Project Number: IST-2000-26417
Project Title: GN1 (GÉANT)

Deliverable D9.6

Report on IPv6 Experiments and Status

Deliverable Type: PU-Public
Contractual Date: 31 August 2002
Actual Date: 11 September 2002
Work Package: WI8.4
Nature of Deliverable: RE - Report

Authors:

Participants in the GÉANT TF-NGN IPv6 Working Group including representatives from Renater, DFN (including WWU-JOIN), UKERNA, SWITCH, CESNET, CERN, POZNAN, RedIRIS, HEAnet, DANTE, RESTENA, UNINETT, ACOnet, INFN-GARR, ULB, ENST, Hungarnet and GRNet. Report edited by Tim Chown (University of Southampton and UKERNA)

Abstract:

In this deliverable we report on the IPv6 activities and experiments undertaken by the participants in the IPv6 Working Group within the GÉANT Task Force for Next Generation Networks (TF-NGN). These experiments focus on three main areas, namely maintenance of an IPv6 testbed centered on a Juniper M5 router operated by Renater in Paris, secondly multicast IPv6 experiments using the PIM-SM protocol and finally evaluation of the Zebra software router. We also describe the status of IPv6 standards development and of IPv6 implementations in hosts and routers, as an assessment of the readiness of IPv6 for production deployment. We summarise the work being undertaken by the NRENs in 6NET, for which there is some overlap with TF-NGN activity. Finally, we describe initial plans for the migration of GEANT to support a dual-stack IPv4-IPv6 service, and identify potential future work items for the IPv6 Working Group.

Keywords:

IPv6, IPv6 multicast, IPv6 transition, GTPv6 testbed, 6NET, IPv6 software routers

Table of Contents

Table of Figures

Executive Summary

In this document we report on the IPv6 work undertaken by the IPv6 Working Group of GÉANT's Task Force Next Generation (TF-NGN).

To some extent, the scope of the work has been reduced by the emergence of the EU-funded 6NET project [6net] since January 2002. Many of the previous TF-NGN IPv6 study items are now explicitly covered in 6NET, where the resources available for the work to be done are much greater. The experiences gained within 6NET will be important for the deployment of IPv6 on GÉANT and therefore there will be continuous feedback between 6NET and TF-NGN. However, 6NET does not cover all IPV6 aspects that are important to GÉANT and the NRENS and there are still many NRENs (over 50%) who are not in 6NET, and for whom the TF-NGN IPv6 WG is a useful forum to exchange ideas and to also gain some feedback from the 6NET partners who attend.

The GTPv6 (GÉANT Test Programme for IPv6) testbed network has for the last few months been focused on a Juniper M5 router running JUNOS, taking address space from the Renater production IPv6 prefix due to its Paris location. Up to ten TF-NGN IPv6 participant NRENs have connected to this router, and there have been no significant interoperability issues reported with Cisco IOS or Zebra router platforms. In the near future a Hitachi GR2000 should be added to the testbed, forming a mini-backbone, at which point the original 6Bone prefix and available ASN may be brought back into use.

A tunnelled IPv6 multicast testbed, known as the M6Bone [m6bone], has also been created, extending to over 20 sites in France, Europe and beyond. The M6Bone has been used for IPv6 videoconferencing (using the IPv6 versions of the popular vic and rat tools) and also for multicast seminar distribution. Future work has been identified in this area, including studies of PIM-SSM, IPv4-IPv6 multicast gateways, and multicast scope.

In CESNET, studies of the performance and reliability of PC-based router platforms have been carried out, in particular with the Zebra router daemon [zebra]. As a backbone router for an early stage of a parallel IPv6 deployment, or as an edge router for a university undertaking initial IPv6 experiments, Zebra appears a promising choice.

Perhaps most important in the next 12 months is the planning for the migration of GÉANT to run in dual-stack IPv4-IPv6 mode. While some lessons can be drawn from SURFnet (who are running dual-stack with Cisco hardware) and Abilene (dual-stack with Cisco now but with Juniper by late 2002), a well-considered plan needs to be drawn up for the introduction of the service into the backbone, and for the deployment of associated support services. Collaboration with networks and organisations such as Internet2 and WIDE will be of benefit in this task.

As reshuffles in the IETF WGs (e.g. from ngtrans to v6ops) reflect the more mainstream and maturing status of IPv6, it is important that the TF-NGN IPv6 WG keeps abreast of the relevant standards and policies, determining best practice for transition and deployment, as well as identifying opportunities for deployment of novel IPv6 services and for methods to bring more users and content to the growing IPv6 world.

1 THE GTPV6 TESTBED NETWORK

The GÉANT Task Force Next Generation (TF-NGN) IPv6 WG [tf-ngn] has run a testbed network for its participants for a number of years, under the banner of the GÉANT Test Programme for IPv6 (GTPv6) [gtpv6]. Up until 2001, this network was centred on a Telebit TBC2000 router, as reported in GÉANT Deliverable D9.3 [geant-d93]. At that time the testbed was focused on ATM technology, such that approximately half of the participants were connected to the TBC2000 over ATM PVCs, using the TEN-155 ATM infrastructure, with the remaining participants connecting via IPv6-in-IPv4 tunnels. The TEN-155 network has now been replaced by the new GÉANT PoS backbone, which is a significant upgrade for the production IPv4 network, but which removes the ability for ATM PVCs to be used as a "native" IPv6 connectivity method.

In this section we review work undertaken in deploying a new testbed network centred on a Juniper M5 router operated by Renater in Paris, and we present in some detail comments on results of trials with PC-based router platforms within CESNET, in particular with the Zebra routing daemon [zebra].

1.1 Routing, topology and addressing

As a result of the progression from ATM, and in the current absence of a deployable and robust method to carry native IPv6 traffic on an end to end basis across GÉANT and NREN networks, the GTPv6 testbed activity has to now be largely tunnelled. In addition to the infrastructure changes, the GTPv6 group has deployed a Juniper M5 router in Paris. This router runs the latest JUNOS software (JUNOS 5.2). One of the goals of the scheduled IPv6 testing was to report on the (production usability) status of IPv6 on the Juniper platform, and to check for interoperability issues with other platforms. Such testing will help in gaining early experience with IPv6 on JUNOS with the migration of GÉANT (which runs on Juniper routers) to dual-stack IPv4 and IPv6 operation in mind.

The GTPv6 group still holds an IPv6 6Bone prefix, inherited from the original QTPv6 testbed (under the Quantum Test Programme and TF-TANT). The prefix, allocated to the entity 'QTPVSIX', is 3ffe:8030::/28. In the Telebit-centred testbed, each participant was allocated a /34 prefix under that top level aggregrate. The testbed also used its own ASN (AS8933), which is held by DANTE.

In evolving from running a testbed using the Telebit to the Juniper M5 testbed, the group chose to run a network with address allocations made from Renater's production RIPE NCC-allocated IPv6 address space (using also Renater's ASN, AS2200), the current allocations as shown in Figure 1.

Partner	IPv6 Prefix	IPv4 tunnel end point
Renater (FR)	2001:0660:1102:4001::/64	Native
UKERNA (UK)	*2001:0660:1102:4002::/64*	-
* CERN (CH)	2001:0660:1102:4003::/64	192.65.185.7
GRnet(GR)	*2001:0660:1102:4004::/64*	-
DANTE (UK)	2001:0660:1102:4005::/64	193.63.211.61
POZNAN (PL)	2001:0660:1102:4006::/64	150.254.210.109
* RedIRIS (ES)	2001:0660:1102:4007::/64	130.205.0.130
* HEAnet (IE)	2001:0660:1102:4008::/64	193.1.195.61
SWITCH (CH)	2001:0660:1102:4009::/64	130.59.32.38
CESNET (CZ)	2001:0660:1102:400a::/64	195.113.156.183
* RESTENA (LU)	2001:0660:1102:400b::/64	158.64.16.21

Figure 1: Address allocations and tunnels to Juniper M5 testbed participants

Those participants in italics are currently not connected to the network. Those highlighted by an asterisk (*) are not members of the 6NET project; this shows the value of the testbed within the GÉANT community given approximately half the participants have no presence in the 6NET project (GTPv6 offers a forum for non-6NET NRENs, and allows cross-fertilisation of work and ideas between the NRENs).

The topology is, as it was under the previous testbed, a simple star topology with all participants directly connected via one hop to the central router. The router interfaces on the testbed have a ::1 postfix at the M5 end of the link, and a ::2 postfix at the partner's end of the link.

1.2 Interoperability

One of the goals of the Juniper testbed was to identify interoperability issues between JUNOS and other IPv6 router implementations. None of the other partners connected using Juniper routers; only the core router ran JUNOS.

The router platforms that were primarily tested were Cisco IOS and the open source Zebra platform [zebra] (which can run for example on Linux or FreeBSD). There is a commercial instance of Zebra, developed by IP Infusion, called ZebOS; the group hopes to trial this platform in the near future, either within GTPv6 or as part of the 6NET project.

An example of one of the partners running Cisco IOS is CERN. CERN's IPv6 router runs IOS 12.2, and has a large number of external BGP4+ peerings, one of which is to the GTPv6 Juniper router [cern-v6]. The CERN configuration includes some extensive controls on transit and prefix-based filtering, which is important where multiple BGP peerings exist. The connection from CERN to the Juniper is tunnelled IPv6-in-IPv4.

CESNET runs an IPv6 router using Zebra. The configuration of this router is discussed in more detail in the section below.

To date the group has not experienced any significant interoperability problems between JUNOS and IOS or between JUNOS and Zebra for regular BGP4+ peering sessions.

1.3 Case study: CESNET

In this section the IPv6 router deployments for the GTPv6 project at CESNET are detailed. CESNET joined the GTPv6 network in January 2002 by connecting to the Juniper M5 router in Paris. In April 2002 a second international connection was established to POZNAN (Poland). Both connections are configured tunnels (GRE and IPv6 over IPv4, respectively) terminated at the same router *prg-v6-gw.cesnet.cz*. This router is a PC with the Linux operating system and USAGI patches for IPv6, running the Zebra routing daemons [zebra], version 0.93a. It is connected with a Gigabit Ethernet interface directly to the CESNET access router to GÉANT.

1.3.1 Exterior routing policy

As was mentioned above, CESNET (AS 2852) has 2 EBGP neighbours in GTPv6:

- 2001:660:1102:400a::1 in AS 2200 (Renater, France)
- 2001:718:0:8000::2 in AS 9112 (PSNC, Poland)

The first one is considered primary and should normally be used for most traffic. Therefore, the following routing policy was configured:

1. For prefixes originated in AS 9112, the direct route to AS 9112 will be preferred.
2. For routes whose AS Path starts with "2200 9112" (i.e., traffic that is routed from AS 2200 to 9112) the shortcut route to AS 9112 will be preferred.
3. In all other cases, the route to AS 2200 will be preferred.

The router *prg-v6-gw.cesnet.cz* itself originates just a single summarized prefix 2001:718::/35. BGP prefixes obtained from one neighbour are all passed to the other one, that is, no filtering or AS Path modification beyond prepending a single 2852 is performed. This also means that CESNET provides full transit in both directions.

1.3.2 CESNET IPv6 backbone

The CESNET IPv6 backbone network currently interconnects seven major CESNET PoPs. It uses dedicated IPv6 routers interconnected by configured tunnels over the production IPv4 backbone. The topology of the IPv6 backbone, shown in Figure 2, nonetheless tries to match the underlying physical topology as closely as possible. In most cases the IPv6 routers are directly connected to the co-located IPv4 backbone router.

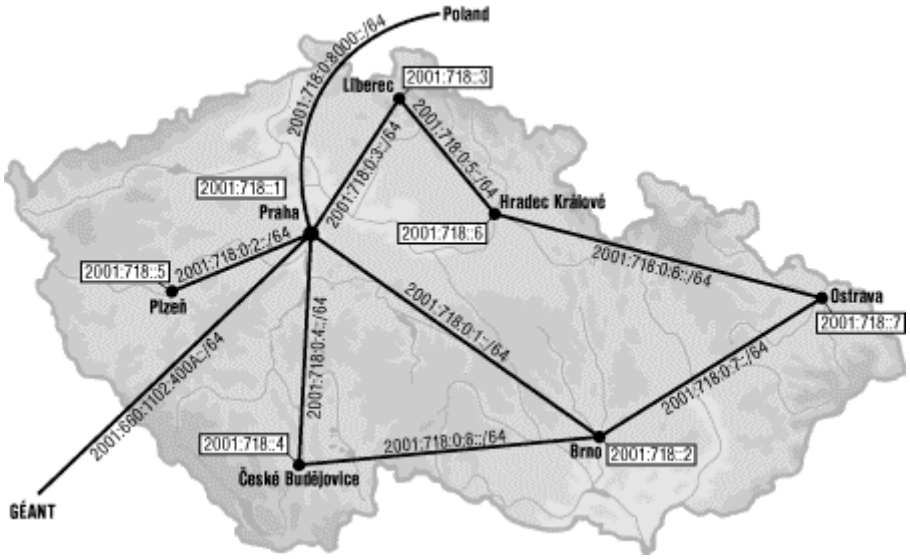

Figure 2: CESNET IPv6 backbone

This way, a future transition to an integrated IPv4/IPv6 network will just amount to combining the IPv4 and IPv6 backbone routers into single dual-stack boxes without any significant changes to the topology of either network.

Virtually all IPv6 enabled routers and hosts in CESNET now use the IPv6 addresses from the block 2001:718::/35 allocated to CESNET by RIPE NCC. The 6Bone prefix 3ffe:803d::/34, which was used in the earlier GTPv6 network in the days of TEN-155, has been phased out and is not supported any more.

Five of the IPv6 backbone routers are PCs with either Linux/USAGI or NetBSD/KAME and Zebra routing daemons. The remaining two routers are Cisco 2600 (Liberec) and 3640 (Hradec Králové). In this mixed environment, the choice of an interior gateway protocol was limited to RIPng, since this is the only IGP supported by both platforms so far. RIPng is used together with IBGP so that the only prefixes propagated in RIPng are those of backbone links and router loopback interfaces. The latter are used as next hops in IBGP so that if any backbone link fails, the traffic will be automatically rerouted (if possible). IBGP is then used in a configuration with two route reflectors located in Praha and Brno.

Another issue of the mixed Cisco/PC router environment was a unified authentication mechanism for controlling access to both platforms. CESNET uses routinely TACACS+ on its routers so the most straightforward option was to enable TACACS+ authentication on PC routers, too. Linux has a flexible authentication mechanism known as PAM (Pluggable Authentication Modules) and a module for TACACS+ was already available. However, some modifications of the existing software were necessary:

- The original *pam_tacacs* module was only capable of authenticating PPP connections to a terminal server. We thus enhanced it to support arbitrary authentications (e.g., SSH). Another modification was also applied in order to be able to combine TACACS+ authentication with other methods in a single box – for example, we don't want the *root* user to authenticate against TACACS+ but rather use the local databases. Our solution is to apply TACACS+ authentication only for users with UID greater than some threshold (which can be supplied as a parameter to the *pam_tacacs* module).

- From the Zebra suite only the integrated shell **vtysh** was capable of PAM authentication. We prefer interacting with individual daemons directly, though, and so we added PAM authentication separately to each Zebra daemon.

1.3.3 Experience with PC routers

In the position of AS border router, as well as in five backbone PoPs, CESNET uses PC routers with either Linux/USAGI or NetBSD/KAME operating system and the Zebra routing protocol suite. We believe this platform is a viable alternative to specialized IPv6 routers, particularly in research networks. It is thus worthwhile to test their interoperability in multi-platform networks.

In terms of performance, there is little doubt that high-end commodity PCs are able to support expected volumes of IPv6 traffic for the next few year. In fact, laboratory tests of IPv4 forwarding performance over Gigabit Ethernet interfaces indicate – see Figure 3 – that the software-based PC router delivers throughput by no means worse than the Cisco 7500 series router, although the latter shows better resilience under excessive load. For a traffic mix whose distribution of packet sizes roughly corresponds to a typical pattern as observed in the CESNET backbone (30% large packets, 30% small packets and the rest uniformly distributed between the two extremes), the maximum load is about 350 Mbit/s.

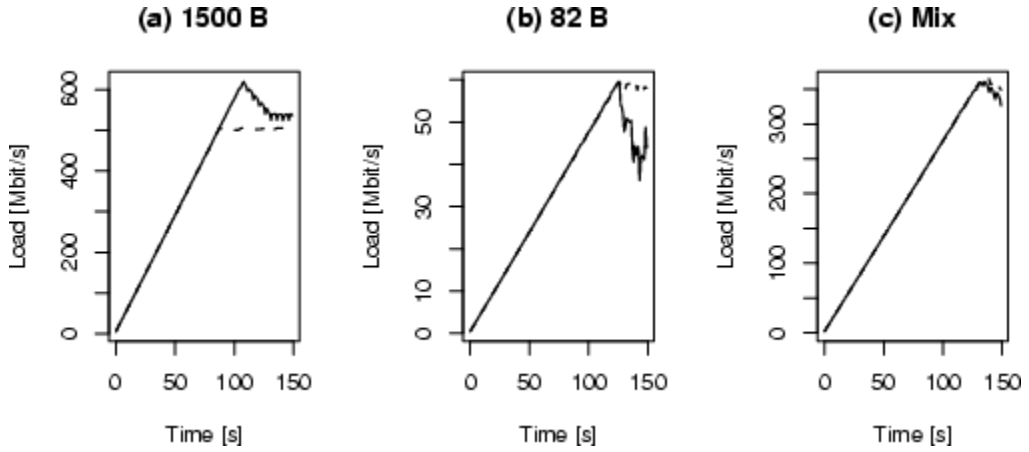

Figure 3: Throughput of a Linux PC router (solid line) and Cisco 7500 Series router (dashed line)

One can certainly expect slightly worse figures for IPv6 since the IPv6 stack is probably not so well tuned yet as its IPv4 counterpart, yet even then the throughput should be sufficient for most real life purposes, at least while IPv6 is in its early deployment stage and low cost, dedicated routers are required for IPv6.

The single most problematic aspect of PC routers is their configuration. Every leading router platform provides the network administrator with a consolidated command line interface where the complete set of necessary parameters, functions and protocols can be configured. Besides of this fundamental text-based interface, additional options might be available, for example web-based interface, SNMP, XML-RPC and so on.

In Unix-based operating systems, any comprehensive network configuration and monitoring (including tunnels, routing, packet filters, QoS etc.) typically involves a number of commands, scripts and files, which moreover differ from one operating system to another, sometimes even among different flavours of the same system. While this is mostly acceptable for a seasoned Unix administrator, it could be very confusing for a general, albeit experienced, network administrator.

In 2002, CESNET started a project *"IPv6 PC routers"* with two specific goals:

1. Design and build a hardware forwarding accelerator that would increase the throughput of PC routers to the order of several Gbit/s.

2. Create a unified and flexible configuration layer based on XML, allowing different configuration front-ends and able to transform the general configuration not only to Unix scripts, but also to other configuration languages (Cisco IOS, Juniper JUNOS).

1.3.4 The Zebra routing daemon

In the past, many Unix-based PC routers were equipped with the GateD routing daemon [gated]. Recently, the Zebra routing suite became a serious contender in most areas of IPv4/v6 routing, the only one notable exception being multicast routing. The advantages of Zebra in comparison to GateD are:

- Better overall software architecture, which results in much easier maintainability and portability. As opposed to the monolithic design of GateD, Zebra is modular – each routing protocol is served by a separate daemon process.

- Free software license – GNU GPL. GateD is now distributed under rather restrictive licensing terms, which not only make it difficult to obtain and deploy but also limit the potential for further software development. On the other hand, Zebra benefits from the well-established open source software model – contributions from many developers and extensive testing.

The commercial ZebOS package [zebos] is based on the Zebra package, and includes a full spectrum of the state-of-the-art routing technology including multicast routing and MPLS.

Apart from GateD and Zebra, there are few other packages, e.g., MRTD or BIRD. In their current though, they still lack a number of necessary features and/or stability.

The command line interface of Zebra is extremely similar to Cisco IOS, except that the network administrator must interact with each of its daemons separately. This usually involves starting several telnet sessions, one for each daemon in use. The Zebra suite contains a program named **vtysh**, which aims at integrating all daemons into one command line shell, but we found it rather problematic and buggy.

In the following subsections we summarize the interesting points, problems and caveats, which we encountered during the configuration of the *prg-v6-gw* router for both intra- and inter-AS routing.

1.3.5 Tunnels

Tunnels are configured in Linux using the **ip** command, for example

```
ip tunnel add M5Paris mode gre local 195.113.156.183 \
      remote 193.51.207.243 ttl 64
ip tunnel add Poznan mode sit local 195.113.156.183 \
      remote 150.254.210.109 ttl 64
```

The first tunnel is assigned the name *M5Paris* and is of type GRE (Generic Route Encapsulation), while the second one - *Poznan* - is a configured IPv6 over IPv4 tunnel. In Linux it is denoted as a SIT (Simple Internet Transition) tunnel.

The only tricky part here was the "`ttl 64`" option. By default, Linux tunnel packets inherit the TTL (time-to-live) from the inner (tunnelled) packets. Since standard EBGP neighbours should be directly connected, the BGP packets are sent with hop limit (the equivalent of TTL in IPv6) set to 1. Consequently, the TTL value of the outer IPv4 tunnel packet is 1 as well and the packet gets dropped at the first hop along the tunnel route. In order to avoid this undesirable behaviour, we thus have to specify an explicit TTL value for the outer tunnel packets.

1.3.6 Loopback interfaces

It is often useful to set up one or more *loopback interfaces* and assign them IP addresses with full masks (/32 for IPv4 and /128 for IPv6). Unlike physical interfaces, these virtual interfaces are usually up as long as the router is operational. Therefore, addresses of loopback interfaces are good candidates for router IDs, next hops in BGP etc. Unix-like operating systems usually have only one such interface (its name in Linux is *lo*), which is automatically assigned addresses 127.0.0.1 (IPv4) and ::1 (IPv6). In principle, it is possible to add other IPv6 addresses to it, yet we think it is much cleaner to use *dummy* interfaces as router loopbacks. If configured as a kernel module, the dummy interface driver may be linked to the kernel several times, thus providing interfaces *dummy0*, *dummy1* etc. Unfortunately, we found IPv6 addresses of *dummyN* (as well as *lo*) interfaces cannot be advertised in RIPng by using the most natural configuration, for example,

```
interface dummy0
 ipv6 address 2001:718::1/128
...
router ripng
 network dummy0
```

This will not insert the required prefix 2001:718::1/128 into RIPng advertisements in the same way as for physical interfaces and tunnels. Instead, we have to use the following configuration:

```
interface dummy0
ipv6 address 2001:718::1/128
...
router ripng
 network Liberec
 network Brno
 ...
 redistribute connected
 distribute-list prefix bb-only out Liberec
 distribute-list prefix bb-only out Brno
 ...
ipv6 prefix-list bb-only seq 5 permit 2001:718::/48
```

The `distribute-list` commands are usually necessary in order to eliminate advertisements of other locally connected interfaces. Also note that this command must be used separately for each of the RIPng interfaces (in our case the named backbone tunnels Liberec, Brno etc.), since after entering the command without the interface name the **ripngd** daemon crashes.

1.3.7 Router advertisements

By default, the **zebra** daemon does not send RA messages out of IPv6 interfaces. It is the desired behaviour for point-to-point links between routers, but for connected local area networks where we want the stateless configuration of hosts to take place, we have to use a configuration as in the the following example:

```
interface eth0
 ip address 195.113.134.217/25
 ipv6 address 2001:718:1:86:240:c7ff:fe97:75d9/64
 no ipv6 nd suppress-ra
 ipv6 nd prefix-advertisement 2001:718:1:86::/64
```

1.3.8 Border Gateway Protocol (BGP)

Configuration of the Zebra **bgpd** daemon can be most of the times copied verbatim from Cisco routers. Thus, the CESNET EBGP routing policy defined above was configured as follows:

```
router bgp 2852
 bgp router-id 195.113.156.183
 no bgp default ipv4-unicast
 ...
 neighbor 2001:660:1102:400a::1 remote-as 2200
 neighbor 2001:660:1102:400a::1 description Renater
 neighbor 2001:718:0:8000::2 remote-as 9112
 neighbor 2001:718:0:8000::2 description Poznan
 ...
!
 address-family ipv6
 network 2001:718::/35
 network 2001:718::/42
 neighbor 2001:660:1102:400a::1 activate
 neighbor 2001:660:1102:400a::1 prefix-list agregus out
 neighbor 2001:660:1102:400a::1 route-map renater-set-locpref in
 neighbor 2001:718:0:8000::2 activate
 neighbor 2001:718:0:8000::2 prefix-list agregus out
 neighbor 2001:718:0:8000::2 route-map poznan-set-locpref in
 ...
 exit-address-family
!
ipv6 prefix-list agregus seq 5 deny 2001:718::/35 ge 36
ipv6 prefix-list agregus seq 10 permit any
...
!
ip as-path access-list 1 permit ^9112$
ip as-path access-list 2 permit ^2200 9112
!
route-map renater-set-locpref permit 10
 match as-path 2
 set local-preference 100
!
route-map renater-set-locpref permit 20
 set local-preference 200
!
route-map poznan-set-locpref permit 10
 match as-path 1
 set local-preference 200
!
route-map poznan-set-locpref permit 20
 set local-preference 100
!
...
```

Note that the 2001:718:: prefix is specified in two network statements with lengths /35 and /42, respectively. The first one is the summarized prefix of the entire CESNET IPv6 network that is

advertised to the EBGP neighbours while the other one is the prefix of the Prague PoP only which must be advertised to the IBGP neighbours. The 2001:718::/42 prefix is filtered out from EBGP advertisements by the `agregus` prefix list. The relative local preferences of different prefixes with respect to their AS Path are enforced by the route maps `renater-set-locpref` and `poznan-set-locpref`.

1.3.9 Summary

Experiments undertaken within CESNET show that PC routers, in particular Zebra, have good potential for the early deployment phase of IPv6, with performance comparing well to mid-range Cisco products such as the 7500 series. In the CESNET experiments, Gigabit Ethernet interfaces were tested. Configuration is made familiar through provision of IOS-like commands. Further experiments will be run with the ZebOS package, which is hoped to be available from IP Infusion for testing shortly.

1.4 Future GTPv6 network deployment

The M5 router used for the current GTPv6 testbed network also forms part of Renater's PlaGE testbed, which includes high-end Cisco routers. Renater has chosen to keep some ATM provision to enable IPv6 native networking in its national IPv6 testbed. The Juniper configuration used in the testbed is listed in Annex A, and the topology and connectivity is shown in Figure 4.

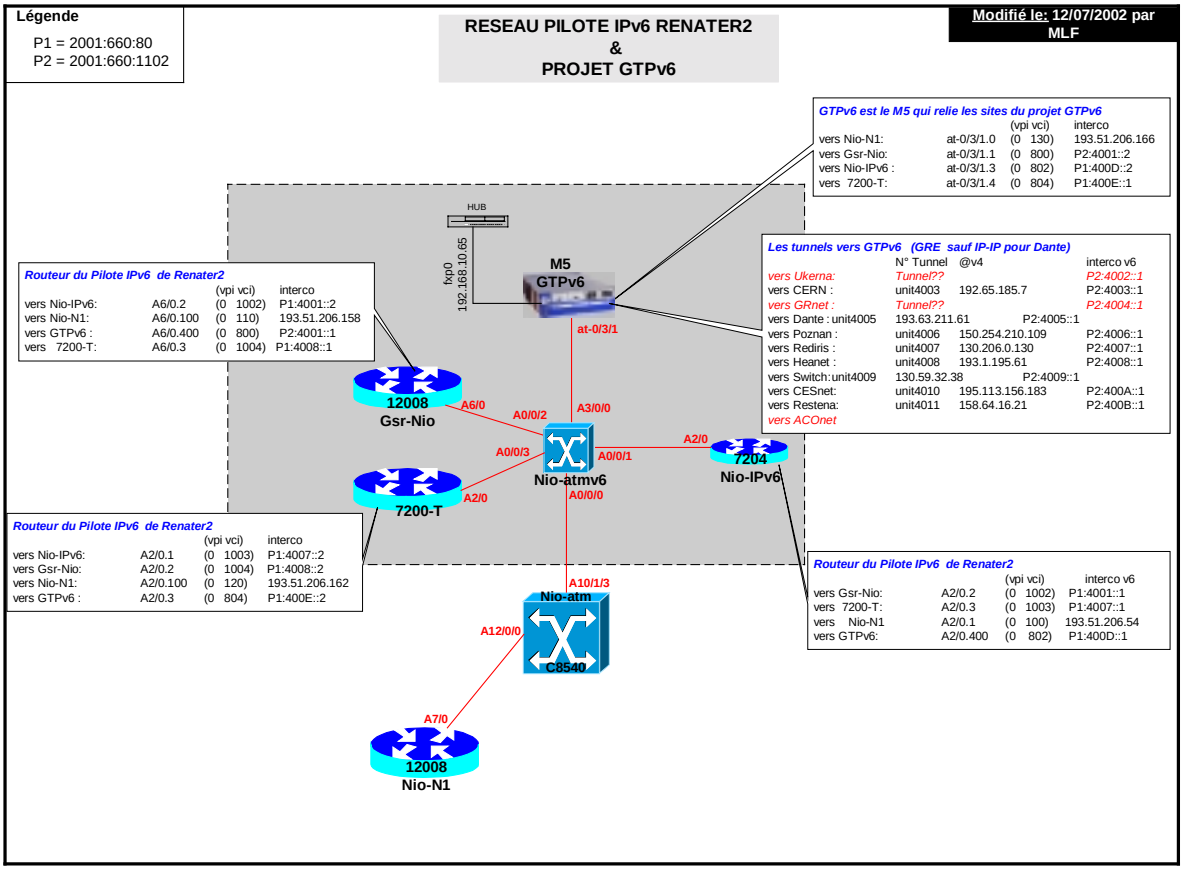

Figure 4: Renater IPv6 testbed showing GTPv6 connectivity (French language)

After discussion within the GTPv6 group, it was felt that while the group had gained a lit of valuable experience with the Telebit TBC2000, the lack of a clear future product path from Telebit, and the TBC2000's general status as an edge device rather than a backbone router, meant that the focus for experiments should be on the Juniper.

In the next phase of the GTPv6 development, it is planned that the use of the QTPVSIX 6Bone prefix will be restored (3ffe:8030::/28), using the available DANTE ASN from the previous testbed work (AS8933). In the initial phase of this development, only the Juniper router will be used as the origin of this network. During this phase it is likely that the Renater prefix will not be used (connecting sites will probably wish to use the 6Bone address space (a /34) allocated from under the QTPVSIX prefix, or, subject to agreement, they may use their own allocated 6Bone or production prefix).

In the secondary phase of this GTPv6 development, the Hitachi GR2000 at Southampton (within the UKERNA network) will be incorporated into the GTPv6 testbed, as illustrated in Figure 5. In this configuration we will be able to

- offer connectivity to different NRENs to different routers (the Hitachi H-GTPv6 or the Juniper J-GTPv6)
- run an iBGP between the two routers (to be identified), and test for interoperability

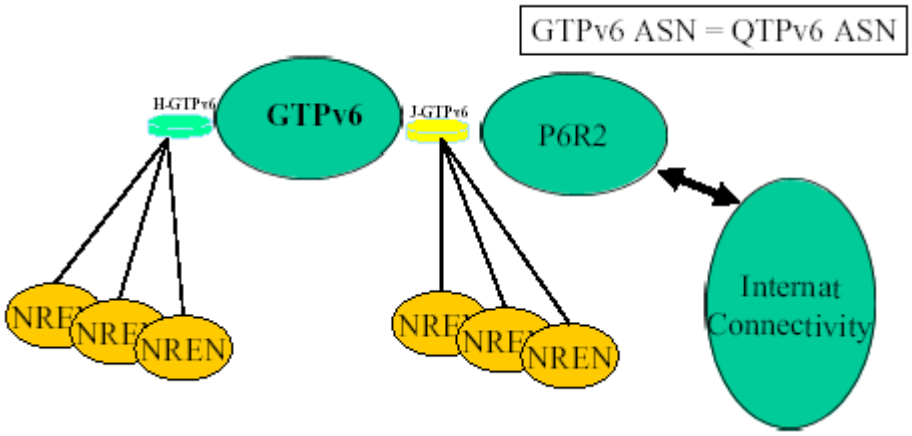

Figure 5: Planned next step for the GTPv6 testbed

The GR2000 has support for IPv6 PIM-SM, which is discussed in the IPv6 multicast chapter below.

1.5 Other TF-NGN IPv6 (GTPv6) activities

In addition to the testbed experiments partners within the TF-NGN IPv6 WG have also reported on some of their own initiatives at the TF-NGN meetings [tf-meets]. The TF-NGN group includes many more NRENs than take part in the 6NET project (although there is some overlap).

Examples of such work include:

- National IPv6 pilot deployments – representatives from NRENs present and discuss their internal IPv6 pilot work, which primarily focuses on the IPv6 infrastructure, whether dual stack or – as is more common at present - as a standalone parallel network. An excellent example of such work lies with the IPv6 deployment in Poland (POZNAN), as illustrated in Figure 6. One of the particularly interesting plans for the growth of this network is to use a separate lambda on the PIONIER Polish optical network for IPv6 traffic. POZNAN is able to offer some native IPv6 links through existing ATM PVCs. One of the POZNAN customers is a local cable TV provider (ICPNET), who holds an IPv6 SubTLA.

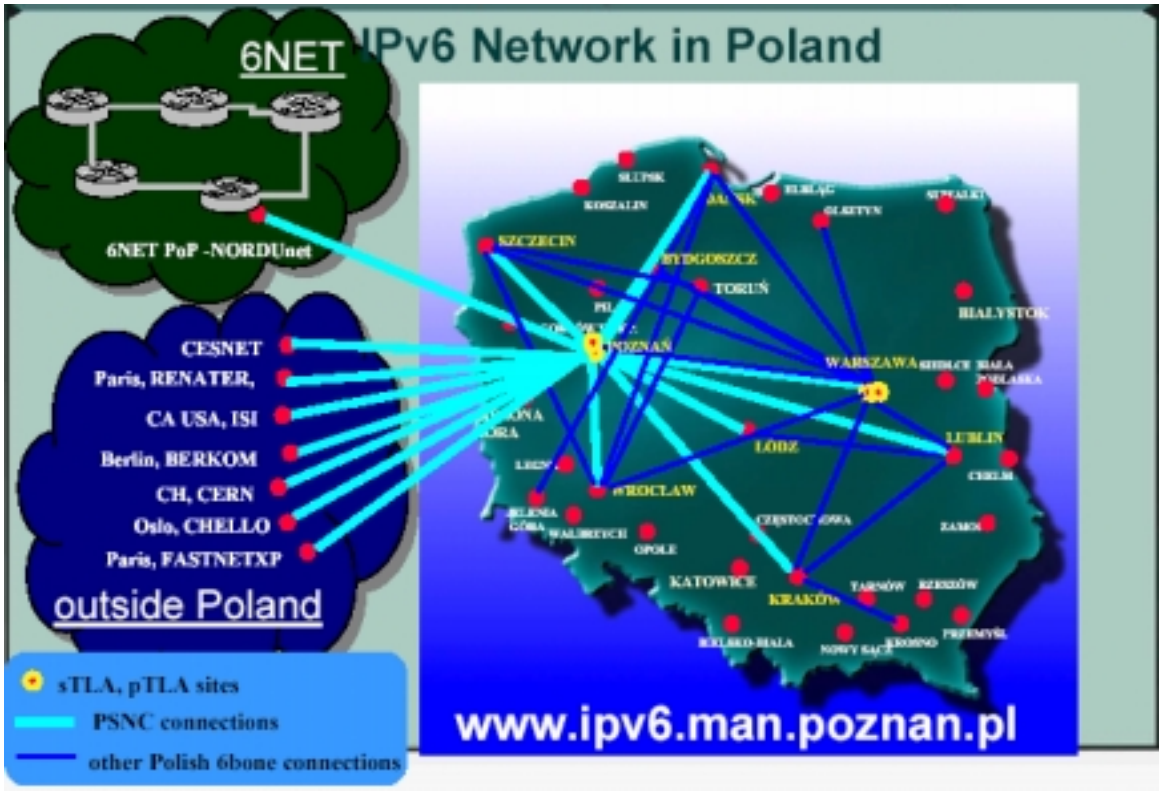

Figure 6: The POZNAN IPv6 network and connectivity

In addition to providing network infrastructure, the Poles have also tested many services including DNS (BIND9) and the Apache 2 web server. IPv6 network monitoring tools include ASpathtree, a looking glass, ping and traceroute reports and MRTG visualisations [poznan-v6].

Examples of other NREN testbeds, including those in Renater and the DFN (through the JOIN project) can be found in the presentations section of the TF-NGN Meeting site [tf-meets]. The NGNLab testbed [ngnlab] connecting ULB and Switzerland has also been presented and discussed within TF-NGN.

- DSTM – the IPv6 team at ENST has a Dual Stack Transition Mechanism (DSTM) implementation [dstm-imp] based on their own IETF ngtrans standards work [dstm]. While the IETF ngtrans WG is currently in a state of flux, the group feels that DSTM is an important transition technique that should be progressed.

- Firewalls – work by Hungarnet and DANTE has reported on the status of IPv6 firewalls, including comments on ipfilter, ip6fw, netfilter and Cisco ACLs. At present there is no robust, complete IPv6 firewall implementation (e.g. that allows processing of all IPv6 Headers), although CheckPoint has recently announced some IPv6 support in Firewall-1. An example of the ACL filtering in the Juniper OS can be seen at the foot of Annex A.

- NGNlab – this is an initiative led by the University of Brussels (ULB), which has established a testbed between Brussels and a partner site in Switzerland.

- Looking glasses – many partners have installed IPv6 Looking Glasses that allow the status of IPv6 links and routing to be inspected remotely, e.g. the POZNAN Looking Glass can be found at: http://www.ipv6.man.poznan.pl

It is expected that the TF-NGN IPv6 WG will continue to be a forum for discussion of such partner activities, complementing the work of 6NET.

2 IPv6 MULTICAST EXPERIMENTS

There are two well-known tunnel-based initiatives on the Internet. One is the MBone [mbone], over which IP multicast traffic is tunnelled between multicast-enabled routers. The other is the 6Bone [6bone], over which IPv6 traffic is tunnelled between IPv6-enabled routers. Both use the production IPv4 network as a medium over which to transport protocols not yet implemented natively.

2.1 The M6Bone

In this section we describe the M6Bone [m6bone] as an initiative to carry IPv6 multicast over native IPv6 networks, as well as tunnelled over both production IPv4 and IPv6 networks. Ideally, a fully native IPv6 network would exist covering all participant countries with all routers supporting a common IPv6 multicast protocol (and where necessary inter-domain multicast protocols). In practice however if we wish to test IPv6 multicast between European (and worldwide) partner sites we have to use tunnelled connectivity for at least the short term.

We believe that the 6NET project will deliver native IPv6 multicast during its lifetime; many of the participants in the GTPv6 group are also active in 6NET. Early experience gained on the M6Bone will be very useful for this goal.

The M6Bone is primarily operated by the G6 group of French IPv6 testers [g6] and uses the Aristote Association [aristote] for seminars to transmit on the network.

This service is based on the IPv6 network of Renater , and benefits from the logistic support of the Aristote association which is involved in the broadcasting of ultra-modern technology seminars and of G6, the French group of IPv6 testers. The first objective is to develop an advanced service on IPv6, in order to participate in the promotion of the protocol.

As the M6Bone grows it is expected that seminars from other sites can be shown. Plans to transmit World Cup matches in July 2002 were not carried through due to licensing issues.

2.2 Topology

There are two general methods to connect to the M6Bone:

- For sites that already have IPv6 connectivity, the tunnel will be an IPv6 (multicast) in IPv6 (unicast) tunnel.
- For sites that only have IPv4 connectivity, the tunnel will be an IPv6 (multicast) in IPv4 tunnel.

The topology is evolving as new sites connect. The centre of the M6Bone is dictated by the location of the single PIM-SM (Protocol Independent Multicast – Sparse Mode)[pim-sm] rendezvous point (RP), which is operated by Renater. This router also acts as the bootstrap router (BSR).

Because PIM-SM uses the unicast routing table to direct multicast traffic, it is necessary to set up routing for the multicast testbed subnet at each participating site to the next upstream site. An initial static route is required to enable subsequent use of RIPng for unicast routing. The testbed prefixes used at each site (usually a plain /64 prefix) are noted such that any unknown prefixes can be filtered at the RP. Southampton (as described below) is using a larger prefix for its tests as it has a hierarchical deployment of PIM-SM routers at its site. Note that multicast applications in this configuration must be run on hosts on the testbed subnets, not on the multicast routers themselves.

In most cases this is the Renater node, but it may be to one or more other nodes, as the current international topology in Figure 7 and local French topology in Figure 8 illustrate. UNINETT for example has three M6Bone links, none of which are directly to the RP in France. A status page is maintained for the network [m6bone], showing the status of the router, tunnel and PIM daemon for each connected site. The multicast routing table of the RP is also available for participants to view.

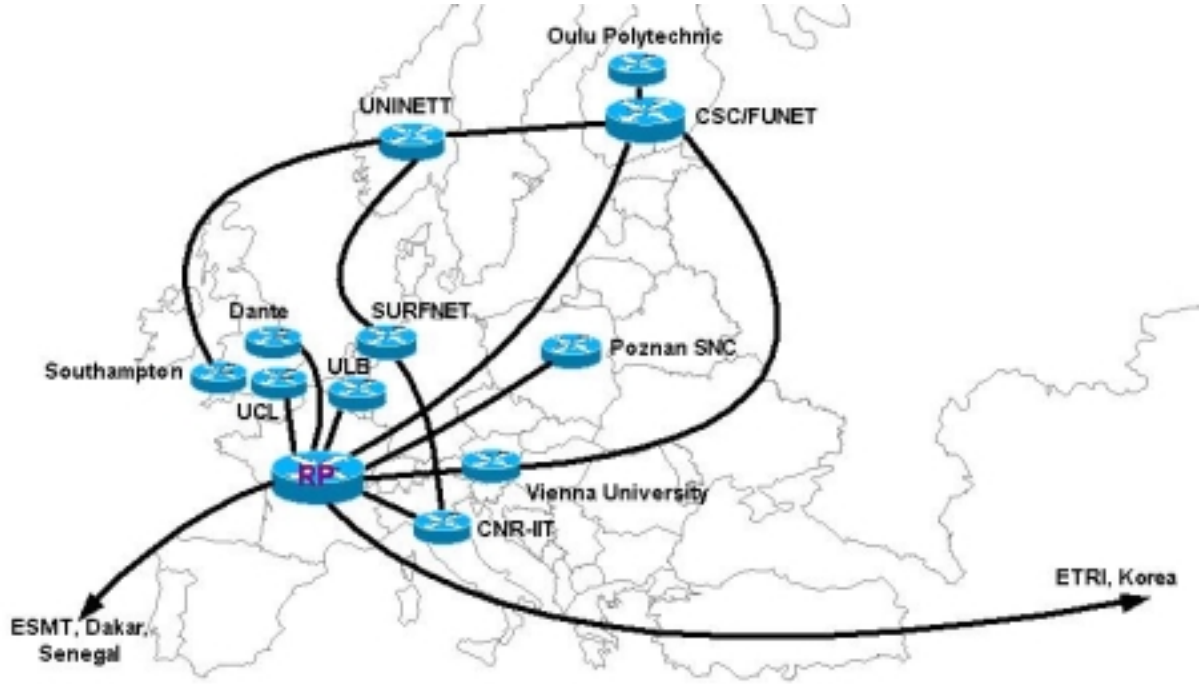

Figure 7: The international M6Bone sites as of August 2002

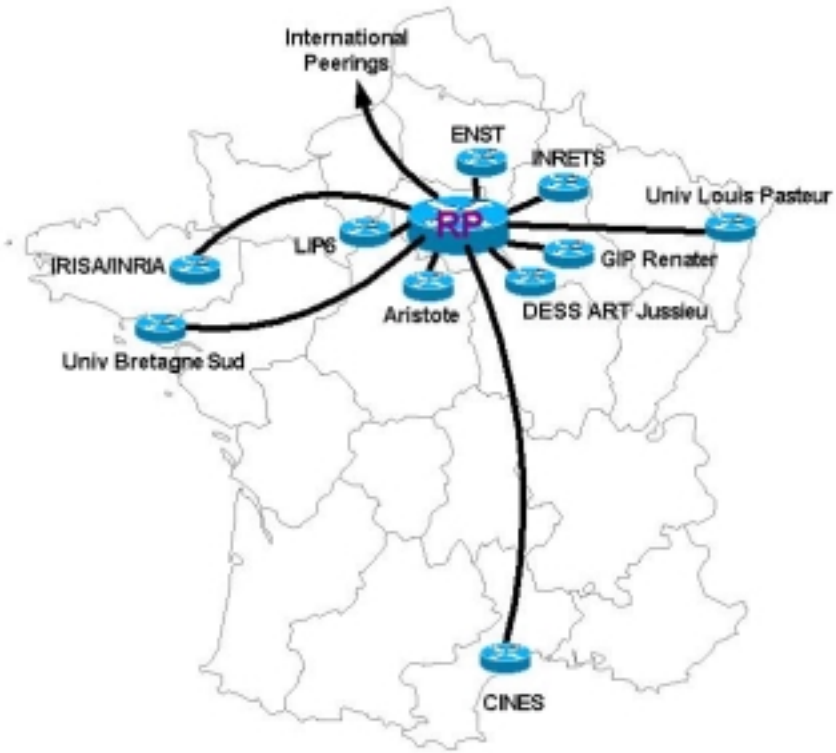

Figure 8: The M6Bone sites on the French network as of August 2002

2.3 Router and host equipment

The network is primarily using FreeBSD [freebsd] with the KAME IPv6 stack [kame]. The RP is currently running FreeBSD 4.5 since March 2002.

The FreeBSD configuration is detailed on the M6Bone web site [m6bone], for both IPv6-in-IPv4 and IPv6-in-IPv6 tunnels, showing settings for the /etc/rc.conf, /etc/rc.local and /etc/pim6sd.conf files.

There have also been tests in France with a 6WINDGate 6200 router [6wind] and early experiments by POZNAN with the Cisco IOS EFT image [cisco] with PIM-SM support. But the bulk of the sites are running FreeBSD (with versions up to FreeBSD 4.6), so at this stage interoperability issues have yet to be explored in earnest for IPv6 implementations of PIM-SM. Note the 6WIND device cannot yet run RIPng in IPv6-in-IPv6 tunnels, and there is a small bug in prune/join messages that means it cannot be used to redistribute multicast traffic.

Southampton plans to deploy PIM-SM on a Hitachi GR2000 router [gr2000] in the near future.

For workstations, the M6Bone participants have to date used FreeBSD, Linux and Windows 2000 [mice-win2k]. Binaries are available for FreeBSD 4.2 through 4.6, while an RPM exists for Linux [linux-rpm] or you can install from the MICE site [mice-install].

2.4 Gatewaying IPv4 and IPv6 multicast

A gateway for IPv4-IPv6 multicast is being developed vy Luc Beurton at the University of South Britany. The architecture, which is in principle relatively simple, is shown in Figure 9.

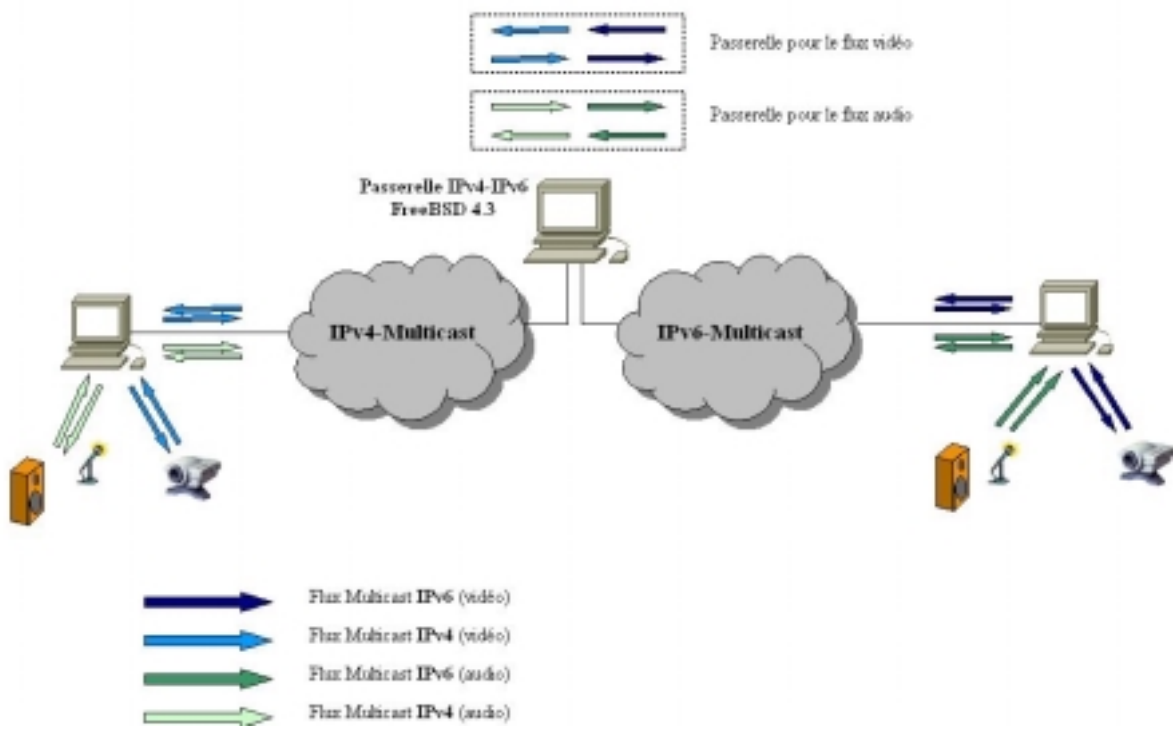

Figure 9: Gatewaying IPv4 and IPv6 multicast

The gateway was used for the Aristote seminar transmitted in June 2002 [mc-conf]. A multicast-to-unicast gateway is also in development from Konstantin Kabassanov of LIP6.

2.5 Multicast application usage

The M6Bone has initially been used for videoconferencing using the IPv6 enabled versions of the vic and rat MICE tools [mice]. This includes the transmission of seminars and conferences on the network [mc-conf]. We expect also to trial other software, e.g. UNINETT has offered an IPv6 multicast "underground radio" service, using royalty-free independent music.

There are new applications being written or ported for IPv6 that can use multicast, e.g. VideoLAN [videolan] for MPEG-2 video and Icecast [icecast] for audio. There are currently over 20 sites connected to the M6Bone.

2.6 M6bone case study: POZNAN (PSNC, Poland)

An IPv6 multicast software router was built based on the FreeBSD 4.5 system. The PC hardware configuration was as follows:

- Processor – MMX 200 MHz
- Memory – 64 MB RAM
- NIC – 2x 3com 3c509TP, 3com509

Additional equipment used during the tests included:

- Media converter – 2x Microsens MS416111
- Cisco Catalyst C924-A and Cisco Catalyst C2924-XL-EN

To build the multicast router the KAME stack compilation was required [kame]. Before starting multicast tests in PSNC, an IPv6 network was configured based on the Cisco C7500 router. On this router the following tunnels to other IPv6 sites were configured:

- Renater
- Nordunet
- Cesnet
- Cern
- Polish IPv6 networks

This router also handles BGP peerings with neighbouring IPv6 sites. The software IPv6 multicast router was connected to the Cisco router via an Ethernet interface. For test purposes we assigned the IPv6 prefix 3ffe:8320:5:0100:/56. The testbed topology is shown in Figure 10.

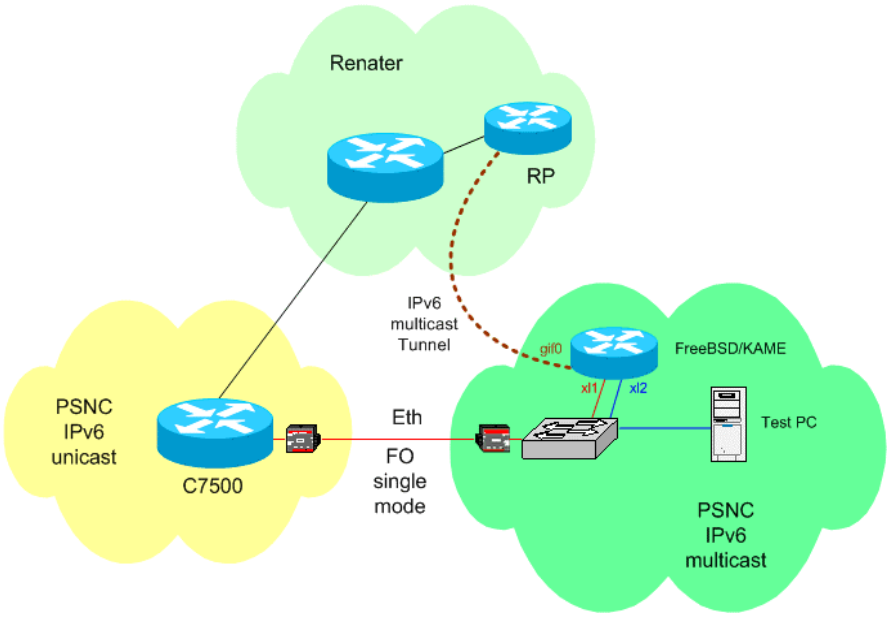

Figure 10: Polish m6bone connectivity

One interface of the software multicast router has an assigned IPv4 address for administrative purposes. Interfaces xl1 and xl2 were connected to a Cisco C2924-XL-EN switch. Additionally two VLANs have been configured on that device: one to connect the software router to the IPv6 network, and a second to connect the PC sending/receiving the data. The IPv6 address 3ffe:8320:5:101::1 was assigned to the xl1 interface. This address is also registered in the DNS as burdock2.m6bone.pl.

RIPng has been used as the unicast routing protocol. This protocol is supported by the route6d routing daemon. It was important for us to launch RIPng over all interfaces but Ethernet interface xl1, which is used to connect to the IPv6 unicast network.

To allow multicast transmission between PSNC and the Renater network an IPv6/IPv6 tunnel has been established. For this purpose we have used virtual interface gif0 configured on our multicast router. The RIPng exchange route update messages arrive only from Renater's RP at the gif0 interface. A default route interface xl1 was configured, in order to allow the multicast router to transmit IPv6 packets via the IPv6 unicast network.

PIM-SM v2 (Protocol Independent Multicast – Sparse Mode version 2) has been used as the multicast routing protocol. This protocol is handled by the pim6sd routing daemon. In the basic case there are no requirements for additional configuration of pim6sd. When pim6sd starts it automatically configures itself to forward packets using all multicast-capable interfaces. The logical testbed topology is shown in Figure 11.

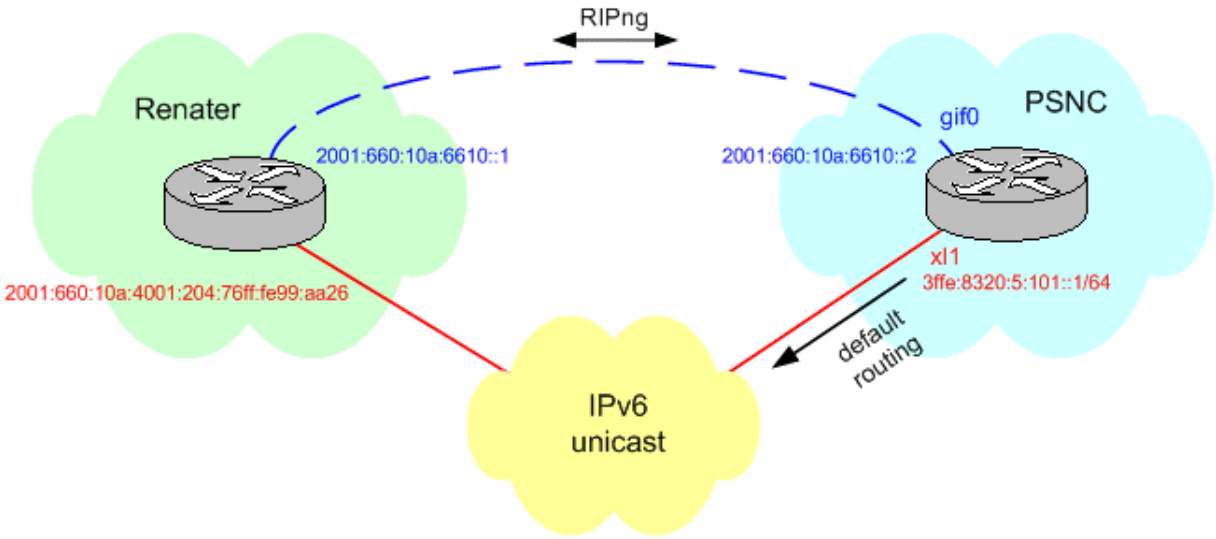

Figure 11: Logical testbed topology between POZNAN and Renater

In order to send a multicast stream we have used the VIC application available at the UCL site [mice]. With that tool we were able to send and receive multicast data between the PSNC and Renater networks as well as other sites connected to the M6bone.

The IPv6 Networking FAQ [ipv6nf] and FreeBSD Handbook [freebsdh] were useful reference material for the configurations.

Currently we are using the Cisco 7500 platform with an experimental IOS that supports IPv6 multicasts. The next step is to check the interoperability between Cisco routers and other software routers, like the one based on FreeBSD and KAME. For that purpose, the Cisco 7500 router will be connected to the M6bone with an additional tunnel.

2.7 M6Bone case study: University of Southampton (UoS)

Southampton is connected to the M6Bone, and is a partner on the 6NET project. Since the 6NET network and access PoPs do not yet support end to end multicast, the only means by which UoS can test multicast tools over remote networks is to use IP-in-IP tunnels to "bypass" those specific routers.

For the tests UoS has chosen to use its existing FreeBSD router hierarchy that provides the bulk of the internal IPv6 site connectivity, as illustrated in overview in Figure 12.

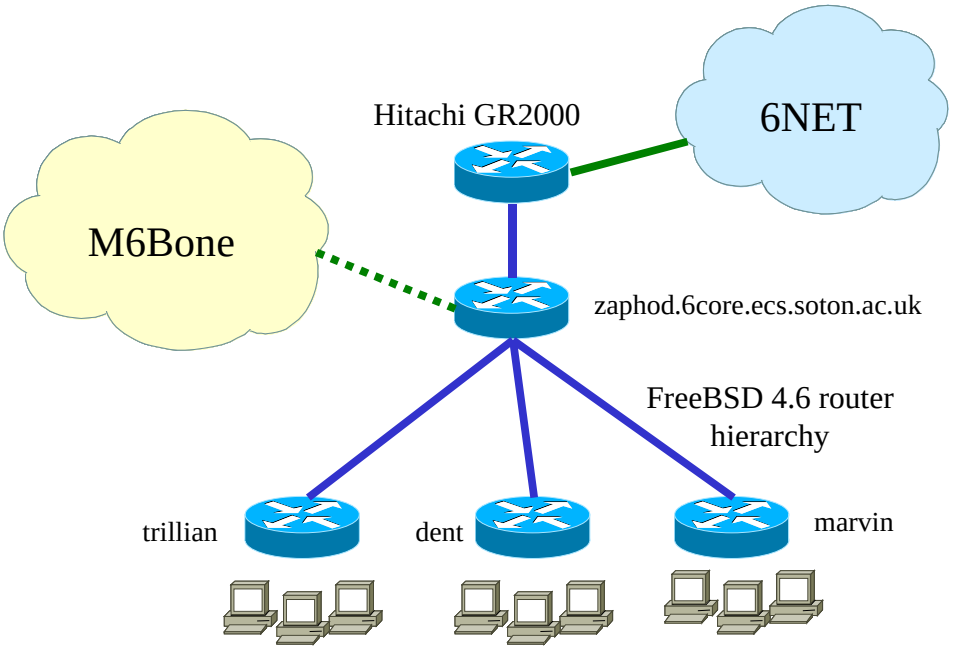

Figure 12: IPv6 multicast hierarchy at Southampton

As a result of earlier work undertaken in projects such as Bermuda [bermuda] and 6INIT [6init], Southampton had deployed various hierarchies of PIM-SM IPv6 multicast routers, with a local RP and BSR. This network has been used to trial typical IPv6 multicast applications such as the vic and rat tools [mice]. The router ‘zaphod’ receives delegation for a /56 prefix from its upstream Hitachi site router, and delegates a /60 prefix to leaf routers that then offer /64 prefixes to the end users (over VLANs on a switched Alcatel network, where we choose to overlay our IPv6 VLANs on our IPv4 VLAN space).

The UoS site prefix is 2001:630:d0::/48, ‘zaphod’ allocates IPv6 subnet from the aggregate 2001:630:d0:100::/56 and , for example, dent offers /64 prefixes from the prefix 2001:630:d0:110::/60. This is not efficient use of the address space, but it allows the establishment of a site router hierarchy. IPv6 /64 prefixes are offered from the leaf routers to many of the research groups of the Electronics and Computer Science Department, to its student teaching laboratories running Linux and Windows XP, and to its 802.11b WLAN network.

To join the M6Bone, the RP status of the router ‘zaphod’, as shown in Figure 12had to be removed, with the aim to establish an IP-in-IP tunnel between ‘zaphod’ (running pim6sd and route6d) and the chosen “upstream” multicast partner (UNINETT). Both pim6-sd and route6d should detect the other router and route both multicast and unicast traffic accordingly. UNINETT in Norway in turn links with Finland, then France which is the sole RP of the test-bed (see Figure 7).

The initial tunnel was IPv6-in-IPv4, and connected to the router ‘marvin’, a 1U-format PC with a D-Link quad FE card running FreeBSD 4.6 with the latest KAME snap-kit (August 2002). Because PIM-SM checks that packets received on an interface from a certain source address are received on the same interface as it would use to send a packet to that source address, the unicast and multicast routing for the test network are tightly coupled.

In enabling route6d on the tunnel interface, the prefix chosen for the tunnel was taken from the UNINETT address space 3ffe:2a00:100:7efe::/64, the UoS end was ::2 and the other ::1, however route6d consistently failed to advertise this route, even though other routes were exchanged between the two networks. The consequence of this was that other hosts on the M6Bone would use the 6NET network to send reply packets (since they had no route for 3ffe:2a00:100:7efe::/64). To solve this problem a different address scheme was used which was covered by routes that were known to be exchanged, i.e. the Southampton prefix is 2001:630:d0::/48. This route was exchanged between the tunnel end points, so by using addresses covered by this prefix the routing would be correct for other hosts. This assumption proved to be correct, and multicast was received correctly from and could be sent to other hosts on the network.

Another issue encountered was that the rtp tools that were being used [fefe] did not set the multicast hop limit, and so it defaulted to 1 hop. Hence when sending multicast data it could only be received on the same Ethernet link as the source. Once the limit was increased pim6sd correctly registered the source. Later the tunnel was migrated to ‘zaphod’ where it now runs in IPv6-in-IPv6 mode. In tests of the multicast tools, we ultimately were able to receive multicast from sites in France through ‘zaphod’ and ‘marvin’ into our student workstations.

The majority of the problems we encountered were related to routing. The pim6sd daemon itself needed no configuration, although the (empty) configuration file must exist for it to run (this file is not empty if the router acts as an RP or BSR). The tool pim6stat can be used for diagnosing the links, rather than using extensive logging which rapidly fills up the /var partitions on the routers.

One issue that needs further discussion is that of filtering the unicast routes on the multicast core router. If more-specific unicast routes are advertised inside the M6Bone, this affects the ability of those sites to route unicast traffic to the partners if filtering is applied. One option is to run without filters, assuming the core router provider is happy to also carry unicast transit traffic between the specific partner networks.

Most of the participants have separated their unicast and multicast routers to avoid this problem. The paradox is that where no multicast is available on all equipment in a network path we have to use different topologies, however where there are no multicast routing tables we can't have different topologies; this is why the separation has been made by most participants.

Our next tasks include enabling PIM-SM in the Hitachi, testing PIM-SM on our 6NET Cisco 7206 access router, and also experiments with site scope for IPv6 multicast. Ideally we should be able to run a local site scope RP on the Hitachi (or on ‘zaphod’), and still have a global scope RP in France.

2.8 Future work

There are many areas of further study for IPv6 multicast trials, in site networks, as part of the M6Bone, and feeding into the 6NET project. These include:

- Testing new code and devices, e.g. the Cisco IOS EFT with PIM-SM, the new 6WINDGate 6200 code, and the Hitachi GR2000. On the host side, Windows XP would be interesting. We have run multicast tools on Familiar Linux on a Compaq iPAQ PDA. WinCE should also support IPv6 soon.

- It would be helpful to have IPv6 beacons running. There are some initial positive results from a Java beacon run by SWITCH.

- Experiments with multicast scopes would be desirable.

- The deployment and testing of PIM-SSM (single source multicast) should be encouraged where available.

- Liaison with the new Internet 2 IPv6 Multicast WG [i2-v6mc] would be mutually beneficial. One area of study could be MBGP, along with multi-RP multicast networks.

- The IPv4-IPv6 multicast gateway is promising. It would be useful to run more tests of such a gateway device as an aid to IPv6 transition and integration. It may be interesting to consider such a gateway in the context of Access Grid nodes, which are heavily based on vic and rat technology.

In parallel with these activities, the M6Bone should be widely advertised and further partners sought, particularly in the TF-NGN IPv6 WG scope.

3 IPv6 STATUS

One of the general activities of the TF-NGN IPv6 WG is tracking of the status of IPv6 standards, implementations and related issues. These are reported briefly in this section.

3.1 IPv6 motivation

One of the common questions that arises is the question of motivation for deployment of IPv6. There are as yet no "killer" IPv6 applications, nor do we probably expect to see any in the immediate future. Without applications, providers are reluctant to deploy infrastructure where no revenue stream is apparent. However, in the NREN environment such pressures are far less; thus IPv6 can be deployed as an enabling technology. Of course IPv6 can be run between islands using an IPv4 infrastructure as the carrier, but native end-to-end deployment is a more attractive proposition where it can be provided.

The applications that IPv6 will favour include those that can benefit from the expanded address space (and thus the absence of NATs) or from improved Mobile IP support in IPv6. This implies peer-to-peer and ad-hoc or mobile computing, but may include other areas such as applications that communicate to devices that would be currently be unreachable (or difficult to reach without IP and/or port forwarding) behind IPv4 NATs (e.g. in home networks, or in student dormitories), or perhaps distributed computing systems where device reachability is an important issue. We can certainly expect many more types of IP-enabled devices to appear, e.g. as embedded systems or as PDAs. Much of the advantage of IPv6 will go hand-in-hand with Wireless LAN, if not with the somewhat delayed 3G and UMTS wireless networks.

The World Wide Web was not realised until many years after the introduction of IPv4. By returning the Internet to the original "always-on", fate-sharing model of over 20 years ago, the thesis is that we will enable new applications, and more equitable access for all.

3.2 Deployment status

The majority of early commercial deployment is happening in Japan, through providers such as IIJ and NTT. The Japanese WIDE network [wide] is also a flagship IPv6 deployment. Deployment in Europe is being led by EU-funded IST projects such as 6NET and Euro6IX [euro6ix]. There is as yet very little commercial service in Europe, although many operators are running pre-service trials. The only NREN to have migrated to dual-stack operation is SURFnet.

In the US a handful of networks have migrated to support dual-stack operation, including ESnet and Abilene. Abilene is running dual-stack on Cisco hardware, and will run dual-stack from the outset when it is upgraded to use Juniper routers later in 2002.

3.3 Standards status

The IETF remains the key standards group for IPv6. IPv6 work is undertaken in many WGs, including *ipv6* (formerly *ipng*), *ngtrans* (soon to become *v6ops*), *mobileip*, *dhc*, *dnsext*, *manet*, *nemo* and *zeroconf*. The IETF is now actively seeking to push out IPv6 requirements across all WGs, with a message that IPv6 is ready for full deployment.

The 3rd Generation Partnership Project (3GPP) [3gpp] is of interest for IPv6 for the potentially massive number of IPv6 devices that would quickly deploy if 3G or UMTS systems adopted IPv6 for the handsets. IPv6 has been accepted for the multimedia component delivery in future versions, but implementation remains some way off. However, Nokia has published its own white papers showing

that it plans dual-stack handsets and has a number of transition and integration scenarios under consideration [nokia-v6].

3.4 Addressing (RIPE NCC)

Address allocations for IPv6 at present come either as production assignments under 2001::/16 from the RIRs – the RIPE NCC in Europe – or as experimental assignments under 3ffe://16 from the 6Bone project. One of the problems faced is the lack of routing stability on the 6Bone, which is not separated from the "production" network.

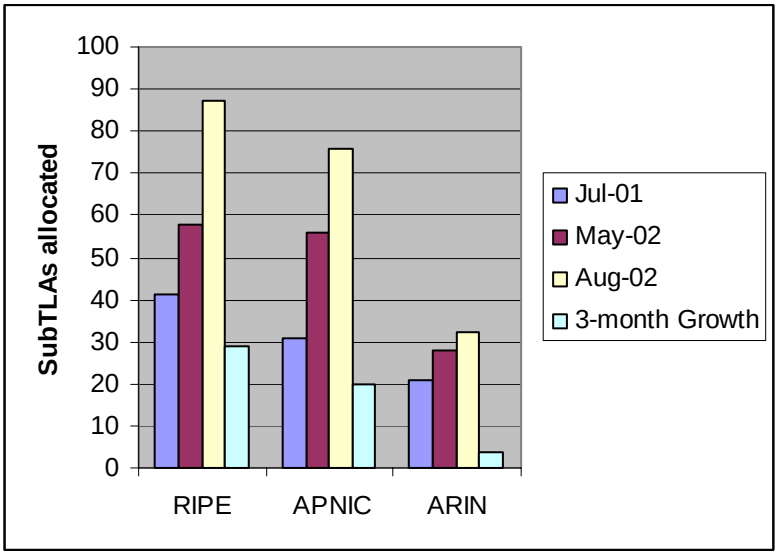

Figure 13: Growth of SubTLA allocations, July 2001 to August 2002

Some groups would like to see the 6Bone restricted to increase stability, but no consensus has yet been reached on how this could be achieved. There is also a proposal on the table to transfer 6Bone address assignment from the (soon to be obsolete) IETF ngtrans WG to the RIRs.

Figure 13 shows the volume of SubTLA allocations for the three RIRs and how they have grown between July 2001 and August 2002 (allocations began in 1999). Most notable is the fact that allocations are growing the fastest in the European region (although RIPE extends beyond Europe), and very slowly in the American region.

The recent rapid growth may reflect the new common RIR address allocation policy, which is far more open and less restrictive than in the past [aa-policy]. This policy also makes the default assignment a /32 prefix, where is used to be a /35.

3.5 Router implementations

The list of router implementations continues to grow and harden. The list includes:

- Cisco
- Juniper
- Hitachi
- FreeBSD (or any *BSD using KAME)
- 6WIND
- Zebra/ZebOS
- Telebit
- GateD

The important general development in the last 12 months has been the hardening of code from Cisco and Juniper (evidenced by networks like SURFnet and Abilene running dual-stack), and early support for IPv6 in hardware.

3.6 Host implementations

Host operating systems supporting IPv6 (to varying degrees) include:

- USAGI Linux
- Various Linux flavours (RedHat, Debian,…)
- FreeBSD (and other *BSD with KAME)
- Windows NT (very experimental), 2000 and XP (and soon .NET and Win CE .NET)
- Solaris 8 and 9
- Familiar Linux for PDAs (e.g. the iPAQ)
- Jaguar (MacOS X v10.2)
- AIX
- HP/UX
- Irix

Not all systems run IPv6 out of the box, e.g. it is an install option on Solaris, and on Windows XP you need to run the "ipv6 install" command at a command prompt.

3.7 Deploying IPv6 Applications

The list of available IPv6 applications continues to grow. It includes:

- DNS – BIND9
- Sendmail
- Apache, Mozilla
- OpenLDAP
- Vic, rat and sdr
- ISABEL conferencing suite
- VideoLAN, Icecast for streaming MPEG-2 or MP3

These applications are being run and tested by GTPv6 participants.

Understanding how to best port software to run in a protocol independent way is an important issue. It is also important that IPv6 comes bundled in applications, not as a bolt-on patch to a slightly old (and possibly vulnerable) version. Some useful online porting references include:

- http://www.kame.net/newsletter/19980604/
- http://www.viagenie.qc.ca/en/ipv6/presentations/IPv6%20porting%20appl_v1.pdf
- http://wwws.sun.com/software/solaris/ipv6/porting_guide_ipv6.pdf
- http://www.iu.hio.no/~mark/lectures/ipprog/
- http://www.tru64unix.compaq.com/internet/ipv6portingassistant/

4 RELATED PROJECTS

There are other projects into which the TF-NGN IPv6 WG work can feed, and also learn from. The current role of the IPv6 WG has evolved to one that both undertakes some level of testbed experimentation, but also acts as a think tank and a forum for input from all NRENs in the GÉANT community.

The main project with which the WG collaborates is the 6NET project, which includes over a third of the total GÉANT NRENs.

4.1 6NET

Many of the activities of the 6NET project arose out of work items previously undertaken in the TF-NGN IPv6 WG, as reported in GÉANT Deliverable D9.3 [geant-d93]. Examples include DNS, routing, multi-homing, transition tools, network monitoring and application porting and testing. However, given the large budget of 6NET (some €15M, covering over 1,100 man months of effort) it has considerably more resources to investigate these issues than the (relatively) ad-hoc TF-NGN IPv6 WG.

The 6NET project has approximately 100 deliverables planned, almost all of which are public. These are all available from the 6NET project web site in the Publications area [6net-pubs]. This includes reports on all the following 6NET Work Packages, there interrelationships of which are shown in Figure 14.

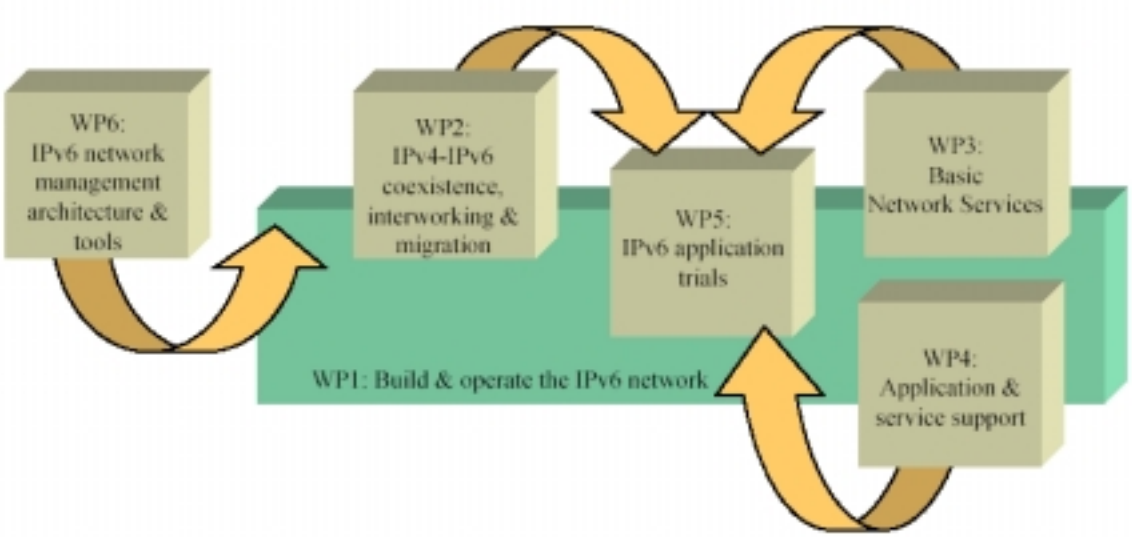

Figure 14: 6NET project work packages

4.1.1 WP1: Network architecture

The initial plan for the 6NET project network backbone is shown in Figure 15. Each backbone node connects via an STM-1 PoS link, and each core node has a national PoP router associated with it, which are generally Cisco GSRs. This network went live in May 2002, using IS-IS as the internal routing protocol.

In year two of the project (from January 2003), at least four of the backbone links will be upgraded to 2.5Gbit/s links, using the GÉANT infrastructure.

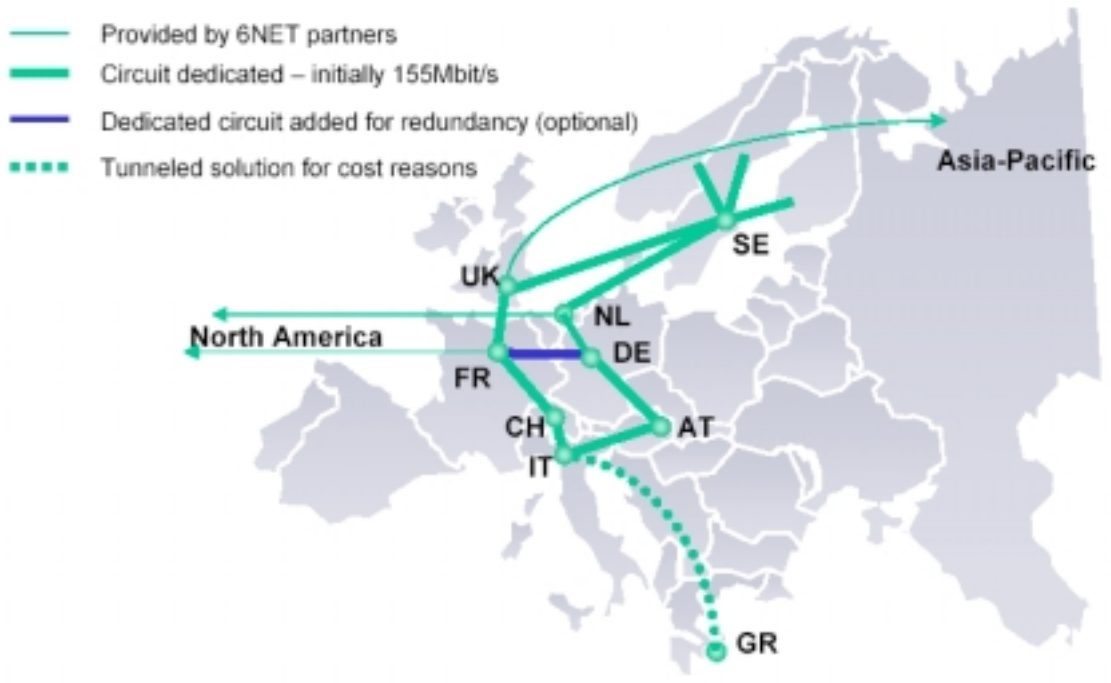

Figure 15: The initial 6NET network backbone topology

4.1.2 WP2: Transition

The transition studies in 6NET are split into three areas: site transition, NREN transition, and backbone transition. There is a considerable amount of overlap in technology between the latter two areas. The focus for the NRENs is to discuss and evaluate the tools that they can deploy in support of the end sites, which includes migrating the national network to dual stack operation, and deploying support services such as tunnel brokers and 6to4 relays. Each NREN is also likely to manage the DNS name space for its network, which needs to support AAAA records and respond to DNS requests via IPv6 transport.

The range of tools for end sites (universities) is much wider. The volume and complexity of the tools is currently a lively topic of debate in the IETF ngtrans WG [ngtrans], where the focus is on defining the transition scenarios and applicability of existing tools. Clearly 6NET can give good input to this process in the academic and research network context. It is expected that ngtrans will be replaced by a new v6ops WG in the near future, as part of the IETF's indication of "readiness" of IPv6.

4.1.3 WP3: Basic network services

The basic network services include such areas as routing and DNS. The 6NET backbone runs IS-IS. DNS servers have been set up to support at least the backbone name space.

4.1.4 WP4: Application and service support

This WP includes advanced services such as QoS and Mobile IPv6, with applications specifically in the area of 802.11b WLAN networks. An evaluation report on existing MIPv6 implementations has already been produced, as well as a report on WLAN-related standards.

4.1.5 WP5: Application porting and development

The project has identified a number of packages that will be ported to support IPv6. This includes a variety of multimedia streaming and conferencing tools, the Globus Toolkit v2, plus e-business and

web services (using WebSphere, in conjunction with IBM). The open source VOCAL VoIP package is also being ported to IPv6, and will be used for communication between 6NET partners.

This WP is also investigating best practice for IPv6 porting, for protocol independent applications.

4.1.6 WP6: Network management and monitoring

Network management and monitoring includes all the operational requirements to manage and run a production quality network. Thus a 6NET NOC has been created with operational procedures. A set of tools has also been identified for porting to IPv6. The WP includes deployment of IPv6-enabled versions of the RIPE Test Traffic servers.

4.2 TERENA Mobility WG

The IETF is actively pushing out IPv6 into working groups not directly related to IPv6 standards development, such that a broad view of the introduction of IPv6 is considered by all WGs. In a similar way, it is important that in the scope of GEANT, the TF-NGN IPv6 WG can reach out to other groups who may benefit from IPv6, or who should consider it in their own work plans.

A good example of this requirement is the TERENA Mobility WG, which is in the process of formation, having had some preliminary scoping meetings. The group is considering mechanisms for an international roaming mechanism for WLAN devices, to include authentication and access control. A natural extension of this work is to include IPv6 mobility in the studies. This has been taken on board by the group, which shares some membership with TF-NGN and 6NET.

4.3 International Collaboration

It is important to recognise the benefits of international collaboration, both in terms of international IPv6 connectivity and also in sharing experiences of IPv6 research and deployment, with a view to establishing joint collaboration projects.

In the past 12 months collaborations have included the following:

- Representatives from France and the UK were invited to the Japanese Gigabit Network conference in Okinawa in 2001. The meeting led to a better understanding of the drivers for IPv6 in Japan, of the current deployment status, of how the funding and promotion of IPv6 occurs, and of possible future joint projects.

- Renater established a direct native IPv6 connection on a link from Paris to Korea (through ETRI), called the Trans Eurasia Information Network (TEIN) [tein]. Collaboration activities include high-quality IPv6 videoconferencing and Multicast IPv6. ETRI is an international partner on the IST project 6WINIT, and should also be joining the 6NET project by the end of 2002.

- A representative from the UK presented the IPv6 work of GÉANT and TF-NGN at the Internet2 Spring Members meeting in Washington DC. This presentation was part of a joint session with Internet2 and WIDE (of Japan). Internet2's imminent dual stack IPv4-IPv6 deployment on its new Juniper-based Abilene network is of particular interest for GÉANT, which also uses the Juniper platform.

- Renater has started some collaboration work with the Russian academic network.

5 GÉANT MIGRATION TO IPV6 PILOT SERVICE

It is expected that GÉANT will migrate to support an IPv6 service, i.e. migrate to running dual stack IPv4 and IPv6 networking over the same links between its Juniper routers, within its lifetime. We note that the Internet 2 Abilene network [abilene] has moved to dual stack mode on its Cisco backbone, and will deploy dual-stack routers in the new Abilene network based on Juniper routers later in 2002.

Some initial considerations for this transition have been reported in 6NET Deliverable D2.1.1 [6net-pubs]. We repeat some of those considerations below:

A first task consists of evaluating the performance and interoperability of various router platforms. The parameters to consider include:

- *Forwarding performance in dual stack mode at line rate*.

  The idea is to measure the level of IPv6 traffic a router can forward (and with which kind of CPU consumption), and compare that to estimated IPv6 traffic load (multiplied by some factor for safety). This evaluation can be done in a laboratory with a simple IPv4/IPv6 setup to get a first idea. Then, a more complex test, which takes into account the other types of forwarding mechanisms, which are currently in production, can be achieved.
  The results expected are that there is no impact on the router's memory and that forwarding performance is close to the estimated traffic load.

- *IGP tests*

  Depending on which IGP the production network is running, tests can be done to measure the performance of the routers in the case where several IGPs are running in parallel for both stacks of protocols and compared with one IGP handling the both stacks.
  The results expected are that there is stability and no impact on the router's memory.

- *Tunnelling methods*

  Tunnels are part of the technical solutions for transitioning the network. One strategy could consist in having only the edge routers enabled dual stack.
  From edge to edge a tunnel IPv6/IPv4 can be established, core routers stay pure IPv4.
  Tunnelling and en/de-capsulation performance have to be evaluated, the parameters to look at are again forwarding performance and memory usage.

- *Interoperability*

  Interoperability tests have to be done for a backbone based on multi-vendor platforms. Tunneling techniques and routing protocols have to interoperate.

Given the initial tests, a design for the dual stack network will be sought that makes the transition as smooth as possible while also offering robustness and stability.

Before any equipment is deployed or routers upgraded, an IPv6 management and monitoring infrastructure needs to be put in place, and a NOC framework established. These issues are covered in WP6 of the 6NET project.

Even with the backbone in place, service upgrades need to be discussed with the NRENs, and where encapsulation is used the tunnel end points need to be determined (it is generally not desirable to terminate tunnels on the core for performance reasons). Many NRENs already have pilot IPv6 services, though only SURFnet has migrated to dual-stack operation on its national backbone.

Dual-stack mode is not the only solution to provisioning IPv6 on the backbone. It is possible to tunnel across the backbone, or to use Layer 2 or Layer 3 VPNs, but the dual-stack approach seems to be widely favoured. Feedback from early adopters should prove valuable to the process of GÉANT migration.

The next year's TF-NGN IPv6 WG activities should support the introduction of IPv6 on the GÉANT backbone.

6 CONCLUSIONS AND FUTURE WORK FOR THE TF-NGN IPV6 WG

The work undertaken in the past 12 months has led to the general conclusion that the support for and readiness of IPv6 for production deployment is rapidly hardening. It is now quite practical to be considering deployment of dual-stack IPv4-IPv6 networking on the GÉANT backbone. The Juniper routers used on that backbone have improving support for IPv6, and the Juniper architecture means that IPv4 and IPv6 are both routed in hardware.

Experiments with the M6Bone show that IPv6 multicast is also becoming more reliable and robust where implementations are available. For those at an early deployment phase, or on a limited budget for trials, we have seen that PC software routers such as Zebra can deliver good performance and features.

While much of the work of the TF-NGN IPv6 WG as it was 12 months ago has migrated into 6NET, there are still a number of areas that the group must continue to study and develop as 6NET does not cover all IPv6 aspects important to GÉANT and NRENs. These include:

- Expanding the M6Bone, getting more content delivered on it, tweaking the topology, trying new equipment. The multicast lessons learnt on this testbed can be fed into 6NET. Other specific items to consider could include PIM-SSM, more trials with IPv4-IPv6 multicast gateways, and tests with code supporting multicast scope.

- Determining the tests, trials and requirements for the migration of GÉANT to dual-stack IPv4-IPv6 operation. Experience can be drawn from the Abilene Juniper deployment, which will go live before the end of 2002.

- There must be a continuation of the GTPv6 network for both 6NET and non-6NET partner NRENs to participate in to study above all interoperability issues between different routing equipment. We plan the addition of a Hitachi GR2000 to the testbed, and to reuse the QTPVSIX 6Bone address space and the available DANTE ASN for the network.

- The group should work in general as a “think tank” for new IPv6 ideas, and to exchange ideas between those TF-NGN members who are in 6NET and those who are not. In 2002, three new TF-NGN members joined 6NET, from Poland, Hungary and the Czech Republic. However, membership of 6NET is not necessarily an end goal for NRENs, rather a better understanding of how to deploy IPv6 is of key importance, within the NREN, and in connecting to a future GÉANT IPv6 service.

- There will be a need to GÉANT to deliver at least four 2.5Gbit/s links to the 6NET project by early 2003. The technology for such a provision could be discussed within TF-NGN, although responsibility for the work lies with DANTE within the 6NET project.

- International collaboration is important. This includes overseas networks such as those operated by Internet2 (Abilene) and in the Pacific Rim, including Japan (WIDE) and Korea (through ETRI). GTPv6 members (from France and the UK) have been invited to present at both Internet2 and Japanese Gigabit Network (JGN) events in the past 12 months. Renater has also been active in establishing collaboration with Russia.

- It is important to get end users of IPv6 into the TF-NGN and 6NET networks. One aspect of this is provision of applications and content, another is making early deployment as cheap as

possible for end sites. For the former case the French initiative to transmit Aristote seminars on the M6Bone is a good example. For the latter case, further investigations of PC-based routers would be useful. Higher-end routers would be used in the higher-capacity backbone networks.

- The group should promote IPv6 to other Task Force groups and WGs, e.g. the TERENA Mobility WG.

- It is important to track standards activities, as well as issues such as addressing, DNS requirements (e.g. ip6.int and ip6.arpa).

There is thus still useful work that can be done in the scope of the TF-NGN IPv6 WG. The potential work items will be discussed at the next TF-NGN meeting in Budapest in October 2002.

It is important to note that the number of NRENs participating in the GTPv6 work and discussions (at least 15 NRENs) is greater than the number of NRENs directly involved in 6NET. It is also reasonable to expect that NRENs who have not yet expressed a direct interest in IPv6 may wish to use TF-NGN as a forum to gain experience from GTPv6 and 6NET participants.

7 REFERENCES

[3gpp] The Third Generation Partnership Project
http://www.3gpp.org/

[6bone] The 6bone
http://www.6bone.net/

[6init] The 6INIT Project
http://www.6init.org/

[6net] The 6NET project, IST-2001-32603
http://www.6net.org/

[6net-pubs] 6NET project publications and deliverables
http://www.6net.org/publications/

[6wind] 6WIND
http://www.6wind.com/

[aa-policy] Common RIR IPv6 Address Allocation Policy
http://www.ripe.net/ripe/docs/ipv6policy.html

[abilene] The Internet 2 Abilene network
http://www.ucaid.org/abilene/

[aristote] Aristote Association
http://www.aristote.asso.fr/

[bermuda] The Bermuda IPv6 Project
http://www.ipv6.ac.uk/bermuda2/

[bird] BIRD
http://bird.network.cz/

[cern-v6] CERN Cisco IOS IPv6 router configuration
http://jmj.home.cern.ch/jmj/qtp/ipv6-atm7-cfg.txt

[cisco] Cisco IOS and IPv6
http://www.cisco.com/warp/public/732/Tech/ipv6/

[dstm] Dual Stack Transition Mechanism: IETF Draft
http://www.ietf.org/internet-drafts/draft-ietf-ngtrans-dstm-08.txt

[dstm-imp] DSTM Implementation by ENST
http://www.ipv6.rennes.enst-bretagne.fr/dstm/

[euro6ix] The Euro6IX Project
http://www.euro6ix.org/

[fefe] FEFE IPv6-enabled RTP tools

http://www.fefe.de/rtp/

[freebsd] FreeBSD
http://www.freebsd.org/

[freebsdh] FreeBSD Handbook
http://www.freebsd.org/handbook/

[g6] G6 Association
http://www.g6.asso.fr/

[gated] GateD Routing Daemon
http://www.gated.org/

[geant-d93] GÉANT Deliverable D9.3
http://www.dante.org.uk/tf-ngn/D9.3.pdf

[gr2000] Hitachi GR2000 router
http://www.v6.hitachi.co.jp/GR2000/

[gtpv6] GÉANT Test Programme for IPv6 (GTPv6)
http://www.ipv6.ac.uk/gtpv6/

[i2-v6mc] Internet 2 IPv6 Multicast Working Group
http://multicast.internet2.edu/mcast-v6.shtml

[icecast] Open Source streaming audio
http://www.icecast.org/

[ietf] The Internet Engineering Task Force (IETF)
http://www.ietf.org/

[ipv6nf] IPv6 Networking FAQ
http://www.netbsd.org/Documentation/network/ipv6/

[kame] The KAME Project
http://www.kame.net/

[linux-rpm] MICE tools for Linux in RPM format
http://www.netcore.fi/pekkas/linux/ipv6/

[m6bone] The M6Bone
http://sem2.renater.fr/m6bone/

[mbone] The IETF MBone Deployment WG Charter
http://www.ietf.org/html.charters/mboned-charter.html

[mc-conf] Transmission of an Aristote Seminar by IPv6 Multicast (6th June 2002)
http://sem2.renater.fr/m6bone/Aristote.pdf (French language)

[mice] The MICE conferencing tools
http://www-mice.cs.ucl.ac.uk/multimedia/software/

[mice-install] Installing the IPv6 MICE tools on Linux
http://domen.uninett.no/~venaas/m6bone-linux/

[mice-win2k] MICE tools for IPv6 on Windows 2000
http://www-rp.lip6.fr/~kabassan/

[mrtd] MRTD
http://www.mrtd.org/

[ngnlab] The NGNLab Project
http://www.ngnlab.org/

[ngtrans] IETF ngtrans WG (to be replaced by the IETF v6ops WG)
http://www.ietf.org/html.charters/ngtrans-charter.html

[nokia-v6] Nokia and IPv6
http://www.nokia.com/ipv6/

[pim-sm] The IETF PIM-SM WG Charter
http://www.ietf.org/html.charters/pim-charter.html

[poznan-v6] POZNAN IPv6 network (Poland)
http://www.ipv6.man.poznan.pl/

[tein] Trans Eurasia Information Network (TEIN)
http://www.transeurasia.org/

[tf-meets] Minutes and presentations of TF-NGN Meetings
http://www.dante.net/tf-ngn/meetings.html

[tf-ngn] GÉANT Task Force Next Generation Networks (TF-NGN)
http://www.dante.net/tf-ngn/

[videolan] VideoLAN - OpenSource Video streaming
http://www.videolan.org/

[wide] WIDE network
http://www.wide.ad.jp/

[zebos] ZebOS Routing Software
http://www.ipinfusion.com/

[zebra] The Zebra router project
http://www.zebra.org/

ANNEX A: JUNIPER M5 CONFIGURATION FOR GTPV6 TESTBED

```
gtpv6# show
version 5.2R2.3;
system {
    host-name gtpv6;
    domain-name cssi.renater.fr;
    backup-router 193.51.206.165;
    time-zone Europe/Paris;
    root-authentication {
        encrypted-password "XXXXXXXXXXXXX"; # SECRET-DATA
    }
    name-server {
        193.51.206.65;
    }
    login {
        class admin {
            idle-timeout 5;
            permissions all;
        }
        class users {
            idle-timeout 5;
            permissions [ admin configure network view ];
        }
        user admres {
            uid 2001;
            class superuser;
            authentication {
                encrypted-password "YYYYYYYYYYYYYYYYY"; # SECRET-DATA
            }
```

```
        }
    }
    services {
        ftp;
        telnet;
    }
    syslog {
        user admres {
            any critical;
        }
        file GTPV6.log {
            any warning;
            authorization any;
            archive size 1000000 files 20 no-world-readable;
        }
    }
}
interfaces {
    gr-0/1/0 {
        unit 4003 {
            description "--  Peering vers CERN (CH) --";
            tunnel {
                source 193.51.207.243;
                destination 192.65.185.7;
            }
            family iso;
            family inet6 {
                address 2001:660:1102:4003::1/64;
            }
        }
```

```
        unit 4006 {
            description "--  Peering vers POZNAN (PL) --";
            tunnel {
                source 193.51.207.243;
                destination 150.254.210.109;
            }
            family iso;
            family inet6 {
                address 2001:660:1102:4006::1/64;

            }
        }
        unit 4007 {
            description "--  Peering vers REDIRIS (SP) --";
            tunnel {
                source 193.51.207.243;
                destination 130.206.0.130;
            }
            family iso;
            family inet6 {
                address 2001:660:1102:4007::1/64;
            }
        }
        unit 4008 {
            description "--  Peering vers HEANET (IE) --";
            tunnel {
                source 193.51.207.243;
                destination 193.1.195.61;
            }
            family iso;
            family inet6 {
```

```
            address 2001:660:1102:4008::1/64;
        }
    }
    unit 4009 {
        description "-- Peering vers Switch (CH) --";
        tunnel {
            source 193.51.207.243;
            destination 130.59.32.38;
        }
        family iso;
        family inet6 {
            address 2001:660:1102:4009::1/64;
        }
    }
    unit 4010 {
        description "--  Peering vers CSnet (CZ) --";
        tunnel {
            source 193.51.207.243;
            destination 195.113.156.183;
        }
        family iso;
        family inet6 {
            address 2001:660:1102:400A::1/64;
        }
    }
    unit 4011 {
        description "--  Peering vers RESTENA (LU) --";
        tunnel {
            source 193.51.207.243;
            destination 158.64.16.21;
        }
```

```
            family inet6 {
                address 2001:660:1102:400B::1/64;
            }
        }
    }
    ip-0/1/0 {
        unit 4005 {
            description "--  Peering vers DANTE (UK) --";
            tunnel {
                source 193.51.207.243;
                destination 193.63.211.61;
            }
            family inet6 {
                address 2001:660:1102:4005::1/64;

            }
        }
    }
    at-0/3/1 {
        description "---- Lien vers Nio-atmv6 A3/0/0 ----";
        atm-options {
            vpi 0 maximum-vcs 1200;
            ilmi;
        }
        unit 0 {
            description "---- Lien v4 vers Nio-N1 A7/0 ----";
            vci 0.130;
            family inet {
                address 193.51.206.166/30;
            }
        }
```

```
        unit 1 {
            description "---- Lien vers Gsr-Nio A1/0 ----";
            vci 0.800;
            family iso;
            family inet6 {
                address 2001:660:1102:4001::2/64;
            }
        }
        unit 3 {
            description "---- Lien vers Nio-IPv6 A2/0 ----";
            vci 0.802;
            family iso;
            family inet6 {
                address 2001:660:80:400D::2/64;
            }
        }
        unit 4 {
            description "---- Lien vers 7200-T ----";
            vci 0.804;
            family iso;
            family inet6 {
                address 2001:660:80:400e::1/64;
            }
        }
    }
    fxp0 {
        unit 0 {
            description "-A-- Administration Locale ---- ";
            family inet {
                address 192.168.10.65/24;
```

```
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 193.51.207.243/32;
            }
            family iso {
                address 49.0001.1930.5120.7243.00;
            }
        }
    }
}
snmp {
    location "CIPB PARIS";
    contact "noc-ipv6@cssi.renater.fr";
    community gNpIc0m {
        authorization read-only;
        clients {

            193.51.206.65/32;
            193.51.206.66/32;
            193.51.206.67/32;
            193.51.206.226/32;
            193.49.160.17/32;
            193.49.160.26/32;
            193.49.160.47/32;
        }
    }
}
```

```
routing-options {
    rib inet6.0 {
        static {
            route 2001:660::/35 discard;
            route ::/0 next-hop 2001:660:1102:4001::1;
        }
    }
    static {
        route 0.0.0.0/0 next-hop 193.51.206.165;
    }
}
protocols {
    bgp {
        advertise-inactive;
        family inet6 {
            unicast;
        }
        local-as 2200;
        group IBGPv6 {
            type internal;
            local-as 2200;
            neighbor 2001:660:1102:4001::1 {
                description "*** IBGP Peering with Gsr-Nio ***";
                peer-as 2200;
            }
            neighbor 2001:660:80:400D::1 {
                description "*** IBGP Peering with Nio-IPv6 ***";
                peer-as 2200;
            }
            neighbor 2001:660:80:400E::2 {
                description "*** IBGP Peering with 7200-T ***";
```

```
                peer-as 2200;
            }
        }
        group EBGPv6 {
            type external;
            neighbor 2001:660:1102:4007::2 {
                description "-- Peering vers REDIRIS (SP)
(miguel.sotos@rediris.es) --";
                import ps-from-INTERNATIONAL;
                export ps-to-INTERNATIONAL;
                peer-as 766;
            }
            neighbor 2001:660:1102:4003::2 {
                description "-- Peering vers CERN (joop.joosten@cern.ch) --
";
                import ps-from-INTERNATIONAL;
                export ps-to-INTERNATIONAL;
                peer-as 513;
            }
            neighbor 2001:660:1102:4008::2 {
                description "-- Peering vers HEANET (IE) (noc@heanet.ie) --
";
                import ps-from-INTERNATIONAL;
                export ps-to-INTERNATIONAL;
                peer-as 1213;
            }
            neighbor 2001:660:1102:4005::2 {
                description "-- Peering vers DANTE (UK)
(janos.mohacsi@dante.org.uk) --";

                export ps-to-INTERNATIONAL;
                peer-as 6683;
            }
```

```
            neighbor 2001:660:1102:4009::2 {

                description "-- Peering vers Switch (CH)
(simon@limmat.switch.ch) --";

                import ps-from-INTERNATIONAL;

                export ps-to-INTERNATIONAL;

                peer-as 559;

            }

            neighbor 2001:660:1102:4006::2 {

                description "-- Peering vers POZNAN (PL) (ipv6-
support@man.poznan.pl) --";

                import ps-from-INTERNATIONAL;

                export ps-to-INTERNATIONAL;

                peer-as 9112;

            }

            neighbor 2001:660:1102:400A::2 {

                description "-- Peering vers CESNET (CZ) (lhotka@cesnet.cz)
--";

                import ps-from-INTERNATIONAL;

                export ps-to-INTERNATIONAL;

                peer-as 2852;

            }

            neighbor 2001:660:1102:400B::2 {

                description "-- Peering vers RESTENA (LU)
(yves.schaaf@restena.lu) --";

                import ps-from-INTERNATIONAL;

                export ps-to-INTERNATIONAL;

                peer-as 2602;

            }

        }

    }

    isis {

        level 1 wide-metrics-only;
```

```
        interface gr-0/1/0.4003 {
            passive;
        }
        interface ip-0/1/0.4005 {
            passive;
        }
        interface gr-0/1/0.4006 {
            passive;
        }
        interface gr-0/1/0.4007 {
            passive;
        }
        interface gr-0/1/0.4008 {
            passive;
        }
        interface gr-0/1/0.4009 {
            passive;
        }
        interface gr-0/1/0.4010 {
            passive;
        }
        interface gr-0/1/0.4011 {
            passive;
        }
        interface lo0.0 {
            passive;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface at-0/3/1.0;
```

```
            interface lo0.0 {
                passive;
            }
        }

    }
}
policy-options {
    prefix-list P6R2-static {
        2001:660::/35;
    }
    prefix-list Accessv4-Accept {
        193.49.160.26/32;
        193.49.160.48/32;
        193.51.206.54/32;
        193.51.206.65/32;
        193.51.206.66/32;
        193.51.206.67/32;
    }
    prefix-list Accessv6-Accept {
        2001:660:10a:4000::/60;
        2001:660:112:4000::/60;
    }
    policy-statement ps-P6R2-static {
        term static {
            from {
                prefix-list P6R2-static;
            }
            then accept;
        }
    }
```

```
policy-statement ps-to-INTERNATIONAL {
    term to-INTERNATIONAL-reject {
        from {
            route-filter ::/0 exact;
            route-filter 2002::/16 exact;
        }
        then reject;
    }
    term to-INTERNATIONAL-accept {
        from {
            route-filter 3ffe::/17 prefix-length-range /24-/24;
            route-filter 3ffe:8000::/17 prefix-length-range /28-/28;
            route-filter 3ffe:4000::/18 prefix-length-range /32-/32;
            route-filter 2000::/3 prefix-length-range /16-/16;
            route-filter 2001::/16 prefix-length-range /29-/35;
        }
        then accept;
    }
    then reject;
}
policy-statement ps-from-INTERNATIONAL {
    term from-INTERNATIONAL-reject {
        from {
            route-filter ::/0 exact;
        }
        then reject;
    }
    term from-INTERNATIONAL-accept {
        from {
            route-filter 2002::/16 exact;
```

```
                route-filter 3FFE::/17 prefix-length-range /24-/24;
                route-filter 3FFE:8000::/17 prefix-length-range /28-/28;
                route-filter 3FFE:4000::/18 prefix-length-range /32-/32;
                route-filter 2000::/3 prefix-length-range /16-/16;
                route-filter 2001::/16 prefix-length-range /29-/35;
            }
            then accept;
        }
        then reject;
    }

}
firewall {
    filter GTPv6-access {
        term Access-Accept {
            from {
                source-prefix-list {
                    Accessv4-Accept;
                }
                destination-port [ telnet ssh ftp ftp-data ];
            }
            then {
                count Access-GTPV6;
                log;
                accept;
            }
        }
        term Tunnel {
            from {
                protocol [ gre ipv6 ];
            }
```

```
            then {
                count Tunnel;
                accept;
            }
        }
        term BGP {
            from {
                protocol tcp;
                port bgp;
            }
            then {
                count BGP;
                accept;
            }
        }
        term Access-Deny {
            then {
                count THE_rest;
                log;
                accept;
            }
        }
    }
}
```