


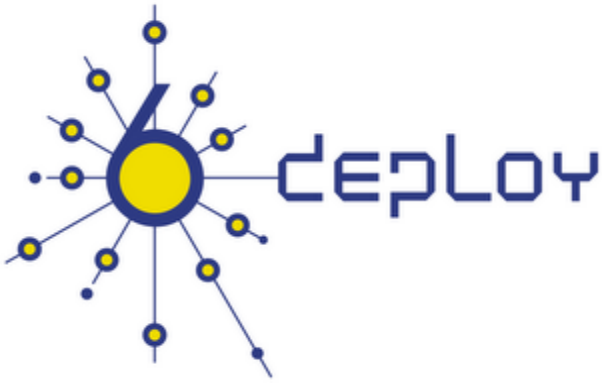

Title:	Document Version:
Deliverable D2.1.3 **Report of 3rd Deployment Case Study**	1.3

Project Number:	Project Acronym:	Project Title:
223794	6DEPLOY	IPv6 Deployment Support

Contractual Delivery Date:	Actual Delivery Date:	Deliverable Type* - Security**:
30th September, 2009	6th January, 2010	R - PU

* Type: P – Prototype, R – Report, D – Demonstrator, O - Other
* Security Class: PU – Public, PP – Restricted to other programme participants (including the Commission Services), RE – Restricted to a group defined by the consortium (including the Commission Services), CO – Confidential, only for members of the consortium (including the Commission Services)

Responsible and Editor/Author:	Organization:	Contributing WP:
Atanas Terziyski	University of Plovdiv / BREN	WP 2

Authors (organizations):

Atanas Terziyski, University of Plovdiv, BREN

Abstract:

This deliverable reports on the process of IPv6 deployment of the campus network at the University of Plovdiv, Bulgaria. It describes the network planning and technical configurations, as well as the experiences gained during the deployment.

Keywords:

IPv6, campus, case study

Disclaimer

Executive Summary

This report describes the process of IPv6 deployment at University of Plovdiv, which is the first IPv6 campus network built within the Bulgarian Research and Education Network [BREN]. The following areas will be discussed in this report:

- Network planning
- Technical configurations
- Network services
- Network monitoring and security
- Promoting of IPv6 at the university campus

Table of Contents

List of Figures

1. NETWORK PLANNING

This section gives an overview of the University of Plovdiv's campus network as well as the discussions that took place regarding the planning and preparation for IPv6 implementation into the existing IPv4 network.

1.1 Overview

The University of Plovdiv represents a small campus network (approximately 2000 IP devices within the campus) with only one line that connects the Border Router to the outside network, thus making the network planning of IPv6 deployment a fairly straightforward process. The entire process took less than one year from the planning stage to the actual IPv6 deployment, and the IPv6 network is currently live.

The IPv6 deployment at the University of Plovdiv needed less time on the initial preparation, such as studying the IPv6 capability for the systems, compared to the planning of a larger scale campus network. Instead, the focus was placed more on the actual deployment itself and in gaining on-site experiences throughout the operation. In this way, we provided both staff and students with the educational background to be able to give IPv6 training themselves and support the testbed activities.

1.2 Existing IPv4 Network

Network Infrastructure

Since the University has only one line that connects the campus network to BREN, our IPv4 network infrastructure was designed to remain simple yet efficient for the demands of the campus community.

The network connections from the different campus locations were grouped into 7 subnets with one additional subnet exclusively for user servers; a PC-based router then connects all the subnets to the PC Border Router. The campus network is then connected to BREN, via BREN to GÉANT and further on to the general Internet.

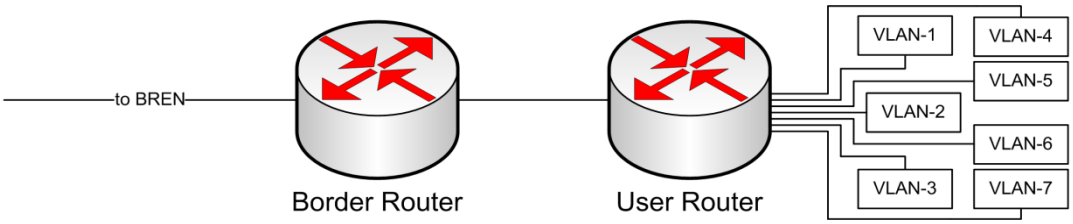

Figure 1 Overview of the University of Plovdiv IPv4 network infrastructure

Due to the University's low financial budget, PC-based routers were chosen for our campus network. Both the Border Router and the User Router run Debian Linux OS and are equipped with two INTEL PRO/1000 MT Quad Port Ethernet adapters plus two 1GE ports on the mainboard.

Network Services

Despite the small size of our network, many registered users (e.g., students) use our network services outside the campus. We provide typical campus network services such as our own mail and web servers, DNS and DHCP for IPv4, wireless access points, Content Management System (i.e., Moodle), hosting space i.e., forum and blog services, live streaming campus radio and IP-Telephony services. All registered users on the University of Plovdiv's campus network are also EDUROAM-enabled.

1.3 IPv4-to-IPv6 Transition

IPv6 has increased in popularity over the last decade based on advantages such as expanded addressing space and added functionalities such as auto configuration and the avoidance of NAT boxes. Being a small campus network, there was no immediate pressure in terms of lack of addressing space to transit from IPv4 to IPv6. However, it was considered to be worthwhile to initiate IPv6 deployment first on a small scale network, in order to obtain the practical experience. This will be advantageous for future network expansion, at the same time providing an educational environment (e.g. training and testbed activities for the staff and students) to acquire and disseminate further knowledge regarding IPv6.

In our campus network, we needed to design the IPv6 network so that the new IPv6 network could coexist with the existing IPv4 network. i.e. IPv6 should be deployed alongside the existing IPv4 infrastructure. In order to achieve this aim, the dual-stack mechanism was used. Using this method, the host or the router is installed with both (IPv4 and IPv6) protocol stacks in the operating system and each such node is configured with both IPv4 and IPv6 addresses. Therefore, it can communicate with all other nodes in the IPv4 and IPv6 network. In this scenario, there is no real transition mechanism to use since it is simply the approach to integrate IPv6 itself. Instead, there are other challenges to consider when using the dual-stack approach, including routing configurations and how to manage the two protocol interactions within the network. Nevertheless, the dual-stack approach was determined to be a cost-saving method of integrating IPv6 into our current campus network.

1.3.1 Address Planning

Step One - Obtaining an IPv6 network prefix

There are 13 Local Internet Registries (LIRs) in Bulgaria that could theoretically provide IPv6 prefixes, but only 4 (including BREN) are currently active. The IPv6 network prefix assigned to the University of Plovdiv by BREN is /42, specifically 2001:4b58:27::/42. It is registered with RIPE and can be found via the RIPE database query:

```
inet6num:       2001:4b58:27::/48
netname:        ISTF-RC-PU-IPv6
descr:          ISTF (NREN - Bulgaria)
```

```
descr:          University of Plovdiv
country:        BG
admin-c:        SE508-RIPE
tech-c:         KSX2-RIPE
status:         ASSIGNED
mnt-by:         AS6802-MNT
mnt-lower:      AS6802-MNT
mnt-routes:     AS6802-MNT
mnt-domains:    AS6802-MNT
source:         RIPE # Filter
```

Step Two - Allocating IPv6 address space

We chose to assign the prefix /64 to all the 8 subnets (including the user server) which exist in the current IPv4 network, in order to apply the IPv6 stateless Router Advertisement (RA) autoconfiguration approach as defined in RFC2462 [RFC2462]. This mechanism allows the IP devices in the network to acquire 'automatically' their IP address, without requiring any intermediate IP support in the form of DHCP. Thus, it saves time and costs compared with the traditional network address configuration process.

Step Three - Assigning the IPv6 address

As mentioned in the previous step, stateless auto configuration for all the subnets was utilized. All the IP devices in the subnets get their IP address without any other intermediate IP support (e.g. DHCP).

As for the two routers, manual configurations are applied because it provides the network administrator better control of the address assignments.

2. TECHNICAL SOLUTIONS

This chapter includes the technical configuration information for the University of Plovdiv's IPv6 network and discusses the network monitoring and security issues.

2.1 Routers

As illustrated in Figure 1, all the subnets of the campus network are connected to the User Router, and then via the Border Router to BREN.

Manual configuration is thus used at the two Ethernet adapters on the Border Router: `eth0` and `eth1`, where `eth0` connects between the Border Router and the User Router, and `eth1` connects the Border Router further to BREN.

Border Router – PC-based router under Debian Linux OS

```
eth0      Link encap:Ethernet  HWaddr 00:0e:0c:4f:35:a0
          inet addr:194.141.27.178  Bcast:194.141.27.179  Mask:255.255.255.252
          inet6 addr: 2001:4b58:27:252::1/126 Scope:Global
          inet6 addr: fe80::20e:cff:fe4f:35a0/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:446035004 errors:0 dropped:0 overruns:61 frame:61
          TX packets:265249959 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1992686436 (1.8 GiB)  TX bytes:678552 (662.6 KiB)

eth1      Link encap:Ethernet  HWaddr 00:0e:0c:4f:2f:09
          inet addr:194.141.252.102  Bcast:194.141.252.103  Mask:255.255.255.252
          inet6 addr: 2001:4b58:acad:252::36/126 Scope:Global
          inet6 addr: fe80::20e:cff:fe4f:2f09/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:295727809 errors:0 dropped:0 overruns:0 frame:0
          TX packets:393822001 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:2097827087 (1.9 GiB)  TX bytes:1891518313 (1.7 GiB)
```

The routing table for the Border Router is thus as follows:

```
Destination                    Next Hop                   Flag Met Ref Use If
2001:4b58:27:252::/126         ::                         U    256 0     1 eth0
2001:4b58:27::/48              2001:4b58:27:252::2        UG   1   0 84881 eth0
2001:4b58:acad:252::34/126     ::                         U    256 0     1 eth1
fe80::/64                      ::                         U    256 0     0 eth0
fe80::/64                      ::                         U    256 0     0 eth1
::/0                           2001:4b58:acad:252::35     UG   1   0192127 eth1
::/0                           ::                         !n   -1  1354577 lo
::1/128                        ::                         Un   0   1 16236 lo
2001:4b58:27:252::/128         ::                         Un   0   1     0 lo
2001:4b58:27:252::1/128        ::                         Un   0   1141156 lo
2001:4b58:acad:252::34/128     ::                         Un   0   1     0 lo
2001:4b58:acad:252::36/128     ::                         Un   0   1  9887 lo
fe80::/128                     ::                         Un   0   1     0 lo
fe80::/128                     ::                         Un   0   1     0 lo
fe80::20e:cff:fe4f:2f09/128    ::                         Un   0   1 28194 lo
fe80::20e:cff:fe4f:35a0/128    ::                         Un   0   1 28589 lo
ff00::/8                       ::                         U    256 0     0 eth0
ff00::/8                       ::                         U    256 0     0 eth1
::/0                           ::                         !n   -1  1354577 lo
```

User Router - PC-based router under Debian Linux OS

The User Router connects all the subnets of the campus network and via `eth8` to the Border Router. All the subnets are configured using the IPv6 stateless address autoconfiguration method that allows the hosts or IP devices to acquire the IP address automatically without requiring of any intermediate IP support. In this manner, the DHCPv4 daemon files were configured on the User Router using the open ISC DHCP [ISCDHCP]. ISC DHCP is a collection of software that implements all aspects of the DHCP suite for connection to a local network. It is a reference implementation of the protocols and also production-grade software that includes a DHCP server, a DHCP client and a DHCP relay agent:

- The DHCP server receives requests and replies to them.
- The DHCP client is bundled with the OS of a client computer and sends requests to the server.
- The DHCP relay agent passes DHCP requests from one LAN to another so that DHCP is not required on every LAN.

The DHCP server answers requests from any client that complies with the protocol standards, and the DHCP client interacts with any server that complies with those standards.

The following shows how DHCP was configured:
In the start-up script file `/etc/rc.local`

```
# starting dhcpd
dhcpd eth0 eth2 eth3 eth10 eth11 eth12 eth16
/etc/init.d/radvd start
```

The command `dhcpd eth0 eth2 eth3 eth10 eth11 eth12 eth16` shows:

```
Internet Software Consortium DHCP Server 2.0pl5
Copyright 1995, 1996, 1997, 1998, 1999 The Internet Software Consortium.
All rights reserved.

Please contribute if you find this software useful.
For info, please visit http://www.isc.org/dhcp-contrib.html

Listening on LPF/eth16/00:15:17:76:82:c3/vlan-7
Sending on   LPF/eth16/00:15:17:76:82:c3/vlan-7
Listening on LPF/eth12/00:0f:3d:ef:49:1a/vlan-6
Sending on   LPF/eth12/00:0f:3d:ef:49:1a/vlan-6
Listening on LPF/eth11/00:0f:3d:ef:49:2d/vlan-5
Sending on   LPF/eth11/00:0f:3d:ef:49:2d/vlan-5
Listening on LPF/eth10/00:0f:3d:ef:49:20/vlan-4
Sending on   LPF/eth10/00:0f:3d:ef:49:20/vlan-4
Listening on LPF/eth3/00:0e:0c:c1:a5:5b/vlan-3
Sending on   LPF/eth3/00:0e:0c:c1:a5:5b/vlan-3
Listening on LPF/eth2/00:0e:0c:c1:a5:5a/vlan-2
Sending on   LPF/eth2/00:0e:0c:c1:a5:5a/vlan-2
Listening on LPF/eth0/00:0e:0c:c1:a5:58/vlan-1
Sending on   LPF/eth0/00:0e:0c:c1:a5:58/vlan-1
Sending on   Socket/fallback/fallback-net
```

Configuration of the DHCP is then as follows:

```
default-lease-time 900;
max-lease-time 7200;
option domain-name-servers 194.141.27.25;
option domain-name "uni-plovdiv.bg";

shared-network vlan-1 { option subnet-mask 255.255.255.0;
#$ descriptoin-1
subnet 194.141.14.0 netmask 255.255.255.0 {
  option broadcast-address 194.141.14.255;
  option routers 194.141.14.254;
  option subnet-mask 255.255.255.0;
```

```
  range 194.141.14.10 194.141.14.253;}
}

shared-network vlan-2 { option subnet-mask 255.255.255.128;
#$ description-2
subnet 194.141.16.0 netmask 255.255.255.128 {
  option broadcast-address 194.141.16.127;
  option routers 194.141.16.126;
  option subnet-mask 255.255.255.128;
  range 194.141.16.1 194.141.16.125;}
}

shared-network vlan-3 { option subnet-mask 255.255.255.128;
#$ description-3
subnet 194.141.17.128 netmask 255.255.255.128 {
  option broadcast-address 194.141.17.255;
  option routers 194.141.17.254;
  option subnet-mask 255.255.255.128;
  range 194.141.17.134 194.141.17.253;}
}

shared-network vlan-4 { option subnet-mask 255.255.255.128;
#$ description-4
subnet 194.141.13.0 netmask 255.255.255.128 {
  option broadcast-address 194.141.13.127;
  option routers 194.141.13.126;
  option subnet-mask 255.255.255.128;
  range 194.141.13.25 194.141.13.125;}
}

shared-network vlan-5 { option subnet-mask 255.255.255.128;
#$ description-5
subnet 194.141.13.128 netmask 255.255.255.128 {
  option broadcast-address 194.141.13.255;
  option routers 194.141.13.254;
  option subnet-mask 255.255.255.128;
  range 194.141.13.141 194.141.13.253;}
}

shared-network vlan-6 { option subnet-mask 255.255.255.128;
#$ description-6
subnet 194.141.16.128 netmask 255.255.255.128 {
  option broadcast-address 194.141.16.255;
  option routers 194.141.16.254;
  option subnet-mask 255.255.255.128;
  range 194.141.16.129 194.141.16.253;}
}

shared-network vlan-7 { option subnet-mask 255.255.255.128;
#$ description-7
subnet 194.141.17.0 netmask 255.255.255.128 {
  option broadcast-address 194.141.17.127;
  option routers 194.141.17.126;
  option subnet-mask 255.255.255.128;
  range 194.141.17.6 194.141.17.125;}
}

host host1 { hardware ethernet 00:03:47:eb:32:29; fixed-address 194.141.x.x; }
host host2 { hardware ethernet 00:13:8f:7c:1a:fa; fixed-address 194.141.x.x; }
host host3 { hardware ethernet 00:13:1f:ca:3f:a3; fixed-address 194.141.x.x; }
```

In the case of University of Plovdiv, hosts with static IP addresses needed to be reserved. Therefore, the full range setting has not been applied. The query tool DHCPStatus [DHCPStatus] is used for obtaining overall as well as detailed configuration information of each subnet from the DHCP files.

The following shows how the addresses are assigned on the User Router:

```
eth0      Link encap:Ethernet  HWaddr 00:0E:0C:C1:A5:58
          inet addr:194.141.14.254  Bcast:194.141.14.255  Mask:255.255.255.0
          inet6 addr: 2001:4b58:27:14::1/64 Scope:Global

eth2      Link encap:Ethernet  HWaddr 00:0E:0C:C1:A5:5A
          inet addr:194.141.16.126  Bcast:194.141.16.127  Mask:255.255.255.128
          inet6 addr: 2001:4b58:27:16a::1/64 Scope:Global

eth3      Link encap:Ethernet  HWaddr 00:0E:0C:C1:A5:5B
          inet addr:194.141.17.254  Bcast:194.141.17.255  Mask:255.255.255.128
          inet6 addr: 2001:4b58:27:17b::1/64 Scope:Global

eth8      Link encap:Ethernet  HWaddr 00:30:48:74:C1:DA
          inet addr:194.141.27.177 Bcast:194.141.27.179  Mask:255.255.255.252
          inet6 addr: 2001:4b58:27::1/64 Scope:Global

eth9      Link encap:Ethernet  HWaddr 00:30:48:74:C1:DB
          inet addr:194.141.27.254  Bcast:194.141.27.255  Mask:255.255.255.192
          inet6 addr: 2001:4b58:27:192::1/64 Scope:Global

eth10     Link encap:Ethernet  HWaddr 00:0F:3D:EF:49:20
          inet addr:194.141.13.126  Bcast:194.141.13.127  Mask:255.255.255.128
          inet6 addr: 2001:4b58:27:13a::1/64 Scope:Global

eth11     Link encap:Ethernet  HWaddr 00:0F:3D:EF:49:2D
          inet addr:194.141.13.254  Bcast:194.141.13.255  Mask:255.255.255.128
          inet6 addr: 2001:4b58:27:13b::1/64 Scope:Global

eth12     Link encap:Ethernet  HWaddr 00:0F:3D:EF:49:1A
          inet addr:194.141.16.254  Bcast:194.141.16.255  Mask:255.255.255.128
          inet6 addr: 2001:4b58:27:16b::1/64 Scope:Global

eth16     Link encap:Ethernet  HWaddr 00:15:17:76:82:C3
          inet addr:194.141.17.126  Bcast:194.141.17.127  Mask:255.255.255.128
          inet6 addr: 2001:4b58:27:17ed::1/64 Scope:Global
```

2.2 Network Services

The network services mentioned in the first chapter (section 1.2) are provided in the IPv4 network. We also enabled some network services in both the IPv4 and IPv6 environments.

The network services enabled for IPv6 on the campus include DNS, web and radio services.

2.2.1 DNS

IPv6 support in the DNS is based on the ability to add an IPv6 record for hosts in the DNS and the ability to communicate with the DNS server using IPv6 transport. The record of IPv6 in the DNS is expressed in the same way as the IPv4 record; the only difference is that IPv6 uses the term “AAAA” whereas IPv4 uses “A”.

The forward DNS record of University of Plovdiv’s IPv6 network is as follows:

```
ns              IN      A       194.141.27.25
ns              IN      AAAA    2001:4b58:0027::9

dell            IN      A       194.141.27.20
torrents        IN      A       194.141.27.20
dell            IN      AAAA    2001:4b58:0027::f

koda            IN      A       194.141.27.21
koda            IN      AAAA    2001:4b58:0027::d

xeon            IN      A       194.141.27.22
xeon            IN      AAAA    2001:4b58:0027::e

tapas           IN      A       194.141.27.25
tapas           IN      AAAA    2001:4b58:0027::9

pirin           IN      A       194.141.27.27
pirin           IN      AAAA    2001:4b58:0027::27

ambit           IN      A       194.141.27.28
ambit           IN      AAAA    2001:4b58:0027::28

neko-int        IN      A       194.141.27.178
neko-int        IN      AAAA    2001:4b58:0027:0252::1

neko-ext        IN      A       194.141.27.181
neko-ext        IN      AAAA    2001:4b58:0038:0252::2
```

Whereas the reserve DNS is similar to IPv4, as is shown following:

```
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.2.5.2.0 IN      PTR     neko-tapas.uni-plovdiv.bg.
2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.2.5.2.0 IN      PTR     tapas-neko.uni-plovdiv.bg.
9.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 IN      PTR     tapas.uni-plovdiv.bg.
d.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 IN      PTR     koda.uni-plovdiv.bg.
e.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 IN      PTR     xeon.uni-plovdiv.bg.
f.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 IN      PTR     dell.uni-plovdiv.bg.
f.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 IN      PTR     forums.acad.bg.
6.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 IN      PTR     slovo.uni-plovdiv.bg.
7.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 IN      PTR     pirin.uni-plovdiv.bg.
8.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 IN      PTR     ambit.uni-plovdiv.bg.
```

2.2.2 Web

The configuration of IPv6 for web servers is a straightforward process compared with other kinds of services. Once IPv6 support is enabled / updated on the web server, no other special configuration is needed. Currently, all the services on our web servers are IPv6-enabled.

We use Apache2, where ‘Listen 80’ or ‘Listen [::]80’ can be chosen for different circumstances, either to listen on a single socket that is receiving IPv6 and IPv4 mapped connections, or to listen for IPv6 only. If needed, ‘Listen 0.0.0.0:80’ can be used to listen for IPv4 only.

2.2.3 Radio

The campus radio at the University of Plovdiv uses Icecast [ICECAST] to broadcast live Internet radio over the campus radio website [RADIONEKO]. It broadcasts under both IPv4 and IPv6 with the following configuration to the icecast.xml file:

```
    <listen-socket>
        <port>8000</port>
        <bind-address>194.141.27.21</bind-address>
    </listen-socket>

    <listen-socket>
        <bind-address>2001:4b58:27::d</bind-address>
        <port>8000</port>
    </listen-socket>
```

Users can listen to the online campus radio using a streaming multimedia player (e.g. VLC media player), without prior registration or time limitation. Live radio broadcasts organized by students are also available on a weekly schedule.

To compile performance statistics of our radio services, one of the important aspects is the number of listeners in real time. We use Cacti [CACTI] to collect and visualize the statistics such as the number of current listeners, number of listeners for the last 24 hours / last week and the number in total. The following shows examples of such statistics:

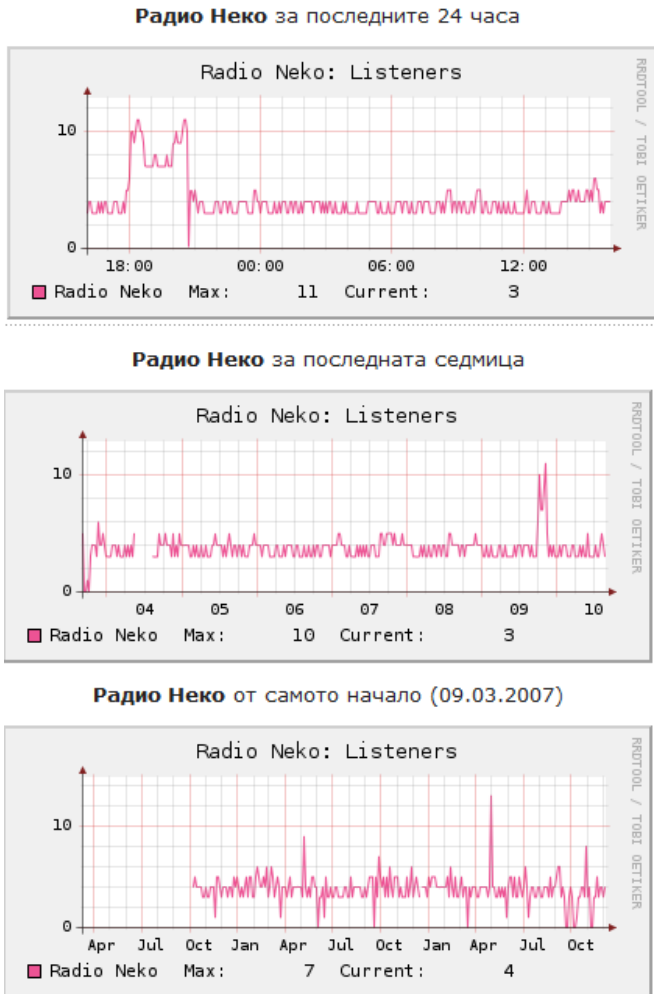

Figure 2 Campus radio live status diagram

2.3 Network Monitoring

It is important to set up a network monitoring system when IPv6 network is implemented as this can help to control both IPv4 and IPv6 networks effectively and measure the network performance in order to assess if any further modifications are required.

Since the dual-stack approach for IPv6 deployment is used, network monitoring tools were chosen that support both IPv4 and IPv6 protocols in a consistent way.

Currently, Cacti [CACTI], ntop [NTOP] and TCPdump [TCPDUMP] are used for our campus network for different purposes of network monitoring. Cacti is used to collect and visualize network traffic information, e.g. the active VLANS shown in the figure below:

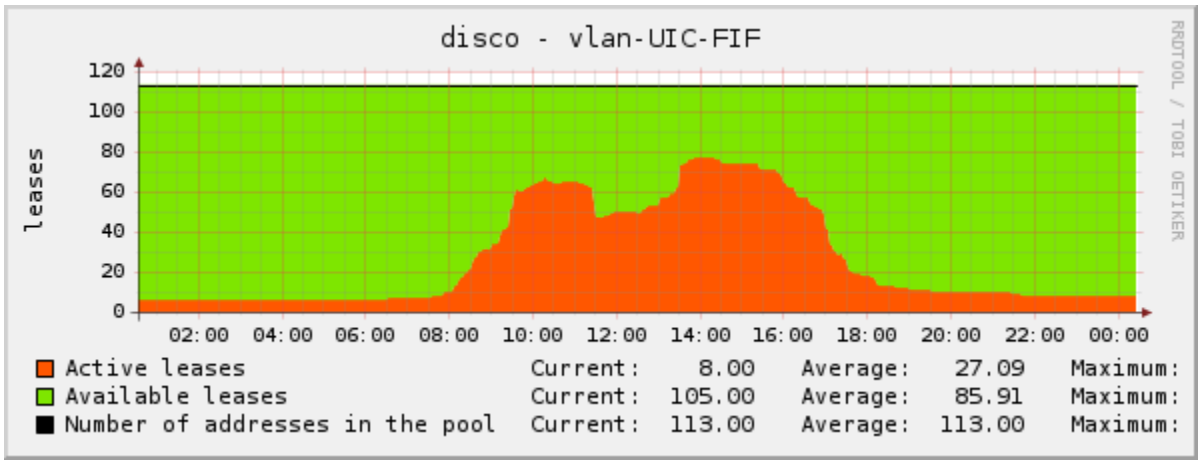

Figure 3 Subnet information collected and displayed by Cacti

Ntop allows navigation through ntop traffic information from a web browser; this also reduces CPU and memory usage in a convenient way.

Tcpdump prints out a description of the contents of packets on a network interface that match the Boolean expression and can also be run with the –w flag that makes it possible to save the packet data to a file.

The following figure shows both the general network information as well as the real-time information on each interface on the User Router. The IPv4 and IPv6 information is not separated.

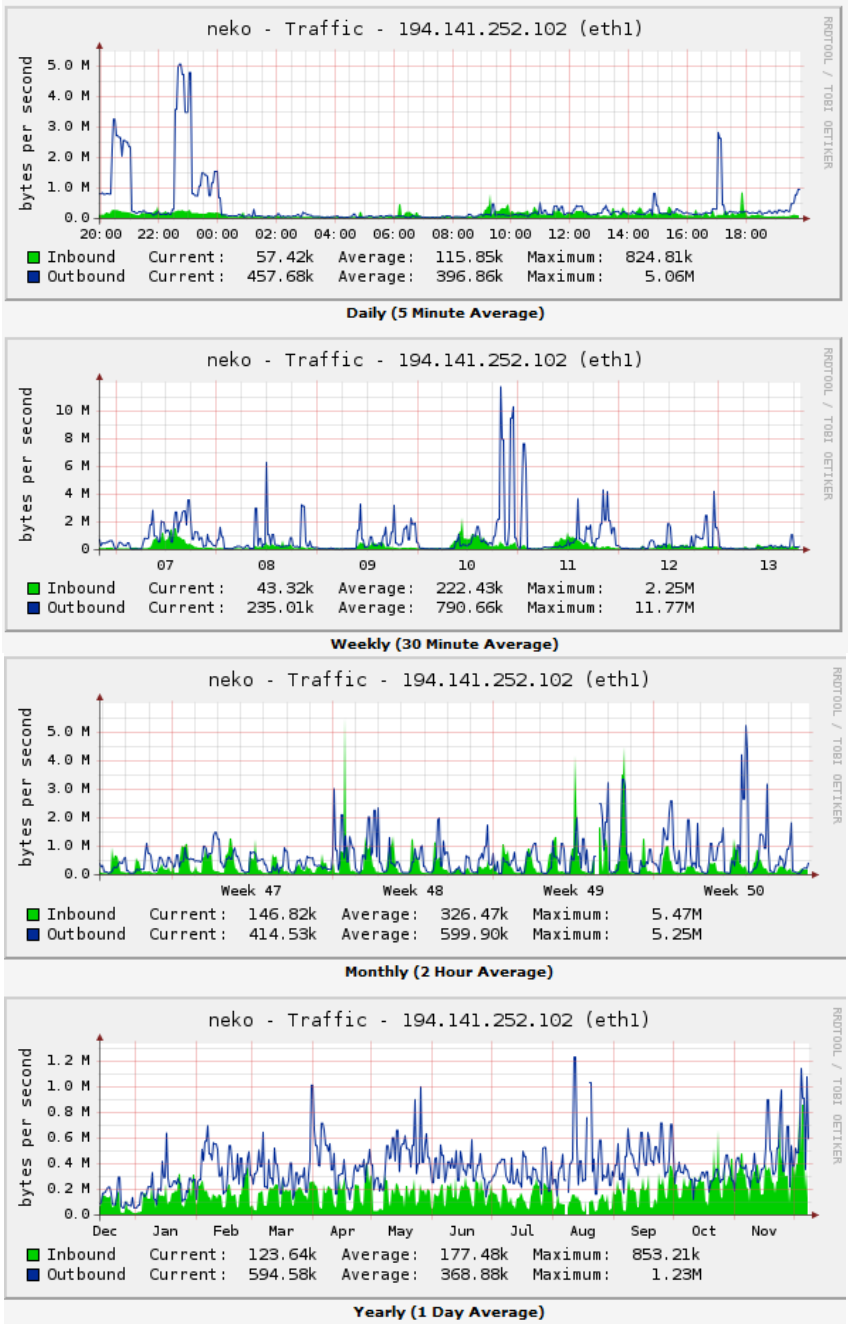

Figure 4 Network information on User Router

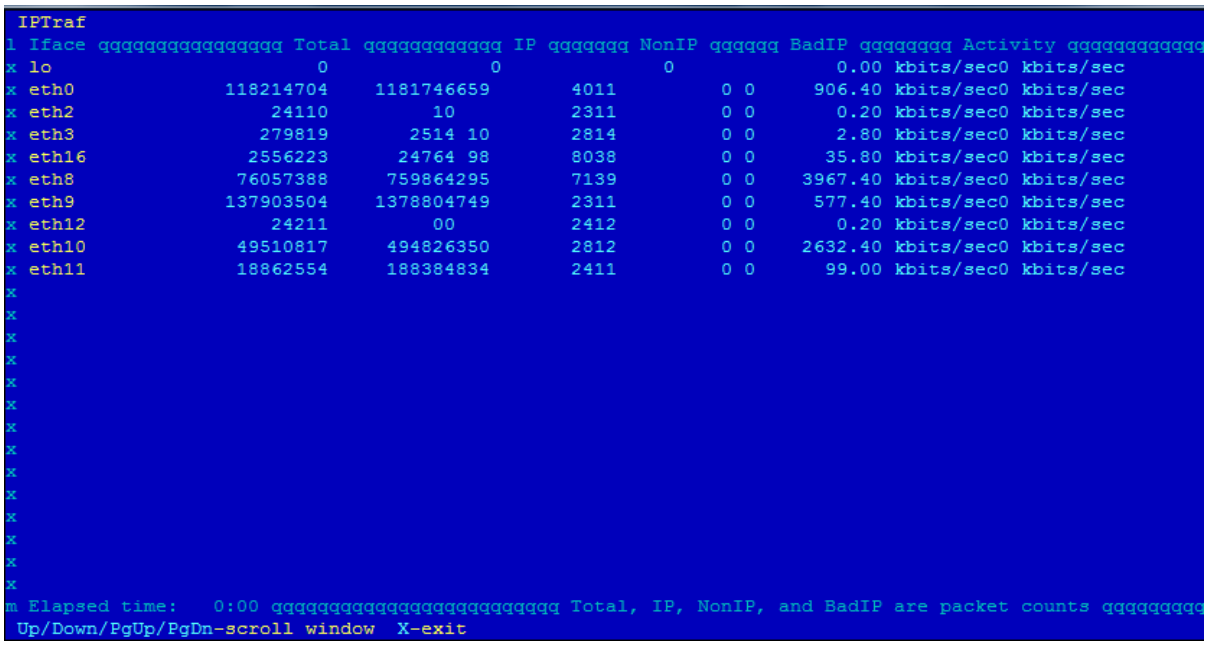

Figure 5 Real-time traffic information

Each user request for an IPv6 address is shown by the modified Router Advertisement daemon in MySQL database on a time/number of requests scale:

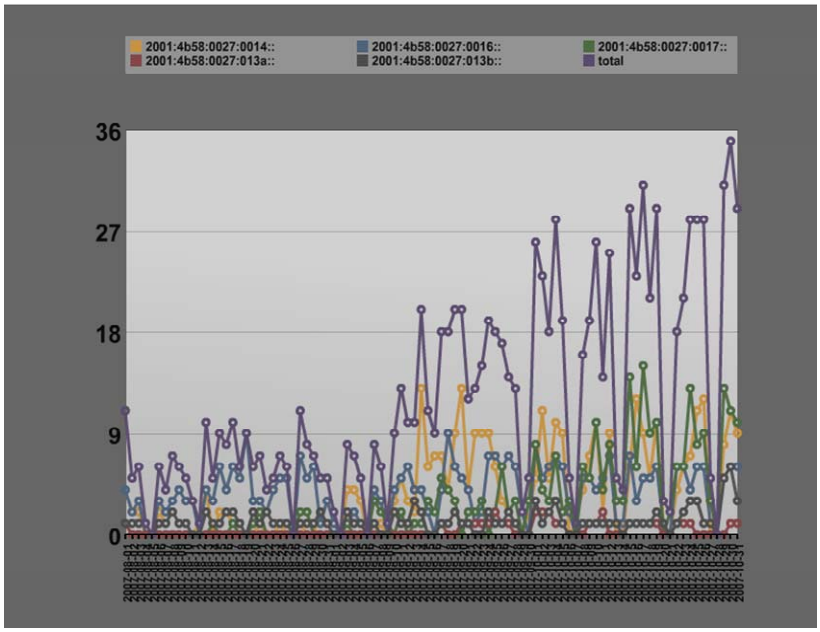

Figure 6 No. of IPv6 requests on User Router

This user request monitoring approach was used when IPv6 was initially deployed at our campus network. It is currently no longer in use due to the fact that the OS, used by the users in the subnets, i.e. Windows Vista or Windows 7, natively support IPv6. Collecting information on the number of IPv6 user requests sent from the subnets is not necessary for our network monitoring.

2.4 Security

The security rules in IPv6 must remain on an equal level with the existing IPv4 network when IPv6 is deployed. This section describes the firewall we use for the IPv6 network.

2.4.1 Firewall

The Linux iptables firewall on the User Router is used for IPv6 filtering. The outgoing traffic of all the subnets is not filtered but the incoming traffic is filtered so that the FTP, web and radio services are allowed. To have the firewall on the User Router also prevents any unauthorized access attempts to our users' computers. The firewall setting is as follows:

```
ip6tables -F
ip6tables -F INPUT
ip6tables -F OUTPUT
ip6tables -F FORWARD

# Accept DNS
ip6tables -A FORWARD -d 2001:4b58:27:252::2 -p udp --dport 53 -j ACCEPT
ip6tables -A FORWARD -d 2001:4b58:27:252::2 -p tcp --dport 53 -j ACCEPT
ip6tables -A FORWARD -d 2001:4b58:27::9 -p udp --dport 53 -j ACCEPT
ip6tables -A FORWARD -d 2001:4b58:27::9 -p tcp --dport 53 -j ACCEPT

# Accept Some User Services
ip6tables -A FORWARD -d 2001:4b58:27::/48 -p tcp -m multiport --dports
20,21,80,443,8000 -j ACCEPT

# Drop the rest
ip6tables -A FORWARD -d 2001:4b58:27::/48 -p tcp --syn -j DROP
```

2.4.2 Further security concerns

Whereas, the security technologies used in IPv4 have been tested over time and are well understood, the same level of understanding with respect to IPV6 security issues does not currently exist. In particular, in the dual-stack approach, the configurations on the routers are different and more complex than on an IPv4-only or IPv6-only router. Therefore, the complexity might also bring higher risk or opportunities for unwanted attacks.

3. PROMOTION ACTIVITIES

IPv6 was comparatively a new term to our campus users at the time that this implementation began. It has therefore been a beneficial exercise for the users to familiarize themselves with IPv6 and the advantages it can bring. A few IPv6 introduction classes have since been given to the students and the response is very encouraging. We have also added IPv6 lectures into our curriculum and provided the students with projects and theses on the same topic. This is giving more people the opportunity to know and understand the advantages and benefits of the IPv6 technology. This promotes the widespread usage of IPv6 for our future network upgrade and expansion.

4. CONCLUSIONS

The overall experience with the IPv6 deployment for our campus network has been very positive. We now have IPv6 deployed on the initial IPv4-only network and most of our network services are enabled as both IPv4 and IPv6 coexisting side by side. These include DNS, web and radio services. Our users are able to have native IPv6 connectivity with the appropriate operating systems, i.e. Windows Vista or Windows 7.

The IPv6 network is now live and there has been no drawback effect or degradation by IPv6 to the existing IPv4 network. During the process of the IPv6 deployment, our network administrators have gained valuable on-site experiences regarding IPv6 configurations and our users are becoming more familiar with using IPv6 as well.

The main challenges relating to the IPv6 deployment were in using the dual-approach which had a higher complexity for configurations compared with an IPv6-only deployment. On the other hand, we did not experience any capability problems with Operating Systems, as the mainstream OSs already support IPv6 natively.

We still need to improve the IPv6 network monitoring, i.e. to separate the IPv4 and IPv6 traffic monitoring, and it is also necessary to obtain more detailed information on IPv6 network observations. What is also needed is to set up the DHCPv6 service, so that the network administrator can statefully assign addresses for better network control over addressing.

5. REFERENCES

[BREN] Bulgarian Research and Education Network, http://bren.bg/en/

[CACTI] Cacti, http://www.cacti.net/

[DHCPSTATUS] DHCPStatus, The query tool for browsing information stored in DHCPD files, http://dhcpstatus.sourceforge.net/

[GEANT] The GÉANT network, http://www.geant.net/pages/home.aspx

[ICECAST] Icecast, http://www.icecast.org/

[ISCDHCP] ISC DHCP, https://www.isc.org/software/dhcp

[NTOP] Ntop, http://www.ntop.org/

[RADIONEKO] Radio Neko, the campus radio of University of Plovdiv, http://radio.uni-plovdiv.bg

[RFC2462] IPv6 Stateless Address Autoconfiguration, IETF RFC 2462, S. Thomson, T. Narten, December 1998

[TCPDUMP] TCPDump, http://www.tcpdump.org/