Journal of Physics: Conference Series

PAPER • OPEN ACCESS

Wi-Fi Service enhancement at CERN

To cite this article: V Ducret *et al* 2017 *J. Phys.: Conf. Ser.* **898** 082010

View the article online for updates and enhancements.

Related content

- Scale out databases for CERN use cases
  Zbigniew Baranowski, Maciej Grzybek, Luca Canali et al.
- Accelerating physics teaching at CERN
- CERN and US renew ties

This content was downloaded from IP address 188.184.3.52 on 18/12/2017 at 13:12

IOP Conf. Series: Journal of Physics: Conf. Series **898** (2017) 082010 doi:10.1088/1742-6596/898/8/082010

Wi-Fi Service enhancement at CERN

V Ducret, A Sosnowski, B Gonzalez Caballero and Q Barrand

IT Department, CERN, Route de Meyrin 385, 1217 Meyrin, Switzerland
vincent.ducret@cern.ch, adam.sosnowski@cern.ch, benito.gonzalez.caballero@cern.ch and quentin.barrand@cern.ch

Abstract. Since the early 2000's, the number of mobile devices connected to CERN's internal network has increased from just a handful to well over 10,000. Wireless access is no longer simply "nice to have" or just for conference and meeting rooms; support for mobility is expected by most, if not all, of the CERN community. In this context, a full renewal of the CERN Wi-Fi network has been launched to deliver a state-of-the-art campus-wide Wi-Fi Infrastructure. We aim to deliver, in more than 200 office buildings with a surface area of over 400,000m^2 and including many high-priority and high-occupation zones, an end-user experience comparable, for most applications, to a wired connection and with seamless mobility support. We describe here the studies and tests performed at CERN to ensure the solution we are deploying can meet these goals as well as delivering a single, simple, flexible and open management platform.

1. Introduction

Wi-Fi deployment at CERN started in the early 2000's, initially to provide basic, easy access to the network in meeting rooms and in some public locations. Given the technology available at the time, the Wi-Fi infrastructure consisted of independent, individually managed access points. This architecture has remained unchanged as the infrastructure has grown in response to the increased demand—demand, above all, for in-office Wi-Fi connectivity—and we are no longer able to meet end-user expectations, notably in terms of mobility and throughput. This document details the different aspects we considered as we planned a modernisation of the Wi-Fi architecture to meet those expectations, notably:

- how to integrate smoothly with the overall campus network,
- how to ensure adequate Wi-Fi coverage,
- how to provide seamless mobility,
- key technology choices, and
- how to deliver an effective management and monitoring infrastructure.

2. An overview of the CERN Campus Network topology

All components of the Wi-Fi infrastructure, in particular the access points, are connected to the CERN campus network. It is important, therefore, to understand the topology of the campus network in order to understand the constraints on, and our choices for, a campus-wide Wi-Fi service.

Content from this work may be used under the terms of the Creative Commons Attribution 3.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.
Published under licence by IOP Publishing Ltd

2.1. Layer 2 and Layer 3 topology

The CERN campus LAN is organised in a multilayer star topology, as shown in Figure 1.

- Two Backbone Routers (full Layer 3) act as the core of the campus network, routing traffic to services in the data centre and the external world;
- two levels of distribution routers (full Layer 3) with "Main Starpoint Routers" aggregating connections to "Starpoint Routers" located in the campus buildings, which support
- Layer 2 Access Switches that provide connectivity for end user devices and other devices, notably Wi-Fi access points.

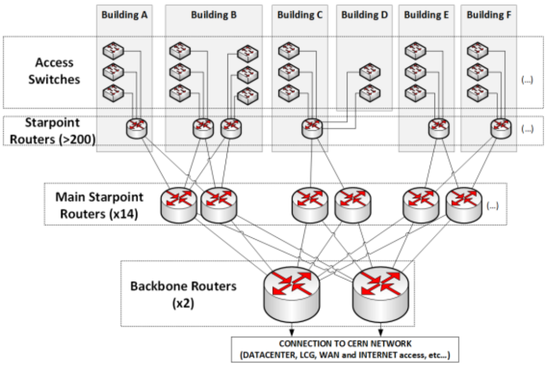

Figure 1. CERN Campus network topology.

2.2. IP address organisation

The access switches have a very basic configuration and all the (dual stack) IP networks are defined in the starpoint routers with, in general, an IPv6 subnet and an IPv4 subnet (/23 to /26, depending on the location) dedicated to a specific access switch. Although this design is highly reliable and scalable, meeting CERN's requirements for simple operational procedures and automated configuration management, a drawback of such a design is that two devices connected to two different access switches within the same building will use different IP networks. While not a problem for wired connections, this leads to limitations for wireless devices, as explained in section 3.4.

3. The Former Wi-Fi infrastructure

3.1. Access points

As mentioned above, Wi-Fi access points (APs) were first deployed in the early 2000's to provide easy network connectivity in meeting rooms, auditoria and a few public locations. Given the design goals, APs were installed in meeting rooms or in corridors and, at that time, there was no option other than using autonomous APs with traffic switched locally. Although the infrastructure has expanded to cover offices in some buildings, the corridor-based AP deployment policy has continued, partly due to concerns from some as to the potential health impacts of in-office AP deployment. Consequently, given the attenuation of signals by the intervening walls, it is almost impossible to deliver, inside an office, the signal level required for clients to be able to negotiate the best modulation and coding. The delivered throughput where office coverage is available is thus usually well below the theoretically achievable rates.

3.2. Service provided

Partly because Wi-Fi was seen simply as an extension of the campus network, but also given the technology limitations of the early 2000's, the Wi-Fi service simply offers connections to the CERN campus network. In particular, there is no "Guest" Wi-Fi service with a lightweight registration procedure offering internet-only access, or restricted access to few on-site networks and services. This is because local traffic switching prevents us from separating guest user traffic from that for fully authorised users.

3.3. Configuration management and monitoring

Autonomous APs have to be configured individually, but with more than 1,200 of them deployed, some automation tools were required. To ease configuration, deployment and operation, scripts had to be developed internally. These rely on PERL and are highly customised, using SNMP to gather monitoring information, and SSH/CLI or SNMP to configure the APs. A Java-based centralised monitoring tool [1] has also been developed, with a GUI presenting information gathered from all the APs (via SNMP) about connected clients, and RF monitoring values (signal quality, retransmissions and so on).

3.4. Limits of the former solution

3.4.1. Roaming issues

As traffic from the autonomous APs is bridged locally by the access switches, Wi-Fi devices obtain an IP address in exactly the same manner as a wired device. Due to the campus network IP organisation, this implies that, when roaming from one AP to another that is connected to a different switch, a Wi-Fi device needs to change its IP address and all network connections are lost or need to be refreshed. Although APs in a meeting room will all be connected to a single switch, this will not be the case if the APs are in different buildings nor even, in some cases, inside a single large building. Roaming is thus essentially impossible for users as they cannot know when a short walk, for example between buildings interconnected via corridors, will lead to their Wi-Fi device changing its IP address.

3.4.2. IPv4 address exhaustion

Another consequence of having Wi-Fi devices obtain an IPv4 address from the subnetwork dedicated to the switch to which the AP is connected is that, for *each* location, we need to allocate sufficient IP addresses to cater for the *maximum* foreseeable number of Wi-Fi clients. For example, we generally reserve a /23 IPv4 network for an auditorium or a large meeting room in order to cater for the maximum capacity of the room—a capacity that has to be measured, moreover, in devices, not simply in seats. The 500 or so IP addresses thus reserved cannot be used elsewhere on the campus, even though, for most of the time, there will be fewer than 500 active devices in the auditorium. On the other hand, on the relatively few occasions when there are more people and devices than expected, the consequences are even worse: some users will not be able to connect to the network because all available IP addresses have been allocated. To avoid an overly-inefficient use of IPv4 addresses, we reduce DHCP lease lifetime to ensure devices simply passing through, say, the main cafeteria, do not block an IPv4 address for many hours. Unfortunately, this mitigation measure aggravates the impression of poor service quality in cases of high occupation as a device is not simply either connected or not; it may gain and lose connectivity depending on lease lifetimes and the behaviour of competing devices nearby.

3.4.3. RF planning

For proper coverage, it is essential to avoid RF interference between APs by adjusting the RF channel and transmit power as appropriate. For autonomous APs, this RF planning is manual, highly time consuming, complex to define correctly in some buildings, and cannot be adapted dynamically if the RF environment changes.

3.4.4. Guest service

As detailed in sections 2.2 and 3.2, there is no way to separate Guest user traffic from that of CERN users, so no Guest service (internet only access) can be provided.

3.4.5. Scalability

As mentioned in section 3.3, some tools have been developed internally to automate the configuration and monitoring of the APs. These, though, because they rely on SSH/CLI and SNMP, have to be modified to handle new devices as well as any change to the CLIs. Maintenance of the configuration management and monitoring tools is resource intensive and significant development would be needed to enable these to cope with the number of APs required for campus-wide coverage, let alone deliver the automated reconfiguration in response to monitoring data that would be required, for example, to increase the signal strength of neighbours of a failing AP in order to provide continuity of coverage.

4. Wi-Fi Service enhancement project goals

4.1. Full campus coverage for Office areas

The main goal of the project is to ensure proper Wi-Fi coverage in offices as well as meeting rooms and public areas. This requires AP deployment in some 200 buildings with a surface area of over 400,000m^2. This office-wide Wi-Fi coverage is not seen simply as a convenience for users but also as the long-term solution for office network access, avoiding the need for large-scale replacement of the structured cabling. A key design aim, therefore, is to ensure coverage and signal quality to support connections that can replace a wired connection for all of the standard applications used in an office environment.

4.2. Global roaming capabilities

The new solution must enable end-users to move around the campus without fear of disconnection. This is particularly important as we foresee an increased use of telephony over Wi-Fi.

4.3. Guest access support

Support for a restricted-access Guest network is a requirement for two reasons. Firstly, to ensure that guests have access only to devices and services that are publically available; malicious visitors to CERN today have unrestricted access to all devices. Secondly, to speed the registration process; if guests have only restricted access, we can offer a lightweight registration mechanism that records only a visitor’s mobile phone number (for traceability requirements), avoiding delays for human approval.

An even more restricted visitor network, allowing only access to public web servers at CERN but not any external access, is required for users that cannot receive an SMS to validate their mobile phone number. This is a request from the CERN Visits Service who wish to provide pointers to on-line in-depth information about exhibits. Many of the over 100,000 visitors to CERN each year do not have roaming access to mobile telephony services; a “walled garden” Wi-Fi service would enable such visitors to access the additional information.

4.4. Configuration and operation automation

The new Wi-Fi infrastructure must be able to integrate with existing tools, provide automated configuration management, and ease installation and operation procedures by providing a central point of configuration and monitoring.

5. Studies, market survey and technical evaluation

5.1. Wi-Fi coverage

A signal level of at least -65dBm is required to meet the coverage and service quality goals set out in section 4.1. As shown in Figure 2, simulations suggest this requires APs to be installed in offices. This conclusion was confirmed by on-site surveys and we plan to deploy one AP per three offices on average.

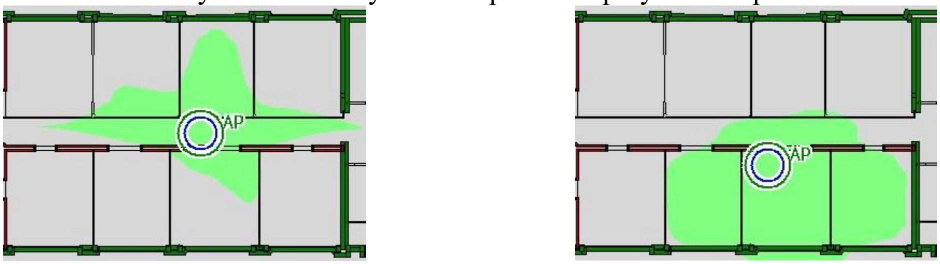

Figure 2. Simulated -65dBm coverage for an AP in a corridor (left) and an office (right).

The target signal level enables clients to use at least the 802.11ac Modulation and Coding Scheme index 6 (MCS6). With 40MHz channels, this corresponds to expected TCP throughputs of:

- ~65 Mbps for most current smartphones/tablets (1 spatial stream; 135 Mbps PHY data-rate);
- ~130 Mbps for most current laptops (2 spatial streams; 270 Mbps PHY data-rate);
- ~200 Mbps for the latest laptop models (3 spatial streams; 405 Mbps PHY data-rate).

These are the minimum targets: clients will be able to negotiate better data rates whenever possible, leading to better TCP throughputs perceived by the end-users. The effectiveness of this negotiation depends on the signal-to-noise ratio, equal to the power (in dBm) of the signal received minus the RF noise (also in dBm) at the client level. MCS6 with 802.11ac and 40MHz channels requires a ratio of between 23 and 27 [2]. As the RF noise observed in our offices is around -90dBm, the target -65dBm signal leads to a SNR of 25, which is within the required range.

5.2. Market Survey

During 2015, a Market Survey was undertaken to improve our understanding of the different technology implementations being offered. We also met IT teams from universities and sites with similar challenges (EPFL, ETHZ, Geneva and Heathrow airports). This study led to three main conclusions.

1. Controller-based architectures are now sufficiently scalable and redundant to match our requirements. This was a great relief as previous studies had concluded that controller-based technologies, although proposed for some years, could not meet our bandwidth needs in a scalable manner. As they deliver a central point of configuration and monitoring, dynamic optimisation of the RF environment, support roaming and enable the isolation of guest user traffic, such an architecture offers an off-the-shelf way of meeting our main design goals. Unfortunately, as the controller/AP protocols are proprietary, such solutions imply a single vendor infrastructure. To mitigate the disadvantages of this vendor lock-in, CERN intends to deploy APs across the site in a relatively short period.
2. APs supporting "Wave 2" of the 802.11ac standard, and so MU-MIMO (Multi-User Multiple Input Multiple Output) capability to send traffic to several "Wave2" clients at the same time, would be available on the required timescale. Although MU-MIMO functionality requires compatible clients, a relative rarity today, we plan to keep the APs for 5 to 8 years so being ready to support such clients when then become widespread without further installation effort is an advantage.
3. Multi-Gigabit Ethernet is not sufficiently mature. Although APs with multi-Gigabit Ethernet ports are available, supporting 2.5 Gbps or 5Gbps on a Cat5e cable, this technology was not fully standardised and switches with compatible ports are still expensive and offer only low port density. We believe, though, that aggregate client load on a single AP is unlikely to require an upstream connection of more than 1 Gbps in the immediate future, and so decided not to opt for multi-Gigabit capable devices for the moment.

5.3. Technical Evaluation

Following the Market Survey, a tender was issued setting out our detailed requirements and, as part of the evaluation process, controller-based configurations from two technically compliant vendors (Cisco and HPE-Aruba) were tested extensively during a two month period. A full setup of each system was installed in a dedicated building enabling detailed validation of their capabilities in terms of

- redundancy and scalability;
- roaming capabilities;
- overall performance (throughput of the clients, both for IPv4 and IPv6);
- centralised configuration management capabilities (via SNMP and APIs);
- ease of integration with our network;
- guest access support; and
- advanced monitoring and debugging capabilities.

In addition to evaluating these features with real clients (including smartphones, laptops and tablets) we used a Wi-Fi traffic generator to test behaviour with multiple clients. This Wi-Fi traffic generator was able, amongst other things, to simulate many tens of clients, perform benchmarks tests for both IPv4 and IPv6 performance, and test roaming delays between two APs. The graphs in Figures 3 and 4 give examples of the results obtained with this test tool.

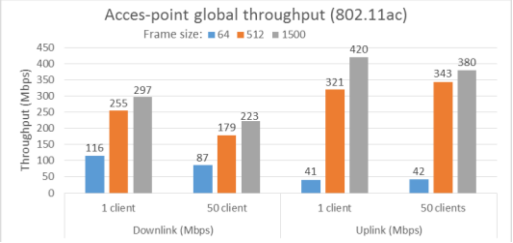

Figure 3. Global AP throughput.

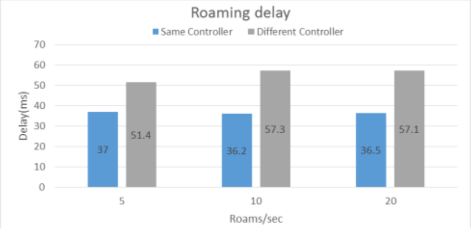

Figure 4. Roaming delay.

6. Pre-deployment work

6.1. RF simulations and site surveys

To establish where to install APs to ensure optimum service quality, both RF simulations and on-site RF surveys were performed for all buildings concerned by the project.

- An RF simulation tool was used to build a 3D model of each building, and, taking into account the RF properties of the various different building elements (such as doors, concrete or plasterboard walls and so on), test different AP configurations in order to define the optimum arrangement to deliver the widest possible coverage at our target -65dBm signal level.
- Unfortunately, these simulations cannot be fully trusted as we never have completely accurate information about the building construction and RF propagation is also affected by furniture and other building contents. We therefore deployed temporary APs (on tripods) at the planned locations in each building and used a laptop with specialist RF survey software to check whether or not the coverage matched that predicted by the simulation.

The combination of simulations and site surveys enabled us to establish the optimum AP placement for each building efficiently and with a minimum of disruption to building occupants. Relying on the simulation alone would have led to sub-optimal placements in some cases whilst avoiding the simulation step would have required multiple site surveys in some buildings to obtain the optimal coverage.

6.2. Cabling

As the current structured cabling for many buildings is more than 20 years old, and as no outlets are available where we plan to install the access-points (inside the office, near the ceiling), we have deployed a dedicated Cat6$_A$ infrastructure. Although we plan to install only one AP per three offices, one dedicated outlet has been installed in each office both to allow flexibility for AP installation and as future technologies—such as 60GHz Wi-Fi or Li-Fi—may require a denser AP deployment. Having a dedicated cabling infrastructure for the access-points will also allow us, if necessary, to upgrade the old structured cabling while providing continued network connectivity for building occupants via Wi-Fi.

7. New Wi-Fi infrastructure design

7.1. Global design

The new Wi-Fi infrastructure, based on HPE-Aruba products, relies on fully redundant, centralised controllers (see Figure 5). Two “mobility master” controllers in an active/backup configuration handle configuration centralisation, license management and radio-frequency planning and optimisation. Wi-Fi client traffic is managed by a central cluster of so-called “local controllers”. All controllers in this cluster are active, with traffic load dynamically balanced across all members of the cluster.

7.2. Roaming and IP address management

All APs build a GRE tunnel (RFC 2784) to the controllers (see Figure 6). User traffic is therefore centralised at the controller level and, from a network point of view, enters where the controllers are connected. All Wi-Fi users share the same IP network, and the IPv4/IPv6 address obtained by a wireless device is independent of its location in the campus.

With this setup:

- clients do not need to change IP address as they move about in the Wi-Fi coverage area, and
- we need only a single global pool of IP addresses, not one pool per building.

This setup thus meets the project goals of delivering seamless roaming for users, better usage of the IP address plan and no risk of IPv4 DHCP pool exhaustion in busy areas.

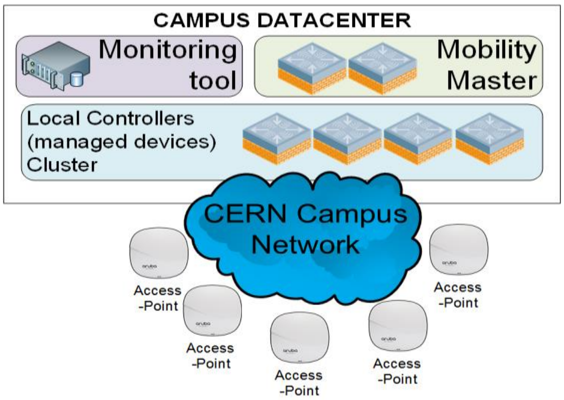

Figure 5. The new Wi-Fi infrastructure.

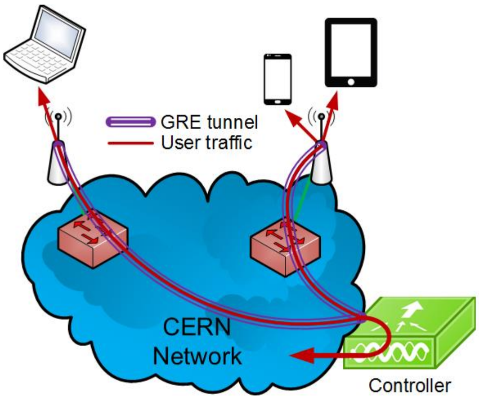

Figure 6. Wi-Fi traffic centralisation.

The drawback of such a setup is that, with so many devices in a single Layer 2 domain, there is an increased risk of a broadcast or multicast storm. However, this risk is mitigated by broadcast/multicast filtering capabilities of the controllers such as the ability to convert broadcast ARP to unicast, to suppress unknown ARP requests to wireless clients, and to optimise IPv6 multicast control messages [3].

In addition to a “production zone” which, with two Mobility Masters and four local controllers, manages Wi-Fi services for almost all of the campus, two additional “zones” have been established, each with their own pair of Mobility Masters and two local controllers:

- an “IT-Pilot” zone, supporting 135 APs deployed across three buildings used by the IT department as well as the nearby restaurant, and
- a “Lab” zone used for testing and debugging purpose;

This overall setup thus allows us to test new features and firmware progressively—first in a controlled environment and then in a limited area but with real users—before a global deployment on the campus.

7.3. Dynamic RF plan calculation

APs send monitoring information to the controllers each 5 minutes. This includes RF signal information which is used to recalculate the best RF settings (channel, channel width, and transmit power) for each AP on a daily basis. Only global settings, such as channel width, the list of allowed channels and the minimum and maximum transmit power, need to be defined. Experience has shown that the RF plan settles down to a stable configuration within two days.

7.4. RF environment optimisation

Once the RF plan has been established, individual APs automatically change their settings in response to interference. Also, as the controllers have a full overview of the RF environment and client/AP association they can, using a combination of IEEE standard features (802.11r, 802.11k and 802.11v) and proprietary algorithms (Aruba Airmatch and ClientMatch [4])

- prevent “sticky client” behaviour by forcing a client to disconnect from a distant AP and to connect to a nearer AP providing a better signal strength;
- distribute clients optimally across the 2.4GHz and 5GHz bands depending on load, available capacity, and channel quality (so-called “band-steering”);
- load-balance clients in a high density environment (e.g. a large auditorium) by forcing clients to move from one AP to another to ensure traffic is distributed equally across all of the APs; and
- optimise the 801.ac Wave2 client distribution by keeping Wave2-capable and –incapable clients on separate APs. This is because, for an AP to operation in MU-MIMO mode, all connected clients must be Wave2 capable. Ensuring MU-MIMO can be used where possible leads to increased system capacity, especially in high density environments.

7.5. Configuration automation

Deployment of new APs is a simple two-step process: first register the details (name, model, MAC address, location…) of the AP in our network database, then simply connect the AP to the network.
Every 15 minutes, an internally developed Python script checks the information in our network database and configures the controllers via their API (the provision of an adequate API was a requirement of our call for tender). With the controllers thus configured, an AP is automatically assigned the appropriate configuration when it appears on the network.

7.6. Advanced monitoring

A proprietary monitoring tool (Aruba Airwave [5]) provides a GUI with access to advanced information about the overall Wi-Fi environment, including status details for each AP and information about clients such as the signal quality and roaming history. Whilst this tool is useful for advanced debugging, integration with our Spectrum-based monitoring infrastructure is required for day-to-day operation. As for the automated configuration, our tender required an API enabling this integration and relevant events send SNMP traps alerting our operators to any problems and generating trouble tickets that ensure incidents can be followed via our existing internal process.

7.7. New service support

With traffic centralised at the controller level, we are able to route traffic differently according to whether the client belongs to an authorised user or to a visitor. Our registration portal has been updated to send a registration password via SMS to users connecting to a dedicated guest SSID and “walled garden” SSIDs will be established to support the needs of our Visitor Service.
In addition to these required new services, localisation (via Wi-Fi or via a Bluetooth Low Energy beacon embedded in the APs) and device tracking services are also supported by the new infrastructure. Evaluation of these features, however, will only take place once the full campus-wide Wi-Fi coverage has been established.

8. Conclusion

To date, 135 APs have been deployed in three office buildings and a nearby restaurant and are managed by a pair of controllers configured for redundancy. This pilot service has demonstrated that we can deploy and manage the access points easily, that we achieve the expected signal coverage with at least -65dBm in the office areas, that the controllers correctly manage the RF settings and actively manage client/AP association where necessary and that the infrastructure delivers the required resilience and redundancy. Up to 400 simultaneous clients are observed during working hours and seamless roaming between the four buildings has been demonstrated to the satisfaction of demanding users. Performance tests with a standard laptop have shown real traffic throughput above 200 Mbps symmetric (two spatial streams).

In conclusion, this pilot deployment demonstrates that we have a Wi-Fi architecture that fully meets our goals. Widespread AP deployment is scheduled to start in Spring 2017 and is expected to be complete by the end of 2018.

References

[1] WIND, HPN-CERN OpenLab project, 2010-2012, Milosz Hulboj, Vlad Lapadatescu
http://cern.ch/openlab-wind
http://openlab.cern/sites/openlab.web.cern.ch/files/technical_documents/WirelessControlAndOptimisation.pdf
[2] MCS Value Achieved by Clients at Various SNR Levels, Andrew Von Nagy, Wirelesslan professional http://www.wlanpros.com/mcs-value-achieved-clients-various-snr-levels-andrew-von-nagy/
[3] Single VLAN architecture for Wireless LAN, Alap Modi, Aruba HPE network
http://community.arubanetworks.com/aruba/attachments/aruba/Aruba-VRDs/74/1/Single%20VLAN%20Architecture%20for%20WLAN.pdf
[4] Airmatch RF optimization and Clientmatch® Technologies, Aruba HPE network
http://www.arubanetworks.com/assets/tg/TB_AirMatch.pdf http://www.arubanetworks.com/pdf/solutions/TB_ClientMatch.pdf
[5] Airwave, Aruba HPE network, http://www.arubanetworks.com/assets/ds/DS_AW.pdf