WLAN network planning and setup

Best Practice Document

Produced by FUNET led working group
on wireless systems and mobility (MobileFunet)
(WLAN network planning)

Author: Wenche Backman
Contributors: Ville Mattila/CSC – IT Center for Science, Tanda Tuovinen/University of Helsinki, Siiri Sipilä/Aalto University, Mikko Laiho/University of Helsinki, formerly University of Jyväskylä, Thomas Backa/Åbo Akademi University

December 2010

Document No: GN3-NA3-T4-WLAN-network-planning
Version / date: 10.12.2010
Original language: Finnish
Original title: "WLAN-verkon suunnittelu ja rakentaminen"
Original version / date: 1.0 of 10.12.2010
Contact: wenche.backman (at) csc.fi

FUNET bears responsibility for the content of this document. The work has been carried out by a FUNET led working group on wireless systems and mobility (MobileFunet) as part of a joint-venture project within the HE sector in Finland.

This translated version is based on the Finnish counterpart approved by the board of IT managers of the Funet-organisations (fin. Funetin työvaliokunta) on 10th of December 2010 after an open consultation period of several weeks.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 23 8875, relating to the project 'Multi-Gigabit European Research and Education Network and Associated Services (GN3)'.

Table of Contents

Executive Summary

The cost-efficiency and reliability of wireless local area networks (WLANs) can be ensured through methodical planning. The planning and setup of WLANs can be divided into various sub-areas: coverage area and radio planning, setup and integration with existing infrastructure as well as testing and optimisation. As for planning, it is recommended that lecture halls, conference rooms, lobbies and corridors are prioritised, and that primary attention is be paid to data rates and secondary attention to signal strength. It is also noted that adding access points randomly solely on the basis of customer feedback can lead to a situation in which network control becomes difficult and obtaining an overall picture of the network becomes extremely problematic. Furthermore, it is noted that supporting several versions of the 802.11 standard may result in a lower network capacity than expected.

1 Introduction

The cost-efficiency and reliability of wireless local area networks (WLANs) can be ensured through methodical planning. While external consultants are generally utilised in the planning of WLANs, the transparency of the planning process could be improved by sharing experiences and any lessons learned. The planning and setup of WLANs can be divided into various sub-areas: coverage area and radio planning, setup and integration with existing infrastructure as well as testing and optimisation. Information security should be taken into consideration in the planning phase. Hardware updates can also be thought to form part of WLAN planning and implementation.

.

2 Coverage area and radio planning

The factors that need to be taken into consideration in the planning phase include required coverage area, capacity and costs. As a rule, an extensive coverage area and a high capacity can only be attained at a high cost. In other words, the lower the transmit power, the smaller the coverage area. On the other hand, this leads to higher total capacity, as access points can be placed closer to each other. In controller-based networks the adjustment of transmit power is generally automatic: transmit power is lowered if access points are located close to each other.

If the resources allocated for the setup of WLANs are scarce, it is sensible primarily to cover lecture halls, conference rooms, corridors, lobbies and other spaces where users are anticipated to use a network connection on their laptops or mobile phones. In other words, offices may be excluded from the coverage area.

In terms of capacity, the rule of thumb is that one access point is at maximum able to serve around 10–15 active users who use browsers and e-mail clients. The association limit, i.e. the maximum number of users that may be connected to the access point simultaneously, is considerably higher (around 30–50 users, depending on the access point model). The older the model of the access point, the lower its capacity.

In principle, coverage area planning can be done for both the 2.4 and 5 GHz frequency bands. In the 5 GHz frequency band the signal fades more strongly as a function of distance and as a result of obstacles than in the 2.4 GHz band. The coverage areas are nearly equal, however, since 5 GHz allows for higher transmit power. However, it should be noted with regard to the differences between the frequency bands that directional antennae may not necessarily work in the 5 GHz band.

There are at least three ways of carrying out coverage area planning: the methodical and orthodox way, testing-based planning, and planning on the basis of customer requests. It should be noted, however, that organisations that have planned their network well have not found themselves compelled to add access points at random locations as a result of users complaining of poor reception.

2.1 Methodical and orthodox coverage area planning

When the methodical and orthodox planning method is used, the first phase should involve the assessment of coverage areas using software to calculate signal strengths at determined locations on the basis of access point parameters and the shape of the building. Input parameters for the calculation include access point transmit power and antenna gain as well as the thickness and material of floors and walls. If detailed information on the buildings is unavailable, simpler models can be used to obtain rough estimates.

The calculation software is generally a commercial product and allows the optimal location of access points to be estimated on the basis of calculation results. However, measurements are generally carried out to complement calculation results. Before measurement, access points for testing purposes should set up with the help of the planning software for measuring their signal strengths and interference. Since the data rate

achieved depends on signal strength and the level of interference, the location of access points should be optimised until a satisfactory data rate is achieved at all locations. The data rate required depends on the number of users and their distribution within the building; for example, lecture halls may have a higher concentration of users requiring a high-quality connection.

SiteSurvey is a tool recommended for network planning. Before calculation, blueprints of the building are imported to the SiteSurvey software with details of the materials used. A few potential access point locations are also determined. SiteSurvey then predicts coverage areas and optimises the location of access points to achieve a maximal coverage area. SiteSurvey allows network planning to be carried out with partial or total coverage. Planning can be done with overall coverage in mind even if initial coverage is to be restricted to certain areas. It should be noted, however, that the locations suggested by SiteSurvey for partial coverage may not be ideal in terms of overall coverage. For this reason, SiteSurvey should be used for optimising all access point locations when extending coverage areas to cover the entire building.

2.2 Coverage area planning using tests

In smaller buildings the need for coverage area planning is not as extensive as described above. If it is not sensible to predict coverage areas using calculation software, it is advisable to obtain a number of the selected or potential access points for testing and measure their coverage areas in practice. One access point is sufficient for testing purposes; the access point should be placed in a location that is estimated to be suitable for an access point, e.g. the ceiling of a lobby or corridor. The coverage area of the access point should be measured by checking the data rate at various distances and in various situations instead of focusing only on signal strength. Data rate is a better indication of network quality from the user perspective than signal strength, as it allows consideration to be given to interference level, for example. To obtain an idea of the coverage area of the access point, data rates should be measured at the following points:

- Close to the access point, e.g. directly under it (point A)
- Close to the access point on the same floor
    - behind a bend in a corridor (point B1)
    - behind a single wall relatively close to the access point (point B2)
    - behind a single wall further away from the access point (point B3)
- Directly above the access point on the floor above (point C) or directly below the access point on the floor below (point D).

If the building is relatively homogeneous, these tests may be enough to provide a sufficient picture of the coverage area attainable with one access point. If the building consists of varied structures, additional measurements are required. For example, it is advisable to measure the impact of fire doors on data rates.

Data rates can be measured using regular terminals. Laptops are well suited to this purpose. If possible, it is advisable to measure data rates using more than one network card, at least at one measuring point. This will give an idea of the compatibility of the access point with various network cards.

Data rates can be measured using the iperf software [1],[2]. Iperf is available for at least Windows, Linux and Mac. Testing requires two computers: one of the computers is connected to the same local area network switch as the access point, while the other is used to measure access point performance. It is advisable to check during the first phase that the data rate bottleneck is in fact on the WLAN.

Locating bottlenecks

In the first phase it should be checked that the bottleneck is in fact on the WLAN. If switch performance has been tested recently, this phase can be skipped.

The test is carried out by connecting both computers to the switch and monitoring throughput between them. The speed of the TCP connection should be checked first: to do this, on both computers go to the directory where iperf is installed and run the following commands:

```
# Start one computer as server
```

```
>iperf -s -w 1M -i 1
# and start one computer as client
>iperf -c <server IP address> -w 1M -l 64K -t 30 -i 1
```

This tests the connection using a TCP window size of 1Mb (should be sufficient; the window should not become a bottleneck since the idea is to establish network performance) and 64kb packets (maximum size of TCP packet, a high-quality network connection will achieve maximum throughput at maximum size). Measurements should be taken for thirty seconds, with results printed on screen each second.

The UDP connection should be checked next:

```
# Start one computer as server
>iperf -s -u -w 1M -i 1 -l 1470
# and start one computer as client
>iperf -c <server IP address> -u -l 1470 -b 300M -i 1 -t 30 -w 1M
```

The packet size (1,470 bytes) is the maximum size that can be sent without fragmentation. The bandwidth of the transmission in the example is 300 Mbps, which is the maximum theoretical rate attainable on WLANs.

For the purpose of WLAN testing, the data rates obtained using TCP and UDP should be noted down and borne in mind during testing. If the data rate obtained during WLAN tests is similar, the bottleneck is likely to be found in the fixed components of the network infrastructure. In other cases the bottleneck is the wireless network card and driver, the air interface, the access point or the connection between the access point and the switch.

Performing WLAN tests

The tests should be performed twice at each of the locations mentioned above (A, B1...B3 and C or D), checking the data rate both from access point to client (downlink) and from client to access point (uplink). In both cases one of the computers should be connected to the same switch as the access point and the other computer connected to the WLAN network provided by the access point. It is advisable to encrypt the wireless network set up for testing purposes, preferably using WPA2/AES encryption, so that the impact of encryption on performance can be taken into consideration in coverage area planning. When testing downlink, the computer connected to the WLAN should be run as an iperf server and the computer connected to the switch as a client. The configuration should be reversed when testing uplink (i.e. the computer on the WLAN as client and the one connected to the switch as server). The tests should be performed using the following commands:

```
#TCP test
>iperf -s -w 1M -i 1
>iperf -c <server IP address> -w 1M -l 64K -t 30 -i 1
#UDP test
>iperf -s -u -w 1M -i 1 -l 1470
>iperf -c <server IP address> -u -l 1470 -b 150M -i 1 -t 30 -w 1M
```

If 150Mbps is successfully transmitted using UDP, bandwidth can be increased. This is unlikely to be necessary, however. If possible, the tests should be performed using both TCP and UDP in order to discover the impact of the protocol used on data rate, at least at the points where the best and worst data rates are anticipated. At other points the measurements can be done using UDP only.

Assessment of the number and location of access points

N.B.: This chapter only focuses on network planning using access points with omnidirectional antennae. If directional antennae are used, their asymmetrical coverage areas should naturally be taken into consideration in network planning.

After the WLAN tests it is possible to assess how densely the access points should be deployed to obtain an extensive coverage area (with respect to the desired data rate). The desired data rate limit can be determined independently, giving due consideration to the desired overall coverage area and the budget. However, results

below 10 Mbps should not be allowed in the planning phase. It should be noted, however, that the actual data rate may in the worst case remain below that. Nothing prevents controller-based access points from being placed more densely if the budget allows this.

In practice, the location of Ethernet and power sockets and, in old buildings, the visibility of the access points must also be taken into consideration in the placement of access points. Where possible, the access points should be placed systematically until the entire building has been covered. However, the following factors need to be taken into consideration to ensure the least possible interference among access points:

- Walls and floors/ceilings prevent access point coverage areas from being fully spherical. If there are no obstructions, the coverage area of an access point can be quite extensive, which may cause problems in corridor areas: if the access points of a controller-based network are located within sight of each other (e.g. in long corridors), they may interfere with each other to such a degree that the controller ends up lowering their transmit power. If the access points are to cover adjacent rooms, which may have thick walls, the actual data rate in the rooms may be considerably lower than desired. When planning to set up a network in a building with long corridors, it is advisable to measure whether a satisfactory coverage area could be achieved by placing the access points inside the rooms at regular intervals along the corridor alternating from one side to the other.
- For multi-storied buildings, coverage area planning needs to be three-dimensional: once the locations of access points on one floor have been determined, the access points on the floors directly above and below should not be placed in the same places. This eliminates unnecessary interference by the signals of access points on other floors.

The reason behind the need for three-dimensional coverage area planning is that it is the distance the signal travels inside structures (walls and floors) that is decisive, not the distance from access point to terminal. This is illustrated in Figure 1: if access points are not placed directly below or above each other, the interfering signal is able to fade as much as possible and does not cause unnecessary interference. The same applies to walls: the overall distance travelled through walls may cause the signal to fade more than expected if there is an access point close to the ceiling on one side of the wall and another one below on the other side of the wall.

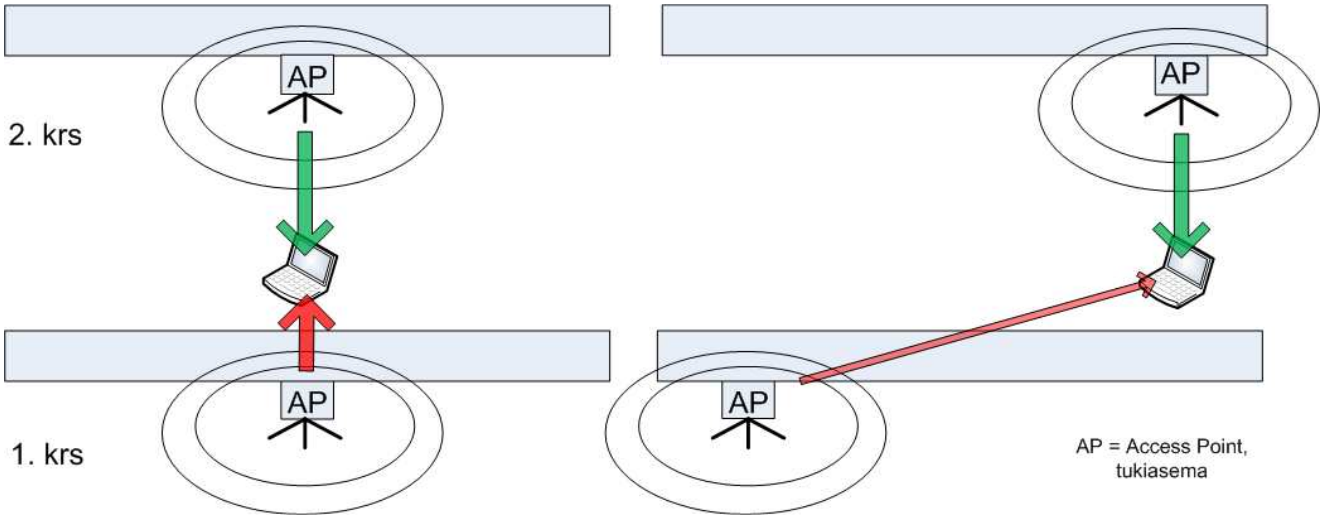

Figure 1. Interference between access points. If access points are placed directly above or below each other, their signals interfere with each other (left). If the access points are located in slightly different places on different floors, the interfering signal fades as it travels through the ceiling, resulting in less interference (right).

An overall plan of the network enables the network to be set up; potential problems can be solved afterwards through network optimisation.

2.3 Coverage area planning on the basis of customer requests

When expanding the first WLAN it may be tempting to set up access points solely on the basis of customer requests. This may even turn out to be a functional solution. However, it should be remembered that it is advisable to completely review access point locations when expanding a network to cover a larger part of the building. If controller-based access points are placed too close to each other, their transmit power is lowered, leading to potential coverage area gaps around the access points. If standalone access points are placed too close to each other, they will cause each other interference, reducing network quality.

2.4 Radio planning

Nearly all networks set up today are controller-based WLANs, which makes maintenance and troubleshooting considerably easier than on networks consisting of standalone access points. Controllers generally involve automatic channel selection and transmit power adjustment, which enable the minimisation of co-channel interference. However, interference is caused by other nearby WLANs and other devices using the same frequency, such as microwave ovens and Bluetooth devices.

The channels selected by the controllers should always be checked. It should also be checked that controllers do not change access point channels too often, as this is an indication of excessive interference. In extreme cases it may be necessary to lock the channels in a certain configuration if the controller is unable to find the optimal configuration.

If the network consists of both controller-based and standalone access points, the channels used by standalone access points should be locked first. It should be noted here that the simultaneous use of standalone and controller-based access points should be avoided as far as possible. Controller-based access points enable the creation of networks with unique SSIDs and settings as necessary. Channels 1, 6 and 11 should be used as they do not cause interference with each other; the channels should be distributed evenly among standalone access points so that nearby access points do not use the same channel. In such an environment the controller selects the suitable channels for its access points, but the result must always be checked.

2.5 Recommendations relating to coverage area and radio planning

With regard to coverage area and radio planning, it is recommended that:

- Lecture halls, conference rooms, lobbies and corridors should be prioritised when setting up WLANs if the entire building or campus area cannot be covered at once.
- At the network planning phase primary attention should be paid to data rates and secondary attention to signal strength.
- WLANs should be set up and extended to cover as extensive an area as possible; the optimisation of access point locations minimises the number of access points required.
- One access point should be presumed to serve around 10–15 active users who use browsers and e-mail clients. The association limit, i.e. the maximum number of users that may be connected to the access point simultaneously, is considerably higher (around 30–50 users, depending on the access point model).

With regard to coverage area and radio planning, the following comment is given:

- Adding access points randomly solely on the basis of customer feedback can lead to a situation in which network control becomes difficult and obtaining an overall picture of the network becomes extremely problematic.

3 WLAN integration with existing infrastructure

WLAN access points and the controller are connected to each other via the building's LAN infrastructure. If the construction of a separate infrastructure for the WLAN is considered, it is advisable to acquire switches supporting the Power over Ethernet (PoE) standard if the maximum output provided by PoE is sufficient for the access points. Access points can be connected to the LAN infrastructure through a number of dedicated patch cabinets in which a the PoE switches are located. PoE is a practical solution as it prevents access points from being disconnected accidentally when the power socket is required for other use. WLAN monitoring has shown that the cause behind problems is often the accidental or wilful disconnection of the access point power cable. If a number of access points are placed far from other access points, their power supply can be ensured using separate power injectors. It should be noted that potential attackers might attempt to access the network using the Ethernet port used by access points, thereby gaining access to VLANs intended for internal use, for example.

More information on WLAN infrastructure will be available in a future Best Practices document.

4 Information security

Information security should be taken into consideration in the network planning phase. More information on WLAN information security can be found in the Best Practices document “WLAN Information Security”.

5 Setting up WLANs

When setting up WLANs, it is extremely important to document the location of all access points by at minimum indicating their room numbers. It is surprisingly easy to forget the location of access points. When setting up controller-based WLANs, it may be sensible to connect all access points to the network before taking them to their actual location; this ensures that they discover the controller and that the software update is straightforward. If the access point has discovered the controller once and the software update has been completed, the access point will be operational in about thirty seconds after it is connected to a power source at its actual location. This enables testing to be performed considerably faster than if the access point was not connected to the network in advance. When setting up WLANs consisting of standalone access points, it is advisable to assign a static IP address to each access point before installation to enable access for control purposes.

As mentioned earlier, the location of Ethernet sockets should be taken into consideration in the placement of access points. If network coverage is desired in places where no Ethernet sockets are available, it is possible to extend coverage using the mesh technology. This means that the access point uses other access points to relay the traffic from the terminal to the fixed network. The University of Helsinki has tested mesh technology with good results: if the traffic volumes are not too high, the method is highly successful. In order to provide as much capacity as possible to the terminals, it is a good practice to offer network access to terminals using the 2.4 GHz frequency and forward the traffic at 5 GHz frequency.

It should also be noted that it is advisable to try to have access point locations taken into consideration when constructing new or renovating old buildings. This can involve the installation of Ethernet sockets and fixtures for mounting hardware to be placed near future access point locations.

With regard to WLAN setup, it is recommended that:

- The location of access points should be documented carefully!
- The location of access points should be taken into consideration when designing new and renovating old buildings.

6 Testing and optimisation

6.1 Coverage area testing

SiteSurvey can be used to measure signal strengths and interference levels using a laptop; SiteSurvey calculates the signal speed on the basis of these parameters. The University of Jyväskylä has used SiteSurvey for both coverage area planning and testing. A signal strength of approximately -70 dBm was considered sufficient. Jyväskylä has also published coverage area maps.

An overall picture of the coverage area can also be obtained by using the ping command or by checking the signal strength bars indicated by the network card software while moving around with a laptop.

As WLANs are set up in a frequency band that is freely available, interference may be quite high in places. Spectrum analysers can be employed to establish the causes for interference. Wi-Spy is a tool that has been tried and tested. It has been noticed, for example, that motion sensors located in office rooms may cause significant interference.

6.2 Optimisation of access point locations

A view of optimisation prospects in the network can be achieved by checking the transmit power levels selected by the controller. This is a fairly straightforward way of detecting opportunities for network optimisation: for example, if the lower transmit power is limited to a number of access points located close to each other, it is probably advisable to move them further away from each other. The optimal scenario is that all access points have the same transmit power.

If it is discovered later that coverage is insufficient or that additional capacity is required in certain places, the addition of access points is simple in controller-based networks, as transmit power and channel are selected in a centralised manner through the controller. It should be remembered, however, that access points should be placed at intervals as even as possible; for example, if additional capacity is required in a certain conference room, it is advisable to place an additional access point at the opposite end of the room instead of next to the other access point.

In conclusion, it should be noted that controller-based networks provide better quality by increasing the transmit power of nearby access points in case of failure by one access point. However, the elimination of coverage area gaps by increasing transmit power requires that the access points are not already using the highest possible transmit power.

6.3 Network parameter optimisation

WLAN performance depends largely on the standards supported. The best network performance can be achieved when the network supports the standard in question only. Networks only supporting the 802.11g standard offer higher data rates than networks that support both the 802.11g and the 802.11b standards. The reason behind this is that in combined networks certain control messages must also be understandable by 802.11b-compatible devices, which do not support faster data transfer. If 802.11b devices are known not to be present on the network, support for the standard can be removed.

Devices compatible with the 802.11n standard are compatible with all previous standards (802.11b and 802.11g in the 2.4 GHz frequency band and 802.11a in the 5 GHz frequency band). In such cases, too, support for older standards lowers network performance [3]. Support for the currently most common standard, 802.11g, must be continued for some time until 802.11n-compatible devices become truly common. The most sensible solution would be to transfer terminals compatible with 802.11n to the 5 GHz frequency or to set up a parallel 802.11n network, in which case the problem would be the sensible distribution of channels between the networks. The best of these alternatives is to transfer 802.11n-compatible terminals to the 5 GHz frequency and to offer a wireless network in the 2.4 GHz frequency to 802.11g (and 802.11b) terminals only.

The WLAN standard provides several opportunities for parameter adjustment. According to the default parameters of the standard, packets are fragmented if their size exceeds 2,346 bytes. Other transmitters are silenced using an RTS/CTS handshake prior to the transmission if the packet size exceeds 2,347 bytes, i.e. only if the fragmentation threshold has been raised. Several other parameters that affect the use of radio resources and power consumption can also be adjusted, including beacon interval, retry limits and listen interval. Although the parameters can be adjusted in various ways, a degree of caution is advised as the consequences of adjustments may be difficult to measure and, occasionally, even surprising. Controllers also generally involve a degree of radio resource management.

As with all networks, WLANs are used for services that set various requirements for the network. VoIP calls require short delay and guaranteed bandwidth, whereas browsing and e-mail are services which are less susceptible to disturbances. WLANs use QoS priority levels Voice, Video, Best Effort and Background. Cisco's controller, for example, enables the specification of data rates and share of bandwidth for all of these. The allocation of a large share of the bandwidth available to certain services and the setting of a high user-specific bandwidth allow improved support for VoIP calls, for example. On the other hand, this leads to a decline in the number of simultaneous users. Depending on the number of users, these adjustments also result in lower data rates for other users when there is voice traffic on the network.

Network service level can also be improved by enabling smooth handover between access points. On controller-based networks, where access points are connected to each other via controllers, this service is easy to implement. Handover between access points can be accelerated by making advance preparations on the controller, eliminating the need for reauthentication by the user.

With regard to network parameters, it is recommended that:

- Use of the 802.11b standard should be monitored. It should also be established whether there are 802.11b/g clients on the network using the 802.11b standard for one reason or another. It is advisable to discontinue support for the 802.11b standard at some point to allow improved utilisation of the air interface through the exclusive use of the 802.11g standard.

7 Updating WLANs

Wireless networks are not static; the technology is developing constantly and users are becoming increasingly demanding. When updating hardware, a network infrastructure consisting of mixed brands and models of devices should be avoided. In other words, as many access points and/or controllers as possible should be updated at the same time.

References

[1] http://en.wikipedia.org/wiki/Iperf

[2] http://sourceforge.net/projects/iperf/

[3] http://www.smallnetbuilder.com/content/view/30224/100/

Glossary

AES	Advanced Encryption Standard
CSC	CSC - IT Center for Science Ltd.
Funet	Finnish University and Research Network
IP	Internet Protocol
MobileFunet	a working group on wireless systems and mobility led by Funet
SSID	Service Set Identification
VLAN	Virtual Local Area Network
WLAN	Wireless Local Area Network
WPA and WPA2	Wi-Fi Protected Access