How to deploy eduroam on-site or on campus (ADVANCED)

- eduroam set up on Campus: IdP and SP
  - eduroam SP
    - Basic deployment considerations for wireless LANs
    - Obligations of eduroam SPs
    - Set up of networking equipment in the network core
    - Set up of eduroam SP RADIUS servers
      - Version information
      - Installation
      - Sample config directory
      - Base configuration / logging / F-Ticks
      - Client definition
      - Request forwarding
      - Goodies
        - Running FreeRADIUS as non-root user
        - F-Ticks
      - CUI for eduroam SP
        - Create a log module
        - Enable modules
        - Client definition
        - Policy
        - Attributes
        - CUI filtering
        - Using policies and modules in your eduroam virtual server
      - Caveats
  - eduroam IdP
    - Obligations for an eduroam IdP
    - Selecting EAP types
      - Choices depending on the Identity Management System
      - Anonymous outer identities
      - Choices depending on the envisaged devices
    - EAP Server certificate considerations
      - Consideration 1: Procuring vs. creating your own server certificate
      - Consideration 2: Recommended certificate properties
      - Consideration 3: Which certificates to send in the EAP exchange
    - Set up of several popular RADIUS servers
      - FreeRADIUS
        - Setting up FreeRADIUS
          - TLS server certificate
          - EAP type configuration
          - Virtual server eduroam: enable EAP, make Operator-Name conditional
          - Virtual server eduroam-inner-tunnel
          - User database: flat file
          - Local authentication for your realm
          - Processing incoming requests
          - CUI for eduroam IdP
          - Goodies
        - More information
      - Radiator
        - Clients
        - Realms and VLAN assignment
        - PROXY example
        - Secure authentication with EAP-TLS
        - EAP-TTLS or EAP-PEAP
      - Microsoft NPS
      - Cisco ACS
      - ClearPass
    - RADIUS Accounting in eduroam
    - Provisioning configuration details of supplicants to end users
      - eduroam CAT
      - Others
  - RADIUS/TLS: Using radsecproxy as an add-on to an existing RADIUS server
    - Goals and Prerequisites
    - Installation
    - Sample configuration
      - Local Logging
      - F-Ticks
      - RADIUS/TLS
    - Client definition and request forwarding
  - Set up of Dynamic Discovery (in federations which have RADIUS/TLS enabled)
    - What is Dynamic Discovery?
    - Adding Dynamic Discovery hints to the IdP's DNS zone
    - NAPTR explained
- Testing and supporting end users
- Wired eduroam
  - Connecting wired infrastructure to eduroam

- End user experience
    - Network indication
    - Support for multiple configurations
    - Encryption of payload

eduroam set up on Campus: IdP and SP

The following sections provide detailed information for the two roles eduroam IdP and eduroam SP, respectively.

The eduroam SP section explains the obligations for an eduroam SP and lists what should be taken into account in a wireless LAN deployment, and the set up of a RADIUS server for use in eduroam.

The eduroam IdP section explains the obligations for an eduroam IdP, the set up of several popular RADIUS servers, and means to provision configuration details of supplicants to end users.

eduroam SP

Basic deployment considerations for wireless LANs

An eduroam wireless network is a wireless network. This sounds trivial, but it is important to keep in mind that

- a poorly managed Wireless LAN won't magically become better by naming it eduroam. Before diving into eduroam-specific configuration, make sure you understand how to manage
    - WiFi coverage
    - bandwidth requirements
    - enough DHCP addresses to accomodate all clients
- by naming the network eduroam, you are becoming part of a world-wide recognised brand. Arriving users will think of this being an eduroam network, with a set of expectations for such networks. If your wireless network fails to deliver in the points mentioned above, users will consider this an eduroam failure and your installation will hurt the global brand eduroam, not only your own site and users.

This section provides general advice regarding eduroam deployment on a wireless LAN. It does not include information on general WLAN network planning and setup, it only covers topics essential to deploying eduroam on an already setup wireless LAN.

Obligations of eduroam SPs

The basic requirement for and eduroam SP is that the underlaying WLAN must be able to support IEEE 802.1X authentications, WPA2/AES support and, if you also want other networks, multi-SSID support. This is usually the case with today's network equipment. If you want to distinguish traffic beloning to the eduroam network from other traffic, you also need to deploy VLANs in your network.

For eduroam, you need to add information of the RADIUS server that you will be using to your WLAN controller (or stand-alone access point). As a pure eduroam SP, the RADIUS server in question is likely the one of your national federation. If you are both an eduroam IdP and an eduroam SP, the RADIUS is your own RADIUS server. You will need to add the IP address of the RADIUS server as well as the shared secret, which is basically a string of characters that has been agreed on by you and the operator of the RADIUS server. You may also have to add information about the ports to use, which are 1812 for authentication and 1813 for accounting.

Once you have added the RADIUS server you need to create the eduroam SSID. This must be a network with 802.1X and WPA2/AES enabled and the SSID must be eduroam and this SSID needs to be broadcasted. For this eduroam network, you still need to define that the RADIUS server defined previously need to be used.

In this wiki it is not possible to keep up-to-date guidelines on how to set up eduroam on all wireless equipment on the market. The best way to set up eduroam on your network is to do the initial setup according to the manufacturer's guidelines and thereafter, check the same guidelines on how to apply the eduroam-spesific settings mentioned above. However, a few guidelines are available through the links below

- Cisco controller
- Aruba
- Lancom

In order to check which ports should be open for the eduroam end users, please check out the eduroam Policy Service Definition document, particularly Chapter 6.3.3.

Set up of networking equipment in the network core

Since an eduroam hotspot always uses the RADIUS protocol to connect to a RADIUS authentication server, your network setup must allow this RADIUS communication. This includes opening firewalls for traffic from the WLAN equipment (AP/Controller) to **UDP port 1812** (do not confuse this with TCP!). The RADIUS protocol can easily create UDP fragments, and will not function fully without **UDP fragmentation support**. Be sure to check your equipment whether forwarding of UDP fragments is supported and allowed. For accounting the **UDP port 1813** also needs to be opened.

If you deploy your own RADIUS server for eduroam SP purposes (see below), also make sure that its own uplinks to your National Roaming Operator are open in the same way.

Set up of eduroam SP RADIUS servers

FreeRADIUS is a very versatile and freely available RADIUS server under the GPL license. Setting up FreeRADIUS as an SP is a rather straightforward task, since it merely needs to forward requests from NASes to other RADIUS servers. In particular, it does not need to authenticate users. The following configuration enables your FreeRADIUS server to be an eduroam SP. At the same time, it is the baseline from which to establish an eduroam IdP configuration, if that is envisaged for a later stage.

Version information

This document is in migration from FreeRADIUS 2 to FreeRADIUS 3. We recommend using the last available version of the stable FreeRADIUS 3 branch. It's easy to compile version 3 (and create packages) if your distribution doesn't provide recent packages. (On Ubuntu/Debian with "make deb" for instance and "rpmbuild -ba redhat/freeradius.spec" should help you on Red Hat based systems.)

Some of the filesystem paths changed between version 2 and 3. The /etc/raddb/modules directory is now split between /etc/raddb/mods-available and /etc /raddb/mods-enabled, plus some of the configuration can be found in /etc/raddb/mods-config. Note that when a module isn't called from the rest of the configuration, placing it in mods-enabled doesn't mean it's active: only that it's available in the rest of your configuration.

Installation

FreeRADIUS is written in C and can be compiled with the usual UNIX compilation sequence. After unpacking the source into a directory of your choice, do

```
./configure --prefix=<your preferred install dir> --sysconfdir=<your preferred configuration base dir>
make
make install
```

In the examples below, we assume the installation is done for --prefix=/usr/local/freeradius/ and the configuration dir is --sysconfdir=/etc

Sample config directory

Base configuration / logging / F-Ticks

The main configuration file is /etc/raddb/radiusd.conf; it does not require many changes from the shipped default.

The following lines are important for eduroam operation: a server status probing mechanism called Status-Server is enabled in the security section. Make sure the config file contains the following security stanza

```
security {
           max_attributes = 200
           reject_delay = 0
           status_server = yes
}

proxy_requests      = yes
```

(From the default distribution, only reject_delay needs to be changed.)

FreeRADIUS is capable of both IPv4 and IPv6. By default, both are enabled in the listen {} section of sites-enabled/default so we'll duplicate them in our new sites-enabled/eduroam configuration. (The listen {} directives used to be in /etc/raddb/radiusd.conf for FreeRADIUS 2.) You can leave out the IPv6 part if your server shouldn't do IPv6.

The logic in the server is defined by activating modules in a certain order. These modules are separately defined in the /etc/raddb/mods-enabled/ subdirectory (and configured in /etc/raddb/mods-config/ where applicable). The order of activation of these modules is defined in so-called virtual servers, which are defined in the /etc/raddb/sites-enabled/ directory. For our eduroam SP purposes, we only need one virtual server "eduroam" and call very few of the modules. It needs to contain as a minimum:

```
server eduroam {

        listen {
                type = "auth"
                ipaddr = *
                port = 0
        }
        listen {
                type = "acct"
                ipaddr = *
                port = 0
        }
        listen {
                type = "auth"
                ipv6addr = ::
                port = 0
        }
        listen {
                type = "acct"
                ipv6addr = ::
                port = 0
        }

        authorize {
                # only use filter_username from version > 3.0.7 on
                filter_username
                                update request {
                        Operator-Name := "1yourdomain.tld"
                                                # the literal number "1" above is an important prefix! Do not
change it!
                }
                # if you want detailed logging
                auth_log
                suffix
        }

        authenticate {
        }

        preacct {
                suffix
        }

        accounting {
        }

        post-auth {
                # if you want detailed logging
                reply_log
                Post-Auth-Type REJECT {
                        reply_log
                }
        }

        pre-proxy {
                # if you want detailed logging
                pre_proxy_log
                if("%{Packet-Type}" != "Accounting-Request") {
                        attr_filter.pre-proxy
                }
        }

        post-proxy {
                # if you want detailed logging
                post_proxy_log
                attr_filter.post-proxy
        }
}
```

The multitude of sections in this above configuration is often confusing to new-comers. The order of execution when proxying a request are:

```
authorize  authenticate  pre-proxy
```

Then, the packet is proxied to an upstream server. When the reply comes back, the execution continues:

```
post-proxy  post-auth
```

Every stanza contains names of modules to be executed. Let's revisit them one after another:

- auth_log: logs the incoming packet to the file system. This is needed to fulfill the eduroam SP logging requirements.
- suffix: inspects the packet to look for an eduroam style realm (separated by the @ sign)
- pre_proxy_log: logs the packet to the file system again. Attributes that were added during the inspection process before are then visible to the administrator - great for debugging
- attr_filter.pre-proxy: strips unwanted attributes off of the request before sending the request to upstream
- post_proxy_log: logs the reply packet to the file system - as received by upstream
- attr_filter.post-proxy: strips unwanted attributes off of the reply, prior to sending it back to the Access Points (VLAN attributes in particular!)
- reply_log: logs the reply packet after attribute filtering to the file system

The paths where the logs are written to, and the files with the list of permitted attributes for filtering, are defined in the corresponding module definitions in /etc/raddb/modules/<name-of-module>.

If attr_filter.pre-proxy is enabled (as per the example above), then by default Operator-Name and Calling-Station-Id are stripped from the proxied request. In order for them not to be removed, add the attributes to /etc/raddb/attrs.pre-proxy (FreeRADIUS 2) or /etc/raddb/mods-config/attr_filter/pre-proxy (FreeRADIUS 3). This is a more sensible default for eduroam:

```
DEFAULT
        User-Name =* ANY,
        EAP-Message =* ANY,
        Message-Authenticator =* ANY,
        NAS-IP-Address =* ANY,
        NAS-Identifier =* ANY,
        State =* ANY,
        Proxy-State =* ANY,
        Calling-Station-Id =* ANY,
        Called-Station-Id =* ANY,
        Operator-Name =* ANY
```

Since the eduroam SP with this configuration will statically use RADIUS to its upstream federation-level server, activation of F-Ticks reporting is not strictly necessary. It is thus described only in the "Goodies" section below.

Client definition

FreeRADIUS defines the connected RADIUS clients in the file /etc/raddb/clients.conf. This file needs to hold all your connected Access Points (and/or wired eduroam-enabled switches, if you have these instead of Access Points). You set a shared secret for each client and define these in the config file as follows:

```
client antarctica-access-point-1 {
   ipaddr                        = 172.25.1.55
   netmask                       = 32
   secret                        = yoursecret12345
   shortname                     = southpole-11g
   virtual_server                = eduroam
   require_message_authenticator = yes
}
```

There are more (optional) settings for clients; please consult the comments in clients.conf for more detail. One option, the "virtual_server" one, enables your RADIUS server to serve more purposes than only eduroam: you can define several other virtual servers for other RADIUS purposes, and link clients to these. That is beyond the scope of this documentation, though.

If you want to connect your clients over IPv6, the syntax is only slightly different:

```
 client antarctica-access-point-2 {
    ipv6addr                      = 2001:db8:1:789::56
    netmask                       = 128
    secret                        = yoursecretABCDE
    shortname                     = southpole-11n
    virtual_server                = eduroam
    require_message_authenticator = yes
}
```

Request forwarding

FreeRADIUS contains a wealth of options to define how requests are forwarded. These options are defined in the file /etc/raddb/proxy.conf. For a single eduroam SP, these may seem overkill, but the required definitions for that purpose are rather static. Assuming you have two upstream servers to forward requests to, the following configuration will set these up - you only need to change the IP addresses and shared secrets in home_server stanzas. The realm NULL will reject authentication requests missing an @ sign, for example Windows always first tries its domain\hostname to authenticate when connecting the first time to eduroam. This authentication would otherwise be sent upstream to the realm "~.+$", which causes delays and is unneeded.

```
proxy server {
        default_fallback        = no
}

home_server antarctica-flr-1 {
        type                    = auth+acct
        ipaddr                  = 172.20.1.2
        port                    = 1812
        secret                  = secretstuff
        status_check            = status-server
}

home_server antarctica-flr-2 {
        type                    = auth+acct
        ipaddr                  = 172.25.9.3
        port                    = 1812
        secret                  = secretstuff
        status_check            = status-server
}

home_server_pool EDUROAM {
        type                    = fail-over
        home_server             = antarctica-flr-1
        home_server             = antarctica-flr-2
}

realm NULL {
    virtual_server  = auth-reject
    nostrip
}

realm "~.+$" {
        pool                    = EDUROAM
        nostrip
}
```

Goodies

Running FreeRADIUS as non-root user

The RADIUS protocol runs on ports >1023, which means it can be started entirely in unprivileged mode on UNIX-like systems. You can easily achieve that by

- creating a user "radiusd" and group "radiusd"
- giving all configuration files in /etc/raddb ownerships for that user radiusd + group radiusd
- changing these two parameters in /etc/raddb/radiusd.conf:

```
user  = radiusd
group = radiusd
```

F-Ticks

F-Ticks is using syslog to deliver user login statistics. You can enable syslog logging for login events by defining a *linelog* module. In the /etc/raddb /modules/ subdirectory, create a new file "f_ticks":

```
linelog f_ticks {
       filename = syslog
       #syslog_facility = local0
       #syslog_severity = info
       format = ""
       reference = "f_ticks.%{%{reply:Packet-Type}:-format}"
       f_ticks {
              Access-Accept = "F-TICKS/eduroam/1.0#REALM=%{Realm}#VISCOUNTRY=YOUR-TLD#VISINST=%{Operator-Name}
#CSI=%{Calling-Station-Id}#RESULT=OK#"
              Access-Reject = "F-TICKS/eduroam/1.0#REALM=%{Realm}#VISCOUNTRY=YOUR-TLD#VISINST=%{Operator-Name}
#CSI=%{Calling-Station-Id}#RESULT=FAIL#"
       }
}
```

Note that you have to adapt VISCOUNTRY to the country you are in (eg. set YOUR-TLD to "LU"), and VISINST to an identifier for your hotspot - which in this example is already set to the Operator-Name attribute. You can set the syslog facility and severity to help forward these ticks to the right place.

You need to enable this new module in the post-auth section of your virtual server eduroam:

```
post-auth {
                # if you want detailed logging
                reply_log
                f_ticks
                Post-Auth-Type REJECT {
                        # if you want detailed logging
                        reply_log
                        f_ticks
                }
        }
```

This way, appropriate loglines will be logged into your local syslog instance. If you want to forward your ticks to the statistics system, please get in touch with your NRO to get to know the syslog destination and configure your syslog daemon to forward the log line correspondingly.

Please note that the file proxy.conf may need your attention: FreeRADIUS' handling of the "DEFAULT" realm changed slightly between 2.1.9 and 2.1.10: previously, it would fill %{Realm} with the actual realm (e.g. "education.lu"), but after the change, it would use the literal "DEFAULT". It is not helpful to generate ticks with REALM=DEFAULT.

If you were using DEFAULT before, and now notice that ticks are sent incorrectly, the mitigation is to use a regular expression instead of DEFAULT - because for realm statements with regular expressions, also the most recent versions still substitute with the actual realm.

You would need to delete the DEFAULT realm and replace it with the following regular expression realm statement *at the end of your proxy.conf*:

```
realm "~.+$" {
...
}
```

CUI for eduroam SP

To use the Chargeable-User-Identity (CUI) you must already use the Operator-Name attribute.

This documentation is only for FreeRADIUS 3.0.X release.

Create a log module

By default the CUI is not logged, you have to use the FreeRADIUS *linelog* module to get a log. In the mods-available/ subdirectory, create a new file "eduroam_cui_log" :

```
linelog cui_log {
#    filename = syslog
    filename = ${logdir}/radius.log
    format = ""
    reference = "auth_log.%{%{reply:Packet-Type}:-format}"
    auth_log {
        Access-Accept = "%t : eduroam-auth#ORG=%{request:Realm}#USER=%{User-Name}#CSI=%{%{Calling-Station-Id}:-
Unknown Caller Id}#NAS=%{%{Called-Station-Id}:-Unknown Access Point}#CUI=%{%{reply:Chargeable-User-Identity}:-
Unknown}#MSG=%{%{EAP-Message}:-No EAP Message}#RESULT=OK#"
        Access-Reject = "%t : eduroam-auth#ORG=%{request:Realm}#USER=%{User-Name}#CSI=%{%{Calling-Station-Id}:-
Unknown Caller Id}#NAS=%{%{Called-Station-Id}:-Unknown Access Point}#CUI=%{%{reply:Chargeable-User-Identity}:-
Unknown}#MSG=%{%{reply:Reply-Message}:-No Failure Reason}#RESULT=FAIL#"
    }
}
```

Enable modules

```
cd mods-enabled; ln -s ../mods-available/eduroam_cui_log; ln -s ../mods-available/cui
```

Client definition

Force parameter 'add_cui' to 'yes' for all your connected clients :

```
client antarctica-access-point-1 {
...
    add_cui = yes
 }
```

Policy

Edit the default policy.d/cui file :

```
...
cui_hash_key = "changeme"         # --> replace with a random string
                                                        # if you use a secondary or backup FreeRADIUS server,
use the same cui_hash_key
                                                        # this allows you to keep the same CUI log even if the
FreeRADIUS server change
cui_require_operator_name = "yes"
...
```

Others values don't need to be changed.

Attributes

Edit mods-config/attr_filter/pre-proxy file, check that attributes Calling-Station-Id, Operator-Name and Chargeable-User-Identity are defined :

```
DEFAULT
...
        Calling-Station-Id =* ANY,
        Operator-Name =* ANY,
        Chargeable-User-Identity =* ANY,
...
```

Edit mods-config/attr_filter/post-proxy file, check that the attributes User-Name and Chargeable-User-Identity are defined :

```
DEFAULT
...
        User-Name =* ANY,
        Chargeable-User-Identity =* ANY,
...
```

CUI filtering

Edit policy.d/filter, add a filter function 'cui_filter'. Simple example :

```
# Filter the Chargeable-User-Identity attribute
cui_filter {
  if (&reply:Chargeable-User-Identity =~ /REPLACE-WITH-CUI-TO-MATCH/) {
                        update request {
                                &Module-Failure-Message += "Rejected: CUI matching '%{reply:Chargeable-User-
Identity}'"
                        }
                        reject
       }
}
```

Using policies and modules in your eduroam virtual server

Add 'cui' in authorize, post-auth and pre-proxy sections. Add 'cui_log' and 'cui_filter' in post-auth section :

```
server eduroam {
...
        authorize {
                # only use filter_username from version > 3.0.7 on
                filter_username
                update request {
                        Operator-Name := "1yourdomain.tld"
                        # the literal number "1" above is an important prefix! Do not change it!
                }
                cui
                                # if you want detailed logging
                auth_log
                suffix
        }
...
        post-auth {
                # if you want detailed logging
                reply_log
                                cui
                                cui_filter
                                cui_log
                Post-Auth-Type REJECT {
                        reply_log
                                                eduroam_log
                }
        }
...
        pre-proxy {
                pre_proxy_log
                cui
                if("%{Packet-Type}" != "Accounting-Request") {
                        attr_filter.pre-proxy
                }
        }
...
}
```

Caveats

Use the most recent version available (3.0.10 at the time of writing) because of known issues in older versions (ranging from filters that prevent people to get online with mixed usernames to TLS-related bugs).

eduroam IdP

Obligations for an eduroam IdP

Selecting EAP types

The decision which EAP type(s) to deploy on your eduroam IdP depends on several factors:

- Capabilities of your Identity management backend
- Types of devices you want to support

Choices depending on the Identity Management System

Regarding the identity management backend, the most fundamental differentiation between EAP types is the type of credential they support.

- Does your identity management backend support X.509 Client Certificates? Then you can use EAP-TLS.
- Does your identity management backend use username/password combinations?
    - Does it store the passwords as either clear text - or - encrypted as NT-Hash? Then you can use EAP-TTLS, PEAP, EAP-FAST, EAP-PWD and more.
    - Does it store the passwords in a different crypt format? Then you can use EAP-TTLS only.

As you see, the decision is largely dependent on your identity management system; so your choices may be limited. As a more concrete advice for some IdM backends:

- Microsoft ActiveDirectory: stores passwords as NT-Hashes.

Anonymous outer identities

Almost all EAP types support the use of anonymous outer identities. The primary use of anonymous outer identities is for better preservation of privacy for your users; a properly configured supplicant will then not even reveal the real username of the user to the visited eduroam SP; instead, the username is replaced with a dummy value.

This feature needs protocol support by the EAP type in question; the basic idea is that there have to be two stages of communicating the client identity:

- one identity, the outer identity, is used to route the user's login request from the eduroam SP via the eduroam RADIUS path to the eduroam IdP
- the second, "inner" identity, is only revealed inside a cryptographically protected tunnel to the IdP

Since the outer identity is only needed for routing purposes towards the IdP, the local username part does not have to be accurate and can be obfuscated. The IETF-suggested way of obfuscating the username is to leave it empty; but it can just as well be replaced with "anonymous", "anon" or similar. However, the realm part (i.e. behind the @ sign) always needs to be accurate because it contains the routing information.

The inner identity always needs to be fully accurate, because it is used to authenticate the user. It does not necessarily have to contain an @ sign at all, because that username is local to the IdP and is only seen and evaluated there.

Example:

- Outer identity: anonymous@restena.lu
- Inner identity: stefan.winter

For eduroam request routing, the part @restena.lu of the outer identity is used to route the request to the restena.lu realm and to establish a secure tunnel; while the real username inside this tunnel which is looked up in a user database is "stefan.winter".

Here is a break-down of anonymous outer identity support for some popular EAP types:

EAP-Type	Support for anonymous outer identites
EAP-TTLS	yes
PEAP	yes
EAP-FAST	yes
EAP-TLS	support in protocol, but not typically available in supplicants
EAP-PWD	no

If the EAP type allows for the use of outer identities, it is a client device configuration option to either make use of them or not; there is little you as an IdP can do to force the use of anonymous outer identities (except for providing and encouraging the use of pre-configured installers which will then make all the necessary settings on the client device automatically).

Choices depending on the envisaged devices

The landscape of wireless-enabled devices is rather heterogenous, and support for EAP types varies. Ideally, you should survey which types of devices you should come to expect among your user base, check the capabilities of these devices, and make an informed decision regarding the EAP type of choice.

However, the EAP protocol is flexible enough to handle multiple EAP types: if your IdM backend can support the use of multiple EAP types, then you can configure all the supported EAP types. In that case, you have to select a "default" EAP type - it should be set to the EAP type with the broadest support in your client base.

Now, assuming you have the option of configuring a range of EAP types *and* your clients support that same range, which of these types should you prefer?

- We suggest the use of PEAP over EAP-TTLS for it does a mild amount of protection of the user password inside the secure tunnel.
- If you cannot support PEAP, consider to allow TTLS-PAP **and** the more unusual variant TTLS-GTC (initially Generic Token Card; also used for passwords which are not savable on the client device). Some older devices (certain Symbian OS builds) support TTLS, but not PAP inside. Enabling TTLS-GTC will allow these devices to connect.

EAP Server certificate considerations

Almost all EAP types in eduroam (with the exception of EAP-PWD) require an X.509 server certificate with which the RADIUS server identifies itself to the end user before the user sends his credentials to the server.

Consideration 1: Procuring vs. creating your own server certificate

In a generic web server context, server certificates are usually required to be procured by a commercial Certification Authority (CA) operator; self-made certificates trigger an "Untrusted Certificate" warning. It makes sense for browsers to have a pre-configured trust store with many well-known CAs because the user may browse to any website; and the operator of that website may have chosen any of those well-known CAs for his website. In an abstract notion, one can say: it is required to have many CAs in the list because the user device does not have all required information for certificate validation contained in its own setup; it misses the information "which CA did the server I am browsing to use to certify the genuinity of his website?".

These considerations are not at all true in an EAP authentication context, such as an eduroam login. Here, the end user device is pre-provisioned with the entire set of information it needs to verify this specific TLS connection: the IdP has a certificate from exactly one CA, and needs to communicate both that CA and the name of his authentication server to the end user. A trust store list from the web browser is thus insignificant in this context; certificates from a commercial CA are as valid for EAP authentications as are self-made certificates or certificates from a small, special-purpose CA. For a commercial CA, the installation of the actual CA file may be superfluous in some client operating systems (particularly those who make their "web browser" trust store also accessible for EAP purposes), but marking that particular CA as trusted for this specific EAP authentication setup still needs to be done.

Note that also root CA certificates have an expiry date. Both for commercial and private CAs please be aware that an exchange of the root CA certificate will require re-configuration of all your end-users' devices to accept the new CA. As a consequence: for commercial CAs, check their root CA's expiry date so you can make an informed decision whether you want to buy the certificate from them or not. For your own private-use CA: choose a very long expiry date for the CA. Especially for commercial CAs, keep in mind that if you ever want to switch to a *different* CA as a trust anchor, all your end-user devices again need to be re-configured for that new root.

Configuration tools like eduroam CAT enable to provision the chosen CA(s) and the expected server name(s) into client devices without user interaction. In that light, it does not make much difference whether to procure a server certificate from a commercial CA or to make your own; either way, configuration steps are necessary on the end-user device to enable and secure your chosen setup. With the conceptual differences being small, a number of secondary factors come into play when making the decision where to get a server certificate from:

- Do you have the necessary expertise to create a self-signed certificate; or to set up a private Certification Authority and issue a server certificate with it? Consider in particular the next "Consideration 2" which imposes some properties onto the certificates you need.
- Does your eduroam NRO operate a special-purpose CA for eduroam purposes, so that you could get a professionally crafted certificate without much hassle?
- Do your end-user devices all verify the exact server identity (issuing CA certificate AND expected server name)?

The third question is particularly important these days because some popular operating systems, particularly early Android versions up until Android 7, do not allow to configure verification of the expected server name in their UI. For such operating systems, using a commercial CA for the server certificate opens up a loophole for fraud: anyone with a valid certificate from this CA, regardless of the name in the certificate, can pretend to be the eduroam authentication server for your end-user; which ultimately means the end-user device will send the user's login credentials to that unauthorised third-party. If you use a self-signed certificate or private CA however, which issues only one/very few certificates, and over which you have full control, then no unauthorised third party will be able to get a certificate in the first place, and thus can't fraud your users.

Another factor to consider when making the decision private vs. commercial CA is that of size and length of the EAP conversation during every login: with a private CA, you will be able to construct a certificate chain without intermediary CA certificates; requiring less bytes to be transmitted inside the EAP conversation (see Consideration 3, below). This results in fewer EAP round-trips and thus a faster authentication.

So, as a general recommendation: if you have the required expertise, it is suggested to set up a private CA exclusively for your IdP's eduroam service. This CA should have a very long lifetime to prevent certificate rollover problems. The CA should issue only server certificates for your eduroam IdP server (s). If you do not have that expertise, you should make use of your NROs special-purpose CA if it exists. If none of these work for you, a certificate from a commercial CA is the third option.

> ⚠️ *With great power comes great responsibility. (Voltaire)*

If you choose to use a private CA and deploy it to your users' devices, there may be side-effects after the installation of the CA. Some devices do not differentiate between a CA which is used for Wi-Fi server authentication purposes and, say, web browser TLS encryption.

Your CA may incidentally yield the power on such client devices of your own user base to inspect their web or other traffic (if you actively abuse it and modify your IT infrastructure to enable this). We do not endorse or encourage this in any way.

Having your own trusted root CA in client devices also makes the protection of the private key to this CA an objective of paramount importance.

We recommend that you inform your users how best to restrict the power of the CA (e.g. with CA installation instructions which point to a dedicate Wi-Fi store [Android 4.3+]; or with the advice not to use browsers which use the built-in CA store of the device [MS Windows]).

Consideration 2: Recommended certificate properties

Various end-user device operating systems impose different requirements on the contents of the server certificate that is being presented. Luckily, these requirements are not mutually exclusive. When creating or procuring a server certificate, you should check with the CA that its certificates satisfy as many of these requirements as possible to ensure broad compatibility with your users' devices. The list below does not include "standard" sanity checks applied to certificates; e.g. well-formedness of the data, validity timestamps etc. These checks are done "as per usual" in every TLS connection.

The most important property of the server certificate is the name of the server. Since this certificate is not for a webserver, there is no necessity to put an actual hostname into the server name. Also, when an Identity Provider uses multiple servers for resilience reasons, then all these servers can and should have a certificate with the same name; and it may well be the identical certificate. Having different names for different servers means that end-user devices must be configured to trust multiple servers, which is more cumbersome than just having to configure one name string.

Some end-user device operating systems might (incorrectly) require the name to be parseable as a hostname; so it is a good idea to use a server name which parses as a *fully-qualified domain name* - the corresponding record does not have to exist in DNS though. The server name should then be both in certificate's Subject field (*Common Name* component) and be a *subjectAltName:DNS* as well.

The following additional certificate properties are non-standard and are of particular interest in the eduroam context:

Property	Content	Remarks
X.509 version	3	The CA certificate should be an X.509v3 certificate.
server name	parses as fully-qualified domain name	Server certificates with spaces, e.g. "RADIUS Service of Foo University" are known to be problematic with some supplicants, one example being **Apple iOS 6.x**.
server name	Subject /CN == SubjectAltName: DNS	Some supplicants only consult the CN when checking the name of an incoming server certificate (**Windows 8 with PEAP**); some check either of the two; some new EAP types such as **TEAP**, and **Linux clients configured by CAT 1.1.2+** will only check SubjectAltName:DNS. Keeping the desired name in both fields in sync is a safe bet for futureproofness. Only use one CN. If you have multiple RADIUS servers, either use the same certificate for all of them (there is no need for the name to match the DNS name of the machine it is running on), or generate multiple certificates, each with one CN /subjectAltName:DNS pair.
server name	not a wildcard name (e.g "*.someidp.tld")	Some supplicants exhibit undefined/buggy behaviour when attempting to parse incoming certificates with a wildcard. **Windows 8 and 8.1** are known to choke on wildcard certificates.
signature algorithm	Minimum: SHA-256 Recommended: SHA-256 or higher	Server certificates signed with the signature algorithm MD5 are considered invalid by many modern operating systems, e.g. ***Apple iOS 6.x and above***. Also **Windows 8.1** and all previous versions of **Windows (8, 7, Vista)** which are on current patch levels will not validate such certificates. Having a server certificate (or an intermediate CA certificate) with MD5 signature will create problems on these operating systems. Apparently, no operating system as of yet has an issue with the root CA being self-signed with MD5. This may change at any point in the future though, so when creating a new CA infrastructure, be sure not to use MD5 as signature algorithm anywhere. The continued use of SHA-1 as a signature algorithm is not recommended, because several operating systems and browser vendors already have a deprecation policy for SHA-1 support. While the deprecation in browser-based scenarios does not have an immediate impact on EAP server usage, it is possible that system libraries and operating system APIs will over time penalise the use of SHA-1. Therefore, for new certificates, SHA-256 is recommended to avoid problems with the certificate in the future.
length of public key	Minimum: 2048 Bit Recommended: 3072 Bit or higher	Server certificates with a length of the public key below 1024 bit are considered invalid by some recent operating systems, e.g. ***Windows 7 and above***. Having a server certificate (or an intermediate CA certificate) with a too small public key will create problems on these operating systems. The continued use of 1024 bit length keys is not recommended, because several operating systems and browser vendors already have a deprecation policy for this key length. While the deprecation in browser-based scenarios does not have an immediate impact on EAP server usage, it is possible that system libraries and operating system APIs will over time penalise the use of short key lengths. 2048 bit is the most popular and default choice these days. However, some applications already suggest 3072 bit or more, and a longer key length does not have an extra cost. So, it is recommended to create new certificates with 3072 bit keys or higher (4096 has been tested and is also unproblematic) to avoid problems with the certificate in the future.
Extension: Extended Key Usage	TLS Web Server Authentication	Even though the certificate is used for EAP purposes, some popular operating systems (i.e. ***Windows XP and above***) require the certificate extension "TLS Web Server Authentication" (OID: 1.3.6.1.5.5.7.3.1) to be present. Having a server certificate without this extension will create problems on these operating systems.

Extension: CRL Distribution Point	HTTP /HTTPS URI pointing to a valid CRL	Few very recent operating systems require this extension to be present; otherwise, the certificate is considered invalid. Currently, ***Windows Phone 8*** is known to require this extension; ***Windows 8*** can be configured to require it. These operating systems appear to only require the extension to be present; they don't actually seem to download the CRL from that URL and check the certificate against it. One might be tempted to fill the certificate extension with a random garbage (or intranet-only) URL which does not actually yield a CRL; however this would make the certificate invalid for all operating systems which do evaluate the extension if present. So the URL should be a valid one.
Extension: BasicConstraint (critical)	CA: FALSE	Server certificates need to be marked as not being a CA. Omitting the BasicConstraint:CA totally is known to cause certificate validation to fail with **Mac OS X 10.8 (Mountain Lion)**; setting it to TRUE is a security issue in itself. Always set the BasicConstraint "CA" to false, and mark the extension as critical.
Certificate Type	Domain-Validated (DV) or Organisation-Validated (OV)	There have been several reports that **ChromeOS** will refuse to accept Extended Validation (EV) certificates. You should avoid these types of certificates if you care about this operating system.
Validity Time	825 days or fewer	**Apple** products as of **macOS 10.15+ and iOS 13+** enforce this limit and consider certificates with a longer lifetime as untrusted. See also this Apple article.

Consideration 3: Which certificates to send in the EAP exchange

End-user devices need to verify the server certificate. They do this by having a known set of trustworthy anchors, the "Trusted Root Certificates". These root certificates need to be available and activated on the device prior to starting the eduroam login. Therefore, it does not serve any useful purpose to send the root CA certificate itself inside the RADIUS/EAP conversation. It is not harmful to send it anyway though, except that it unnecessarily inflates the data exchange, which means more round-trips during eduroam authentication, and in turn a slower login experience. One possible exception is: there are reports of certain Blackberry devices for which it is advantqageous to send the root CA certificate nontheless; if you expect you need/want to support Blackberry devices, sending the root CA may be of help.

During the EAP conversation, the eduroam IdP RADIUS server always needs to send its server certificate.

One question needs an administrative decision: if there is one or more intermediate CAs between the root CA and the server certificate (such as is the case with, for example, the TERENA Certificate Service (TCS) and many commercial CAs), should the intermediate CA certificates be sent to the end user device during the EAP conversation, or should the devices pre-install the intermediate CAs along with the root certificate?

In any case, for a successful verification of the server certificate, the end-user's device must have the full set of CA certificates available. It does not matter whether the intermediate CAs have been pre-provisioned or are sent during the login phase; but if any one intermediate CA is missing, the verification of the server certificate will fail.

Pre-provisioning the intermediate CAs has the advantage of a relatively small amount of data being sent during the EAP authentication, which means fewer round-trips between the end-user's device and the eduroam IdP RADIUS server. The downsides of this approach are that any changes to intermediate CAs (re-issue, rollover) will also need to be pushed to end-user devices. Also, if end-user devices are not under administrative control of the IdP, there is a danger that some end users do not follow the advice to install all intermediate CAs even though they should, and end up in a situation where the server certificate can not be validated.

Sending the intermediate CAs during the login phase means a longer time to authenticate due to more round-trips, but means that it is sufficient for client devices to install the root CA certificate; if intermediate CAs change, the new ones will always become available to the device during the next authentication data exchange.

For most deployments, it probably makes more sense to include the intermediate CA certificates during the RADIUS/EAP conversation.

Set up of several popular RADIUS servers

FreeRADIUS

Setting up FreeRADIUS

This section describes how to set up FreeRADIUS for an IdP. It assumes that you have already executed the configuration steps for the eduroam SP configuration of FreeRADIUS. We will expand that configuration to turn FreeRADIUS into a simple IdP. N.B.: even if you are going to have an IdP-only installation, the eduroam SP configuration for FreeRADIUS is still the exact same. You just don't define any own Access Point clients in clients.conf.

Adding IdP support in FreeRADIUS needs several steps to be executed:

- a TLS server certificate needs to be created for EAP methods to work
- the desired EAP types need to be configured.
- the virtual server eduroam needs to be instructed to do tunneled EAP authentication
- a user database needs to be linked to the FreeRADIUS instance to authenticate the users
- a realm needs to be marked as to-be-authenticated-locally in the configuration
- the server needs to be prepared to process incoming requests *from* the upstream FLR server

These steps are explained in detail below. For the user database, this example will use a "flat file" with usernames and passwords. The Goodies section contains examples for MySQL and other types of backend databases.

TLS server certificate

While it is possible to buy and install a commercial TLS certificate, this is neither necessary (the trust settings of web-browser stores don't apply for EAP, so there are no "recognised" CAs) nor prudent (a commercial CA issues many certificates, and uncautious users might be tempted to accept other certificates from that same CA).

We suggest to create an own certificate. FreeRADIUS makes this very easy by providing an automatic script for that purpose. Execute the

```
/etc/raddb/certs/bootstrap
```

script. It will generate certificates which are suited for EAP authentication, and name them so that the server can find them immediately without further configuration. Later, for the supplicant configuration, you will need to include the generated CA certificate into your supplicant configurations.

EAP type configuration

The file /etc/raddb/eap.conf defines how EAP authentication is to be executed. The shipped configuration file is not adequate for eduroam use; it enabled EAP-MD5 and LEAP, for example; which are not suitable as eduroam EAP types. Use the following content for eap.conf instead. It enables PEAP and TTLS:

```
eap {
                default_eap_type = peap
                timer_expire     = 60
                ignore_unknown_eap_types = no
                cisco_accounting_username_bug = no

                tls {
                        certdir = ${confdir}/certs
                        cadir = ${confdir}/certs
                        private_key_password = whatever
                        private_key_file = ${certdir}/server.key
                        certificate_file = ${certdir}/server.pem
                        CA_file = ${cadir}/ca.pem
                        dh_file = ${certdir}/dh
                        random_file = /dev/urandom
                        fragment_size = 1024
                        include_length = yes
                        check_crl = no
                        cipher_list = "DEFAULT"
                }

                ttls {
                        default_eap_type = mschapv2
                        copy_request_to_tunnel = yes
                        use_tunneled_reply = yes
                        virtual_server = "eduroam-inner-tunnel"
                }

                peap {
                        default_eap_type = mschapv2
                        copy_request_to_tunnel = yes
                        use_tunneled_reply = yes
                        virtual_server = "eduroam-inner-tunnel"
                }

                mschapv2 {
                }

        }
```

A common question regarding this definition is: "why is TLS also configured? I don't want it, can I disable it?" The answer is: the TTLS and PEAP sections depend on the tls stanza for the definition of which server certificates to use. You cannot delete the stanza, but that doesn't mean you can't effectively disable TLS: the tls stanza contains the ca_file parameter. Only clients with a TLS client certificate from this CA will be accepted. We have just created a brand-new CA with the "bootstrap" script. Simply don't issue nor distribute any client certificates from this CA, then nobody will be able to log in with EAP-TLS. Note that in newer versions of FreeRADIUS (>3.0.14) there is a new tls-config section that allows you to configure the common TLS configuration without configuring the TLS EAP type. The config above is backwards compatable, but if you want to take advantage of the new section you can replace the name of the "tls" block above with "tls-config tls-common" and then reference it from each EAP type with "tls = tls-common" (the example eap config shows you how to do this).

Another question is regarding the mschapv2 section. For all practical purposes, the easy answer is that it is a piece of magic and needs to be there for PEAP to work. If you are curious regarding the gory details, please let us know.

Note that one parameter for both the ttls and peap stanza is "virtual_server = eduroam-inner-tunnel". This means that the inner EAP authentication will be carried out in this other virtual server, which we will define later.

Virtual server eduroam: enable EAP, make Operator-Name conditional

Compared to the eduroam SP config, you need to additionally mention the "eap" module in both the authorize and authenticate stanza of the file /etc/raddb/sites-enabled/eduroam so that your server can process EAP requests from your own userbase.

You should also make sure to only tag those incoming requests with the Operator-Name attribute which actually originate from your own WiFi gear - as an IdP, your own users roaming elsewhere will also be processed, but they should not carry your own Operator-Name. For the purposes of this wiki, let's assume that you are connected to one FLR server, and it is defined in your clients.conf with the shortname "antarctica-flr-1" (see below for the exact definition).

It will then look like the following:

```
authorize {
                if ("%{client:shortname}" != "antarctica-flr-1") {
                   update request {
                           Operator-Name := "1yourdomain.tld"
                            # the literal number "1" above is an important prefix! Do not change it!
                   }
                }
                auth_log
                suffix
                eap
        }

authenticate {
       eap
}
```

Virtual server eduroam-inner-tunnel

When the eap module has started with an authentication, it will first establish a TLS tunnel; this is done by enabling the module in the previous "eduroam" virtual server. After the TLS tunnel is established, the content (i.e. the tunneled authentication) is processed separately in this new virtual server. Create the file in /etc/raddb/sites-enabled/eduroam-inner-tunnel and give it the following content:

```
server eduroam-inner-tunnel {

authorize {
        auth_log
        eap
        files
        mschap
        pap
}

authenticate {
        Auth-Type PAP {
                pap
        }
        Auth-Type MS-CHAP {
                mschap
        }
        eap
}

post-auth {
        reply_log
        Post-Auth-Type REJECT {
                reply_log
        }

}
}
```

Let's revisit the modules which this virtual server executes one after another:

- auth_log: logs the incoming packet to the file system. This is needed to fulfill the eduroam SP logging requirements. Note that this log *may* contain the user's cleartext password if TTLS-PAP is used. You can log the packet with omitted User-Password attribute if you prefer; see the "Goodies" section for more details).
- eap: if the EAP authentication contains another EAP instance inside, the module will decode it. This is the case for PEAP.
- files: this module tries to find out the authoritative password for the user by looking up the username in the file
- mschap: this module is in effect only if PEAP-MSCHAPv2 or TTLS-MSCHAPv2 is used. It will mark the packet as to be authenticated with MS-CHAP algorithms later.
- pap: this module is in effect only if TTLS-PAP is used. It will mark the packet as to be authenticated with PAP alogrithms later.
- reply_log: logs the reply packet to the file system

User database: flat file

By default, the "files" module will use information in the file

```
/etc/raddb/users
```

for authenticating users. This file has a straightforward format

```
icecold@group1.aq       Cleartext-Password := "snowwhite"

otheruser@group1.aq     Cleartext-Password := "swordfish"
```

Local authentication for your realm

In the SP configuration, all requests were unconditionally forwarded to upstream. We will need to revisit the file "proxy.conf" and mark one realm to NOT proxy. In this example, we will use "@group1.aq" as the local authentication realm. Simply add the following stanza immediately preceeding the "DEFAULT" realm:

```
realm group1.aq {
        nostrip
}
```

Since the stanza doesn't contain a server pool to proxy to, this realm won't be proxied and instead authenticated locally. This stanza works only for users who correctly use the full username format "user123@group1.aq" for their eduroam login.

If the IdP and SP are colocated, it is possible to *locally* also accept users who erronuously omitted their realm (just "user123"). This is NOT permitted by the eduroam policy (read 6.3.2 bullet 6 under AAA Servers of the current service definition document: "The outer EAP identities (and with it, RADIUS User-Name attributes) for the IdP MUST be in the format of arbitrary@realm"). Allowing this also requires further configuration and it is strongly discouraged, because it will give such users a "halfways-working" experience: they will be able to use eduroam when on their own IdP's campus, because no routing information needs to be evaluated, but their account will fail at all other locations. Therefore, this guide does not include instructions for that kind of setup.

Processing incoming requests

As an eduroam IdP, your users can go to other eduroam hotspots around the globe. They will still be authenticated at your server. In these roaming cases, your upstream FLR servers will send Access-Requests to your server. Surprisingly, it is very simple to configure that: these upstream servers are simply clients - just like an Access Point. So, simply add client stanzas for your FLR servers into clients.conf:

```
 client antarctica-flr-1 {
        ipaddr                          = 172.20.1.2
        netmask                         = 32
        secret                          = secretstuff
        require_message_authenticator   = yes
        shortname                       = antarctica-flr-1
        nastype                         = other
        virtual_server                  = eduroam
}
```

CUI for eduroam IdP

To use the Chargeable-User-Identity (CUI) you must already use the Operator-Name attribute. This documentation is only for FreeRADIUS 3.0.X release.

Modify the log module

Edit "eduroam_cui_log" file in the mods-available/ subdirectory and add the following lines to your virtual inner server :

```
...
linelog cui_inner_log {
#    filename = syslog
    filename = ${logdir}/radius.log
    format = ""
    reference = "inner_auth_log.%{%{reply:Packet-Type}:-format}"
    inner_auth_log {
        Access-Accept = "%t : eduroam-inner-auth#VISINST=%{request:Operator-Name}#USER=%{User-Name}#CSI=%{%
{Calling-Station-Id}:-Unknown Caller Id}#NAS=%{%{Called-Station-Id}:-Unknown Access Point}#CUI=%{%{%{reply:
Chargeable-User-Identity}:-%{outer.reply:Chargeable-User-Identity}}:-Local User}#RESULT=OK#"
        Access-Reject = "%t : eduroam-inner-auth#VISINST=%{request:Operator-Name}#USER=%{User-Name}#CSI=%{%
{Calling-Station-Id}:-Unknown Caller Id}#NAS=%{%{Called-Station-Id}:-Unknown Access Point}#CUI=%{%{%{reply:
Chargeable-User-Identity}:-%{outer.reply:Chargeable-User-Identity}}:-Local User}#RESULT=FAIL#"
    }
}
```

Use policy and module in your eduroam-inner-tunnel virtual server

Add 'cui-inner' (policy already defined, you don't need to change it) and 'cui_inner_log' in post-auth section :

```
server eduroam-inner-tunnel {
...
                post-auth {
                                reply_log
                                cui_inner_log
                                cui-inner
                                Post-Auth-Type REJECT {
                                        reply_log
                                        cui_inner_log
                                }
                }
...
}
```

That's it! Now your server is prepared for eduroam IdP operation! You can add users to your "database" by amending the "users" file; if you do, you will unfortunately have to restart FreeRADIUS so that it picks up the change.

Goodies

Omitting User-Password in inner authentication logs

By default, the "detail" modules log every attribute as it was received. For inner authentications with TTLS-PAP, this means that the attribute "User-Password" with the user's perceived password will be logged. This is often considered harmful. You can deactivate it by blacklisting the attribute in the auth_log module in /etc/raddb/modules/auth_log:

```
detail auth_detail {
  ...
  suppress {
    User-Password
  }}
```

adding VLAN assignment attributes

You can mark every user with a VLAN where he should be put into. This is done by assigning "reply items" to the user in the authentication database. In our flat file example, reply attributes are in a separate line, indented by a tab. To put our two example users into VLANs 17 and 42, respectively, the entries would look like the following:

```
icecold@group1.aq       Cleartext-Password := "snowwhite"
                        Tunnel-Type             := VLAN,
                        Tunnel-Medium-Type      := IEEE-802,
                        Tunnel-Private-Group-ID := 17

otheruser@group1.aq     Cleartext-Password := "swordfish"
                        Tunnel-Type             := VLAN,
                        Tunnel-Medium-Type      := IEEE-802,
                        Tunnel-Private-Group-ID := 42
```

Using SQL as user database backend

Using a flat file as in our example scales very poorly. You can use arbitrary database backends instead; the FreeRADIUS documentation can give you an overview. If you wish to use SQL, changing our example configuration is very easy: simply replace the "files" line in eduroam-inner-tunnel:authorize with "sql". You'll need to specify the connection details for your SQL backend in the corresponding module (/etc/raddb/modules/sql).

The schema which FreeRADIUS uses to store user information is similarly structured to the "users" file: a table radcheck holds the check items (i.e. the username and password), and the radreply table contains the reply items (for example VLAN memberships, as explained above).

Mandating or forbidding use of anonymous outer identity

eduroam at large supports anonymous outer identities for user logins. It is at the discretion of eduroam IdPs whether they want to

- mandate that their users use an anonymous outer identity
- forbid their users to use an anonymous outer identity
- be agnostic in that respect

Configuring any one of the three choices is done with only a few lines of configuration. The easiest choice is being agnostic: no configuration is necessary, since there is no link between the inner and outer User-Name attribute in FreeRADIUS.

If you want to mandate the use of anonymous outer identities, the recommended way is using the identity "@realm" (i.e. the part left of the @ sign should be empty). You can enforce that only this outer User-Name is allowed to proceed to EAP authentication by adding the following to the authenticate section:

```
if ( User-Name != "@realm" ) {
      fail
}
```

If you want to forbid usage of anonymous outer identities, you can do this by comparing the two presented User-Name attributes of the outer and inner authentication. You can only do this in the eduroam-inner-tunnel virtual server obviously, since only that server has access to the inner identity. Put the following into the "authenticate" section of eduroam-inner-tunnel:

```
if ( User-Name != outer.User-Name ) {
     fail
}
```

More information

Eduroam-in-a-box web configuration tool:http://eduroam.sourceforge.net

Radiator

One popular RADIUS server is Open Systems Consultant's "Radiator. This section details its configuration.

Because of the EAP authentication within RADIUS, a (small) PKI is required. If there is no PKI available, you could create the required key and certificate with, for instance, TinyCA. TinyCA (http://tinyca.sm-zone.net/) is a simple graphical interface on top of OpenSSL. It is possible to use OpenSSL directly (but instructions to do so are outside the scope of this document).

There is also a bootable CD available based on Knoppix that runs TinyCA, the roCA (read-only CA) that can be found at http://www.intrusion-lab.net/roca/.

Depending on the EAP-type used, client certificates may also be needed.

Within the Radiator distribution there are also simple scripts available to create certificates for testing purposes.

The Radiator RADIUS server needs the configuration file /etc/radiator/radius.cfg.

This configuration file can be created with the editor of choice, for example

vi /etc/radiator/radius.cfg
or
pico /etc/radiator/radius.cfg

In the following examples there are two kinds of EAP that are configured at "institution":

- EAP-TLS based on client-certificates.
- EAP-TTLS and EAP-PEAP that do not require client certificates but use the traditional mechanism of username/password authentication instead.

Clients

RADIUS is based on a client-server model. The NAS-devices (Access Points, switches etc.) forward credentials to a RADIUS server, i.e. act as a client, and therefore need to be defined on the RADIUS server. Other RADIUS servers can act as a client as well, so every kind of RADIUS-request can be forwarded to another server.

The clients are configured within Radiator using the <Client>-clause:

```
<Client 192.168.10.200>
            Secret 6.6obaFkm&RNs666
            Identifier ACCESSPOINT1
            IdenticalClients 192.168.10.201
                        RequireMessageAuthenticator
</Client>
```

In this example there is a client definition for 192.168.10.200, an Access-Point. The "secret" is a series of (at best 16) characters that are used to encrypt the credentials sent in the RADIUS-request.

It is of course recommended to create a secret that cannot be guessed easily, otherwise the RADIUS-message can be decrypted. This is not an issue with EAP-authentication using 802.1X, since the credentials are also transmitted over a SSL-encrypted tunnel between the client and the final authentication server. However, with regular credentials (like those used with Web-based redirection authentication) this is sensitive information that might be captured, therefore a reasonably complex secret and an SSL tunnel is recommended.

The Identifier in the Client-definition can be used later on in the Radiator configuration to filter a specific request.

If more then one Client is to use this same secret and identifier definition, the IdenticalClients statement can be used. If there are many clients with different IP-addresses, there is also the possibility for a "catch-all" client. This is the default client that is used after all other client definitions didn't match. Define this client as:

```
<Client DEFAULT>
```

If this kind of configuration is used, it is worth filtering with firewall-rules on RADIUS packets. There are only a few places where a RADIUS-request should come from; the management VLAN with the NAS-devices (switches and access-points), and the RADIUS infrastructure where unknown requests can be sent to.

Realms and VLAN assignment

The processing of authentication and accounting requests is done by linear processing of the present <Realm>- or <Handler>-clauses in the Radiator configuration file. Handler-clauses are more potent than Realm clauses in terms of filtering anything besides realms, and are therefore the preferred method. A realm is the part behind a username to indicate the origin of a user. With RADIUS, the username is usually separated from the realm with a "@" so the complete username looks like a regular e-mail address.

A <Handler>-clause is terminated with a </Handler>.

Within a Handler many mechanisms can be configured that define what to do with the RADIUS request.

PROXY example

The simplest Handler for proxying the request to another server uses the "AuthBy RADIUS" definition within this Handler.

In this example a proxy-configuration is shown. First we have a Handler that matches on any request, as long as it does not come from the client with the identifier "Proxy-Identifier". This is to prevent a proxy loop. When a request comes from a proxy-server, it should never be forwarded back to that proxy-server.

Another important part is the hostname to which the request should be forwarded. Multiple hostnames can be defined here for redundancy reasons: if the first host does not respond within three seconds, the second one is tried instead. The UDP ports to which the RADIUS-request should be forwarded can be defined in this "AuthBy RADIUS" clause as well.

```
<Handler Client-Identifier=/^(?!Proxy-Identifier$)/>
         <AuthBy RADIUS>
                    Host            192.87.36.3
                    Secret         super_secret!
                    AuthPort     1812
                    AcctPort     1813
                    StripFromReply Tunnel-Type, Tunnel-Medium-Type, Tunnel-Private-Group-ID
                    AddToReply Tunnel-Type=1:VLAN, Tunnel-Medium-Type=1:Ether_802, Tunnel-Private-Group-ID=1:909
          </AuthBy>
</Handler>
```

For a "Host", both the IP-address and FQDN can be used. The choice is more or less a personal preference of the RADIUS administrator, but be aware that the hostnames are only looked up once at the Radiator (re)start. If the lookup fails, the Host cannot be used until the next restart. This can represent a problem at a power outage, where for instance the DNS server is not yet available even though Radiator is.

While by using hostnames one benefits from the administrative ease when an IP-address is changed, it is still necessary to restart the RADIUS server.

The last part in this <AuthBy RADIUS>-definition shows the addition of RADIUS-attributes to the RADIUS- response. These attributes can be used to define a VLAN that will be assigned to users that are authenticated using this Handler. With StripFromReply, the attributes that came from the proxy-server are stripped first to prevent malicious VLAN-assignments, afterwards the attributes are added with the proper values for the local network design. In this case, VLAN 909 is used for guests.

Secure authentication with EAP-TLS

EAP-TLS requires both server and client certificates. Rolling out such certificates is a sometimes daunting administrative process, and is out of the scope of this document. The remainder of this section assumes that client certificates have been issued to the users already.

In this example the AuthBy-definition is outside the Handler, and is referred to using the Identifier. (This is useful if the AuthBy-definition is reused in another Handler, for instance.)

```
<AuthBy  FILE>
             Identifier ID4-TLS
             Filename %D/TLS-users
             EAPType TLS
             EAPTLS_CAFile %D/cert/institution-ca-chain.pem
             EAPTLS_CertificateFile %D/cert/radius-server-cert.pem
             EAPTLS_CertificateType PEM
             EAPTLS_PrivateKeyFile %D/cert/radius-server-key.pem
             EAPTLS_PrivateKeyPassword (the secret that secures the private-key)
             EAPTLS_MaxFragmentSize 1024
             AutoMPPEKeys
             SSLeayTrace 1
             StripFromReply Tunnel-Type,Tunnel-Medium-Type,Tunnel-Private-Group-ID
             AddToReply Tunnel-Type=1:VLAN,Tunnel-Medium-Type=1:Ether_802,Tunnel- Private-Group-ID=1:909,User-
Name=%u
</AuthBy>
```

In this AuthBy-clause there is an EAPTLS file defined that lists every employee. In this way, the users that can access the infrastructure using EAP-TLS are controlled.

The definitions that follow determine what to do with the EAP-request. First the "EAPType TLS" limits the use of this AuthBy-definition for TLS-only. Here regular password authentication is not desired, just certificates. Next, the certificate files are configured and the secret that secures the private-key file can be provided. If there is no secret for the private key, this can be omitted.

The next part defines in what size blocks the EAP-messages should be fragmented. Since some parts of the EAP-TLS challenge are too big to fit in a RADIUS request, the packets should be fragmented.

The MPPE-keys (Microsoft Point to Point Encryption, the protocol for encrypting the data across links) portion is important for wireless access. With 802.1 X, encryption occurs at the edge of the network, between the Access- Point and the client. To provide this secure encryption, a unique key is created and encrypted using the MPPE- keys that are derived from the SSL-challenge. This can be done at the Access-Point and the Client end so that there is no need to transfer the WEP-key in plain text over the air. This, and the fact that the key can be rotated within a period defined by either the Access-Point or the RADIUS server, provides 802.1x users with a good level of security.

The last part of the AuthBy-definition shows how to assign a proper VLAN.

```
<Handler Realm=instituion.cc, EAP-Message=/.+/>
             AuthBy   ID4-TLS
</Handler>
```

The Handler above shows the referral to the AuthBy-definition and some filtering mechanisms to filter out the proper requests. If more things need to be filtered on, they can be added to this handler, as follows:

```
NAS-Port-Type=/^(Wireless-IEEE-802-11|Ethernet)$/
```

In this way, only requests with the proper NAS-Port-Types are allowed. For Accounting purposes, a new handler should be defined in this case, that filters on:

```
"Request-Type=Accounting-Request"
```

since the request does not match the Handler that filters on the EAP-Message.

EAP-TTLS or EAP-PEAP

When issuing end user certificates is not an option, the EAP-mechanisms PEAP and TTLS can be used.

These two mechanisms look the same in that they both set up a TLS tunnel on which the credentials can be transported. They vary in the supported password encryption schemes.

Virtually all implementations of PEAP encrypt the user's password as an NT hash exclusively. TTLS implementations typically offer plain text transport of the password, called TTLS-PAP (the outer TLS tunnels makes sure the password cannot be eavesdropped) and sometimes other encryption schemes like MS- CHAPv2.

Administratively, the choice whether to use PEAP or TTLS can be challenging, since TTLS is not supported out of the box in the Microsoft Windows environment, therefore third party supplicant needs to be installed by users in case of a TTLS deployment.

Technically, three backend cases need to be considered for deployment:

Backend stores passwords in...	PEAP-MSCHAPv2?	TTLS?
plain text or reversibly encrypted	Yes	Yes (TTLS-PAP, TTLS-MSCHAPv2)
NT-Hash	Yes	Yes (TTLS-PAP, TTLS-MSCHAPv2)
other irreversible encryption	No	Yes (TTLS-PAP)

Where both options are possible, we suggest the following order of preference: TTLS-MSCHAPv2, PEAP- MSCHAPv2, TTLS-PAP (in descending order of preference).

Instead of a flat file, a more flexible backend for user accounts is a database like MySQL, or LDAP.

```
<Handler TunnelledByPEAP=1, Realm=tunnelled.institution.cc>
         <AuthBy FILE>
                   Filename %D/peap-users
                   EAPType MSCHAP-V2
         </AuthBy>
</Handler>

<Handler TunnelledByTTLS=1, Realm=tunneled.institution.cc>
         <AuthBy FILE>
                   Filename %D/ttls-users
         </AuthBy>
</Handler>
```

In these Handlers, the filtering options "TunneledByPEAP" and "TunnelledByTTLS" define that the tunnelled authentication (with the username and password in it) is handled here.

The "outer authentication", where the SSL tunnel is set up, looks like the TLS handler.

```
<Handler Realm=group_1>
          <AuthBy FILE>
                    Filename %D/users
                    EAPType TTLS, PEAP
                    EAPTLS_CAFile %D/root.pem
                    EAPTLS_CertificateFile %D/server.pem
                    EAPTLS_CertificateType PEM
                    EAPTLS_PrivateKeyFile %D/server.pem
                    EAPTLS_PrivateKeyPassword serverkey
                    EAPTLS_MaxFragmentSize 1024
                    EAPAnonymous anonymous@group1
                    AutoMPPEKeys
           </AuthBy>
</Handler>
```

Microsoft NPS

The GEANT Campus Best Practice (CBP) activity has created a very valuable configuration instruction document "Using Windows NPS as RADIUS in eduroam" which was available from the now defunct CBP website.

Cisco ACS

This section begs for community input.

ClearPass

This section begs for community input.

RADIUS Accounting in eduroam

eduroam, as a not-for-profit roaming consortium, is in the comfortable position of not having to rely on accurate RADIUS Accounting data.

This is quite fortunate because RADIUS Accounting, even though specified in RFC2866, is very underspecified and has many interoperability issues - especially in a world-wide roaming environment with numerous vendors and firmware versions of WiFi access points, which may all report with RADIUS accounting in a slightly different way. Since it provides no tangible benefit to eduroam operations, RADIUS Accounting is generally frowned upon. For many countries, particularly for European countries, eduroam Operations expects National Roaming Operators to acknowledge and discard Accounting packets destined to an eduroam IdP outside their own national federation; i.e. to confine Accounting, if it is really desired, to their own country.

Some countries do use Accounting information inside their country for various reasons, and eduroam Operations does not interfere with that.

For eduroam SPs, the consequence is that they can safely turn off RADIUS Accounting unless their National Roaming Operator insists on accounting records being sent.

As an eduroam IdP, the consequence is that you may receive accounting records from your own users when they roam to a different hotspot for which RADIUS Accounting has been turned on. Note that the number and content of attributes in the Accounting packets varies greatly due to the underspecification in RFC2866; you can not rely on any single Accounting attribute being present. It is possible, and probably the best option, to simply discard Accounting packets which cannot be correctly understood by your RADIUS server.

Provisioning configuration details of supplicants to end users

Once the chosen EAP method is configured and the IdP RADIUS server is connected to the authentication backend, the next step is to provision the access configuration to the actual end users.

Many operating systems support IEEE 802.1X and EAP authentication, but the user interfaces in supplicants differ significantly. For some supplicants, manually clicking through a series of GUI pages is the only option. This is sometimes tedious for end users.

If possible, an IdP administrator should prepare pre-configured packages which contain the necessary information to securely connect to eduroam:

- the SSID: "eduroam"
- the crypto setting: WPA2/AES
- the EAP type setting
- the CA that issued the eduroam IdP server's EAP server certificate
- the Common Name in the eduroam IdP server's EAP server certificate

There are tools that can be used to create such auto-installers. The use of one these windows 10 drivers update is recommended, because it will likely have a positive effect on user uptake, and reduce helpdesk load.

eduroam CAT

eduroam CAT has been created with the sole purpose to ease eduroam installation in many different client platforms through the use of auto-installers. The IdP administrator enters the information listed in the bullets above, after which installers are created for all kinds of platforms for the end users of the IdP. Please see the documentation; or visit the production website at https://cat.eduroam.org.

Others

In addition to eduroam CAT, there are other tools as well, e.g. su1x and XpressConnect (Cloudpath).

RADIUS/TLS: Using radsecproxy as an add-on to an existing RADIUS server

Goals and Prerequisites

This section describes a trimmed-down installation of radsecproxy which enables it to augment any RADIUS server with RADIUS/TLS transport, in order to enhance security. More precisely, with this setup radsecproxy will:

- Accept requests from a RADIUS server running on the same host
- Unconditionally forward the requests via RADIUS/TLS to an upstream server (typically a federation-level or a regional proxy server)
- Accept requests via RADIUS/TLS from the upstream server
- Unconditionally forward these requests to a RADIUS server running on the same host

The prerequisites for this deployment are:

- radsecproxy version 1.4.2 or higher
- The local RADIUS server's DEFAULT realm is configured to forward requests to the radsecproxy port on localhost.
- The local RADIUS server has configured localhost as a client (this is typically the case).
- The deployment requires a server certificate and a private key for that certificate to establish the RadSec connection which designates the server as an eduroam SP and IdP. For further information regarding eduroam certificates see this section.

Installation

On UNIX-like systems, the installation is very simple:

1. Download the code from GitHub https://radsecproxy.github.io/.
2. Unpack the code.
3. Navigate into the unpacked directory (the base directory)
4. type the usual UNIX compilation sequence:

```
./configure
make
make check
make install
```

4. After compiling and installing, the executable

```
radsecproxy
```

is in the installed directory. Execution of the installed binaries does not require root rights.

5. Copy the template configuration file below into

```
/etc/radsecproxy.conf
```

6. Create the directory */etc/radsecproxy/certs/ca/*. The template configuration file requires this directory to contain the accredited CA root certificates and the corresponding Certificate Revocation Lists (CRLs) in their OpenSSL hash form. See this section for information about the CA download
7. Fill the lines marked with *_STUFF_* with the required information as explained below.
8. Start radsecproxy and enjoy (for first-time use, starting it with the -f option is recommended, it will start radsecproxy in the foreground and show some verbose startup messages).

Sample configuration

Most of the radsecproxy configuration file is static. This section goes through the template radsecproxy.conf line by line and explains the meaning of each stanza.

```
ListenUDP                       localhost:11812
ListenTCP                       *:2083
```

Since there is a RADIUS server on the same host that occupies UDP/1812, radscproxy has to listen on a nonstandard port. It only needs to listen on the loopback device since it will only communicate with the RADIUS server on the same machine. The choice of 11812 is arbitrary and can be adapted if that port is in use. Since radsecproxy will also accept requests from an upstream RadSec-enabled server, it listens on the default TCP port for RadSec (2083) for requests from outside (the * meaning: all interfaces). If you want radsecproxy to listen only on specific interfaces, enter the interface names here. Beware: in this case you may also have to set the more exotic option SourceTCP (see the man page of radsecproxy for details).

Local Logging

A logging level of 3 is the default and recommended log level. Radsecproxy will then log successful and failed authentications on one line each. The log destination is the local syslog destination.

```
LogLevel                3
LogDestination          x-syslog:///LOG_LOCAL0
```

radsecproxy features a semi-automatic prevention of routing loops for RADIUS connections. If you define a client and server block (see below) and give them the same descriptive name, the proxy will prevent proxying from the client to that same server. Turn this feature on with:

```
LoopPrevention          On
```

F-Ticks

If you use Radsecproxy, you should send basic statistical information about the number of logins for national and international roaming to the eduroam Operations Team. The system to do that is "F-Ticks". radsecproxy has built-in support for F-Ticks: you simply add an option to all client { } definitions for which you know the country they are physically located in. That typically means all your connected institutions' RADIUS clients, at the national level, but excludes the international roaming top-level servers (e.g. the European Top-Level RADIUS Servers). For an institution it means all your WLAN controller connections. The client definition examples below assume that you do use F-Ticks.

When the client definitions are set-up, the following options enable F-Ticks and send the syslog messages in a privacy-preserving way (by hashing parts of the connecting end-user device's MAC address:

```
FTicksReporting Full
FTicksMAC VendorKeyHashed
FTicksKey arandomsalt
```

The ticks will end up in your local syslog daemon; they are NOT automatically sent forward to eduroam Operations. It will depend on your syslog configuration how to achieve forwarding of the messages. For "rsyslog", a popular recent syslog daemon, the following settings will make it work:

```
# radsecproxy

if      ($programname == 'radsecproxy') and ($msg contains 'F-TICKS') \
then    @192.0.2.204
&       ~
```

As usual, the IP address above is NOT the actual destination for the eduroam Operations F-Ticks server. Please contact eduroam OT for the the IP address of their server. Also keep your own server's IP address handy, because the F-Ticks server is firewalled to accept ticks only from known sources.

RADIUS/TLS

The following two sections define which TLS certificates should be used by default. This installation of radsecproxy always uses the same certificates, so this is the only TLS section. CACertificatePath contains the eduroam-accredited CA certificates with filenames in the OpenSSL hash form. The parameters below need to be adapted to point to your server certificate in PEM format, the private key for this certificate and the password for this private key if needed, respectively. If no password is needed for the private key, you can comment this line (precede it with a # sign). The option CRLCheck validates certificates against the Certificate Revocation List (CRL) of the CAs. It requires a valid CRL in place, or else the certificate validation will fail. Therefore, it is important to regularly update the CRLs by re-downloading them as described above.

Right now, checking CRLs is discouraged due to a pending bug in OpenSSL regarding CRL reloading.

Replace your paths to the certificate files as necessary - please refer to the "Certificates" section for details on how to obtain and manage RADIUS/TLS certificates.

```
tls defaultClient {
    CACertificatePath                   /etc/radsecproxy/certs/ca/
    CertificateFile                     /etc/radsecproxy/certs/CERT_PEM__
    CertificateKeyFile                  /etc/radsecproxy/certs/CERT_KEY__
    CertificateKeyPassword              __CERT_PASS__
    policyOID                           1.3.6.1.4.1.25178.3.1.1
#    CRLCheck                            On
}

tls defaultServer {
    CACertificatePath                   /etc/radsecproxy/certs/ca/
    CertificateFile                     /etc/radsecproxy/certs/CERT_PEM__
    CertificateKeyFile                  /etc/radsecproxy/certs/CERT_KEY__
    CertificateKeyPassword              __CERT_PASS__
    policyOID                           1.3.6.1.4.1.25178.3.1.2
#    CRLCheck                            On
}
```

The following section deletes attributes from RADIUS requests that convey VLAN assignment information. If VLAN information is sent inadvertently, it can cause a degraded or non-existent service for the end user because he might be put into the wrong VLAN. Connected service providers should filter this attribute on their own. Connected Identity Providers should not send this attribute at all. Checking for the existence of these attributes on your server is just an optional additional safety layer. If you do have a roaming use for cross-institution VLAN assignment, you may want to delete this stanza.

```
rewrite defaultClient {
     removeAttribute                       64
     removeAttribute                       65
     removeAttribute                       81
}
```

Client definition and request forwarding

To accept requests from the RADIUS server on localhost, this server needs to be listed as a client. That server may either use IPv4 or IPv6 for its communication, so both variants are defined. If your system is not IPv6-enabled, simply delete the stanza about client ::1. The *_LOCAL_SECRET_* must match the secret which you configured in your RADIUS server for the server catering the DEFAULT realm.

```
client localhost {
        host        127.0.0.1
        type        udp
        secret      __LOCAL_SECRET__
}

client ::1 {
        type        udp
        secret      __LOCAL_SECRET__
}
```

To accept requests from the upstream RADIUS/TLS-enabled server, this other server needs to be listed. Replace *_RADSEC_PEER_* with the hostname of your upstream server. The traditional RADIUS shared secret has no meaning in RADIUS/TLS, and must instead be statically set to "radsec" throughout eduroam. This has no security implications.

```
client _RADSEC_PEER_ {
        type        TLS
        secret     radsec
}
```

To deliver requests to the local RADIUS server, this server needs to be configured. See above for the parameter *_LOCAL_SECRET_*.

```
server localhost {
        type         UDP
        port         1812
        secret       __LOCAL_SECRET__
}
```

This server needs to be configured to deliver requests to the upstream RadSec-enabled server. See above for the configuration item *_RADSEC_PEER_*.

```
server _RADSEC_PEER_ {
        type         TLS
        secret      radsec
        statusserver on
}
```

There are some known client-side misconfigurations. If they were not already caught by the local RADIUS server, it does not make sense for the proxy to send these requests up to the eduroam infrastructure. These requests are immediately rejected.

Note: if you need to blacklist an existing realm for some reason, you can follow the myabc.com example, copying and replacing it with the realm to be blacklisted.

```
realm /myabc\.com$/ {
         replymessage "Misconfigured client: default realm of Intel PRO/Wireless supplicant!"
}

realm /^$/ {
          replymessage "Misconfigured client: empty realm!"
}
```

Requests that are coming in from upstream and are supposed to be handled by the own RADIUS server are listed in *_OWN_REALM_*. Create multiples of these stanzas if your local server serves more than one realm.

```
realm /_OWN_REALM_$/ {
             server localhost
}
```

The configuration stanza above will terminate all requests that end in *_OWN_REALM_* but which were not caught by previous realm definitions, which prevents a possible loop. This is achieved by having the client name = localhost and the server name = localhost, and LoopPrevention enabled, which checks for this condition.

Finally, all other realms which are not to be authenticated locally are sent to the upstream RADIUS/TLS peer:

```
realm * {
          server    __RADSEC_PEER__
}
```

Set up of Dynamic Discovery (in federations which have RADIUS/TLS enabled)

What is Dynamic Discovery?

eduroam has traditionally used a hierarchy of RADIUS servers. All national roaming authentication traffic was aggregated into a national proxy server; all international roaming traffic was aggregated into a set of international proxy servers.

While this works quite well under most circumstances, there are some drawbacks in efficiency, and a rather unpleasant inflexibility when it comes to routing realms which do not fit into the national aggregations model because they do not use the national .TLD ending of their federation (e.g. realms in ".net", ".org", etc.).

Dynamic Discovery places routing hints towards the responsible authentication server or national proxy into DNS, making routing more efficient.

As an IdP, you do not have to know much about the mechanics behind this - the only required step to make your realm dynamically discoverable is by adding a single resource record into your domain's DNS zone.

While adding this DNS record is optional, it has advantages for you in that it reduces the time it takes to authenticate your users when roaming internationally, so eduroam Operations RECOMMENDS to add these records if your national federation supports dynamic discovery.

For realms in generic top-level domains like .net, .org, .com etc. it is also RECOMMENDED to add these entries; and it may become mandatory at a later point in time.

Adding Dynamic Discovery hints to the IdP's DNS zone

To add Dynamic Discovery hints to your zone, you must first contact your national eduroam operator to determine which target name they have set up on the national proxy server; because you need to enter that discovery target in your DNS entry. In this documentation example, let's assume your national operator told you that the target name in your federation "Antarctica" is "_radsec._tcp.eduroam.aq". Let's further assume that your realm for eduroam and DNS domain is "greatidp.aq".

The DNS entry is of resource record type **N**etwork **A**uthority **P**oin**T**e**R** (NAPTR). These records look quite complex, but eduroam's deployment profile of the NAPTR is making it simple to understand. The full entry as required for eduroam dynamic discovery in your DNS zone is:

```
greatidp.aq.           43200   IN      NAPTR   100 10 "s" "x-eduroam:radius.tls" "" _radsec._tcp.eduroam.aq.
```

This is all! This entry says, paraphrased, "The eduroam authentication for the realm greatidp.aq works over RADIUS with TLS encryption and is handled by the service target "_radsec._tcp.eduroam.aq". Note that this does not replace your normal RADIUS uplink to your national server; this is only an additional hint to streamline *international* roaming.

Don't worry, RADIUS software knows how to interpret this further 🙂 If you are curious though, the next section explains what all these entries mean.

NAPTR explained

NAPTR records are more complex than, say, "A" records - the A record has only one piece of information to convey, namely the IP address which belongs to a name.

In contrast to that, NAPTR records are a generic entry to any kind of service. As such, it needs to specify which service this particular NAPTR entry is for, how that service is handled, and finally, who is handling it. It also provides basic failover and load-balancing mechanisms; there can be multiple NAPTR entries for the same service, with different priority and different weighting.

So let's take a look at the parts of the above entry:

Entry	Meaning
greatidp.aq.	This is the zone name (label) for which the NAPTR entry is defined

43200	DNS caching lifetime of the entry (just like any other DNS resource record)
IN	This entry is meant for consumption in the INternet (just like any other DNS resource record)
NAPTR	This entry is a Network Authority PoinTeR
100	Order: if multiple NAPTR entries are defined for the label, prefer lower order number over higher ones (Note: since eduroam requires only one single entry, any number is fine here, unless your national federation operator instructs you otherwise)
10	Preference: if multiple NAPTR entries with the same Order are defined for this label, alternate between all those entries when resolving names (Note: since eduroam requires only one single entry, any number is fine here, unless your national federation operator instructs you otherwise)
"s"	This NAPTR entry should be resolved to hostnames by doing a subsequent SRV lookup on the target label (Note: eduroam only works with "s" labels; it is a configuration error to use "a" or "u" targets)
"x-eduroam:radius.tls"	This is the service; only resolve the later target name if you want to use the service - otherwise ignore the NAPTR response (Note: this string is fixed in eduroam, as the roaming service with Dynamic Discovery is exclusively defined for RADIUS/TLS)
""	Regular Expression: some very advanced uses of NAPTR records allow transformation of target names according to regular expressions. (Note: eduroam does not make use of this feature. The regular expression field MUST be the empty string; it is a configuration error to speciffy anything else)
_radsec._tcp.eduroam.aq	The target: please contact this server (after resolving its IP addresses and port numbers) if you want to use the "x-eduroam" service

At this point you may wonder: so how does this eventually yield an IP address of my national authentication server?

The answer is: this is a first step of DNS resolution (and the only one you need to actively help with). The later steps are:

1) The server which queried for this NAPTR record will get the reply that he should resolve an SRV record (remember the "s" ?) of the target name "_radsec._tcp.eduroam.aq".

2) The querying server will then query the label _radsec._tcp.eduroam.aq and check for SRV records. The eduroam.aq zone is managed by your national eduroam admins, and the reply could look like the following:

```
_radsec._tcp.eduroam.aq. 43200  IN      SRV     10 0 2083 tld2.eduroam.
aq.


_radsec._tcp.eduroam.aq. 43200  IN      SRV     0 0 2083 tld1.eduroam.aq.
```

As you see, this reply contains two hostnames of the national eduroam servers, and also the port number to connect to (2083).

Finally, the querying server will then either ask for A or AAAA records to get to the IP address of the responsible server - and the discovery process is complete.

Testing and supporting end users

In order to test the eduroam setup from an end user perspective, please check out the Section How to offer support to end users. Also use this for helping end users get on to the newly-established eduroam network on campus.

Wired eduroam

There are use cases for securing access to a wired port in the same way as to a wireless Access Point. This use case is often not primarily intended for roaming though, as it is much more difficult for an incoming guest to locate a physical port in a building than it is to pick up a wireless broadcast signal from an access point. Much rather, an IdP who has already turned his users into eduroam users might want to benefit from the existing configuration in his end-users' devices to secure dormitories and similar locations, where finding the attachment point is less of an issue.

Connecting wired infrastructure to eduroam

The eduroam consortium is media-agnostic and welcomes wired eduroam. Wired eduroam works a lot like wireless - instead of an Access Point with 802.1 X support, an eduroam SP needs a wired switch with that same 802.1X support. The SP configures it to authenticate the physical port users to the same RADIUS server that his access points do, too (i.e. the switch becomes a RADIUS client for the eduroam SP RADIUS server) - and that's all, the switch is then part of the eduroam infrastructure proper without further changes needed.

End user experience

Experience for **end users** is somwhat grim compared to wireless.The IEEE 802.1X revisions of 2001 and 2004 lack some of the features eduroam users are used to in the wireless world where IEEE 802.11 augments the missing features of IEEE 802.1X. The following text describes the shortcomings of versions prior to IEEE 802.1X-2010, because as of this writing, no single switch and/or supplicant is known which would support the new features of IEEE 802.1X-2010. Deployers of wired eduroam should be aware of the following differences for end users.

Network indication

When connecting to eduroam, users are usually made aware that the network they are connecting to is an eduroam network - simply by observing the SSID "eduroam" occuring on their computing device. In wired IEEE 802.1X networks, the concept of SSIDs does not exist. The user needs to plug in his device and try to connect in the usual way (using his supplicant, the usual EAP configuration and his eduroam credentials) - which will work on the configured eduroam ports, but not at any other network ports. eduroam Service Providers are advised to give clear indication which ports in a building are eduroam ports and which are not. One means to achieve this is by putting an eduroam logo besides the ports in question, or announcing the existence of eduroam on a building's wired ports on a signpost near the entrance to the building.

Improvements with IEEE 802.1X-2010

With IEEE 802.1X-2010, so-called "Network Announcements" are supported. After plugging in the device, the user's supplicant is presented with a list of strings which identify the network. This feature provides the equivalent of SSIDs from wireless networks and can help end users identify usable wired network plugs.

Support for multiple configurations

Computing devices typically allow multiple independent configurations for the wireless network interface; the configurations are usually seperated by the SSID they belong to. That way, users can configure their eduroam credentials for the "eduroam" SSID, and can have any number of other configurations for other networks at work or at home. In wired IEEE 802.1X, supplicants typically only allow "the" IEEE 802.1X configuration (likely due to the missing network indication as described above, which makes it impossible for the supplicant to determine which configuration is to be used for the wired connection). Users who have a need for multiple wired IEEE 802.1X configurations will probably be dissatisfied with that situation.

Improvements with IEEE 802.1X-2010

With IEEE 802.1X-2010's network announcements, it becomes easy for supplicants to distinguish between different configurations. It is expected that, once IEEE 802.1X-2010 support on the attachment points becomes common, supplicant implementations will follow suit and enable storing and choosing different configurations per network identifier. This is implementation-specific per supplicant though, and support for this can and will vary.

Encryption of payload

In IEEE 802.11 enterprise wireless, the end-user's network payload data is encrypted between his device and the Access Point. This happens automatically when using EAP and RADIUS, and is also highly desirable for the user's security because it prevents snooping of his private data by bystanders in the broadcast domain of the access point.

In wired networks, user payload is not broadcast, but is instead directed on a copper or fibre wire directly to the switch. This makes encryption of the payload a less pressing problem. Consequently, IEEE 802.1X does not foresee encryption of user payload data.

This is not particularly alarming because of the point-to-point nature of wires, but is a difference to the wireless use case deployers of IEEE 802.1X should be aware of.

Improvements with IEEE 802.1X-2010

IEEE 802.1X-2010 includes support for a previously standalone feature called "MACsec" (previously IEEE 802.1AE). This feature allows to encrypt payload data on the wire between the end-user device and the switch. Implementations of supplicants and switches supporting this feature have not yet been observed in the wild though.