Monitoring of RADIUS Infrastructure

Best Practice Document

Produced by the AMRES-led working group
on Network Monitoring
(AMRES BPD 111)

Authors: Jovana Palibrk, Ivan Ivanović, Esad Saitović, Marina Vermezović, Marko Stojaković

February, 2013

© TERENA 2010. All rights reserved.

Document No: GN3-NA3-T4-AMRES-BPD-111
Version / date: February 2013
Original language: Serbian
Original title: “Nadgledanje RADIUS infrastrukture”
Original version / date: Revision 1 (of the document dated 10 November 2012) 28 February 2013
Contact: jovana.palibrk@amres.ac.rs, ivan.ivanovic@rcub.bg.ac.rs

AMRES/RCUB is responsible for the contents of this document. The document was developed by the AMRES-led working group on Network Monitoring with the purpose of implementing joint activities on the development and dissemination of documents encompassing technical guidelines and recommendations for network services in higher-education and research institutions in Serbia.

Parts of the report may be freely copied, unaltered, provided that the original source is acknowledged and copyright preserved.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 238875, relating to the project 'Multi-Gigabit European Research and Education Network and Associated Services (GN3)'.

Table of Contents

Summary

This document describes the implementation of the system used for monitoring a complex server authentication hierarchy based on the RADIUS (Remote Authentication Dial In User Service) protocol. The solution presented herein has been developed within the *eduroam*®[1] service of the Academic Network of Republic of Serbia (AMRES), which at the time of writing has 60 *eduroam*® access points across Serbia. The *eduroam*® authentication infrastructure requires a suitable monitoring system, which enables testing the functionalities of all the RADIUS servers this service comprises. The monitoring system has been designed to provide a sufficiently detailed insight into the state of the RADIUS infrastructure, while not infringing upon user privacy as required under the *eduroam*® policy.

The monitoring of the infrastructure of the RADIUS-based server on the AMRES network is conducted to check the availability of RADIUS servers through the network, as well as to establish whether the RADIUS servers are processing client authentication requests in the appropriate manner. This document presents a simple and scalable monitoring solution and this solution can also be used in other environments where services rely on the RADIUS protocol.

[1] eduroam is a registered trademark of TERENA, the Trans-European Research and Education Networking Association

Introduction

eduroam® (EDUcation ROAMing) is a service developed in an international academic environment by the TERENA Task Force, TF Mobility and it has been realised through the GÉANT project. The purpose of this service is to provide users of academic institutions throughout Europe, and recently even further afield, with safe, fast and simple access to wireless Internet at all points across the world where *eduroam*® is implemented. Whether they are accessing the *eduroam*® service at their home institution or at an institution they are visiting, users can use the credentials they have received from their home institution.

The basic security principle of the *eduroam*® service is that user authentication is performed through the user's home institution using the EAP (Extensible Authentication Protocol) authentication methods implemented by that institution, regardless of the point of access to the network. The authorisation to access the Internet and other network services provided by local network resources is granted by the accessing network.

In order to enable the secure exchange of authentication messages between the user visiting an institution and the RADIUS server of his/her home institution. The RADIUS architecture supporting *eduroam*® services is structured hierarchically on three levels:

- the highest (top) level includes the European RADIUS servers – ETLR (European Top-Level RADIUS), which contains a list of national domains;
- the second level is comprised of the national FTLR (Federation Top-Level RADIUS) servers, which contain a list of institutional domains within a certain country. The FTLR servers are managed by the national roaming operator, which is AMRES in the case of Serbia;
- the third level includes the RADIUS servers of end institutions – the identity and/or resource providers. Specifically, each institution can participate in the *eduroam*® service by providing the *eduroam*® service to its users (identity provider) and/or by providing access to the Internet through the *eduroam*® service (resource provider).

This document introduces the system for monitoring the national RADIUS hierarchy, which is comprised of the individual RADIUS servers of institutions and a FTLR server. The document begins with a presentation of the RADIUS infrastructure of the AMRES network, for whose purposes this solution has been developed. The document further describes, in detail, the manner of implementation of the monitoring system. Finally, the document provides concrete examples of the testing of RADIUS servers and the results obtained on the AMRES network.

.

1 The Hierarchical Structure of the RADIUS Infrastructure

On the AMRES network, the RADIUS hierarchy is comprised of the RADIUS servers located at the member-institutions of the *eduroam*® service; these RADIUS servers are directly linked to the FTLR server. The RADIUS servers of the institutions are responsible for authenticating their users, whether they are trying to use the *eduroam*® service at their home institution or an institution they are visiting. The home institutions are also responsible for maintaining the data and credentials of their users, which are usually kept in a database used by the RADIUS server in the process of authentication.

The username is in the form of *username@domain*, where the *domain* is the DNS (Domain Name Server) name of the institution. The RADIUS servers use the information about the institution's domain in order to decide whether a request should be processed locally (the server itself is responsible for the said domain) or if it should be routed through the hierarchical structure to the RADIUS server of the user's home institution.

By using the appropriate EAP authentication methods, a secure tunnel is established from the device through which the user accesses the network (e.g., a computer or a mobile phone) to the authentication server at his/her home institution, which transmits the information relevant for authentication of the user. The AMRES *eduroam*® service uses the EAP-TTLS (EAP Tunnelled Transport Layer Security)[i] or the PEAP (Protected EAP) protocols[ii]. These authentication methods serve as a basis for establishing secure TLS (Transport Layer Security) sessions (tunnels) from the end user's device to the authentication server at his/her home institution so that the user's sensitive information is protected from any interception. An example of setting up an EAP-TTLS secure tunnel within the *eduroam*® services is shown in Figure 1.

The types of information exchanged within the TLS tunnel depend on the protocol used for authenticating the user's identity, which can be the PAP (Password Authentication Protocol), CHAP (Challenge Handshake Auth Protocol), MSCHAP (Microsoft CHAP), EAP-GTC (Generic Token Card), or MD5-Challenge protocols.

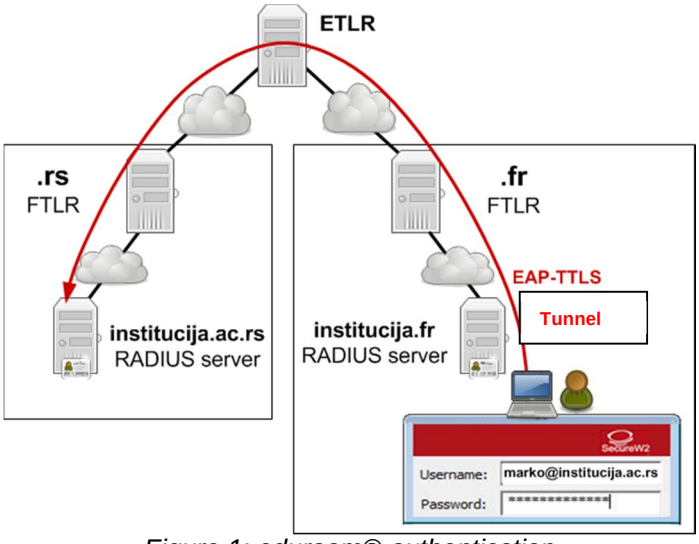

Figure 1: eduroam® authentication

2 The Monitoring System of the RADIUS Infrastructure

The system for monitoring the RADIUS infrastructure in the Academic Network of the Republic of Serbia has been implemented as part of the centralized monitoring system of the network – the NetIIS (Networking Information and Monitoring System), developed at the Computer Centre of the University of Belgrade[iii].

The NetIIS system enables passive and active monitoring of the equipment and services in the entire AMRES network. The monitoring tools can monitor different parameters and a time interval can be configured for each type of monitoring that is to be performed. The elements of the NetIIS system, through which the monitoring is performed, are called monitors. These monitors can be assigned a graphic overview of the time statistics, relying on the RRD (Round-Robin Database) and MRTG (Multi Router Traffic Grapher) systems. Likewise, each monitor can be assigned alarms that will notify the user about any event related to the monitor.

The monitors that test the availability of the RADIUS servers through the network use the ICMP protocol, and the service has been implemented in order to test the availability of the RADIUS servers in the NetIIS system. However, even if a server is available via the network, it does not necessarily mean that it is operational. The RADIUS process may be disabled for some reason, so that the server is not actually functional, despite its availability. Therefore, the operability of the RADIUS service on the server is also tested.

The functionality of the RADIUS service is tested using the eapol_test program[iv]. This program is a part of the wpa_supplicant software[v] and represents the WPA (Wi-Fi Protected Access) supplicant – the software on the user's device that provides support for the WPA and WPA2 (Wi-Fi Protected Access II) protocols.

The eapol_test program is used to simulate an access device by sending Access-Request messages to the RADIUS server in order to authenticate the user who wants to access the network (Figure 2). The input parameters for this test are the IP address of the RADIUS server that needs to authenticate the user or reroute the request to another RADIUS server that will perform authentication, the protocol used for client authentication, and the test account (username and password). Each member institution of the Academic Network of the Republic of Serbia that participates in the *eduroam*® service needs to create a special test user account that will be used exclusively for testing the RADIUS server. If the test account is already in the database, along with other user accounts, the communication between the RADIUS server and the Access-Request message, the user database will be checked along with the operability of the RADIUS server. The Access-Request message containing the user credentials from the test account is sent via the eapol_test program to the RADIUS server being tested. In this way, EAP-TTLS or PEAP secure connection (tunnel) is established.

A shell script used for running the eapol_test is created on the server where the NetIIS system is installed. The NetIIS monitors run the script periodically, i.e., at five-minute intervals. Since the NetIIS supports integration with the Nagios monitors, the monitoring script is designed to show the results of the operation in the Nagios format. If the RADIUS server responds to a request with the Access-Accept message, the monitor provides information that everything is in order. Otherwise, an alarm is activated within the monitor and email notifications are sent to a defined group of users (the IT staff of the institution) about a problem in the

functioning of the RADIUS server. Once the problem is solved and the monitor value is changed, the monitor alarm will again be activated and the same group of users will be notified about the proper functioning of their RADIUS server. In this way, the notifications are sent through the NetIIS system to the IT staff of the institution each time the monitor value changes.

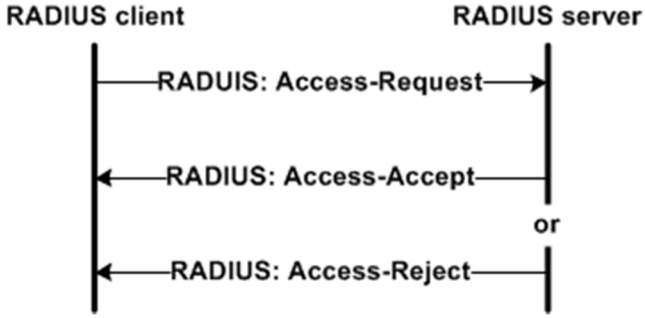

Figure 2: The flow of exchange of RADIUS messages during authentication

The NetIIS system sends requests to the tested RADIUS servers by running the shell script shown below.

```
#!/bin/bash
#Input variables
a1=$2
a2=$4
a3=$6
a4=$8
shift 2
a5=$8
shift 2
a6=$8

#Defining of the eap.conf.randomnumber file that contains the data for the
#eapol_test script
xy=`cat /dev/urandom | tr -dc _A-Z-a-z-0-9 | head -c8`
touch ./eap.conf.$xy

#Defining the authentication protocol used within TLS tunnel.
#With EAP-TTLS protocol PAP protocol is used, and in case of PEAP protocol
#MSCHAPv2 is used
if [ "$a6" == "TTLS" ]; then
phase2="auth=PAP"
else
phase2="autheap=MSCHAPV2"
fi

#Population of eap.conf.randomnumber file
echo network={                          >./eap.conf.$xy
echo        ssid=\"eduroam\"            >>./eap.conf.$xy
echo        key_mgmt=WPA-EAP            >>./eap.conf.$xy
echo        eap=$a6                     >>./eap.conf.$xy
echo        identity=\"$a2\"            >>./eap.conf.$xy
echo        anonymous_identity=\"$a3\"  >>./eap.conf.$xy
```

```
echo        password=\"$a4\"            >>./eap.conf.$xy
echo        phase2=\"$phase2\"          >>./eap.conf.$xy
echo        }                           >>./eap.conf.$xy
#Time variable t1 defines the beginning of the eapol_test script execution
t1=`date +%s%N`
# Execution of eapol_test script
a=`eapol_test -c eap.conf.$xy -a $a1 -s $a5 -r 1 -t 10 2>/dev/null| tail -1`
#Time variable t2 defines the end of the eapol_test script execution
t2=`date +%s%N`
#dtN defines how long eapol_test script takes to execute
dtN=`expr $t2 - $t1`
dt=`expr $dtN / 1000000`

rm -fr ./eap.conf.$xy
#Creating the Nagios output form
if [ "$a" == "SUCCESS" ]; then
        echo "OK" "|""value=1;Response time:$dt ms"
else
        echo "False" "|""value=0;Response time:$dt ms;Radius virtual server
credential error or eduroam user doesn't exist or radius down!"
fi
```

The input parameters required by the script are:

- a1 – the IP address of the RADIUS server being tested. This can be the RADIUS server of the identity-provider institution and/or the resource-provider institution and the FTLR server;
- a2 – the username of the test account;
- a3 – anonymous identity;
- a4 – the password of the test account;
- a5 – a shared secret key used so that the tested RADIUS server can accept requests sent from the NetIIS server. The same key also needs to be defined on the RADIUS server side;
- a6 – the protocol used for client authentication. The EAP TTLS/PAP or the PEAP/MSCHAPv2 protocols are implemented in the AMRES network.

The input parameters are entered through the NetIIS web interface. The input values are entered in the appropriate variables within the script. The eap.conf configuration file is created from a part of variables (a2, a3, a4 and a6). The configuration file has a predefined structure, and it is used along with the other data (a1 and a5), when running the eapol_test program. The result of the execution of the eapol_test program is assigned to variable a, and its value can be either SUCCESS or FAILURE. Based on the result obtained, the NetIIS monitor that runs the script receives the corresponding value: 1 if the message received is Access-Accept, or 0 if the message received is Access-Reject.

The time intervals before the running of the eapol_test program are recorded at the script, as well as the time after receiving an eapol_test response. Recorded timestamps are used in order to calculate response-time of the tested Radius server.

The script is optimised for running from the NetIIS system. However, the script is simple and can easily be

adapted to any monitoring system with sufficient flexibility and configurability to enable the creation and running of the Linux or Nagios scripts. Even in situations when there is no computer network monitoring system, the functions for sending email messages to notify the institution's IT staff of any change of status of their RADIUS servers can be added to the script and run periodically by using the Linux cron service.

3 Monitored Parameters

Two groups of monitors are configured on the NetIIS system in order to monitor the RADIUS infrastructure. One of them tests the availability of RADIUS servers through the network, while the other is designed to test the operability and measure performance of the RADIUS servers.

3.1 Testing the Availability of RADIUS Servers

A ping monitor is configured in order to check the availability of servers. This monitor sends ICMP requests to the IP address of the RADIUS server. Based on the response received, graphs are generated that show packet delay (Figure 3a) and packet loss (Figure 3b) within a selected time interval.

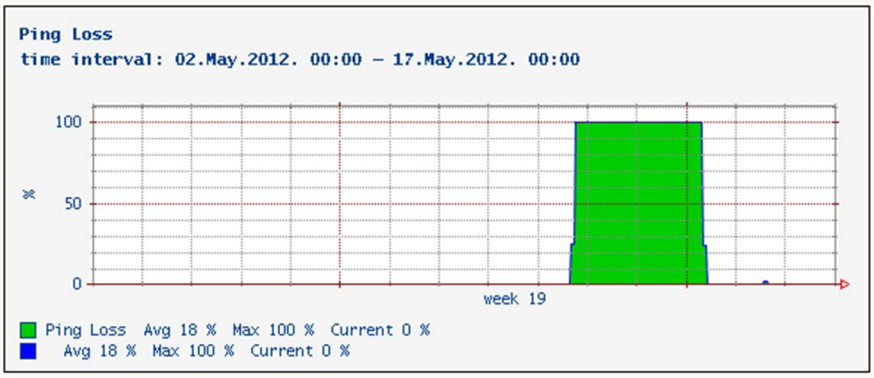

Figure 3a: The delay of ICMP packets exchanged with the RADIUS server over a fifteen-day period

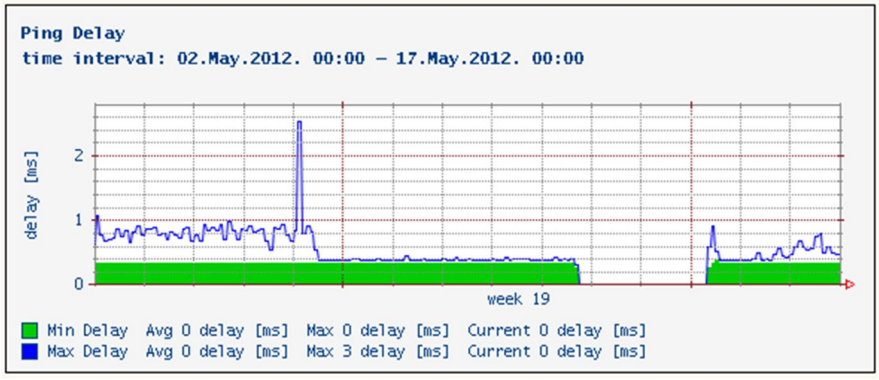

Figure 3b: The loss of ICMP packets exchanged with the RADIUS server over a fifteen-day period

If the server is unavailable, the packet loss will be 100%. If there are any problems in the network, e.g., if the network is congested, this will be shown as a certain packet loss percentage, which will be less than 100%.

Therefore, a problem in the network on the way to the RADIUS server can be detected based on these results. As the RADIUS relies on the UDP protocol, this information may be very important. Figures 3a and 3b show that the RADIUS server was unavailable for three days (from 11 until 14 May) since the ICMP packet loss during that period was 100% and no values for packet delay were recorded.

3.2 Testing the Operability of RADIUS Servers

The functionality of an institution's RADIUS server differs when the institution is the identity-provider and when it is the resource-provider. The RADIUS server of the identity-provider institution processes the received requests for authentication of local users. If the institution is also the resource-provider, the RADIUS server should be configured to forward the requests of users of other institutions to the FTLR server, which will then forward them to the RADIUS server of the users' home institution. If the institution is only the resource provider, its RADIUS server will forward all received requests to the FTLR server. For that reason, four possible scenarios have been taken into account during the configuration of the monitors that test the operability of the RADIUS servers:

- Scenario 1 simulates the situation in which an *eduroam*® user uses the *eduroam*® service at his/her home institution. The corresponding monitor checks how the tested RADIUS server processes authentication requests received directly from local users.

- Scenario 2 addresses the situation in which a user of the institution whose RADIUS server is being tested uses the *eduroam*® service at another institution. The corresponding monitor checks how the tested RADIUS server processes the authentication request of a local user received from the FTLR server.

  The first two scenarios are used for testing the functionality of the RADIUS server of an identity-provider institution.

- Scenario 3 simulates the situation in which a user of another institution is using the *eduroam*® service at the institution whose RADIUS server is being tested. The corresponding monitor is used to check the functionality of the RADIUS server of the resource-provider institution. Specifically, the test is performed in order to establish whether the RADIUS server will forward the authentication requests of users from another institution to the FTLR server.

  The monitors where these tests have been implemented for the AMRES institutions that are members of the *eduroam*® service can be found on the following link: http://netiis.rcub.bg.ac.rs/netiis/NetIIS?service=main&ID=group.35865.

- Scenario 4 is used for testing whether the FTLR server successfully forwards user authentication requests.

  The monitors testing the functionality of the AMRES FTLR server can be found on the following link: http://netiis.rcub.bg.ac.rs/netiis/NetIIS?service=main&ID=group.37378.

The combination of the above tests can help locate problems in the operation of the RADIUS infrastructure. If the test conducted for Scenario 1 has been successful and the test conducted for Scenario 2 has not, it can be concluded that the RADIUS server can authenticate local users, but that there is a problem between the tested server and the FTLR server. In that case, it is likely that one of the following problems has occurred:

- An error has occurred during the configuration of the FTLR server and as a result, it cannot forward the request to the tested RADIUS server.

- There is a problem in the communication between the RADIUS server of the institution and the FTLR server (for example, the connection parameters are configured incorrectly, or the RADIUS servers are not available to each other through the network).

- The FTLR server is not operational, which will be indicated by the monitor used for testing Scenario 4.

If the tests for both Scenario 1 and Scenario 2 have been unsuccessful, it can be concluded that the RADIUS server of the institution is not functioning properly. This, of course, does not eliminate the possibility of a problem with the operability of the FTLR server. That is why the Scenario 4 monitor is used, which indicates whether the FTLR server functions properly.

The failure of the monitor used in Scenario 3 can be caused by a problem in the functioning of either the tested RADIUS server or the FTLR server. In that case, the status of the Scenario 4 monitor is used in order to locate the problem that needs to be solved.

3.2.1 Scenario 1

The corresponding monitor in the NetIIS system uses the script to run the eapol_test, which simulates a client sending a request to the RADIUS server and establishes the direct EAP TTLS or PEAP tunnel between them, depending on the protocol implemented at the institution whose RADIUS server is being tested (Figure 4). The credentials from the test account are used for authentication. The RADIUS server processes the received request, and based on the result obtained, NetIIS generates a graph showing the status of the server (Figure 5a). If the message received is Access-Accept, the value of the status is 1. Otherwise, the value of the status is 0. The graph showing the response delay is also obtained (Figure 5b). If everything is functional, the delay has a constant value over a certain time interval. If the value of the delay varies over time, it can be concluded that there are problems in the operation of the server itself, e.g., the server resources could be overloaded, which can point to the possibility of a DoS (Denial-of-Service) attack on the RADIUS server. The delay of messages can also be a consequence of a problem in the communication between the RADIUS server and the user database (in cases when the test account is kept in the user database) or it may come as a result of problems in the network on the way to the RADIUS server.

Note: if the response does not reach the NetIIS system, the delay graph will display the default value of 10s, which is the maximum time allowed to get the response (the timeout parameter) defined in the script at the moment of running the eapol_test program.

Figure 4: The establishment of the EAP tunnel directly between the client and the RADIUS server

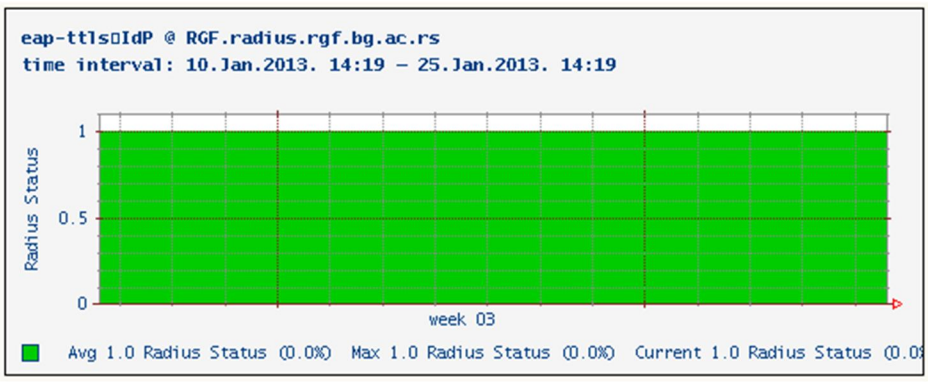

Figure 5a: The status of the RADIUS server in a 15-day period: Scenario 1

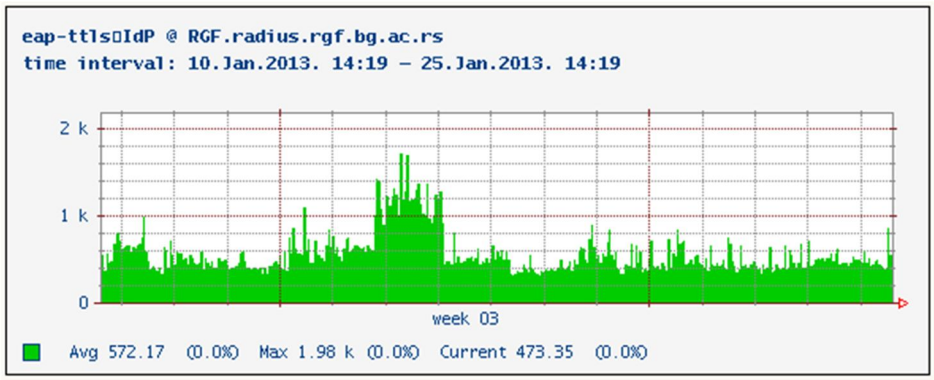

Figure 5b: The delay in the response to the RADIUS server over a fifteen-day period: Scenario 1

Figure 5a shows that in the displayed fifteen-day period, the monitored RADIUS server successfully processed all received requests because the status of the server is 1, while Figure 5b indicates that in the course of one day (16 January), the value of message delay increased, due to either server overload or potential interruptions in the network.

3.2.2 Scenario 2

In order to monitor the EAP tunnel established between the client's device and the RADIUS server through the FTLR server, the monitor implemented in Scenario 2 sends a request through the eapol_test program using the same user credentials applied in the previous scenario, but this time it sends them to the IP address of the FTLR server, which, if everything is functioning properly, forwards this request to the tested RADIUS server (Figure 6).

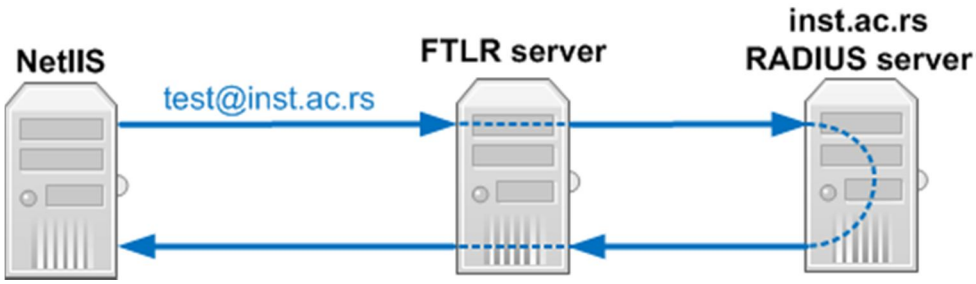

Figure 6: The establishment of the EAP tunnel between the client and the RADIUS server through the FTLR server

Based on the results obtained, graphs are generated showing the status of the server (Figure 7a) and the response delay (Figure 7b).

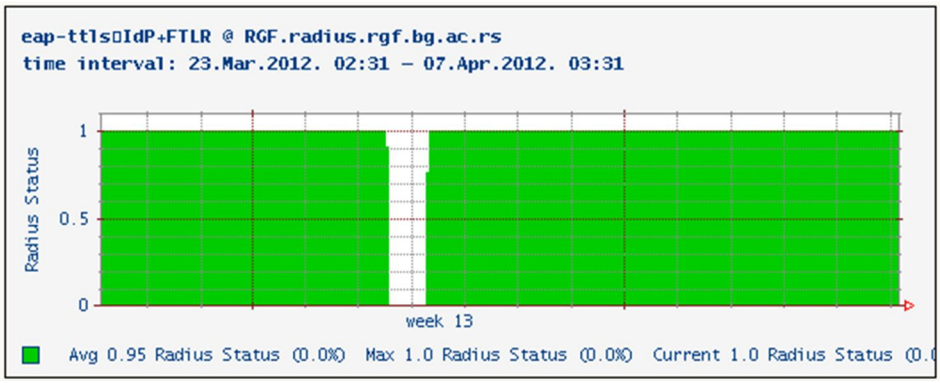

Figure 7a: The status of the RADIUS server over a fifteen-day period: Scenario 2

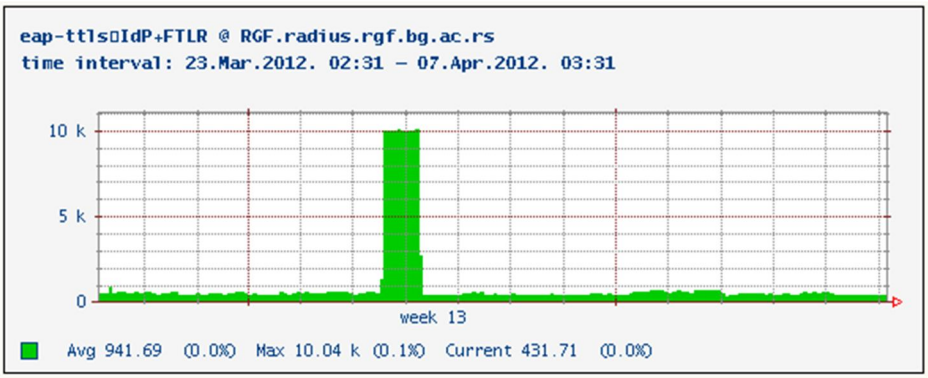

Figure 7b: The response delay of the RADIUS server over a 15-day period: Scenario 2

Figure 7a shows that the NetIIS system did not receive any responses from the tested RADIUS server in the period from 28 to 29 March 2012. Therefore, the value of delay presented in the graph shown on Figure 7b is 10s, which represents the value of the timeout parameter defined at the moment of running the eapol_test program.

3.2.3 Scenario 3

In order to test the operability of the RADIUS server that should send user authentication requests to the FTLR server (Figure 8), an additional test RADIUS server is installed. This test RADIUS server is used only for the purpose of monitoring. The NetIIS monitor sends a request through the eapol_test program directly to the RADIUS server being tested. The authentication data of the test account with the test domain defined at the test RADIUS server are used in this case. The RADIUS server forwards the received request to the FTLR server, which then forwards the request to the test RADIUS server. Of course, the FTLR server is configured to communicate with the test RADIUS server. Based on the results obtained, graphs are generated showing the server status (Figure 9a) and the response delay (Figure 9b).

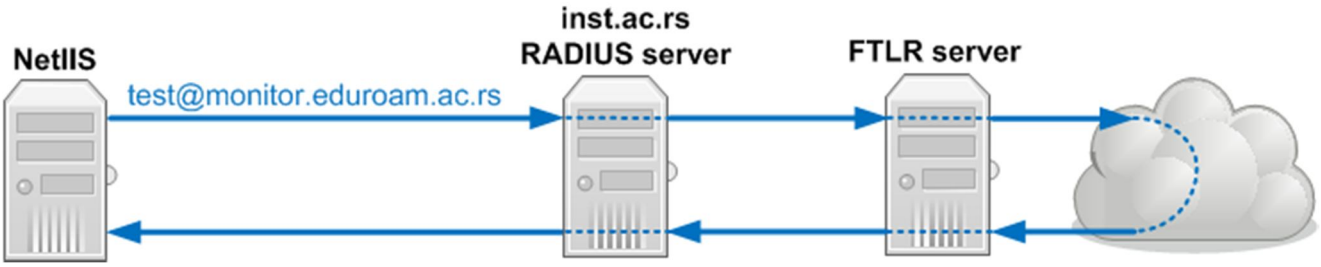

Figure 8: The establishment of the EAP tunnel directly between the client and the RADIUS server of his/her home institution through the tested RADIUS server

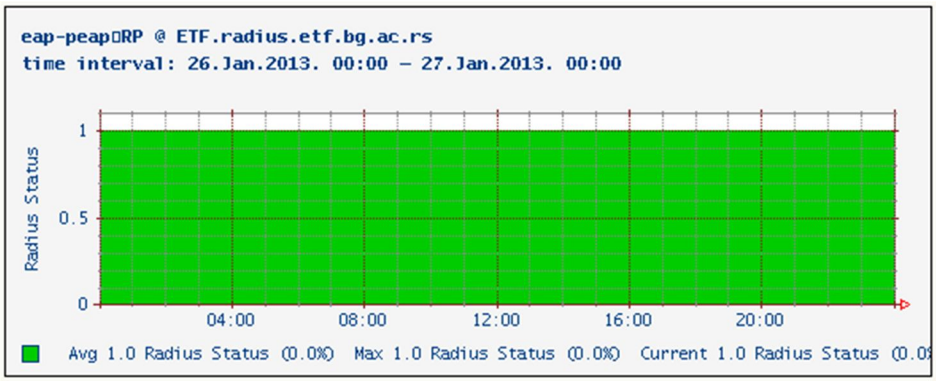

Figure 9a: The status of the RADIUS server in the one-day period: Scenario 3

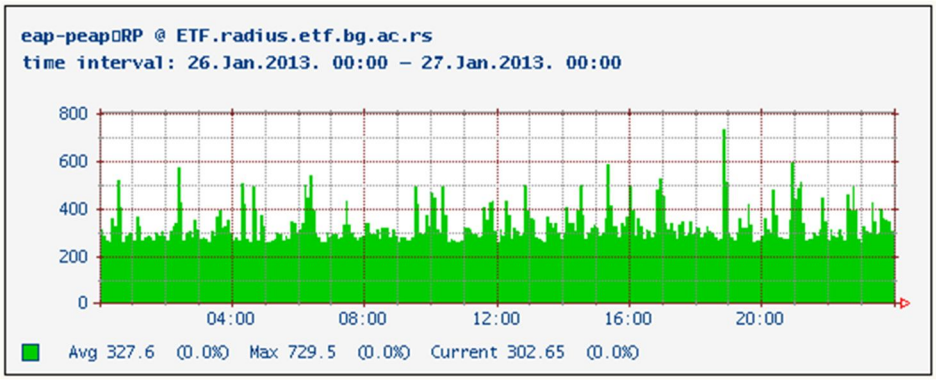

Figure 9b: The response delay of the RADIUS server in the one-day period: Scenario 3

Figures 9a and 9b show examples when the monitored RADIUS server functioned properly. The value of the RADIUS status is 1, while the delay is almost constant, with an average value of 330 ms.

3.2.4 Scenario 4

The test RADIUS server from the previous scenario is also used when testing the functionality of the FTLR server. A request for authentication of a test user is sent through the eapol_test program directly to the FTLR server. The FTLR server forwards the request to the test RADIUS server, whose response is sent back to the NetIIS system by the FTLR server (Figure 10). Based on the results obtained, a graph is generated showing the status of the server.

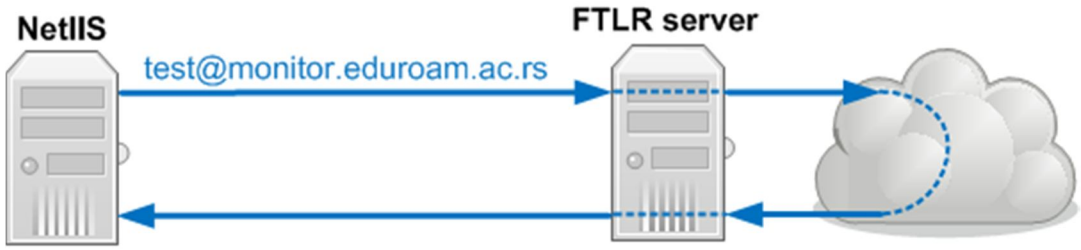

Figure 10: The establishment of the EAP tunnel through the FTLR server

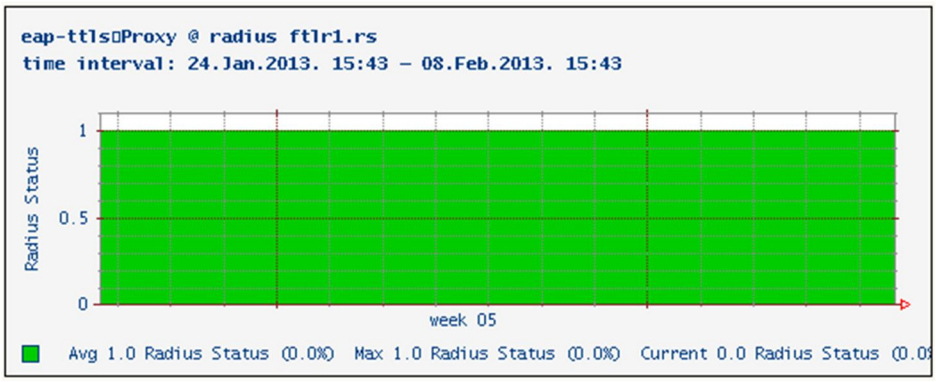

Figure 11: The status of the FTLR server over a fifteen-day period: Scenario 4

This completes the system for monitoring the national RADIUS infrastructure, which is based on a simple, scalable, and centralised solution. The system of monitors is implemented to provide a representation of the functionality and availability of the national RADIUS server hierarchy. The alarm system is also implemented in order to notify the IT staff about anomalies in the functioning of the service. The principles of monitoring the RADIUS servers presented above are also applicable to less complex configurations of authentication infrastructure, and also to other services operating within telecommunication networks.

Glossary

RADIUS	Remote Authentication Dial In User Service
FTLR	Federation Top Level RADIUS
ETLR	European Top Level RADIUS
EAP	Extensible Authentication Protocol
EAP-TTLS	EAP Tunnelled Transport Layer Security
PEAP	Protected EAP
TLS	Transport Layer Security
PAP	Password Authentication Protocol
MSCHAPv2	Microsoft Challenge Handshake Auth Protocol version 2
NetIIS	Networking Information and Monitoring System
MRTG	Multi Router Traffic Grapher
WPA	Wi-Fi Protected Access

References

[i] P. Funk, S. Blake-Wilson, “Extensible Authentication Protocol Tunneled Transport Layer Security Authenticated Protocol Version 0 (EAP TTLSv0)”, RFC 5281, August 2008.
[ii] A. Palekar, D. Simon, G. Zorn, S. Josefsson “Protected EAP protocol (PEAP)”, INTERNET-DRAFT, March 2003
[iii] S. Gajin, D. Pajin, D. Novaković, “Sistem za nadgledanje računarske mreže - NetIIS”, YuInfo, 2006 http://www.e-drustvo.org/proceedings/YuInfo2006/html/pdf/149.pdf
[iv] Deploying RADIUS - http://deployingradius.com/scripts/eapol_test/
[v] Linux WPA/WPA2/IEEE 802.1X Supplicant - http://hostap.epitest.fi/wpa_supplicant/