


IST-2000-32603	Deliverable D 2.2.3	
----------------	---------------------	---

Project Number:	IST-2001-32603
Project Title:	6NET
CEC Deliverable Number:	32603/UOS/DS/2.2.3/A1
Contractual Date of Delivery to the CEC:	31 st December 2003
Actual Date of Delivery to the CEC:	25 th May 2004
Title of Deliverable:	D2.2.3: Updated IPv4 to IPv6 transition Cookbook for organisational/ISP (NREN) and backbone networks.
Work package contributing to Deliverable:	WP2
Version:	1.03 (24 May 2004)
Type of Deliverable*:	R
Deliverable Security Class**:	PU
Editors:	Tim Chown (University of Southampton)
Reviewers:	Tina Strauf, Christian Strauf (Muenster/JOIN)
Contributors:	Erik-Jan Bos, Niels den Otter (SURFnet), Pekka Savola (CSC/Funet), Gabriella Paolini, Valentino R. Carcione (GARR), Jérôme Durand (Renater), Ladislav Lhotka (CESNET), Christian Schild (WWU), Rob Evans, Duncan Rogerson, Rina Samani (JANET NOSC and UKERNA), Dimitrios Kalogeras (NTUA), Stig Venaas, Trond Skjesol (UNINETT), Fotis Karayannis, Athanassios Liakopoulos, Chrysostomos Tziouvaras (GRNET), Janos Mohacsi (HUNGARNET), Simon Leinen (SWITCH), Carlos Friacas (FCCN), Jan P. Sorensen (UNI-C), Wilfried Woeber (ACONet), Bartek Gajda (POZNAN)

* Type: P - Prototype, R - Report, D - Demonstrator, O - Other

** Security Class: PU- Public, PP- Restricted to other programme participants (including the Commission), RE - Restricted to a group defined by the consortium (including the Commission), CO - Confidential, only for members of the consortium (including the Commission)

Abstract:

We describe the IPv6 transition mechanisms available to the National Research and Education Networks (NRENs) who are part of the 6NET project. The mechanisms need to operate to complement those that provide an IPv6 service to the end users in the universities. We review the mechanisms, state the current usage of those mechanisms, and describe some of the scenarios for NREN transition. This “cookbook” of transition mechanisms and experience will be updated throughout the duration of the 6NET project.

Keywords:

IPv6 NREN transition, IPv6 ISP transition, Dual-stack Networks

Executive Summary

The 6NET project involves the participation of some 15 European National Research and Education Network (NRENs). While the early focus of the project lay on deploying a native, IPv6-only backbone network spanning these NRENs, provision of end-to-end IPv6 services between universities and other sites requires IPv6 support in the national networks.

The NRENs provide connectivity to universities directly, or to MANs which then connect the universities. These NRENs will have existing production IPv4 networks, the efficient and reliable running of which are paramount to their operation. Methods of introducing IPv6 into these networks are required that both enable IPv6 at the same level of performance as IPv4, but also that do not adversely impact the performance of the production IPv4 network.

In this report, which is the third version of an ongoing NREN IPv6 Transition Cookbook development process, we describe the scenarios for IPv6 transition for NRENs (drawing comparison to IETF work in the ISP transition area), we describe the theory behind the transition techniques applicable to NRENs (some of which was reported previously in 6NET Deliverable D2.2.1 and D2.2.2), and complete the document with descriptions of how some of these techniques have been applied in the 6NET NRENs to date, along with configuration and related examples. We also offer a summary of a survey of NREN deployment status and experience drawn in late March 2004.

The methods described in this report are applicable to general ISPs, although not to all their operations (e.g. NRENs do not generally deliver IP services to home networks). The focus here is on backbone network services. Thus the report also embraces the backbone transition elements reported previously in D2.1.1, but which were then merged into D2.2.2. Due to the similarities of techniques with NREN and backbone transition, and because the 6NET backbone started IPv6-only (and no transition is thus possible), the backbone transition coverage now appears in the NREN reporting under 6NET Activity A2.2.

The cookbook will be updated throughout the lifetime of the project, culminating in a final version with D2.3.4 in December 2004.

The authors would be very happy to receive feedback and suggestions for improvements to the content of this cookbook, with a view to improving its usefulness and applicability. Please send such comments to the editor, Tim Chown at tjc@ecs.soton.ac.uk.

Table of Contents

1. INTRODUCTION.....	7
2. SCENARIOS FOR NREN/ISP TRANSITION	9
2.1. TRANSITION SCENARIOS FOR NRENS	9
2.1.1. <i>Layer 2 transport protocol</i>	9
2.1.2. <i>Routing protocol</i>	9
2.1.3. <i>Additional service support</i>	10
2.2. COMPARISON TO IETF ISP TRANSITION SCENARIOS DOCUMENT	10
2.3. REQUIREMENTS FOR TRANSITION	12
2.4. SOLUTION SPACE FOR IPV6 TRANSITION FOR NRENS.....	12
2.4.1. <i>PoS scenario</i>	12
2.4.2. <i>MPLS scenario</i>	13
2.4.3. <i>ATM scenario</i>	13
3. REVIEW OF NREN/ISP TRANSITION MECHANISMS	14
3.1. GENERAL APPROACH	14
3.2. DUAL-STACK.....	15
3.3. GENERAL TUNNELS	15
3.4. IPV6 OVER MPLS	16
3.4.1. <i>6PE – theory of operation</i>	17
3.4.2. <i>Comparison to other transition technologies</i>	19
3.4.3. <i>IPv6 over MPLS on HUNGARNET</i>	21
3.5. IPV6 OVER ATM	22
3.5.1. <i>Permanent and Switched Virtual Circuits</i>	22
3.5.2. <i>Tunnel setup and tunnel overhead</i>	22
3.5.3. <i>IPv6 transition with ATM</i>	22
3.6. DEPLOYING A PARALLEL IPV6-ONLY NETWORK	23
3.7. SUPPORT MECHANISMS: TUNNEL BROKER	24
3.8. SUPPORT MECHANISMS: 6TO4 AND 6TO4 RELAY.....	25
3.8.1. <i>Architecture</i>	25
3.8.2. <i>Considerations for 6to4</i>	26
3.8.3. <i>Operational and security issues</i>	27
3.8.4. <i>Configuration example for 6to4</i>	27
4. GÉANT IPV6 TRANSITION OVERVIEW	30
4.1. DUAL STACK CONSIDERATIONS.....	30
4.2. ROUTING PROTOCOLS	30
4.3. IS-IS AND OSPFV3 CONSIDERATIONS	32
4.4. TRIALS AND CONSIDERATIONS PRIOR TO TRANSITION	32
4.4.1. <i>Router performance</i>	33
4.5. MIGRATING THE IGP	33
4.5.1. <i>Monitoring the transition</i>	33
4.6. TOWARDS AN IPV6 SERVICE.....	34
5. DUAL-STACK NREN DEPLOYMENT CASE STUDIES	35
5.1. SURFNET CASE STUDY (NETHERLANDS)	35

5.1.1.	<i>Introduction.....</i>	35
5.1.2.	<i>The SURFnet5 dual stack network.....</i>	35
5.1.3.	<i>Customer connections.....</i>	36
5.1.4.	<i>Addressing plan.....</i>	37
5.1.5.	<i>Routing.....</i>	39
5.1.6.	<i>Network management and Monitoring</i>	39
5.1.7.	<i>Other services</i>	39
5.1.8.	<i>Contact information.....</i>	41
5.2.	FUNET CASE STUDY (FINLAND)	41
5.2.1.	<i>Overview.....</i>	41
5.2.2.	<i>History.....</i>	41
5.2.3.	<i>Addressing plan.....</i>	42
5.2.4.	<i>Configuration details</i>	44
5.2.5.	<i>Monitoring</i>	45
5.2.6.	<i>Other services</i>	46
5.3.	RENATER CAST STUDY (FRANCE)	46
5.3.1.	<i>Overview.....</i>	46
5.3.2.	<i>Native support.....</i>	46
5.3.3.	<i>Addressing and naming.....</i>	46
5.3.4.	<i>Connecting to Renater 3</i>	47
5.3.5.	<i>The regional networks.....</i>	47
5.3.6.	<i>International connections</i>	47
5.3.7.	<i>IPv6 Multicast.....</i>	48
6.	OTHER CASE STUDIES FOR NREN TRANSITION	49
6.1.	6WiN: INTRODUCTION OF A PARALLEL IPv6 NETWORK (DFN: GERMANY)	49
6.1.1.	<i>Internal connectivity and setup.....</i>	49
6.1.2.	<i>External connectivity.....</i>	52
6.1.3.	<i>Addressing.....</i>	52
6.1.4.	<i>Internal Routing</i>	53
6.1.5.	<i>Multicast.....</i>	53
6.1.6.	<i>Configuration example.....</i>	53
6.1.7.	<i>Future usage of 6WiN.....</i>	54
6.2.	PILOT NREN IPv6 SERVICE (UKERNA: UK).....	54
6.2.1.	<i>Current services.....</i>	54
6.2.2.	<i>Future services.....</i>	55
6.2.3.	<i>Applying for Connection.....</i>	55
6.2.4.	<i>Support Issues</i>	55
6.3.	NREN MIGRATING FROM 6BONE TRIALS (GRNET: GREECE).....	55
6.3.1.	<i>IPv6 Service in GRNET.....</i>	55
6.3.2.	<i>Services in production.....</i>	55
6.3.3.	<i>Services planned for production</i>	57
7.	SURVEY OF NREN DEPLOYMENT EXPERIENCE	58
7.1.	OVERVIEW OF IPv6 SERVICE	59
7.1.1.	<i>Nature of service.....</i>	59
7.1.2.	<i>Equipment used.....</i>	59
7.1.3.	<i>IPv6 addressing plan for core.....</i>	59
7.1.4.	<i>IPv6 addressing plan for universities</i>	60

7.1.5.	<i>Routing protocols used</i>	60
7.1.6.	<i>Any special considerations with IPv4 routing</i>	60
7.1.7.	<i>IPv6 performance issues</i>	60
7.2.	SUPPORT MECHANISMS	61
7.2.1.	<i>6to4 relay service deployment</i>	61
7.2.2.	<i>Tunnel broker deployment</i>	61
7.2.3.	<i>Other transition/access methods</i>	61
7.3.	OPERATIONAL NOTES	61
7.3.1.	<i>Universities connected to the IPv6 NREN service</i>	61
7.3.2.	<i>NREN dual-stack services</i>	62
7.3.3.	<i>IPv6 transport for nameservers</i>	63
7.3.4.	<i>Main network management and monitoring tools</i>	63
7.3.5.	<i>Special security concerns in deploying IPv6</i>	63
7.4.	DEPLOYMENT EXPERIENCE	64
7.4.1.	<i>Mistakes and lessons learnt</i>	64
7.4.2.	<i>Unexpected results from transition</i>	64
7.4.3.	<i>Technology gaps</i>	64
7.4.4.	<i>Standardisation gaps</i>	65
8.	CONCLUSIONS	66
9.	REFERENCES	67
10.	APPENDIX A: HUNGARNET'S IPV6 OVER MPLS CONFIGURATION EXAMPLE (CISCO)	69
11.	APPENDIX B: FUNET'S CORE NETWORK CONFIGURATION EXAMPLE (JUNIPER)	75
12.	APPENDIX C: 6WIN CONFIGURATION EXAMPLES (CISCO)	78
13.	APPENDIX D: SWITCH CONFIGURATION EXAMPLES (CISCO)	82

Table of Figures

FIGURE 3-1: MPLS ROUTING HIERARCHY	17
FIGURE 3-2: BGP AND LABEL ADVERTISEMENT OF A 6PE ROUTER	18
FIGURE 3-3: STRUCTURE OF A MPLS PACKET SENT FROM PE-1 TO P-1	19
FIGURE 3-4: STRUCTURE OF AN MPLS PACKET	20
FIGURE 3-5: STRUCTURE OF A GRE TUNNEL PACKET	20
FIGURE 3-6: STRUCTURE OF AN IPV6 PACKET ENCAPSULATED IN IPV4	20
FIGURE 3-7: DEPLOYMENT OF 6PE IN A PRODUCTION MPLS NETWORK	21
FIGURE 3-8: TUNNEL BROKER OPERATION	24
FIGURE 3-9: THE 6TO4 ADDRESS FORMAT	25
FIGURE 3-10: TYPICAL EXAMPLE OF THE USAGE OF THE 6TO4 MECHANISM.....	26
FIGURE 4-1: GEANT NETWORK TOPOLOGY, DECEMBER 2002.....	31
FIGURE 5-1: LOGICAL TOPOLOGY FOR SURFNET5	36
FIGURE 5-2: SURFNET PREFIXES PER POP	39
FIGURE 5-3: ANONYMOUS-FTP OVER IPV6 VOLUME.	40
FIGURE 5-4: THE FUNET IPV6 NETWORK.....	42
FIGURE 5-5: THE RENATER-3 NETWORK.....	48
FIGURE 6-1: THE 6WIN NETWORK, INCLUDING MÜNSTER	50
FIGURE 6-2: 6WIN AND SITES IN 6WIN	51
FIGURE 6-3: 6WIN CONNECTIONS	52
FIGURE 6-4: 6WIN ADDRESSING OVERVIEW	53
FIGURE 6-5: GRNET CONNECTIVITY.....	56
FIGURE 6-6: THE GRNET MPLS-ENABLED PILOT	57

1. Introduction

In this document we outline the candidate mechanisms that may be used by the European National Research and Education Networks (NRENs) as they plan the introduction of IPv6 services. A concurrent document, D2.3.3, describes the mechanisms available for site (university) networks. Together, the two cookbooks provide a combined theoretical and hands-on guide for the introduction of IPv6 services in NRENs and universities.

In Europe, there are over 25 NRENs who operate production networking for their national academic networks. These are interconnected via the GÉANT (production) backbone network. 6NET, as a research project, has deployed a separate interconnecting (experimental) native IPv6 backbone between up to 15 of the NRENs.

Initially 6NET was the IPv6 backbone between the NRENs for IPv6. As 6NET has progressed, the NRENs have transitioned towards IPv6 services in their national networks, while GÉANT has also transitioned to support an IPv6 service (dual-stack, alongside IPv4). Many NRENs now use GÉANT for production IPv6 traffic, as a result of the experience gained in the 6NET project, while the 6NET backbone is now used for bleeding edge experiments, e.g. new IPv6 Multicast techniques.

The document is an update to the previous 6NET transition scoping report deliverable for NRENs and core networks (D2.2.2) [6NET-D222], which in turn was an update and merger of the previous scoping reports for NREN [6NET-D221] and backbone [6NET-D211] networks..

The document repeats content D2.2.2, but adds new areas including:

- Scenarios for NREN deployment;
- Comparison with IETF work in the ISP transition area;
- Updates to NREN case studies (in particular SURFnet and 6WiN);
- Summary of survey results of NREN deployment status and experience;
- Summary of mistakes made and lessons learnt in that survey.

6NET has many goals. The deployment of an IPv6-only network core, with PoPs located in a large number of NRENs, was the early focus of the work. But to deliver services to the end users in the universities, the NRENs need to deploy mechanisms to allow the integration of new IPv6 services. Most of, if not all, the NRENs already have some IPv6 deployment.

An important goal of 6NET is to foster development of those national IPv6 deployments, to bring the end users online, and to enable end-to-end IPv6 application usage across Europe (and beyond). The earliest of these initiatives have been reported in the IPv6 activities of GÉANT's Task Force: Next Generation Networks (TF-NGN) working group, as Deliverable 9.3 [D9.3] and subsequently Deliverable D9.6 [D9.6] of the GÉANT project. More recent examples are given in Sections 4, 5 and 6 of this cookbook.

The mechanisms that may be used depend on the nature of the existing IPv4 infrastructure, and the goals of the particular NREN in question. These scenarios are discussed in Section 2. Where ATM or MPLS is already deployed, P_v6 can be deployed on top of those technologies, as described in Section 3 of this document. However, one would not normally consider deploying ATM or MPLS to enable IPv6 deployment, unless hardware considerations enforced it. Indeed,

most NRENs have now moved from ATM to PoS and more “modern” technologies, thus use of ATM for IPv6 is expected to be limited, and MPLS is also currently only used by a relatively small number of NRENs. That said, MPLS (6PE) has been used in situations where otherwise (expensive) hardware upgrades would be required to deliver line rate IPv6 (an example is given for SURFnet in Section 5). Being dual-stack does not necessarily mean that IPv4 and IPv6 are handled in hardware, e.g. older Cisco line cards would do IPv6 only in software (then when your IPv6 traffic becomes high, IPv4 service may be affected).

The simplest early transition technique is to deploy an IPv6 infrastructure tunnelled over the existing IPv4 network, leveraging the IPv4 network routing topology and performance, but this is only an intermediate aid to transition. Most NRENs now have a dual-stack native infrastructure deployed, as a direct result of the experience gained through 6NET.

We believe that this dual-stack IPv4 and IPv6 approach on the NREN backbone routers will be adopted by all the NRENs in due course. In a dual-stack deployment those routers hold IPv4 and IPv6 routing tables, possibly share a common IGP (in particular IS-IS), and both protocols run natively on the wire. This is the basis for the examples and case studies in Section 5.

One alternative, in use currently by DFN in the German 6WiN network, is the introduction of a parallel IPv6 infrastructure, as described in Section 6. The parallel infrastructure may include separate links, or just separate routers tunnelling over existing IPv4 links. The choice between dual-stack and a parallel infrastructure has a number of tradeoffs, in terms of factors such as cost of equipment, performance, independence of the networks, and management (tradeoffs which vary with forwarding performance required and traffic levels observed in each protocol). However, the ultimate goal is almost universally a common infrastructure, as it is for those currently also using MPLS (called “6PE” on Cisco hardware) as an interim solution.

NRENs do not generally have to operate translation mechanisms; such mechanisms are applied at the edges of the network, at the university border routers or within the university networks. It would be expected that any university operating IPv6-only network elements would introduce its own translation mechanisms (if translation were the adopted interoperability/integration approach). Translation can occur at a variety of layers, e.g. the network layer (NAT-PT), the transport layer (through transport layer relays), or the application layers (through ALGs such as a web proxy).

Ultimately the exit strategy for transition is an IPv6-only network carrying IPv4 in tunnels, but we expect that scenario to be a distant one, certainly beyond the timeframe of 6NET (which currently is set to conclude in December 2004). This document does not discuss transition to an IPv6-only service (with encapsulated IPv4).

NRENs may deploy support services for IPv6 sites and users, e.g. they may deploy a tunnel broker or a 6to4 relay service, and they may also make a number of their services dual-stack enabled (e.g. DNS, FTP sites, web resources). Such support deployment is discussed in the survey reported in Section 7, in which the status of NREN deployments as of March 2004 is summarised from the responses and feedback of fourteen of the NRENs.

Where new or updated case studies become available, updated versions of this cookbook will be released, particularly where new examples or scenarios may evolve (though this is not expected).

2. Scenarios for NREN/ISP transition

In this cookbook, the transition mechanisms are presented on the basis of examples of transition strategies that have been adopted by the NRENs. As such, the style of this cookbook is “what the NRENs did” towards transition rather than “how a generic ISP would transition”. In the case of the NRENs, the end customers are the university and college site networks. The NRENs do not generally deploy broadband or DSL services out individual end users, and thus the focus of this document is the transport of IPv6 across the NREN networks, rather than (for example) how an IPv6 DSL service would be enabled. The fact that the end customer is typically a large site affects the choice of transition mechanism for that site (e.g. manually configured IPv6 in IPv4 tunnels are used almost universally in preference to tunnel brokers or 6to4, which are more geared for the home/SOHO user).

2.1. Transition scenarios for NRENs

There are a number of different scenarios that an NREN transitioning to support IPv6 may find itself in. In the scope of the NRENs in the 6NET project, these can generally be categorised by the combination of the Layer 2 technology being used, and the routing protocol being used to route the existing IPv4 traffic.

In addition, the NREN will need to provide supporting services for or during transition, either to enable IPv6 connectivity over existing IPv4 networks (e.g. where dual-stack is not yet possible) or to make existing services (e.g. DNS) available over IPv6 transport in addition to the existing IPv4 transport.

2.1.1. Layer 2 transport protocol

There are three basic transport types that may be considered for NRENs:

- PoS (Packet over SONET)
- MPLS (Multiprotocol Label Switch)
- ATM (Asynchronous Transfer Mode)

The most common deployed technology is plain PoS. However, the use of MPLS is becoming more common as NRENs adopt the technology for traffic engineering and QoS purposes. ATM was commonplace prior to the deployment of the GÉANT backbone in Europe (pre-2001) but is now very rare, due to its prohibitive cost and limitations at high line rates.

2.1.2. Routing protocol

The commonly used internal NREN routing protocols for IPv4 are

- OSPFv2
- IS-IS
- iBGP
- RIP (very rare)

OSPF is most common due to its more mature status than IS-IS. However, given there are IPv6 capable versions of IS-IS and OSPFv3, the question for the NREN is whether to use the deployment of IPv6 as an opportunity (or reason) to (for example) migrate from OSPFv2 for IPv4 to use IS-IS for both protocols, or whether a separate protocol should be used for each IP version. Some static routes may be used of course, outside or in addition to the routing protocols.

2.1.3. Additional service support

The NREN is likely to run other services for IPv4, e.g. DNS. While we report on transition status in Section 7 of this document for the supporting services, these are generally considered in other 6NET deliverables (e.g. under WP6 for network management and monitoring [6NET-D632], and the security of transition mechanisms [6NET-D622]).

2.2. Comparison to IETF ISP transition scenarios document

The IETF IPv6 Operations WG [V6OPS] is undertaking a study of scenarios and analysis for IPv6 transition for ISPs [Lind01]. This document is currently at the Internet Draft (I-D) state. As stated above, the scope of the IETF work is broader than the 6NET NREN scope, due to the differing focus of the services deployed by NRENs and commercial ISPs.

The IETF I-D begins by describing stages of transition that an ISP network may go through during transition, and notes that an ISP may transition different elements of its network at different times (e.g. it may start by offering IPv6 in IPv4 tunnels to early adopter customers ahead of a full dual-stack service). The I-D identifies four stages:

- 1: Launch: the ISP obtains an IPv6 SubTLA
- 2a: Backbone: the ISP upgrades its backbone to carry IPv6 traffic
- 2b: Customer connection: the ISP customer(s) connection is upgraded to carry IPv6 traffic
- 3: Complete: IPv6 is available through the ISP and customer network

The I-D suggests that the ISP backbone upgrade can happen as new procurements are made. This is largely true. However some NRENs in 6NET have noted that having dual stack on certain hardware only recently bought can be a problem, and that additional expensive upgrades may be required sooner than expected; as a result such an NREN may choose to use MPLS on existing hardware to carry IPv6 at line rate rather than doing a “double upgrade” (see Section 5 for an example of this in SURFnet).

Clearly a recommendation to procure IPv6-capable equipment is wise, but some platforms have variations in line card capabilities that may be subtle to the buyer.

The I-D considers different backbone technologies:

- Regular: in this case, tunnels may be used initially to be replaced with dual-stack networking as it becomes possible.
- MPLS: this can be deployed as native IPv6-over-MPLS, or as the IPv6-over-IPv4/MPLS method described in [Ooms01]. Ideally MPLS networks should deploy native IPv6 routing and forwarding, and/or use IPv6 LSPs, but using tunnelling over IPv4 LSPs or through the [Ooms01] approach (aka. 6PE) offers an interim step.

MPLS in the NREN context is discussed in Section 3 below. ATM is not discussed in the I-D. Given its general absence from NREN networks now, this seems appropriate.

The I-D considers OSPFv2 and IS-IS as the possible IPv4 IGPs before transition. RIPv2 and iBGP are not considered beyond point-to-point routing. The choice presented in the I-D is the same as the NRENs have generally faced however: OSPFv3 or IS-IS for IPv6. The important question then is whether to have separate routing processes for each protocol (IPv4 and IPv6). Separate processes have more overhead, but offer separation and thus resilience between IPv4 and IPv6 stability. The possible combinations are:

- OSPFv2 for IPv4, IS-IS for IPv6
- OSPFv2 for IPv4, OSPFv3 for IPv6
- IS-IS for IPv4, OSPFv3 for IPv6
- IS-IS for both IPv4 and IPv6 (same process)

The deployment and usage of these protocols is discussed for NRENs in [6NET-D312].

The I-D recommends IS-IS for IPv4 and IPv6, a method already used by three of the NRENs (routing protocol usage is discussed in the summarised survey report in Section 7).

Multicast is somewhat dismissed by the I-D, but has formed a major part of the work of 6NET. For discussion on IPv6 Multicast routing, see [6NET-D312].

The I-D then discusses customer connection, but because ISP customers tend to be home or SOHO sites, the techniques discussed are generally out of scope for NRENs. However, an NREN should offer some transition support for home users (e.g. academics or students on broadband) who may have no commercial ISP IPv6 support and thus who may seek a tunnel broker or 6to4 relay to run their favoured tunnelling solution. This is discussed in Sections 3.7 and 7.2. Specific details of end-site oriented tools are reported in [6NET-D233], e.g. tunnel broker, 6to4, Teredo and NAT traversal.

The I-D does include consideration of large end sites, but doesn't comment in detail. The NREN experience is that large end sites (universities) connect using statically configured tunnels where a dual-stack service to the university PoP is not available (e.g. due to an intervening regional network that is as yet not dual-stack). This turns out to be manageable because the rate of change (addition, deletion or change of endpoint IPv4 address) is quite low.

The I-D cites three example scenarios for transition analysis. The first is an xDSL provider. This is not an NREN scenario, however the recommendations, e.g., to deploy a 6to4 relay for optimised customer 6to4 performance, do map to the NREN scenario (e.g. for home users in a non-IPv6 capable ISP or users in a non-IPv6 capable campus).

The second example is an IPv4 MPLS network; this scenario is in the NREN scope, and is discussed later in this document.

The third example is for a transit provider, which would be equivalent to GÉANT as the inter-NREN backbone, which is presented here in Section 4.

Finally, the I-D closes with security considerations, in three areas:

- Generic best practice security;
- Security concerns through complexity of transition tools (e.g. open 6to4 relays);
- Complexity in managing dual-stack, e.g. ACLs, and in being able to trace customers (e.g. where stateless address autoconfiguration is used, or perhaps more importantly RFC3041 addresses).

In 6NET, we report on security issues in [6NET-D312] and [6NET-D622].

Overall, the IETF I-D contains many similar issues to this document. It does not include specific theory, or case studies (at the same level as in this guide), or issues such as address allocation plans. It of course does not report current usage or configuration examples. However, the overlap is good, and one thus feels the IETF direction is in line with the direction the NRENs are actually taking in practice. [Note that members of 6NET have contributed to the I-D.]

2.3. Requirements for transition

There are a number of requirements that can be identified for an NREN wishing to introduce an IPv6 service:

- The existing IPv4 service should not be adversely disrupted (e.g. as it might be by router loading of encapsulating IPv6 in IPv4 for tunnels);
- The IPv6 service should perform as well as the IPv4 service (i.e. at the IPv4 line rate);
- The service must be manageable and be able to be monitored (thus tools should be available for IPv6 as they are for IPv4);
- The security of the network should not be compromised, due to the additional protocol itself or a weakness of any transition mechanism used;
- An IPv6 address allocation plan must be drawn up (falling under the production SubTLA obtained by the NREN from the RIPE NCC in Europe).

2.4. Solution space for IPv6 transition for NRENs

We discuss the theory behind the potential transition mechanisms for NRENs in the next section, after which specific case studies are given that match a number of the above scenarios.

2.4.1. PoS scenario

In the case of a “plain” IPv4 over PoS network, one might take two different approaches:

- The NREN may see dual stack operation on the NREN core routers as a natural path forward, given appropriate robust code and appropriate studies into the management and operational implications (e.g. as presented in [6NET-D622]). Examples are given in Section 5 for SURFnet, Funet and Renater.
- The NREN may wish to deploy a parallel, non-disruptive IPv6 network. An example is given in Section 5 for the GWiN

Ideally in dual-stack both IPv4 and IPv6 are handled natively at line rate, but this may not be the case. It depends on the vendor/platform architecture.

If moving to dual-stack, the NREN should also consider how long term its dual stack mode of operation would be, if it took that path, and what its exit strategy would be to run IPv6 with IPv4 carried as tunnelled traffic in the IPv6-only national backbone. That scenario is however some time away (probably five years, at the very least).

2.4.2. MPLS scenario

The NREN may have an existing MPLS network. There are some specific options for running IPv6 over MPLS (this is a scenario that has been used by Cisco with 6PE [Ooms01] for example). Discussion of this method is given in Section 3. The GRNET and CESNET NRENs initially considered this method. MPLS would not normally be deployed purely to enable IPv6. However if the upgrade path for hardware to enable IPv6 at line rate is expensive, MPLS is an option to carry IPv6 natively at line rate; for this reason SURFnet then deployed 6PE.

2.4.3. ATM scenario

The NREN may have an ATM network, and thus be able to run IPv6 over parallel PVCs. Discussion is given in Section 3, but no NRENs are currently using this method (outside of very specific testbeds that would not be part of a production network). ATM would not be deployed purely to enable IPv6.

Solutions are discussed in more detail via case studies in Sections 4 (for the GÉANT “transit” backbone), 5 (dual-stack) and 6 (parallel infrastructure).

3. Review of NREN/ISP transition mechanisms

In this section we discuss the range of transition and integration techniques available to NRENs, including some of the support mechanisms that an NREN may operate for universities (e.g. a tunnel broker or a 6to4 service with 6to4 relay).

3.1. General Approach

A general approach for an ISP looking to transition is given in the IETF v6ops WG document [Lind01].

In the case of a 6NET NREN this approach may be broken down into broad stages as follows.

1. Obtain IPv6 address space: most likely a /32 SubTLA (from RIPE NCC in Europe);
2. Devise an IPv6 address allocation plan for the NRENs network and the end-site universities (see the examples in the case studies in this document);
3. Study the available tools for network management and monitoring (see [6NET-D632]) and establish any required new operational procedures;
4. Select the appropriate transition path for IPv6 transport over the NREN network infrastructure (the theory for which is discussed below in this Section, but may include direct transition to dual-stack, use of MPLS/6PE, or perhaps ATM PVCs);
5. Select the appropriate IPv6 routing protocol (see [6NET-D312]) and decide the routing policy (which may be the same as IPv4);
6. Deploy any necessary transition aids (e.g. tunnel broker or 6to4 relay, see [6NET-D233]);
7. IPv6-enable any required services (e.g. DNS, QoS or Multicast see [6NET-D312]);
8. Follow the best practice for secured transition mechanism deployment (see [6NET-D622]);
9. Apply steps [#2-#8] for any regional networks attached to the NREN backbone between the NREN and end sites (universities);
10. Enable the equipment in the end-site premises;

If a native dual-stack approach is not enabled initially, the interim method (e.g. tunnelling or 6PE) should be upgraded to IPv6 native when procurement allows.

It would be expected that the resilience and robustness of the platform used would be tested in a testbed environment before production deployment (this has been a very valuable feature of the 6NET testbed itself, of course). An example of this is described in Section 4.

The specific transition path may depend on existing technology and equipment/software available, as well as non-technical issues such as available budget.

3.2. Dual-stack

The term "dual-stack" is a broad term, and can be used to mean many things. When performing dual-stack transition at the NREN scope, the ultimate goal is to make the same production routers route and forward both IPv4 and IPv6 traffic (and thus not create any kind of virtual overlay network, e.g. via 6PE, AToM or ATM PVCs).

The steps in deploying a dual-stack IPv4-IPv6 network may include:

- Create and operate a test network with IPv4-tunnelled (or even MPLS or ATM) connections to gain perspective on the operation of IPv6.
- Evaluate the router software versions in the test environment to see if they are stable and robust enough to be used in the main network with IPv4 and IPv6 together, and how IPv6 affects IPv4 performance.
- If they are stable, start upgrading production routers to IPv4/IPv6, and enable IPv6 on the links that are used. Usually the network topology will be the same as with IPv4.
- If problems (e.g. severe bugs affecting production services) arise, either try to fix or avoid them or drop back to an IPv4-only operation.

In this document we report on three of the 6NET participating NRENs that have completed the process (SURFnet, Funet and Renater) in Section 5. The European NRENs in the 6NET project are amongst the vanguard of production IPv6 deployment.

Some router vendors have now been shipping IPv6 capability in their production software for some time, which has gained in stability for core features as a result of experience gained in deployment in the 6NET NRENs. This is a big advantage for emerging subsequent commercial ISP deployment.

In order to break the "chicken and egg" status, it would be desirable to deploy IPv6 in advance of heavy demand, as there will not be demand until the service is well-supported by the NRENs. Indeed, in academic networks, a commercial case for deployment is not generally required – the service is provisioned to enable research. By introducing dual-stack networking throughout the core (6NET and GÉANT) and NREN networks, provision and deployment issues are pushed to the edge, such that they become a per-university issue.

The advantage of dual-stack operation is that the network is the same for IPv4 and IPv6: there need not be new routers for IPv6, and there is no need to maintain a potentially complex overlay network. Similarly, this is also a disadvantage: because the network is the same, the problems (especially software bugs), if such arose, could also affect IPv4 services, which would probably not be the case if the network was separated. One also has to be aware of the performance impact of running IPv6 in the IPv4 service, especially if the dual-stack implementation is not in hardware and IPv4 routers are encapsulating IPv6 in software.

3.3. General tunnels

There are two classes of general tunnelling techniques, IPv6-in-IPv4 tunnels, and Layer 2 tunnels (including encapsulation methods such as AToM, CCC or UTI). Such tunnels can be considered to belong to the same family as MPLS or ATM, only the "link-layer" technology is different.

There is an advantage to using IPv6 tunnels over the existing IPv4 infrastructure – namely that the infrastructure is already well-tuned by the NREN to perform well; thus even with the tunnelling overhead, the IPv6 overlay should perform sufficiently well. However, where the tunnels are configured manually, it is quite possible that the tunnels do not always take an optimal path

between sites, where one IPv6 hop may span many IPv4 hops. Automatic (6to4) tunnelling is described in a later part of Section 3, but is very rarely used by the NRENs for connection of large university sites (where static tunnels are used if native connectivity is not available).

The dependence on the existing IPv4 infrastructure may be a weakness, e.g. software problems, denial of service attacks against routers, etc, would also affect the IPv6 service. That said, one would expect the production IPv4 service to be well supported, so such issues ought to be rare.

In IPv6, path MTU discovery and management, fragmentation and reassembly is handled by the end hosts, not intermediate routers. This has a number of implications. For example, in Ethernet networks, on Fast Ethernet interfaces, the MTU is 1500. Using tunnelling on the interfaces where the MTU is 1500 reduces the usual path MTU to 1480 bytes, which will add some latency as path MTU discovery is initiated.

3.4. IPv6 over MPLS

We do not expect IPv6 provisioning over Multi Protocol Label Switching (MPLS) to be a common method for the European NRENs. To date, only three NRENs (GRNET, HUNGARNET and CESNET) have reported on studies of IPv6 over MPLS. There are two likely usage scenarios:

- Where an MPLS network exists, the method may be most appropriate for the NREN;
- If hardware upgrade for line rate IPv6 is expensive, running IPv6-over-IPv4/MPLS (6PE) may be a good compromise for line rate IPv6 (avoiding encapsulation in software tunnels), as described in the SURFnet case study in Section 5.

Backbone networks that have already deployed MPLS might consider several IPv4-IPv6 migration strategies:

- *Native IPv6 over MPLS*: In this scenario, IPv6 transport over an MPLS network is completely symmetric to the IPv4 case. It requires that all routers in the MPLS network become dual-stack and use IPv6 routing protocols (both interior and exterior) together with IPv6-enabled Label Distribution Protocol (LDP);
- *L2 tunnelling over MPLS*: The underlying technology is based on the IETF draft [Martini02]. Entire L2 frames (e.g., Ethernet with IEEE 802.1q encapsulation, ATM AAL5 etc.) are switched across the MPLS core, hence the L3 protocol is completely transparent. This feature is available, in one form or another, on most major routing platforms including Cisco IOS and Juniper JunOS;
- *IPv6 over IPv4/MPLS core*: This method is a special case of BGP tunnelling as described in IETF draft [Ooms01]. It relies on the distribution of IPv6 prefixes (and corresponding labels) among the edge Label-Switching Routers (LSR) using standard BGPv4 over IPv4, where the Next Hop is identified by an IPv4 address. Cisco Systems implements this functionality under the name *6PE (IPv6 Provider Edge Router)*, which has now been applied generically to the draft [Ooms01].

From the above approaches, only the latter one is of immediate interest to 6NET partners. Native IPv6 over MPLS is currently available on minority or development platforms (ZebOS, AYAME) while major router vendors seem to have no plans for adding IPv6 support to LDP in a foreseeable future (although RFC 3036 contains a specification, there is little demand for it at present). On the other hand, L2 tunnelling over MPLS brings nothing new from the perspective of IPv6 and, moreover, is not really suitable for wide-area networks.

3.4.1. 6PE – theory of operation

The concept of 6PE stems from the canonical routing hierarchy of an MPLS network, which is illustrated in Figure 3-1.




Figure 3-1: MPLS routing hierarchy

In this hierarchy, the core of the network consists of so called P (Provider) routers that switch MPLS packets and thus, for the most part, do not parse the L3 header. At the edge of the MPLS core we find PE (Provider Edge) routers. They receive standard IP packets from CE (Customer Edge) routers, impose an MPLS label¹ according to their MPLS forwarding table and send the packet to the appropriate P router. Therefore, MPLS packets travel only across PE-P and P-P links (thick lines in Figure 3-1). P and PE routers are together denoted as Label Switching Routers (LSR). Routing is performed in three relatively independent levels:

1. Between PE and CE routers, any of the common routing protocols may be configured (RIP, OSPF, BGP or even static routing). Using this routing protocol, the PE router learns the prefixes that are reachable through each CE router.
2. PE routers exchange these prefixes among each other via IBGP sessions. Depending on the situation, either a full mesh of BGP sessions may be established or route reflectors [RFC2796] may be used. In any case, each PE router advertises the (summarised) prefixes learned from attached CE routers to all other PEs as NLRI in BGP and inserts itself as the next hop for these prefixes.
3. Consequently, each PE router must also be able to determine the route to each potential BGP next hop (another PE). This is accomplished by an interior gateway protocol like IS-IS or OSPF. This protocol involves exactly all P and PE routers and its routing database usually forms the basis of the MPLS forwarding table. In other words, Label-Switched Paths (LSP) between PE routers are initially constructed from the information provided by this IGP. In

¹ Unless the destination is behind a CE router that is attached to the same PE router.

order to achieve stability of label assignments in the network core, only host routes to the P and PE routers should be included in this IGP.

As opposed to the typical configuration of IGP and IBGP in an autonomous system, in this setup the P routers are completely unaware of any external routing information so that

- synchronisation between IGP and IBGP, which is suggested by [RFC1772], must be turned off – the sets of prefixes in the two routing protocols are practically disjoint.
- P routers can not perform the usual “hot potato” IP routing to external prefixes – MPLS is the only method they can use for forwarding traffic to external destinations.

The 6PE concept leaves the MPLS core (P routers) intact and assumes that PE routers become dual-stack. CE routers may be dual-stack or IPv6-only and use again any of IPv6 routing protocols for advertising local IPv6 prefixes to the PE router they are directly connected to. PE routers readvertise this reachability information into IBGP, this time under IPv6 address family. IBGP sessions are transported over TCP/IPv4 as before and the backbone OSPF process is also unchanged. Each PE router thus identifies itself as next hop using its IPv4 address. However, the next hop field in BGP UPDATE messages must be of the same address family as the NLRI, which is IPv6 in this case. Therefore, the IPv4 next hop address is encoded as “IPv4-mapped IPv6 address” [RFC2373]. In the same BGP UPDATE message, the PE router also includes the MPLS label associated with the prefix using the method of [RFC3107]. The ingress PE router uses this label as the inner MPLS label – the outer one in the standard label learned from IGP, which is used for forwarding traffic towards the egress PE (the next hop for the destination).

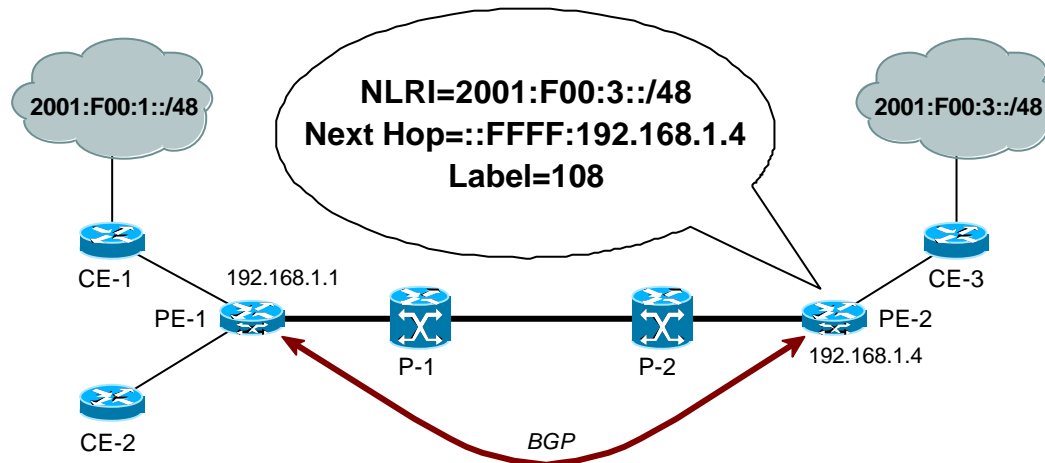


Figure 3-2: BGP and label advertisement of a 6PE router

An example is shown in Figure 3-2. Router PE-2 advertises to PE-1 the reachability of destinations under $2001:F00:3::/48$ and announces itself as the next hop for these destinations using the IPv4-mapped address $::FFFF:192.168.1.4$ and also advertises the binding of label 108 to the prefix.

Now assume PE-1 receives a datagram, e.g., from CE-1, with destination in $2001:F00:3::/48$. Using the above information, it finds the BGP next hop and looks up the information for *IPv4 address* $192.168.1.4$ in its MPLS forwarding table, obtaining thus the outgoing interface, the

IGP next hop (P-1) and the outer label, say 55. The MPLS packet sent to P-1 then has the structure shown in Figure 3-3.




Figure 3-3: Structure of a MPLS packet sent from PE-1 to P-1

This packet is then switched to P-2 in a normal MPLS way where only the outer MPLS header is inspected and its label rewritten. P-2 then performs *penultimate hop popping (PHP)*, i.e., removes the outer MPLS header so that PE-2 receives the packet with a single label. PE-2 removes this label and performs standard forwarding based on the IPv6 destination address etc.

The above scenario is quite similar to packet forwarding in MPLS/BGP VPNs [RFC 2547]. The difference is that in the latter case the inner label carries the information about the VPN Routing/Forwarding (VRF) instance the packet belongs to. For the basic 6PE it is not the case and so the question naturally arises whether the second level of MPLS labels is really necessary if the egress PE router does IPv6 header lookup anyway. Indeed, the inner label is not required for this mechanism to work, it only helps to keep the MPLS core unaffected. In particular, without the inner label the “penultimate hop” P router would have to be able to forward a plain IPv6 datagram to the egress PE router. Moreover, a VPN-based extension of the 6PE mechanism is currently under discussion in IETF and then the inner label will carry important forwarding information.

3.4.2. Comparison to other transition technologies

The 6PE approach differs from other IPv6 tunnelling techniques in two main aspects:

- tunnel set up
- tunnel overhead

Of course, the most significant difference is the whole MPLS control plane but we have to assume it is in place anyway – it would probably not make sense to deploy MPLS only for the sake of 6PE unless hardware constraints dictated otherwise, e.g. one NREN has a backbone running at 10 Gbit/s at 15 PoPs and no Cisco Engine 5 line cards are available yet from Cisco – thus there is no solution for becoming wire speed for IPv6 in hardware and 6PE becomes viable for that purpose (although a far from ideal solution, OC-192c POS line cards are not cheap!).

3.4.2.1. Tunnel set up

The 6PE mechanism relies on MPLS to create tunnels through the IPv4/MPLS core. The basic variant, where MPLS tunnels follow the IGP shortest paths, is almost equivalent to the 6to4 technique [RFC3056], except that the latter needs special IPv6 addresses. However, 6PE can also use MPLS tunnels (LSPs) configured by other means, e.g., traffic engineering.

3.4.2.2. Tunnel overhead

Tunnel encapsulations by definition add new protocol header and thus increase the transmission overhead. We can directly compare the overheads of the 6PE MPLS encapsulation and the two most common tunnelling methods:

- We saw previously that the 6PE encapsulation involves two labels, which are – at least in the common cases of Ethernet or PoS link layers – contained in two “shim” headers, each 4 bytes long. The structure of the entire packet is shown in Figure 3-4.




Figure 3-4: Structure of an MPLS packet

- One option for configured IPv6 tunnels is the Generic Routing Encapsulation (GRE) defined in [RFC2784]. In this case, the overhead consists of the GRE header (4 bytes) and outer IPv4 header (mostly 20 bytes) as shown in Figure 3-5.



Figure 3-5: Structure of a GRE tunnel packet

- In the IPv6 world, the most common tunnelling technique is the direct encapsulation of IPv6 datagrams in IPv4 [RFC2893], which is used for both configured tunnels and automatic 6to4 tunnels [RFC3056]. Figure 3-6 indicates that the overhead is just the IPv4 header of 20 bytes (again assuming no IP header options).



Figure 3-6: Structure of an IPv6 packet encapsulated in IPv4

Table 3-1 shows overhead percentages for different lengths of the tunnelled IPv6 datagram. For small packets the 6PE encapsulation is clearly superior whereas for packet sizes close to the MTU of common backbone media types (Gigabit Ethernet or PoS) the difference becomes essentially negligible.

Datagram size [Bytes]	6PE/MPLS	GRE	IPv6 in IPv4
40	20.0%	60.0%	50.0%
200	4.0%	12.0%	10.0%
750	1.0%	3.2%	2.7%
1500	0.5%	1.6%	1.3%
3000	0.3%	0.8%	0.7%
4470	0.2%	0.5%	0.4%

Table 3-1: Overhead ratio for different encapsulation methods

We already mentioned that the 6PE technology by itself does not justify the deployment of the MPLS control plane. However, backbone networks that already use MPLS in their core might

benefit from 6PE, since, apart from the lower overhead, it fits very well into the general MPLS philosophy, being actually one of very few instances of “multiprotocolness” of MPLS towards the network layer.

As long as the 6PE features are included only in experimental releases of routing software, it would be rather risky to install it on production PE routers. Instead, one could use the configuration shown in Figure 3-7.




Figure 3-7: Deployment of 6PE in a production MPLS network

In this case, the original IPv4/MPLS infrastructure requires no hardware or software modification and the probability of 6PE routers interfering with the production network functions is reasonably small. Actually, the added 6PE routers can induce problems to the production backbone only in the limited area of MPLS forwarding and LDP sessions between adjacent 6PE and P routers.

3.4.3. IPv6 over MPLS on HUNGARNET

HUNGARNET has deployed IPv6 over MPLS. The key points of the configuration are:

- IPv6 connectivity provided via Metropolitan Ethernet VLANs, 6PE and IPv4 tunnels.
- Use of OSPFv2 for distributing IPv4 address of loopback interfaces of PE routers.
- Use of ISIS to exchange IPv6 loopback addresses and ISIS passive interfaces. There is no Globally Aggregatable Address assigned to backbone point-to-point interfaces. All loopback interfaces, interfaces connecting to the HUNGARNET members, IPv4 tunnels and LAN stub interfaces were made passive.
- A BGP route reflector is being used.
- 6PE is used to propagate IPv6 reachability over the IPv4/MPLS core of HUNGARNET backbone. TDP is used for propagating labels related to IPv4 prefixes. Labels related to IPv6 NLRI are exchanged between 6PE routers through BGP adding


```
neighbor {peer-group-name | ipv6-address} send-label.
```
- An IPv6 source interface was also configured on the 6PE for IPv6 traffic generated locally.

3.4.3.1. Configuration example

The configuration example on *pecs.6net.hbone.hu*, the 6PE router, is given in Appendix A.

3.5. IPv6 over ATM

Using ATM for transition to IPv6 is very similar to using MPLS. Indeed, one could say MPLS technology was derived from ATM. Both use an overlay network model, where the core network elements (be they ATM switches or MPLS routers) do not need to know anything about IPv6 to support encapsulating IPv6 packets to ATM/MPLS. Only the edge network devices (in both cases, certain routers) need to be IPv6-aware.

In [6NET-D211] we also showed an example of IPv6 over ATM over MPLS, but that is not reproduced here, being a somewhat specialist technique.

Many networks have traditionally been built on ATM infrastructure. As with MPLS, it is not reasonable to build an ATM network just for a transition to IPv6, but if an ATM network does exist, it can be used in the early phase. However, with IPv6 support appearing in many vendor hardware products (e.g. Cisco, Juniper, Hitachi), the length of that early phase for new adopters is shortening, and performance is becoming a secondary issue to those such as management, addressing and routing. So much so, that currently no NREN is using IPv6 over ATM for transition (thus this technique may be removed from the final version of this cookbook).

3.5.1. Permanent and Switched Virtual Circuits

ATM provides PVC (Permanent Virtual Circuit) and SVC (Switched Virtual Circuit) capabilities. The former provides a facility for a network administrator to create a statically configured channel between (usually) two endpoints connecting to the ATM network. The latter provides a mechanism for ATM-capable network devices to create and delete channels automatically, when needed. This is a more complex operation, and has not really been implemented with IPv6; however, it typically isn't generally that useful, as usually one only wants to set up relatively long-lived circuits in the ATM network. Something in between is called "soft PVC" or sPVC: this is a permanent circuit that need not be configured in all the ATM devices between the two endpoints; it is only specified which ATM endpoints are used, and the ATM network protocols find the shortest route and reroute if necessary. This can simplify the set up of ATM connections.

Fixed PVCs provide the capability for an endpoint to perform an on-demand connection to a remote IPv6 node using an ATM address (NSAP or E.164) and employing a mechanism provided by UNI 3.0/3.1/4.0 in the ATM Forum. A lot of telephony and ISP networks have been implemented using ATM infrastructure, outside of the academic world. A switched ATM infrastructure does not generally match a provider's needs, which are usually fulfilled by PVCs. An SVC infrastructure is better suited in site and enterprise topologies where connections tend to have a temporary (transient) nature.

3.5.2. Tunnel setup and tunnel overhead

IPv6 over ATM relies on ATM infrastructure to create tunnels using either ATM PVCs, Soft PVCs or SVCs. Tunnels appear like virtual interfaces with the property that traffic between the end points is transported using AAL5 encapsulation. Currently only AAL5-SNAP encapsulation is supported. The overall overhead for the ATM header plus AAL5 encapsulation header is approximately 22–24%.

3.5.3. IPv6 transition with ATM

As a transition technique, an ATM network should, as with MPLS, not be built solely to provide an approach for IPv6 transition. If it does exist, it can be employed to provide an affordable (short-term) solution. Regarding the operation of IGP and EGP in such a transition mechanism, all the

customers of NRENs appear as physically connected to an IPv6-capable router. There is not any need for a particular IGP protocol, while an ordinary MP-BGP with IPv6 unicast NLRI family capabilities is required to provide EGP routing between the connected customers.

Many NRENs ran ATM networks in the timeframe of TEN-155, but as GÉANT has deployed 2.5Gbit/s (and faster) PoS, the NRENs have also mostly phased out ATM in favour of PoS networks, in part for relative simplicity of management, but also for cost reasons. However, some NRENs, e.g. Renater, have kept some ATM infrastructure specifically for IPv6 testing, given the ability to set up a native IPv6 link over an ATM PVC (although Renater has since gone dual-stack in its new Renater3 deployment – see Section 5).

Routers that have ATM connectivity can be incrementally transitioned to IPv6 by upgrading them to software that supports IPv4/IPv6 dual-stack operation. IPv6 connections to other IPv6-routers can then be made through the ATM network using PVC's or sPVC's. Such routers can then provide IPv6 connectivity to the connecting organizations with ATM or some other means (for example, IPv4 and IPv6 on the same Ethernet links).

In practice, creating ATM circuits is analogous to creating IPv6-in-IPv4 tunnels, but one just uses ATM infrastructure for this, which is not necessarily or often dependent on IPv4. Another feature is that the encapsulation does not decrease the path MTU as the packet size in ATM networks is more than 1500 bytes.

In summary, ATM (where available and feasible) can be used as a good IPv6-in-IPv4 tunnel replacement at least in the early phases where dual-stack protocols “on link” is not yet a realistic option. The advantages are that ATM does not depend on IPv4 as tunnelling does and it has higher IP-level packet size. The disadvantage is that ATM technology does not support really high-speed interfaces, and it is technologically a dead-end (it is not realistic to invest money in it), but is still very usable for years to come.

3.6. Deploying a parallel IPv6-only network

Deployment of a separate IPv6 network may also be an option for some NRENs. In some environments, e.g. where there are tight service-level agreements in IPv4, or there would have to be a great deal of IPv6 education for IPv4 operators, deploying a parallel infrastructure may be an appropriate early method.

However, given there is a cost to having two sets of routers, and two sets of links if the network is fully parallel, a more detailed cost analysis would be necessary to see whether this makes financial sense. In some cases, it might (e.g. if IPv4 routers are all upgraded and the older ones are left unused, and they would be sufficient for IPv6, and if line costs would not be too much during a low bandwidth early adopter period). Building such a network would, however, practically require that it would be used considerably, so people could see the justification for the costs. This may or may not happen. Also, it seems probable that the users, being used to high-speed connections, would not settle for anything significantly less than with IPv4, e.g. “low cost” serial lines would probably be out of the question.

In an early deployment, with low levels of IPv6 traffic, a PC-based parallel infrastructure (e.g. based on BSD and Zebra) may suffice. However, this is not a long-term solution where IPv6 traffic levels grow. This scenario is discussed in Section 6 considering the example of the DFN 6WiN network. A similar method can also be used for IPv6 deployment in site networks, with a

parallel IPv6 routed hierarchy feeding IPv6 prefixes into shared IPv4 VLANs for dual-stack networking (see [6NET-D233]).

3.7. Support mechanisms: Tunnel Broker

In addition to the transition mechanisms already mentioned in this section, there are also supporting mechanisms that, in the absence of such services on a site (university), may be operated by an NREN for the benefit of its community. Two examples of such services are the tunnel broker, discussed here, and 6to4 (with an associated 6to4 relay), discussed in the next section.

The tunnel broker (described in [6NET-D233]) allows a user on a dual-stack host to connect to a web server to request a tunnel connection from their (dual-stack) workstation to the broker's IPv6 network (and wider IPv6 internet). After filling in information on their IPv4 address and operating system (the tunnel creation commands on the user's host would be OS-specific), the user can download a script that can then be run to establish an IPv6 connection (an IPv6-in-IPv4 tunnel) to the tunnel broker's tunnel server.

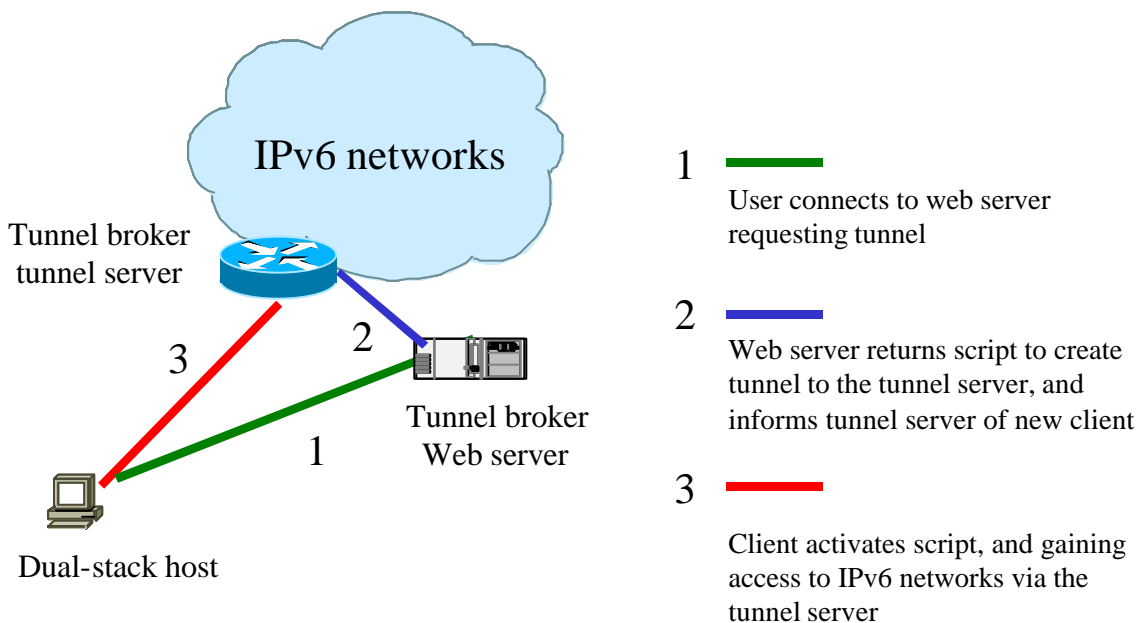


Figure 3-8: Tunnel broker operation

The tunnel broker setup process is illustrated in Figure 3-8. The key requirement for the broker is that the broker has a mechanism to remotely set up tunnel end points on the tunnel server. In a PC-based implementation, the server and broker could be co-located.

At an early stage of IPv6 deployment within an NREN, rather than having each university manage its own tunnel broker, the NREN could operate a tunnel broker for the benefit of its member universities. This would be preferable to the users seeking tunnels from overseas networks (e.g. popular tunnel brokers such as Freenet6 in Canada), where routing would be far from optimal (due to the significantly long first hop).

In [6NET-D233] tunnel brokers from Lancaster University (Microsoft tool based), the University of Southampton (open source based using Apache2, OpenLDAP and ssh) and the University of Muenster/JOIN project (using OpenVPN) are described.

3.8. Support mechanisms: 6to4 and 6to4 relay

The 6to4 transition mechanism [RFC 3056], also described in [6NET-D233], is a flexible mechanism that enables communication between IPv6 islands over the IPv4 Internet. Its usage is expected to be most common during the medium-term phase of the transition process to IPv6, when there are many IPv6 islands, but no wide scale native IPv6 connectivity exists.

3.8.1. Architecture

The 6to4 mechanism has been assigned the 2002::/16 IPv6 prefix. Any organization that wants to enable the IPv6 protocol can use this prefix in order to gain its own /48 IPv6 prefix without needing to request production address space (under 2001::/16) from its associated NREN or the RIPE NCC; only a single global IPv4 address is needed.

Construction of the 6to4 IPv6 prefix is made by the concatenation of the 2002::/16 prefix and the global IPv4 address. The format of the 6to4 IPv6 address is illustrated in Figure 3-9.




Figure 3-9: The 6to4 address format

When connecting to another 6to4 site, because the IPv4 address of the far tunnel endpoint is encoded inside the IPv6 destination address, no direct configuration is needed, since the connection to the remote site can be established as an “automatic” tunnel when required. In most cases the IPv4 address that is used is the address of the interface that connects the organization to the Internet, but if only part of the site is IPv6-enabled, it could be an internal site router. Inside the organization’s network the IPv6 prefix can be used like any other IPv6 prefix for subnetting and autoconfiguration tasks (because the general allocation for any site is a /48 prefix, whether from 6to4 or production SubTLA address space).

By using the 6to4 mechanism any site can fully deploy the IPv6 protocol, and any 6to4 site can easily communicate with other 6to4 sites. However, communication with the native IPv6 world, and IPv6 sites under non-6to4 address space (i.e the production 2001::/16 space or the 6bone 3ffe::/16 space) is achieved only with the deployment of a 6to4 relay router. The relay router is a router connected to the non-6to4 IPv6 network that has been enabled for the 6to4 mechanism, i.e. it can communicate with both 6to4 and non-6to4 IPv6 sites.

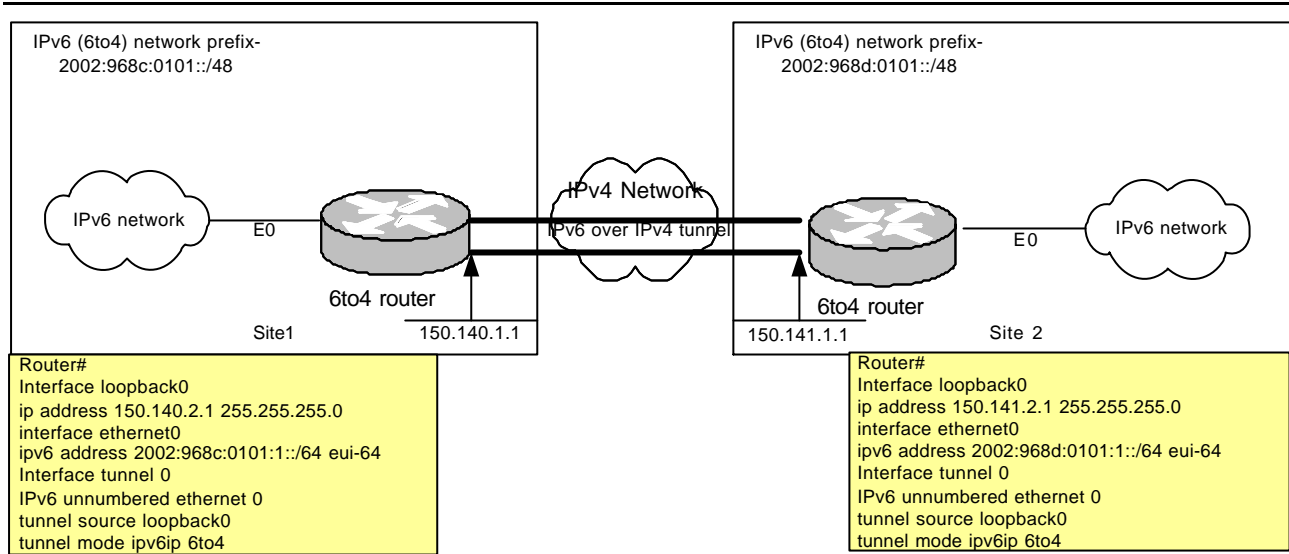


Figure 3-10: Typical example of the usage of the 6to4 mechanism

A typical example that demonstrates the usage of the 6to4 mechanism is shown in Figure 3-10. Also in the same figure the configuration that is needed for the two routers (Cisco-like) is shown.

In order for Site 1 and Site 2 to communicate with the IPv6 world beyond other 6to4 sites, they must deploy a 6to4 relay router either in their premises or use an external (public) 6to4 relay service. It is that 6to4 relay router that the NREN can deploy as a supporting mechanism to the university sites.

The relay advertises 2002::/16 to the native IPv6 internet, and is reachable from 6to4 routers inside its “domain” by use of a well-known anycast address (although the 6to4 router may be manually configured with the 6to4 relay’s real IPv4 address). The 6to4 relay router can be deployed on an IPv4 anycast address, allowing multiple relay routers to be available.

Regarding 6to4 relay reachability, the 192.88.99.0/24 block is allocated for use as 6to4 relay anycast addresses, according to [RFC3068]. In 6NET, NRENs deploying 6to4 relays are using 192.88.99.1 as the advertised relay (anycast) address.

3.8.2. Considerations for 6to4

The advantages of deploying 6to4 services can be summarised as follows:

- The 6to4 mechanism is a very powerful tool that enables IPv6 deployment in a corporate network with relatively low resource requirements. Connections to other 6to4 sites are established by automatic tunnels, while other connections are directed to a 6to4 relay router.
- Any site by deploying the 6to4 mechanism can easily enable the IPv6 protocol inside its network independently of its ISP support to IPv6 (i.e. if the site has no native IPv6 service offered, 6to4 is a practical solution).
- It is a very scalable mechanism that can be deployed from a single host case to an entire network; 6to4 can be used for a whole university, a student household, or just a single host.
- By deploying the 6to4 mechanism any organization gains a full functional IPv6 prefix without having to request one from its NREN or the registries.

For a university, the immediate alternative to 6to4 is a manually configured tunnel to a router operated by the NREN. This would have the advantage of using production IPv6 address space, which would be used anyway when the link became native IPv6. If this choice is adopted by the university, the NREN's 6to4 relay router is still useful for the university to reach 6to4 sites within the same country (or beyond). In this case, the relay router can advertise the 2002::/16 prefix, since it is able to reach any 6to4 site.

The reality of deployment at present in 6NET NRENs is that manually configured tunnels are used for university connections. The use of 6to4 is constrained to home/SOHO users, or to universities where only individuals (rather than the administrators) have a big interest in IPv6.

3.8.3. Operational and security issues

The IPv4 address should be a static long-lived address, since a dynamic address implicitly renumbers the entire IPv6-island (IPv6-prefix).

Also, 6to4 only works in the (today's) world where IPv4 connectivity is complete, since 2002::/16 can be announced by a million different relays. To guarantee connectivity, 6to4 sites have to accept traffic from any relay. That's where 6to4's big security problem comes from - there is nothing to stop "evil" people from creating bogus relays generating DDoS attacks. The NREN should be careful in checking who can use the relay router.

Security concerns with 6to4 and 6to4 relays are also described in [6NET-D6.2.2].

One should also be careful about using RFC 1918 addresses since 6to4 makes no sense with non-routable addresses. However, 6to4 can work between IPv4 NAT gateways where each NAT has a global IPv4 address, and devices behind the device then become dual-stack (IPv4 NAT and IPv6 global with a 6to4 prefix).

3.8.4. Configuration example for 6to4

In this subsection we show the example code for a 6to4 router; in the next section we describe more fully examples for 6to4 relays.

3.8.4.1. Cisco IOS: 6to4 router

The following example configuration enables 6to4 on IOS:

```
interface Tunnel0
no ip address
no ip redirects
ipv6 unnumbered FastEthernet0
tunnel source FastEthernet0
tunnel mode ipv6ip 6to4
!
ipv6 route 2002::/16 Tunnel0
```

3.8.4.2. BSD 6to4 relay configuration, with Zebra

On BSD platforms running Zebra, 6to4 configuration, with the relay on the anycast address of 192.88.99.1, can be done using the following example configuration:

In /etc/rc.conf (relevant parts):

```
ifconfig_xl0_alias0="inet 192.88.99.1 netmask 0xffffffff00"
stf_interface_ipv4addr="192.88.99.1"
stf_interface_ipv6_ifid="::"
ipv6_gateway_enable="YES"
```

In /etc/rc.local (relevant parts):

```
--8<--
# add anycast flag to 6to4 anycast
ifconfig stf0 inet6 2002:c058:6301:: prefixlen 16 anycast
# start zebra and bgpd.
/usr/local/sbin/zebra -d
/usr/local/sbin/bgpd -d
/usr/local/sbin/ospfd -d
/usr/local/bin/vtysh -b
--8<--
```

In /usr/local/etc/Zebra.conf (relevant parts):

```
--8<--
router bgp 1741
  no bgp default ipv4-unicast
  neighbor 2001:708::2 remote-as 1741
  neighbor 2001:708::2 description 6net-rtr
  neighbor 2001:708:0:1::625 remote-as 1741
  neighbor 2001:708:0:1::625 description v6-rtr
!
address-family ipv6
  network 2002::/16
  neighbor 2001:708::2 activate
  neighbor 2001:708::2 soft-reconfiguration inbound
  neighbor 2001:708:0:1::625 activate
  neighbor 2001:708:0:1::625 soft-reconfiguration inbound
  exit-address-family
!
```

```
router ospf
  ospf router-id 128.214.231.106
  network 128.214.231.104/29 area 3248883125
  network 192.88.99.0/24 area 3248883125
!
```

Future versions of the cookbook will report on 6to4 operational experience (either here, or in the university cookbook document [6NET-D233]).

4. GÉANT IPv6 transition overview

GÉANT is the existing production IPv4 network that connects the European NRENs. It is an IST project in its own right, and operated by DANTE on behalf of the GÉANT consortium. It was committed by the project's GN1 programme to deliver an IPv6 service by November 2004.

Due in a large part to the assistance of the 6NET community, GÉANT was able to connect its first set of NRENs in March 2003 for native IPv6 accesses, and offer a production service by the end of 2003. In this section we review the processes and design choices in the transition.

At the time of writing there are 18 of the 25+ NRENs connected to GÉANT's IPv6 service natively.

This overview is presented here in the absence of a 6NET backbone transition, because the 6NET backbone is IPv6-only from the start, thus a dual-stack transition is not applicable.

4.1. Dual stack considerations

A dual stack network is a network that can forward at least two different types of IP packets. The technical decision to go to a dual stack network or to an "encapsulated model" depends on the performance of the routers and the design of the network.

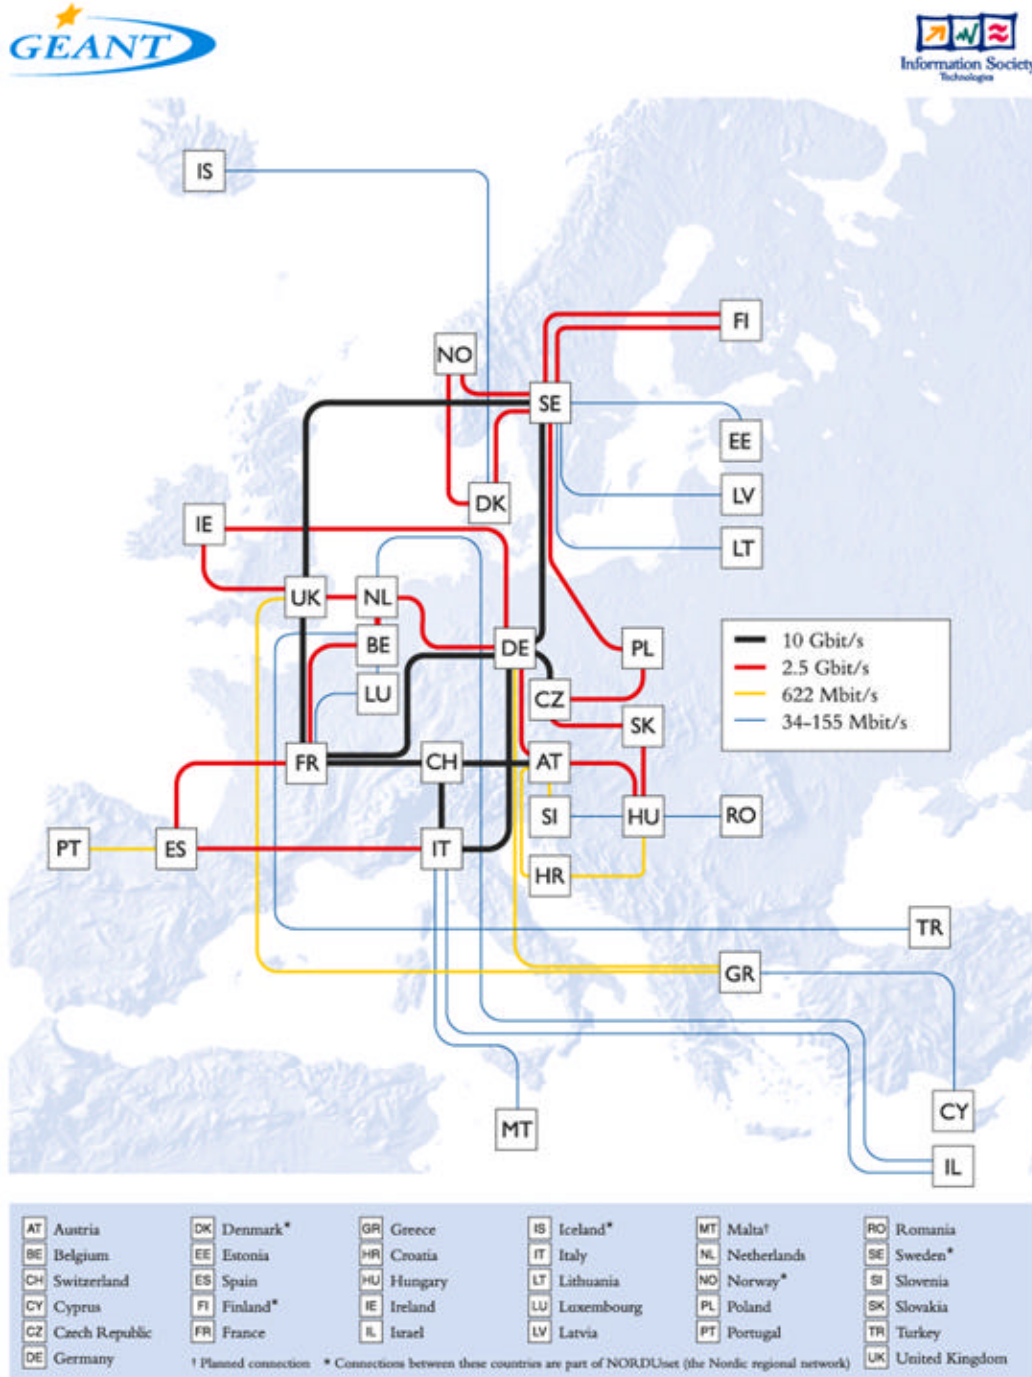
In the case of GÉANT, 90% of the routers were deployed from the last tender undertaken in 2001. The resulting Juniper platforms were recent enough to support the forwarding of IPv4 and IPv6 packets and the dynamic protocols needed for the forwarding decision. Also, GÉANT was designed in a way that there is no notion of core routers and edge routers. Each router in GÉANT was a core and edge device, as illustrated in Figure 4-1. Therefore, one could not deploy an encapsulated model where the edge routers would have been dual stack while the core routers could have stayed IPv4-only, crossed by a mesh of tunnels (MPLS or GRE) landing on the edge routers.

Based on these parameters it was decided to implement a dual stack model. However the IGP deployed in GÉANT was OSPFv2 which doesn't support IPv6. Therefore another IGP had to be implemented to support the routing of IPv6 packets inside the core. IS-IS was chosen, and GÉANT migrated successfully in the beginning of December 2002 from OSPFv2 to IS-IS.

4.2. Routing protocols

To discover the topology and route traffic across the network, routing protocols have to be implemented for both stacks of protocols, IPv4 and IPv6. Generally, a core backbone has two types of routing protocols, the Interior Gateway Protocols (IGPs) and the Exterior Gateway Protocols (EGPs).

The IGP is used to learn the topology of the network and forward traffic across the backbone through the best path. Currently IS-IS, OSPF(v2) and RIP are the classical routing protocols used by IPv4. The EGP (e.g. BGP) is used for connectivity to external networks. BGPv4 is the external protocol to route IP traffic between domains. It has been enhanced with IPv6 NLRI and is available in BGP4+ after upgrade of the image of the router. Nothing special has to be done, IPv4 and IPv6 BGP peerings can coexist. However vendor specific requirements and interoperability have to be taken into account.



Multi-Gigabit pan-European Research Network
Backbone Topology December 2002



Figure 4-1: GEANT network topology, December 2002

IS-IS and RIPng are available for routing IPv6 while OSPF version 3, which is a complete rewrite of OSPFv2 for IPv4, is now becoming available for the key platforms.

Depending on which IGP is implemented in the core backbone, the efforts required to migrate the network from an IPv4 network to a dual stack network will be different.

A core backbone installed with IS-IS, after the dual stack upgrade of the routers, will have one IS-IS process building one database in which the links have different attributes according to which IP protocols (IPv4 or IPv6) are used for the forwarding. That means that a link which has IPv4 and IPv6 addresses will have different TLVs in IS-IS describing it. Moreover, if a service provider runs IS-IS for IPv4 and wants to use it for IPv6 but doesn't intend to align both topologies, then a multi-topology IS-IS is needed.

A core backbone with OSPFv2 implemented can not build a database for IPv6 and needs, in parallel, a second IGP which can be IS-IS or OSPFv3.

In some cases, if it is not possible to run an "IPv6-only" IGP database - to avoid having IPv4 routes in the second IGP either - two strategies have to be studied:

- The core backbone can run in parallel two IGPs.
- The core backbone has to change its IGP before running in dual stack mode.

The choice that was made for GÉANT is discussed in the next section.

4.3. IS-IS and OSPFv3 considerations

There were potentially two choices for the new IGP for GÉANT, either the OSI protocol IS-IS or the new version of OSPF, OSPFv3.

IS-IS was chosen because this protocol has been used in many ISP networks for a number of years, and because IS-IS can handle within one single process/database IPv4 and IPv6 routes.

OSPFv3 is a new protocol and is dedicated for IPv6 only, thus OSPFv2 would still have been needed in the core to route IPv4 packets. The two protocols OSPFv2 and OSPFv3 could cohabit on the routers and behave as a "ship in the night" model. The routers might be able to run two IGPs in parallel with good performance. However, having two IGPs in parallel will make it difficult to maintain consistency for both databases of the IPv4 traffic. While IPv6 will be handled only by the second IGP, IPv4 will fill the databases of both IGPs and the administrator will have to configure preferences among the two link state protocols. This might be difficult to manage and may create inconsistencies.

Therefore, when OSPFv3 becomes hardened and proven for production networks, people will have the choice between one IGP that can handle several IP protocols, or one IGP per IP protocol. However, OSPFv2 could still be run in parallel with IS-IS for IPv6.

4.4. Trials and considerations prior to transition

These considerations were first reported in [6NET-D211]. These notes describe the plans for transition, and trials undertaken for IPv6 validation.

4.4.1. Router performance

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

According to 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.

- Interoperability

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

Once the product evaluation is achieved a laboratory can be set up to evaluate the best design for the production network.

Based on the current architecture, several designs can be evaluated for provisioning the IPv6 service and transitioning the network.

The best design will be the one that achieves the smoothest transition and provides the best stability.

4.5. Migrating the IGP

The migration of GÉANT was based on a sound preparation with the followings steps:

- 1) Training on IS-IS
- 2) Designing the IS-IS configuration of the network.
- 3) Validation of the design and procedures for the migration
 - a. With IS-IS experts and experienced people
 - b. Tests in laboratories
- 4) Then migration

4.5.1. Monitoring the transition

The introduction of IPv6 on the backbone implies having monitoring and troubleshooting tools in place. The basic tool set would be a DNS server for the resolution of IPv6 addresses and router names. Telnet (ssh), Ping and Traceroute have to be available for IPv6 on the routers.

In the perspective of having a pure IPv6 backbone, monitoring tool platforms like HP-Openview, Infovista etc. have to be made available and evaluated in dual stack mode and pure IPv6. Other useful applications like Netflow monitoring and SNMP polling (Cricket, MRTG) have to be evaluated. In addition, TFTP/FTP is desirable on the routers.

These issues are covered in 6NET WP6, and reported in [6NET-D632].

4.6. Towards an IPv6 Service

Other items that were considered for the network to deliver IPv4 and IPv6 services included:

- 1) The configuration of the core network with a flexible IPv6 addressing plan
- 2) The routing policy to implement for E-BGP and I-BGP
- 3) DNS duties if the LIR is providing IPv6 delegation.
- 4) Establishment of operational procedures for the new service
- 5) The Monitoring Infrastructure

These items will be reported in more detail in future versions of this cookbook, in other 6NET related Work Packages (e.g. in [6NET-D312], [6NET-D632] and [6NET-D622]) as well as in reports of the GÉANT project itself.

5. Dual-stack NREN deployment case studies

The most common method to introduce IPv6 services into IPv4 networks will be through dual-stack networking on the NREN. This complements the backbone transition and pushes the issue of deployment to the edge, i.e. to the universities.

In the scope of 6NET, many NRENs have already migrated to dual stack; the specific experiences of SURFnet, Funet and Renater are reported here. A summary of a survey of the status of other NREN deployments is given in Section 7.

The SURFnet case study (updated from [6NET-D222]) is particularly interesting because there SURFnet changed to use 6PE because it wished to avoid expensive line card upgrade costs just to have line rate IPv6 (in the previous dual-stack mode up until March 2004, IPv4 was handled in software).

5.1. SURFnet Case Study (Netherlands)

5.1.1. Introduction

SURFnet is the national computer network for higher education and research in the Netherlands. SURFnet connects the networks of universities, colleges, research centers, academic hospitals and scientific libraries to one another and to other networks in Europe and the rest of the world. The SURFnet network enables its users to communicate with other network users and to consult the Internet from their office or their home PC.

In the summer of 2001 the national research infrastructure of SURFnet, called SURFnet5, was constructed as part of the GigaPort project and in partnership with BT Nederland² and Cisco Systems.³ After a public procurement process these partners were awarded a six year contract at the very end of 1999, with BT Nederland supplying the infrastructure and Cisco Systems providing the router equipment.

5.1.2. The SURFnet5 dual stack network

The SURFnet5 network consists of a core that is situated at two locations in Amsterdam, at the Points of Presence (PoPs) called “Amsterdam1”⁴ and “Amsterdam2”⁵. Two Cisco 12416 routers are installed at each location implementing the core functionality. Fifteen PoPs are each connected to both core locations over two separate unprotected SONET/SDH framed lambdas running at 10 Gbit/s. The engineering of the lambdas is such that it ensures that one connection is always maintained in case of a single transmission or core router failure. At each of the fifteen PoPs a

² The original contract was signed with Telfort, who later became BT Ignite The Netherlands; the current name is BT Nederland.

³ See also the press release: “GigaPort, BT Ignite, Cisco Systems and SURFnet launch the world's most advanced research network” at http://www.gigaport.nl/publicaties/pers/en_pers270601.html.

⁴ This PoP is located at SARA in the eastern part of Amsterdam.

⁵ This PoP is located at Hemptpoint, a BT Nederland co-lo facility, in the western part of Amsterdam.

Cisco 12416 is installed, as the concentrator router, together with a Cisco 7507 router for the low speed, legacy connections at speeds up to 155 Mbit/s. The backbone of SURFnet5 makes use of IP-over-DWDM, using POS framing. Figure 5-1 shows the logical topology of SURFnet5.




Figure 5-1: Logical topology for SURFnet5

From the start of SURFnet5 all routers running Cisco's IOS have IPv6 implemented. During the early days of the network, IOS versions in the 12.0ST train were used in the GSR routers. During 2002 the transition to 12.0S was made. The current version, in March 2004, in the GSR routers is "IOS Version 12.0(26)S1" and in the 7500 routers is "IOS Version 12.2(14)S2".

SURFnet5 started as a dual-stack network and has been running as such until March 2004. During March 2004 IPv6 routing within the core of the network was migrated from dual-stack to 6PE. The most important reason for this migration was to achieve line rate IPv6 forwarding in the SURFnet5 network. SURFnet5 is a 10 Gbit/s network largely built on Engine4 line cards and these cards handle IPv4 on the fast path while IPv6 is handled by the processor of the line card. After a large replacement in 2003 of Engine2 line cards with Engine3 line cards, the edges of SURFnet5 became capable of handling IPv6 traffic at line rate. By implementing 6PE in the core of the network it was ensured that IPv6 traffic was also handled at line rate in the core and potential bottlenecks were removed.

The routing set-up as implemented and running today runs without problems today. A number of features are on our requirements list, of which IPv6 multicast is high on the list. We already have the first hands-on experience with IPv6 multicast in the 6net context using external MBGP for IPv6.

5.1.3. Customer connections

SURFnet5 supports customer connections on IPv4 as well as IPv6. For IPv4 unicast and multicast are supported, while for IPv6 this is currently unicast only. IPv6 can be delivered towards SURFnet's (approximately) 200 customers in the field of research and higher education using either a native or tunnelled approach

Native IPv6 connections are implemented either “dual stack” or “on a dedicated port”. Dual stack implies that the customer connection runs IPv4 as well as IPv6 on the same local loop. A few customers, however, prefer IPv6 on a dedicated port, next to their IPv4 port for the simple reason that they have a dedicated router for IPv6 on their LAN.


Tunnelled connections to customers are only allowed in case the customer is not able to implement IPv6 in dual stack mode or on a dedicated port. In case we implement a tunnel connection with a customer, we always ask the customer to make plans for going native.

5.1.4. Addressing plan

SURFnet received an official IPv6 prefix from the RIPE NCC in August 1999 of length /35, which was enlarged to a /32 during the summer of 2002:

2001:0610::/32⁶

This prefix is split up as follows:



Only the first /35 is currently used. The other 7 /35s are for future use right now. For each PoP a /40 prefix (NLA 1) is reserved and each /40 is again split into 256 /48 prefixes (NLA 2) to number the backbone links and customer networks. The break-down of SURFnet’s prefix on a per-PoP basis is shown in detail below.

The prefix 2001:0610::/40 is in use as the carrier network for SURFnet5, in which all links reside. For the links we use /64 prefixes, in the format:

2001:0610:00xx:xyyy::zzz

In this format, xxx denotes the third octet of the link’s or LAN’s IPv4 prefix and yyy denotes the fourth octet of the link’s or LAN’s IPv4 prefix. The zzz denotes the actual address. An example of this is shown below in Figure 5-2:

```
Cisco#sh int po0/0 | inc Internet
Internet address is 145.145.160.5/30
Cisco#sh ipv6 int po0/0 | inc subnet
2001:610:16:4::5, subnet is 2001:610:16:4::/64
```

Prefix	NLA 1	PoP
2001:0610:0000::/40	0	(not used yet)

⁶ For registration details on this prefix, see: <http://www.ripe.net/cgi-bin/whois?2001:0610::/32>.

2001:0610:0100::/40	1	Amsterdam1
2001:0610:0200::/40	2	Amsterdam2
2001:0610:0300::/40	3	Hilversum
2001:0610:0400::/40	4	Leiden1
2001:0610:0500::/40	5	Utrecht
2001:0610:0600::/40	6	Chicago, IL, USA
2001:0610:0700::/40	7	(not used yet)
2001:0610:0800::/40	8	(not used yet)
2001:0610:0900::/40	9	Delft
2001:0610:0A00::/40	10	(not used yet)
2001:0610:0B00::/40	11	DenHaag
2001:0610:0C00::/40	12	Rotterdam
2001:0610:0D00::/40	13	(not used yet)
2001:0610:0E00::/40	14	(not used yet)
2001:0610:0F00::/40	15	(not used yet)
2001:0610:1000::/40	16	(not used yet)
2001:0610:1100::/40	17	Eindhoven
2001:0610:1200::/40	18	Maastricht
2001:0610:1300::/40	19	Nijmegen
2001:0610:1400::/40	20	Tilburg
2001:0610:1500::/40	21	(not used yet)
2001:0610:1600::/40	22	(not used yet)
2001:0610:1700::/40	23	(not used yet)
2001:0610:1800::/40	24	(not used yet)
2001:0610:1900::/40	25	Enschede
2001:0610:1A00::/40	26	Groningen
2001:0610:1B00::/40	27	(not used yet)
2001:0610:1C00::/40	28	(not used yet)
2001:0610:1D00::/40	29	Wageningen

2001:0610:1E00::/40	30	Zolle
2001:0610:1F00::/40	31	(not used yet)

Figure 5-2: SURFnet prefixes per POP

5.1.5. Routing

During the build-up of SURFnet5, in the summer of 2001, IS-IS was used for IPv4 and IPv6. Since the transmission of SURFnet5, the 10G lambdas, is unprotected we used MPLS Traffic Engineering's extension Fast ReRoute as the protocol to protect the network against failures by making available an alternative path well within the 50 msec time frame. However, MPLS TE was only supporting IPv4, yielding a situation in which the IPv4 topology was different from the IPv6 topology. At that time, IS-IS was not able to handle different topologies for the two protocols, and we had to move away from IS-IS for IPv6. We temporarily transitioned to RIPng for IPv6, while we stayed at IS-IS in combination with MPLS TE FRR for IPv4. In the mean time Cisco developed multi-topology IS-IS, which enabled SURFnet to move away from RIPng back to the original plan. Before we went back to IS-IS for IPv6, we decided to abandon the MPLS TE FRR ship for network management reasons, as we found the operations and maintaining of FRR too complex. At the same time, enhancements to IS-IS made it possible to fully rely on the routing protocols at the IP layer for the resilience in SURFnet5. During March 2004 the dual-stack approach was migrated to a 6PE implementation. The four core routers act as BGP route-reflectors in both the IPv4 as IPv6/6PE routing. All fifteen concentrator routers and border routers are BGP route-reflector clients hanging of these.

5.1.6. Network management and Monitoring

As a ground rule, SURFnet treats IPv6 on the network level equal to IPv4, as much as possible. This implies that procedures between SURFnet's NOC and SURFnet Network Services are streamlined to support this. Also, the external peering policy towards other non-European NRENs such as Abilene and CA*net 4 and towards the commodity Internet through SURFnet's upstream providers and through peerings at the Amsterdam Internet Exchange (AMS-IX) is the same for IPv4 and IPv6. At the AMS-IX it means that IPv6 peering requests are handled exactly like IPv4 peering requests, as SURFnet's peering policy on the Exchange is identical for both protocols. In March 2004, SURFnet has 54 IPv6 peering sessions enabled over the AMS-IX.

While we strive to be able to perform network management over IPv6 as well, today this is done only partly at this moment. We monitor the availability of the IPv6 customer connections as well as external connections. Also where software allows us to (like SSH, Nagios and Rancid) we use IPv6 to manage the router equipment.

5.1.7. Other services

The following actions on applications and the like in the area of IPv6 are being or have been undertaken:

- The anonymous-FTP archive of SURFnet and NLUUG was enabled for IPv6 in the second quarter of 2002. See Figure 5-3 for the IPv6 volume transferred by this server.
- SURFnet's Stratum 1 Network Time Protocol (NTP) servers `chime3.surfnet.nl` and `chime4.surfnet.nl` are up and running and serving time over IPv6. This server is available for the 6net community to use and synchronize on.

- SURFnet's DNS server `ns3.surfnet.nl` is IPv6 aware as well as reachable over IPv6.
- For NNTP 1 production news feeder runs IPv6 and has an external peering with Switch and other smaller peerings. In the testbed two news reader machines are used with IPv6 connectivity.
- The main web site of SURFnet⁷ is enabled for IPv6.
- Several experimental services like videostreaming (unicast, multicast) and internet telephony (SIP) run over IPv6 today.
- All SURFnet services, currently available over IPv4, will be made available over IPv6 as well.

For new application services that are being developed and for which equipment or software is procured, we asked the potential suppliers on their readiness for IPv6.

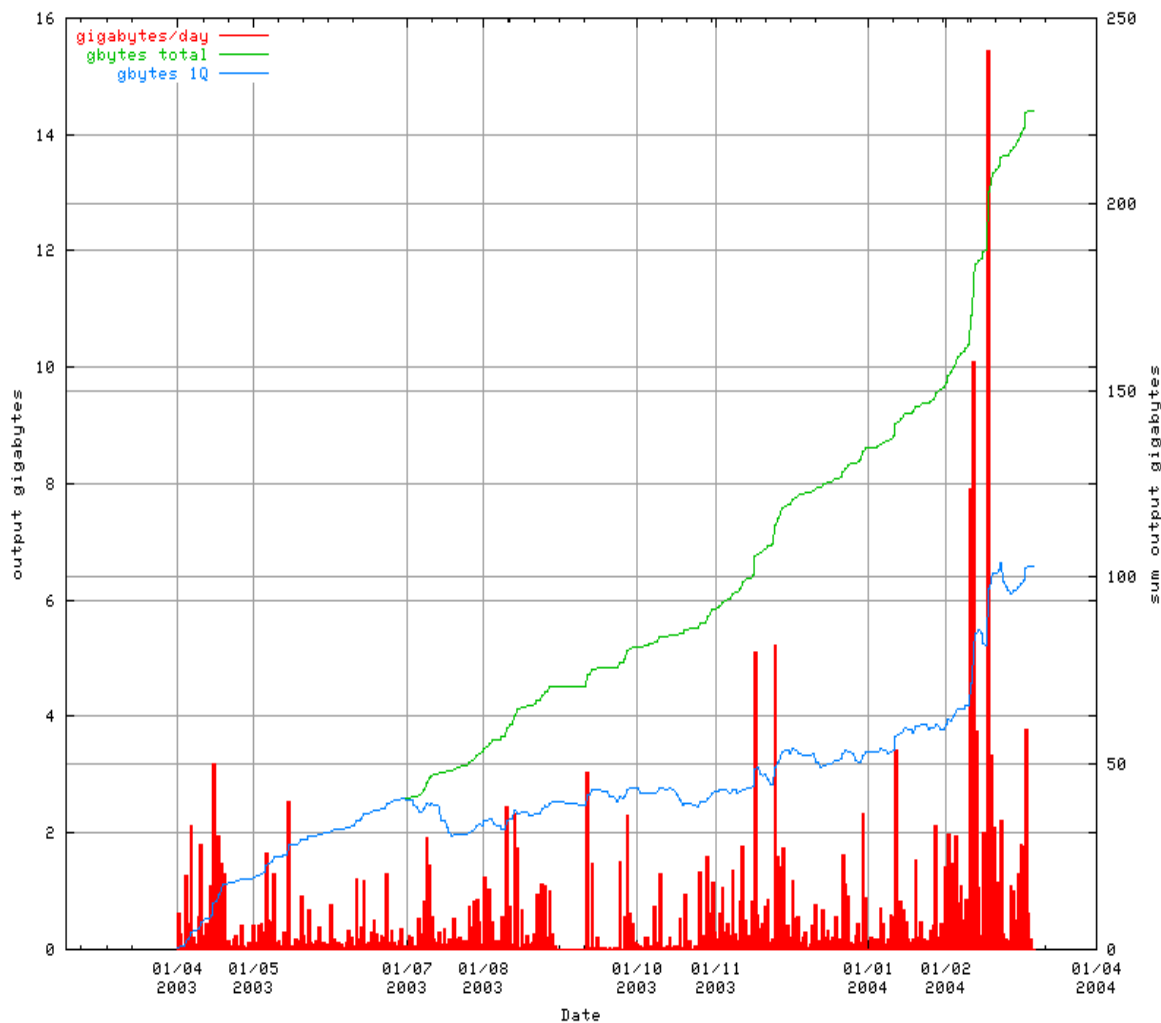


Figure 5-3: Anonymous-FTP over IPv6 volume.

⁷ The English version is available at: <http://www.surfnet.nl/en/>.

5.1.8. Contact information

For more information on this case study, the reader can contact SURFnet Network Services.⁸

5.2. Funet Case Study (Finland)

5.2.1. Overview

In 2001, Funet's core network was upgraded to 2.5 Gbit/s PoS links. Six Juniper routers (M20) were added to the previously all-Cisco network. There was not so much user demand for native IPv6 at this stage, but by the end of 2001, some plans for enabling dual-stack on these Juniper routers had been made.

During Q1/2002 various tests and trials were run.

After fixing all the issues, the Funet core had transitioned to complete dual-stack deployment in the core networks by Q2/2002

5.2.2. History

The Funet experience involved some bleeding edge IPv6 deployment, leading to a subsequently stable service.

In Q2/2002, the minimum IPv6 working version 5.2R2 was installed on the Juniper routers. Later, versions 5.3R2, 5.3R3, and 5.4R3 were also used. Since 5.4R3 was deployed, no further issues have been noticed.

The IPv4 network used OSPFv2 and BGP for routing, but OSPFv3 for IPv6 was not ready. Funet didn't want to change the routing protocol, and for clarity, they wanted clearly separate IGP's for IPv4 and IPv6: OSPF and IS-IS. In addition, if IS-IS had been deployed for IPv4 and IPv6, multiple topologies would have to have been supported as IPv6 routing would have been different (in parts) from IPv4. So, IPv6-only IS-IS was clearly the best (and only, discounting RIPng) choice.

Unfortunately, there were problems in the Juniper and Cisco devices with IPv6-only IS-IS. For example, the Juniper platform would not support IS-IS only for IPv6. The first, and worst, problem was that the Juniper routers would always also advertise IPv4 addresses used in loopbacks and point-to-point links. Cisco's IOS, unless you enabled IS-IS for IPv4 too (which was a non-starter for Funet), would discard all such attempts to form adjacencies: a total inoperability problem. Cisco's IS-IS implementation has an option 'no adjacency-check' to override this; however, an undocumented fact was that it would only work (at least in this case), using level-2-only IS-IS circuit-type (which was not the default). A first step in interoperability was gained when these were enabled in IOS.

Some problems continued. IS-IS route advertisements from Ciscos to Junipers were accepted in the route database, but not put to the Junipers' routing table: this was caused by the above mentioned problem with adjacencies; this was reported, and fixed, in 5.2R2; a minimum workable version for Funet to use. The issue with Juniper always advertising IPv4 addresses in IS-IS was fixed, as an undocumented feature 'no-ipv4-routing' in 5.4R1. Also, one could not redistribute static discard routes to IS-IS (to generate a default route) until this was fixed in 5.3R3. You also could not set a

⁸ SURFnet, Department of Network Services, P.O. Box 19035, 3501DA Utrecht, The Netherlands, Email: netmaster@surfnet.nl, Tel: +31 30 2305305

metric when advertising a default route other than by redistributing a static discard route and applying a route-map in Cisco.

Fortunately, the rigorous tests in the lab network were enough to expose all of the above problems, which were fixed.

Also in 2002, most of Funet's Cisco routers were replaced by Junipers in an upgrade of the network.

In early 2003, Funet noticed significant bandwidth bottlenecks (in the order of 10's of megabits/second) with their remaining Cisco equipment -- 7200's, 7500's, GSR's -- regarding IPv6 forwarding capabilities, but the situation improved tremendously when software supporting CEFv6 came out later in 2003. This is very important for Funet as the IPv6 traffic level is 30+ Mbit/s.

Currently, Funet has almost all IPv6 features they need. Only IPv6 multicast is really missing, and the support for that is in progress (see [6NET-D312]).

The Funet IPv6 network is shown in Figure 5-4.

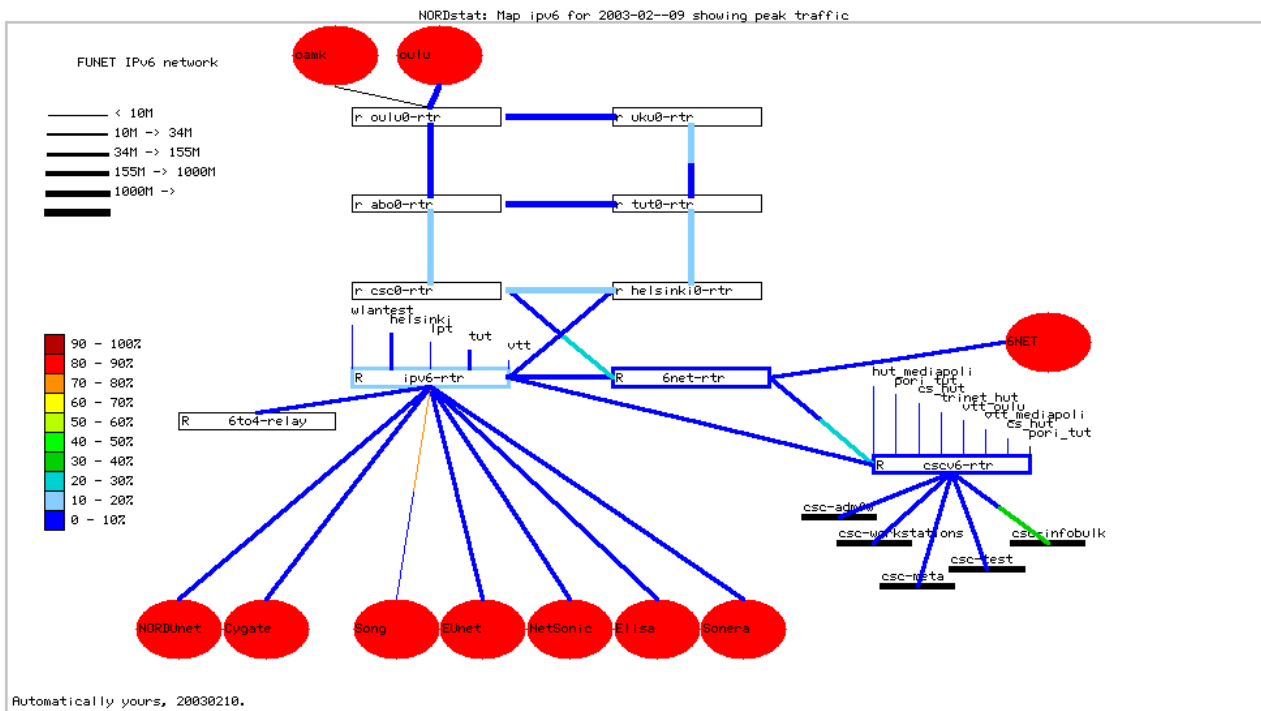


Figure 5-4: The Funet IPv6 network

5.2.3. Addressing plan

Funet has an IPv6 SubTLA address block of 2001:708::/32.

There are two different aspects to addressing: the customers and the network infrastructure.

5.2.3.1. Customers

There are six "SuperPoP" core routers; these are used when giving out /48 prefixes to customers, e.g.:

2001:708:0KLM::/48

"K" here is the number of the regional SuperPoP. This is not meant to be used for aggregation purposes, rather than to just create some mild hierarchy for the addresses, rather than allocate them in a sequential fashion.

"L" is the sequential number of customers within the SuperPoP. "0" is reserved.

"M" is reserved for expanding the number of customers in the SuperPoP and/or increasing the size of the assignment to the customers. To maintain future flexibility, this is still zero.

When assigning the prefix to the customers, Funet recommends they keep the first four bits (the first nibble) zero for now.

An example of a customer assignment is 2001:708:510::/48.

5.2.3.2. Network Infrastructure

Here, "KLM" has a slightly different meaning. "K" still means the SuperPoP, "L" means the PoP acting under the SuperPoP in an access network, and "M" is an identifier for the router in the PoP.

Loopback addresses are taken from the prefix:

2001:708:0:10:KLM::/112

So, a few loopback addresses will be:

tut0-rtr.funet.fi: 2001:708:0:10:300::1

tut1-rtr.funet.fi: 2001:708:0:10:301::1

uta3-rtr.funet.fi: 2001:708:0:10:313::1

These are configured using /128 prefix length. Note that the identifier of the router is also reflected in the naming ("uta_3_-rtr").

Point-to-point addresses in the core and the access networks are taken from one block - all addresses come from under a single /64:

2001:708:0:F000::/64

In particular:

2001:0708:0:F000::klmn:KLMm:z/112

klm and KLM identify the routers at the end of the point-to-point links, taken from the loopback addresses; in above, these would have been "300", "301", and "313". SuperPoP's or the smallest number goes first as klm. "n" and "m" are sequential numbers, used when necessary - for example if there are multiple links between routers which need to be numbered - defaulting to zero. "z" is the end-point of the point-to-point link: always "1" or "2". The same SuperPoP or smallest first rule applies here too. So, in consequence, the addressing becomes like:

uku0-jyu3: 2001:708:0:F000::4000:3230:[12]/112

uku0-oulu0: 2001:708:0:F000::4000:5000:[12]/112

uta3-jyu3: 2001:708:0:F000::3130:3230:[12]/112

The point-to-point links toward customers are always numbered from the customer's addresses, due to simplicity and policy reasons.

For peerings and miscellaneous use, a block of:

2001:708:0:F001::/64

is reserved.

In addition, some special use addresses are used inside 2001:708::/48, for example 2001:708::{1,2} (for a few routers), 2001:708::123 (NTP), 2001:708::53 (DNS) etc.

5.2.4. Configuration details

In this section, core and customer (edge) configuration examples are listed.

5.2.4.1. Configuring the Core

The configuration of the Juniper routers is given in Appendix B.

As can be seen, the model is such that all the routes of the router are redistributed in the IS-IS. An alternative approach would be to include all the interfaces in the IS-IS as passive interfaces.

This is not considered to have serious drawbacks, as none of these redistributed routes are advertised outside the autonomous system: the advertisement includes the aggregates only.

5.2.4.2. Connecting the Customers (edge)

Customers are connected using static routes. The configuration is very simple, like the below on Juniper:

```
interfaces {
  fe-X/X/X {
    unit Y {
      [...]
      family inet6 {
        rpf-check fail-filter RPF_FAIL_IPV6;
        address 2001:708:KLM:xxxx::2/64; # from the customer
      }
    }
  }
}

routing-options {
  rib inet6.0 {
    static {
      route 2001:708:KLM::/48 next-hop 2001:708:KLM:xxxx::1;
    }
  }
}
```

```
}

firewall {
  family inet6 {
    filter RPF_FAIL_IPV6 {
      term DEFAULT {
        then {
          count count-rpf-fail-ipv6;
          log;
          discard;
        }
      }
    }
  }
}
```

And on Cisco this could look like:

```
interface Fa0/0
[...]
ipv6 address 2001:708:KLM:xxxx::2/64
!
ipv6 route 2001:708:KLM::/48 2001:708:KLM:xxxx::1
```

As one can see, using static routes is a very simple operation. Customers connected to Junipers have IPv6 unicast-RPF enabled; this feature is not available on IOS at the time of writing (in 2003).

5.2.5. Monitoring

The addresses of all loopback and point-to-point addresses are entered into DNS in a special format. A script configured to allow zone-transfer of "ipv6.funet.fi" zone fetches this information and digs out the IP addresses which should be in use. The pinger periodically (once in five minutes) checks that the links (including the links to customers and the peers) are up and responding; if not, it sends an alert.

BGP and IS-IS adjacencies are also monitored using a tool which collects syslog warnings sent from routers to a central syslog server. If adjacencies or sessions flap, this can be noted in the monitoring page.

All routers and links are collected to a custom network map/monitoring tool, where the traffic levels and similar can be monitored easily.

5.2.6. Other services

Funet has been using and advertising a 6to4 relay to everywhere (openly) for over a year now. This includes the advertisement of 2002::/16 and 192.88.99.0/24. The router in question is a FreeBSD system running zebra OSPF and BGP routing protocols. A more detailed description is available in the 6NET Site Transition Cookbook (Deliverable D2.3.2).

IPv6 multicast is also being used. For this purpose, a different FreeBSD router is being used.

A TCP/UDP relay (faith in FreeBSD) is used on server-side to experimentally enable IPv6 access to a few IPv4-only services.

An IPv6 newsfeed service is IPv6-enabled, and is generating 30+ Mbit/s of steady IPv6 traffic. Also, ftp.ipv6.funet.fi is also IPv6-operational.

5.3. Renater Cast Study (France)

5.3.1. Overview

Within the framework of a pilot project by GIP Renater, carried out by G6, the infrastructure of Renater2bis has been used to set up an IPv6 pilot network. The aim for the GIP Renater was to begin to establish the means and mechanisms required to allocate the necessary resources to connect the test sites to the IPv6 pilot (NLA-ID allocation, reverse DNS delegation, IPv4 - IPv6 transition mechanisms, etc.)

The 2001:660::/35 prefix was used for addressing the pilot and connected academic sites. All industrial partners were addressed in the 3ffe:300::/24 address space. Dedicated ATM PVCs were used to transport IPv6 between the different IPv6 POPs. The original /35 prefix has been expanded to a /32 by the RIPE NCC (as part of common RIR policy) since the prefix was originally allocated.

5.3.2. Native support

Thanks to this experience, IPv6 is now offered as a native service on Renater3's backbone.

All Renater3's points of presence offer global IP connectivity to the regional networks and to the sites which contain both IPv4 unicast, IPv4 multicast and IPv6 unicast.

Both traffic types (IPv4 and IPv6) are carried in the backbone without any distinction, offering equal performance, availability, supervision and support levels. The Renater NOC was IPv6 trained to be able to achieve the same service for IPv6 and IPv4.

5.3.3. Addressing and naming

The addressing of the whole network is designed in a hierarchic way: this enables the aggregation of all routes, so reducing the routing table's size. Each point of presence of Renater3's backbone is allocated a /40 IPv6 prefix. Sites connected to a POP receive a /48 prefix derived from the /40 of the POP.

The GIP Renater manages the delegation of reverse zones of Renater3's SubTLA prefix (2001:0660::/32). It delegates to each site the reverse zone of the NLA-ID (/48) allocated to the site.

AFNIC, the French Network Information Center, is managing the .fr top-level domain name. They are connected to Renater3's Internet exchange point (SFINX) that supports IPv6.

5.3.4. Connecting to Renater 3

Using the experience gained with the IPv6 pilot of Renater2, procedures for connecting to Renater3 were designed and the teams were trained to be aware of the new processes. All the sites connected to the pilot have to be moved to Renater3, and be allocated a new prefix in the new address space. There was no D-day between the pilot and Renater3 as connectivity was not shut down for people connected through the pilot, to let them have time to do the procedures to connect to Renater3.

The procedures defined for IPv6 are very close to the ones defined for IPv4. This implies that a site can connect only if the network administrator fills some forms. An issue is that all these administrators are not IPv6 aware, and that some people in the site they manage need IPv6 connectivity. We have examples of labs in universities that are connected to the IPv6 pilot to have IPv6 connectivity. If those labs want to move to Renater3, they need the agreement of the network administrator. The prefix given to the site is aimed at addressing the whole network, and the administrator has to delegate a part of this prefix to the lab, which implies some IPv6 deployment forecast. This is not easy if the administrator is not IPv6 aware. This can lead to long delays to connect some sites to Renater3.

5.3.5. The regional networks

Renater3 is a national backbone with at least one POP in every French region. To connect to this POP, the sites use some access network (regional or metropolitan network). At the beginning of Renater3, none of these access networks were IPv6 enabled, meaning that the connection between the Renater3 POPs and the sites was done using tunnels or dedicated links (ATM PVCs, serial links, etc). As the core backbone is a Cisco GSR infrastructure, the choice was made not to set up tunnels on the core routers. Some dedicated equipment was deployed to concentrate the tunnels in the regions. Some of these routers are the ones used for the IPv6 pilot.

Six months after the deployment of Renater3, already some regional networks have started their IPv6 deployment. At this stage, they are still in the conception phase. There are many different approaches for this IPv6 deployment: 6PE, VLANs, fully dual-stack, PVC ATM, tunnels, etc, and all these deployments are done in straight collaboration with Renater.

5.3.6. International connections

Renater3 offers IPv6 connectivity to national (RNRT) and Europeans (IST) projects.

Renater is connected to the IPv6 world with some international connections, to North America (6TAP), to Europe (6Net core), and to Asia (TEIN).

The IPv6 service is extended to the SFINX (Service for French Internet eXchange), which offers IP actors an interconnection point that carries both IPv4 and IPv6. IPv6 is exchanged in dedicated VLANs. This makes it possible to manage IPv6 traffic more easily.

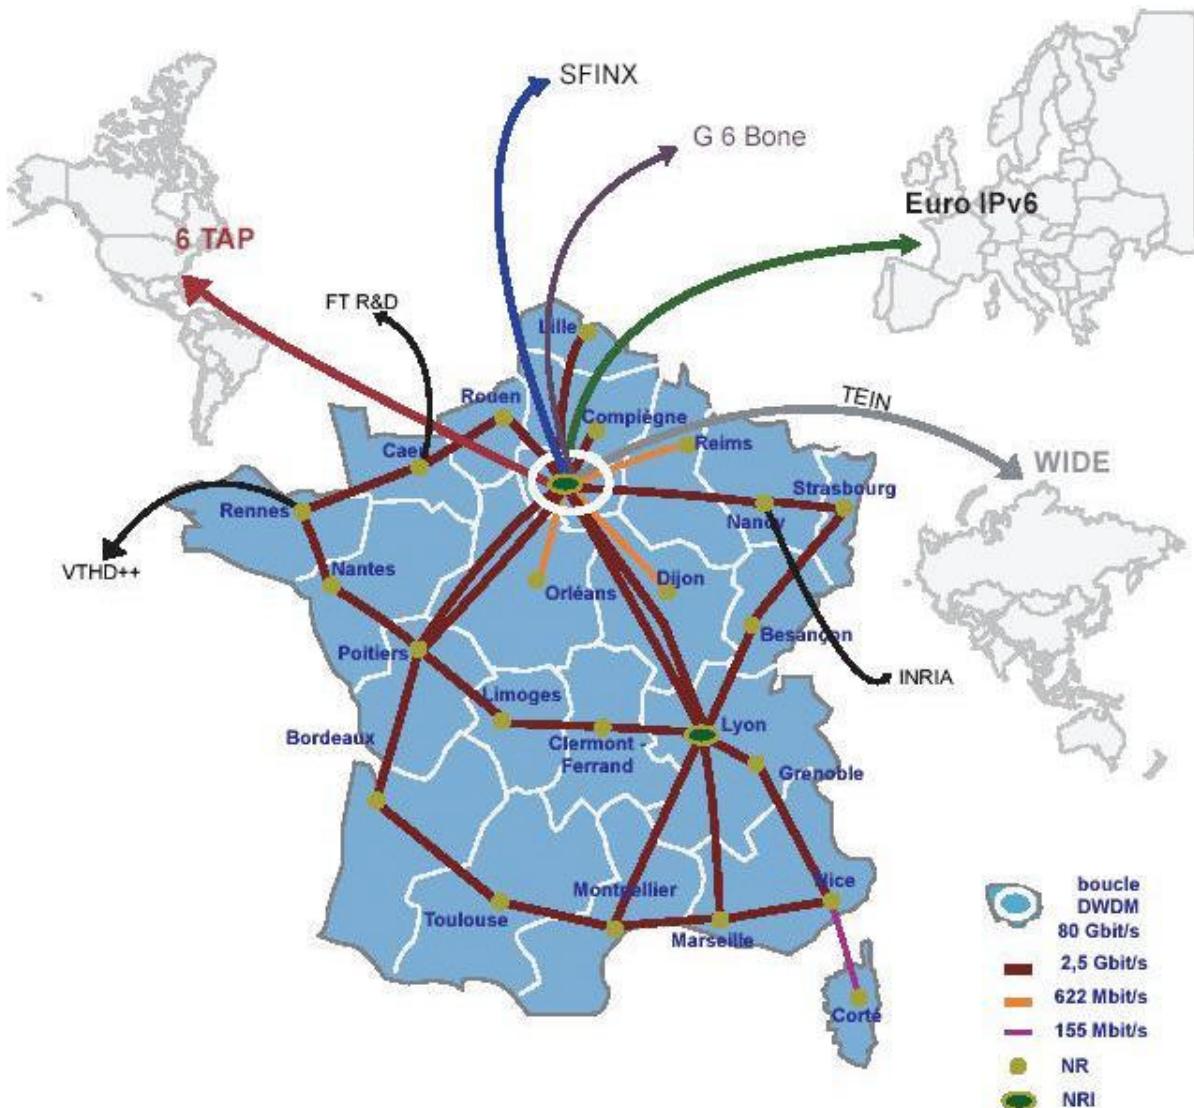


Figure 5-5: The Renater-3 network

5.3.7. IPv6 Multicast

An experimental IPv6 multicast network (M6Bone) is running on the Renater3 infrastructure. It allows the connection of lots of sites, all over the world. This network allows all the sites connected to test and develop IPv6 multicasting. This network is connecting today over 50 sites in Europe, Asia and Africa, making the network one of the most advanced multicast IPv6 network in the world. The work of m6bone is feeding into the leading IPv6 Multicast work on the 6NET backbone (see [6NET-D312]).

6. Other case studies for NREN Transition

In this section we report other transition plans of interest for NRENs in the 6NET project. In particular, the German 6WiN parallel IPv6 deployment, which was first reported in [6NET-D221], and is an interesting case study of an alternative (early step) to full dual-stack deployment. The 6WiN network is expected to run for at least one more year, and then be reviewed.

We also include a brief overview of the UKERNA status and plans for IPv6 deployment on the JANET network, and GRNet's experience of migrating from 6bone to production IPv6 networking.

6.1. 6WiN: Introduction of a parallel IPv6 network (DFN: Germany)

In Germany most of the research and educational facilities are connected over the G-WiN (Gigabit Wissenschafts-Netz) of DFN (Deutsches ForschungsNetz). G-WiN is a pure IPv4 network. In the past IPv6 connectivity was established solely to the JOIN project at the University of Münster over IPv6-in-IPv4 tunnels.

To integrate IPv6 into the services of DFN a dedicated network called 6WiN was developed. The goal was to get native IPv6 connectivity as close to the R&E facilities as possible without compromising the stability and reliability of the IPv4 network. For this reason and because (at the time) there was no production IPv6 software with the full set of needed features for the G-WiN core routers, a separate network with dedicated routers and dedicated connections was set up. As the G-WiN is a very large network with more than 30 routers, it was not possible to establish a full duplicate of the network for IPv6, and only a few PoPs are used for the creation of the 6WiN.

6.1.1. Internal connectivity and setup

The core of 6WiN is a set of five Cisco 7206 routers spread throughout Germany that are connected with dedicated serial lines (E3/34MBit). The 6WiN routers are located in the same locations as the G-WiN routers. Every 6WiN router is connected to the G-WiN via a Fast Ethernet interface to the IPv4 world.




Figure 6-1: The 6WiN network, including Münster

There is a sixth Cisco 7206 located in Münster. This is a special case since this router is not conceptually a part of the 6WiN core. But as it has external peerings to other networks and therefore resides in the same IS-IS- and BGP-cloud like the other routers, it belongs technically to the core.

The connection between the router in Münster and the 6WiN is a 34MBit E3 line, too. Like all internal lines of the core it is native and IPv6-only. It is also intended to attach DFN's customers with native connections. This is often not possible, because a native line often means additional expenses for the R&E facility. So most of them are connected over IPv6-in-IPv4 tunnels. To avoid high delay of packets traversing the IPv4 network, every facility is connected to the 6WiN PoP, which is topologically closest in IPv4. Currently 28 facilities are connected to 6WiN. In four facilities that reside at the same location as a 6WiN PoP, or where an additional line was available for little or no additional cost, it was possible to connect them natively (Fraunhofer Gesellschaft in Berlin, University of Erlangen, University of Essen and University of Münster).

6.1.2. External connectivity

There are several external peers to the 6WiN. All of them are peerings with (E)BGP.

- 6bone: The JOIN Project in Münster still hosts a large 6bone backbone node, so there is a connection to the 6bone in Münster.
- Deutsche Telekom: Deutsche Telekom (resp. T-Systems Nova Berkomp) is active in the Eurosix project and has Pops as a part of that network. In 2002 JOIN and Berkomp had a cooperation project (“IPv6 Showcase”), during which 6WiN and a similar network at Berkomp was build. That project has ended now, but the native line to Berkomp is still in existence.
- 6NET: Of course 6WiN is connected to the 6NET in Frankfurt.
- GÉANT: Since late 2003 GÉANT has offered IPv6 connectivity to every NREN. As described above, in this case 6WiN is a dedicated IPv6-only network, next to the IPv4-only network G-WiN. As GÉANT is only directly connected to G-WiN there is no direct line from GÉANT to 6WiN. Therefore the only way to connect 6WiN to GÉANT was over a IPv6-in-IPv4-tunnel. This is the only external tunnel in 6WiN.
- DECIXv6: DECIX offers an IPv6 IX in Frankfurt in the same building as the 6WiN PoP. 6WiN has a presence in that IX and has currently a peering with nine industrial partners there. Some of these peerings are used as a transit peering, exchanging a full BGP table.
- BCIX: A similar IX exists in Berlin. 6WiN is also connected there, but until now there are only two peering partners connected.

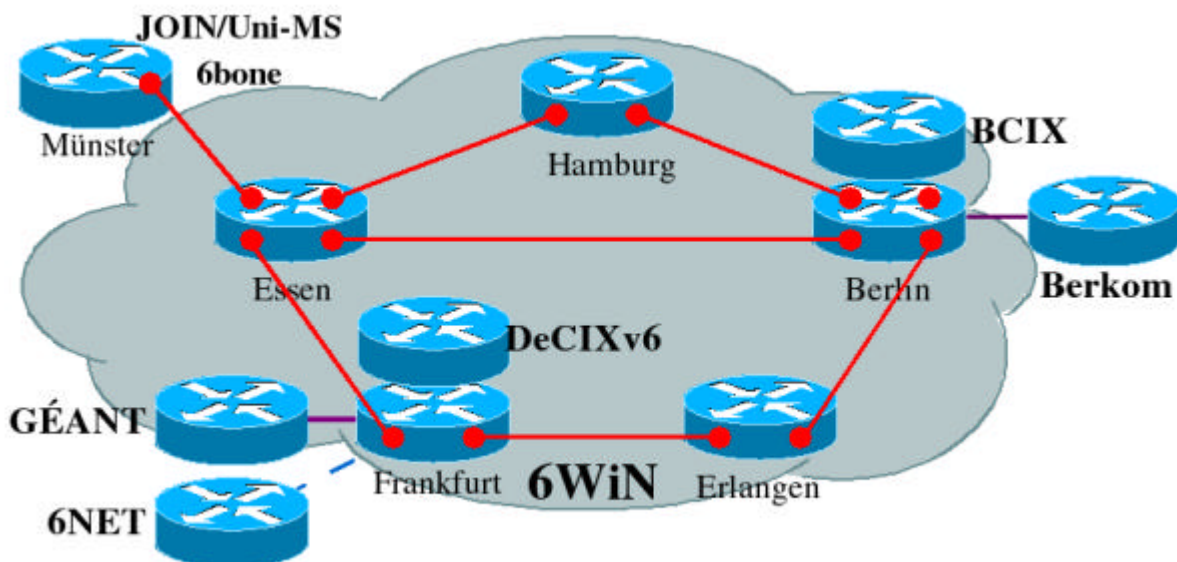


Figure 6-3: 6WiN connections

6.1.3. Addressing

DFN's original prefix 2001:638::/35 got split into several blocks of /40-prefixes. Every one of the five core routers was assigned two of these blocks, e.g. 2001:638:400::/40 and 2001:638:500::/40 to 'Essen'. From these blocks a standard /48-Prefix was assigned to every facility that is connected to this 6WiN PoP, e.g. 2001:638:500::/48 to Münster.

The prefix 2001:638:0::/48 was reserved for 6WiN backbone addressing. Every 6WiN PoP got a /56-prefix within this prefix, e.g. 2001:638:0:500::/56 for Essen.

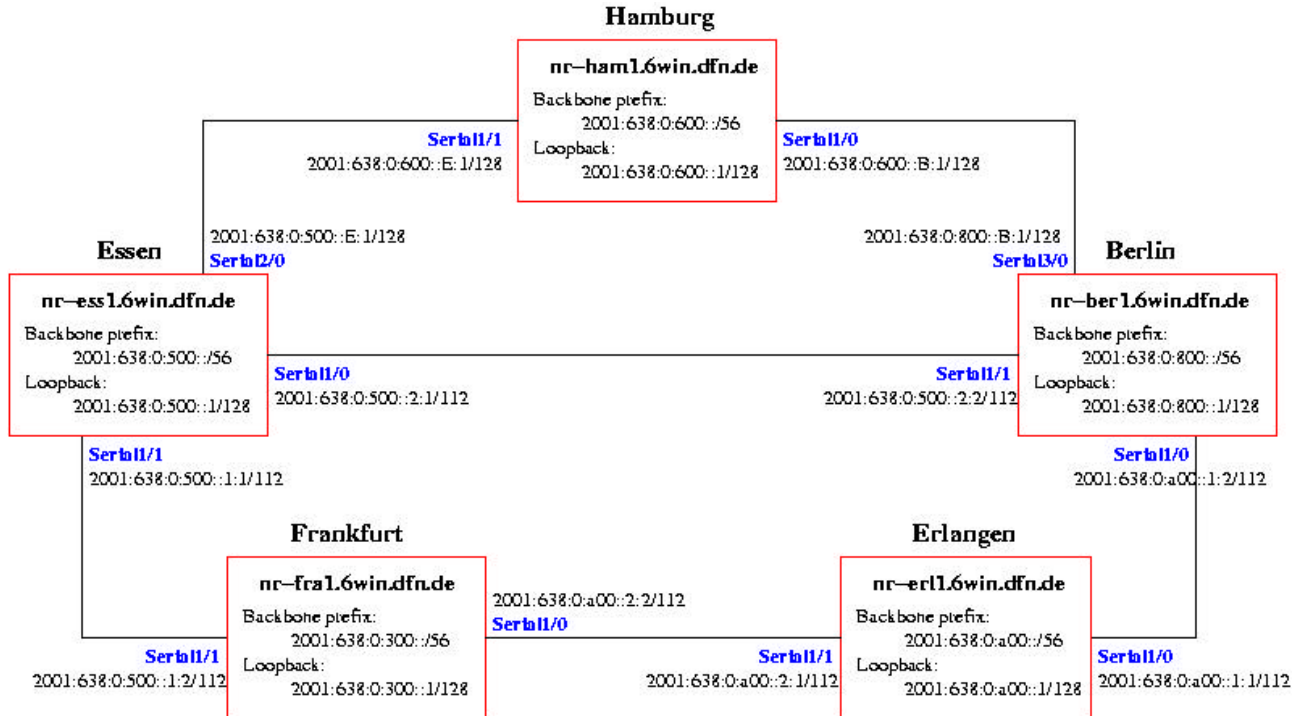


Figure 6-4: 6WiN addressing overview

6.1.4. Internal Routing

Routing inside the 6WiN is done with IS-ISv6 and iBGP. All six routers (core plus Münster) share a level-2 IS-IS cloud. The five core routers are full-meshed iBGP, the router in Münster is integrated in this BGP network via a route-reflector in Essen.

All backbone routes are propagated through IS-IS, while customer prefixes and external routes are propagated through iBGP. To reduce the size of the routing tables, the backbone /56-prefix on every router is aggregated in IS-IS, and the customer prefixes are aggregated in BGP to the routers' /40-prefix block. Only the /35-prefix is announced to external peers.

6.1.5. Multicast

At the end of 2003 multicast functionalities became available in production versions of Cisco IOS. After some tests Multicast (PIM, etc.) was activated and 6WiN was configured to use 6NET's multicast network "m6net" for global IPv6 multicast connectivity. This multicast traffic is offered to all 6WiN customers, but so far only two facilities are using it.

6.1.6. Configuration example

As mentioned above we use Cisco7206 routers running IOS 12.3(4). Configuration examples for this router are given in Appendix C.

6.1.7. Future usage of 6WiN

Despite the fact that there are already customers connected, the 6WiN is still a test network. It can be used for tests in the future, e.g. OSPF, advanced multicast techniques like BSR or embedded RPs, or for different addressing schemes.

6.2. Pilot NREN IPv6 Service (UKERNA: UK)

Since 1999 UKERNA has supported early IPv6 trials, in particular providing management for the Bermuda 2 IPv6 project [Bermuda]. The Bermuda project was able to utilise native IPv6 links over the JANET ATM Managed Bandwidth Service. However, the latest instance of JANET no longer supports ATM (the network has moved to high-capacity 10Gbit/s PoS), which means that most connectivity has dropped back to IPv6-in-IPv4 tunnels for the time being.

In May 2002, UKERNA launched an IPv6 Experimental Service on JANET, to which UK universities and colleges can connect. The service primarily supports research activities within the UK, and is built upon a star topology network of IPv6 in IPv4 tunnels (an informal service had been available for about five years prior to this).

Organisations connecting to the service receive a /48 prefix from the JANET IPv6 space 2001:630::/32, and an IPv6 over IPv4 tunnel to the centrally managed Cisco 7505 router which forms the hub of the pilot network. This router is located at the University of London Computer Centre (ULCC), and is managed by the JANET NOSC, the same team who operate JANET's IPv4 core network. This router also hosts JANET's connection to the 6bone, and has native links to two border routers, one servicing 6NET, and the other connecting to UK6X and the IPv6 trials on the LINX (London INternet eXchange).

The aim of the experimental service is allow UKERNA and JANET connected organisations to gain early experience in operating IPv6 based networks and services, and to focus the JANET community's IPv6 efforts.

This pilot network continues at the time of writing. The JANET became dual-stack over the summer of 2003. The existing OSPFv2 IGP was replaced with IS-IS, such that both protocols now share a common IGP.

6.2.1. Current services

The services currently offered under the JANET IPv6 Experimental Service include:

- *Native IPv6 connections:* It is not expected that many organisations will initially connect via native links, however where this is possible, UKERNA endeavours to support such requests.
- *Manually configured IPv6 in IPv4 tunnels:* The most common initial connection method. A number of JANET sites are already connected to the service in this way.
- *6to4 tunnels and 6to4 relay* [RFC3068]: A 6to4 relay router is already available within JANET on the "well-known" anycast address (192.88.99.1).
- *IPv6 Routing Options:* Standard routing between the end organisation and the JANET IPv6 Experimental service is currently provided using static routes. RIPng and BGP are available on request, as are further routing protocols, as and when Cisco IOS software supports them. However, continued support of all available routing protocols to end organisations may not be available beyond the pilot phase.

6.2.2. Future services

Services that are currently being introduced include:

- IPv6 tunnel broker/server: A tunnel broker and server are planned to be made available, to allow individual users to establish tunnelled connectivity from dual-stack hosts.
- IPv6 looking glass: A looking glass allows connected organisations and peers to view the state of IPv6 routing on the JANET network via a web browser. This will allow read-only access to the JANET BGP IPv6 routing table, and other troubleshooting information. In the interim, ASPathTree is being used to provide a visual representation of the BGP4+ routing table.
- IPv6 DNS: UKERNA plans to study IPv6 DNS issues for .ac.uk, and IPv6 access to JANET central services. Reverse DNS for the JANET prefix, 2001:630::/32, is delegated, and best-effort supported by the JANET NOSC.

The planning and provision of these pilot IPv6 services on JANET will give UKERNA results and experience that will be fed into the 6NET project.

The current challenge is to encourage the 17 regional networks connected to the national backbone to deploy IPv6 services. The LeNSE network, which connects Southampton, is deploying 6PE for this purpose in the Summer of 2004.

6.2.3. Applying for Connection

To join the JANET IPv6 Experimental Service an application form must be completed by the JANET connected organisation, and forwarded to the UKERNA JANET Customer Service Desk.

On joining, the organisation will receive a /48 prefix from the JANET production IPv6 space, 2001:630::/32. The JANET NOSC will then handle the connection to the pilot IPv6 network.

6.2.4. Support Issues

Self-support is available from the ipv6-users@jiscmail.ac.uk mail list, which is monitored by UKERNA, JANET NOSC and experienced IPv6 users from the Universities involved in the Bermuda project.

6.3. NREN migrating from 6bone trials (GRNET: Greece)

Here we describe the experience in GRNET of moving from 6bone towards production IPv6 services.

6.3.1. IPv6 Service in GRNET

GRNET has run internal IPv6 pilot projects investigating IPv6 migration and IPv6 applications since 2000. GRNET participates in the 6bone testbed and has been allocated the pTLA 3FFE:2D00::/24 which will be used in order to cover the address needs of most of the Mediterranean countries and their Service Providers. Each Mediterranean country will be allocated a /27 address space. From the 3FFE:2D00::/24 address pool, GRNET supplies a /48 to every University or Technical Educational Institute that gets connected to 6bone.

6.3.2. Services in production

GRNET currently offers the following IPv6 services:

1. Connection to the 6bone by means of static tunnels.
2. Allocation of a range of addresses via the pTLA that it holds.
3. DNS services: forward and reverse zone data.

GRNET uses a dual stack router (Cisco 4500) that is connected to the 6bone. Using this router, IPv6 (over IPv4) tunnels are established with GRNET client sites. GRNET can provide both of the IPv6 addresses that are necessary for the tunnels edges, in particular from the address range 3FFE:2D00:1::/48.

GRNET also undertakes the creation of the necessary records to its DNS servers, regarding the forward and reverse data, for both tunnels edges, that are named as follows:

`<remote_site>-GRNET.ipv6.GRNET.gr` and `GRNET-<remote_site>.ipv6.GRNET.gr`

e.g.

`upatras-GRNET.ipv6.GRNET` and `GRNET-upatras.ipv6.GRNET.gr`

The routing method that is used for the tunnels is either static routing or BGP4+.

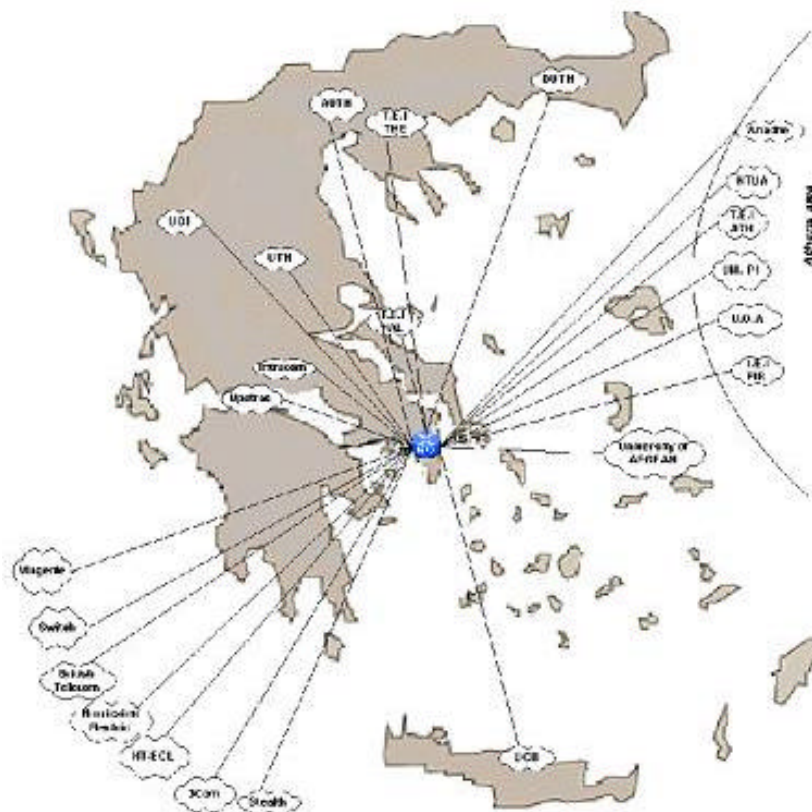


Figure 6-5: GRNET connectivity

Regarding DNS, GRNET undertakes all necessary actions regarding the delegation of reverse zone data to peripheral servers. Given the pilot nature of the network, requests of clients that do not dispose their own servers and demand primary server services regarding forward and reverse data, are examined on an individual basis. For the moment, GRNET servers do not offer IPv6 secondary name services.

Figure 6-5 demonstrates the established IPv6 tunnels between the GRNET router and the existing client sites.

6.3.3. Services planned for production

GRNET is currently in the process of investigating new services and applications for its IPv6 community. Some of these include:

- IPv6 over MPLS. GRNET has already implemented an MPLS enabled IPv6 testbed (see Figure 6-6) so as to investigate issues concerning IPv6 over MPLS deployment.
- Tunnel broker. A tunnel broker will be set up on a workstation Sun Ultra 1 running Solaris 8. The ipv6tb implementation of TILAB will be used.
- Setup of network monitoring tools for the IPv6 network. The monitoring tools that will be used are ASpath tree, weathermap and MRTGs. The IPv6 community of GRNET will be able to real time monitor the IPv6 network through the web.
- A NAT/PT server will be set up to allow communication between IPv4-only and IPv6-only services.

Figure 6-6: The GRNET MPLS-enabled pilot

7. Survey of NREN deployment experience

During March 2004, a survey of NREN IPv6 deployments was carried out within the 6NET project scope. The goal of the survey was to understand the current status of the NRENs, and to gather issue-specific notes for the benefit of the NREN community.

The questions and issues covered by the survey were:

- Equipment used
- IPv6 addressing plan for core
- IPv6 addressing plan for universities
- Routing protocols used
- Any special considerations with IPv4 routing
- IPv6 performance issues
- Running a 6to4 relay service? Is it open to externals?
- Running a tunnel broker?
- Other transition/access methods? Manual tunnels? 6PE?
- How many universities are connected to your IPv6 NREN service
- NREN running dual-stack services? DNS? Mail? Web? FTP? NTP? Other?
- IPv6 transport to academic nameservers?
- Main network management and monitoring tools used
- Any special security concerns in deploying IPv6
- Mistakes made?
- Lessons learnt for others deploying?
- Unexpected results from transition work
- Technology gaps (e.g. vendor features)
- Standardisation gaps (e.g. in IETF)

Responses were received from:

FCCN (Portugal), Hungarnet (Hungary), SWITCH (Switzerland), UNINETT (Norway), CESNET (Czech Republic), DFN (Germany, via JOIN), UNI-C (Denmark), SURFnet (Netherlands), UKERNA (UK), AConet (Austria), FUNet (Finland), GRNet (Greece), Poznan (Poland) and GARR (Italy).

The responses are gathered and summarised in the following sections.

7.1. Overview of IPv6 service

Information on the general nature of the IPv6 service was requested.

7.1.1. Nature of service

By far the most common answer was dual-stack.

Two NRENs run a parallel infrastructure for IPv6, but two of those are for IPv6 multicast only, where the multicast features are not supported on the regular backbone routers. The 6WiN, in the DFN, uses six 34Mbit lines between five IPv6 PoPs. The other NREN uses dedicated VLANs for IPv6 in a switched Gigabit Ethernet backbone.

Two NRENs are using MPLS, one of which (SURFnet) only introduced MPLS to obtain line rate IPv6 without an expensive set of line card upgrades on their (quite new) Cisco equipment.

For connecting sites, statically configured tunnels are used in almost all cases where native connectivity is not possible.

7.1.2. Equipment used

The router platforms used by the NRENs are all Cisco or Juniper, predominantly Cisco.

The most common equipment is Cisco 12000/GSR (most frequent is the 12400 series), 7200 and 7500/7600 series, with the 7206 VXR named by many.

The Juniper m5, m10(i), m20 and T320 routers are also in use, with the m10 most common.

7.1.3. IPv6 addressing plan for core

Each NREN has a /32 prefix allocated from RIPE-NCC in Europe. All NRENs offer /48 prefixes to end sites.

Allocations to regional networks (PoPs) are cited at the /40 boundary quite commonly, depending on the size of the region. For example, in SURFnet each PoP gets a /40 and each /40 is again split into 256 /48 prefixes to number the backbone links and customer networks.

Addressing within the core network itself tends to use a single /48, with a /56 allocated to each region in the PoP hierarchy. For example, GwiN (DFN) uses /56's for the level 1 PoPs and /60's for the level 2 PoPs, with a /64 for addressing the customer's link (their site gets a whole, and different, /48 prefix).

Most NRENs use /64 prefixes to address links (including point-to-point), but in some cases point-to-points and loopbacks are taken from a single /64 prefix, where (in one case) the bits inside the prefix indicating PoPs, lines, router numbers, etc., so they are quite informative, and there the point-to-point addresses use a /112 prefix. In other cases /126's are used for point-to-point links. However, use of /64's is most common.

Some specific core network addressing plans have been cited at:

- <http://www.ipv6-tf.com.pt> (FCCN)
- <http://www.ces.net/project/ipv6/addressing.html> (CESNET)

7.1.4. IPv6 addressing plan for universities

An NREN with a /32 allocation can address 2^{16} (over 65,000) end sites. The most common allocation scheme is a /40 per PoP and a /48 per university end site from that regional PoP (i.e. up to 256 regions of 256 networks).

It is possible to apply some “clever” coding tricks to such allocations. For example, in the GWiN, each level 1 PoP is assigned a /40 prefix corresponding to the SLA it has as routing prefix. That means if a level 1 pop is assigned the (core) prefix 2001:638:0:500::/56 the /40 prefix out of which the connected universities are addressed is 2001:638:500::/40.

In most NREN plans, the /48-prefix allocation leaves no space in between the allocations for university prefix length growth. It is not expected that universities will need more than a /48. If they nevertheless present a strong case in favour of getting more address space they with either get an additional /48 (without aggregation) or receive a longer prefix from a reserved address space higher up in the regional /40.

However, some NRENs are leaving “gaps” for bigger universities to allow them to grow into a large (aggregated) prefix later. In one case, the NREN is reserving a /44 for each, and is also recommending that the universities only use (for now) the first /52 of their prefix (though such an internal site guideline seems unnecessary).

Some specific end site network addressing plans have been cited at:

- <http://www.ipv6-tf.com.pt> (FCCN)
- <http://www.ces.net/project/ipv6/addressing.html> (CESNET)

7.1.5. Routing protocols used

The most common solution in the NRENs is IS-IS for IPv4 and IPv6, but it is not a majority solution because of the range of solutions/combinations possible. The OSPFv2 and OSPFv3 combination was a close second. The lesson from the responses is that the NRENs have a wide range of solutions deployed, often for historical reasons.

Other responses included OSPFv3 with IS-IS + OSPFv2 for IPv4+6PE, and OSPFv2 for IPv4 with IS-IS for IPv6 only, so all four major combinations of OSPF and IS-IS are in use somewhere in a 6NET NREN, although the first two combinations stand out.

7.1.6. Any special considerations with IPv4 routing

None were cited.

7.1.7. IPv6 performance issues

Most responses cited no issues. However, some NRENs reported that lack of hardware support for packet forwarding in equipment was an important issue.

One comment suggested performance was not an issue while the traffic level is generally very low; problems with non-hardware solutions will only arise when traffic levels increase.

The 6WiN parallel network only uses 34Mbit/s lines where sometimes there are peaks that are capped at 34Mbit/s. This is the reason why for example JOIN can not receive full NNTP-feeds via those lines which would need up to 40Mbit/s alone.

One NREN commented that by using 6PE they hope to stay away from performance issues.

Another suggested it's hard for the end-users to distinguish whether they are using IPv6 or IPv4, and thus complain if/when performance is poor.

Overall, performance issues were not a concern, though it is not clear in how many cases the performance has been really stretched.

7.2. Support mechanisms

Here we present the summary of responses for IPv6 support methods.

7.2.1. 6to4 relay service deployment

Around 75% of the 6NET NRENs run a 6to4 relay service in support of their users, and of those the majority of relays are run openly without restriction (only two have restrictions, such that only the NREN's own sites can use the relay).

A number of NRENs are considering applying further restrictions on their open 6to4 relays, but presumably have not yet detected any abuse that would lead them to bring those plans forward.

Thus 6to4 appears popular in the NRENs. Actual usage levels would be interesting to investigate.

7.2.2. Tunnel broker deployment

In contrast, only around 25% of the NRENs run a tunnel broker, of which one is only in an early test phase (using the Hexago commercial broker) and another is in its infancy. One NREN (UNINETT) used to run a broker but has since turned it off.

One respondent said they would run a broker if an open source solution was made readily available and was easy to deploy. This may be one reason why brokers are currently less popular than 6to4 (along with the higher ongoing manpower support costs).

This is a little surprising as a broker appears to be a reasonable, managed way to build an IPv6 community for users at sites currently not deploying IPv6 natively, or even home broadband users whose ISP have no ISP support (and who would get good service from a broker in the same country).

7.2.3. Other transition/access methods

By far the most common transition mechanism for end site universities is manually configured tunnels from the NREN to the university.

The only other response from more than one NREN was on the use of 6PE, which is used by three of the respondents.

7.3. Operational notes

The operational responses are summarised here.

7.3.1. Universities connected to the IPv6 NREN service

The reported count of connected sites in NRENs is as follows:

NREN	Sites
------	-------

FCCN	5
HUNGARNET	9
UNINETT	4
CESNET	14
UNI-C/Denmark	2
ACONet	9
GRNet	7
Poznan	0
GARR	14
JOIN/DFN	29
SURFnet	22
Total = 11 NRENs	115 university sites

The DFN lists (sites and prefixes) can be found here:

- http://www.6win.de/6WiN/list_6win_active_tunnels.php?lang=en
- http://www.6win.de/6WiN/list_DFN_Prefixes_complete.php?lang=en

7.3.2. NREN dual-stack services

The most common services supported, in decreasing order of number of responses received, are:

- DNS
- FTP (including large FTP archive: SWITCHmirror)
- Web
- NNTP/INND (Usenet News)
- NTP
- Mail (SMTP), IMAP
- Dual-stack web proxy
- Media streaming, Darwin streaming server (experimental)
- SIP
- SSH
- IRC
- Quake server
- LDAP

The first three responses were quoted rather more commonly than the other answers. Of the rest, only NNTP and NTP had more than one reply.

7.3.3. IPv6 transport for nameservers

IPv6 transport for DNS was enabled in only a minority of the NREN networks, including Funet (everywhere), GARR (in its own primary DNS server), HUNGARNET (in parts), FCCN (the secondary DNS server), CESNET and SWITCH (in the name servers for both CH and LI).

In most countries, the academic domains are not distinguishable from other domains, indeed only the .uk and .at domains have .ac for academic sites/subdomains.

FCCN has enabled any registrar to register DNS domains under .PT with IPv6 glue records. SWITCH plans to do likewise in the near future.

7.3.4. Main network management and monitoring tools

There is a review of tools available in [6NET-D632].

In this survey, the responses in declining order of frequency included (aside from ping and traceroute):

- Rancid
- ASpathtree
- Nagios
- Looking glass
- Argus
- Netsaint
- Cricket
- Netflow
- HP Openview (IPv4 transport monitoring of IPv6 routers)
- MRTG
- Weathermap
- Ethereal
- Multicast Beacon
- Net-snmp
- PCHAR
- Big Brother

The first four responses were quite a way ahead of the others.

7.3.5. Special security concerns in deploying IPv6

The general response was to follow IPv4 best practice for IPv6.

Although obvious, users/sites should be aware that if they enable IPv6 on a host, their service may be available via IPv6 also, and thus potentially open.

One NREN uses a tool that translates higher-level filter policy descriptions into access lists (ACLs) that can be installed on routers. This tool was made IPv6-aware within the 6NET project.

Transition security issues are reported in [6NET-D622].

7.4. Deployment experience

NREN feedback on deployment experience is given in this section.

7.4.1. Mistakes and lessons learnt

Almost all NRENs had no mistakes to confess to, or lessons to report.

One said it felt that MPLS was not abandoned in time (i.e., together with ATM).

Another with a parallel infrastructure (DFN/JOIN) said that while having a separate infrastructure for IPv6 in the beginning is nice for tests we regard it as not the ideal way to go for NREN IPv6 deployment, because customers need a new line to connect natively for which most don't like to pay. Instead we use IPv6-in-IPv4 tunnels which are stable and topologically sound, but of course these do impose extra overhead. For DFN itself the main problem maintaining a parallel network of course also lies in the additional work required for management and the additional money necessary to pay for the lines etc. So GWiN should become dual-stack as soon as possible. Another drawback is also that we don't yet have experience running both IPv6 and IPv4 at once on GWiN. Those tests will still need to be performed.

Another NREN commented on connectivity issues. Isolated IPv6 islands will first try to use the IPv6 address, then try the IPv4 address of our web/mail services. This gives slow response, so one should make sure not to run IPv6 (and add AAAA entries to the DNS) if connectivity to the core is not available.

Finally, one NREN commented that it can take some energy to un-deploy some methods, such as injecting IPv6 using VLANs (when you want to enable IPv6 in the IPv4 first-hop routers that is).

7.4.2. Unexpected results from transition

Almost all NRENs responded with no comment here.

The only comment received was that DNS issues were more complex than expected initially, due to the many possible combinations of datatypes in query/response (A/AAAA/NS) and transport (IPv4/IPv6), and tree-walk from the root of the name-space.

7.4.3. Technology gaps

A number of gaps were reported in the responses:

- Lack of availability of IPv6 Cisco PIX firewall
- Lack of VRRP/HSRP for IPv6;
- IPv6 support on lower end devices/switches, e.g Catalyst 35xx and Catalyst 37xx;
- The domain registration software used for some TLDs currently isn't capable of registering IPv6 "glue" and putting it in the TLD zones (the IPv6-related capability required for a DNS registry);

- Problems in deploying IPv6 multicast on old GSR cards. Problems with support on Catalyst 4000, 35xx and 37xx. Problems with some older 6500 units with SUP1 cards, mainly related to IPv6 multicast support - both in routers and switches;
- In terms of managing the network via IPv6 we are still waiting for SNMP with IPv6 transport;
- DHCPv6 relay agents are another feature needed but probably not within the core;
- P routers in current 6PE environment don't respond with ICMP hello's for IPv6 which makes traceroute output less clear and makes troubleshooting more difficult. We understand that IOS 12.2S should have the features to enable this, but this is currently not available for the Cisco 12000 platform;
- A growing number of SURFnet customers use a Catalyst 3750 to connect to the network. IPv6 support on this platform should make it easier to implement IPv6 in their networks;
- Bugs found in Cisco IOS: SNMP over IPv6, ability to poll for IPv6 traffic rates separately from IPv4. Bugs were found during deployment of dual stack. IOS CLI is inconsistent in many cases when used for IPv6 instead of IPv4. Inconsistent vendor documentation was also found;
- Lack of NetFlow export for IPv6 on the Catalyst 6500/7600 OSR.

The SNMP and MIB issue was mentioned by more NRENs than any other issue.

7.4.4. Standardisation gaps

The following issues were cited:

- Detailed analysis of IPv6 deployment in campus/enterprise environments
- Interdomain multicast
- Multihoming solutions
- RPSLng
- IPFIX
- Missing or unimplemented MIBs

6NET is working on enterprise scenarios in [6NET-D233].

8. Conclusions

This cookbook has presented a summary of IPv6 transition tools and case studies applicable to the NREN scope in Europe. By far the most common method for introduction of production quality (to the level of existing IPv4 networks) IPv6 services is via transition to dual-stack networking on the national networks. As the core network also becomes dual-stack, the deployment issues for IPv6 are then pushed to the edge (and these issues are discussed in the site/university cookbook, [6NET-D233]).

We have shown three examples of dual-stack IPv4-IPv6 NREN deployments in this cookbook, and other NRENs have already followed suit with their own dual-stack deployments (see the responses in Section 7, and case studies in Section 5). Other transition methods may be used in a small number of cases in the shorter term, but these are expected to be replaced by dual-stack in the longer term. IPv6 over MPLS (6PE) may be used also where hardware upgrades are otherwise required to deliver IPv6 over multi-gigabit links, as shown by the SURFnet case study in Section 5. A parallel infrastructure is also possible, as shown by the 6WiN case study in Section 6.

The selection of routing protocol in the NRENs is not so clear-cut, though IS-IS for both IPv4 and IPv6 is generally considered the better technical solution. However, running a different protocol for each IP version is also possible (see [6NET-D312]).

In the early to middle transition phase NRENs can be expected to offer supporting services where universities have not yet deployed them. Two examples would be the provision of a tunnel broker and a 6to4 service (a relay). These are both discussed in Section 3. Other services should also be considered, e.g. IPv6 transport for the top-level academic DNS domains, IPv6 accessibility to common NREN services (e.g. FTP mirrors, news feeds). Network management and monitoring is also very important, but these aspects are covered in [6NET-D632]. Similarly IPv6 Multicast is reported in [6NET-D312].

The survey of NREN IPv6 statuses in Section 7 raised some interesting statistics, e.g.

- 75% run a 6to4 relay (mainly openly), while only 25% run a tunnel broker;
- Rancid, Nagios, ASpathtree and a looking glass are the most commonly used monitoring and management tools;
- No IPv6 performance issues were being reported;
- Separate OSPF (v2 and v3) and single IS-IS for IPv4 and IPv6 are the two common routing protocol combinations;
- Manually configured tunnels are the de facto way to offer IPv6 to non-native end sites;
- There are at least 115 IPv6-connected universities across 11 NRENs.

Security is an important issue, and is covered in detail in [6NET-D622].

The next major version of this cookbook – Deliverable D2.2.4 – is due in December 2004; until then updated versions will be available from the 6NET project web site as new sections are added or updated.

The editor (Tim Chown) welcomes feedback via email to tjc@ecs.soton.ac.uk.

9. References

- [6NET-D211] “*Backbone network transition scoping report*”, 6NET Project deliverable D2.1.1, <http://www.6net.org/publications/deliverables/D2.1.1.pdf>
- [6NET-D221] “*NREN network transition scoping report*”, 6NET Project deliverable D2.2.1, <http://www.6net.org/publications/deliverables/D2.2.1.pdf>
- [6NET-D222] “*Initial IPv4 to IPv6 migration Cookbook for organisational/ISP (NREN) and backbone networks*”, 6NET Project deliverable D2.2.2, <http://www.6net.org/publications/deliverables/D2.2.2.pdf>
- [6NET-D233] “*Updated IPv4 to IPv6 transition Cookbook for end site networks/universities*”, 6NET Project deliverable D2.2.2, <http://www.6net.org/publications/deliverables/D2.3.3.pdf>
- [6NET-D312] “*IPv6 cookbook for routing, DNS, intra-domain multicast, inter-domain multicast and security*”, 6NET Project deliverable D3.1.2, <http://www.6net.org/publications/deliverables/D3.1.2.pdf>
- [6NET-D622] “*Operational procedures for secured management with transition mechanisms*”, 6NET Project deliverable D6.2.2, <http://www.6net.org/publications/deliverables/D6.2.2.pdf>
- [6NET-D632] “*Interim report on the implementation of tools and operational procedures*”, 6NET Project deliverable D6.3.2, <http://www.6net.org/publications/deliverables/D6.3.2.pdf>
- [Bermuda] The Bermuda 2 IPv6 Project, <http://www.ipv6.ac.uk/bermuda2/>.
- [D9.3] GÉANT Deliverable D9.3: “*IPv6 Testing*”, <http://www.dante.net/tf-ngn/D9.3.pdf>
- [D9.6] GÉANT Deliverable D9.6: “*Report on IPv6 Experiments and Status*”, http://www.dante.net/tf-ngn/D9.6-IPv6_test.pdf
- [Lind01] “*Scenarios and Analysis for Introducing IPv6 into ISP Networks*”, IETF Internet Draft, draft-ietf-v6ops-isp-scenarios-analysis-02, April 2004, Work in Progress
- [Martini01] “*Transport of Layer 2 Frames Over MPLS*”, L. Martini et al., IETF Internet Draft, draft-martini-l2circuit-trans-mpls-13, December 2003, Work in Progress.
- [Martini02] “*Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*”, L. Martini et al., IETF Internet Draft, draft-martini-l2circuit-encap-mpls-06, November 2003, Work in Progress.
- [Martini03] “*Encapsulation Methods for Transport of ATM Cells/Frame Over IP and MPLS Networks*”, L. Martini et al., IETF Internet Draft, draft-martini-atm-encap-mpls-01, June 2002.

-
- [Ooms01] J. DeClercq, D. Ooms, S. Prevost, F. Le Faucheur, “*Connecting IPv6 Islands over IPv4 MPLS using IPv6 Provider Edge Routers (6PE)*”, IETF Internet Draft, draft-ooms-v6ops-bgp-tunnel-03, April 2004, Work in Progress.
- [RFC1483] “*Multiprotocol Encapsulation over ATM Adaptation Layer 5*”, J. Heinanen, IETF RFC, July 1993
- [RFC1772] “*Application of the Border Gateway Protocol in the Internet*”, Y. Rekhter, P. Gross, IETF RFC, 1995.
- [RFC2373] “*IP Version 6 Addressing Architecture*”, R. Hinden, S. Deering, IETF RFC, 1998.
- [RFC2492] “*IPv6 over ATM Networks*”, G. Armitage, P. Schuler, M. Jork, IETF RFC, Jan 1999.
- [RFC2547] “*BGP/MPLS VPNs*”, E. Rosen, Y. Rekhter, IETF RFC, 1999.
- [RFC2784] “*Generic Routing Encapsulation (GRE)*”, D. Farinacci, T. Li, S. Hanks, D. Meyer, P. Traina, IETF RFC, 2000.
- [RFC2796] “*BGP Route Reflection – An Alternative to Full Mesh IBGP*”, T. Bates, R. Chandra, E. Chen, IETF RFC, 2000.
- [RFC2858] “*Multiprotocol Extensions for BGP-4*”, T. Bates, Y. Rekhter, R. Chandra, D. Katz, IETF RFC, 2000.
- [RFC2893] “*Transition Mechanisms for IPv6 Hosts and Routers*”, R. Gilligan, E. Nordmark, IETF RFC, 2000, under update as draft-ietf-v6ops-mech-v2-02, February 2004.
- [RFC3056] “*Connection of IPv6 Domains via IPv4 Clouds*”, B. Carpenter, K. Moore, IETF RFC, February 2001
- [RFC3107] “*Carrying Label Information in BGP-4*”, Y. Rekhter, E. Rosen, IETF RFC, 2001.
- [V6OPS] IETF IPv6 Operations WG, <http://www.ietf.org/html.charters/v6ops-charter.html>

10. Appendix A: Hungarnet's IPv6 over MPLS Configuration Example (Cisco)

```
version 12.2
service timestamps debug datetime show-timezone
service timestamps log datetime show-timezone
service password-encryption
no service dhcp
hostname pecs.6net.hbone.hu
!
!enable secret <removed>
!
ip subnet-zero
no ip source-route
ip cef
!
no ip bootp server
ipv6 unicast-routing
ipv6 cef
mpls ldp logging neighbor-changes
tag-switching tdp router-id Loopback0
call rsvp-sync
!
!
interface Loopback0
 ip address 195.111.97.230 255.255.255.255
 ipv6 address 2001:738:7800::1/128
!
interface Ethernet0/0
 description - IPv4 Management Port -
 ip address 195.111.98.134 255.255.255.240
 duplex full
 no cdp enable
!
interface GigabitEthernet0/0
 description -- Towards Pecs MPLS P router --
 ip address 195.111.103.250 255.255.255.252
 ip ospf network point-to-point
 duplex full
```

```
speed 1000
media-type gbic
negotiation auto
mpls label protocol tdp
tag-switching ip
no cdp enable
!
interface FastEthernet2/0
  ipv6 address 2001:0738:7801::1/64
  duplex auto
  speed auto
  no cdp enable
!
router ospf 100
  log-adjacency-changes
  passive-interface default
  no passive-interface GigabitEthernet0/0
  network 195.111.97.230 0.0.0.0 area 0
  network 195.111.103.248 0.0.0.3 area 0
!
router bgp 1955
  no bgp default ipv4-unicast
  bgp log-neighbor-changes
  neighbor 6PE peer-group
  neighbor 6PE remote-as 1955
  neighbor 6PE description IPV6 over MPLS peering peer-group
  neighbor 6PE update-source Loopback0
neighbor 195.111.97.229 peer-group 6PE
!
  address-family ipv4
  no auto-summary
  no synchronization
  exit-address-family
!
  address-family ipv4 multicast
  no auto-summary
  exit-address-family
!
  address-family ipv6
```

```
neighbor 6PE activate
neighbor 6PE send-label
neighbor 195.111.97.229 peer-group 6PE
network 2001:738:7800::/48
redistribute connected
no synchronization
exit-address-family
!
...
```

On cntrl.6net.hbone.hu: 6PE and BGP route reflector router:

```
version 12.2
service timestamps debug datetime show-timezone
service timestamps log datetime show-timezone
service password-encryption
!
hostname cntrl.6net.hbone.hu
ip cef
!
!
no ip bootp server
ipv6 unicast-routing
ipv6 cef
mpls ldp logging neighbor-changes
tag-switching tdp router-id Loopback0
no call rsvp-sync
!
Interface Loopback0
 ip address 195.111.97.229 255.255.255.255
 ipv6 address 2001:738::1/128
!
Interface GigabitEthernet0/0
 description -- Towards central switch --
 no ip address
 no ip redirects
 no ip proxy-arp
 ip route-cache flow
 no ip mroute-cache
```

```
duplex full
speed 1000
media-type gbic
negotiation auto
no cdp enable
!
interface GigabitEthernet0/0.35
description -- server segment --
encapsulation dot1Q 35
no ip redirects
no ip proxy-arp
ipv6 address 2001:738:0:402::1/64
ipv6 enable
no cdp enable
!
interface GigabitEthernet0/0.807
description cntrl.6net.hbone.hu - OSPFv2/MPLS segment
encapsulation dot1Q 807
ip address 195.111.96.26 255.255.255.252
no ip redirects
no ip proxy-arp
mpls label protocol tdp
tag-switching ip
no cdp enable
!
router ospf 100
log-adjacency-changes
passive-interface Loopback0
network 195.111.96.24 0.0.0.3 area 0
network 195.111.97.229 0.0.0.0 area 0
!
router isis
net 49.0001.0000.0000.0001.00
passive-interface GigabitEthernet0/0.35
passive-interface POS2/0
passive-interface Loopback0
passive-interface Tunnel3
!
address-family ipv6
```

```
no adjacency-check
exit-address-family
is-type level-2-only
metric-style wide
max-lsp-lifetime 65535
lsp-refresh-interval 65000
spf-interval 1 1 10
prc-interval 1 1 10
lsp-gen-interval 5 1 50
no hello padding
log-adjacency-changes all
!
router bgp 1955
  bgp router-id 195.111.96.22
  no bgp default ipv4-unicast
  bgp cluster-id 3278856218
  bgp log-neighbor-changes
  neighbor 6PE peer-group
  neighbor 6PE remote-as 1955
  neighbor 6PE description IPV6 over MPLS peering peer-group
  neighbor 6PE update-source Loopback0
  neighbor 195.111.97.230 peer-group 6PE
  !
  address-family ipv4
    no auto-summary
    no synchronization
    exit-address-family
  !
  address-family ipv4 multicast
    no auto-summary
    exit-address-family
  !
  address-family ipv6
    neighbor 6PE activate
    neighbor 6PE route-reflector-client
    neighbor 6PE send-community
    neighbor 6PE soft-reconfiguration inbound
    neighbor 6PE send-label
    neighbor 195.111.97.230 peer-group 6PE
```

```
redistribute connected route-map redistribute-loopback
network 2001:738::/32
no synchronization
exit-address-family
!
```

11. Appendix B: Funet's Core Network Configuration Example (Juniper)

```
interfaces {
  lo0 {
    unit 0 {
      family iso {
        address 49.0001.1931.6600.5180.00;    "IS-IS address from IPv4"
      }
      family inet6 {
        address 2001:708:0:BB:eeee:ffff:0000:1111/128; "Loopback"
      }
      family inet {
        [...]
      }
    }
  }

  so-X/X/X {
    unit 0 {
      [...]
      family iso;                                "For IS-IS"
      family inet6 {
        address 2001:708:0:BB:aaaa:bbbb:cccc:dddd/112; "Core"
      }
      family inet {
        [...]
      }
    }
  }
}

protocols {
  isis {
    no-ipv4-routing;
    export ipv6-to-isis;
    level 1 disable;
    interface so-X/X/X.0 {                      "Core connections"
      level 2 metric 2;
    }
  }
}
```

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

policy-options {
    policy-statement ipv6-to-isis {
        from {
            protocol [ direct local static isis ];
            family inet6;
        }
        then {
            accept;
        }
    }
}
```

And respectively on Cisco:

```
--8<--
interface POSX/Y
[...]
    ipv6 address 2001:708:0:BB:aaaa:bbbb:cccc:dddd/112
    ipv6 router isis
    isis circuit-type level-2-only
    isis metric 2
!
router isis
    passive-interface Loopback0
!
    address-family ipv6
    redistribute static
    redistribute connected
    no adjacency-check
    exit-address-family
    is-type level-2-only
```

net 49.0001.1931.6600.5181.00

metric-style transition

log-adjacency-changes

!

12. Appendix C: 6WiN Configuration Examples (Cisco)

6to4-"Sink" with corresponding route:

```
interface Tunnel64
  no ip address
  no ip redirects
  ipv6 address 2002:C1AE:4BF9::/48 eui-64
  tunnel source Loopback0
  tunnel mode ipv6ip 6to4
  !
  ipv6 route 2002::/16 Tunnel64
```

Configuration of a few interfaces (2 core, one to 6net):

```
interface Serial1/0
  description NR-Erlangen
  no ip address
  ipv6 address 2001:638:0:A00::2:2/112
  no ipv6 mld router
  ipv6 router isis
  framing g751
  dsu bandwidth 34010
  down-when-looped
  serial restart-delay 0
  !
interface Serial1/1
  description NR-Essen
  no ip address
  ipv6 address 2001:638:0:500::1:2/112
  no ipv6 mld router
  ipv6 router isis
  framing g751
  dsu bandwidth 34010
  down-when-looped
  serial restart-delay 0
```

```
!  
interface POS2/0  
  description peering-to-6NET  
  no ip address  
no ip redirects  
  ipv6 address 2001:798:14:200::2/64  
  ipv6 enable  
  pos scramble-atm  
  pos flag c2 22  
!
```

IS-IS configuration on one of our routers:

```
router isis  
  net 49.0001.0300.0000.0001.00  
  is-type level-2-only  
  passive-interface Loopback6  
  passive-interface Tunnel64  
  !  
  address-family ipv6  
  summary-prefix 2001:638:0:300::/56  
  summary-prefix 2002::/16  
  exit-address-family  
!
```

BGP configuration for the core (router without any other exits):

```
router bgp 680  
  bgp log-neighbor-changes  
  neighbor internalv6 peer-group  
  neighbor internalv6 remote-as 680  
  neighbor internalv6 password * ****  
  neighbor internalv6 update-source Loopback6  
  neighbor 2001:638:0:300::1 peer-group internalv6  
  neighbor 2001:638:0:300::1 description ERL-to-FFM  
  neighbor 2001:638:0:500::1 peer-group internalv6  
  neighbor 2001:638:0:500::1 description ERL-to-E
```

```
neighbor 2001:638:0:600::1 peer-group internalv6
neighbor 2001:638:0:600::1 description ERL-to-HH
neighbor 2001:638:0:800::1 peer-group internalv6
neighbor 2001:638:0:800::1 description ERL-to-B
!
address-family ipv4
neighbor internalv6 activate
no neighbor 2001:638:0:300::1 activate
no neighbor 2001:638:0:500::1 activate
no neighbor 2001:638:0:600::1 activate
no neighbor 2001:638:0:800::1 activate
no auto-summary
no synchronization
exit-address-family
!
address-family ipv6
neighbor internalv6 activate
neighbor internalv6 send-community
neighbor 2001:638:0:300::1 peer-group internalv6
neighbor 2001:638:0:500::1 peer-group internalv6
neighbor 2001:638:0:600::1 peer-group internalv6
neighbor 2001:638:0:800::1 peer-group internalv6
network 2001:638:100::/40
network 2001:638:A00::/40
exit-address-family
!
address-family ipv6 multicast
neighbor internalv6 activate
neighbor internalv6 next-hop-self
neighbor internalv6 send-community
neighbor 2001:638:0:300::1 peer-group internalv6
neighbor 2001:638:0:500::1 peer-group internalv6
neighbor 2001:638:0:600::1 peer-group internalv6
neighbor 2001:638:0:800::1 peer-group internalv6
exit-address-family
!
```

Static routes on core router with exit to 6net:

```
ipv6 route 2001:638:0:300::/56 Null0 254
ipv6 route 2001:638:201:10::/64 2001:638:0:300::201:2
ipv6 route 2001:638:201::/48 2001:638:0:300::201:2
ipv6 route 2001:638:202::/48 2001:638:0:300::202:2
ipv6 route 2001:638:204::/48 2001:638:0:300::204:2
ipv6 route 2001:638:208::/48 2001:638:0:300::208:2
ipv6 route 2001:638:200::/40 Null0 254
ipv6 route 2001:638:300::/40 Null0 254
ipv6 route 2001:638::/32 Null0
ipv6 route 2002::/16 Tunnel64
```

Multicast configuration:

```
ipv6 pim rp-address 2001:610:14:5145::145 rpm6net
ipv6 pim rp-address 2001:660:3007:300:1::
ipv6 pim rp-address 2001:700:E000:501::2 rpmcgw
!
ipv6 access-list rpm6net
 permit ipv6 any FF0B::/16
 permit ipv6 any FF1B::/16
 permit ipv6 any FF3B::/16
!
ipv6 access-list rpmcgw
 permit ipv6 any FF3E:30:2001:700:1:FFFF::/96
```

!

13. Appendix D: SWITCH Configuration Examples (Cisco)

This section contains three configuration snippets that were used to add IPv6 to a router that was IPv4-only until its software was upgraded recently. The first two snippets go on the new router, the third one on SWITCH's iBGP (for IPv6) Route Reflectors.

```
!!
ipv6 unicast-routing
ipv6 cef
interface Loopback0
  ipv6 address 2001:620:0:C000::13/128
  ipv6 ospf 1 area 0
ipv6 prefix-list static/connected-to-igp seq 5 permit 2001:620::/48 le 128
ipv6 prefix-list static/connected-to-igp seq 15 permit 3FFE:2000::/48 le 64
ipv6 prefix-list static/connected-to-igp seq 20 permit 2002::/16 le 64
ipv6 prefix-list static/connected-to-igp seq 25 permit ::/0
route-map static/connected-to-igp permit 10
  match ipv6 address prefix-list static/connected-to-igp
ipv6 router ospf 1
  log-adjacency-changes
  redistribute connected metric-type 1 route-map static/connected-to-igp
  redistribute static metric-type 1 route-map static/connected-to-igp
ipv6 ospf name-lookup
interface GigabitEthernet0/1
  description bidir single fibre CWDM GBIC to swiEZ2
  ipv6 address 2001:620:0:C01F::2/64
  ipv6 ospf 1 area 0
  ipv6 ospf network point-to-point
  ipv6 ospf cost 1000
interface GigabitEthernet0/3
  description bidir single fibre CWDM GBIC to swiFF2
  ipv6 address 2001:620:0:C020::1/64
  ipv6 ospf 1 area 0
  ipv6 ospf network point-to-point
  ipv6 ospf cost 1300
end

!!
```

```
router bgp 559
 neighbor RRV6 peer-group
 neighbor RRV6 remote-as 559
 neighbor RRV6 password 7 10780C1729153B3A210C19273C1D23010818061F28
 neighbor RRV6 update-source Loopback0
 neighbor RRV6u peer-group
 neighbor RRV6u remote-as 559
 neighbor RRV6u password 7 10780C1729153B3A210C19273C1D23010818061F28
 neighbor RRV6u update-source Loopback0
 neighbor 2001:620:0:C000::1 peer-group RRV6u
 neighbor 2001:620:0:C000::1 description swi6netCE1
 neighbor 2001:620:0:C000::2 peer-group RRV6u
 neighbor 2001:620:0:C000::2 description swi6T1
 address-family ipv4
 no neighbor RRV6 activate
 neighbor RRV6u activate
 address-family ipv6
 neighbor RRV6 activate
 neighbor RRV6 next-hop-self
 neighbor RRV6u activate
 neighbor RRV6u next-hop-self
 neighbor 2001:620:0:C000::1 peer-group RRV6u
 neighbor 2001:620:0:C000::2 peer-group RRV6u
end

!!
!! Add iBGP route reflector clients
router bgp 559
 neighbor 2001:620:0:C000::5 peer-group CLIENTSv6u
 neighbor 2001:620:0:C000::5 description swiCS4
 neighbor 2001:620:0:C000::9 peer-group CLIENTSv6u
 neighbor 2001:620:0:C000::9 description swiIX1
 neighbor 2001:620:0:C000::A peer-group CLIENTSv6u
 neighbor 2001:620:0:C000::A description swiCE2
 neighbor 2001:620:0:C000::B peer-group CLIENTSv6u
 neighbor 2001:620:0:C000::B description swiEZ2
 neighbor 2001:620:0:C000::C peer-group CLIENTSv6u
 neighbor 2001:620:0:C000::C description swiCE3
 neighbor 2001:620:0:C000::D peer-group CLIENTSv6u
```

```
neighbor 2001:620:0:C000::D description swiEL2
neighbor 2001:620:0:C000::E peer-group CLIENTSv6u
neighbor 2001:620:0:C000::E description swiLS2
neighbor 2001:620:0:C000::F peer-group CLIENTSv6u
neighbor 2001:620:0:C000::F description swiZH2
neighbor 2001:620:0:C000::10 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::10 description swiBE2
neighbor 2001:620:0:C000::11 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::11 description swiBA2
neighbor 2001:620:0:C000::12 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::12 description swiPS2
neighbor 2001:620:0:C000::13 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::13 description swiWI2
neighbor 2001:620:0:C000::14 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::14 description swiFF2
address-family ipv6
neighbor 2001:620:0:C000::5 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::9 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::A peer-group CLIENTSv6u
neighbor 2001:620:0:C000::B peer-group CLIENTSv6u
neighbor 2001:620:0:C000::C peer-group CLIENTSv6u
neighbor 2001:620:0:C000::D peer-group CLIENTSv6u
neighbor 2001:620:0:C000::E peer-group CLIENTSv6u
neighbor 2001:620:0:C000::F peer-group CLIENTSv6u
neighbor 2001:620:0:C000::10 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::11 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::12 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::13 peer-group CLIENTSv6u
neighbor 2001:620:0:C000::14 peer-group CLIENTSv6u
end
```