

CESNET Technical Report 2/2009

Overlapping eduroam Networks Operated by Different Organizations

JAN FÜRMAN

Received 31.3.2009

Abstract

This paper describes one of the most problematic part of the **eduroam** network deployment in heterogeneous environment and its possible solution. The problem described in this paper may occur whenever two or more organizations providing the **eduroam** wireless network cover the same physical space and their radio networks overlap. This well known issue is also mentioned in the European roaming policy. The aim of this article is to describe the general technical solution - not to provide the detailed configuration procedure. This would be just a useless replication of manual pages.

Keywords: eduroam, wireless controller, campus WiFi solution

1 Introduction

eduroam which stands for Education Roaming, is a RADIUS-based infrastructure that uses 802.1x security technology to allow for inter-institutional roaming.



Figure 1. The **eduroam** logo

Having an **eduroam** account allows users visiting other institutions connected to the **eduroam** infrastructure to log on to the visited WLANs using the same credentials which they use at their home institution. Depending on the local policies at the visited institutions, **eduroam** participants may also have additional resources at their disposal. All this with a minimum administrative overhead. See project web pages for more details.¹

This text is focused on description how to deploy radio networks in campuses where more than one independent academic institution is located on limited area. This specific environment determines the requirements on the WiFi radio networks of each organization. The interferences and signal distortion between the networks should be avoided while providing the signal coverage of the campus area as good as possible and without signal gaps. Unfortunately, these requirements are in contradiction and in practice it is usually impossible to avoid the interference between radio networks of different organizations completely.

¹ <http://www.eduroam.org>

The efficiency of deploying radio networks (not only the expenses but also the radio bandwidth utilization) should be taken into account as well. Having two independent radio networks in an area where only one is sufficient is obviously a waste of resources. With accent to user comfort, the best solution seems to be to interconnect all radio networks in the campus and to ensure transparent roaming between them. From the user's point of view, all borders between radio networks would be hidden and whole campus would be considered by user as one unified radio network.

The **eduroam** system can ensure unified network environment in campuses where more than one university/organization is located. It allows transparent network access for users within the whole campus but it also arises couple of technical issues, namely radio signal overlapping and interferences between radio networks of different organizations. To minimize these effects, some complex radio channel and transmit power settings are required. This work is not easy and needs cooperation between network operators of all involved organizations. Of course, the more Access Points are in the game, the worse the problem is. The management is difficult by itself if some institution has tens of APs (nothing uncommon in these days) and thanks to interferences with other radio networks the serious problems can arise. As it was mentioned above, the biggest of the problems is overlapping of two or more radio networks with the same SSID.

2 Problem Description and Possible Solutions

The problem comes up when the radio signals of two or more networks operated by different organizations but using the same network identifier (**eduroam**) are accessible at one physical location. A computer located at this place can then keep flipping between access points (AP) operated by different organizations (see Figure 2).

Although these networks share the same SSID (the radio network identifier "eduroam", they provide access to different networks with different IP address spaces, different DNS servers, different filtration policies etc. On every re-association to a different network the client computer must re-authenticate, obtain a new IP address and a whole network configuration. This process is time consuming (it may take several seconds) and makes users unhappy because of temporary connection breakdowns. The computer network is becoming useless for users if these "jumps" occurs frequently.

In theory, there are several possible solutions to this problem. For instance, it is possible to change the SSID of each organization (e.g., by adding a suffix which is describing an organization: `eduroam_cesnet`, `eduroam_vscht`, ...). This solution is mentioned in the **eduroam** roaming policy document (see [1]). It is easy to implement and completely solves the problem of flapping clients but it has also couple of disadvantages. It's necessary to change the configuration on client's device and the beauty of transparent roaming between different organizations is totally lost – client must reconnect to different SSID during changing his/her position between institutions.

Another possible solution is to deploy a single VLAN across whole campus.

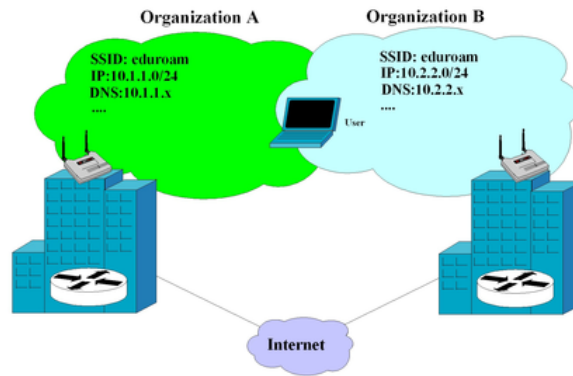


Figure 2. Campus WiFi deployment

APs of all organizations are connected to this single VLAN and they are operated in a shared IP address space. The whole network would appear as one although it would consist of access points of different organizations. The problem of flapping clients would not be solved completely although it would improve a bit. The client would still have to re-authenticate after a “jump” to an AP belonging to a different organization but the IP address it would get would be the same (from the same IP address space, from the same DHCP server) and all open network connections would stay up. The client would observe just a short network malfunction during re-authentication. However, this solution has many disadvantages. Management of the APs would be very complex in such a wide network, it would not be clear who is responsible for what (either the local network administrator or the administrator from the organization operating the central VLAN). It also imply very complex and unclear troubleshooting (who should the user contact in case of network malfunction?). Last but not least, this solution requires one big IP address space which must provide IP addresses for every AP in the campus. To obtain and to manage such a large IP address space is much more complicated than for many small IP spaces.

The best solution of the flapping clients problem seems to be the deployment of a wireless controller.

3 A Solution Based on WLC (Wireless Lan Controller)

Deployment of wireless controllers can very effectively eliminate the problem described above. Each organization in the physical area operates its own controller (this is necessary for preservation of full control over their radio networks). All APs are connected to the WLC via L3 (or less frequently via L2) tunnel which transports all data and control traffic between the AP and the WLC. The tunnel protocol is known as LWAPP (Light Weight Access Point Protocol) – a CISCO proprietary protocol). Its standardized equivalent (CAPWAP) that would allow interoperability between devices of different vendors is being prepared by IETF. All traffic from radio network is transported to the WLC via the LWAPP tunnel; the AP acts just as a “remote radio” without any sophisticated logic. All data processing is done

on the controller. The roaming of the client (which is associated to an AP of any particular WLC) to the AP connected to the same WLC takes only couple of milliseconds. All authentication data is stored on the WLC and the client does not have to re-authenticate to the network again (also DHCP assignment of the IP address is omitted).

The controller-based solution has also other additional benefits like dynamic control of the radio characteristics or significant simplification of configuration which is apparent mainly with big number of access points, etc. The entire configuration setup is executed centrally on the WLC. That brings significant simplification (the individual autonomous APs do not have to be reconfigured manually any more). The WLC ensures automatic control of radio parameters (such as transmitting power, channel selection, etc) and thereby completely eliminate complex manual configuration of the APs. Transmitting power and radio channel of any particular AP is set by the WLC based on information of signal strength from the neighbour APs (the WLC collects this information from every AP). The WLC tries to cover whole area around APs as well as possible (without any signal gaps) and simultaneously without wasted interferences between APs. Another interesting and important feature is the rogue AP detection. The WLC can uncover (and in limited manner also locate) malicious APs which are located on the site but which are not part of our radio network.

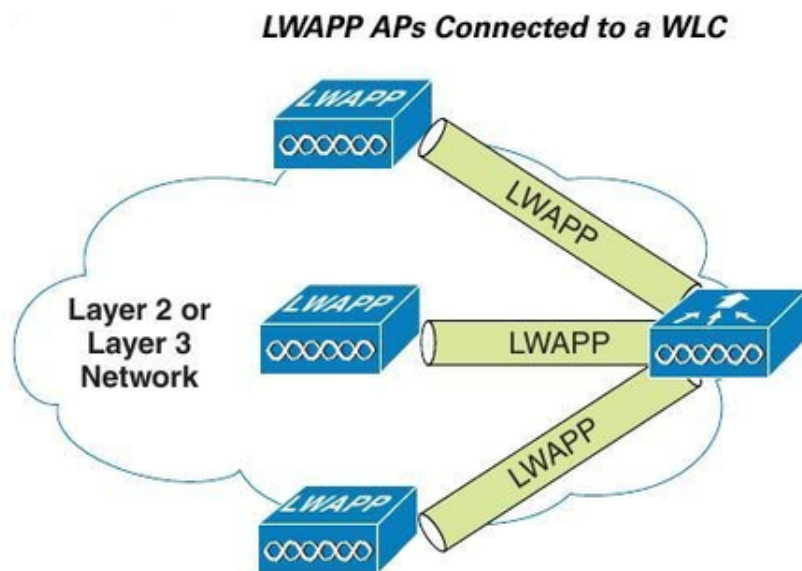


Figure 3. LWAPP access points

The detection is based on data provided by every AP connected to the WLC. This is very important instrument for network operator during network maintenance and especially for solving security incidents.

The described architecture has its drawbacks too. It is relatively costly and proprietary (but this should change in near future with the CAPWAP protocol deployment). Only LWAPP capable APs can be connected to the WLC. But most widespread APs like CISCO 1230 or 1240 series do not support LWAPP. Fortunately,

these APs can be very simply converted to the LWAPP mode – just by upgrading the firmware.




Figure 4. Client roaming

The fast roaming works well not only within the scope of a single WLC but it can be also configured for roaming between APs connected to different WLCs. For this, the controllers should be connected into one “mobility group” (see [2] for more info). All WLCs exchange important data (about radio parameters of their APs and about associated clients) with other WLCs within the scope of the mobility group. As a result, one global management of radio parameters is maintained within the mobility group (coordinated control of channel allocation and transmitting power settings). At the same time, fast roaming between APs under control of different WLCs is possible. Although the WLCs are connected to a single mobility group, their autonomy is preserved and each controls the APs connected to it thus keeping them under the control of their owners.

Let’s assume that a client is connected to the AP B (which is connected to the WLC1) for the first time (see Figure 4). After successful authentication to the network, the client gets its IP address from the DHCP server. The whole client’s IP payload goes from the client via the AP B to the WLC 1 and then to the wired network. During the first association of the client, the WLC 1 sends information about the client (its MAC address, IP address, credentials) to all controllers within the scope of the mobility group. If the client roams to another AP which is under control of another WLC, e.g., to the AP C, the lengthy authentication and IP address assignment does not occur again because the WLC2 already knows the client and it can immediately start transferring the data. Once any WLC finds out a connection of a client which had been associated to any other controller before (i.e., to an AP under control of any other WLC), it makes bidirectional tunnel to transport user’s payload. This tunnel is established between a WLC where the client associated for the first time (its “anchor controller”) and the WLC where is the client associated currently. The whole user payload is carried through this tunnel to the network. This means that the routing is not affected by roaming and the client can keep its original IP address (assigned under the WLC1) in spite of the WLC2 is operating

with a completely different IP address space. All APs of an organization are connected (at the IP layer) to the organization's controller via the Light Weight Access Point Protocol giving the controller the full control over them. If the client "jumps" between the APs of different organizations, both involved controllers immediately use a tunnel (created during the initial configuration of the mobility group) to transport the client's payload to the controller to which the user was associated for the first time. The data stream is then decapsulated and routed to the wired network as if the client was connected directly to this "initial" controller.




Figure 5. Mobility group

The concept described above is used for solving the problem with interferences of many radio networks (with the same SSID) within single campus area. Each organization runs its own autonomous WLC and these controllers are interconnected to single common mobility group. Client then can move around the whole campus and during roaming between APs still keep the same IP address (the one it has got during the first association to the network). No re-authentication neither changing of the client's IP address is required and the user does not observe any connection breakdown. The time for roaming between two APs (no matter if they are connected to the same controller or to different controllers in one mobility group) is assured to be shorter than 30ms which is sufficient even for Voice over IP applications.

4 Case Study – CTU Campus

This technical solution is successfully operated at the CTU (Czech Technical University) campus in Dejvice, Prague, Czech Republic, where six academic organizations are located very closely in one area. Each of the organizations operates its own

autonomous radio network which is connected to the **eduroam** roaming system. Our solution is based on the CISCO platform because all radio equipment (access points) used by organizations in the campus are from this vendor. However, this solution is generic and can be based also on products of other manufacturers (unfortunately no other solutions was tested because the equipment by other vendors was not available).

This solution provides optimal signal coverage almost in the whole campus. Interferences are minimized because radio parameters (channels assignment and transmitting power) of all APs are controlled centrally by the WLCs which are exchanging data about radio parameters of their APs. The autonomy of each radio network is still preserved because each organization run its own controller and have all its APs under its full control. Each organization usually transmits also its own SSID which has only local relevance and is used only by users local to the organization. The only common SSID in the whole campus is **eduroam**.

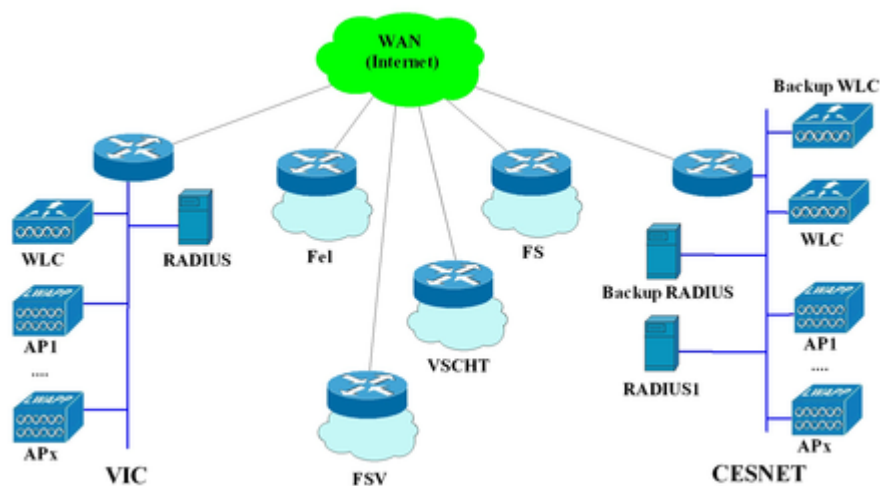


Figure 6. Campus scheme

Because a controller itself is potentially a single point of failure it is necessary to provide an appropriate backup. If the WLC fails the whole radio network becomes unusable – the APs can not work without a controller. One backup controller for all organizations is implemented in the campus. If any controller fails, its APs automatically connect to the backup WLC and continue working (just under little bit restricted mode). The backup WLC is configured to support only the **eduroam** SSID and the private IP address space 10.20.0.0/16 (with NAT). Local SSIDs of the participating organizations neither their IP address spaces are not available when using the backup WLC. This setup is sufficient for potentially supporting all client devices within the campus in a case of a failure.

5 Configuration Description

The basic configuration of the WLC is relatively easy and is quite well described in the original CISCO documentation (e.g., in [4]). The first step to deploy the architecture described above is the configuration of the radio network with SSID **eduroam**. An example configuration of the **eduroam** network can look like this:

```
(Cisco Controller) >show wlan 1
WLAN Identifier..... 1
Profile Name..... eduroam
Network Name (SSID)..... eduroam
Status..... Enabled
MAC Filtering..... Disabled
Broadcast SSID..... Enabled
AAA Policy Override..... Disabled
Number of Active Clients..... 14
Exclusionlist Timeout..... 60 seconds
Session Timeout..... 1800 seconds
Interface..... eduroam-int
WLAN ACL..... unconfigured
DHCP Server..... Default
DHCP Address Assignment Required..... Disabled
Quality of Service..... Silver (best effort)
WMM..... Allowed
CCX - AironetIe Support..... Enabled
CCX - Gratuitous ProbeResponse (GPR)..... Disabled
CCX - Diagnostics Channel Capability..... Disabled
Dot11-Phone Mode (7920)..... Disabled
Wired Protocol..... None
IPv6 Support..... Disabled
Peer-to-Peer Blocking Action..... Disabled
Radio Policy..... All
Radius Servers
  Authentication..... 10.20.30.40 1812
  Authentication..... 10.20.30.41 1812
  Accounting..... 10.20.30.40 1813
  Accounting..... 10.20.30.41 1813
Local EAP Authentication..... Disabled
Security
  802.11 Authentication:..... Open System
  Static WEP Keys..... Disabled
  802.1X..... Disabled
  Wi-Fi Protected Access (WPA/WPA2)..... Enabled
    WPA (SSN IE)..... Enabled
      TKIP Cipher..... Enabled
      AES Cipher..... Enabled
    WPA2 (RSN IE)..... Enabled
```

```

    TKIP Cipher..... Enabled
    AES Cipher..... Enabled
Auth Key Management
    802.1x..... Enabled
    PSK..... Disabled
    CCKM..... Disabled
    CKIP ..... Disabled
    IP Security..... Disabled
    IP Security Passthru..... Disabled
    Web Based Authentication..... Disabled
    Web-Passthrough..... Disabled
    Conditional Web Redirect..... Disabled
    Auto Anchor..... Disabled
    Cranite Passthru..... Disabled
    Fortress Passthru..... Disabled
    H-REAP Local Switching..... Disabled
    Infrastructure MFP protection..... Enabled
    (Global Infrastructure MFP Disabled)
    Client MFP..... Optional
    Tkip MIC Countermeasure Hold-down Timer..... 60
Mobility Anchor List
WLAN ID  IP Address  Status
-----  -

```

The next configuration step is the interconnection of all controllers into one common mobility group. This ensures the exchanging of important data about associated clients and radio parameters among all the WLCs within the campus. The configuration is easy and well documented in the WLC manual. Each WLC within the mobility group is identified by its IP and MAC address. It's also necessary ensure free transport of data between controllers on eventual intermediated firewalls. The example of the mobility group follows:

```

(Cisco Controller) >show mobility summary
Symmetric Mobility Tunneling (current) ..... Enabled
Symmetric Mobility Tunneling (after reboot) .... Enabled
Mobility Protocol Port..... 16666
Mobility Security Mode..... Disabled
Default Mobility Domain..... DEJVICE
Mobility Keepalive interval..... 10
Mobility Keepalive count..... 3
Mobility Group members configured..... 8
Controllers configured in the Mobility Group
MAC Address  IP Address  Group Name  Status
11:22:33:44:55:66  1.2.3.4  DEJVICE  Up
22:33:44:55:66:77  2.3.4.5  DEJVICE  Up
33:44:55:66:77:88  3.4.5.6  DEJVICE  Up

```

```

44:55:66:77:88:99  4.5.6.7  DEJVICE  Up
55:66:77:88:99:11  5.6.7.8  DEJVICE  Up
66:77:88:99:11:22  6.7.8.9  DEJVICE  Up
77:88:99:11:22:33  7.8.9.1  DEJVICE  Up
88:99:11:22:33:44  8.9.1.2  DEJVICE  Up

```

The next important (but not necessary) configuration step is the restriction of AP connection to any particular WLC. An organization can define which APs can connect to its controller. Typically, only its own APs are allowed to connect the organization's controller. The backup WLC is open for any AP from the whole campus. These restrictions help to prevent a very unpleasant effect that might occur during a failure of a WLC of any particular organization. If an AP loses the connection to its primary WLC, it can theoretically connect to any other WLC (if it has free resources) within the mobility group. Protecting their WLCs against unauthorized connection from foreign APs is in interest of every organization. The problem described above really occurred in our practice.

Other configuration steps are standard and well described in vendor documentation.

6 Conclusions

After having operated the described architecture for almost a year, we can say that our experience is good. All the problems caused by channels overlapping have been solved by deploying the WLCs.

However, the advertised roaming speed in milliseconds works only for clients supporting Cisco Compatible Extensions (CCX) version 4 and higher. This requirement prevents all Linux based computers from fast roaming as they do not support the CCX at all. These machines must re-authenticate against a RADIUS server during each "jump". This operation takes about 3 seconds. After the successful authentication the whole mechanism works the same way as in the "fast roaming" case. It means that the client device keeps its initial IP address, and in the case of roaming between two WLCs all user payload is transported via special tunnel to the initial (anchor) WLC. Thus the client's routing is preserved and its opened network connections are not lost. The short time network drop-out during the actual roaming time is acceptable for majority of common applications like email reading or web browsing. It is, however, unpleasant for applications sensitive to packet loss like IP telephony.

Unfortunately, the "fast roaming" does not work also with Intel WiFi cards under the Windows XP system which fully supports CCXv4. What is the cause of the problem is not clear – it might be an error in the WLC software. Because this combination of OS and WiFi card is very common, the problem is serious. With the CISCO wireless card and an appropriate supplicant, the "fast roaming" between controllers works as expected. The issue might be caused by any of the WLC, the Windows system, the supplicant, or the wireless card driver. We are working hard trying to find a solution of this problem in cooperation with vendors.

Unfortunately, we can't expect support of the CCX in Linux. As an alternative, the PKC (Proactive Key Caching) might be used in the future. The PMK (Pairwise

Master Key) is cached on the WLC during the initial client authentication. When the client then roams to any other AP, this stored key is used without requiring the client to full authentication against the RADIUS server again. This principle is similar to CISCO solution but it is not proprietary and dependent on the CCX extensions.

In despite of the problems described above, the WLC deployment seems to be a good choice and we can recommend it to other campuses with overlapped radio networks with the same SSID.

7 Acknowledgements

Some pictures in this article were adopted from web pages of Cisco Systems, Inc. **eduroam** is a registered trademark of TERENA.

References

- [1] SIMONSEN, D. et al. *Roaming Policy and Legal Framework Document – Part 2*. GN2 Project, Deliverable DJ 5.1.3,2, 2006. Available online².
- [2] CISCO SYSTEMS. *Cisco Unified Wireless Network: Introduction*. Available online³.
- [3] CISCO SYSTEMS. *Cisco Wireless LAN Controller Configuration Guide, Release 5.2*. Available online⁴.

² http://www.geant2.net/upload/pdf/GN2-06-080v5-Deliverable_DJ5-1-3_2_Roaming_Policy_and_Legal_Framework-Part2_20061108094807.pdf

³ http://www.cisco.com/en/US/netsol/ns340/ns394/ns348/ns337/networking_solutions_package.html

⁴ <http://www.cisco.com/en/US/docs/wireless/controller/5.2/configuration/guide/Controller52CG.html>