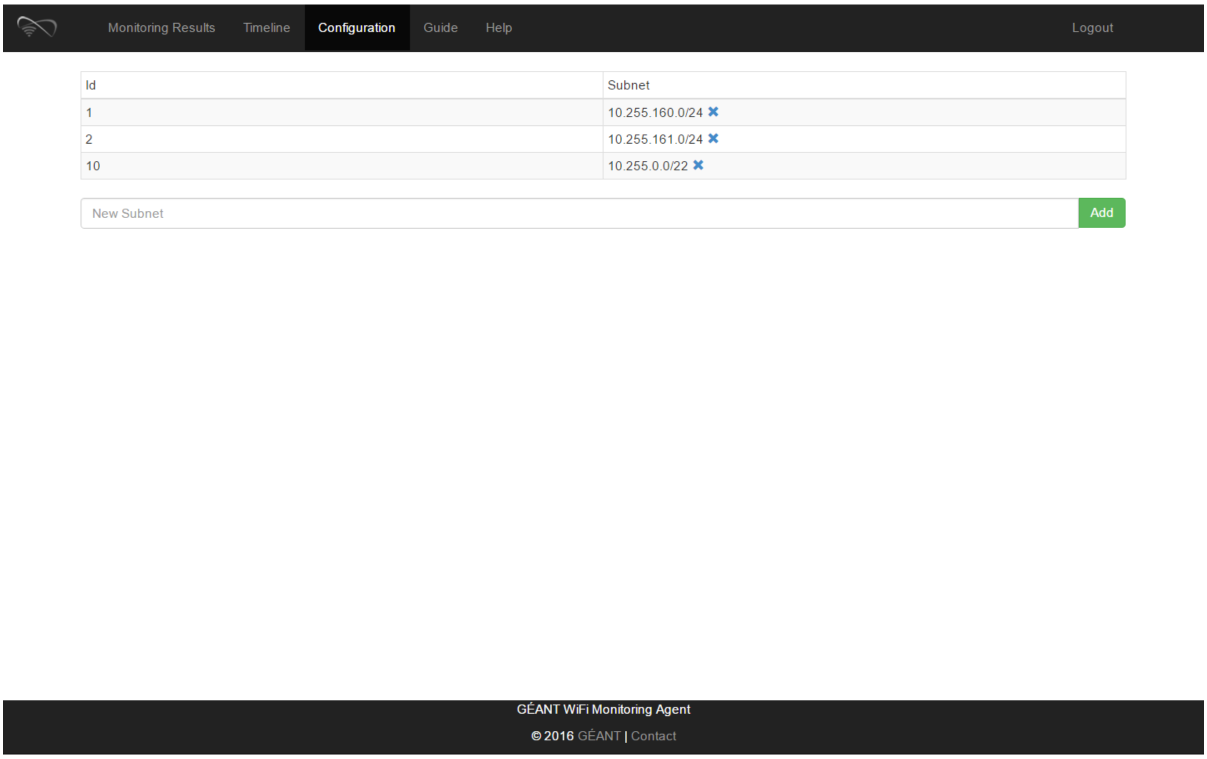

Id	Subnet
1	10.255.160.0/24 ✖
2	10.255.161.0/24 ✖
10	10.255.0.0/22 ✖

Figure 3.7: Screenshot from WiFiMon UI: Configuration tab

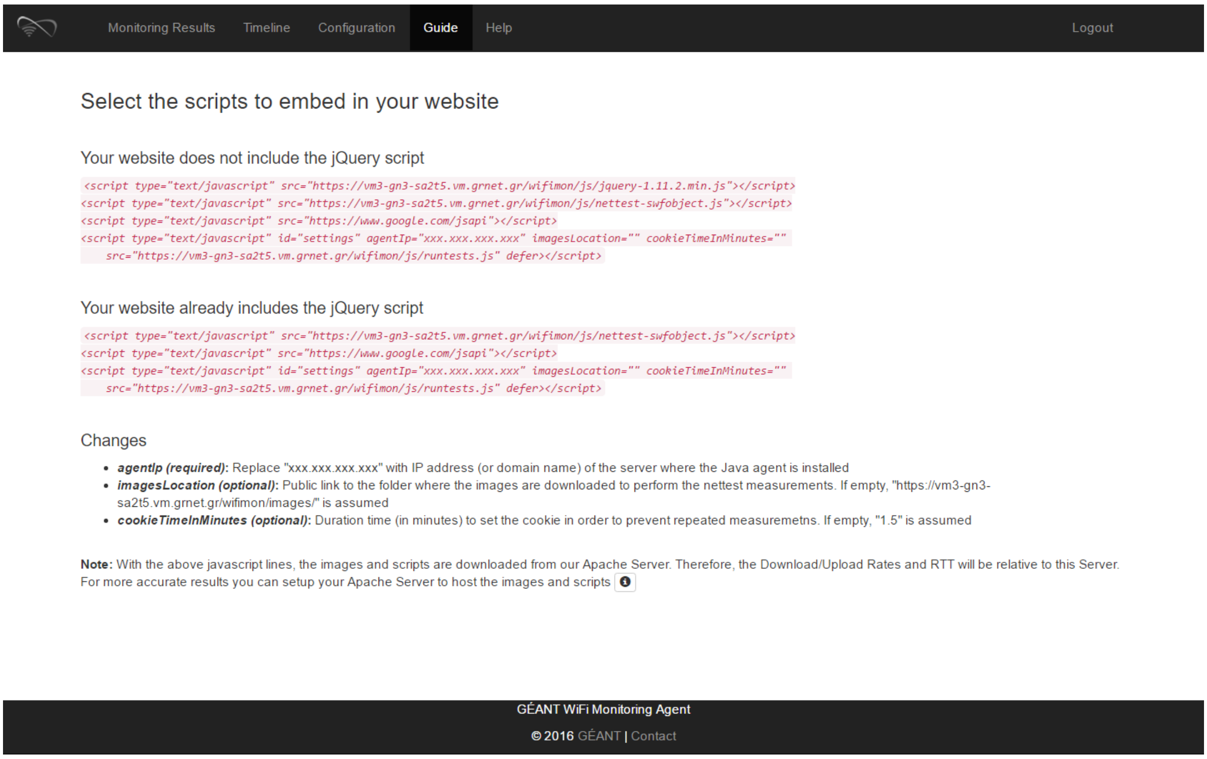

Select the scripts to embed in your website

Your website does not include the jQuery script

```
<script type="text/javascript" src="https://vm3-gn3-sa2t5.vm.grnet.gr/wifimon/js/jquery-1.11.2.min.js"></script>
<script type="text/javascript" src="https://vm3-gn3-sa2t5.vm.grnet.gr/wifimon/js/nettest-swfobject.js"></script>
<script type="text/javascript" src="https://www.google.com/jsapi"></script>
<script type="text/javascript" id="settings" agentIp="xxx.xxx.xxx.xxx" imagesLocation="" cookieTimeInMinutes=""
    src="https://vm3-gn3-sa2t5.vm.grnet.gr/wifimon/js/runtests.js" defer></script>
```

Your website already includes the jQuery script

```
<script type="text/javascript" src="https://vm3-gn3-sa2t5.vm.grnet.gr/wifimon/js/nettest-swfobject.js"></script>
<script type="text/javascript" src="https://www.google.com/jsapi"></script>
<script type="text/javascript" id="settings" agentIp="xxx.xxx.xxx.xxx" imagesLocation="" cookieTimeInMinutes=""
    src="https://vm3-gn3-sa2t5.vm.grnet.gr/wifimon/js/runtests.js" defer></script>
```

Changes

- ***agentIp (required)*:** Replace "xxx.xxx.xxx.xxx" with IP address (or domain name) of the server where the Java agent is installed
- ***imagesLocation (optional)*:** Public link to the folder where the images are downloaded to perform the nettest measurements. If empty, "https://vm3-gn3-sa2t5.vm.grnet.gr/wifimon/images/" is assumed
- ***cookieTimeInMinutes (optional)*:** Duration time (in minutes) to set the cookie in order to prevent repeated measuremetns. If empty, "1.5" is assumed

Note: With the above javascript lines, the images and scripts are downloaded from our Apache Server. Therefore, the Download/Upload Rates and RTT will be relative to this Server. For more accurate results you can setup your Apache Server to host the images and scripts 🛈

Figure 3.8: Screenshot from WiFiMon UI: Guide tab

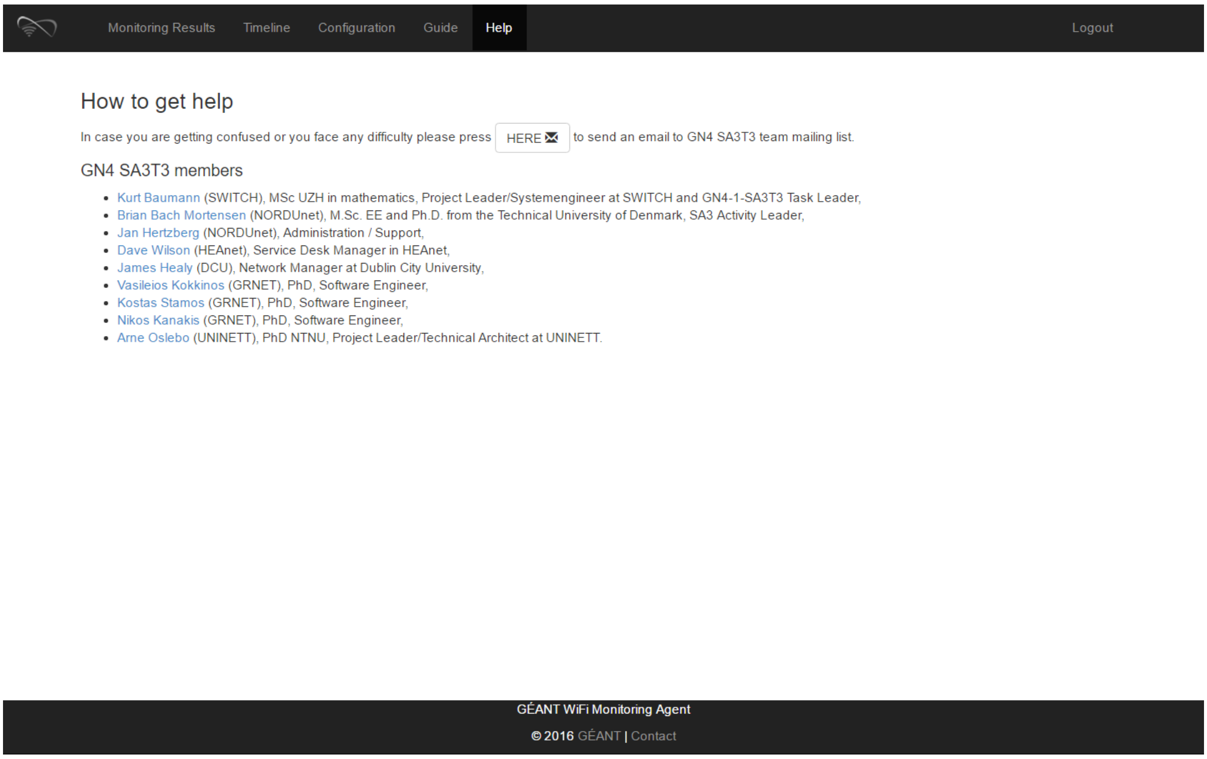

How to get help

In case you are getting confused or you face any difficulty please press HERE ✉ to send an email to GN4 SA3T3 team mailing list.

GN4 SA3T3 members

- Kurt Baumann (SWITCH), MSc UZH in mathematics, Project Leader/Systemengineer at SWITCH and GN4-1-SA3T3 Task Leader,
- Brian Bach Mortensen (NORDUnet), M.Sc. EE and Ph.D. from the Technical University of Denmark, SA3 Activity Leader,
- Jan Hertzberg (NORDUnet), Administration / Support,
- Dave Wilson (HEAnet), Service Desk Manager in HEAnet,
- James Healy (DCU), Network Manager at Dublin City University,
- Vasileios Kokkinos (GRNET), PhD, Software Engineer,
- Kostas Stamos (GRNET), PhD, Software Engineer,
- Nikos Kanakis (GRNET), PhD, Software Engineer,
- Arne Oslebo (UNINETT), PhD NTNU, Project Leader/Technical Architect at UNINETT.

GÉANT WiFi Monitoring Agent

© 2016 GÉANT | Contact

Figure 3.9: Screenshot from WiFiMon UI: Help

More detailed information about this testbed will be made available on the SA3 T3 Confluence wiki in due course.

4 Wireless PM&V Demonstrations

This chapter discusses relevant forms of implementation and valid test cases of WPM&V on various campuses and conference locations around the world. Valid test cases include: the DCU campus, the implementation at IT Carlow, HEAnet and conferences such as TNC2015, APAN40 and NORDUnet technical workshops [APAN40], [NORDUNETTW]. Data is measured and analysed from: each conference/campus log-files; values from defined collectors, such as DHCP logs, RADIUS logs (if available) (see Section 3.2.1); and AP-IDs in a triangulation - with JavaScripts used NetTest on most frequent web-sources.

4.1 Dublin City University (DCU)

The WiFi infrastructure at DCU comprises a pair of Motorola RFS7000 wireless controllers, with a range of Motorola-dependant and lightweight access points. DCU uses FreeRADIUS running on a Linux platform, with an ISC DHCP server on Linux. There are over 800 Wireless Access Points (WAPs), across multiple campuses. The solution is implemented to use non-tunnelled bridge mode, where client devices are placed on a local switch VLAN after authentication based on a radius attribute, which distinguishes categories of users. This allows authenticated staff to be dropped into staff VLANs, students into student VLANs and external (visiting) eduroam users into external user VLANs. eduroam is the mechanism for all WiFi authentication on the DCU campus(s), including all local users. RADIUS logs contain an AP identifier, which can be matched to a room/corridor description on the controller. This allows location based performance verification.

On Wednseday, 20 May 2015, DCU performed pilot tests to determine whether it was possible to measure performance metrics, such as the download and upload rates and round-trip time, of the wireless network, via JavaScript, and whether these measurements can be correlated with the information contained in the RADIUS and DHCP logs. At the same time, there was the challenge of distributed locations, so to enrol the measurement schema over multiple locations.

4.1.1 Procedure

Measurement tests: While roaming, a number of clients executed the NetTest on a number of occasions over a period of time from different locations across the campus. Each time NetTest was executed, information about the wireless network performance was stored on the PostgreSQL database (the relational database block of the architecture (see Section 3.2.2). In the initial installation at DCU, geolocation information from end users could also be extracted and stored to the database. The complete set of data stored at the database can be found at the following link [DATA]. A query was triggered in order to automatically populate the data (timestamp, download rate, upload rate, ping time, latitude, longitude, client IP) from the individual measurements (the data sources block of the architecture (see Section 3.2.1) to the PostgreSQL database.

Log files: Data collected from DCU involved a manual extraction from a FreeRADIUS server of the authN/Z-detail file along with the dhcp.log file from the DHCP server. The next step concerned the extraction of the information of the Radius and DHCP logs. To this direction two distinct bash scripts were developed that enabled the transformation of these logs into SQL files. Next, these files were manually inserted in a PostgreSQL database, and are now available at the following links [FILES]. It is worth mentioning that the timestamps of logs needed to be modified to ensure that entries corresponded to UTC time.

Correlation: The correlation of the measurements data with the Radius and DHCP logs was performed with a query (the query and report generation block of the architecture (see Section 3.2.4). The query actually joined the three individual database tables, taking into account the following constrains:

- The client IP in the measurements data should match the client IP from the DHCP logs.
- The client MAC address in the DHCP logs should match the client MAC in the RADIUS logs.
- Regarding the timestamps of the entries, the following condition should be true after taking into account the entries time sequence: RADIUS_timestamp < DHCP_timestamp < Measurement_timestamp.

4.1.2 Results

The results after the correlation are available at [RESULTS], while some initial charts are available at [INITALCHARTS]. Detailed presentation of the results is included in the following sections.

4.1.2.1 *Measurement Statistics*

A total of 154 performance tests were recorded where it was possible to associate with an access point on the DCU campus. These measurements were only triggered by GN4-1 SA3 Task 3 members (three devices in total) after they visited a test page that had the NetTest JavaScripts embedded. End-Users roaming in the DCU campus executed the test measurements on a number of occasions, over a period of time, from different locations across the campus.

Some statistics regarding the number of measurements performed by each client IP, client MAC and AP MAC are presented in the following figures.

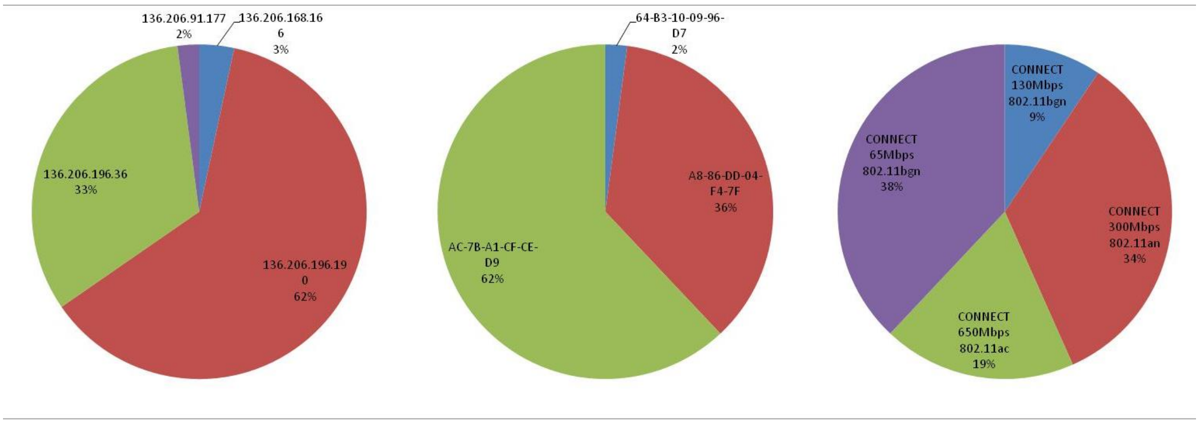

Figure 4.1: DCU - Measurement statistics per Client IP (left), Client MAC (middle), Connect Info (right)

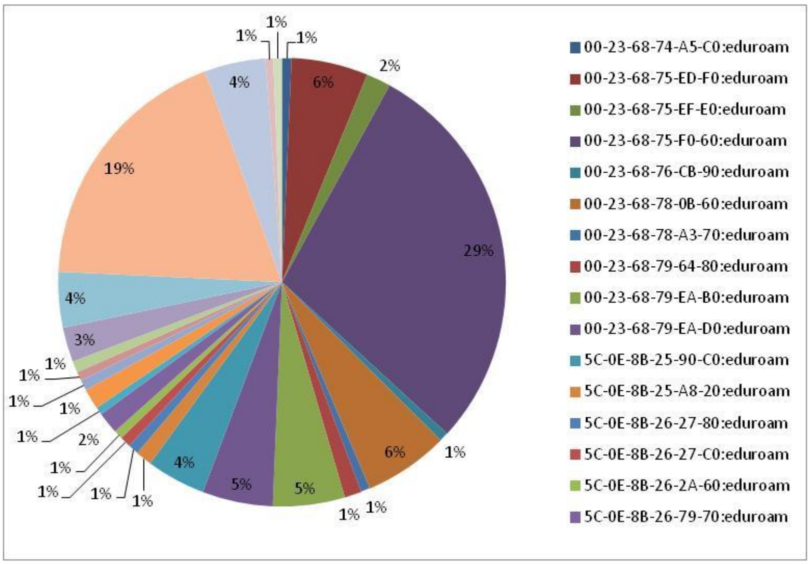

Figure 4.2: DCU - Measurement statistics per AP MAC

4.1.2.2 *Google Maps with Measurements in DCU*

The geolocation option enabled in the initial deployment provided an opportunity to visualise the measurements performed in the DCU campus. Figure 4.3 presents a screenshot with the performance tests that took place in DCU. It is clear that the measurements were triggered from different locations across the campus.

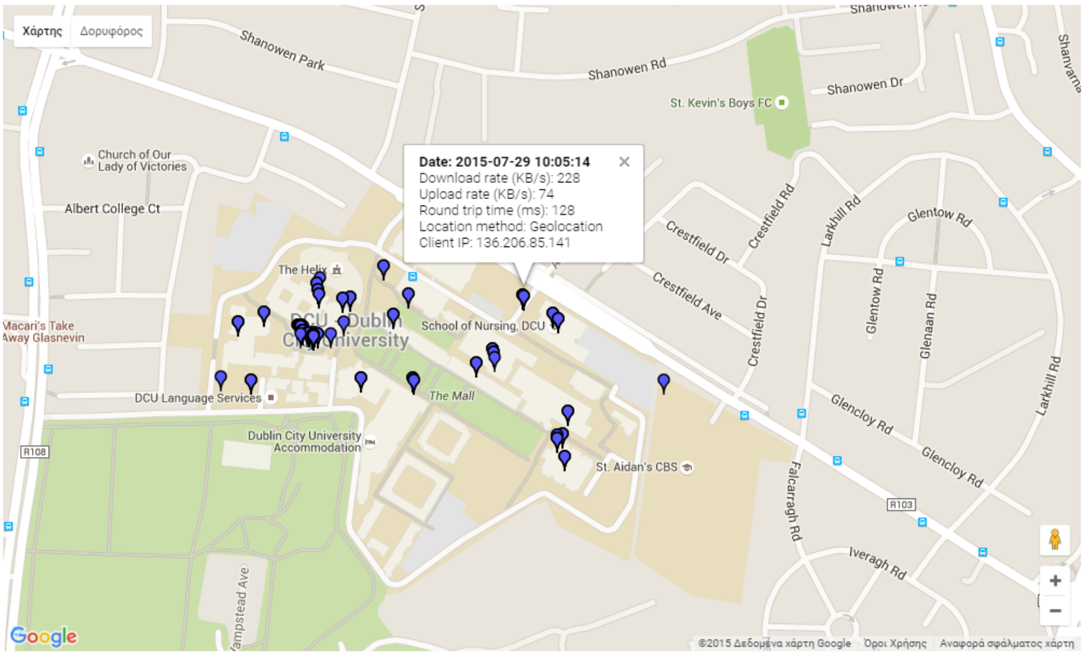

Figure 4.3: DCU - Map with measurement tests

4.1.2.3 *Correlation of Measurements*

This section includes the figures after the correlation of the download rate, upload rate and RTT with the Client IP, Client MAC, AP MAC, and connection speed and technology information (Connect Info). Generally, the download and upload rates show great variation, ranging from 16 KB/s to 9300 KB/s in the case of download and 16 KB/s to 3800 KB/s for the upload. However, these variations may be due to the wireless technology (e.g. 300Mbps 802.11an, 650Mbps 802.11ac, 130Mbps 802.11bgn) and the user's distance from the AP during the measurement.

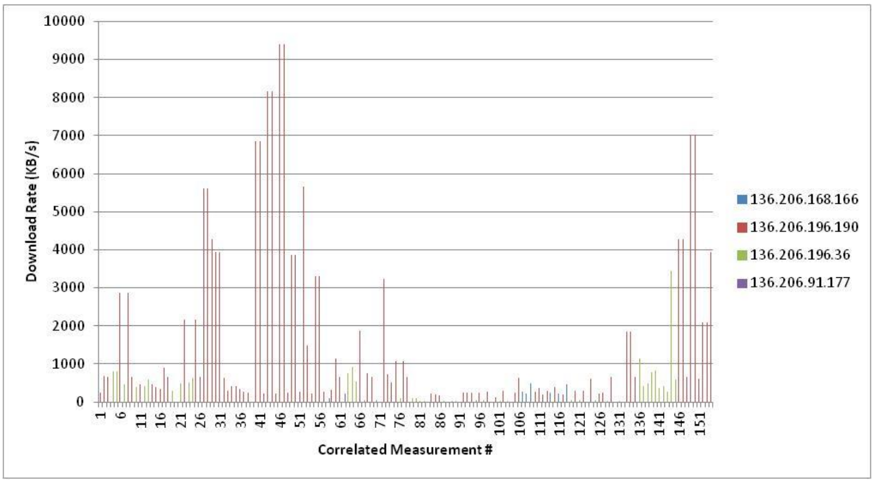

Figure 4.4: DCU - Correlation of Download rate with Client IP

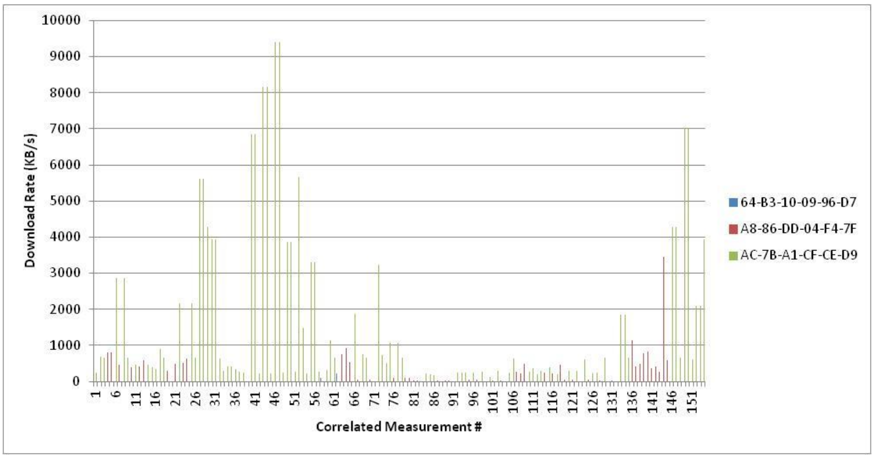

Figure 4.5: DCU - Correlation of Download rate with Client MAC

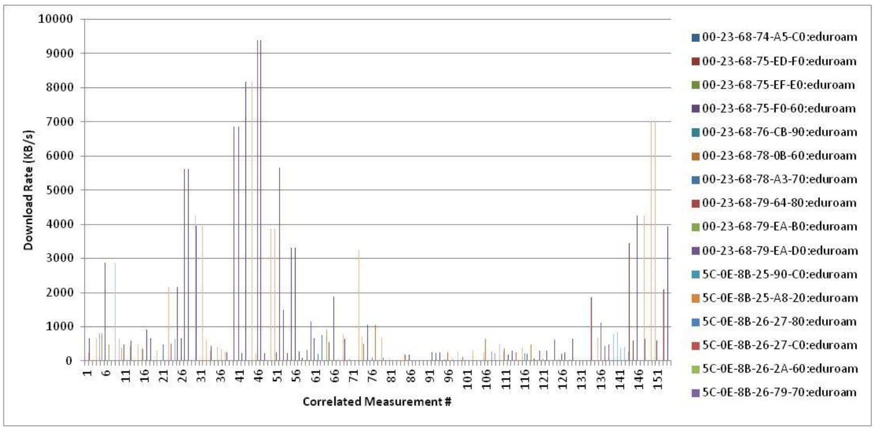

Figure 4.6: DCU - Correlation of Download rate with AP MAC

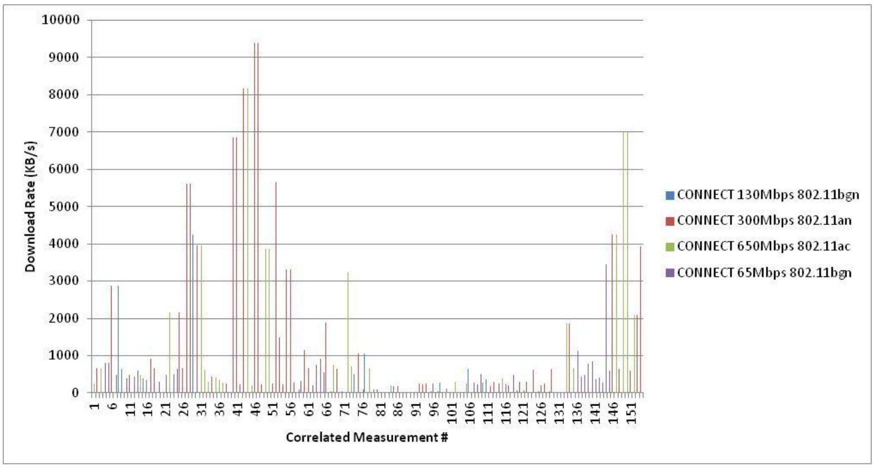

Figure 4.7: DCU - Correlation of Download rate with Connect Info

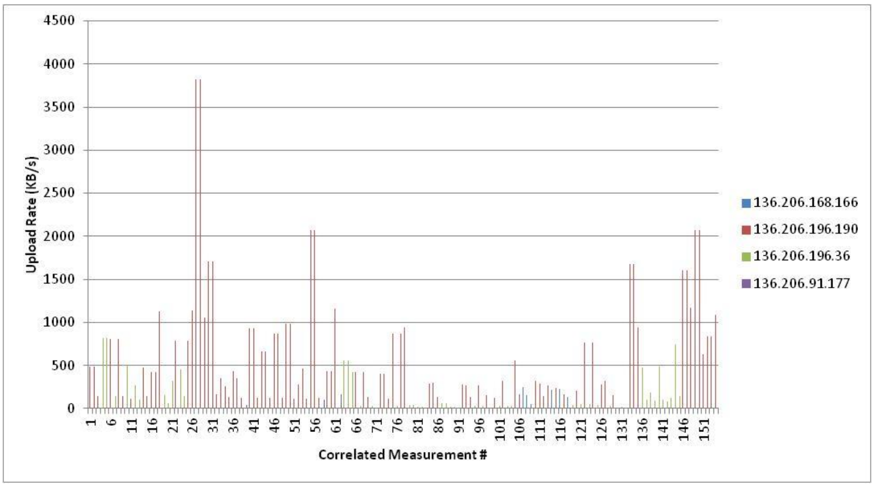

Figure 4.8: DCU - Correlation of Upload rate with Client IP

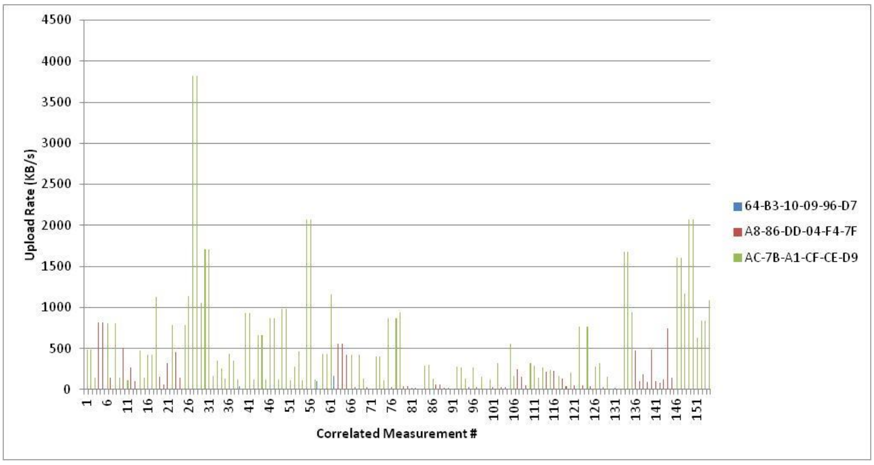

Figure 4.9: DCU - Correlation of Upload rate with Client MAC

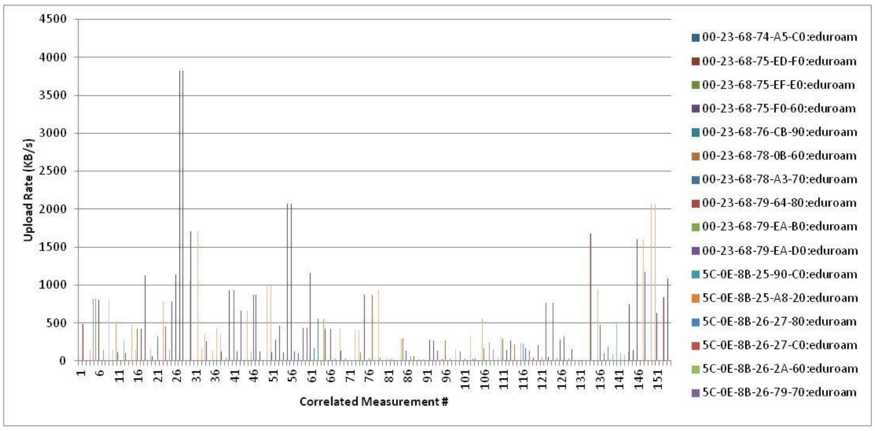

Figure 4.10: DCU - Correlation of Upload rate with AP MAC

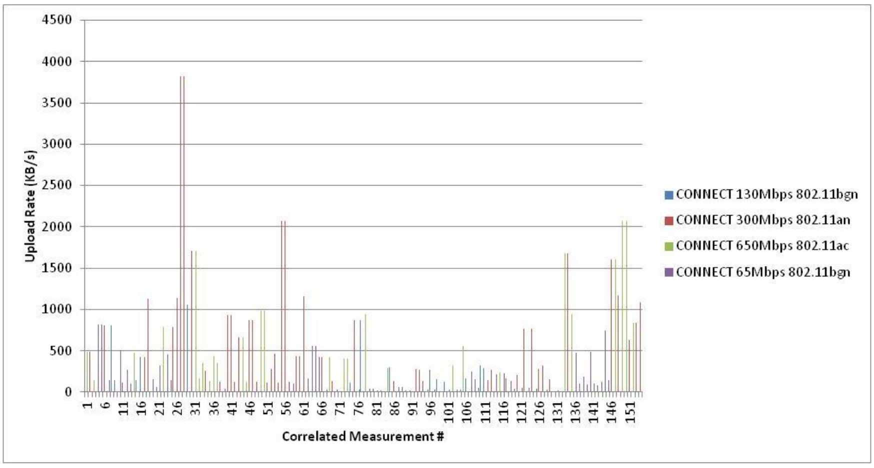

Figure 4.11: DCU - Correlation of Upload rate with Connect Info

One of the most important parts of the measurements is the correlation of the round-trip time (RTT) with the information provided by the log files. From the following figures we see that the RTT ranges between 31.5 ms and 170.5 ms, where the highest values are observed for measurements where both the download and the upload rates are relatively low. In addition, the figures reveal that the technology of the wireless network has a direct impact on the RTT performance. Indeed, the majority of the high RTT values were observed when the user was connected to a low speed AP, i.e. “65Mbit/s 802.11bgn”. On the other hand, since the measurements server is located in Athens, a temporary problem on the path could also significantly affect the RTT measurements.

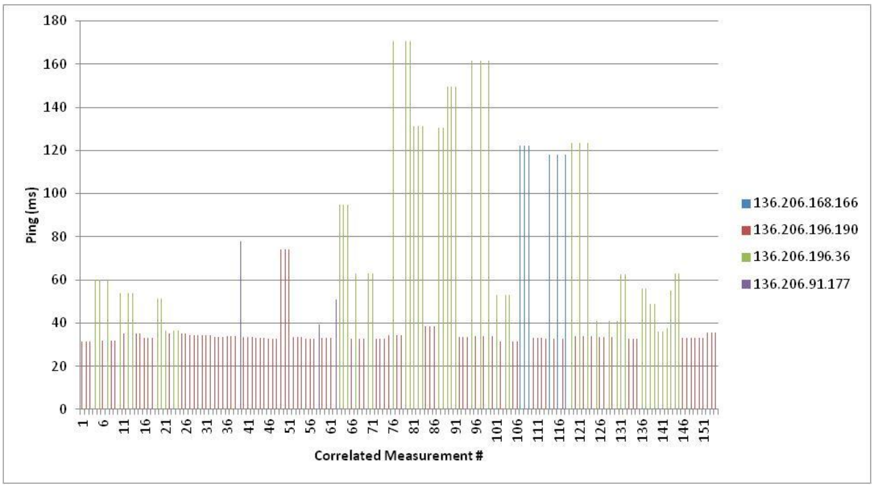

Figure 4.12: DCU - Correlation of RTT with Client IP

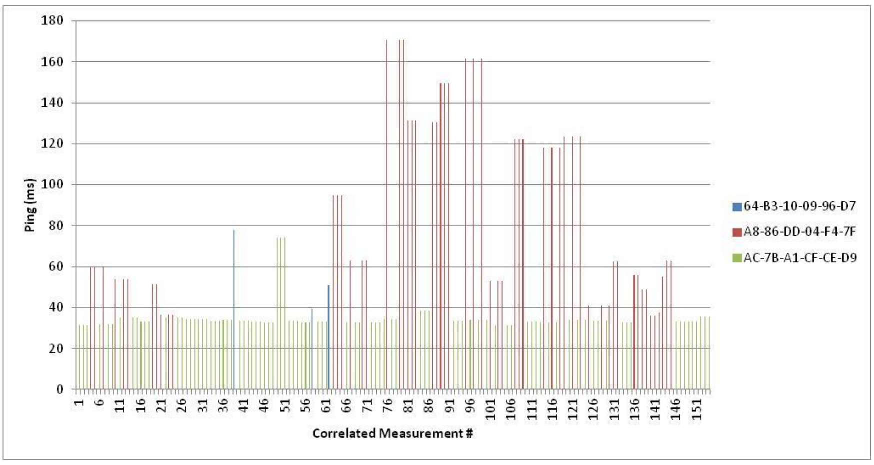

Figure 4.13: DCU - Correlation of RTT with Client MAC

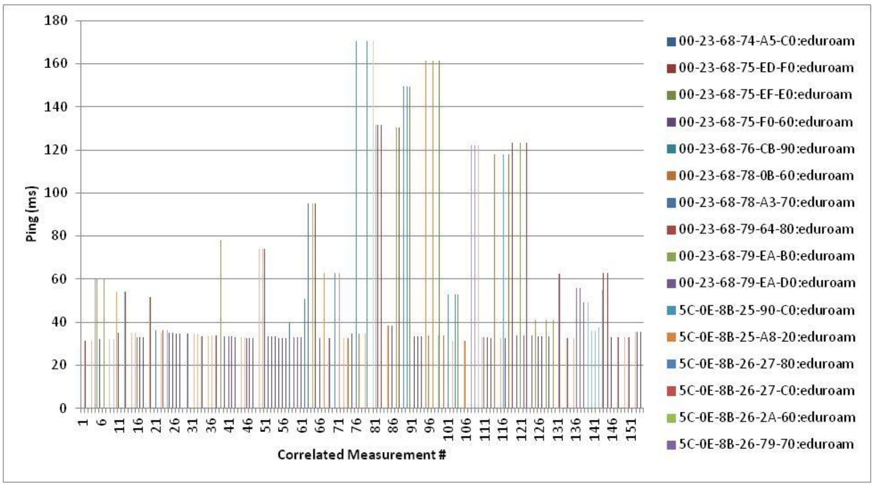

Figure 4.14: DCU - Correlation of RTT with AP MAC

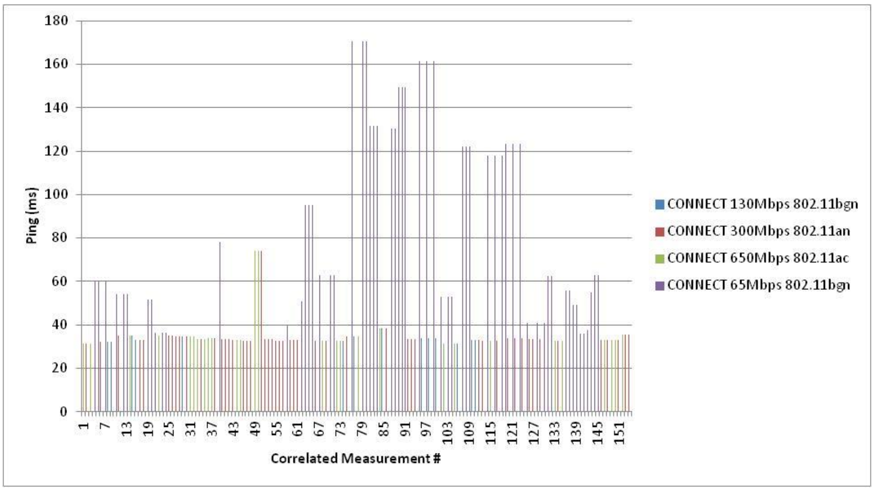

Figure 4.15: DCU - Correlation of RTT with Connect Info

4.1.3 Lessons learned, Hypothesis and next Steps

The initial pilot measurements in DCU were in general successful, in the sense that they revealed that it is possible: (i) to perform wireless network performance tests though JavaScript and (ii) to correlate those tests with the information provided in the RADIUS and DHCP logs in order to extract useful information.

Yet, there were some points that needed attention in order to embed those JavaScript on a website (e.g. conference or University website) that is frequently visited by network users to run those performance tests:

- Geolocation information: Even though the geolocation option offered the opportunity to visualise the measurements performed, it requires user intervention – i.e. users have to enable the geolocation in their browsers (see examples of transparency in Section 3.5).

Next step: The geolocation option has to be removed in order to avoid user intervention.

- Number of tests: For the pilot tests in DCU, when a user visited the test page that had the NetTest JavaScripts embedded, there was no restriction regarding the number of tests performed. This means, that if the user refreshed the test page, a new test was triggered and the sample images (with size up to 5MB) were downloaded again. For deployments in conference venues or in University campuses, the number of tests should be restricted in order to avoid the wireless network overloading.

Next step: NetTest should set a cookie with a configurable expiry time (e.g. 1 hour) to make sure that the same client doesn't run the test repeatedly (see subsection 4.2).

- Additional Information: During each measurement, the following parameters were stored in our database: measurement timestamp, download rate, upload rate, RTT, user geolocation information, client ip. Apart from the above parameters, information regarding the name, version and platform of the browser could provide a useful insight in the performance tests.

Next step: Include userAgent property to store the value of the user-agent header sent by the browser to the server.

There may be a time-line problem in a campus environment that may confound some of the measurements. The scenario being that a client may associate with many different access points on a campus over the course of a day and create many log entries without ever running the NetTest; thus the client may then be triggered to run the NetTest. Only at that point we are interested in the data from our logs, and we need to extract only the most recent (relevant) entries from those logs and to correlate them.

A user/client associates with an AP that is set to challenge with authentication (as per eduroam 802.1x and Radius) - on successful authentication (or not) RADIUS log entries are created at that point and no further entries if the user remains connected. Also, as part of this process and immediately following the user/client requests an IP address and handshake occurs and the client receives an IP address (a series of DHCP log entries created for that client at that point in time). Depending on the lease time, no further entries are created for a period of time. The next time an entry is created is either at the end of the DHCP lease period or when a client disconnects and then reconnects.

The missing piece may be the accuracy of the AP location identifier (called station id) that the user is attached to in a roaming scenario before the NetTest is running - but perhaps that issue will never be fully solved, other than there will be a correct entry in enough cases (the key phrase being 'enough cases') to aggregate the results into meaningful information. What is meant by roaming is that enterprise wireless solutions allow a user to roam without the need for re-authentication or initiating a new DHCP request. This is the case for the solution implemented in DCU. If re-authentication is occurring as a user roams we get better accuracy – e.g. IP address is the same and we get an accurate Radius entry for the new AP location. Investigations in other campus environments may allow us to further clarify these points.

4.2 TNC2015

When a user reports a performance problem with WiFi, the data needed to troubleshoot is not often available. If the problem affects a subset of users, or occurs only in particular circumstances (however frequent), then our usual approach of attempting to replicate the problem often fails. The wireless infrastructure is set-up on the IEEE802.11abgn.

4.2.1 Procedure

Our procedure comes in two parts:

1. Use JavaScript on a website that is frequently visited by a network's users to run non-intrusive performance tests of that network.
2. Separate the tests by access point using RADIUS data or other means.

We then use this data to try to build a picture of the performance experienced by users on the network, broken down by access point (so that problems can be traced to an individual room or AP.)

Our objective at TNC2015 was to try to validate this procedure by attempting to get the tests run by a number of participants at the conference, and then coordinate these with their respective access points. If this was successful, data would be obtained that could be analysed to understand what results were reasonable to expect, and what the next steps might be.

4.2.2 Results

Between 10.15am on Friday, 12 June and 10.40 on Friday, 19 June, a total of 1713 performance tests were recorded that were able to be associated with an access point on the TNC15 site network. An initial, rudimentary analysis was provided during the lightning talks session on Wednesday, 17 June [ANALYSE]. The detailed results are included below and an overview is presented in the CONNECT article "Wireless Drowedsourced Performance Monitoring and Verification, A First Trial: https" [CONNECT].

4.2.2.1 *Measurement Statistics*

Some statistics regarding the number of measurements performed by each AP MAC is presented in the following figure.

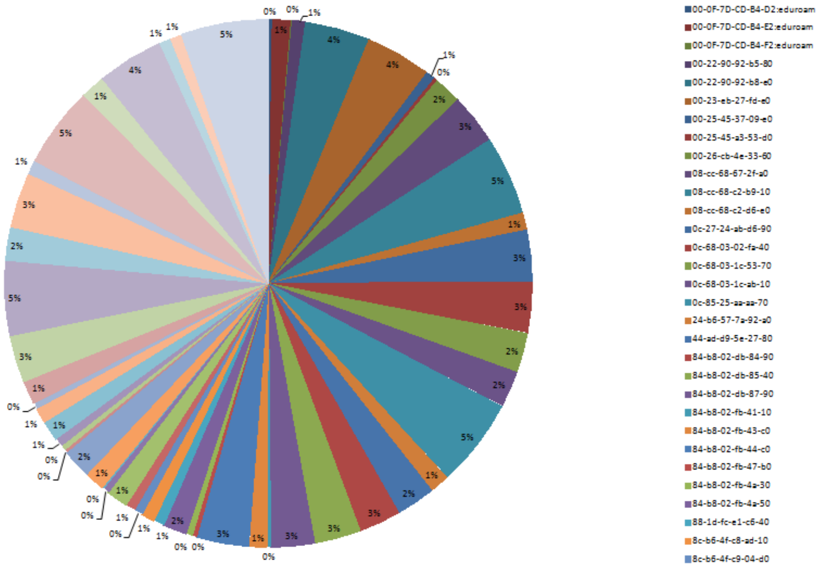

Figure 4.16: TNC2015 - Measurement statistics per AP MAC

4.2.2.2 *Correlation of Measurements*

According to Figure 4.17 and Figure 4.18, the download and upload speed test results were very fluctuating, ranging from 0 to 1736 KB/s (mean download: 662.9 KB/s) and from 0 KB/s to 1068 KB/s respectively (mean upload: 406.5 KB/s). This was however expected as users might be very close or very far away from the APs.

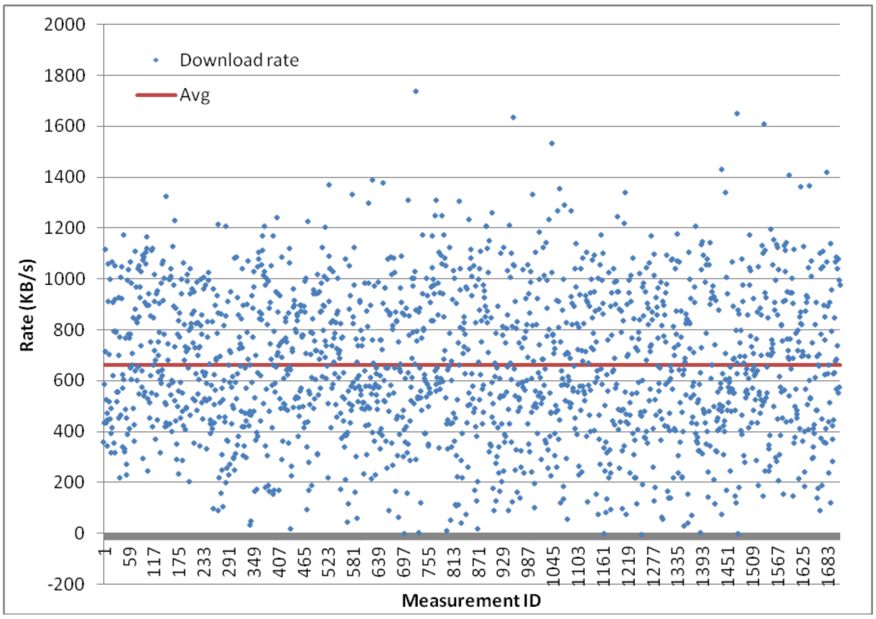

Figure 4.17: TNC2015 - Download rate measurements

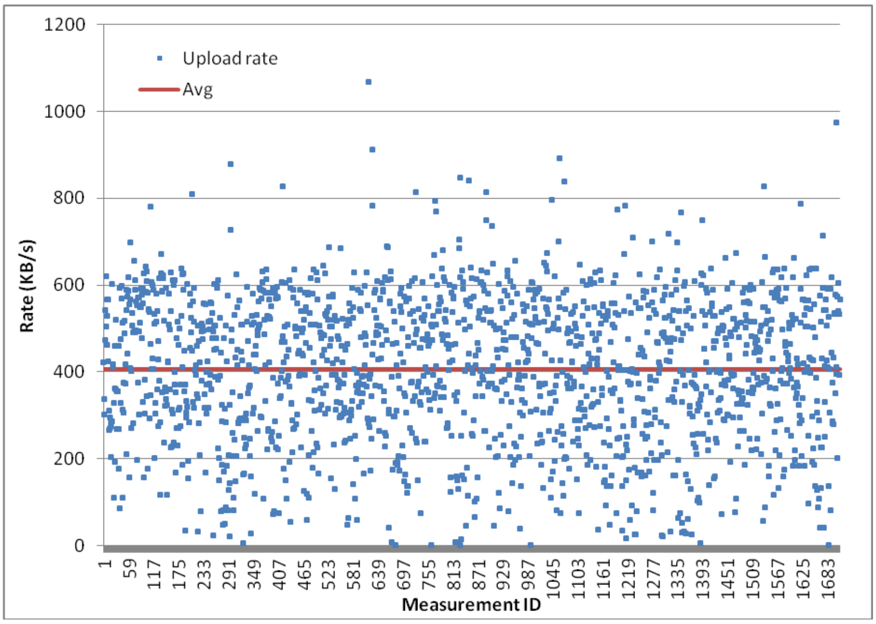

Figure 4.18: TNC2015 - Upload rate measurements

On the other hand, the ping results appeared to be remarkably consistent (Figure 4.19). Most of them were in the range of 40-90 ms and only some spikes were observed that may be caused by temporary network problems.

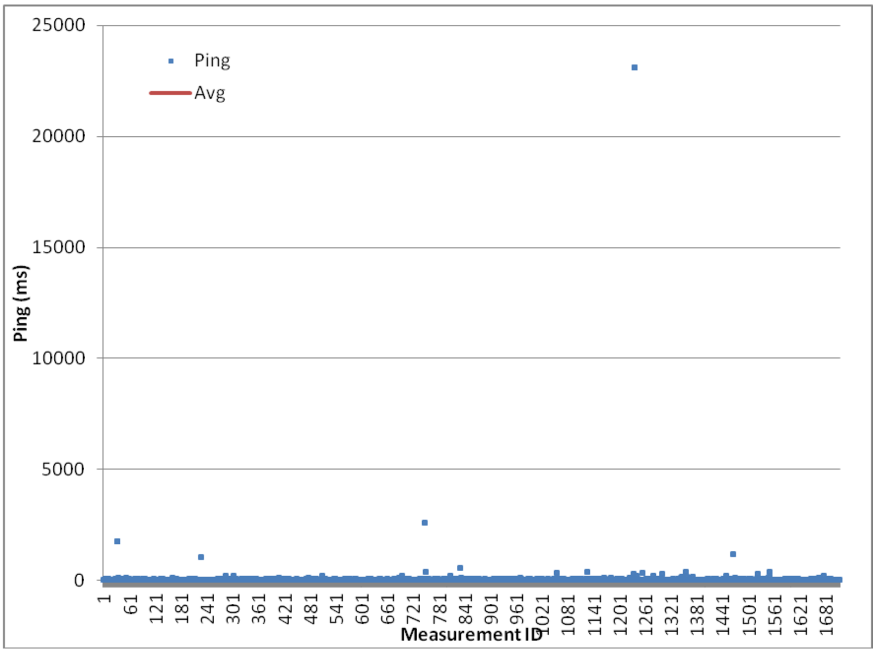

Figure 4.19: TNC2015 - Ping measurements

Below are some figures detailing the correlation of the download rate, upload rate and RTT with the AP MAC and User Agent.

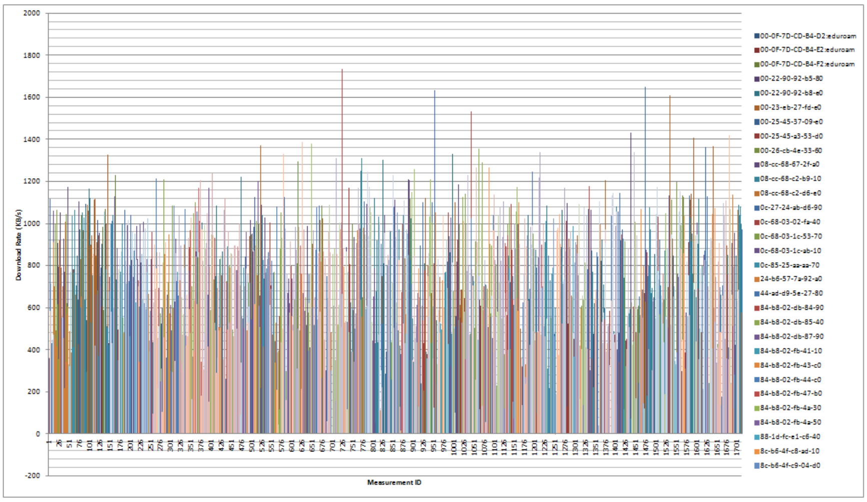

Figure 4.20: TNC2015 - Correlation of Download rate with AP MAC

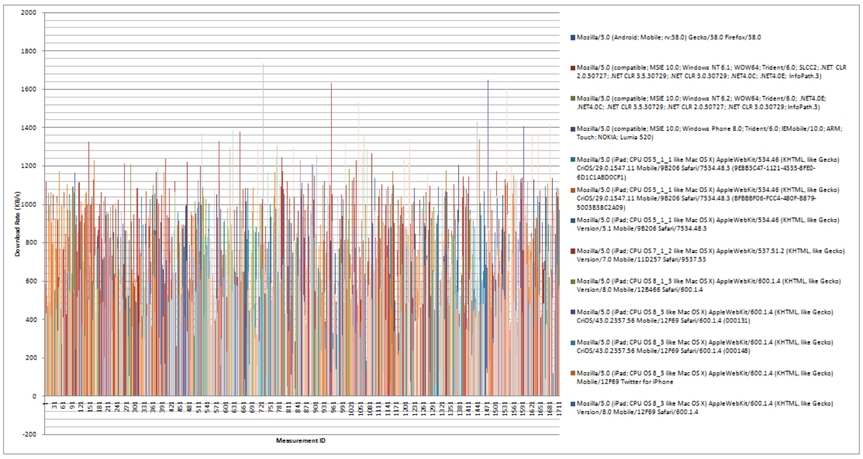

Figure 4.21: TNC2015 - Correlation of Download rate with User Agent

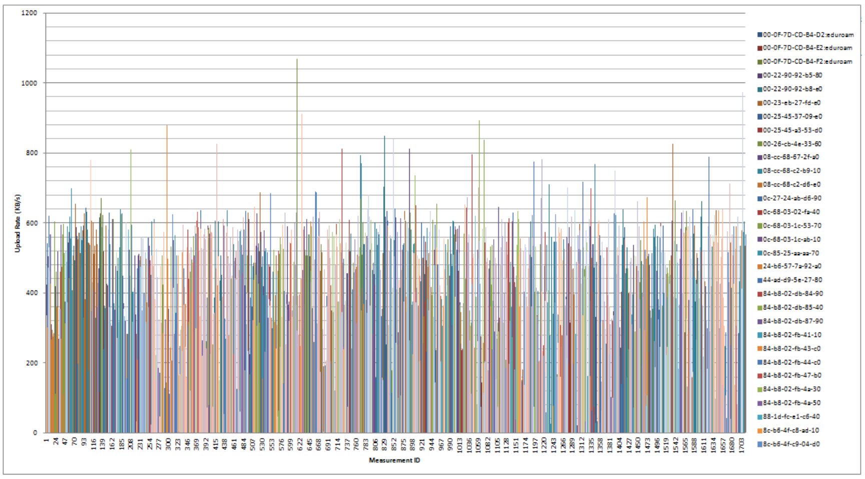

Figure 4.22: TNC2015 - Correlation of Upload rate with AP MAC

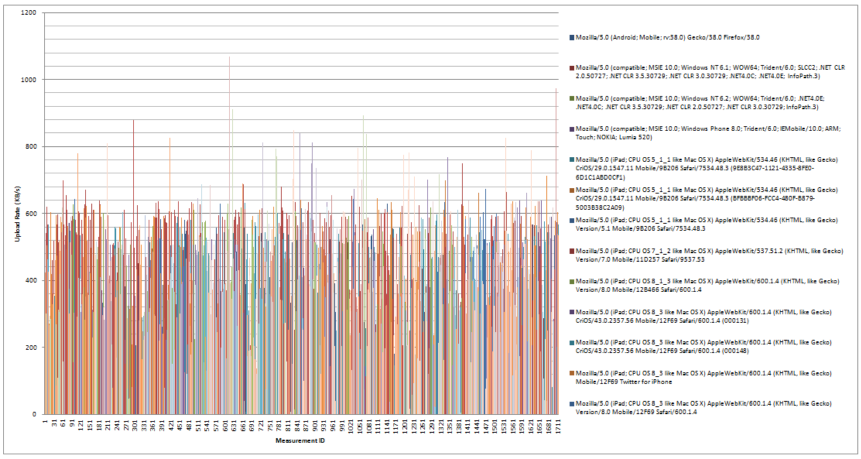

Figure 4.23: TNC2015 - Correlation of Upload rate with User Agent

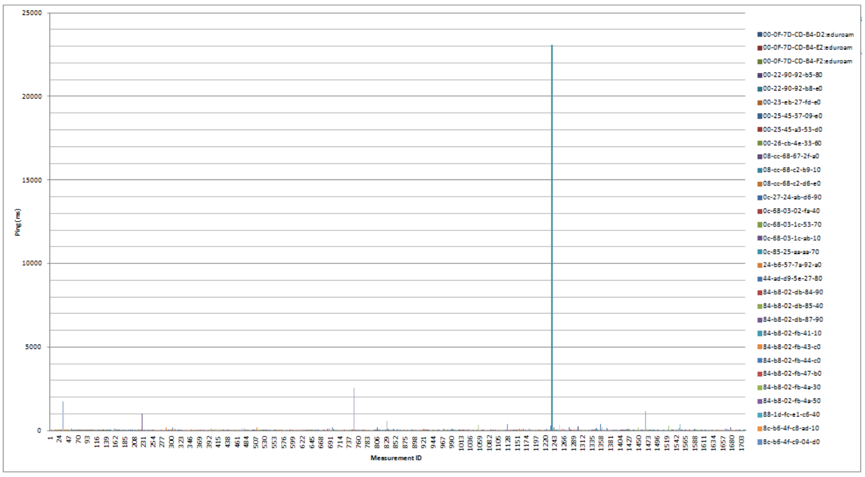

Figure 4.24: TNC2015 - Correlation of RTT with AP MAC

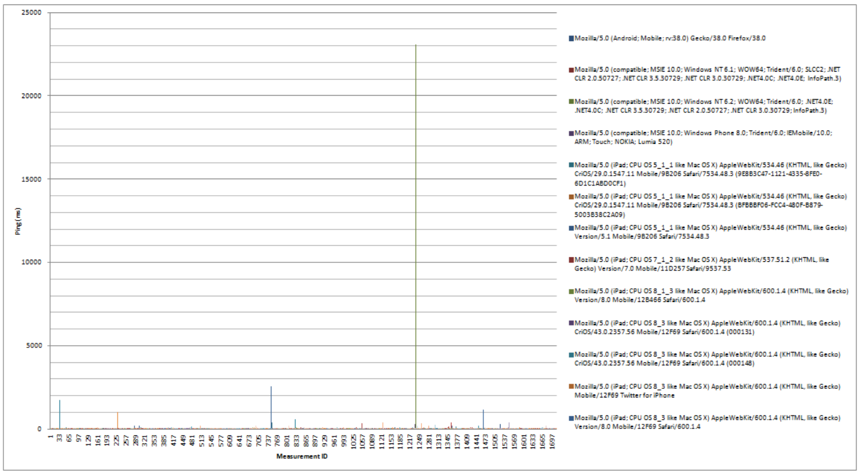

Figure 4.25: TNC2015 - Correlation of RTT with User Agent

4.2.3 Lessons Learned, Hypothesis, Next Steps

Network monitoring utility NetTest sets a cookie with a configurable expiry time (e.g. 1 hour) to make sure that the same client doesn't repeatedly run a test [NETTEST]. Because the JavaScript was loaded from a third-party server - even though the cookie is only used locally and never read by the server - some browsers refused the cookie and results were repeated. Hosting the code locally on the server is entirely possible as long as we are comfortable that it can't be changed without the cooperation of the website owner.

- ***Hypothesis 1:*** *NetTest/Boomerang code could be stored directly on the host website* [BOOMERANG].

The download speed test results were very scattered and it's not clear how accurate they were. We were very conservative with the load caused by the download speed test (1MB images) in case this inadvertently caused problems. But we were not aware of any performance problems caused by load, and the restriction to the host network by IP seemed to work so increasing the load in future should be considered. Boomerang runs tests sequentially, repeatedly, and increasing in load until it either reaches a timeout or completes. Using this information provides an accuracy estimate for the download speeds. A hybrid[1] setup, with both probes and user-provided data, could be self-reinforcing (the probes validate the user results, and vice versa) and extremely powerful.

- ***Hypothesis 2:*** *The load frequency of the download tests can be increased for greater accuracy.*

- ***Hypothesis 3:*** *Objective measurements (e.g. from probes) can be arranged for comparison with the browser results.*

Adding the JavaScript to the TNC2015 website caused two unexpected problems. One was that the version of jquery we loaded conflicted with the different version of jquery already in use on the TNC2015 website. This broke some of the functionality on the website. Fortunately, the NetTest turned out to be compatible with this older version of jquery, and the website admins were willing to reinstate the test with the additional jquery reference removed. Where necessary, it might be possible to load separate versions of jquery and scope them appropriately to avoid conflicts.

- ***Hypothesis 4:*** *Java libraries already loaded on the websites should be checked for conflicts.*

- ***Hypothesis 5:*** *The potential of loading jquery in a scope local to NetTest/Boomerang should be investigated*.

The way the NetTest library was initially loaded on the TNC2015 website caused browsers to block rendering of the page until the library had been loaded and execution completed. This caused unnecessary delay on slower connections, even if the device was outside the IP range that was allowed to run the tests, because the script still had to be downloaded (from a separate site) and executed. After this was discovered, the DEFER parameter to the <SCRIPT> tag was used to ensure that the rest of the page was completed before the performance test was run. However, there was not an opportunity to test if this had an impact on the results. (Boomerang does not require this tag as it operates by waiting for the page to complete loading.)

- ***Hypothesis 6:*** *Options should be explored to ensure that any additions are non-blocking and check the impact of those measures on test results.*

There was a good mix of reports from most popular platforms, but very few from Android. This is suspicious. It is possible that Android users just weren't hitting the website, but we should find out if something else was preventing them from executing the tests. We should also validate the data beyond the initial rudimentary analysis.

- ***Hypothesis 7****: Find out if the number of visits to the TNC2015 website using various platforms from the venue is consistent with the test results*

[1] Hybrid in WPM&V means a correlation between the collected date of HW probes and the user-feedback (JavaScript collected data from web-sources).

- ***Hypothesis 8:** We should analyze the data should be analysed further (e.g. compare Radius START times with performance tests) to see if we can have confidence in the association between IP address and access point.*

TNC2015 used the Conference4me app to distribute the conference program and its used was encouraged by session chairs [CONFERENCE4ME]. There's good reason to believe that, in conference and university situations, user behaviour is skewing away from the web and toward specific mobile apps. Asking users to install a dedicated performance-testing app would defeat our goal of gathering data in a non-intrusive manner. However, where an app already exists, we should look at the possibility of embedding performance tests inside.

- ***Hypothesis 9**: In parallel with browser-based results, the potential for developing code that can be embedded into mobile apps such as Conference4me should be investigated.*

While the download speed results seemed very scattered, the latency results appeared to be remarkably consistent. No results were lower than about 40ms, and most results stayed within 10ms or so of this lower bound. The absolute values are questionable – a ping RTT to the GRNET server from the venue was about 80ms, so it is unknown why NetTest reported a value lower than this – once the cause of this is better understood, it may be that HTTP request latency could be a more appropriate proxy for network performance than download speed.

- ***Hypothesis 10:** HTTP request latency could be a better proxy for network performance than download speed*.

While the test results were loaded into a web-based database in real time, the information to identify the access point for each test had to be extracted from RADIUS accounting logs. Scripts were used to process this data in a batch manner on Monday (at 19:00) and Tuesday (at 16:00, to give enough time to create visualisations for the lightning talk on Wednesday morning.) In an operational environment, where a performance issue needs to be investigated while it is ongoing, it would be necessary to have this information available in real time.

- ***Hypothesis 11**: Operational troubleshooting will need to get the access point data in real time.*

There were two WiFi networks operating at TNC2015, eduroam and tnc15-porto. It was only possible to extract access point data from eduroam, because tnc15-porto didn't use RADIUS. Also, while RADIUS provided an opportunity (quick way) to run tests, on this specific network, not all installations log an IP-address-to-access-point-mapping needed to correlate performance tests with an access point.

It is also likely to encounter networks where IPv4 is NATted, which would lose the IP information if the results server is outside the NAT. Where a network supports IPv6, it may be adequate to test on IPv6 alone; IPv6 can be used to test the local WiFi, while other existing monitoring systems report protocol-specific measurements.

- ***Hypothesis 12**: We need other ways of extracting the IP<=>AP data, including on IPv6.*

- ***Hypothesis 13:*** *It may be adequate to test on IPv6 alone, avoiding the address being obscured by NAT.*

4.3 APAN40

The project team was looking for a further test case, a conference for confirmation the functionality of procedures, collecting and analysing data of the conference WiFi infrastructure. From discussions with SA5 Task 7, it was possible to introduce the measurement schema, the JavaScript deployment on the most frequent web-source, and the main and subpages of the 40th Asia Pacific Advanced Network (APAN40) meeting. The APAN40 meeting took place 10–14 August 2015, at University of Malaya, Kuala Lumpur, Malaysia (UM). The problem statement was still the same as at the TNC2015, not to have the right data for network troubleshooting.

4.3.1 Procedure

The used procedure for collecting the data was the same as at the TNC2015. The webcode to be embedded on the conference website to run NetTest measurements was the same as that used at TNC2015. However, compared to the TNC2015 case, the size of the download images was increased from 1.3MB to 5MB in order to improve accuracy (see Hypothesis 2 in Section 4.2.3).

The main issues encountered were with communications, as we were not onsite. Moreover, after several email exchanges with the responsible persons at UM and after they checked the UM eduroam RADIUS server, only transaction/data from 23 August onwards was captured. Therefore, no correlation could be performed between the measurements and the log files, making it impossible to draw network-related conclusions.

4.3.2 Results

Since the conference organisers did not provide the log files, this section presents charts regarding the measurement tests.

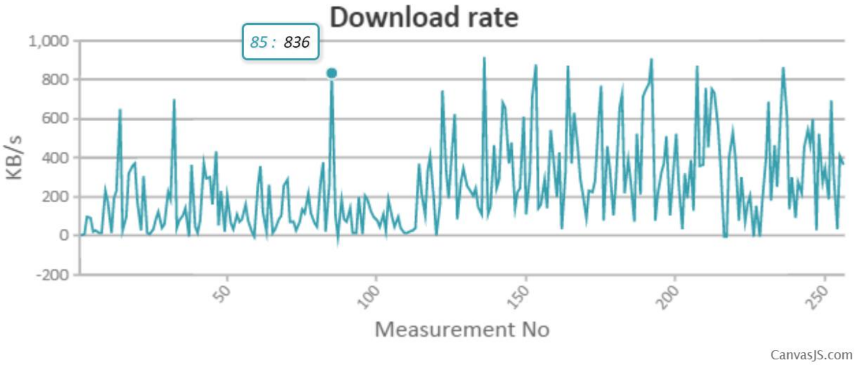

Figure 4.26: APAN40 – Download rate measurements

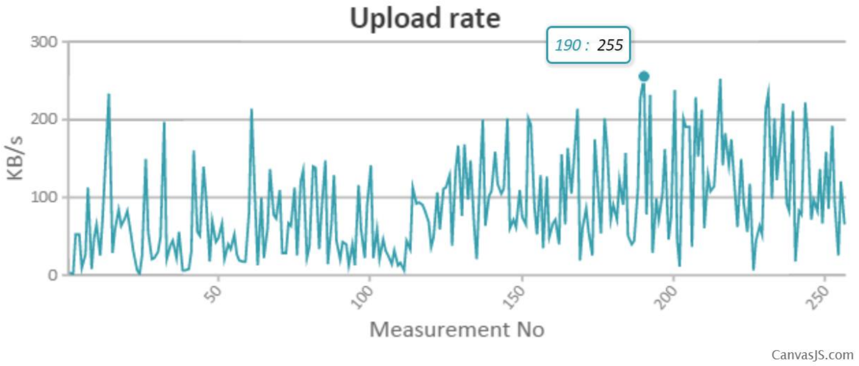

Figure 4.27: APAN40 – Upload rate measurements

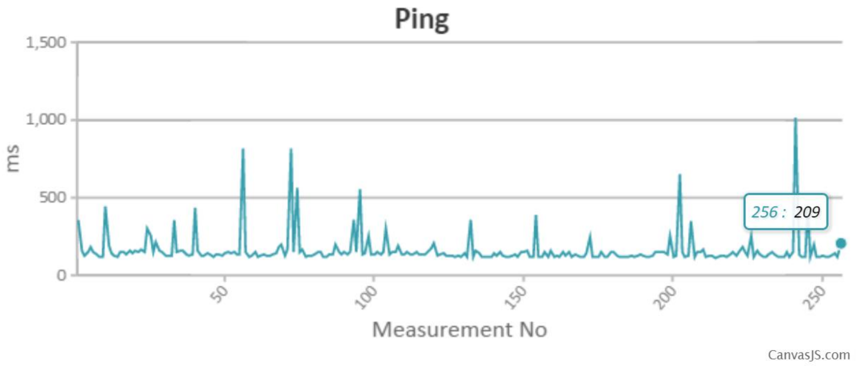

Figure 4.28: APAN40 – Ping measurements

The fact that the conference location was in Malaysia, while the server where the test images were downloaded was located in Athens, explains the fact that the download and upload rates are very scattered, and lower, compared to the previous cases. The latency results are not that scattered, and most of the results stayed close to 200ms. Apart from the distance (and possible bottlenecks in the path), without the log files, it is impossible to draw network-related conclusions.

4.3.3 Lessons learned, Hypothesis, Next Steps

RADIUS Log files are essential when drawing any conclusions on network performance.

4.4 NORDUnet Technical Workshop

The NORDUnet technical workshop (15–17 September 2015) is targetting NRENs and campuses/universities in the Scandinavian/Nordic countries. As part of the piloting and dissemination for WiFiMon the NORDUnet workshop coordinators were contacted.

4.4.1 Procedure

The procedure to complete the live pilot at the workshop was quite simple, and similar to the TNC2015 case:

- Send the JavaScript block that needs to be installed at the workshop pages.
- Find out the IP address range for the WIFI installation in order to configure the access list at the WIFIMon back-end server.
- Get access to Radius/DHCP logs from eduroam and AP controller.
- Test that the JavaScript are properly installed and correlate results from the WIFIMon portal.

There were some differences compared to the TNC2015 case:

- The JavaScript files were installed on the same server that hosts the website, which avoided the third-party cookie problem (see Hypothesis 1 in Section 4.2.3).
- The load of the download images was increased from 1.3MB to 5MB in order to gain accuracy (see Hypothesis 2 in Section 4.2.3).
- The code to trigger the JavaScripts, and therefore the test, was only embedded in the main page and not in subpages, resulting in only a few measurements.
- The IPv4 was NATted and therefore hosts were mapped to one publicly exposed IP address. This resulted in all measurements to be performed by the same IPs, making it impossible to correlate the data (see Hypothesis 12 in Section 4.2.3).

4.4.2 Results

Due to the time presure at the workshop, eduroam was not installed at the venue. This meant that no RADIUS log files were collected. However, a decision was made to continue the testing, even if there was not the chance to correlate the data with AP controller logs. Without proper log files, it is impossible to determine the AP where users are connected. However, any measurements carried out will still indicate basic parameters such as bandwidth (upload/download speed of the JavaScript image 1MB size) and delay (ping - RTTs).

In detail, during the workshop, 105 measurements took place, all from the same public IP. Considering that the people attending might number more than 100, that the JavaScript was only installed at the conference home page, and that users could not repeat the test for 1 hour (cookie had a 1-hour expiry time) which was reasonable. The download and upload speed test results fluctuated from 0 to 2584 KB/s and from 23 KB/s to 808 KB/s, respectively (see Figure 4.29 and Figure 4.30). However, this was expected, as users might be very close or very far away from the APs.

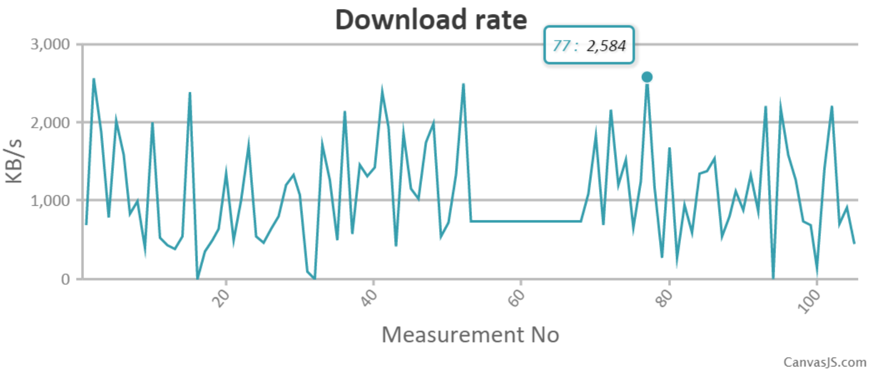

Figure 4.29: NORDUnet – Download rate measurements

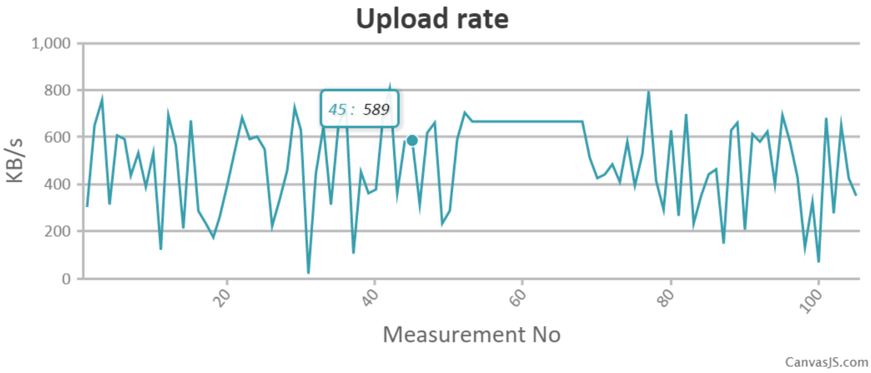

Figure 4.30: NORDUnet – Upload rate measurements

The ping results, on the other hand, appeared to be remarkably consistent (Figure 4.31). Most of them were about 40ms and only some spikes were observed that could have been caused by temporary network problems.

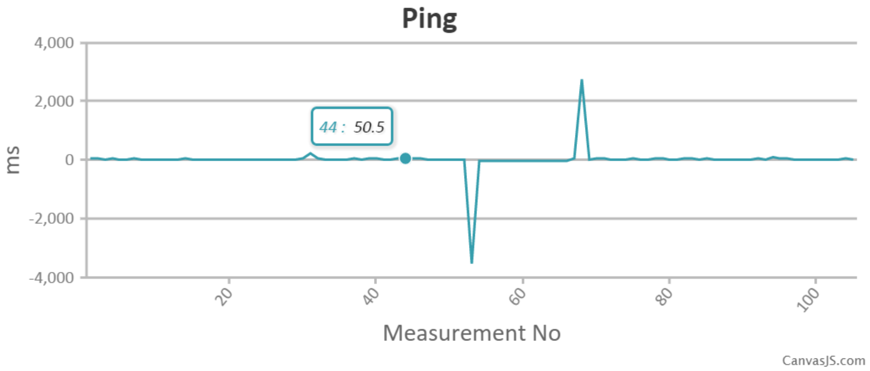

Figure 4.31: NORDUnet - Ping measurements

In order to detect any anomalies and potential service disruptions, it might be best to look at the statistical properties of the crowd sourced measurements rather than the individual measurements.

As a Wi-Fi network operator, it might be useful to have some historical data, including: distributions, mean, and standard deviation. This can be used to tell if problems from users are part of a bigger problem, or only related to the individual client. Furthermore, if a particular client reports problems and is willing to share certain properties of the client (browser, OS etc.) the local support desk may use the measurement to more clearly identify the problem.

4.4.3 Lessons Learned, Hypothesis and Next Steps

The technical coordinators at the NORDUnet workshop were very responsive and rapidly adapted requests for system implementation. However, it was later revealed that they did not have full read/write access to the conference schedule pages, which may have resulted in a significantly lower number of measurements made. In addition, a decrease of the cookie timeout should be considered, in order to obtain measurements more often than the current default of once per hour.

Further, without the log files it is impossible to draw network-related conclusions.

4.5 HEAnet Workshop

The annual HEAnet Conference took place on 12th and 13th November 2015 in Cork, Ireland. The conference was attended by more than 250 delegates from educational institutions across Ireland, a mixture of IT staff/managers and librarians.

WiFiMon was given a 30-minute slot in the main room, and was very well attended. It is clear that there is very strong interest from universities and institutes of technology in this problem, and in our way of solving it. Simultaneously with the talk, we had the WiFiMon-javascript on the HEAnet conference web page and so were able to take measurements from the venue. Since we had the live analysis up and running, we were able to show the results live during the talk itself. The session was quite early in the conference – at noon on the first day – but we still had 42 measurements by the time of the talk itself, which showed the expected spread of download/upload speeds, but very consistent latencies with notable (and investigable) exceptions. This number continued to grow over the course of the event.

Not only was the talk well attended, but it was also exceptionally well rated by the attendees. The two speakers received the second- and third-highest ratings of the conference, out of 59 speakers. The audience ratings for one of the speakers was: 41% Excellent, 30% V Good, 29% Good, 0% Fair, 0% Poor.

The conference also sought comments on the best sessions, asking attendees which session they derived the most benefit from and why. Comments on the WiFiMon session are shown in Appendix E.

The session was recorded and may be found at [HEAnetCONF].

5 Conclusions

The section reflects the concept described in Section 3, and the expertise made by PoCs on the DCU campus, TNC2015 conference, the APAN40 Conference, the NORDUnet technical workshop and the implementation of the WPM&V at a “real” customer, the HEAnet workshop detailed in Section 4. The conclusion also references standardisation from an IETF-LMAP perspective. Finally, a recommendation and future steps for the GN4-1 community in WPM&V closes the Milestone.

5.1 Demonstrations

A number of demonstrations were improved, step by step, by the concept described in Section 3 and Section 4 as PoCs. The implementation of the JavaScript deployment on essential, highly frequented web sources, running the NetTest/boomerang server directly on them and having eduroam in place simplifies performance measurement and verification. Access to needed log files, RADIUS authN,Z/DHCP logs and having AP’s IDs allows correlations and the ability to manifest a quality statement of the performance behaviour on the wireless campus/conference network environment. Without availability of the mentioned log files, it is impossible to draw network-related conclusions. From a statistical point of view, the crowd of end-users and, as a consequence, the huge amount of measurements allows an accurate validation of collected data sets. To introduce a hybrid approach would complete the WPM&V picture on multiple locations/campuses.

5.2 Standards: IETF-LMAP and GN4-1 WiFiMon

One of the big motivations for using LMAP is to be able to standardise how measurements are carried out so that it is possible to compare results from different systems. Defining a standardised control protocol also allows equipment from different vendors to be used in the same measurement infrastructure.

If a hybrid solution that includes monitoring probes is selected as the optimal means of measurement, LMAP should be considered as a protocol for controlling the measurements. Much work has gone into the development of LMAP to ensure that the protocol can be used to control all of the different measurement types. LMAP is designed for large-scale measurements, and contains many features for making sure that the measurements do not negatively impact regular network traffic.

Researchers from Jacobs University in Germany are working on a reference implementation of LMAP. This is still in early stages. So if we decide to use LMAP we might have to do some development ourselves. While a full implementation of LMAP will be challenging to implement, a simple implementation that only implements the features needed by WiFiMon should not be too much work.

Even if LMAP is not selected, the work on the measurement registry should be closely followed. If any metrics are listed in this registry, effort should be made to make sure metrics are measured according to the RFCs, where the metrics are defined. If the collected metrics are not included in the registry, we should consider if the metrics are general and important enough to try and standardise in IETF [RFC7398].

5.3 Recommendation

The experience gained so far – mainly through the implementation of our WPM&V on various campuses and conference locations – showed that it is possible to measure specific parameters of a wireless network through JavaScript, and to correlate these measurements with the information provided by the various log files. However, several steps should be made in order to provide a complete and automated service for monitoring and validating the wireless networks' performance. The recommended steps have been grouped into the following categories:

- **Measurements verification:** It is necessary to verify whether the measurements obtained through JavaScripts embedded in frequently visited pages are accurate. This could be accomplished by installing monitoring probes in different rooms of a campus/conference and at different distances from the APs. In this hybrid approach, the monitoring probes could measure the quality of the wireless network in parallel with the measurements that take place automatically through JavaScript. It is worth mentioning that a first attempt towards this direction was made during the NORDUnet technical workshop. Unfortunately, this attempt was not successful, because all the users were served by the wireless network of the hotel that hosted the workshop, which did not support eduroam. In every case, the verification of the accuracy results accuracy will be one of the first steps needed in future work.

- **Service automation:** In order to have a useful, attractive and easy-to-use service, the service should require minimum or even no intervention by the wireless network administrators, i.e. to be as automated as possible. For the time being, the only automated process is the storage of the user's performance measurements in the Relational Database. Alternatively, the extraction of the information contained in the log files involves a manual procedure (see Section 4.1.1), while the correlation of the measurements the log files involves a semi-automated process where the administrator has to choose the corresponding measurements and log file tables from the Relational Database (from a list of all available log files). It is therefore recommended to fully automate the data retrieval/storage to the RDB process and the correlation process. Some steps have already been made towards automating the service. In detail, the WiFiMon team has investigated several alternatives for data ingestion, real time analytics, batch processing, and data export, with Spring XD being the most promising candidate.

- **GUI development:** Network administrators will be the end-users of the service, and therefore a Web-UI should be implemented to allow them investigate the collected performance data, and in turn, to check the status of their wireless network. A first version of this webGUI [WEBGUI] has already been implemented, see screenshots in Appendix D. It is recommended to enhance the functionality of the implemented Web-GUI with new features that will enable, for example:

  - Federated user login via eduGAIN.
  - Display min-max-mean values in measurements.
  - Specific AP discovery.
  - Location selection and graphical display of historical performance over time.

- **Mobile devices/smartphone apps deployment:** Experience gained from the TNC2015 saw end-users satisfy their communication needs through the use of smart devices/phones. Most of this communication was not carried out using browser-based functions, but instead, with conference applications. The solution to reach the end-users could be about app deployment on Android and Apple-related products. As a first step, further investigation is recommended to investigate Android-based app deployments for the collection of data sets.

Appendix A PM&V Process Description

This section presents further description of the four (4) main functions (outlined in Section 1):

1. **Definition Metrics** – INPUT: Granularity of measurement scope / OUTPUT: Key Operational Performance Indicators (KOPIs)

- ○ The KOPIs have to be defined according the service topology. As a minimal set-up of KOPI will be used in SA3, measurements will be grouped into active- and passive monitoring metrics, as listed in Figure 1.1. Numbers of KOPIs can vary from use case to use case and priorities, coming up in an Operational Level Agreements (OLAs). The definition of metrics is essential for collecting appropriate measurement data and its verification, which is a negotiation with the Service Owners (SO) and the Operations (Ops).

2. **DESIGN PM&V Architecture** – INPUT: Defined Metrics in SW, HW / OUTPUT: PM&V proposal Metrics in SW/HW to the service topology

- ○ This step includes the design of how performance will be measured, using the metrics defined in the previous step. The output of the PM&V proposal describes the design of the network architecture building blocks solving functional and non-functional requirements, as negotiated with the Ops. Essential indicators can be Measurement Points (MPs), scalability and robots for testing purposes. This function is an iterative process with the SO and Ops.

3. **ESTABLISH PM&V Architecture** – INPUT: PM&V proposal / OUTPUT: PM&V Architecture, metrics implemented in SW/HW to the service topology

- ○ This step includes engineering and deployment tasks, as well as the establishment of the PM&V proposal. The implementation form of the metrics can be pure SW based (e.g. JavaScripts and analyse procedures on web resources) or on HW probes as described Section 2.1.1 UNINETT activities. This function is as collaboration between SO and Ops and operations.

4. **TEST PM&V metrics** – INPUT: PM&V architecture / OUTPUT: Test sequences on metrics, and validation to the OLAs

- ○ This step includes the scenarios (sequences) for testing and validation of the metrics of the PM&V architecture. Test sequences can be implemented manually or automated, and show the accuracy of the metrics to negotiated OALs. Test cases can be fully automated by robots. In this process step we have the flexibility of metrics revision regarding changes of the service delivery: TEST SUCCESSFUL: TESTs are input to HANDOVER process, which is not a part of this document. FAILD; Identify/verify unaccounted REQUIRED operational performance data, modification of this metrics set-up (e.g. extension of KOPIs), set-up revised metrics into the PM&V architecture, and prepare for further tests.

The output of this testing process facilitates the development of an operative performance measurement and verification service, in our case the “Wireless Performance Monitoring and Verification as a Service” (WPM&VaaS).

Appendix B Measurement Metrics

Many different metrics are available and a configuration script specifies which metrics are actually measured and exported for a given probe. The available metrics are:

Layer 1

- cells_tota - total number of WiFi cells detected
- cells_2ghz_edur - number of WiFi cells advertising the eduroam SSID in the 2.4Ghz band
- cells_5ghz_edur - number of WiFi cells advertising the eduroam SSID in the 5Ghz band
- cells_edur - number of WiFi cells advertising the eduroam SSID
- cells_2ghz - number of WiFi cells in the 2.4GHz band
- cells_5ghz - number of WiFi cells in the 5GHz band
- cells_uniq_ssid - number of unique SSID's seen
- cells_uniq_radi - number of unique AP radio interfaces seen
- cells_uniq_chan - number of unique AP Channels seen
- cells_num_mast - number of APs in Master mode
- cells_2g_freq - number of unique 2.4G channels used
- cells_5g_freq - umber of unique 5G channels used
- ssid_list - list of SSIDs seen
- wifi_freq - frequency
- wifi_apid - access point ID
- wifi_sign - signal level
- wifi_qual - link quality
- wifi_rate - bit rate

Layer 2

- wifi_asso - time it takes to associate with eduroam SSID
- dhcp_time - time it takes to receive IPv4 address
- ipv6 - true if IPv6 is available

Layer 3

All these metrics are available for both IPv4 and IPv6

- dns - time it takes to resolve an IP address.
- http - time it takes to get response from web server.

- owj - jitter
- owpl - packet loss
- bwup - maximum upload speed
- bwdo - maximum download speed
- rtt_min - minimum round trip time
- rtt_avg - average round trip time
- rtt_min - minimum round trip time
- myipv4 - assigned IPv4 address
- myipv6 - assigned IPv6 address

Possible future metrics

- Measurements per frequency band
- Separate monitoring for each 802.11 standard
- Application level measurements, e.g. SIP

Appendix C Terms and Terminology Architecture

Term	Definition
Data Source Collector	The layer responsible for: (i) Generating information regarding the performance of the wireless network, and (ii) exporting the data to the Relational Database.
Relational Database	A database where the raw data from the Data Source layer is collected.
Analytics Engine	The architecture block responsible for examining/analysing of the RDB's raw data.
Raw Data	Referes to the data collected from the data source layer, not subjected to processing or any other manipulation.
Performance	The aggregated throughput or mean round-trip time (RTT) of all clients connecting to a specific AP. An individual client / mobile device throughput or RTT.
Embedded test	JavaScript code embedded in websites that enables users to run performance tests automatically, without interaction.

Appendix D WiFiMon Web-GUI screenshots

This Appendix contains screenshots of the first version of the WiFiMon Web-GUI.

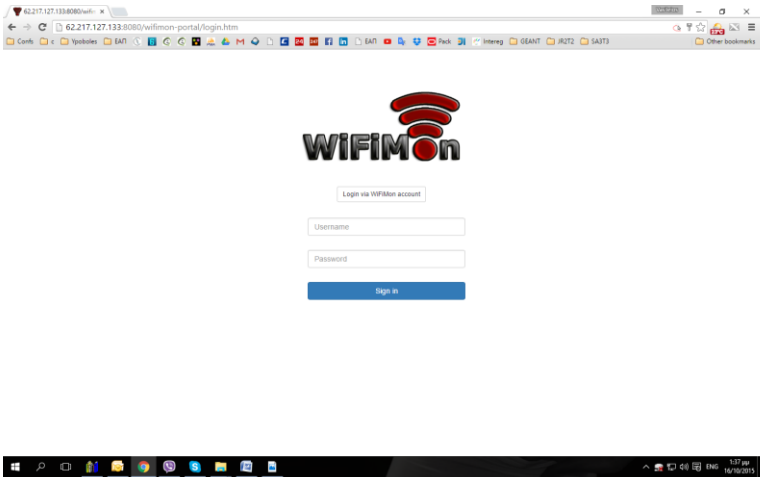
62.217.127.133:8080/wifimon-portal/login.htm
WiFiMon
Login via WiFiMon account
Username
Password
Sign in

Figure D.1: WiFiMon portal – Login page

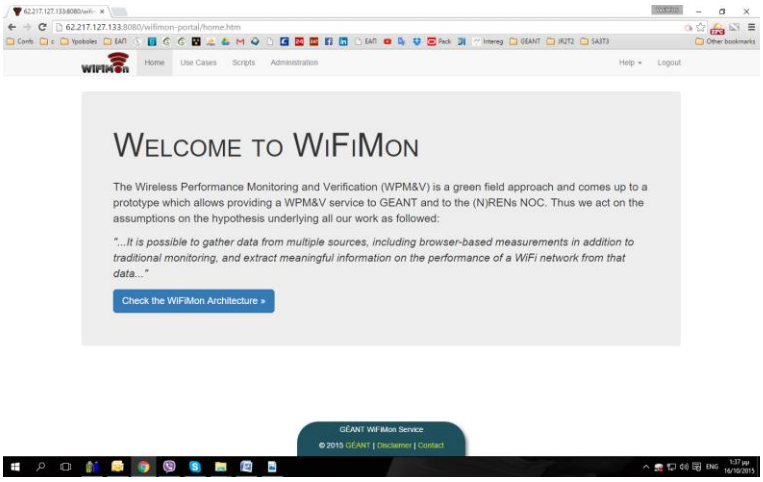

WELCOME TO WIFIMON

The Wireless Performance Monitoring and Verification (WPM&V) is a green field approach and comes up to a prototype which allows providing a WPM&V service to GEANT and to the (N)RENs NOC. Thus we act on the assumptions on the hypothesis underlying all our work as followed:

"...It is possible to gather data from multiple sources, including browser-based measurements in addition to traditional monitoring, and extract meaningful information on the performance of a WiFi network from that data..."

Check the WiFiMon Architecture »

GÉANT WiFiMon Service
© 2015 GÉANT | Disclaimer | Contact

Figure D.2: WiFiMon portal – Admin homepage

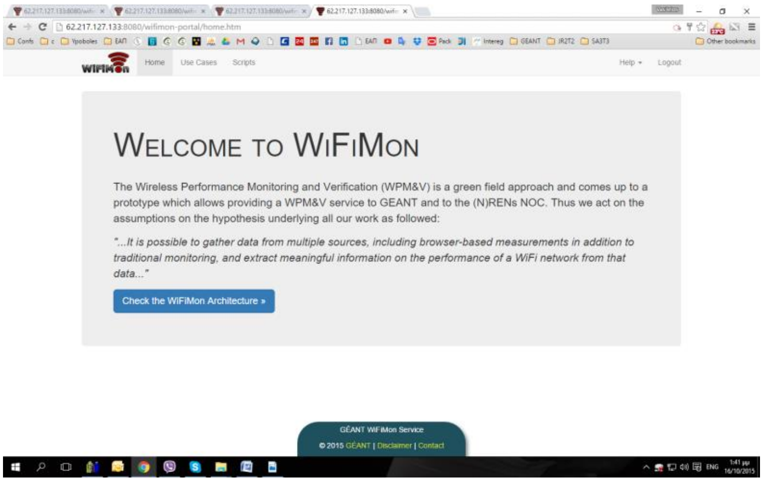

WELCOME TO WIFIMON

The Wireless Performance Monitoring and Verification (WPM&V) is a green field approach and comes up to a prototype which allows providing a WPM&V service to GEANT and to the (N)RENs NOC. Thus we act on the assumptions on the hypothesis underlying all our work as followed:

"...It is possible to gather data from multiple sources, including browser-based measurements in addition to traditional monitoring, and extract meaningful information on the performance of a WiFi network from that data..."

Check the WiFiMon Architecture »
GÉANT WiFiMon Service
© 2015 GÉANT | Disclaimer | Contact

Figure D.3: WiFiMon portal – Simple user homepage

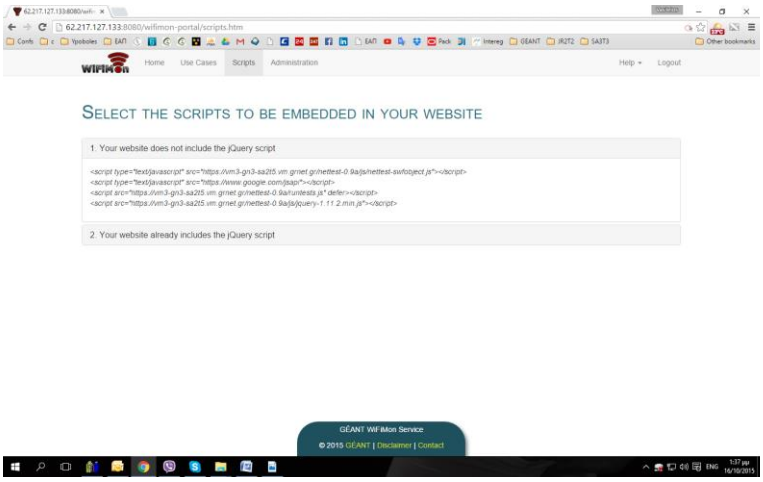

62.217.127.133:8080/wifimon-portal/scripts.htm
Home
Use Cases
Scripts
Administration
Help
Logout

SELECT THE SCRIPTS TO BE EMBEDDED IN YOUR WEBSITE

1. Your website does not include the jQuery script

```
<script type="text/javascript" src="https://vm3-gn3-sa2t5.vm.grnet.gr/nettest-0.9a/js/nettest-swfobject.js"></script>
<script type="text/javascript" src="https://www.google.com/jsapi"></script>
<script src="https://vm3-gn3-sa2t5.vm.grnet.gr/nettest-0.9a/runtests.js" defer></script>
<script src="https://vm3-gn3-sa2t5.vm.grnet.gr/nettest-0.9a/js/jquery-1.11.2.min.js"></script>
```

2. Your website already includes the jQuery script

GÉANT WiFiMon Service
© 2015 GÉANT | Disclaimer | Contact

Figure D.4: WiFiMon portal – Scripts for websites (access: all)

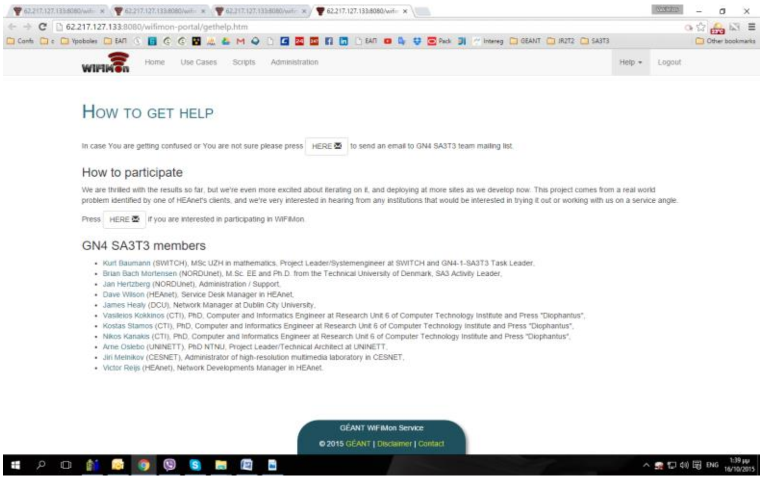

HOW TO GET HELP

In case You are getting confused or You are not sure please press HERE ✉ to send an email to GN4 SA3T3 team mailing list.

How to participate

We are thrilled with the results so far, but we're even more excited about iterating on it, and deploying at more sites as we develop now. This project comes from a real world problem identified by one of HEAnet's clients, and we're very interested in hearing from any institutions that would be interested in trying it out or working with us on a service angle.

Press HERE ✉ if you are interested in participating in WiFiMon.

GN4 SA3T3 members

- Kurt Baumann (SWITCH), MSc UZH in mathematics, Project Leader/Systemengineer at SWITCH and GN4-1-SA3T3 Task Leader,
- Brian Bach Mortensen (NORDUnet), M.Sc. EE and Ph.D. from the Technical University of Denmark, SA3 Activity Leader,
- Jan Hertzberg (NORDUnet), Administration / Support,
- Dave Wilson (HEAnet), Service Desk Manager in HEAnet,
- James Healy (DCU), Network Manager at Dublin City University,
- Vasileios Kokkinos (CTI), PhD, Computer and Informatics Engineer at Research Unit 6 of Computer Technology Institute and Press "Diophantus",
- Kostas Stamos (CTI), PhD, Computer and Informatics Engineer at Research Unit 6 of Computer Technology Institute and Press "Diophantus",
- Nikos Kanakis (CTI), PhD, Computer and Informatics Engineer at Research Unit 6 of Computer Technology Institute and Press "Diophantus",
- Arne Oslebo (UNINETT), PhD NTNU, Project Leader/Technical Architect at UNINETT,
- Jiri Melnikov (CESNET), Administrator of high-resolution multimedia laboratory in CESNET,
- Victor Reijs (HEAnet), Network Developments Manager in HEAnet.

Figure D.5: WiFiMon portal – Get Help page (access: all)

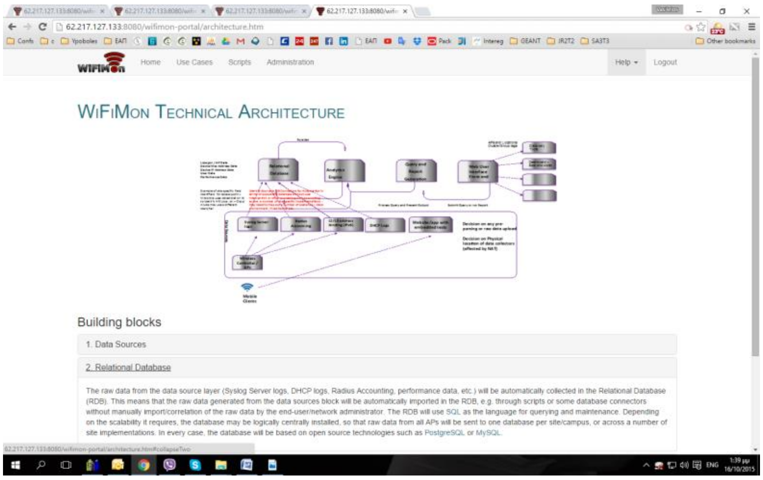

WiFiMon Technical Architecture

Building blocks

1. Data Sources

2. Relational Database

The raw data from the data source layer (Syslog Server logs, DHCP logs, Radius Accounting, performance data, etc.) will be automatically collected in the Relational Database (RDB). This means that the raw data generated from the data sources block will be automatically imported in the RDB, e.g. through scripts or some database connectors without manually import/correlation of the raw data by the end-user/network administrator. The RDB will use SQL as the language for querying and maintenance. Depending on the scalability it requires, the database may be logically centrally installed, so that raw data from all APs will be sent to one database per site/campus, or across a number of site implementations. In every case, the database will be based on open source technologies such as PostgreSQL or MySQL.

Figure D.6: WiFiMon portal – Architecture page (access: all)

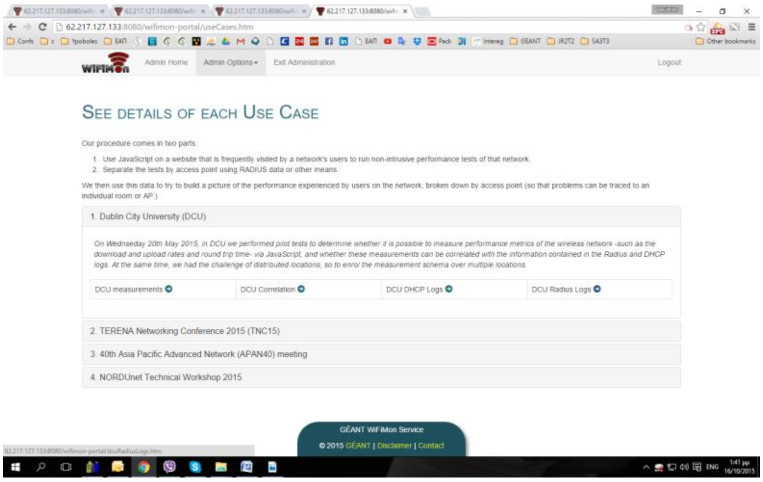
62.217.127.133:8080/wifimon-portal/useCases.htm
Admin Home
Admin Options
Exit Administration
Logout
GÉANT WiFiMon Service
© 2015 GÉANT | Disclaimer | Contact

See details of each Use Case

Our procedure comes in two parts:

1. Use JavaScript on a website that is frequently visited by a network's users to run non-intrusive performance tests of that network.
2. Separate the tests by access point using RADIUS data or other means.

We then use this data to try to build a picture of the performance experienced by users on the network, broken down by access point (so that problems can be traced to an individual room or AP.)

1. Dublin City University (DCU)

On Wednseday 20th May 2015, in DCU we performed pilot tests to determine whether it is possible to measure performance metrics of the wireless network -such as the download and upload rates and round trip time- via JavaScript, and whether these measurements can be correlated with the information contained in the Radius and DHCP logs. At the same time, we had the challenge of distributed locations, so to enrol the measurement schema over multiple locations.

DCU measurements ●	DCU Correlation ●	DCU DHCP Logs ●	DCU Radius Logs ●

2. TERENA Networking Conference 2015 (TNC15)

3. 40th Asia Pacific Advanced Network (APAN40) meeting

4. NORDUnet Technical Workshop 2015

Figure D.7: WiFiMon portal – Use cases main page (access: admin)

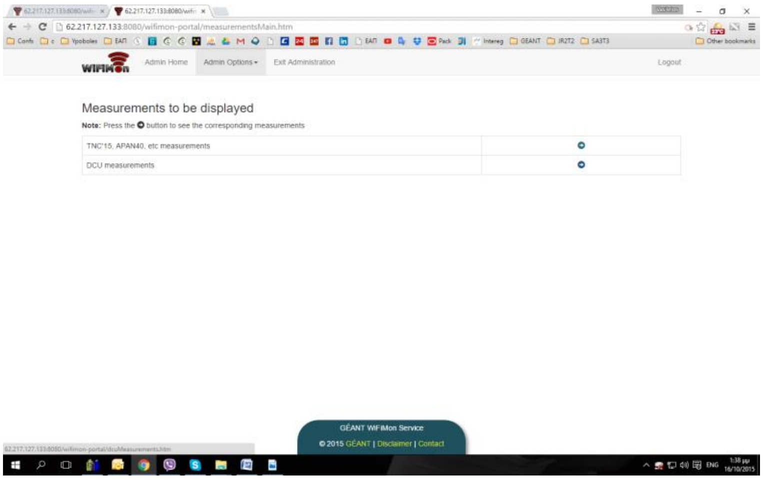
62.217.127.133:8080/wifimon-portal/measurementsMain.htm
Admin Home
Admin Options
Exit Administration
Logout
GÉANT WiFiMon Service
© 2015 GÉANT | Disclaimer | Contact

Measurements to be displayed

Note: Press the ● button to see the corresponding measurements

TNC'15, APAN40, etc measurements	●
DCU measurements	●

Figure D.8: WiFiMon portal – Measurements main page (access: admin)

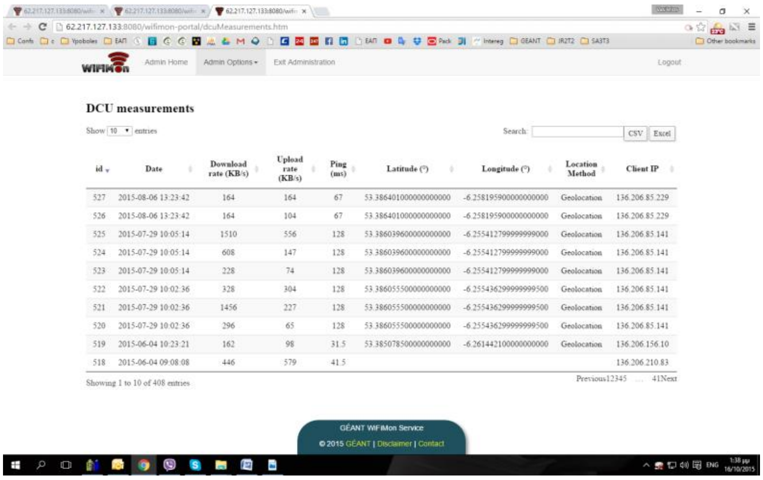

DCU measurements

Show 10 ▾ entries

Search: CSV Excel

id	Date	Download rate (KB/s)	Upload rate (KB/s)	Ping (ms)	Latitude (°)	Longitude (°)	Location Method	Client IP
527	2015-08-06 13:23:42	164	164	67	53.386401000000000000	-6.2581959000000000000	Geolocation	136.206.85.229
526	2015-08-06 13:23:42	164	104	67	53.386401000000000000	-6.2581959000000000000	Geolocation	136.206.85.229
525	2015-07-29 10:05:14	1510	556	128	53.386039600000000000	-6.2554127999999999000	Geolocation	136.206.85.141
524	2015-07-29 10:05:14	608	147	128	53.386039600000000000	-6.2554127999999999000	Geolocation	136.206.85.141
523	2015-07-29 10:05:14	228	74	128	53.386039600000000000	-6.2554127999999999000	Geolocation	136.206.85.141
522	2015-07-29 10:02:36	328	304	128	53.386055500000000000	-6.2554362999999999500	Geolocation	136.206.85.141
521	2015-07-29 10:02:36	1456	227	128	53.386055500000000000	-6.2554362999999999500	Geolocation	136.206.85.141
520	2015-07-29 10:02:36	296	65	128	53.386055500000000000	-6.2554362999999999500	Geolocation	136.206.85.141
519	2015-06-04 10:23:21	162	98	31.5	53.385078500000000000	-6.2614421000000000000	Geolocation	136.206.156.10
518	2015-06-04 09:08:08	446	579	41.5				136.206.210.83

Showing 1 to 10 of 408 entries

Previous12345 … 41Next

GÉANT WiFiMon Service

© 2015 GÉANT | Disclaimer | Contact

Figure D.9: WiFiMon portal – DCU measurements (access: admin)

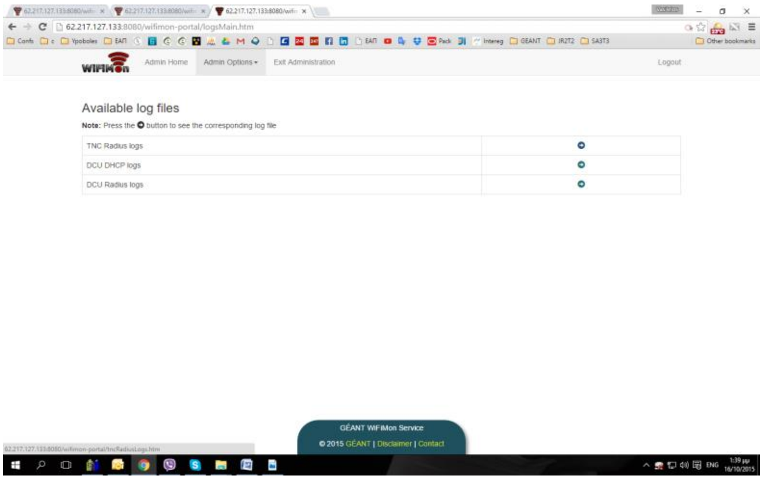

Available log files

Note: Press the ● button to see the corresponding log file

TNC Radius logs	●
DCU DHCP logs	●
DCU Radius logs	●

GÉANT WiFiMon Service

© 2015 GÉANT | Disclaimer | Contact

Figure D.10: WiFiMon portal – Logs main page (access: admin)

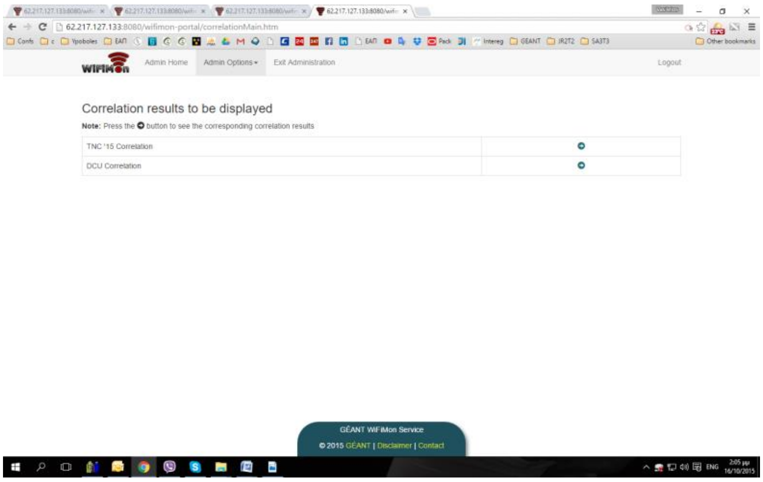

Correlation results to be displayed

Note: Press the ● button to see the corresponding correlation results

TNC '15 Correlation	●
DCU Correlation	●

GÉANT WiFiMon Service

© 2015 GÉANT | Disclaimer | Contact

Figure D.11: WiFiMon portal – Correlation main page (access: admin)

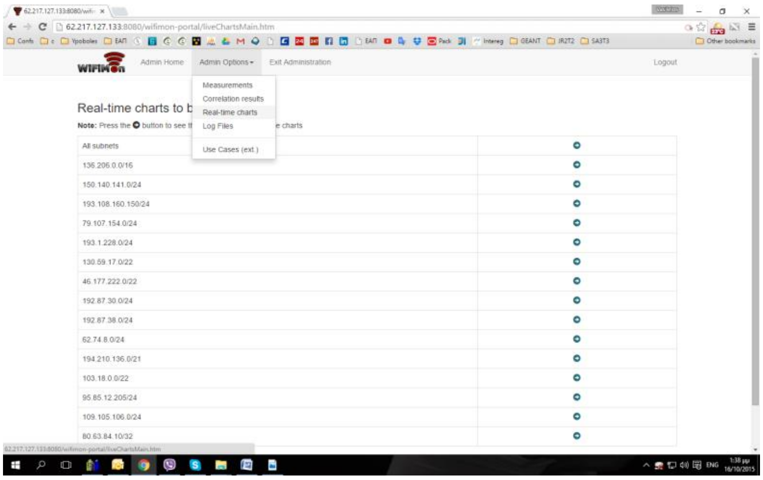

Real-time charts to b

Note: Press the ● button to see th e charts

All subnets	●
136.206.0.0/16	●
150.140.141.0/24	●
193.108.160.150/24	●
79.107.154.0/24	●
193.1.228.0/24	●
130.59.17.0/22	●
46.177.222.0/22	●
192.87.30.0/24	●
192.87.38.0/24	●
62.74.8.0/24	●
194.210.136.0/21	●
103.18.0.0/22	●
95.85.12.205/24	●
109.105.106.0/24	●
80.63.84.10/32	●

Figure D.12: WiFiMon portal – Real-time charts main page (access: admin)

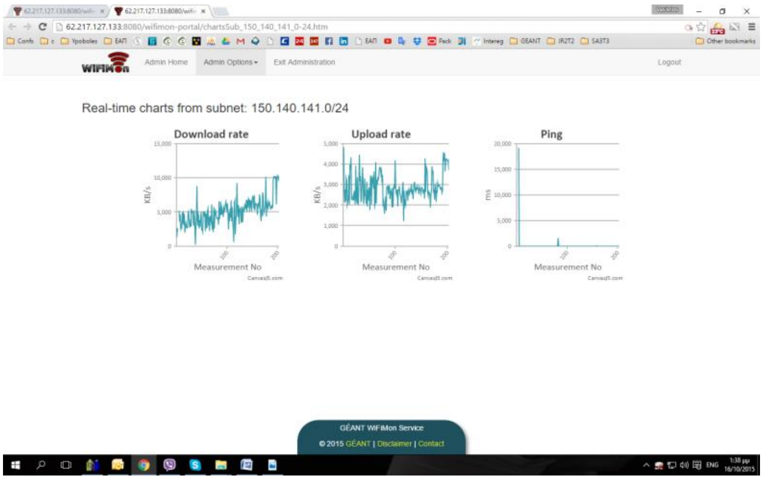

Real-time charts from subnet: 150.140.141.0/24

GÉANT WiFiMon Service

© 2015 GÉANT | Disclaimer | Contact

Figure D.13: WiFiMon portal – Real-time charts from subnet 150.141.121.0/24 (access: admin)

Appendix E HEAnet Conference: Comments on the Best Sessions

"In answer to the Question, which Parallel A session did you derive the most benefit from, and why?"

Answers:

- I really enjoyed the WiFi with crowdsourced monitoring talk by Dave Wilson and James Healy. I think it is really interesting what they are doing and i thought the delivered the presentation in an excellent way;
- How to Fix WiFi with crowdsourced Monitoring - very interesting topic and it was delivered excellently;
- Dave and James talking about the WiFi crowdsourcing was excellent. I just found the idea to be a unique take on an old problem. Well done to the two guys;
- How to Fix Wifi with crowdsourcing was both fascinating and enjoyable, WiFi issues are regularly occurring problems and this gave insight into how other Institutes are tackling the problem;
- How to fix WiFi in the cloud. I thought James and Dave made and excellent team. They spoke so well together and made their presentation fun and interesting. They also explained there project very clearly and the results were great. I really enjoyed this talk. With the single negative comment:
- How to fix Wifi - an interesting topic and an on-going challenge, not sure of the proposed fix though (we'll look at it internally).

The session was recorded and may be found at [HEAnetCONF].

References

[AIROS] https://www.ubnt.com/airmax/airos7/

[ANALYSE] Analyse TNC2015: https://tnc15.terena.org/core/presentation/206

[APAN40] Kuala Lumpur, Malaysia: http://www.apan.net/meetings/KualaLumpur2015/

[BOOMERANG] http://www.lognormal.com/boomerang/doc/; https://github.com/yahoo/boomerang

[CAMPUSBP] “Campus network monitoring and Security” organised by the GEANT plus Project together with Campus-BP in April 2014 - https://www.cesnet.cz/cesnet/events/campus-network-ws/?lang=en

[CESNET] CESNET: https://www.cesnet.cz/?lang=en

[CONFERENCE4ME] https://tnc15.terena.org/web/media/news/id/3893

[CONNECT] Wireless Drowedsourced Performance Monitoring and Verification, A First Trial: https://intranet.geant.org/gn4/1/Activities/SA3/T3/Documents/SA3T3_WirelessMon/GoingPublic/CONNECT_Article_WiFiMON_final_20150909.docx?Web=1

[DATA] https://vm3-gn3-sa2t5.vm.grnet.gr/nettest-0.9a/dcu_results.php

[EDUPERT] eduPERT monthly call of July 2015: http://services.geant.net/edupert/Resources/Documents/edupert-wifimon-introduction.pdf

[EDUROAM] www.eduroam.org

[EDUROAMSP] eduroam service policy: https://www.eduroam.org/downloads/docs/GN3-12-192_eduroam-policy-service-definition_ver28_26072012.pdf

[FILES] https://vm3-gn3-sa2t5.vm.grnet.gr/nettest-0.9a/dcu_radius_logs.php / https://vm3-gn3-sa2t5.vm.grnet.gr/nettest-0.9a/dcu_dhcp_logs.php (Note: Currently there is no automated procedure that populates the useful information from the DHCP and/or Radius logs to the Relational Database. The manual procedure described above will be replaced by an automated one, which will enable the population of the log files information to the Relational Database without user’s intervention.)

[FLASH] Flash is a multimedia and software platform used for creating vector graphics, animation, browser games, rich Internet applications, desktop applications, mobile applications and mobile games - https://en.wikipedia.org/wiki/Adobe_Flash

[FLUKEAIRCHECK] AirCheck Wi-Fi Tester - Fluke Networks." 2011. 14 Aug. 2015 http://www.flukenetworks.com/enterprise-network/network-testing/AirCheck-Wi-Fi-Tester

[HEAnetCONF] http://www.heanet.ie/conferences/2015/talks/id/165

[IETFLMAP] IETF LMAP charter: https://datatracker.ietf.org/wg/lmap/charter

[INITIALCHARTS] https://vm3-gn3-sa2t5.vm.grnet.gr/nettest-0.9a/dcu_correlation_charts.php

[IPERF] https://iperf.fr/

[IPPM] https://tools.ietf.org/wg/ippm/

[**JAVASCRIPT**] JavaScript is the programming language of HTML and the Web - http://www.w3schools.com/js/

[**LMAP**] Private Draft:http:/datatracker.ietf.org/doc/draft-deng-lmap-collaboration/%20

[**MOBILEWORLD**] http://www.slideshare.net/bge20/2013-11-mobile-eating-the-world

[**MOBILEDATA**] http://www.asymco.com/2014/07/08/late-late-majority/

[MySQL] https://www.mysql.com/

[NETTEST] https://code.google.com/p/nettest/

[NAGIOS] https://www.nagios.org/

[NORDUNETTW] https://www.nordu.net/content/nordunet-technical-workshop

[OOKLA] http://www.ookla.com/

[PERFORMANCE] https://vm3-gn3-sa2t5.vm.grnet.gr/nettest-0.9a/results.php

[PRIME] http://www.cisco.com/c/en/us/products/cloud-systems-management/prime-infrastructure/index.html

[POSTGRESQL] http://www.postgresql.org/

[RASPBERRYPI] https://www.raspberrypi.org/

[RASPBERRYPIB+] Raspberry Pi 1 Model B+." 2015. 19 Aug. 2015 https://www.raspberrypi.org/products/model-b-plus/

[RASPBIAN] https://www.raspbian.org/

[REFERENCE] http://whatis.techtarget.com/definition/reference-data

[RESULTS] https://vm3-gn3-sa2t5.vm.grnet.gr/nettest-0.9a/dcu_correlation.php

[RFC7398] https://datatracker.ietf.org/doc/rfc7398/

[RFC7536] https://tools.ietf.org/html/rfc7536

[RTT] http://searchnetworking.techtarget.com/definition/round-trip-time

[SPEEDTEST] http://speedtest.cesnet.cz

[SPRINGXD] http://docs.spring.io/spring-xd/docs/current/reference/pdf/spring-xd-reference.pdf

[SQL] https://en.wikipedia.org/wiki/SQL

[TESTBED] The URL for the testbed UI is https://10.255.0.25:8443/login. However, since access is restricted to team members only, interested readers may log in to an open-access WiFiMon installation: https://62.217.125.88:8443/login (Username: admin; password: gn4). Note that this installation does not include the eduroam configuration.

[THROUGHPUT] http://searchnetworking.techtarget.com/definition/throughput

[UNINETT] https://www.uninett.no/

[WEBGUI] http://62.217.127.133:8080/wifimon-portal/login.htm

[WIRELESSPERF] http://searchnetworking.techtarget.com/tip/Three-ways-to-measure-wireless-network-performance

Glossary

ACS	Access Control Server
AE	Analytic Engine
AMeN	Active Monitoring eduroam Node
AP	Access Point
CAPWAP	Control and Provisioning of Wireless Access Points
DCU	Dublin City University
GN3plus	(GÉANT Network 3 plus), a project part-funded from the EC's Seventh Framework Programme under Grant Agreement No.605243
GN4-1	(GÉANT Network 4, Phase 1), a project part-funded from the EC's Horizon 2020 research and innovation programme under Grant Agreement No.691567
HW	Hardware
IdP	Identity Provider
IETF	Internet Engineering Task Force
IOS	(originally iPhone) Operating System
IPPM	Internet Protocol Performance Metrics
ISP	Internet Service Provider
KOPI	Key Operational Performance Indicators
LMAP	Large-Scale Measurement of Broadband Performance
MA	Measurement Agent
MAC	Media Access Control
MP	Measurement Points
ms	millisecond
NAT	Network Address Translation
NCC	Network Coordination Centre
NREN	National Research and Education Network
NRO	National Research Organisations
PM&V	Performance Measurement and Verification
PoC	Proof of Concept
QaRG	Query and Report Generator
OLA	Operational Level Agreement
Ops	Operations
OS	Operating System
RADIUS	Remote Authentication Dial-In User Service
RDB	Relational Database
REST	Representational State Transfer
RFC	Request for Comments
RTT	Round-Trip Time

SO	Service Owners
SP	Service Provider
SW	Software
SSID	Service Set Identifier
UM	University of Malaya, Kuala Lumpur, Malaysia
WAP	Wireless Access Point
Web-UI	Web-User interface
WLAN	Wireless Local Area Network
WPM&VaaS	Wireless PM&V as a Service