Project Number:	**IST-2001-32603**
Project Title:	**6NET**
CEC Deliverable Number:	**32603/UOS/DS/2.2.1/A1**
Contractual Date of Delivery to the CEC:	30th June 2002
Actual Date of Delivery to the CEC:	2nd August 2002
Title of Deliverable:	D2.2.1: IPv4 to IPv6 migration scoping report for organisational (NREN) networks.
Work package contributing to Deliverable:	WP2
Type of Deliverable*:	R
Deliverable Security Class**:	PU
Editors:	Tim Chown (University of Southampton)
Contributors:	Ladislav Lhotka (CESNET), Pekka Savola (CSC), Christian Schild (WWU), Rob Evans, Duncan Rogerson, Rina Samani (JANET NOSC and UKERNA), Dimitrios Kalogeras (NTUA), Fotis Karayannis, Chrysostomos Tziouvaras (GRNET)

* 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 may apply in the core IPv6 network, and most importantly to 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.

Keywords:

IPv6 NREN transition, IPv6 ISP transition

Executive Summary

This document is the first deliverable in a series of IPv6 NREN transition deliverables produced as part of the 6NET project. In this document we:

- Describe the available technologies for NREN IPv4-IPv6 transition
- Discuss some candidate transition scenarios, and the applicability of different mechanisms
- Give an overview of participant status and plans

Following on from this scoping report, the project will produce an NREN transition cookbook for M12, M24 and M36 of the project. The guide will offer a blend of theory, practical advice and considerations for NRENs wishing to deploy IPv6 services.

At this stage little deployment experience is included, though we do present a summarised case study of the 6WiN deployment in Germany. We expect the first cookbook report to include initial deployment results and sample configurations.

Table of Contents

Table of Figures

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. Concurrent documents, D2.1.1 and D2.3.1, describe the mechanisms available for core (backbone) and site (university) networks respectively. Together, the three scoping documents provide an overview of the initial work planned within the 6NET project.

6NET has many goals. The deployment of an IPv6-only network core, with PoPs located in a large number of NRENs, is 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 migration and integration of new IPv6 services. A number of NRENs already have some IPv6 deployment, which has been reported in the IPv6 activities of GÉANT's Task Force: Next Generation Networks (TF-NGN) working group, as deliverable 9.3 [D9.3] of the GÉANT project. 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 mechanisms that may be used depend on the nature of the existing IPv4 infrastructure, and the goals of the particular NREN in question. Where ATM or MPLS is already deployed, IPv6 can be deployed on top of those technologies, as described in this document. However, one would not normally consider deploying ATM or MPLS to enable IPv6 deployment. The simplest early transition technique is to deploy an IPv6 infrastructure tunneled over the existing IPv4 network, leveraging the IPv4 network routing topology and performance.

We believe the most common transition will involve running dual-stack IPv4 and IPv6 on the NREN backbone routers, such that those routers hold IPv4 and IPv6 routing tables, and both protocols run natively on the wire. One alternative, in use currently by DFN in the German 6WiN network, is the introduction of a parallel IPv6 infrastructure. The choice between dual-stack and a parallel infrastructure has a number of tradeoffs, in terms of factors such as cost of equipment, performance, and management – we plan to investigate and report on those issues here in WP2.

NRENs do not generally have to operate translation mechanisms; such mechanisms are applied at the edges of the network, at the university border routers. 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).

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 (to December 2004).

2. Review of NREN transition mechanisms

In this section we discuss the broad transition 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).

2.1. Dual-stack

The term "dual-stack" is a broad term, and can be used to mean many things. It's assumed that in a dual stack deployment phase, all routers also support IPv4, but are not necessarily configured as active IPv4 routers. When performing dual-stack transition, 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:

1. Create and operate a test network with tunneled/MPLS/ATM connections to gain perspective on the operation of IPv6.
2. 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.
3. If they are, and other possible concerns (e.g. whether it seems there is sufficient demand for IPv6 in the main network) are met, 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.
4. If problems (e.g. severe bugs affecting production services) arise, either try to fix or avoid them or degrade back to an IPv4-only operation.

Most NRENs following this path are currently in steps 1 or 2.

Some router vendors have already, some time ago, shipped IPv6 in their production software, and it has gained a reasonable amount of stability. If other concerns (e.g. is the use of or demand for IPv6 sufficient to warrant it) are met, in some networks the gradual transition to dual-stack could be achieved in a relatively short time. In order to break the “chicken and egg” status, it would be preferable to deploy IPv6 in advance of heavy demand, as there will not be demand until the service is well supported by the NRENs.

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 might 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 The timing is important: going dual-stack too early may decrease the overall robustness, but going too late may increase the complexity of the separate IPv6 network.

2.2. General tunnels

There are two classes of general tunneling 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 – that infrastructure is already well-tuned by the NREN to perform well; thus even with the tunneling 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 underneath be many IPv4 hops.

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 Ethernet networks, on Fast Ethernet interfaces, the MTU is 1500. Using tunneling 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.

2.3. IPv6 over MPLS

Backbone networks that have already deployed Multi Protocol Label Switching (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 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 [Clercq02]. 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)*.

From the above approaches, only the latter one is of immediate interest to WP2. 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. 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.

2.3.1. 6PE – theory of operation

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

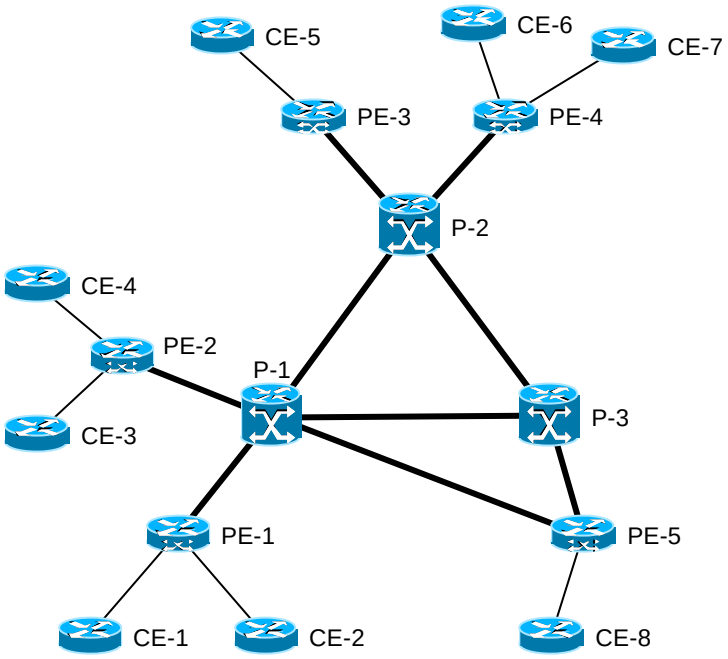

Figure 1: MPLS routing hierarchy

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 label[1] 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 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 infomation into IBGP, this time under `IP6` 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 `IP6` 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).

An example is shown in Figure 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.

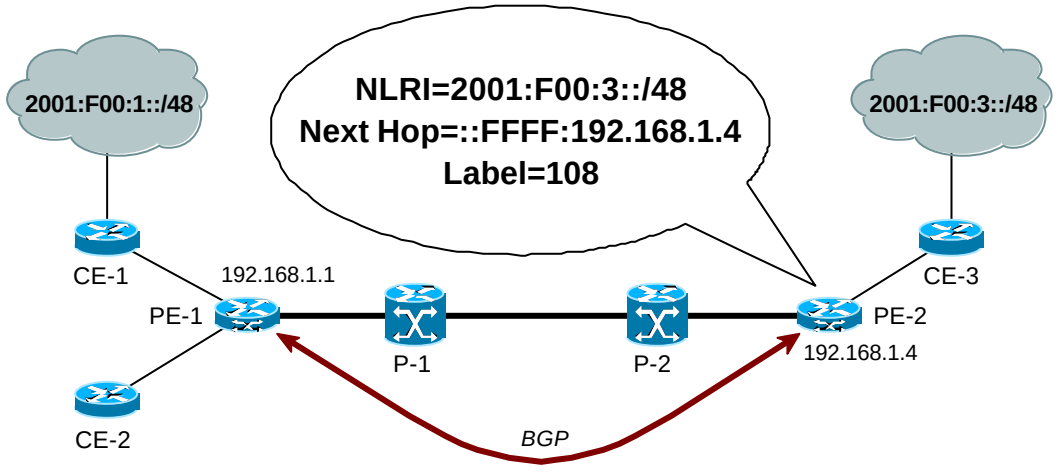

Figure 2: BGP and label advertisement of a 6PE router

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.

Outer label = 55	Inner label = 108	IPv6 datagram

Figure 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.

2.3.2. 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.

2.3.2.1 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 [RFC2529], except that the latter needs special IPv6 addresses. However, 6PE can also use MPLS tunnels (LSPs) configured by other means, e.g., traffic engineering.

2.3.2.2 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 4.

4 bytes	4 bytes	
Outer MPLS shim header	Inner MPLS shim header	IPv6 datagram

Figure 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 5.

Figure 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 [RFC2529]. Figure 6 indicates that the overhead is just the IPv4 header of 20 bytes (again assuming no IP header options).

Figure 6: Structure of an IPv6 packet encapsulated in IPv4

Table 1 shows overhead percentages for different lengths of the tunneled 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 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 7.

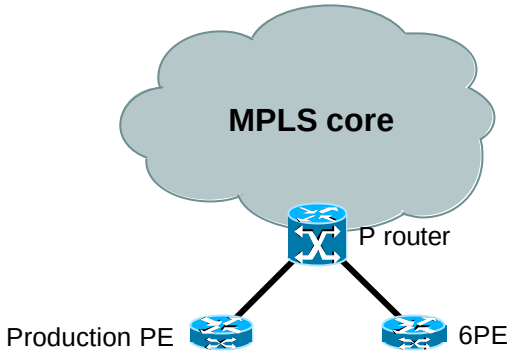

Figure 7: Deployment of 6PE in a production MPLS network

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.

2.4. 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. Later in this section we show an example of IPv6 over ATM over MPLS.

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 very well be used in the early phase, usually the first couple of years when it is not yet feasible to enable IPv6 in all routers. 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.

2.4.1. 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 PTT and ISP networks have been implemented

using ATM infrastructure. 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. It is common that site topologies may use a LAN Emulation (LANE) environment. For those sites, where hosts are equipped with a dual stack OS, transition to IPv6 requires a native IPv6 over ATM connection [RFC2492]. The later alternative to PVCs is soft PVCs. Soft PVCs provide the dynamic nature of SVCs while preserving the simplicity of a PVC to end stations. End stations are aware of the terminating PVC characteristics while the ATM network dynamically routes ATM cells to appropriate destinations, thus end hosts are not aware of any ATM addresses. This simplifies the set up of ATM connections.

2.4.2. 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 %. From the operational view a tunnel setup requires similar effort to the 6to4 tunnels.

2.4.3. 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 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.

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.

2.5. IPv6 over ATM over MPLS

MPLS is a natural evolution of ATM networks mainly due to the integrated approach of traffic engineering and topology based routing protocols. MPLS, by means of 6PE, can provide a transition scenario for NRENs. It is possible that 6PE is not available or that the 6PE solution does not meet the stability requirements of an NREN, which would like to deploy one IPv6-only router, where customers can terminate all their IPv6 traffic. A truly smooth transition plan allows for access devices which use the services of the backbone to remain untouched. This section describes a service allowing delivery of non-IPv4 traffic over ATM with a MPLS incapable of native IPv6 capabilities.

2.5.1. Description

A typical scenario of a MPLS backbone with customers having ATM attached routers capable of IPv6 is shown in Figure 8.

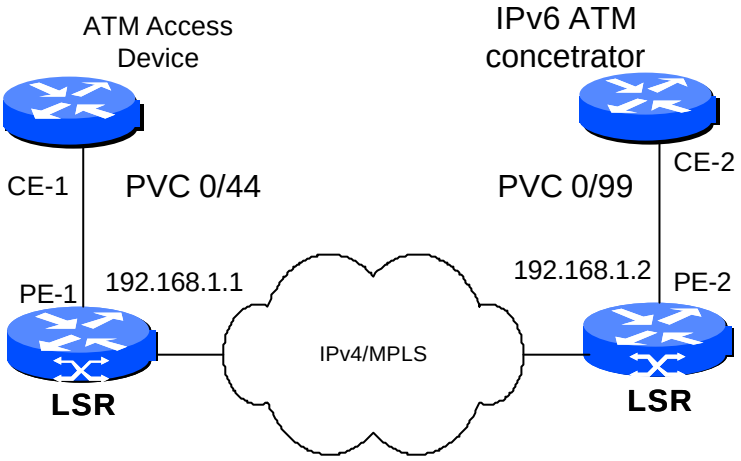

Figure 8: A Typical scenario IPv6 over ATM over MPLS

The ATM Access Device and the NREN IPv6 concentrator operate as a CE devices in the MPLS terminology. In this MPLS network, it is possible to carry the Protocol Data Units (PDUs) of layer 2 protocols, like ATM cells, by prepending an MPLS label stack to these PDUs [Martini01].

2.5.2. Operation

Suppose it is desired to transport ATM cells from ingress LSR PE-1 to egress LSR PE-2, across an intervening MPLS network. It is assumed that there is an LSP from PE-1 to PE-2. In that sense PE-1 can cause a packet to be delivered to PE-2 by pushing some label onto the packet and sending the result to one of its adjacencies. Call this label the "tunnel label", and the corresponding LSP the "tunnel LSP". The tunnel LSP merely gets packets from PE-1 to PE-2, the corresponding label doesn't tell PE-2 what to do with the payload, and in fact if penultimate hop popping is used, PE-2 may never even see the corresponding label. (If PE-1 itself is the penultimate hop, a tunnel label may not even get pushed on.) Thus if the payload is a non-IPv4 packet, there must be a label, which becomes visible to PE-2, that tells PE-2 how to treat the received packet. Call this label the "VC label". So when PE-1 sends a layer 2 ATM cell to PE-2, it first pushes a VC label on its label stack, and then (if PE-1 is not adjacent to PE-2) pushes on a tunnel label. The tunnel label gets the

MPLS packet from PE-1 to PE-2; the VC label is not visible until the MPLS packet reaches PE-2. PE-2's disposition of the packet is based on the VC label.

Note that the VC label must always be at the bottom of the label stack, and the tunnel label, if present, must be immediately above the VC label. Of course, as the packet is transported across the MPLS network, additional labels may be pushed on (and then popped off) as needed. Even PE-1 itself may push on additional labels above the tunnel label. This operation is depicted in Figure 9.

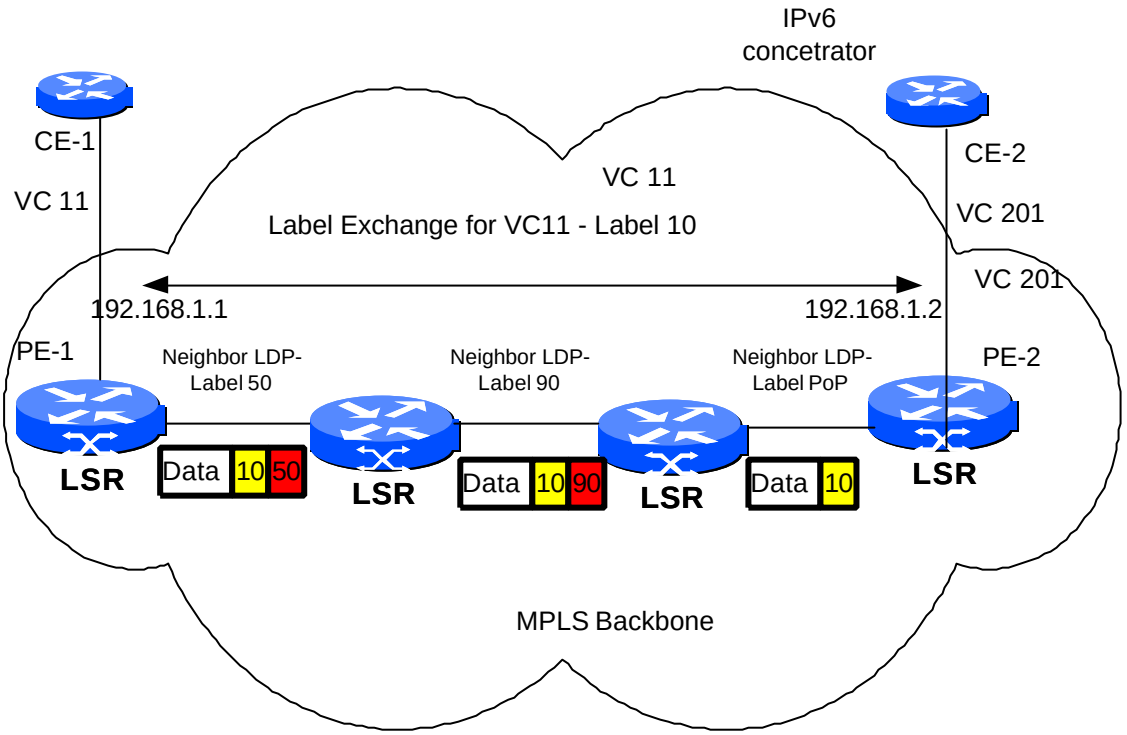

Figure 9: Label advertisement of ATM encapsulation

An additional improvement of the aforementioned operation becomes available, considering that the IP and IPv6 packets are encapsulated in AAL5 headers [Martini03]. It is then possible for the ingress LSR to perform reassembly of the AAL5 payload (AAL5-CPCS) and transmit it in an MPLS payload. The egress LSR must regenerate the AAL5 trailer and PAD to transmit it to CE-2.

The operation of the ordinary IGP and EGP routing protocols using this transition scheme, is similar to that of IPv6 over ATM.

2.5.3. Comparison to other technologies

The IPv6 over ATM over MPLS differs from other IPv6 tunneling techniques in two main aspects: tunnel set up and tunnel overhead.

2.5.3.1 Tunnel set up

The IPv6 over ATM over MPLS like 6PE relies on MPLS to create tunnels. The tunnels are like virtual interfaces with the property that traffic between the end points is transported using mpls encapsulation and travels over the same LSP(s) that would be used for IP packets between the end points. Tunnel creation in that sense follows the IGP. It is also possible for a tunnel to follow a traffic engineered tunnel, built using the RSVP-TE mechanism.

2.5.3.2 Tunnel overhead

IPv6 over ATM over MPLS provides the efficiency of a unified backbone at the cost of extra MPLS and cell headers required for proper forwarding of ATM cells on top of MPLS layer. Typically the overhead of MPLS encapsulation is two mpls shim headers, which count 8 bytes in total. Variable overhead results in the case of cell relay or AAL5 transport service. In the former case the overhead is 12 bytes for every ATM cell while in the later the overhead is 12 bytes for the AAL5 payload. The ATM overhead is 5 bytes plus a variable overhead varying with the packet size and depending on the AAL5 header adaptation. If the overall overhead should be evaluated the reader should consider the ATM overhead which is approximately 15% for the AAL5 VC-MUX encapsulation [RFC1483] and 22 % for the AAL-SNAP encapsulation.

It appears that the tunnel overhead can be affordable only in the case of ALL5 service and for large datagram sized packets.

2.6. Deploying a parallel IPv6-only network

Some have argued that a new, separate IPv6 network may also be an option. At least in some environments, where there are tight service-level agreements in IPv4 or there would have to be a great deal of IPv6 education for IPv4 operators, some feel this is the easiest and cheapest way: build everything from scratch, and keep the networks separate.

However, given there is a cost to having two sets of routers, and two sets of links, a more detailed cost analysis would be necessary to see whether this makes financially 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). 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. serial lines would probably be out of the question.

This scenario is discussed in a later section considering the example of the DFN 6WiN network.

2.7. Support mechanisms

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, and 6to4 (with 6to4 relay).

2.7.1. Tunnel broker

The tunnel broker (described in Deliverable D2.3.1 in section 1.2.7), allows a user on a dual-stack host to connect to a web server to download a script that can be run to then establish an IPv6 connection (an IPv6-in-IPv4 tunnel) to the tunnel broker's tunnel server. 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).

2.7.2. 6to4 and 6to4 relay

The 6to4 mechanism has been described in Deliverable D2.3.1 in Section 1.2.3, but we present a slightly different description here.

The 6to4 transition mechanism [RFC 3056] 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 backbone exists.

2.7.2.1 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.

Construction of the 6to4 IPv6 prefix is made by the concatenation of the 2002::/16 prefix and the hexadecimal representation of a single globally routable IPv4 address that belongs to the organization. The format of the 6to4 IPv6 address is illustrated in Figure 10.

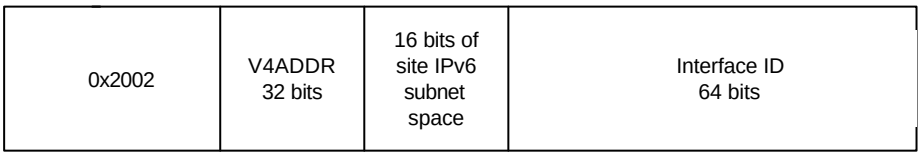

Figure 10: 6to4 address format

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, since 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 sites.

A typical example that demonstrates the usage of the 6to4 mechanism is shown in Figure 11. Also in the same figure the configuration that is needed for the two routers (Cisco-like) is shown.

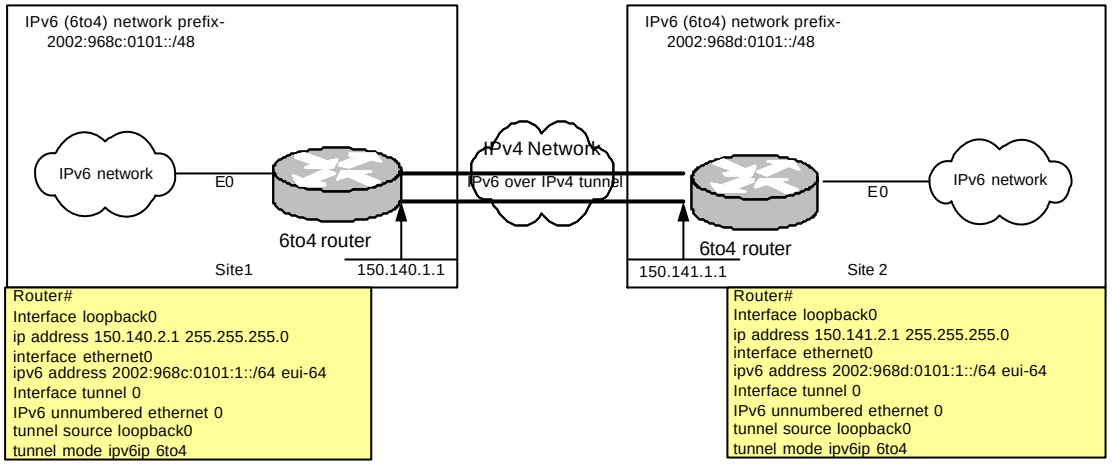

Figure 11: Typical example of the usage of the 6to4 mechanism

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 6to4 relay router can be deployed on an IPv4 anycast address, allowing multiple relay routers to be available.

2.7.2.2 Conclusions

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 NREN should be careful in checking who can use the relay router.

3. Current NREN transition mechanism deployment status

3.1. Partner status and future plans

	MM	Dual stack	General tunnel	IPv6 over MPLS	IPv6 over ATM	Parallel IPv6 network	General support mechanisms
CISCO	2	-	-	-	-	-	-
IBM	2	-	-	-	-	-	-
RENATER	2	-	-	-	D	-	-
UKERNA	4	Y	D	-	D	-	Y
DFN	1	-	-	-	-	-	-
ACONET	2	-	-	-	-	-	-
GRNET	4	D	D	Y	Y	Y	D
INFN-GARR	3	-	-	-	-	-	-
UoS	5	-	-	-	D	-	D
CSC (FUNET)	2	-	-	-	-	-	-
UNI OULU	2	-	-	-	-	-	-
WWU (JOIN)	6	-	-	-	-	D	D
CESNET	3	-	D	Y	-	-	-
UNINETT	0	Y	D	-	D	Y	-

Table 2: Planned NREN transition studies by project partners

The entries in Table 2 have the following meaning:

D = NREN has (or had) some deployment already

Y = NREN has plans to deploy the mechanism

- = NREN may deploy, but has no plans at this stage

This Table is currently only lightly committed to. The NREN partners will firm up their plans in time for the M12 deliverable (the NREN transition cookbook). There are seven NRENs represented in this activity; the other partners are undertaking supporting work, assisting with trials, developing code or doing theoretical studies.

4. Scenarios for NREN Transition

There are a number of different scenarios that an NREN that wants to migrate to support IPv6 may find itself in:

- It will have an IPv4 service, and may see dual stack operation on the NREN core routers as a natural path forward, given appropriate robust code and appropriate studies into the mangement and operational implications.
- It may have an MPLS network. There are some specific options for running IPv6 over MPLS (this is a scenario that has been used by Cisco).
- It may have an ATM network, and thus be able to run IPv6 over parallel PVCs.

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 tunneled traffic in the IPv6-only national backbone.

In the following sections, we describe two case studies of early IPv6 deployment undertaken to date in NREN networks. Further examples are described in [D9.3].

4.1. 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.

4.1.1. Internal connectivity and setup

The core of 6WiN is a set of five (resp. six) 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.

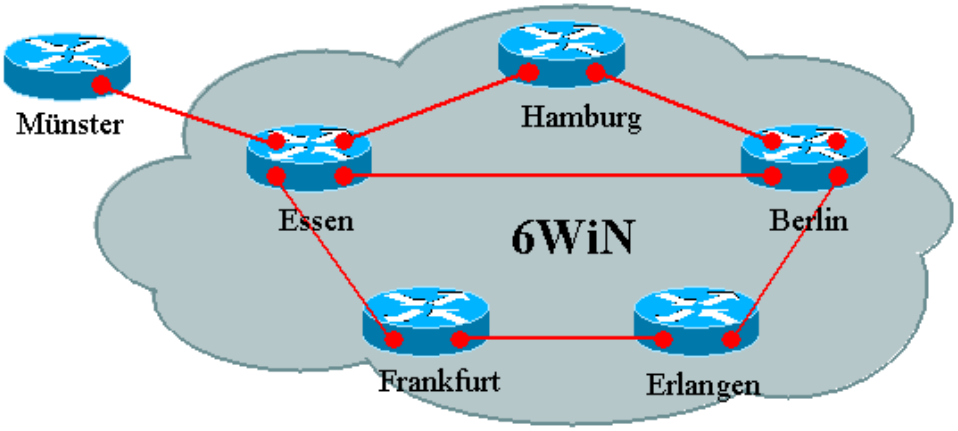

Figure 12: The 6WiN network, including Münster

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 twelve facilities are connected to 6WiN, two of which are natively connected.

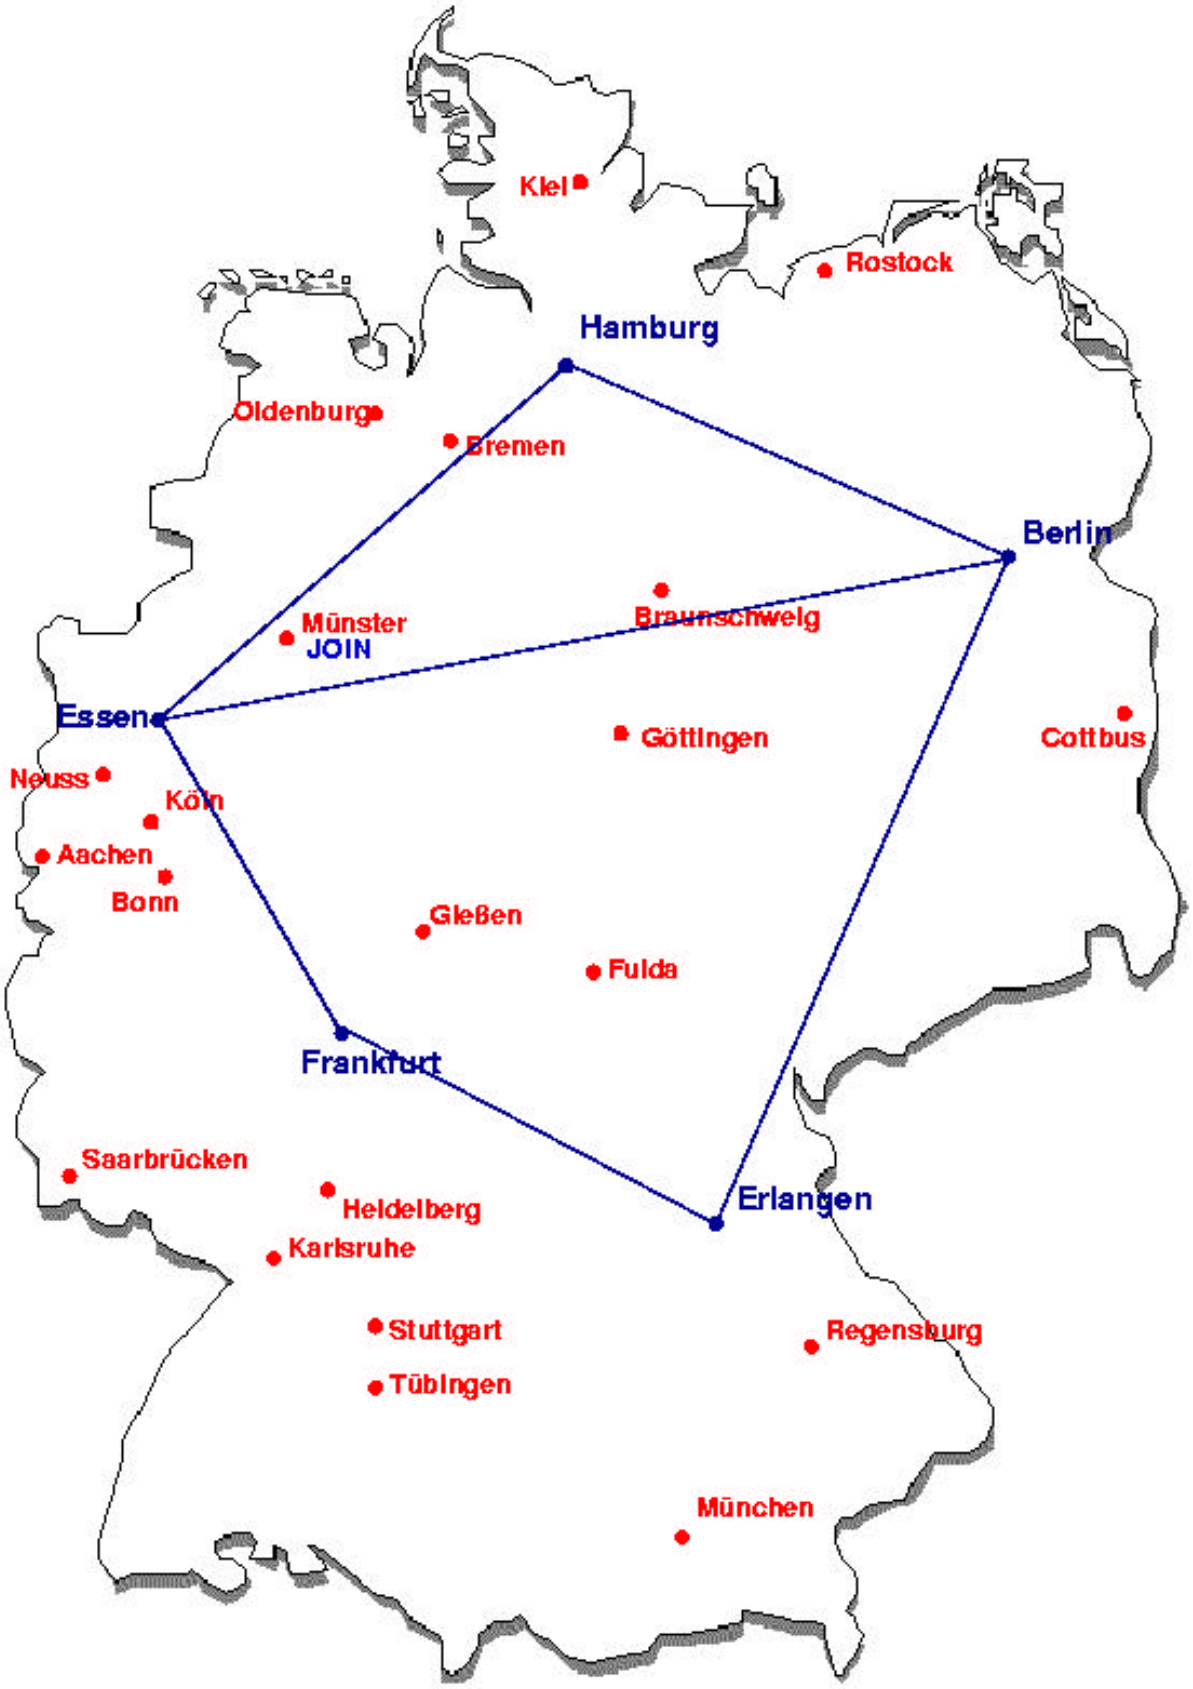

Figure 13: The full 6WiN scope across Germany

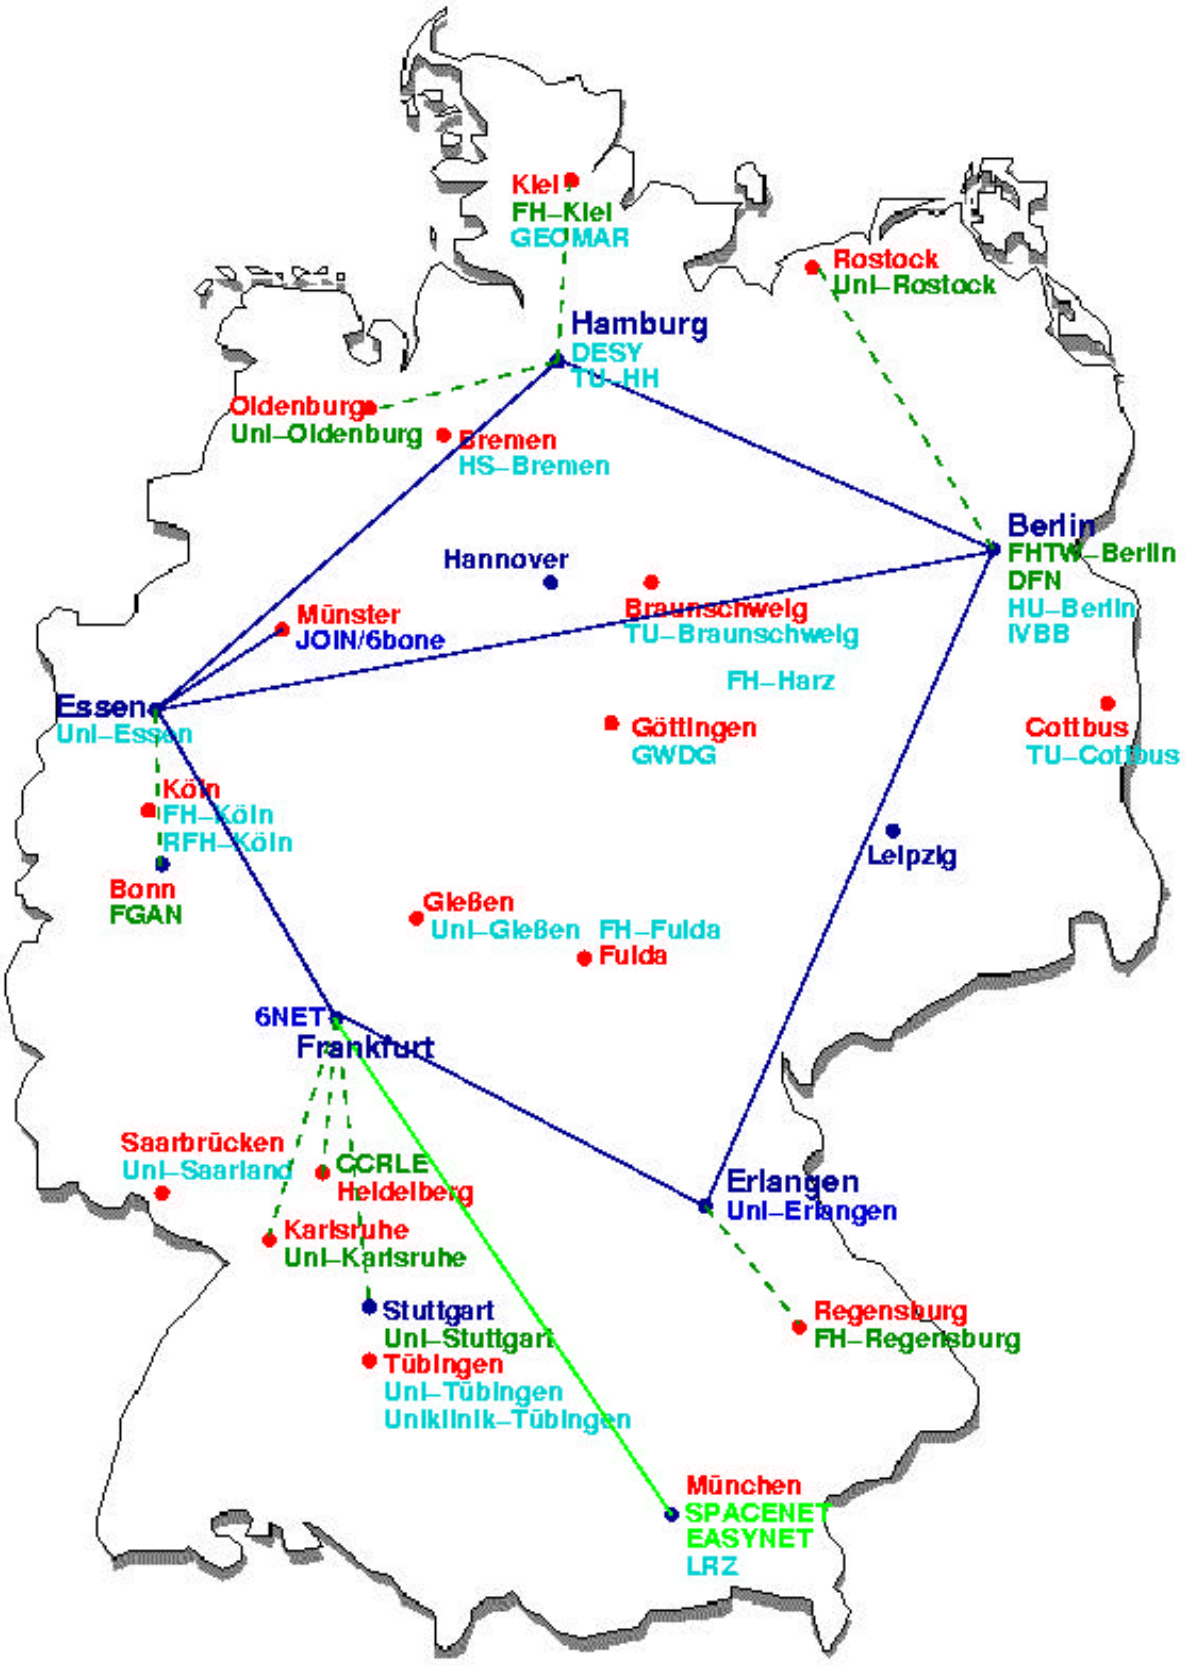

Figure 14: 6WiN and sites in 6WiN

4.1.2. External connectivity

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

- 6bone: The JOIN Project im Münster still hosts a large 6bone backbone node, so there is a connection to the 6bone in Münster.
- Deutsche Telekom: Deutsche Telekom has a similar test network like DFN. These two networks - 6WiN and DT - are connected in Münster and Berlin. At this point in time only the connection in Berlin is up and running. As the network of Deutsche Telekom is only a temporary project these connections may vanish later.
- German industry: There are a few ISPs in Germany who already offer IPv6 services. The 6WiN is connected to two of them - SpaceNet and EasyNet – over IPv6-in-IPv4 tunnels. Again, only those networks, that are topologically close in the IPv4 network, are allowed peering partners. In this case the ISPs have to have a direct peering contract with DFN.
- 6NET: Of course 6WiN is connected to the 6NET in Frankfurt.
- DECIXv6: DECIX offers an IPv6 IX in Frankfurt in the same building as the 6WiN PoP. It is intended to eventually connect to that IX, but it is not done yet. If 6WiN gets connected to DECIXv6, the tunnels to SpaceNet and EasyNet will become obsolete and will be shut down.

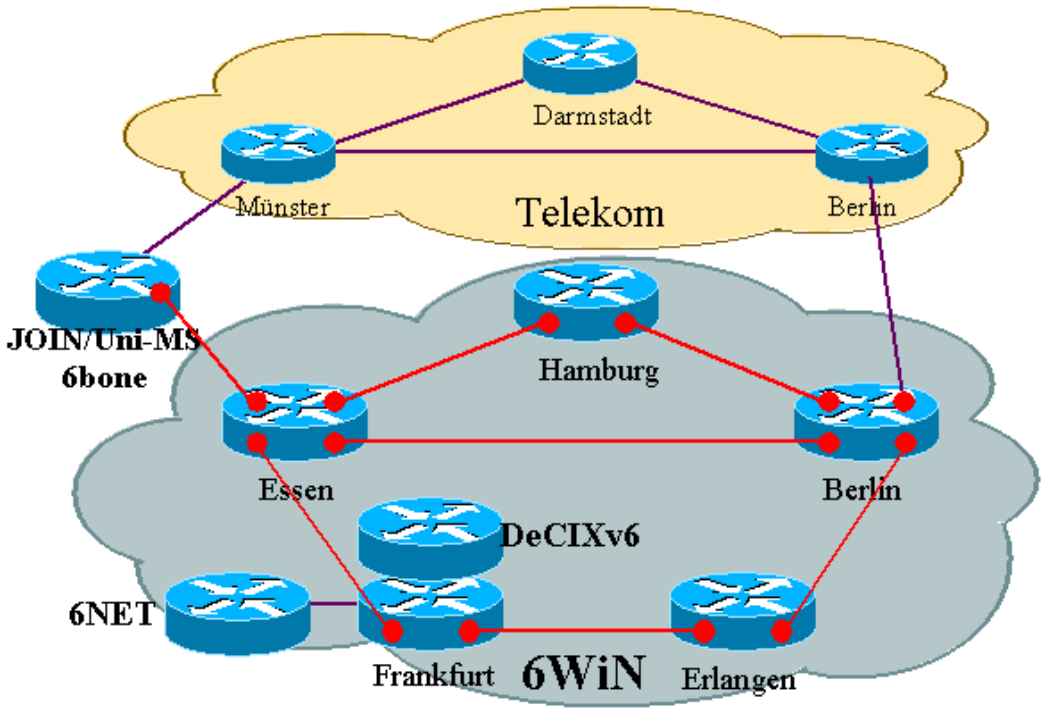

Figure 15: 6WiN connections

4.1.3. Addressing

DFN's 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.

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.

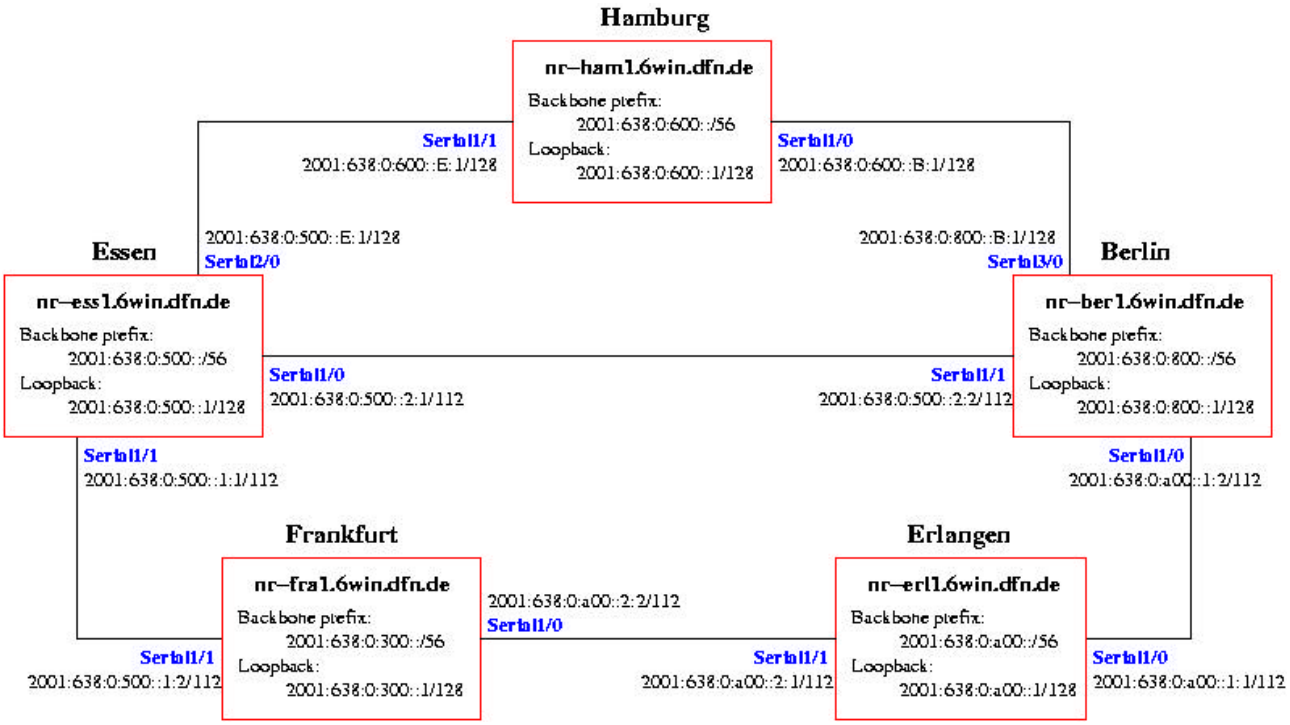

Figure 16: 6WiN addressing overview

4.1.4. 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.

4.1.5. 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 or multicast, when these protocols become available for IPv6, or for different addressing schemes.

4.2. 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).

This pilot network will run until the end of December 2002, with a review in the November to plan subsequent phases, which will include the schedule for migration to dual-stack services on JANET's core routers. 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.

4.2.1. 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 will endeavour 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). This will be advertised to the community once sufficient documentation has been prepared.
- *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.

4.2.2. Future services

Services that will be introduced within the coming months 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 is planned, to allow 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 plan 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.

4.2.3. 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.

4.2.4. 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.

4.2.5. Futures

UKERNA think it unlikely that a full and complete migration from an IPv4 based network to a pure IPv6 based network will happen for the foreseeable future. However, the feasibility of operating a dual-stack production core network, including the issues of interim connectivity between IPv4 and IPv6 only hosts, is being investigated. In the meantime, the Experimental Service will continue to provide IPv6 connectivity to interested JANET connected organisations.

UKERNA will also establish a JANET IPv6 Working group, to assist in the further deployment of IPv6 in the JANET community.

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

GRNET has run an internal IPv6 pilot project investigating IPv6 migration and IPv6 applications since 2000. GRNET has participated in the 6bone testbed and has been allocated the pTLA 3FFE:2D00::/24 which is being used 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 the 6bone

Initially GRNET applied and was granted a pTLA for experimenting with IPv6 networking in the 6bone and was allocated 3FFE:2D00::/24. Later GRNET was allocated a SubTLA (sTLA) from RIPE, which is 2001:648::/35 (which under the new global policy adopted in July 2002 can now grow to a /32 prefix). GRNET is now planning to migrate the IPv6 addressing in its region to use the RIPE production address space.

GRNET is connected to 6NET through a L2 VPN (CCC Juniper) to 6NET, although this is not active at the time of writing. GRNET plans to natively connect to 6NET next year (in 2003).

4.3.1. Services in production

GRNET currently offers the following IPv6 services:

- Connection to the 6bone by means of static tunnels.
- Allocation of a range of addresses via the pTLA that it holds.
- 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 tunnel edges.

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+.

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, but this could be considered.

4.3.2. 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 next figure) 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.
- *Transition*: a NAT/PT server will be set up so as the IPv4 clients to have access to some basic IPv6 services (and vice versa). This is not intended as an NREN service, but to gain experience such that GRNET can give advice to universities setting up such a service.

5. Conclusions

In this document we have described IPv6 transition tools available to NRENs within the European academic and research community. The choice of which tools to adopt is a matter for each NREN to decide, but the tools should complement the requirements of the sites (universities) and the core network (in the context of the project, 6NET, and in the context of European networking, GÉANT), as described in deliverables D2.3.1 and D2.1.1 respectively.

It would be expected that unless technology is already in place that allows a specific solution, e.g. MPLS or ATM, that NRENs will start with an initial phase of a pilot tunneled IPv6 service offered to universities from a small core IPv6 router presence. Such a deployment could include 6to4 (and 6to4 relay) services, and a tunnel broker. As and when native links can be made available, they would be deployed and utilised. A second phase would then see migration to dual stack operation in the NREN backbone network. We would expect that such a mode of operation would continue for many years (ten or more) until the core networks run IPv6 only, and carry IPv4 tunneled in IPv6 (although radical advances in optical technology may alter such plans). NRENs would most likely leave translation techniques (where used at the egdes of IPv6-only networks) as a responsibility for the university sites (see Deliverable D2.3.1).

NREN IPv6 service offerings beyond basic connectivity to universities include the support services like 6to4 and tunnel brokers, but also other services are required, e.g. IPv6 access and support in the NREN's national DNS (e.g. .ac.uk in the UK). We will report on results of experiments, with configuration examples and case studies in the first (M12, December 2002) NREN transition cookbook (Deliverable D2.2.2).

6. References

[Bermuda] The Bermuda 2 IPv6 Project, http://www.ipv6.ac.uk/bermuda2/.

[Clercq02] De Clercq J., Gastaud D. Nguyen T., Ooms D., Prevost S. and Le Faucheur F. "Connecting IPv6 Islands across IPv4 Clouds with BGP". IETF, work in progress, 2002.

[D9.3] GÉANT Deliverable D9.3: IPv6 Testing, http://www.dante.net/tf-ngn/D9.3.pdf

[Martini01] "Transport of Layer 2 Frames Over MPLS", L. Martini et al., IETF Internet Draft, April 2002.

[Martini02] "Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks", L. Martini et al. , IETF Internet Draft, November 2001

[Martini03] "Encapsulation Methods for Transport of ATM Cells/Frame Over IP and MPLS Networks", L. Martini et al., IETF Internet Draft, June 2002.

[RFC1483] "Multiprotocol Encapsulation over ATM Adaptation Layer 5", J. Heinanenm, 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. Schulter, M. Jork, IETF RFC, January 1999.

[RFC2529] "Transmission of IPv6 over IPv4 domains without Explicit Tunnels", B. Carpenter, C. Jung, IETF RFC, 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.

[RFC3107] "Carrying Label Information in BGP-4", Y. Rekhter, E. Rosen, IETF RFC, 2001.