An IPv6 Deployment Guide

Editor: Martin Dunmore

An IPv6 Deployment Guide

6NET was a three-year European IST project to demonstrate that continued growth of the Internet can be met using new IPv6 technology. The project built and operated a pan-European native IPv6 network connecting sixteen countries in order to gain experience of IPv6 deployment and the migration from existing IPv4-based networks.

6NET involved thirty-five partners from the commercial, research and academic sectors and represented a total investment of €18 million; €7 million of which came from the project partners themselves, and €11 million from the Information Society Technologies Programme of the European Commission. The project commenced on 1st January 2002 and officially finished on 30th June 2005. The network itself was decommissioned in January 2005, handing over the reigns of pan-European native IPv6 connectivity to GÉANT.

When we began 6NET, IPv6 code was in the form of early beta releases from most commercial companies. The 6BONE had been built but was only using tunnels; there were very few native IPv6 networks and none of these ran production traffic. One thing we strived for in the early days of 6NET was developing a pan-European testbed that had as much native IPv6 connectivity as was affordable. This gave everyone involved the chance to really exercise the IPv6 protocol developments we planned without the added complexity of tunnels that might detract from the real work. Soon we were able to peer with other IPv6 networks in the US (Abilene, 6TAP), Japan (NTT) and S.Korea (KOREN) to provide global IPv6 connectivity. The final stages of the project moved into exploiting the protocol and providing demonstrations that IPv6 was ready for full production service.

The information contained in this book is taken from the project's deployment cookbooks and other deliverables. Since each cookbook/deliverable generally concentrates only on specific IPv6 features or deployment scenarios (e.g. site transition, multicast, mobility, DHCP, routing etc.), we believe that providing all the important information in a single reference book is much more preferable to the reader than negotiating our multitude of project deliverables.

PART II

Case Studies

Chapter 13

IPv6 in the Campus/Enterprise

In this chapter we present case studies of IPv6 in the Campus/Enterprise. Since the vast majority of the 6NET partners were academic related, this case studies in this chapter are indeed related to Universities and academic departments. Nevertheless, there are many similarities between University/Campus based deployments and Enterprise deployments.

First, we we look at the Campus IPv6 deployment at the University of Münster. Next we describe two deployments at small and large academic departments (Tromsø and Southampton University respectively). A second University Campus deployment case study is given for Lancaster University and finally, we briefly describe some other deployment scenarios relating to the Campus/Enterprise class of network.

13.1 Campus IPv6 Deployment (University of Münster, Germany)

As the University of Münster is quite large, with a widespread network and is using at large set of different hardware and network techniques, several considerations had to be taken into account.

If one wants to integrate IPv6 in the network, the most desirable form of integration is always to run in dual-stack mode on each and every interface and node. However, while nowadays support for IPv6 is present in nearly every new product, there are still older hardware and technologies that do not easily support IPv6 capabilities or don't support them at all.

Especially in large sites, that have been in place for a long time, the network infrastructure has evolved over a number of years. Such networks often have a modern core, but still use old technology in some areas and on internal "stubby" edges. In such environments it is practically impossible to run full dual-stack mode. Several of the transition methods described in this cookbook can be used to reach such areas.

In addition, network administrators often hesitate to introduce IPv6, because they fear that they will destabilise their IPv4 infrastructure or because they are unfamiliar with IPv6 and IPv6 management. To overcome these fears it is helpful to start with IPv6 just in a few parts of the network and to leave the IPv4 infrastructure untouched.

A good method for this is using VLAN technology (802.1q). VLANs are very common and often used in modern networks, and it is especially easy to integrate IPv6 in these networks. If a dedicated IPv6 router is used, it can get access to only those VLANs where IPv6 is desired. So the IPv4 network remains unchanged, and all IPv6 traffic is routed and managed over a different set of hardware.. If no additional hardware is available, it might be sufficient to use only a small set of the existing routers to do IPv6 routing.

Since VLANs are spread throughout the whole University, it is possible to give IPv6 access to various areas. Still, there are some drawbacks. In those areas where no VLAN technology is available, but older remnants of e.g. ATM or FDDI infrastructure exist, other methods are needed to give IPv6 access to the hosts. This can be achieved with various tunnel technologies or a tunnel broker. Also, if there is a "secured" area, one should consider carefully if IPv6-access should be added to such a VLAN, because when bypassing the IPv4 infrastructure those security mechanisms might get bypassed. For example if there are ACL rules in use for IPv4, these should be applied also for IPv6. This is not always possible, because the two routing topologies are different. If IPv4 ACL rules rely on some kind of hierarchical routing infrastructure, they probably cannot be rebuilt directly for IPv6 in this case, or at least not easily so.

13.1.1 IPv4

In the University of Münster the "Centre for Information Processing", or "Zentrum für Informationsverarbeitung" (ZIV), has a key role as the Network Operation Centre (NOC) not only for the University itself, but also other facilities, e.g. the medical hospital and other educational institutes in Münster. That said, while ZIV is part of the university, it operates the network for several parties and is a kind of provider to multiple users. Therefore the network is not just a core network, but is more complex as it has several core networks and a transit network between them. Some smaller core networks connected to the transit network are not operated by the ZIV-NOC.

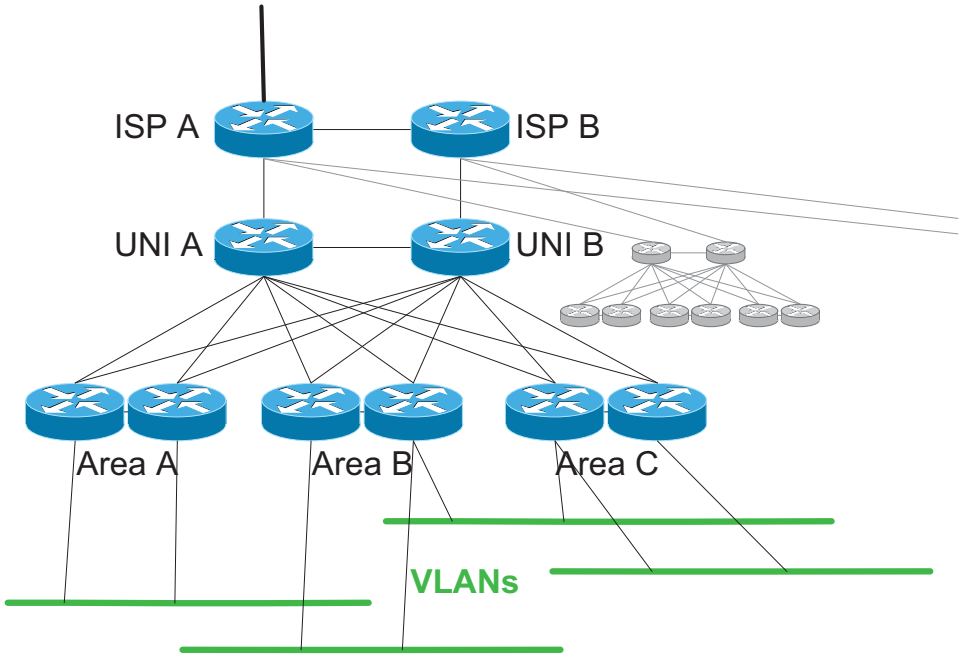

Figure 13-1 Ideal Overview of University's IPv4 Network

Figure 13-1 gives a rough overview of the current IPv4 network. All the routers shown are Cisco Catalyst 6509's with different kinds of engines (SUP2 or SUP720). Usually the connections shown are aggregated Gigabit Ethernet or 10Gigabit Ethernet connections. Redundancy is a key part of the whole

network infrastructure and you will find redundant routers and structures in every core part of the network.

At the top of Figure 13-1 you see the transit routers ISPA and ISPB. One of them is connected to our upstream provider DFN. The network infrastructure below these two routers is the idealised form of the core network of the University. UNIA and UNIB act as border routers with redundant connections to the transit routers and likewise as the core routers of the internal network.

The University of Münster doesn't have a centralised campus but its buildings are spread out throughout the whole city centre. Thus the network is more like a metropolitan area network (MAN) than just a simple LAN. Several areas are equipped with a set of access switches that connect to the central core routers UNIA and UNIB. Reliability is provided by running the HSRP protocol between the routers.

VLANs are not necessarily limited to a single area but may span several access router located elsewhere.

Next to the core network of the university the ZIV manages a similar core network for the clinical facilities which is organised likewise (see the right side of Figure 13-1). Other sites connected to the transit routers are the Fachhochschule Münster and soon the Max-Planck-Institute. Other site networks may follow later on.

The scenario described here is an ideal one. As the University has quite a lot of areas, it is not affordable to allocate two high end routers to every location. To realise the described infrastructure we use virtualised router technology. (The reason for this extended virtualization is to create a hierarchical routing structure. This allows a strict correlation between a double set of routers and an area/site, which allows defining a security infrastructure and limits access based on the fact where a node is located in the network).

In addition, not every building is connected via this newer network infrastructure. There still exist remnants of old technologies used in the past, namely an ATM network and even some last parts of an FDDI ring. Some far off site buildings use ISDN or DSL to connect to the University's network.

Of course the management of this network is quite sophisticated. The most important components are Tivoli Netview management, an Oracle Database which contains every router, host, cable, number and data that is used for this network, and a web-based front-end system to update information in this database easily.

13.1.2 IPv6

The IPv6 infrastructure we started with took advantage of the fact that the University's core network uses VLAN technology extensively. This way we were able to add an IPv6 router independent from the rest of the IPv4 infrastructure. Therefore no harm could be done to running IPv4, be it the danger of higher resource usage (like CPU power) or caused by the necessity to use experimental IOS versions.

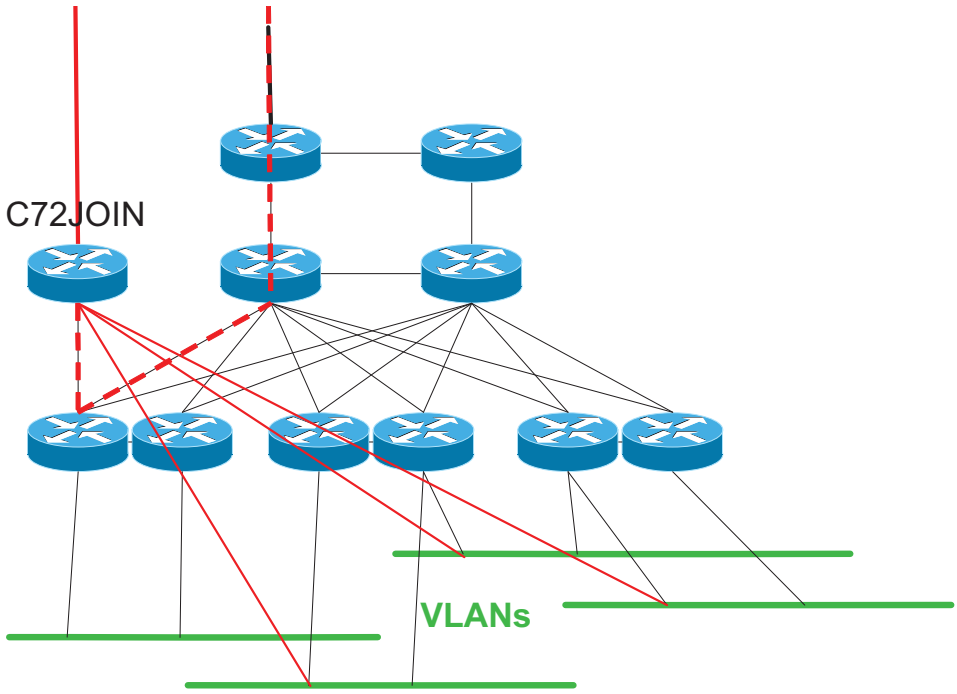

Figure 13-2 IPv6 Test Network

So we placed a dedicated Cisco 7206 router next the IPv4 infrastructure (see Figure 13-2). It is connected to the upstream network 6WiN via an IPv6-in-IPv4-Tunnel. It gets access to every VLAN we want to provide IPv6 services to. This IPv6 traffic is routed through the dedicated router and IPv4 uses the production network.

Management of the IPv6 network is limited. However, as it mainly consists of a single router, it is not so complex to monitor. We use mainly Argus and MRTG to obtain and visualize the network status.

Only a few VLANs are integrated into the IPv6 network. To give the user the opportunity to use IPv6, we offer most of the key services (DNS, Web, FTP, SMTP, NTP, NNTP, jabber, etc.) with dedicated servers we setup over time during the projects lifetime.

13.1.3 IPv6 Pilot

While being simple and effective, this method to distribute IPv6 doesn't match the requirements and features of the neighbouring IPv4 network at all. There is no redundancy and in comparison it is hardly managed.

A good start, to try to deploy an IPv6 pilot network, is to extend the IPv6 network to the current IPv4-only nodes and to run the network in dual-stack mode. With the goal to use this pilot for production usage one day, it is not only necessary to deploy IPv6 as dual-stack, but also to integrate IPv6 into all management procedures and to require that all concepts that are current practice in IPv4 will also be applied for IPv6. This means especially the requirement for the demand for redundancy, reliability and security.

Looking at the well-established procedures and methods used for the IPv4 network this is not so easy to accomplish and will require a lot of work, adaptation and education of NOC personnel.

13.1.3.1 Goal

The overall goal is to evolve the centralised single router IPv6 test network to a well managed IPv6 pilot that can be renamed to a highly reliable production network later on. The best way to achieve this is to establish a dual-stack network that uses the currently existing IPv4 network infrastructure for both IPv4 and IPv6 likewise.

The brief sketch in the chapter above might give you a hint the University's network is quite complex in its IPv4 structures and tries to serve complex needs. Taking into account that there is also a lot of old hardware in use, one can imagine that it is illusionary to add dual-stack to the whole network in one step.

13.1.3.2 Chosen hardware

We decided for a slow-start behaviour when introducing IPv6 into the production network. So in the beginning IPv6 will not be activated on every core router but only the most important ones, namely the transit routers. As some of the routers are only virtual and the virtual technology is quite new in the University's network, we will try to refrain from using these routes.

Still, we want to offer IPv6 in every VLAN that wants to use it. Which VLAN is allowed to receive IPv6 traffic is limited by the level of security that is active in that VLAN. As high security areas are secured with special (IPv4)hardware we cannot achieve the same level for IPv6 with current means.

In summary, in some VLANs the IPv4 router will not physically be the same as the IPv6 router. So the IPv4 network will differ from the IPv6 network, the NOC has to handle two different networks and we cannot achieve a full dual-stack mode from day one.

13.1.3.3 Setup

Currently the two transit routers ISPA and ISPB are the only Catalysts with a SUP720 engine. In contrast to SUP2, this engine is capable to forward IPv6 in hardware. So in the beginning we decided to use these routers as the core routers for the IPv6 network and to create the VLAN interfaces on these routers instead of using the routers at the customers end. Eventually the engines in others routers will be replaced with a SUP720 and every time this happens, we will extend the IPv6 network to these routers and will create the proper VLANs on those stub routers. In fact, currently one area is already equipped with such engines and the IPv6 core network spans these routers. In future, more and more of the older engines will be replaced by SUP720 engines.

13.1.3.4 External connectivity

For the time being, ISPA will take over the IPv6-to-IPv4-tunnel to 6WiN. Hopefully DFN is transforming 6WiN into a network that uses their IPv4 network. There are plans going on to use the G-WiN MPLS core to achieve this. This way, we could use the native IPv4-only connection for both IPv4 and IPv6. In addition we will add a second native connection to the IPv6 pilot project of T-Com in the near future.

13.1.3.5 Internal connectivity

As mentioned above the Catalyst 6500's with SUP720 will be used to form the IPv6 core network. In addition we have two Cisco 7206 routers at hand that can be used.

Most of the physical connection will use the same lines as for the IPv4 network. As the whole core network is trunked we have the option to use transit VLANs if necessary.

To connect the clients we use:

- Dual-stack in the core and for the VLANs
- A tunnel broker for satellite subnets without native IPv6 connectivity
- A tunnel broker and ISATAP for satellite clients without native IPv6 connectivity

ISATAP

The ISATAP server is already running on one of the Cisco 7206 routers. Clients have to configure the proper target IP or - if capable - can use the name isatap.uni-muenster.de.

Tunnel broker

The tunnel broker is an OpenVPN system as described in Chapter 5. Currently hosts get a /64 prefix (if more than one host is connected) or a /128 address (if only that host is connected). As soon as the university gets an extended address space it is intended to assign proper /48-prefixes to end sites, to meet the requirements of RFC3177.

Currently the used addresses and prefixes are hard configured in the script we deliver with the clients OpenVPN files. In the future we will use DHCPv6 prefix delegation to assign the /48 prefixes to clients.

IGP

The IPv4 network uses OSPFv2 as internal routing protocol. So far there was no need for an IGP for IPv6 as we only had one router with multiple VLANs attached. But for the extended IPv6 network we need an IPv6 IGP. We had some experience from running the 6WiN where we used IS-IS, so now there are several combinations what IGP's to use for these two networks.

- OSPFv2/OSPFv3
- OSPFv2/ISIS
- IS-IS single topology
- IS-IS multi topology

Running IS-IS with a single topology is not an option, as we do not intend to run IPv4 and IPv6 on the same topology. In any case, running IS-IS only would require that we substitute the current OSPFv2 with IS-IS. We feel that this is too much of an effort, at least for our size of network.

So the choice is only between OSPFv3 and IS-IS for IPv6. While we had some experience with IS-IS in the project, the University's NOC has more experience with OSPF. As OSPFv2 and OSPFv3 don't differ much, we felt it was best to go with the OSPFv2/OSPFv3 combination.

DNS

So far the standard University's nameservers (running BIND 8.3.x) are not used for IPv6, they serve IPv4 hosts only. Within the JOIN project we run a dedicated IPv6 nameserver (running BIND 9.3.1) that stores all AAAA resource records. It hosts a subdomain of the University domain uni-muenster.de. The University's main nameserver delegates the zone ipv6.uni-muenster.de to the dedicated IPv6 nameserver. This nameserver also hosts the reverse delegation for the IPv6 addresses. This is the default behaviour; there are a few (manual) exceptions.

As a result, you have to use different names when you want to access something via IPv4 or via IPv6. There is an additional domain (join.uni-muenster.de) which contains both AAAA and A records. Most times we use names out of this domain.

In order to deploy IPv6 in the University we have to add AAAA record to the now standard names in domain uni-muenster.de. This requires two prerequisites. First, we have to upgrade the University's nameservers to a newer version. While the currently running servers are able to store IPv6 records, they do not talk IPv6 transport and thus cannot answer to a request coming from an IPv6 host via an IPv6 packet.

To achieve this, the already running BIND 9.3.1 nameserver will become a secondary nameserver to the University's main servers. In addition, two additional new servers will be set up and the old server will get replaced eventually. Secondly, we have to add an AAAA record next to the A records in the domain uni-muenster.de and get rid of the two subdomains. This is the more complex part, as the zonefiles are generated with some scripts reading from a large database. The database and scripts have to be adapted.

Database

The main support tool our NOC has is the very extensive Oracle database. As mentioned before it contains every detail of the network, e.g. routers, hosts, cables and any kind of configuration data, especially IPv4 addresses. The NOC people feed the database via web pages that use homegrown ASP scripts. They retrieve and check data likewise. The database is also used to generate several statistics and lists. For our case it is important that the name server and the DHCP server are configured via scripts that read from this database.

To introduce IPv6 in the University's network it is necessary to integrate IPv6 addresses in this database, so that the University's NOC people can manage IPv6 as easily as IPv4. The database itself needs new fields to hold IPv6 addresses and prefixes and new pointers to relate these to hosts and routers. Web pages are required to manage IPv6 address space and to add IPv6 addresses. The scripts that generate lists need to get updated to use these new fields and to generate proper lists.

DHCPv6

All hosts in the University use DHCP to get an IPv4 address. All users have to register their hosts with their MAC address to the NOC and with this information the database and then the DHCP server is configured. With this policy, the NOC is in a position to maintain maximum control over IP addresses and who is using them. They want to keep this level of control for IPv6 too.

There are two methods to assign an address to a host in IPv6 automatically, via Router Advertisements (RA) or via stateful DHCPv6. As the interface ID that is used when configuring addresses with RAs is derived from the MAC address (EUI-64 format) we could use the data that is already available in the database. But, the fact the MAC address is visible in the network and that a hardware component of a host is identifiable, was a major concern of our NOC people. They had the strong wish to use other means to form an interface ID than the EUI-64 format. There are several ways to achieve this:

- configure IPv6 address manually
- use privacy addresses (RFC3041)
- use stateful DHCPv6 with non-EUI-64 address formats

The first option is clearly unacceptable as it disables autoconfiguration and complicates management. If we would use the second option, the IPv6 addresses are randomized and we would loose control of the addresses and there is no longer a relation between user and IP address. This is something our NOC people strongly rejected.

The last available option for us was to use stateful DHCPv6. Unfortunately as of today there is no production stateful DHCPv6 server available for us. So the decision was to start with RAs and addresses in EUI-64 format (this was still disliked, but the least painful option) and change to stateful DHCPv6 later on, when mature products are available.

Usage of SubnetID

As recommended for any enterprise site, the University of Münster received a /48 prefix from its provider DFN. We discussed a lot how to use the available address space and the SubnetID with the technical personnel of ZIV-NOC, based on our own discussion paper.

We came to the conclusion that a single /48 prefix is insufficient for our needs. There are multiple reasons. In the first place our University is very large and has ten thousands of students and personnel. Then the ZIV is not only giving connectivity to the University but is also managing the network and is giving broadband access to other parties, like the medical facilities and e.g. the Fachhochschule.

Thus ZIV-NOC is more a provider than just a University's NOC. Finally ZIV-NOC offers all students and personnel Internet access via ISDN. Following the rules in RFC 3177 [RFC3177] we should allocate a /48 to most of these users. With more than 50,000 students we would need more than a /32 prefix to achieve this. Something similar applies for VPN customers using the IPv6 tunnel broker system.

We tried to obtain more address space from our provider DFN. After lengthy discussions with DFN and RIPE, DFN was not able to fulfil our needs. The solution was to become a RIPE member, an LIR, and apply for our own /32 prefix. This way we would be allocated enough address space and additionally solve any multihoming issues, by having what amounts to Provider Independent (PI) IPv6 address space (we would be the provider).

At the time of writing, we have become RIPE members, but haven't been allocated a prefix so far. Therefore we will use the standard IPv6 documentary prefix 2001:db8::/32 (see RFC 3849 [RFC3849]) for examples in this section.

Usage of the provider /32 prefix

We split the 16 bit provider address range from bit 33 to 48 into four large chunks and defined a Format Prefix (FP) for each address range:

< 2 bits Format Prefix >< 14 bits free for use >

The Format Prefix is used in this way:

FP 00:	2001:db8:0000::/48 - 2001:db8:3fff::/48:	Standard broadband customer allocations
FP 01:	2001:db8:4000::/48 - 2001:db8:7fff::/48:	VPN access, e.g. OpenVPN tunnelbroker system
FP 10:	2001:db8:8000::/48 - 2001:db8:bfff::/48:	Remote Access Service (RAS), e.g. ISDN dial-in customers
FP 11:	2001:db8:c000::/48 - 2001:db8:ffff::/48:	Reservation for RAS or other

Usage of the last three blocks is hopefully very simple. Except for the first block, prefixes will become allocated consecutively starting with Zero (of each block). Currently we have about 8,000 RAS customers and about 15 OpenVPN tunnel broker customers. While the amount of RAS customers rises slowly, we expect a significant increase in the OpenVPN tunnel broker customers, once the tunnel broker system is integrated in the management plane of ZIV-NOC and we are able to give production IPv6 access via that system. The FP 11 block can be used for further RAS users if the FP 10 is ever exceeded or for any other usage if it arises. The only Format Prefix block that needs further subdividing is FP 00.

At least the addressing plan of the University's /48 prefix is very tight and it is highly likely that it will be exhausted at some point in time in the future. So instead of addressing consecutively like in the other three Format Prefix blocks, making reservations makes sense. Instead of making sparse reservations we choose to be generous here, as we do not expect to see that many broadband users for

this Format Prefix block anyway. We split up the FP 00 address range into 64 chunks of /40 prefixes. Every customer will get a /48 prefix out of this chunk and the rest will be reserved for future use. Only half of the available FP 00 block will be used in this way, the second half (which equals a /35 prefix) will remain untouched and will be used only if this allocation method is too generous.

In addition, we spared the first /40 block for special prefixes, such as a /48 prefix for addressing of the backbone and a /48 prefix for specially needed subnets (like DNS servers and RPs).

In summary, the address plan for FP 00 looks like this:

2001:db8:0000::/40:	Assigned for special purposes
2001:db8:0000::/48	Used for addressing of the backbone
2001:db8:0001::/48 - 2001:db8:000f::/48	Reserved for backbone addressing
2001:db8:0010::/48	Used for special subnets (DNS, RP, etc.)
2001:db8:0011::/48 - 2001:db8:001f::/48	Reserved for further special subnets
2001:db8:0020::/48 - 2001:db8:00ff::/48	Free
2001:db8:0100::/40:	Assigned to University of Muenster (WWU)
2001:db8:0100::/48	Allocated to WWU network
2001:db8:0101::/48 - 2001:db8:01ff::/48	Reserved for further extension
2001:db8:0200::/40:	Assigned to Medical Clinic of Muenster (UKM)
2001:db8:0200::/48	Allocated to UKM network
2001:db8:0201::/48 - 2001:db8:02ff::/48	Reserved for further extension
2001:db8:0300::/40:	Assigned to Fachhochschule Muenster (FH)
2001:db8:0300::/48	Allocated to Fachhochschule network
2001:db8:0301::/48 - 2001:db8:03ff::/48	Reserved for further extension
2001:db8:0400::/40:	Assigned to Max-Planck-Institut Muenster (MPI)
2001:db8:0400::/48 - 2001:db8::4ff::/48	Allocated to MPI network
2001:db8:0500::/40:	Assigned to Studentenwerk Muenster
2001:db8:0500::/48 - 2001:db8::5ff::/48	Reserved
2001:db8:0600::/40:	Assigned to School network Muenster
2001:db8:0600::/48 - 2001:db8::06ff::/48	Reserved
2001:db8:0700::/48 - 2001:db8:1fff::/48	Free
2001:db8:2000::/48 - 2001:db8:3fff::/48	Unused

Usage of the University's /48 prefix

The address plan used for the /48 prefix for the University's network is older than our plans to apply for a /32 prefix. Nevertheless, as seen from a providers point of view, the University is still a single customer and a /48 should be enough. From this point of view we kept the /48 address plan and just freed and moved some reserved prefixes - those that are used to address the providers backbone - to other address ranges (see above).

In our University we have a sophisticated organisation structure for who is responsible for which institute or department in terms of computer and network support. Those units are named IVV (Informations-Verarbeitungs-Versorgungseinheit) and each of them takes care of one or more faculties.

The network partially reflects this organisation structure, so we decided to arrange the network addressing structure likewise. We refrained from dictating too strict an addressing framework, but decided to just use that IVV structuring and leave large blocks of the SubnetID range untouched for future use and let the IVV management decide how to use their SubnetID range.

Apart from this really simple method, we specifically refrained from using any kind of mapping any kind of information into the SubnetID field (like IPv4 addresses or VLAN IDs). The ZIV-NOC people hat strong concerns against such an approach, as change in the mapped structure would require renumbering of the network to keep the address plan intact. That was not accepted at all. One could say that using IVV unit IDs is the same kind of method, which is true. But in contrast, this organisational structure is vital for our network and our management to work, and it hasn't changed ever since it was created.

We ended up with the following address scheme for the SubnetID address range:

< 3 bits FP >< 4 bits IVV >< 9 bits free for use >

We defined a Format Prefix again. This time we used 3 bits yielding 8 different address blocks, but only the first 2 Format Prefixes are used. The last 6 FPs are spared for any future addressing plan, or if the current address range is too small. Currently we have 10 IVV units, so 4 bits - giving room for 16 of such units - will be more than enough. It is highly unlikely that there will be more IVV units in the future, much more likely that the number will be reduced.

How each IVV will use their remaining 9 bits of SubnetID address range depends on the administrators of each IVV. We decided to leave that decision in their hands. 9 bits for each IVV will be sufficient. Currently we serve about 600 subnets in the whole network. With 9 bits address range, each IVV can address 512 subnets. Even some hierarchical structure is possible. Some of the IVV assist more than one faculty, so our generic advice is to use the first 2 bits of the IVV address range to indicate the faculty. But that is not a requirement, just an advice.

The address plan in detail:

FP 000:	2001:db8:100:0000::/51:	Used for addressing
	2001:db8:100:0000::/55:	IVV 1
	2001:db8:100:0000::/57:	History/Philosophy
	2001:db0:100:0080::/57:	Philology
	2001:db8:100:0100::/57:	Speech Centre
	2001:db8:100:0180::/57:	Free
	2001:db8:100:0200::/55:	IVV 2 (Economic Science)
	2001:db8:100:0400::/55:	IVV 3 (Law)
	2001:db8:100:0600::/55:	IVV 4
	2001:db8:100:0600::/57:	Physics
	2001:db0:100:0680::/57:	Chemistry/Pharmacy
	2001:db8:100:0700::/57:	Biology
	2001:db8:100:0780::/57:	Free
	2001:db8:100:0800::/55:	IVV 5

2001:db8:100:0800::/57:	Mathematics/Computer Science
2001:db0:100:0880::/57:	Psychology/sports science
2001:db8:100:0900::/56:	Free
2001:db8:100:0a00::/55:	IVV 6 (Geology)
2001:db8:100:0c00::/55:	IVV 7
2001:db8:100:0c00::/57:	Evangelic Theology
2001:db0:100:0c80::/57:	Catholic Theology
2001:db8:100:0d00::/57:	Educational/Social Science
2001:db8:100:0d80::/57:	Free
2001:db8:100:0e00::/55:	IVV 8 (Medical Science/UKM)
2001:db8:100:1000::/55:	IVV 9 (University Administration)
2001:db8:100:1200::/55:	IVV 10 (ZIV)
2001:db8:100:1400::/48 - 2001:db8:100:1fff:/48:	Unused
FP 001: 2001:db8:100:2000::/51:	Reserved for extension of addressing range
FP 010: 2001:db8:100:4000::/51:	Unused
FP 011: 2001:db8:100:6000::/51:	Unused
FP 100: 2001:db8:100:8000::/51:	Unused
FP 101: 2001:db8:100:a000::/51:	Unused
FP 110: 2001:db8:100:c000::/51:	Unused
FP 111: 2001:db8:100:e000::/51:	Unused

This current address plan is still under discussion. With a /32 and probably an extension of the /48 prefix at hand there is no need to be that sparing. For example, as mentioned above, we already swapped out some address ranges to external /48 prefixes. It is in question if we still really need 3 bits for Format Prefixes, but should instead use only 2 bits and give 10 bits to the IVV units.

Usage of the Medical Clinic's /48 prefix

ZIV-NOC is also responsible for the Medical Clinic's network and will also create the address plan for them. This has to be coordinated with clinic administration. It is most likely that they will get a similar addressing plan to the one above, because they have a similar organisational structure. So probably with some bits shifted, we will apply the same method. This is still under discussion.

Usage of the other /48 prefixes

Further /48 prefixes like those for the Fachhochschule are not in the responsibility of ZIV-NOC, as these are independent institutions and manage their network themselves.

Management

The main tool for management of the IPv4 network that is used by the University's NOC is Tivoli Netview. So far we use smaller tools like Argus and MRTG to monitor the IPv6 network. These smaller tools would be insufficient for a larger IPv6 network and it would require the University's NOC to learn and use a second tool next to Netview. It is better and more ergonomic to stick with one tool that can handle both, IPv4 and IPv6.

Unfortunately Netview does not support IPv6 yet. Still, it is adaptable and we can write own scripts to probe the IPv6 network, which can be inserted in the Netview command environment. So far, we

haven't evaluated the amount of work to be done here, as we can use the "smaller" tools, as long as we run only a pilot network. So we still wait and hopefully there is then more time for IBM to react and integrate IPv6 functionality into Netview.

Services

During the lifetime of the 6NET project, JOIN helped to integrate IPv6 in several applications and set up such services to test them and to make all common services available over IPv6. After some time we now offer the following services within our University:

- Nameservice (DNS)
- Webservice (HTTP)
- Fileservice (FTP)
- Mail (SMTP)
- Time (NTP)
- News (NNTP)
- Jabber server (IMPP)
- Conference server (OpenMCU/Asterisk)

In order to launch a pilot we are in the lucky position to already have all major applications at hand with IPv6 support. A slight disadvantage might be that - except for the FTP service - all services run on dedicated hardware and are managed separately from the standard IPv4 application services in the University. While it is good to have all services at hand, it is a drawback to have to maintain two sets of hardware. One piece of advice might be not to set up IPv6 services on different hardware, but to try to integrate IPv6 in the already running applications/services from the start. This should be possible for nearly all modern applications. This could prove difficult or at least need additional work though, as one has to integrate IPv6 in the management structure of these applications as well (see our efforts to try to integrate IPv6 in the name service).

13.1.4 Summary

In order to integrate IPv6 in the University's premises we have to work on the following items:

- External connectivity
- Internal connectivity
    - Select/purchase Hardware
    - Select IGP
    - Choose VLANs to add IPv6 to
- Invent a suitable addressing scheme
- Integrate IPv6 in database
- Adapt DNS
- Add DHCPv6
- Adapt Monitoring Tools
- Educate NOC

Some of these items are easy to achieve, some of them are tough and long term work. One could say that integrating IPv6 in a larger network is not easy to do and needs a lot of work. Lots of resources - financial and personnel - are needed. If you do not have the proper equipment, one has to buy new hardware. Most of the time the latest equipment is already IPv6 capable one way or the other, but you might also need additional equipment if you plan to run dedicated hardware for IPv6 or if you plan to use some transition mechanisms (e.g. hosts for ALG gateways and/or tunnel broker systems). Personnel costs might be even higher, because you need manpower for programming, education and debugging.

In summary, depending on the size and the level of complexity of your network and depending on how elaborate the tools you use to monitor and manage that network, integration of IPv6 might become an time-consuming experience.

13.2 Small academic department, IPv6-only (Tromsø, Norway)

This section describes what has been done by Telenor Research in cooperation with the University of Tromsø and Invenia Innovation.

13.2.1 Transitioning Unmanaged Networks

This section has its focus on transitioning so-called unmanaged networks. Unmanaged networks are relatively small and simple networks that are not administered by competent technical personnel. Instead, such networks are setup and administered by their users. An example of an unmanaged network is the network in the home of a private person. Another example is the network of a small business. Such an unmanaged network typically consists of a single subnet connecting the equipment of the user. This subnet is connected via a small dedicated router to an Internet Service Provider. This router (also known as gateway) is normally not actively managed and it has been supplied and initially configured by the ISP. In some cases, however, the ISP may actively manage the gateway. In this setting, the home user has no control over the gateway. The user neither supplies nor configures nor manages the gateway that connects his network to the ISP. This restriction on the user may result in that the user adds an additional router between his home network and the gateway supplied by the ISP in order to gain control or add additional functionality. In the remainder, we will use the term gateway to refer to the router/gateway supplied by the ISP that provides the user with upstream connectivity.

Unmanaged networks were discussed by the last of the so-called deployment teams within the IETF v6ops working group. This deployment team has identified application requirements and relevant scenarios for transitioning unmanaged networks in “Unmanaged Networks Transition Scope”. Their work meanwhile reached information standard status and was published as RFC 3904 [RFC3904]. This is a short summary:

The main application requirements with regards to introducing IPv6 in an unmanaged network can be summarized by:

1. Applications running fine over IPv4 should continue to run fine after IPv6 is introduced in the network.
2. New applications that are hard to realize over IPv4 should be able to benefit from the introduction of IPv6 in the network.
3. Deploying new IPv6 applications should be simple and not create problems.

The unmanaged network topology includes 3 parts: the ISP, the gateway and user equipment. Each of these can either support IPv4, both IPv4 and IPv6 or can support IPv6 only. This gives 27 (3*3*3) combinations of IPv6 support in unmanaged networks. The IETF deployment team selected 4 out of these 27 scenarios for further consideration:

A. Gateway supports IPv4 only.
B. ISP and gateway support IPv4 and IPv6, i.e. they are dual-stack.
C. Gateway supports IPv6, but the ISP supports only IPv4.
D. ISP supports IPv6 only.

We refer to the draft for a more elaborate description of these scenarios.

13.2.2 Implementation of a Pilot Network

Telenor Research in Tromsø, Norway designed and implemented an IPv6 network mostly based on radio links in cooperation with the Department of Computer Science at the University of Tromsø and Invenia Innovation AS. It covers most of the city. The purpose of establishing such a large pilot network is to experiment with an infrastructure that in design will be identical to what an ISP will have in the last phase of migration from IPv4 to IPv6. That is, an infrastructure where the ISP only supports IPv6 inside the network, but where customers will have legacy applications that can not run on anything but IPv4. Customers will have nodes running IPv6-only, IPv4-only and dual-stack.

In this experiment, the users have been staff from Telenor Research, faculty and technical staff from the Department of Computer Science, and students. The level of user competence has varied tremendously, from seasoned IPv6 administrators to complete novices with regards to IPv6 or (network) technology in general. The unmanaged networks were the home networks of these users. This experiment is therefore an example of Scenario D identified by the IETF deployment team on unmanaged networks.

The network has been designed to mimic the situation that an ISP will face and problems related to traffic at the user level is not in our focus. This leads to the following requirements:

- An IPv6-only core.
- All home networks must support a mix of IPv6-only, dual stack and IPv4-only hosts.
- IPv6-only should not need any modification.
- IPv6-only nodes must have access to services that run IPv4 only (outside of our network)
- Nodes running IPv4 only should have a functionality and performance as if the ISP supported IPv4 native.
- Nodes running IPv4-only must be given access to services on IPv6-only.
- The machinery for transition from IPv4 to IPv6 must be transparent.

13.2.2.1 The core of the network

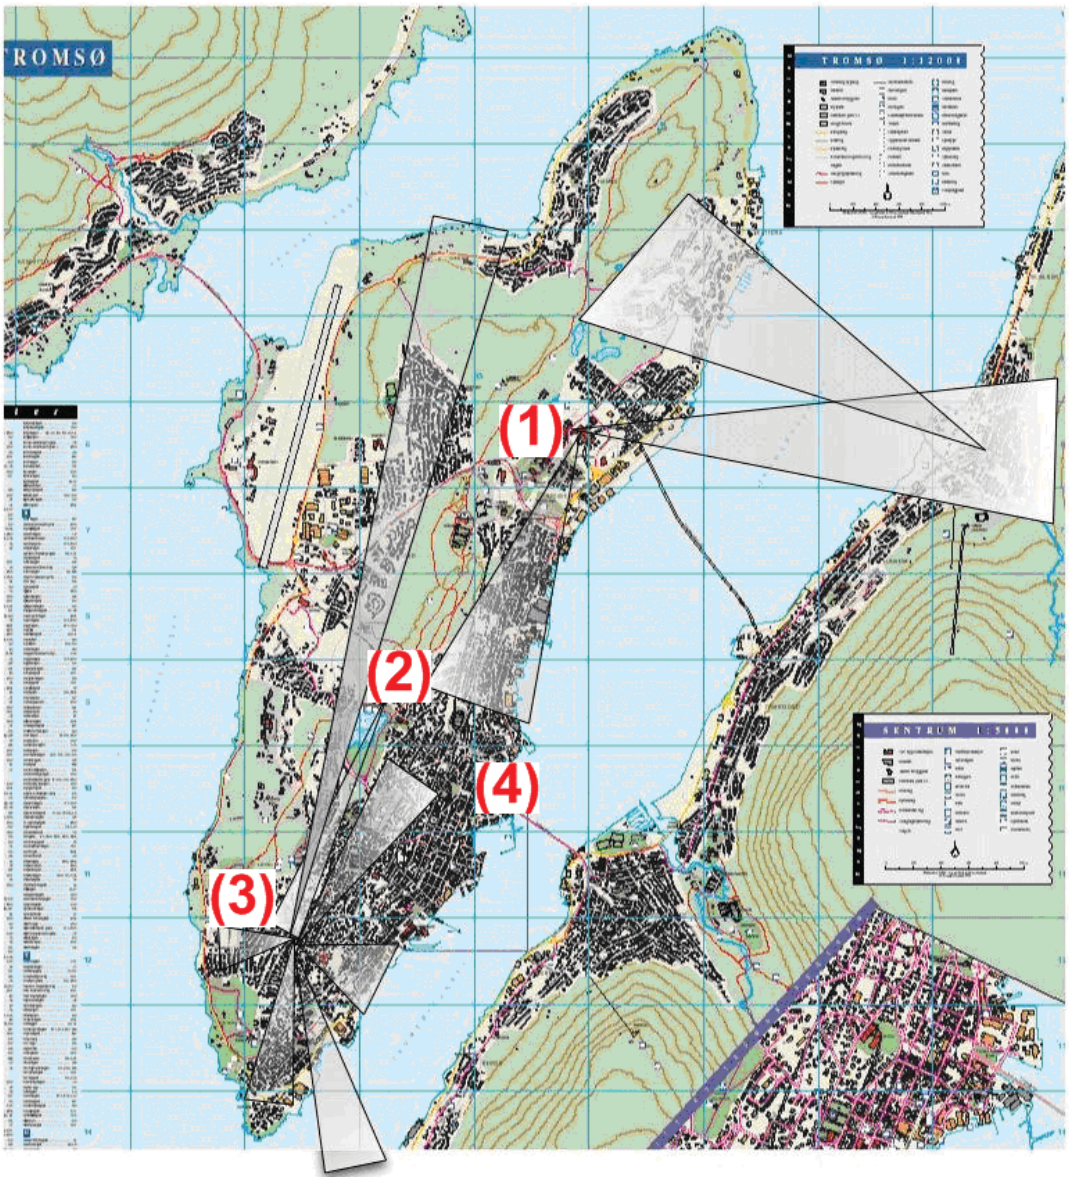

Figure 13-3 Overview of Tromsø and Transmission Points

Three main transmission points have been established in Tromsø; they are both connected to the 6NET IPv6-only network. From these three main points, a total of 12 new networks have been built. As of writing there are somewhere between 15 and 20 home networks spread out over these 12 networks, with an unknown number of machines connected.

In practice, connectivity to 6NET is provided by Uninett and the first transmission point is located at the University of Tromsø. This point is marked (1) in Figure 13-3. From there a fibre carries the traffic to the second transmission point (marked (2)). From there a single beam carries the traffic across town to a high mast (marked (3)) from where most of the city centre is covered. In addition, the

Student House ("Driv") downtown was connected to the network by means of a fibre; this point is marked (4).

A variety of antennas are in use, from large 24dBi dishes to small 6dBi omni-directional. The network is fully routed with a mix of Cisco and NetBSD routers. In addition, all home networks are routed and prefixes are allocated to homes as needed. The core is IPv6-only in order to ensure that no traffic "leaks" out of the home networks.

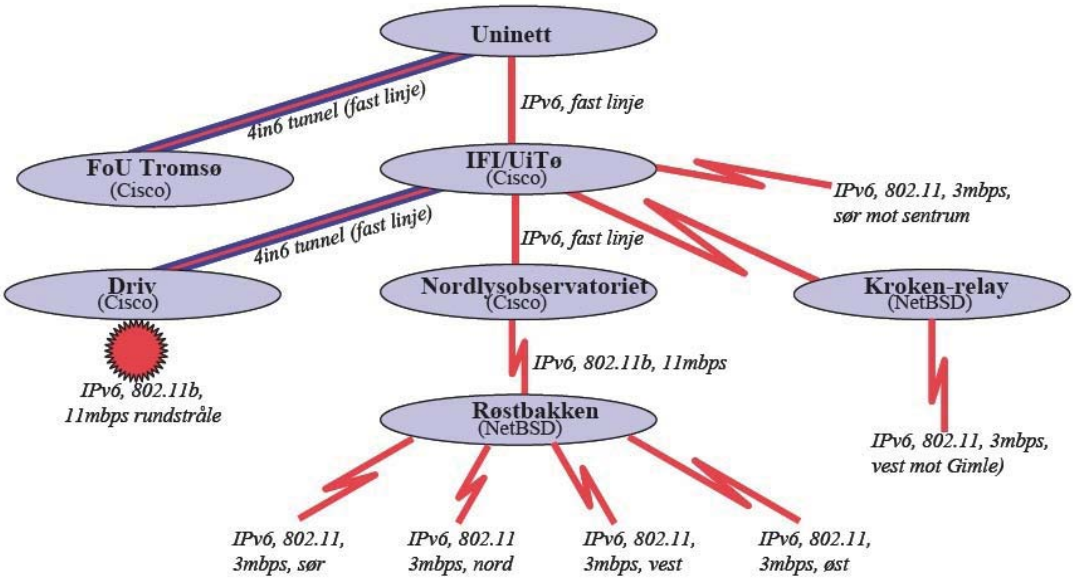

Figure 13-4 Topology Overview

13.2.2.2 The Home networks

All home networks are connected to the core by 802.11 WLAN. All networks use a small PC with two network cards that runs the NetBSD operating system as home gateway. Some of the home gateways run RIP, but most have default routing hard coded into their initial configuration. A prefix is manually allocated to each home network and routed to the NetBSD machine. The home networks use a /56 prefix. Almost all home networks are dual stack (on the inside, that is). A subset of the home networks has been used to experiment with complete IPv6-only operation.

All necessary software to facilitate the requirements of the network must be run on the NetBSD gateway. In particular, traffic of all kinds, including DNS that arrives at the router must be dealt with in a transparent manner. For the address allocation and routing of IPv6 packets the standard IPv6 autoconfiguration mechanisms were used. The gateway thus supplies user machines with the network prefix and the default route for IPv6. In addition, the gateway runs DHCPv4 which is used to assign non-routable private IPv4 addresses and IPv4 default route to user equipment. DHCPv4 is also used to configure the DNS server to use by user equipment. The DNS server configured is the non-routable IPv4-address of the home gateway. DHCPv4 was used to lacking DHCPv6 support or other DNS discovery mechanisms. Pure IPv6-only nodes that do not use or implement DHCPv4 therefore need to have their (IPv6) DNS server address configured manually.

On each home router a DNS proxy was run which would forward the DNS requests (over IPv6) to a real nameserver running at Telenor Research Laboratories in Tromsø. Each home network was given a

domain beneath tft.tele.no. The DNS proxy is the totd proxy designed and implemented by Invenia Innovation within the 6NET project.

13.2.2.3 Interoperability

In this section, we describe how traffic makes its way from an application to the Internet. How it is treated depends on its type (IPv6, IPv4, TCP or UDP), its source, and its destination.

IPv6 from home network to IPv6 Internet

This traffic passes unaltered through the network. All home networks have global 2001::/16 prefix addresses and routing works as expected. There is no special treatment or translation of packets.

IPv6 from home network to IPv4 Internet

The IPv6 traffic to IPv4 services must be translated. We have opted for FAITH and NAT-PT. The former is a transport level connection translator, while the latter is a network level packet translator. While both translators are installed and operational, in practice, the transport level translator is used for daily use. The available NAT-PT implementation has a variety of non-trivial problems and does not seem ready for production use. It is installed and used in order to try to get a better understanding of its problems.

The problem with the faith implementation is that it only supports TCP. However, we have noticed that although there has not been support for UDPv4, no user request for it has been received. A lot of effort was saved this way. It is possible to add UDP support to FAITH and Invenia Innovation and the University of Tromsø have made an experimental implementation of such UDP support. This experiment was a successful proof-of-concept but further development needed for actual deployment was not performed due to limited interest compared to the effort involved and the number of problems to resolve. The most important application using UDP is DNS which is handled by the DNS proxy (Application Level Gateway). Most other UDP applications in use are multiparty games. At the moment there are few games with IPv6-capable clients that want to interoperate with their IPv4 counterparts.

In any way, DNS must be dealt with somehow in order for faith and NAT-PT to work transparently to the user. When a request is made, that results in an IPv4 address, the IPv6-only node would not know what to do with the result. To that end, we use the totd (Trick Or Treat) 6NET DNS proxy. This proxy needs to be accessible over both IPv4 and IPv6. Since the core is IPv6 only, this translation must occur at the “upper” edge of the network, i.e. at each home gateway.

Each home gateway then runs a totd DNS proxy, configured with the FAITH and/or NAT-PT network prefixes and the address of the ‘real’ recursive nameservers. In addition, the gateway is configured with a static route for these prefixes such that the traffic that requires translation is properly routed to the machines hosting the translators.

IPv4 from home network to IPv6 Internet

For HTTP/HTTPS and FTP connections we use the www6to4 ALG developed by Invenia Innovation AS. This simple ALG does not ‘understand’ the FTP protocol and can only forward it to another proxy that does. In our case, it forwards all FTP connections to a single central ftp proxy on a dual-stack machine located at the Telenor Research Lab. HTTP(S) traffic is handled directly by the www6to4 proxy itself. A more complex proxy could be developed or used that does understand ftp itself. However, the main advantage of www6to4 is its extremely small size, its simplicity of configuration and it has no complicated failure modes. As it is installed on each home gateway, these are important considerations.

In addition, transport level relays are used to translate connections for specific applications between specific source IPv4 address and IPv6 destination address. These relays are configured statically in

each home gateway on user request. These relays are used mostly to allow IPv4-only email clients to connect to specific (IPv6) POP, IMAP or SMTP servers.

IPv4 from home network to IPv6 Internet

There is no direct IPv4 connectivity from the home network to the IPv4 Internet, and no network or transport layer translators are available currently for IPv4 to IPv6 translation. For this traffic two translations are required. First, as already explained the traffic is translated into IPv6 and routed to the main networking server. There it is again translated into IPv4 as described above.

IPv4 from Internet to IPv6 service at home

Support for this has been in a rather ad-hoc manner. Those who have servers at home need to give notice, and by means of 6tunnel (which is a Transport Relay Translator) incoming connections are then forwarded over IPv6 to the server on the home network. It works by reserving an address/port pair for each service on a dual-stack machine at the ISP. That is, the port is used to differentiate between the different sites (home networks) on the inside. This is not an elegant solution, but it works for simple protocols such as HTTP (not for FTP for example). There are obvious problems related to DNS since the target (where the service runs) and the relay do not have the same DNS entry (reverse lookup mismatch). There has been very little traffic over this mechanism and it is available mostly for completeness.

An alternative approach is to use the ALG proxy support (in apache 2, for example) to relay connections from the ISP to servers at the home networks. We have successfully tested this approach (although we inadvertently created a spam-relay for awhile due to inappropriate access control), but did not use it extensively.

IPv4 from Internet to IPv4 service at home

IPv4-to-IPv4 traffic is dealt with in the same manner as above, but with an additional translator running on the home gateway performing the translation from IPv6 to IPv4.

IPv6 from Internet to IPv4 service at home

This is dealt with by the same Transport Relay Translator discussed above.

IPv6 from Internet to IPv6 service at home

No special action is needed.

13.2.2.4 Summary

Only non-routable (private) IPv4 addresses are used on the home networks. This triggers the use of translators of various kinds. That the core only runs IPv6 complicates the matter further, but not beyond what can be maintained and run.

The components for IPv6 to IPv4 are:

- Translating IP directly (NAT-PT), and copying the user's data from one connection to another (FAITH). Cooperation with the totd DNS proxy is needed to get the connections needing translation to the translators.
- TRT (Transport Relay Translator) and one-to-one translation (6tunnel).
- Application specific solutions (for FTP).

The components for IPv4 to IPv6 are:

- TRT (Transport Relay Translator), 6tunnel

- Application specific translators (HTTP proxy and FTP)

13.2.3 Evaluation of the Pilot Network

The pilot network has been operational for more than two years now. In this section we describe our experiences and some problems we encountered with regards to the implementation and management of such a network.

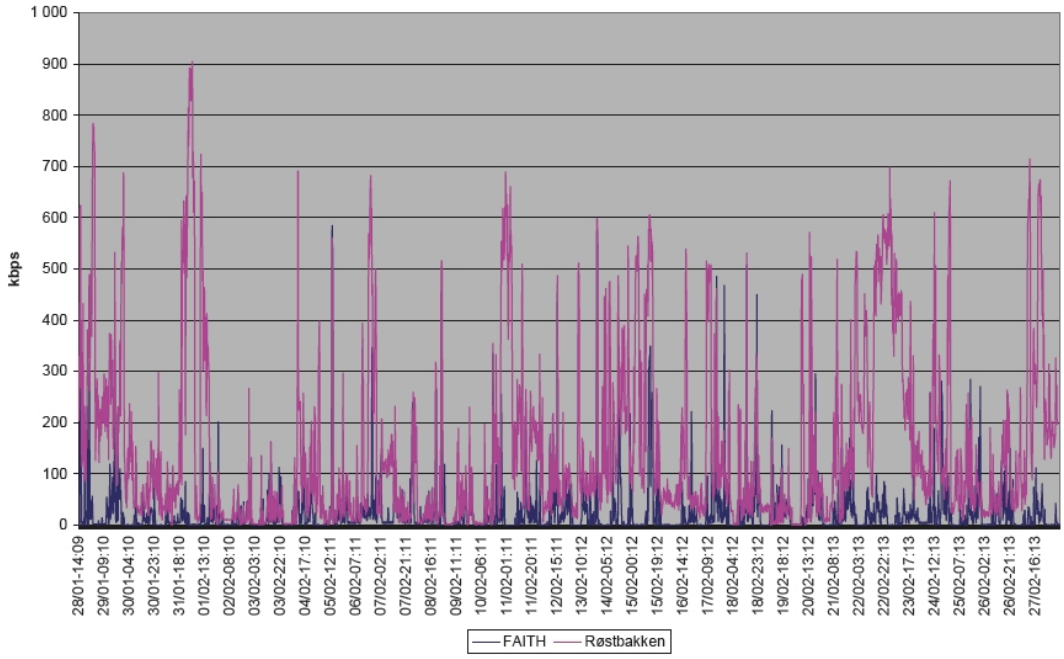

Figure 13-5 Traffic Statistics of Røstbakken Tower and faithd

The figure above gives an impression of the traffic patterns encountered in the core of our metropolitan network. The figure shows traffic measured at Røstbakken tower which is a central connection point for the majority of the home networks. In addition, the figure shows traffic measured at the most used IPv6 to IPv4 translator, namely the machine doing faithd TCP relaying that is located in the Telenor Research Laboratory. This does not give a complete image of the traffic on the whole network. Not only does it not capture the traffic from the northern part of the network, it also does not show the (IPv6-only) traffic produced by the IPv6-only network of the University of Tromsø that uses the same infrastructure.

13.2.3.1 NAT-PT

The NAT-PT implementation from KAME for FreeBSD that we used is not as mature as most other software from by the KAME project. Problems were exposed by HTTP connections to some websites failing when going through NAT-PT. We have not analysed the problems in more detail, nor extended the experiment to other applications. This decision was mainly motivated by the impression that the KAME project itself had abandoned further development of the NAT-PT code.

13.2.3.2 Faithd

Faithd only supports TCP connections and for each TCP port a separate process needs to be started. We have experimented with adding UDP support by making a python prototype (proof of concept). We did not continue to develop a production quality C implementation due to lack of interest by the users. In addition, we modified the original faithd to relay TCP connections for any port. This simplifies its configuration dramatically. Access control is then performed using firewall filtering. However, this modified faithd requires a small patch to the kernel code (NetBSD in our case). As our patches where not accepted by the NetBSD project team, kernel patching is required to use our faithd. This does not outweigh the advantage of easier configuration, and we therefore abandoned this modified faithd too.

13.2.3.3 6tunnel

Using 6tunnel for connecting to POP and IMAP servers does not allow SSL to be used. This is because unmodified SSL implementations use source and destination addresses as endpoint identifiers in their security associations. With 6tunnel in the middle modifying endpoint addresses, this fails.

13.2.3.4 DNS

Occasional problems were discovered where totd failed to deal with particular DNS server implementations. Some problems where dealt with in subsequent totd versions. One remaining problem is that some old nameserver implementations return a 'Server Failed' instead of the prescribed 'No Such Domain' response when asked for an IPv6 address record (that these nameservers have no knowledge of). Of course, these nameservers should be fixed, but in the meantime newer versions of totd work around the problem by always asking a nameserver for an IPv4 record even when it returns 'Server Failed'.

Another DNS-related problem, that totd cannot solve is that some services use hardcoded IPv4 addresses instead of DNS names to save some DNS lookups. This is one of the main reasons why we preferred using HTTP ALGs running on dual-stack machines over translating HTTP connections on the network or transport layer. A few popular sites like www.hotmail.com and some advertisement suppliers used hardcoded IPv4 addresses in URLs. This generated many user-generated problem reports.

Finally, some servers advertise an IPv6 address in DNS but offer services that are IPv4-only. Due to totd, applications will use IPv4 through a translator when no IPv6 address is found, but they can/will not fallback to IPv4 through a translator when the IPv6 connection fails for some reason. An example of such a server was mail.netbsd.org which was a dual-stack machine with IPv6 address in DNS, but ran an IPv4-only sendmail configuration. Sending email from an IPv6-only mail server to mail.netbsd.org therefore failed. This problem was fixed some time after the NetBSD maintainers were notified of the problem, but many sites are not so responsive to complaints from a (still) small user population.

13.2.3.5 Hardware Failures

We experienced several times problems with hardware used for IPv6 which worked fine when used for IPv4 only! Some problems could be fixed when identified like the one where a network card was too slow to initialize multicast reception to pick up replies to router solicitation it sent on startup. However, others remained a mystery and we had to 'fix' them by switching to other hardware. Luckily the failures were with low-end Ethernet cards that were cheap to replace.

13.2.3.6 IPv4-only applications at home

This required the most support and configuration due to the lack of a generic translation mechanism or ALG. ALGs for the most popular applications like Web browsing and ftp are available, but is often lacking for increasingly popular applications like multiparty gaming and P2P filesharing applications. A generic solution is to add tunnelling of IPv4 in IPv6. But in some cases the combination of reduced MTU due to tunnelling and non-routable addresses (going through one or more NATs) caused some applications to fail. However, in most cases, tunnelling works fine and transparent to the applications. We are therefore now introducing tunnelling at the time the experimental network is turning into a more permanent service to its users.

13.2.3.7 Security

We experienced the hard way, how several transition mechanisms can be (ab)used by third parties, mostly by spammers. We inadvertently created a spam relay twice. These relays where discovered quickly by spammers and used to forward almost 300,000 junk mails before we discovered what was happening. Proper access control, which takes into account the 'backdoors' created, is easily overlooked when introducing transition mechanisms. The most notable example of this was experienced when we started adding IPv4 over IPv6 tunnels from the home networks to the ISP (by the end of the pilot project), and a home gateway running an IPv4-to-IPv6 transport relay suddenly became globally accessible over IPv4. As formerly no global IPv4 addresses where used in the home network, no access control was deemed needed at the time the transport relay was setup. And, when a global IPv4 address was added later, the missing access control to the transport relay was overlooked creating a 'backdoor' to the main mail servers which believed the traffic through it to be from authorized home users.

13.2.3.8 Scalability

A variety of functionality to make IPv6 transitioning for a large number of unmanaged networks scalable and feasible is missing. Most notably we missed:

1. Dynamic DNS updating for autoconfigured addresses.
2. Automatic configuration on clients of IPv6 DNS resolver addresses (e.g. using DHCPv6, SLP or some well-known address).
3. Automatic configuration of transition mechanisms, especially those used on the home gateways.
4. Automatic configuration of the network prefix to use inside a home network.

The latter was sorely missed on a wireless based network with many links. It can be non-trivial to find out what antenna/link you are actually connecting too and what IPv6 addresses belong to it. We needed to define fairly elaborate technical and administrative procedures for new users to get their home gateway connected. We used the RIP routing protocol on home gateway routers and the routers of the core network. The main reason was not routing, but network debugging and making it easier to assist users remotely in getting connected.

13.2.4 Conclusions

Configuration of network equipment against a pure IPv6 ISP is still too complicated and out of reach of non-technical users. This is mostly due to the lack of autoconfiguration mechanisms for DNS localisation, prefix delegation and autoconfiguration of transition mechanisms.

The quality of service of the traditional IPv4 applications is significantly reduced because generic easy-to-setup transition support for many applications is still missing.

Our conclusion is then that we have shown that it is possible for ordinary users to live with an IPv6-only ISP, but at the cost of a significant reduction in user friendliness and quality of service. This means that it is still too early for an ISP to transition from IPv4 to IPv6-only. Currently, the best approach is a dual-stack ISP or maybe IPv6-only with IPv4 in IPv6 tunnelling. In our experience it seems that transition mechanisms are ok for centralised use, but that for IPv6-only ISP operation we should push for porting of applications to IPv6 instead of putting our hopes on better and easier to configure future implementation of translators.

13.3 Large Academic Department (University of Southampton)

In this section we begin by describing the systems components in this scenario of a large "departmental" network (1,500+ users with around 1,000 hosts) that wishes to transition to support IPv6. We describe the elements that need to be considered for the transition.

Southampton has been running IPv6 since 1996, but has only in the last two years adopted IPv6 as a production service in its network.

This scenario assumes no IPv6 is deployed beforehand, although in reality the transition at Southampton is already underway, in fact it is at the time of writing in a reasonably advanced stage. Thus after a review of the systems components, we present an overview of the status to date, current plans and next steps, and also the major remaining obstacles that have been identified.

Our motivation to deploy was in support of teaching, research projects, and communication with IPv6 networks (including those used by overseas students), while also encouraging innovation from staff and students – indeed already we have seen new streamed radio (Surge/Virgin[1]) and multicast video services (ECS-TV[2]) emerge.

This work has contributed to the IETF v6ops WG enterprise scenario descriptions and analysis. We have documented our campus experience in an Internet Draft [Cho04a] and also documented the IEEE 802.1q VLAN approach to introducing IPv6 [Cho04b] that has also been used at Muenster (see above). The [Cho04a] draft was built from a similar analysis as appears here.

13.3.1 Systems Components

The components fall into the following categories:

- Network components
- Address allocation components
- Services
- Host and device platforms
- User tools

We discuss these categories below.

In the light of the IETF v6ops WG activity on studying IPv6 network renumbering [RFC4192], we also cite components where hardcoded IP(v4) addresses may be found, that may need consideration in IPv6 networks. Further renumbering reporting can be found in D3.6.1 [D3.6.1] and D3.6.2 [D3.6.2].

13.3.1.1 Network

Physical connectivity (Layer 2)

The technologies used are:

- Switched Ethernet
- Gigabit Ethernet
- Wireless networking (802.11b)

[1] http://www.ipv6.ecs.soton.ac.uk/virginradio

[2] http://www.ecstv.ecs.soton.ac.uk

There is no use of ATM, FDDI or other "older" technologies. The network is purely Ethernet. IEEE 802.1q VLANs are supported by the network equipment.

Routing and Logical subnets (Layer 3)

The hybrid Layer 2/3 routing equipment is composed of Alcatel OSR and Omnicore L2/L3, with approximately 15 internal IPv4 subnets (in effect, routed VLANs). There is no specific internal routing protocol used. There is a static route via the site firewall to the main upstream provider (academic) running at 1Gbit/s, and there is also a static route to the secondary (low bandwidth) link offsite (commercial).

Hard coded IP information:

- The IPv4 address space assigned by academic provider
- There is hardcoded IP subnet information
- IP addresses for static route targets

Firewall

The firewall is currently a CheckPoint Firewall-1 solution running on a Nokia IP740 hardware platform. There is one internal facing interface, one external facing interface, and two "DMZ" interfaces, one for wired hosts and one for the Wireless LAN provision.

Hard coded IP information:

- Names resolved to IP addresses in FW-1 at "compilation" time
- IP addresses in remote firewalls allowing access to remote services
- IP-based authentication in remote systems allowing access to online bibliographic resources

Intrusion Detection (IDS)

The Snort package is used for intrusion detection.

Management

Some network management is performed by SNMP; there is no specific package for this. There is a greater emphasis on monitoring than explicitly in management.

Monitoring

A number of tools are used, to monitor network usage as well as systems availability, e.g. nocol, nagios and MRTG. The IBM AWM tool is used for network performance testing, along with iperf, rude and crude.

Remote Access

The components supporting remote access are:

- Livingston Portmaster 56K/ISDN dialup
- RADIUS server
- (Microsoft) VPN server

IPv6 access (e.g. for local testbed)

A native IPv6 service from the regional network (LeNSE) to the JANET (NREN) is used. This is facilitated by use of 6PE over MPLS.

Hard coded IP information:

- IP endpoints of upstream connectivity interface

13.3.1.2 Address allocation

The department receives its IPv4 and IPv6 address allocations from the University. For IPv4, the University has a /16 allocation which is not aggregated under the JANET NREN. For IPv6, the University receives its allocation as a /48 site prefix from JANET.

IPv6 network prefix allocation

The department currently has approximately 10 (non-contiguous) /24 IPv4 prefixes allocated to it by the campus computing services department (ISS).

For IPv6, JANET has the prefix 2001:630::/32 from RIPE NCC, as the national academic ISP in the UK. The University has been allocated 2001:630:d0::/48 by JANET. The department has been allocated a /52 size prefix 2001:630:d0::/52.

The department is a RIPE member and could obtain LIR status (as Muenster has done). However, we have not applied for an IPv6 prefix at this time.

In the initial deployment, we expect that IPv4 and IPv6 subnets will be congruent (and thus share and run over the same VLANs). We feel in an initial deployment that this approach simplifies management. We numbered our IPv6 links by topology, depending on which links were attached to which subrouter (in effect, an arbitrary scheme, given our small number of links – 15 or so).

The advantage for IPv6 is that subnets will not need to be resized to conserve or efficiently utilise address space as is the case currently for IPv4 (as subnet host counts rise and fall for administrative or research group growth/decline reasons).

Hard coded IP information:

- The IP prefix allocation from the university

IPv6 Address allocation

It is expected that the network devices will use a combination of address allocation mechanisms:

- Manually configured addresses (in some servers)
- Stateful DHCPv6 (probably in fixed, wired devices and some servers)
- Stateless address autoconfiguration (probably in wireless and mobile devices)
- RFC3041 privacy addresses (in some client devices)

For devices using stateless or RFC3041 mechanisms, a Stateless DHCPv6 server will be required for other (non-address) configuration options, e.g. DNS and NTP servers.

There is some concern over the use of RFC3041 addresses, due to the complexities it causes for tracking devices and knowing which network accesses are actually made by the same node over time.

13.3.1.3 Services

The component services hosted by the departmental network are:

Email

There are three MX hosts for inbound email, and two main internal mail servers. Sendmail is the MTA. POP and IMAP (and their secure versions) are used for mail access, using the UW-IMAP open

source code. There is an MS Exchange server used by up to 200 users (generally those wanting shared access to mail spools, e.g. professors and secretaries).

MailScanner is used for anti-spam/anti-virus. This uses external services including various RBLs for part of its anti-spam checking.

Successful reverse DNS lookup is required for sendmail to accept internal SMTP connections for delivery.

Hard coded IP information:

- Blocked SMTP servers (spam sources)

Web hosting

Web content hosting is provided either with Apache 1.3.x (open source) or Microsoft IIS 5.0. Common components used to build systems with are MySQL, PHP 4 and Perl 5; these enable local tools such as Wikis to be run.

Hard coded IP information:

- .htaccess and remote access controls
- Apache "Listen" directive on given IP address

Databases

All database systems are presented via a web interface, including the financial systems. In some cases, e.g. student records, ODBC-like access is required to/from from the department systems to/from the campus systems. Databases include: finance records, people, projects and publications (offered using ePrints).

Directory services

A number of directory service tools are in use:

- NIS (6 servers, all Solaris)
- LDAP
- Active Directory
- RADIUS

DNS

The three DNS servers have recently been upgraded to BIND9. A DNS secondary is held at another UK university site.

PKI

The department has at least 10 SSL certificates from Thawte, including Web-signing certificates. No personal certificates are supported by the department (though users may have their own).

NTP

The JANET NREN offers several stratum 0 NTP servers. The department also has a GPS-based NTP server built-in to its own RIPE NCC test traffic server.

USENET news

The news feed is delivered using dnews.

Multicast

There is PIM-SM IPv4 multicast via a dedicated Cisco 7206 router. This supports applications including the IPv4 AccessGrid conferencing system. A number of bugs in the Alcatel equipment prevent heavy use of IPv4 multicast within the department network (thus an IPv6 multicast solution is highly desirable). An IPv4 multicast beacon is used for monitoring multicast.

Remote login

Remote login access is offered via ssh, with sftp for file transfer. Remote use of telnet and ftp is denied by the firewall.

File serving

The main file servers are SGI systems, hosting large (multi-TB) standalone RAID arrays. The files are offered via NFS and Samba to client systems.

The content distribution server is hosted on such a system (e.g. containing MS software licensed under the Campus Agreement).

13.3.1.4 Host and device platforms

Server platforms

The following server platforms are in use in the department:

- Windows 2003 server
- Windows 2000 server
- Windows NT
- Solaris 8
- Solaris 9
- RedHat Linux
- SGI Origin 300 (Irix 6.5.x)

Desktop/laptop platforms

The following client platforms are in use in the department:

- Windows 98, 2000, ME, XP
- Linux (various flavours)
- MacOS/X
- BSD (various flavours)

PDA platforms

The following PDA platforms are in use in the department:

- Windows CE/.NET
- PalmOS
- Familiar Linux
- Zaurus

13.3.1.5 User tools/systems

The following tools or systems are used by the department's user base.

Hardware

Various dedicated systems, for example:

- Networked printers
- Networked webcams

Mail client

Various, including:

- Outlook (various versions)
- Eudora
- Mutt
- Pine

Web browser

Various, including:

- MS Internet Explorer
- Mozilla
- Safari
- Opera

Conferencing systems

The following conferencing tools are in regular use:

- AccessGrid
- A dedicated H.323 system
- MS Netmeeting

Other collaboration tools

Collaboration tools in regular use include:

- IRC
- Jabber
- MSN Messenger
- cvs

USENET news client

Various, including:

- nn
- Mozilla

Host communications

Specific tools for remote host communications include:

- X11
- VNC
- PC Anywhere

13.3.2 Transition Status

Having described the components, we now outline the steps already taken towards transition at the site. The focus here is to provide increasing IPv6 functionality in a dual-stack environment, with the goal of allowing IPv6-only devices to be introduced and to operate successfully using only IPv6 transport (that does not mean they have to interoperate with all 'legacy' services, but they should be able to use DNS, NTP, SMTP and similar basic services).

Because the Alcatel switch/router equipment does not route IPv6, an alternative method was required to deliver IPv6 on the wire to existing IPv4 subnets. To enable this, IPv6 router advertisements were delivered using a parallel IPv6 routing infrastructure. These IPv6-only routers support IEEE 802.1q VLAN tagging; the BSD routers can inject a different IPv6 prefix onto each IPv4 subnet, using congruent VLANs. The router tags the packets travelling towards the internal network with a configured VLAN ID depending on the destination IPv6 prefix/link. This VLAN method is described by the authors in [Cho04b], and illustrated in Figure 13-6.

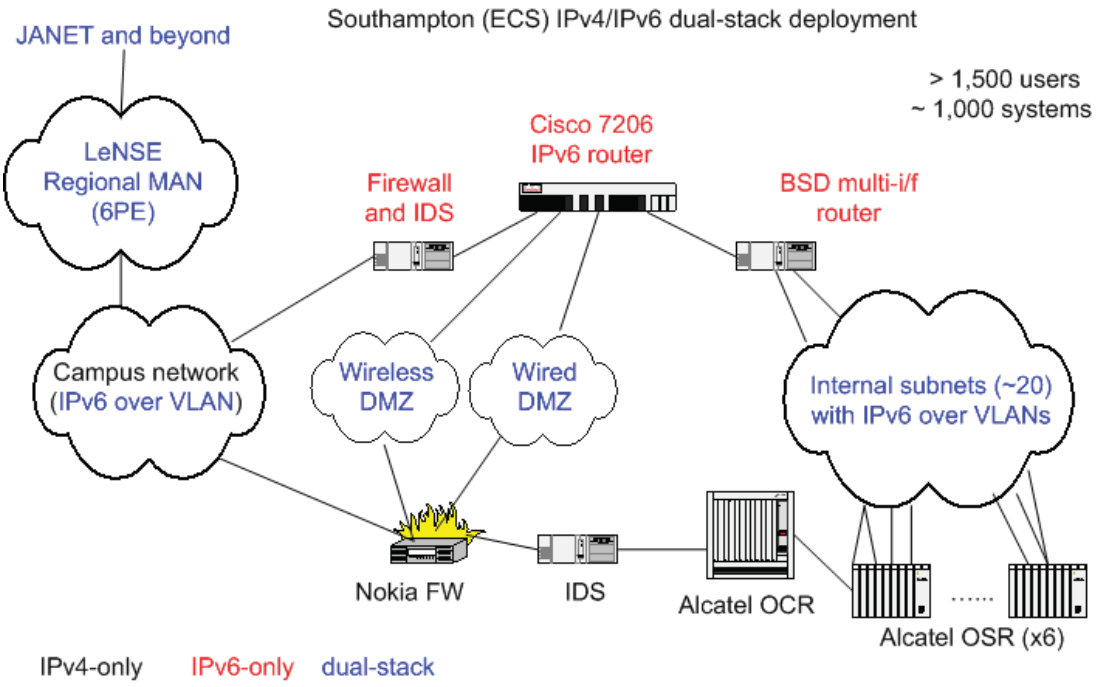

Figure 13-6 Use of IPv6 VLANs at Southampton

As traffic in the site grows, multiple routers can be dedicated to this task for internal routing, or a router with multiple interfaces. We currently have four routers with quad Fast Ethernet interfaces. BSD allows multiple-VLAN tagging per interface, so in light traffic conditions the interface count can

be collapsed. We have found, however, that our BSD systems are beginning to struggle under the growing load of IPv6 traffic.

External IPv6 connectivity was configured using a Cisco 7206, for unicast and multicast (SSM and PIM-SM), with IPv6 multicast routed internally onto the VLANs using the BSD IPv6 multicast support. The connection to the JANET IPv6 service is delivered natively through the LeNSE regional network.

The longer term plan is to use IPv6 firewalling on the Nokia IP740; until then the firewall is an additional BSD system, on which ports are blocked by default. This is a partially stateful firewall.

Two IPv6-only DNS servers have been run in the past; now the main servers network/subnet is IPv6-enabled the department's three primary BIND9 DNS servers have been IPv6-enabled. This includes reverse delegation of our prefix under 0.d.0.0.0.3.6.0.1.0.0.2.ip6.int and 0.d.0.0.0.3.6.0.1.0.0.2.ip6.arpa (the .int is being phased out). The new DNS server information is currently as follows:

```
ns0.ecs.soton.ac.uk.    390     IN      A       152.78.70.1
ns0.ecs.soton.ac.uk.    390     IN      AAAA    2001:630:d0:116::53
ns1.ecs.soton.ac.uk.    390     IN      A       152.78.68.1
ns1.ecs.soton.ac.uk.    390     IN      AAAA    2001:630:d0:117::53
ns2.ecs.soton.ac.uk.    390     IN      A       152.78.71.1
ns2.ecs.soton.ac.uk.    390     IN      AAAA    2001:630:d0:121::53
```

The main Linux login server is IPv6-enabled, with ssh logins and sftp file transfer available through the firewall. Once IPv6 is present on the wire, all that was needed was the firewall hole to be opened up for the service, an IPv6 AAAA DNS entry added for the login server, and the sshd daemon with IPv6 support turned on. Offering only secure protocols (and not plain ftp or telnet) can be easier to do when starting afresh with a new protocol.

NTP has been provisioned for IPv6 by use of both the RIPE TTM server as an NTP server, and also a dedicated NTP server from Meinberg, that supports both IPv4 and IPv6.

Our SMTP and MX servers now exchange external email over IPv6. IPv6 DNS records were added for the hosts that provide these services. If the sending or receiving node we are communicating with supports IPv6, IPv6 transport for email is usually preferred.

Almost all of our web servers/sites have been made available using Apache 2, e.g. the main department web site at www.ecs.soton.ac.uk, and many of the 100 or so hosted domains that we run, e.g. the IST IPv6 Cluster site, as operated for the 6LINK project, or the IPv6 Forum web site www.ipv6forum.org.

The department's Wireless LAN (over 30 access points) is IPv6 enabled. Some Mobile IPv6 has been deployed and tested between the WLAN and the local community wireless network (SOWN), using the MIPL code (which lacks security elements, but is usable). The more advanced WLAN network we deployed uses 802.1x based access control, which is IP version neutral and thus can be used to secure the IPv4 and IPv6 WLAN access.

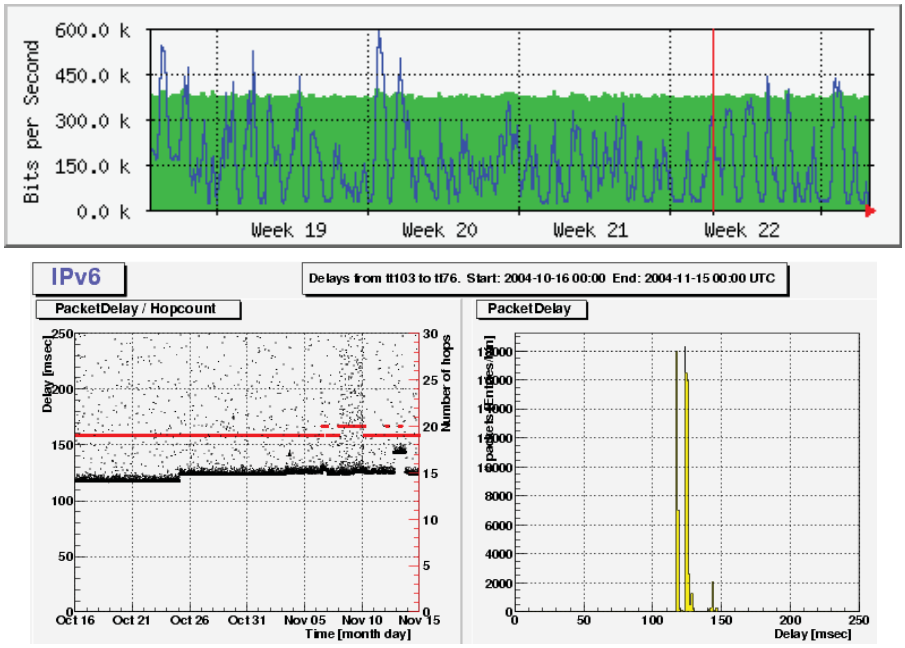

Figure 13-7 MRTG Monitoring Surge Radio Node (top) and RIPE TTM view (bottom)

Monitoring is achieved by a mixture of Netflow v9, Nagios, our RIPE TTM server and MRTG. An example of MRTG doing host monitoring and RIPE TTM output is shown in Figure 13-7.

The latest versions of Radiator and FreeRADIUS allow IPv6 transport for RADIUS.

Dual-stack Jabber and Internet Relay Chat (irc) servers are deployed.

An H.323 IPv6 conferencing system has been tested (GnomeMeeting for Linux), and we host a dual-stack Open H.323 MCU server for videoconferencing.

A key point to emphasise is that in making the transition to support IPv6 dual-stack pervasively in our network, we have not seen any adverse affect on IPv4 services. Making our DNS, MX and web servers all able to serve data over IPv4 or IPv6 has not had any noticeable adverse impact on robustness or reliability.

13.3.3 Supporting Remote Users

We offer a tunnel broker and a 6to4 relay for remote users, e.g. students in home networks or staff in wireless hotspots at conferences. At present tunnel broker users connecting get either a single IPv6 address or a /64 prefix. To offer more, in particular a /48 from our tunnel broker, we would need to use our RIPE membership to obtain LIR status and get a larger 'ISP' prefix.

We are assisting UKERNA to deploy a national academic IPv6 tunnel broker, using the Hexago broker which includes TSP support.

13.3.4 Next Steps for the Transition

Much has been achieved, but there is also still work to be done. Some imminent next steps include the following:

- A DHCPv6 service is required; implementations will be tested in the near future (from Cisco and Lucent). This will enable IPv6 DNS resolver discovery for hosts, but also be used for address allocation.
- Work on modifications to Snort to allow IPv6 IDS functions to be used.
- A dynamic DNS service is required for statelessly configuring hosts.

We are also deploying new network infrastructure over the summer of 2005, with a Cisco solution selected. One Cisco 6509 and around 30 Cisco 3750's will be deployed, with full dual-stack IPv6 support for unicast and multicast. We have already performed tests, e.g. for MLD snooping for IPv6 on the 3750's (with an advanced EFT image).

13.3.5 IPv6 Transition Missing Components

In the study for transition, we have identified a number of missing (unavailable) components for IPv6 transition, including:

1. No IPv6 Layer 3 functionality on the Alcatel OSR/Omnicore equipment (this will be resolved when the new Cisco IPv6-capable equipment is deployed in July 2005);
2. Lack of NFS/Samba IPv6 support;
3. Lack of MS Exchange, Outlook or Eudora IPv6 support;
4. No IPv6 intrusion detection system;
5. No IPv6 support for Active Directory;
6. No IPv6 dnews, so one would have to use inn as a Usenet news server;
7. Lack of supported IPv6 for Windows 98/2000/ME;
8. Lack of supported IPv6 for Irix;
9. Lack of supported IPv6 for various PDA platforms;
10. No method available to offer reverse IPv6 DNS for sendmail to verify autoconfiguring hosts (prepopulating a 64 bit subnet space is a problem, some wildcard method is required);
11. No available IPv6-enabled X11 (there is an xfree but it is encumbered by an unpopular copyright statement that most distributors find unacceptable).

This list is by no means a show-stopper for IPv6 deployment. Indeed, we take the view that by deploying IPv6 we are enabling a new environment. Some nodes will be able to take advantage of the new environment, and new applications that support IPv6, while others will happily still use 'legacy' IPv4 for basic functions such as email and web interactions.

The overall transition experience has been very positive to date.

13.4 University Deployment Analysis (Lancaster University)

IPv6 has been deployed within the Computing Department at Lancaster University for several years as part of our ongoing research activities and during that time, external connectivity has been provided via configured tunnels over the UK academic network to ULCC. However, as part of a recently initiated network upgrade process, Lancaster University's network operations group, ISS (Information System Services), decided to begin the native deployment of IPv6 across the main campus with the assistance of the Computing department. Indeed, one important motivator for IPv6 deployment at this time was due to the Computing Departments involvement in the 6NET project which represented a good opportunity to get useful additional support throughout the process.

Following the 6NET IPv6 management seminar in October 2004, which was attended by members of ISS, we began the collaborative process of planning the deployment of a dual stack IPv6 network. The goal of this was initially to provide native IPv6 connectivity to the Computing Department in the new InfoLab21 building before extending it to offer the service on a campus-wide basis. This type of widespread IPv6 deployment impacts on all aspects of the network from the underlying infrastructure to the services and applications run over it and so one of our first tasks was to conduct an analysis of what technologies will be affected by the deployment of IPv6 and what new or existing services will be needed. Once that was completed, the initial IPv6 deployment began and now that this too has been completed, we are ready to begin considering moving on to the next stage in the deployment process.

The aims in writing this section are therefore both to document the progress that has been made to date and to highlight the major issues that have been encountered and what solutions have been used to overcome them. In many cases, the issues that arose were the result of missing IPv6 support within deployed hardware or software and so we can (hopefully) expect this to be largely resolved in the near future as the wider deployment process continues. They are all currently however important considerations for those contemplating IPv6 deployments in the Enterprise, and particularly the campus scope, at this point in time.

13.4.1 IPv6 Deployment Analysis

This section presents a summary of the analysis that was conducted prior to IPv6 deployment at Lancaster and will begin with a brief discussion of our starting point in terms of the existing IPv6 deployment and production network. Thereafter, it will summarise the components of a production IPv6 network that were considered as essential for the deployment in the University campus. This includes areas such as addressing, network services, OS support and IPv6 application status.

13.4.1.1 Starting Point

As discussed briefly in the introduction, IPv6 has been used within the Computing Department at Lancaster University since around 1995 -1997. While this was primarily used as a research tool, it was also available both in the lab and office networks and over the departmental wireless network on a semi-production basis. The address space used for this was a subnet of the university prefix allocated to R&D activities. External connectivity was provided via configured tunnels to the 6Bone and on to other interested UK parties, including ULCC, UCL and the University of Southampton. This was coordinated under the Bermuda2 collaboration (www.ipv6.ac.uk/bermuda2/).

This deployment was still largely in place prior to the formal deployment of IPv6 across the University campus and so it is anticipated that when this occurs, the configured tunnels will be closed and the renumbering to production addresses (where appropriate) will take place.

At the outset of this project, the regional MAN, C&NLMAN (Cumbria & North Lancashire MAN) (http://www.canlman.net.uk/), and the University itself were both IPv4-only and the JANET Access Point had yet to be upgraded. As such, the most obvious IPv6 deployment approach was for it to begin at the innermost point in the network and be rolled-out to the edge. This gives us quite an interesting perspective to IPv6 deployment to see the progress from core to edge.

The deployment of IPv6 in the University campus was seen by ISS as part of a much larger network upgrade process moving away from a legacy 'flat' (ATM) network to a hierarchical subnet-based layout using VPNs. As such, the aim is that IPv6 deployment will ultimately be deployed in parallel to the new (IPv4) network giving a modern dual stack layout.

13.4.1.2 Addressing Considerations

One of the first aspects of the IPv6 deployment to be tackled was addressing both in terms of the allocation plan and the mechanism used. While an allocation plan had been in place for some years, the opportunity to refine this was taken and the format was changed to reflect a more logical allocation method over a physical plan based on buildings or departments to simplify management. A summary of the new plan is presented below.

Basic Configuration

Prefix allocated by JANET:	2001:0630:0080::/48
Special Addresses:	2001:630:80:0000::/64 - Reserved (DNS, services, etc.)
	2001:630:80::1 Router
	2001:630:80::4 DNS - primary
	2001:630:80::7 DNS - secondary

General Address Format

Other than the above special addresses, Lancaster University IPv6 addresses observe the following format:

<48 UNI> <3 Res> <1 Site> <12 Subnet> <64 Host>

where:

<48 UNI>	- 48 bit University prefix 2001:630:80::/48
<3 Res>	- 3 bits reserved for future aggregation
<1 Site>	- 1 bit identifying the site of the network (provisional)
<12 Subnet>	- 12 bits for site subnet
<64 Host>	- 64 bit host identifier

IPv6 Subnet Format

One issue that was of particular concern for ISS throughout was that flexibility be incorporated into the design in order to minimise the need for extensive or complicated reallocation in the future. As such, in addition to the high-order bits reserved for future use, the subnet allocation methodology follows whereby bits closest to the boundary between groups are allocated last in order to maximise the potential for reallocation.

Production subnets will use their 12 bits subnet allocation in the following way:

<4 Cat> <8 Physical>

The category group makes up the first 4 bits of the subnet identifier and is reserved for allocating subnets based broadly according to the roles they perform. An example of this would be to allocate subnets for office, public or wireless networks which can each then be further divided in the physical group.

The physical group (8 bits) can then be used to add geographical identifiers to the category. These will, for example, define the different colleges, academic departments and physical buildings that make up the university campus. In practice of course, there are likely to be multiple subnets allocated to larger buildings or departments as necessary.

Address Allocation Mechanism (DHCPv6 vs. SLAAC)

The DHCPv6/SLAAC (stateless address autoconfiguration) issue is still one that is a matter of continued debate among the community at large. After consultation with other 6NET members however, it was determined that there is a strong case to justify the ISS position that stateful DHCPv6 is the most suitable candidate for managed IPv6 address allocation in an enterprise-type network. The benefits of a centralised address allocation and service discovery method for network management make DHCPv6 a better choice than SLAAC for this type of network.

As such, the majority of Lancaster's IPv6 address allocation requirements will be handled via stateful DHCPv6 but there will be limited use of manually configured addresses (as in the special cases, see above) and SLAAC in certain cases and environments that favour it, such as wireless or mobile networks.

13.4.1.3 Operating System Considerations

Lancaster University uses a variety of Operating Systems including Microsoft Windows, Linux and Unix variants. In all cases, various versions are in use ranging from the most current, incorporating native IPv6 to older versions that will need to be upgraded as production IPv6 is not and will not be supported.

Microsoft

A range of Microsoft Operating Systems are in use ranging from NT4.0 to XP and Server 2003 and as such the degree of IPv6 support is obviously variable. Windows NT4.0 has no IPv6 support, Windows 2000 has no commercial IPv6 support but XP and Server 2003 does and so some upgrading may be necessary in this case.

Linux

The main versions of Linux in use at Lancaster are Redhat 7.3 – 9.0 and Fedora and so IPv6 support is generally available in most cases out of the box. IPv6 support was introduced in Linux from kernel version 2.4 onwards and since the versions in use range from 2.4 – 2.6, this should introduce no major problems.

Solaris

Solaris is the main UNIX variant used at Lancaster from version 2.5.1 right through to 9, but mainly version 7. Newer versions will not be seriously affected since IPv6 support was formally introduced from version 8 onwards. In older versions however, despite IPv6 patches being available, it would be highly preferable to upgrade to newer versions before IPv6 support is considered.

13.4.1.4 Core IPv6 Services

DNS

The Lancaster University DNS servers now run BIND9 and so can be used to provide native IPv6 transport (http://www.isc.org/index.pl?/sw/bind/). The deployment of IPv6 enabled DNS is now in progress and should be completed in the near future.

DHCPv6

As already discussed, ISS has made the decision to deploy stateful DHCPv6 to handle IPv6 address allocation in parallel with the existing DHCP service. Given the relatively 'new' nature of DHCPv6, the level of support is currently low but there are some implementations available including Sourceforge (open source), NEC, Dibbler, Cisco and HP. However, there are issues that arise with each and in any case ubiquitous DHCPv6 client support (in MS Windows for example) is some way off.

This is an important part of our IPv6 deployment plan and so we expect to begin experimentation and trialling of the various DHCPv6 implementations in the near future.

13.4.1.5 Network Services

Email

The University runs a number of email services based primarily on Exim, MS Exchange and Pine IMAP. Unfortunately, the availability of IPv6 over these services is variable. Exim has included IPv6 support for some time and all versions that are in use at Lancaster are IPv6 enabled. Pine IMAP has rudimentary IPv6 support provided via a patch but should not be considered production level. MS Exchange has no IPv6 support at this time.

Web Services

Lancaster utilises two key web services, Apache servers and web proxying via Squid. The newer versions of Apache, 2.x onwards, feature native IPv6 support but 1.3.x is also still in use and only supports IPv6 via a non-production patch (http://www.apache.org/). As such, some upgrading may be necessary to deploy IPv6 fully.

The Squid developers have enabled IPv6 in the current version, 2.5 as a patch against the latest version of the CVS sources (http://www.squid-cache.org/). While this is readily available, there has been little work in developing native IPv6 functionality in Squid and so this represents a significant issue that may necessitate further consideration.

File Sharing

Lancaster primarily uses MS Windows-based file sharing using Active Directory. Unfortunately this is not currently IPv6 enabled and so again introduces a significant issue when IPv6 deployment reaches this stage. There are no plans to upgrade our file sharing system in the immediate future but until this issue is resolved, it represents a significant problem for IPv6 deployment is this area.

In addition to the Microsoft service, Samba and NFS are available for non-Windows systems including Linux, UNIX and Solaris. To date IPv6 support is not included in either of these services.

Other Services

The above obviously only represents a small range of the services and applications deployed at Lancaster University and once the initial IPv6 deployments are complete a more thorough survey will be completed prior to the production roll out across campus.

Areas still to be addressed at this stage include security, management and remote access to name a few. There are now however many useful sources that list the current status of IPv6 applications including the deepspace6 site:

http://www.deepspace6.net/docs/ipv6_status_page_apps.html

13.4.2 IPv6 Deployment Status

This section will discuss in detail the IPv6 deployment efforts undertaken to date. This is concerned with first stage of IPv6 deployment which covers the rollout of native IPv6 connectivity from the JANET core network, across C&NLMAN, and out over the University backbone. As such, the majority of this section will be concerned with the issues encountered at the network level, enabling native IPv6 connectivity down to the Computing Department. The final part of this section however will summarise all the issues that have been encountered to date in our IPv6 deployment efforts.

13.4.2.1 Native Connectivity to the Computing Department

Lancaster University is connected to the UK academic network (JANET) via the C&NLMAN regional network, also administered in Lancaster. At the outset, this did not offer native IPv6 but was intended to be made dual stack as a precursor to a deployment within the University. This presented two distinct sets of problems, enabling IPv6 support within C&NLMAN and the subsequent deployment in the University campus. Since the focus of this section is the Enterprise/campus scope, this section focuses primarily on the later except in cases where the deployment of IPv6 in the provider network has implications for the campus deployment. We will first describe the basic topology of the network to outline how this was affected by the deployment of IPv6 and what modifications were necessary to achieve it.

Starting Point

The situation at Lancaster differs from that in other Enterprise networks because in our case ISS is responsible for the management and operation of both the regional MAN and the University network, thus reducing the number of parties involved in the process to two, the Computing Department and ISS. We exclude UKERNA here (who manage the JANET network) because while they are responsible for UK academic IPv6 prefix allocation and core deployments, they were not directly involved in the planning of our IPv6 deployment. This obviously has many advantages and has undoubtedly led to a reduction in the amount and time and effort spent during the planning the deployment process.

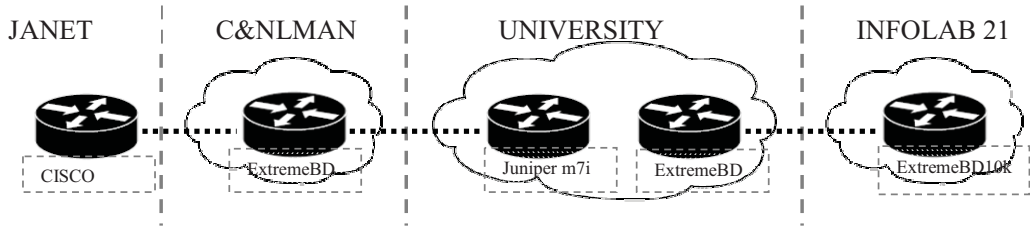

Figure 13-8 Basic Configuration of the Upgrade Path

As such, while the deployments in C&NLMAN and the University were considered separately (including addressing and routing considerations), the work was conducted almost in parallel, again leading to a reduction in the amount of time taken during the deployment. The original configuration

of the network from the UKERNA access point to the InfoLab21 entry-point is shown below (Figure 13-8):

In our case, hardware is supplied from a number of vendors including Cisco, Juniper and Extreme and each device was first tested in a sandbox environment to both determine the extent of IPv6 functionality offered and to allow ISS to familiarise themselves with this aspect of their operation. Unfortunately, while the Juniper and Cisco equipment proved very capable of supporting IPv6, the Extreme devices did not, introducing problems at both the MAN and the University level. In the case of the Computing Department, a BlackDiamond 10k has recently been deployed on which limited IPv6 support is available and full support is promised shortly, unfortunately however the older BlackDiamond 6k devices deployed across the rest of the infrastructure do not currently offer IPv6 support and are unlikely to do so in the immediate future.

13.4.2.2 The IPv6 Deployment Process

An incremental approach was adopted by ISS for IPv6 deployment which involved the upgrade of first the UKERNA Access Point and the C&NLMAN infrastructure before moving on to address the University network.

Provider Network Deployment

With the UKERNA equipment being beyond our control (and upgraded separately: http://www.ja.net/development/ipv6/statustable.html), the most pressing issue was that of C&NLMAN connectivity. ISS determined that the most effective way to overcome the problem given the issues with Extreme networking equipment was to modify their existing configuration by introducing an additional device upstream of the BlackDiamond 6k to handle L3 (i.e. routing) functionality leaving the former as purely a L2 switch for which it is better suited in any case. The new device, a Juniper m7i, now handles the IP routing and is well suited to this role (with hardware support) thereby neatly sidestepping the problem. This process was completed as of December 2004 and so from that date native IPv6 connectivity was available into the University network.

University Backbone Deployment

The University upgrade followed a roughly parallel path to that of C&NLMAN as the infrastructure is similar. As such, the main issue faced was with the Computing Department employing a BlackDiamond 10k and so while the L3/L2 approach was followed across the majority of the University backbone, this device could not necessarily be circumvented in the same way.

Ironically, the BlackDiamond 10k claimed to have full IPv6 functionality but when this was explored by ISS, it soon became clear that a dual stack configuration opens significant security holes in the IPv4 stack and so was not a viable option (IPv6-only operation however was fine). An updated version of the firmware for this device is currently with ISS who are in the process of testing it ahead of the release and use within the network.

As a result, it was decided that the best solution to this was a temporary direct link into the Computing Department based on a Cisco 7206 that was available. This was connected directly to the Juniper m7i at the university BAR (Border Access Router) to circumvent the existing infrastructure until the remainder of the network (including the BlackDiamond 10k) has been upgraded and production IPv6 can be supported. This was achieved recently and as of April 2005 native IPv6 via the new infrastructure is now available into the Computing Department on a research basis. In the near future this will be used to replace the existing tunnelled connectivity, and be made available department-wide. Figure 13-9 shows the current configuration of the university infrastructure.

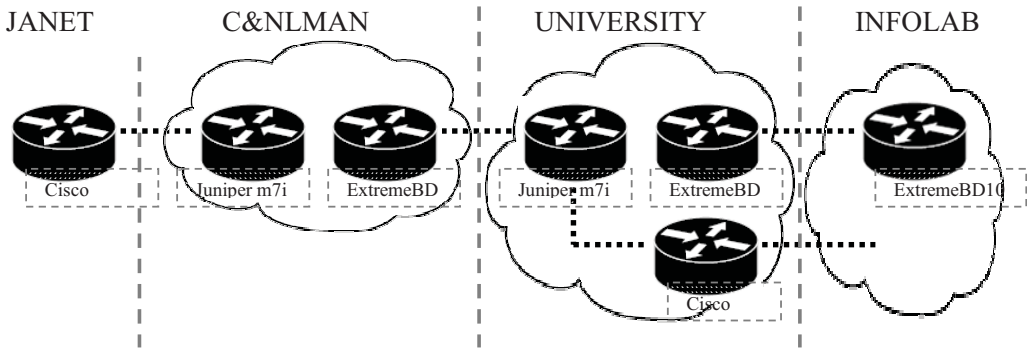

Figure 13-9 Alternative Configuration Providing Native IPv6 to the Computing Dept

Finally, as part of the continued upgrade process, the BlackDiamond 6k within the University network has now been replaced by a 10k. As such, once suitable firmware is available for the BlackDiamond 10k devices, the university infrastructure will be fully capable of supporting production dual stack from end-to-end across the backbone.

13.4.2.3 Outstanding Components

In order to complete an IPv6 deployment that is equivalent to the existing IPv4 infrastructure, a number of missing 'pieces' in the IPv6 deployment need to be upgraded to support IPv6. Unfortunately, this lack of support applies to every level of the deployment, from hardware to network services and applications. This section therefore provides a summary of the outstanding network components that need to be enabled with IPv6 or circumvented to allow the campus deployment process to continue.

In terms of hardware, the lack of IPv6 support on the BlackDiamond equipment was the most significant problem that we faced. This was largely overcome through the use of available Cisco and Juniper equipment to perform L3 routing and upgrades in the BlackDiamond 10k firmware. As a result, this problem has largely now been overcome.

The core services, DNS and DCHPv6 are an essential part of the IPv6 deployment process and as such the lack of proper DHCPv6 represents a significant problem. IPv6 enabled DNS deployment is now underway within the University but the lack of widespread DHCPv6 support cannot be overcome at this time. This should again 'fix' itself as more implementations become available and proper support is incorporated into more products over time but at present, this represents perhaps the most significant immediate issue to IPv6 deployment at Lancaster.

Other significant issues include the lack of IPv6 support in other important network services and software. This level of support is mixed but perhaps the most pressing issue is the lack of proper IPv6 support in mainstream Microsoft products such as MS Exchange and Active Directory. In general, while the level of IPv6 support is quite solid in Microsoft Operation Systems, the applications and services that run across them are much less so. Other issues include the lack of file sharing / NFS support and the lack of IPv6 support in web caching software such as Squid among others.

13.4.3 Next Steps

There is still a long way to go until a 'full' IPv6 deployment has been completed at Lancaster, both in terms of supporting the key network services that would be expected in such a network and providing

connectivity to the number of users necessary to make IPv6 available at an Enterprise level. As such, this section will outline the major steps that still need to be completed in order to complete the IPv6 deployment process.

This will begin by summarising the current state of IPv6 deployment and highlighting the, as yet, missing pieces of the deployment either due to one of the issues discussed in the previous section or because it has simply yet to be addressed. Thereafter we will focus on discussing our ongoing deployment from the immediate next steps to the ultimate long-term goals.

13.4.3.1 Current Status Summary

The current state of IPv6 the deployment at Lancaster is that native IPv6 has now been enabled all the way down to the Computing Department and is now being used in a number of R&D activities. This is significant not least because it has necessitated upgrading the connectivity all the way from the regional network core, through the University infrastructure down to the Computing Department.

With this now complete, the continued rollout of native IPv6 over the campus is now possible and it is also being used to replace the existing tunnelled infrastructure in the Computing department. Additionally, IPv6 services are expected to come online shortly starting with DNS (and DHCPv6 when possible) and moving out from there.

13.4.3.2 Deployment Goals

With the initial IPv6 deployment now completed, the next steps will be to continue to rollout IPv6 over the rest of the campus and extend its functionality to provide a more complete set of services. This will proceed in three parts, the first of which will complete the deployment of IPv6 in the Computing Department, making it widely available in both the office and research networks. The second will continue to upgrade the University's IPv6 deployment both in terms of the physical deployment over campus to interested groups and departments beyond ourselves and upgrading network services to make them fully IPv6 compatible. Finally, as a long term goal, there must be an effort to offer IPv6 in a ubiquitous manner across the entire campus network and use this to deploy new IPv6 services such as multicast, QoS and mobility.

Immediate Goals

The immediate goals are those related to completing the next stages of IPv6 deployment and include the wider deployment of production IPv6 in the Computing department and the provision of the core IPv6 services such as DNS. With the general deployment of IPv6 in our department, the first use of production IPv6 address space is now occurring and as such represents a significant deployment milestone.

Completing the Campus Rollout

This stage of IPv6 deployment will involve completing the deployment of IPv6 dual stack over the campus backbone by resolving the outstanding issues (particularly with the BlackDiamond 10k) along with the upgrade of essential networking services such as DNS to support the deployment, this has already been largely completed. This provides a basic level of connectivity and will allow production access to the IPv6 Internet. Thereafter, there are will be two parts of the continued deployment, the provision of more IPv6-enabled services and the continued rollout of IPv6 connectivity to other interested groups and departments.

Now that the core IPv6 infrastructure has been deployed, connectivity can be extended on an 'as necessary' basis, initially only to those who request it. There is no obvious limit to the potential rollout of IPv6 to the rest of the campus and so it becomes a matter of policy as to when this actually takes place. The continued provision of IPv6 services will initially include those outlined such as email,

www access, file sharing, etc. The goal here will be to continually dual stack enable networking services over time to eventually provide an IPv6 service equivalent to the existing IPv4 infrastructure.

Further IPv6 rollout

With the initial rollout of IPv6 now well underway in the Computing Department, the next logical goal for IPv6 deployment within the University will be to try to improve the IPv6 user-base by increasing the number of IPv6-enabled hosts within the network. There are two logical ways in which this could potentially be pursued in our case: either by enabling dual stack IPv6 on the student network (ResNet) or in the University computer laboratories. Either (or both) of these can be done relatively easily depending on the requirements of ISS for managed IPv6 deployment.

If it is determined that initial IPv6 deployments should be managed then the lab networks are the natural choice as they are heavily managed by ISS whereas the student networks are much less so and so IPv6 can be enabled but left to the students to use IPv6 if they wish. Since the majority of hosts on the both networks will be MS Windows-based, the majority of new IPv6 users here are likely to be Windows XP-based but in the student network case ISS should be prepared for both Linux and Unix support also. Either of these deployments would potentially introduce several hundred new IPv6-enabled hosts to the network (at a conservative estimate) and so would greatly increase the scale of IPv6 deployment within the University campus.

Protocol Selection Policy Definition

With a significant IPv6 deployment in place, an issue that gains relevance at this point is in defining how IPv6 is used within IPv6-enabled hosts. The conventional choice within IPv6 early-adopters is currently not to use IPv6 as the first-choice protocol except in explicit special cases where IPv4 is not supported. This should certainly still be the case in the majority of our deployments where the network is dual stack and IPv6 is not yet able to offer an equivalent service to the existing IPv4 network.

Ultimate Goals

Once this is complete, IPv6 will be available across the entire network and so the ‘deployment’ will be largely complete. As such, it will then be possible to consider what other services could potentially be deployed across the infrastructure and while it is too early to define these with a great degree of confidence, multicast and mobility are both strong candidates for this as they feature strongly as IPv6-only applications that would showcase the potential of the new protocol.

13.4.3.3 Other Considerations

Beyond the deployment of IPv6 connectivity and services at Lancaster, there are a number of other considerations that must be made alongside it. The prime example of this is the operation and management aspects of the network that are necessary to ensure the safe and secure operation of both the IPv6 and IPv4 networks. In addition to this, another consideration that has yet to be fully addressed are the transitioning requirements of the IPv6 deployment and this will be considered here also.

Network Operations and Management Services

This area of the deployment actually encompasses a wide range of networking applications and technologies that support the management and safe operation of the network. Obviously, the impact of IPv6 deployment will be significant in this area and so it is necessary to consider each aspect of this in turn. This has yet to be fully assessed from the perspective of IPv6 deployment but summary of these areas include:

Firewalling

The addition of IPv6 on the campus backbone will necessitate the provision of IPv6 firewalling in addition to the existing IPv4 firewalls, this is already underway.

Security

This includes both intrusion detection and snooping software and both will need to be upgraded to support IPv6.

Management and Monitoring

A range of management and monitoring services are employed within the University and these will need to be upgraded to support IPv6. This process is also currently in progress.

Secure Remote Access

Remote access to the university is provided via a VPN server that is available to both staff and students. While the campus is still dual stack, this will in all likelihood not be upgraded but an IPv6 VPN service may be provided in the future.

Transitioning Approach

Due to the IPv6 deployment at Lancaster being dual stack, specific transitioning solutions are not necessary at this point in time. Indeed, our existing tunnelled connections will shortly be replaced with native production links thanks to the deployment efforts made to date and since this is going to be deployed from end-to-end, interoperation mechanisms are also largely unnecessary at this stage.

There has however been some discussion of specific scenarios where IPv6-only connections may be needed at some point in the future and as part of our deployment study it would be wise not to rule this out as a possibility. As such, we will briefly outline the mechanisms that may be considered for deployment at Lancaster in the future, either to offer connectivity to isolated IPv6 subnets or to provide IPv4 interoperation to IPv6-only subnets.

To provide IPv6 connectivity, a number of tunnelling mechanisms may be deployed depending on the circumstances. Internal tunnelling will be largely unnecessary since the campus supports dual stack but external connectivity may be supported via 6to4 or perhaps tunnel brokers. Likewise, since IPv6-only networks have not been deployed to date, interoperation mechanisms are not necessary but in the event that they do become deployed, TRT (with a suitable range of ALGs) or DSTM seem the most obvious solutions.

13.5 Other Scenarios

The above scenarios and 'case studies' are reports of ongoing IPv6 deployments at 6NET partner sites. In this section we give brief commentary on other academic-related scenarios, to allow the reader to gain some insight into how we feel solutions will be chosen from the quite large toolbox that is available to a network administrator when faced with these situations.

13.5.1 Early IPv6 Testbed on a Campus

There are two variants of a scenario where a campus wishes to gain early IPv6 experience before a wider adoption. We assume the campus does not have native connectivity to its NREN or regional network (upstream provider).

13.5.1.1 An IPv6 laboratory or single IPv6 subnet

In this case, a single router can be deployed in the testbed subnet and a tunnel used for connectivity from the testbed to an upstream IPv6 provider.

The choice for the site is which tunnel mechanism to use. The options include:

- Manually configured tunnel
- Tunnel broker
- 6to4

While 6to4 is convenient and automatic, it can lack reliability (depending on the 6to4 relays being used). It also means the site does not use production address space, or its own allocated address space. A manual tunnel or tunnel broker would generally be preferred.

If the site is using IPv4 NAT, a tunnel can still be established, but may need specific forwarding of (for example) Protocol 41 on the NAT device, or use of a protocol such as TSP to establish a NAT-friendly tunnel method such as UDP tunnelling.

13.5.1.2 IPv6 hosts scattered around the campus

As per the above case, a router should be deployed with tunnelled uplink/external connectivity.

The decision then is how to connect the sparsely deployed hosts. Where VLANs are available, they could be used to distribute IPv6 by Layer 2 segregation to specific physical ports on the network edge. An alternative is to use ISATAP as an automatic tunnelling solution, or an internal tunnel broker. The tunnel broker and VLAN solutions are more manageable, e.g. the broker can include authentication, and thus may be preferred where a more controlled deployment is desirable.

13.5.1.3 IPv6-only testbed

If the testbed is IPv6-only, then a different approach is required. The reader is referred to the Tromso scenario described earlier in this section.

If there are dual-stack nodes in the IPv6-only testbed, DSTM may be an appropriate solution.

13.5.2 School Deployment of IPv6 to Complement IPv4+NAT

This scenario has been covered in depth by the 6NET report on the IPv6 Deployment in the Greek School Network [D5.14]. Here, IPv6 with global addressing is adding new possibilities to an infrastructure previously focused on IPv4 + NAT.

13.5.3 IPv6 Access for Home Users

Here we consider home user networks (staff, students) wanting connectivity, possibly behind a NAT, possibly with just a single machine, possibly with a dynamic IPv4 address.

The solutions available include

- Teredo
- Tunnel broker
- 6to4

Some tunnel broker implementations, especially with TSP, can support IPv4 NAT traversal and/or dynamic IPv4 addresses. As described above, 6to4 can have reliability issues dependent on the relay location (but is strong when used to connect just to other 6to4 sites, e.g. other home networks using it). Teredo is a method of 'last resort', designed for use behind IPv4 NATs. The broker or 6to4 would generally be preferred to Teredo.

13.6 Summary of Unexpected Results and Unforeseen Difficulties

The case studies above discuss the general scenarios and deployment procedures that were followed in each case. In each case, the challenges in deployment were generally technical but understood, rather than 'exploratory'. Thus very few, if any, unexpected results or difficulties were encountered. Quite the opposite – the deployments described have gone very well and have been very positive.

Here we repeat a small summary of the unforeseen issues captured in the scenarios above:

- A surprising level of fear, uncertainty and doubt (FUD) about the IPv6 technology by certain administrators that we had to interact with.
- The fact that if you want to run a tunnel broker service to your staff or students, a /48 prefix is not big enough for a campus unless the broker users only get a /64 prefix, which is not ideal.
- Lack of VLAN capability where it was desired, preventing the VLAN method being used where desired in all cases.
- Variable performance in software driven VLAN tagging; Southampton's BSD routers hit a forwarding limit with the specific DLINK cards used.
- A desire by network administrators (managers) to use DHCPv6 even when IPv6 supports stateless autoconfiguration. Sites seem used to managing their IP address assets.
- A desire to disable IPv6 Privacy Addresses to make management and monitoring simpler (device identification simpler).
- Unexpected impact on IPv4 security when IPv6 is turned on. This happened in the Lancaster case, but is resolved via a firmware upgrade
- The relative high performance of DSTM versus NAT-PT.

- Various issues with totd at Tromso – interactions with certain DNS implementations, with use of IP literals, and systems with AAAA records that do not offer all services over IPv6 transport.
- Some hardware problems to run IPv6 in old PCs, caused by lack of required multicast support in old Ethernet cards. These can be replaced cheaply.
- Abuse of some transitioning relay services. The spammers will find you.

Although these have been overcome, they may represent generic issues that may resurface in similar situations.

For a discussion of issues and challenges remaining for IPv6 deployment, see [D2.5.3].

13.7 Summary of tradeoffs made in solutions chosen

Many tradeoffs are related to the desirable properties of the mechanisms being chosen from, for example:

- Which IPv4 and IPv6 functionality/connectivity is required?
- The cost of deploying dual-stack, versus IPv6-only with translation to communicate with IPv4 devices
- The cost to 'the Internet architecture' of a certain method, e.g. are relays expected to be deployed by every ISP, or even every site? How do these relays interact?
- Security requirements – is the method open to relay abuse?
- Ease of management – how is the transition mechanism managed?
- Whether dynamic IPv4 addresses on the client need to be handled
- Whether IPv4 NAT traversal is needed from the client
- Whether just hosts or networks need to be connected
- How the solution scales – is there just one point of failure? Is load-sharing or balancing supported?
- Resilience in the solution
- Who owns the transition components (ISP or end site)
- How the service is discovered or configured, if 'plug and play' is desirable
- Whether specific services are enabled, like multicast
- How optimal the routing is (or isn't), e.g. whether there is any route optimisation method
- What performance penalty is suffered from encapsulation/translation methods
- Whether reverse DNS lookups are available for the method
- Accountability – how is usage reported and monitored, or identified?

There is no single, 'best' solution for all scenarios. Factors such as those listed above will determine which is best for which scenario.

Chapter 14

IPv6 on the Move

This chapter describes three case studies of 6NET partners who deployed and trialled Mobile IPv6 testbeds. We first look at the testbed at Fraunhofer Fokus and after that we describe the testbeds at Lancaster University and the University of Oulu.

14.1 Fraunhofer Fokus

The Fokus Mobile IPv6 testbed has undergone several changes throughout the 6NET project. The current environment does not longer distinguish between an internal and an external part as before when there was a local, experimental testbed not constantly connected to the 6NET network and another part with continous provision of native IPv6 connectivity to the outer 6NET world. Furthermore the test environment was completed by several components providing for VoIP and video conferencing capabilities.

The Mobile IPv6 testbed consists of the components illustrated in Figure 14-1.

Connected to the central router “adrahil” there are basically two different networks with prefixes 2001:638:806:2002::/64 and 2001:638:806:2001::/64. Attached to those networks are the corresponding Home Agents for the MIPL- and the KAME-Mobile IP implementation, respectively. A MCU connected to the “2001”-network was used as Corresponding Node for Mobile IPv6 functionality- and interoperability testing: Using a MN with video conferencing equipment, i.e. the GnomeMeeting application with camera and headset, this scenario allowed for establishing a connection to the CN-MCU and subsequently changing the IPv6 network point of attachment while maintaining the connection to the MCU.

Currently the Fokus testbed is made up of the different components:

- End systems with different operating systems

  At the leaves of the network, standard PCs with different operating systems (Windows Server 2003, Windows XP, Linux, FreeBSD) are installed. They are used for testing IPv6 network applications like video conferencing, web surfing, downloading audio and video streams, IP telephony applications etc.

- Home Agents

  Mobile IPv6 Home Agents as offered by the MIPL- and KAME project.

- IP softphone

The original BonePhone application was complemented for performance reasons by the SIP-based KPhone telephony application (also developed in the context of WP5 of the 6NET project) for use as MN.

- FreeBit hardphone for use as MN.
- GnomeMeeting video conference client for use as MN.
- H.323 OpenMCU for use as CN.

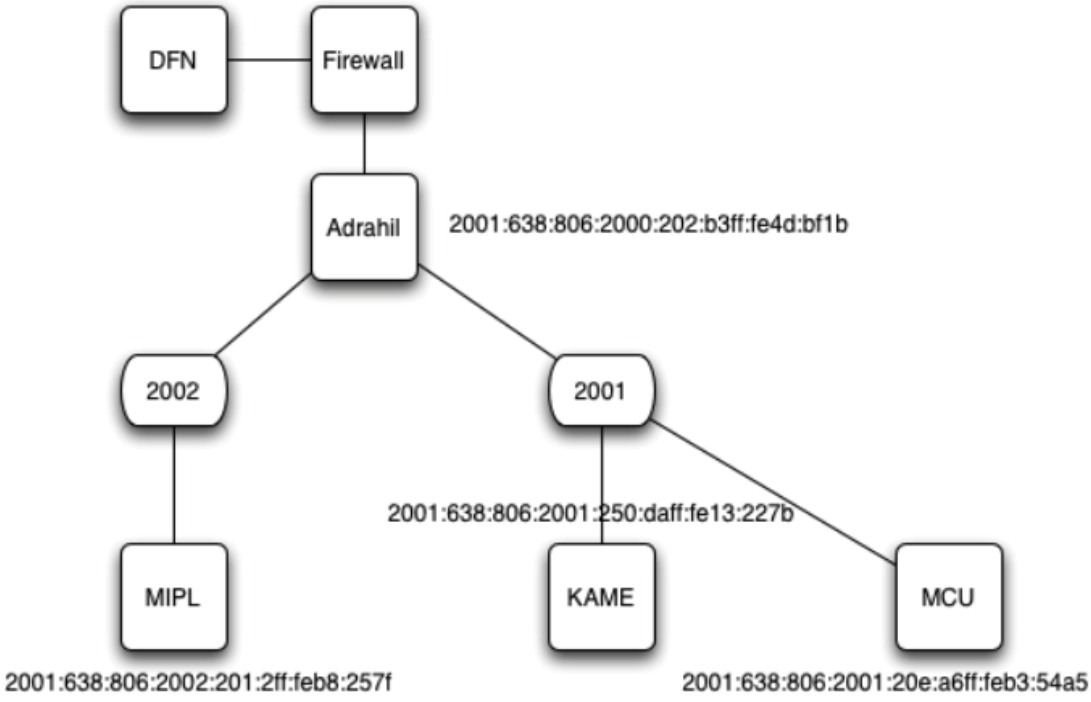

Figure 14-1 Schematic Representation of the Testbed Setup

14.1.1 MIPL-HA

The Linux HA is running version 1.1 of the MIPL implementation. The HA has the address: 2001:638:806:2002:201:2ff:feb8:257f. The home address in the network has the prefix: 2001:638:806:2002::/64.

14.1.2 Kame-HA

The KAME HA is running a snapshot dated 22/03/2004. Due to compatibility reasons you need at least version 1.1 of the MIPL implementation if you want to use the KAME HA in combination with a Linux mobile node. The HA has the address: 2001:638:806:2001:250:daff:fe13:227b.The home address in the network has the prefix: 2001:638:806:2001::/64.

14.1.3 MCU-CN

The MCU is running version 1.1 of the MIPL implementation as well, the address is 2001:638:806:2001:20e:a6ff:feb3:54a5. There are two rooms, the default room is “room101”. In order to enable checking of local audio and video settings there is a loopback room simply called “loopback”. In case of audio or video problems the MCU room “loopback” can be used to check proper working of audio and video peripherals.

14.1.4 IPSec

IPSec was turned off at the time of MIPv6 testing due to compatibility reasons. Only the KAME implementation supported it at that time.

14.2 Testbed Components

Table 14-1 Fokus MIPv6 Testbed Components

Host	System	OS	Type	IPv6 Address
Trolloc	Cisco 7206VXR[1]	Version 12.2(8)T4	Router	DFN - 6WIN - IPv6 - Network prefix usage for FhG Fokus: 2001:0638:0806::/48
adrahil	PIII/500 MHz	Debian GNU/Linux Rel. 2.4.18, dual stack	Access router IPv6	2001:638:806:2000:202:b3ff:fe4d:bf1b 2001:638:806:2001:202:b3ff:fe4d:c8e0 2001:638:806:2002:250:4ff:fe64:eb78
Mipl.lab6.fokus.fraunhofer.de			HA (MIPL)	2001:638:806:2002:201:2ff:feb8:257f
kame.lab6.fokus.fraunhofer.de			HA (KAME)	2001:638:806:2001:250:daff:fe13:227b
mcu.lab6.fokus.fraunhofer.de (MCU)		GNU/Linux 2.4.26 i686	CN	2001:638:806:2001:20e:a6ff:feb3:54a5

One can think of a scenario in which a Mobile Node is roaming only between foreign networks, never arriving home.

For offering this kind of service one needs a HA connected to the public IPv6 infrastructure. The MN’s home address has to have the same prefix as the HA’s one.

The Fokus testbed offers this kind of use of MIPv6 Home Agents for functionality and interoperability tests.

The required configuration parameters for the concerning MN are:

In case you want to use the Linux HA: MIPL HA-address: mipl.lab6.fokus.fraunhofer.de (2001:638:806:2002:201:2ff:feb8:257f) Home address in the network: 2001:638:806:2002::/64

In case you want to use the Kame HA: KAME HA-address: kame.lab6.fokus.fraunhofer.de (2001:638:806:2001:250:daff:fe13:227b) Home address in the network: 2001:638:806:2001::/64 (Note: To use the Kame HA with a Linux Mobile Node (MN) you need at least version 1.1 of the MIPL-implementation.)

[1] Cisco 7206VXR, 6-slot chassis, 1 AC Supply w/IP Software, 7200VXR NPE-400 (128MB default memory), 256 MB Memory for NPE-400 in 7200 Series, 1-Port Packet/SONET OC3c/STM1 Singlemode (IR) Port Adapter, Cisco 7200 Input/Output Controller with Dual 10/100 Ethernet, Cisco 7200 I/O PCMCIA Flash Disk, 128 MB Option, Cisco 72009 Series IOS IP

In order to support simple MN-CN connectivity/roaming tests an IPv6 H.323-MCU was connected to the “2001”-network: The address of the MCU is: mcu.lab6.fokus.fraunhofer.de (2001:638:806:2001:20e:a6ff:feb3:54a5) There are two conferencing rooms, the default room is "room101". For the convenience of checking local audio settings an audioloopback room called "loopback" was set up. If you encounter any kind of audio problems in "room101", please use the room "loopback" first to check if your local microphone and speaker settings work correctly.

There is no security association or any other restriction whatsoever. Any Mobile Node in the network above should be able to use these HAs.

14.3 Lancaster University

Lancaster University has had a Mobile IPv6 testbed in one form or another since 1997. In those early years we used Lancaster University's own implementation running on Linux. Since then, more extensive research efforts (some with the help of Lancaster University to port their code to different platforms) have spawned other implementations which are much more up to date and robust than the Lancaster University implementation. In our MIPv6 testbed for 6NET we have thus not considered the Lancaster University implementation since it is now not interoperable with more recent implementations.

14.3.1 The Testbed

As simplified diagram of the Lancaster University MIPv6 testbed is shown in Figure 14-2.

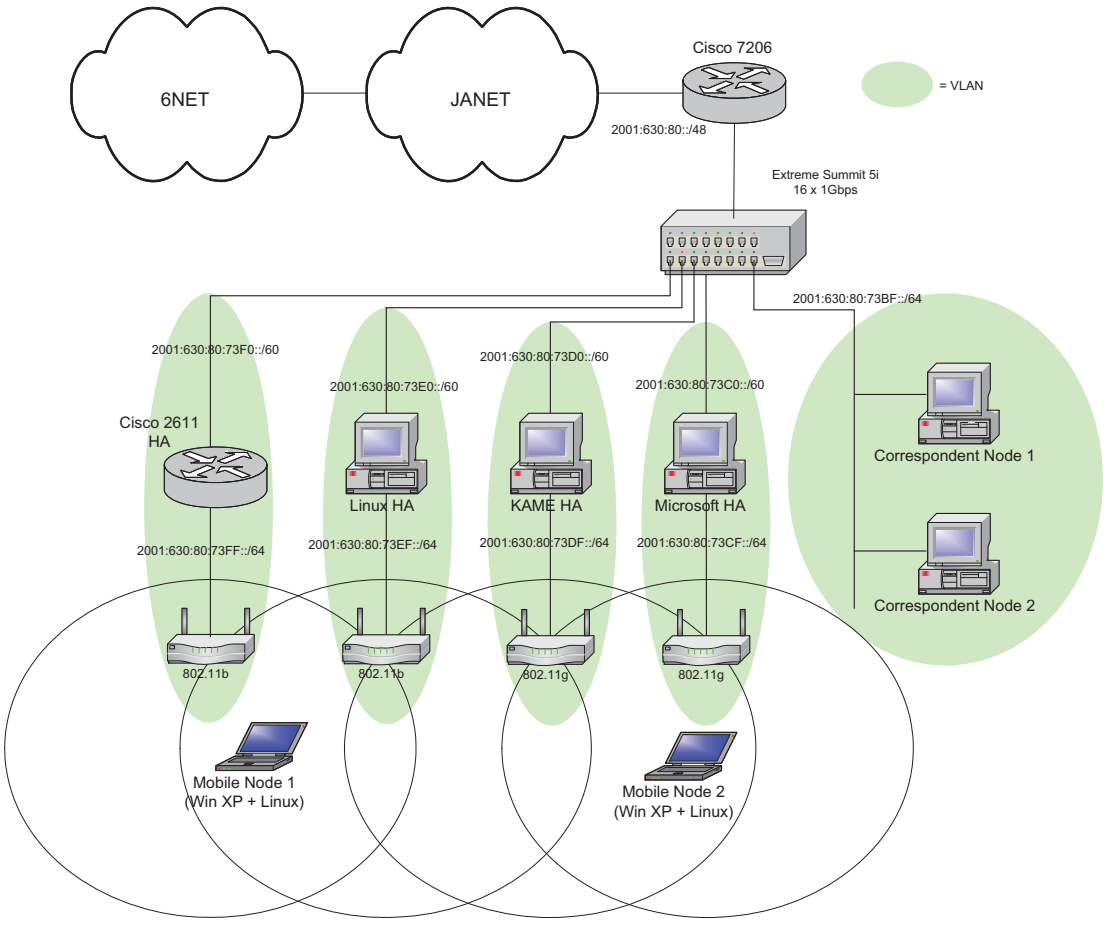

Figure 14-2 Lancaster University MIPv6 Testbed

Connectivity to 6NET (via SuperJANET) is achieved through a Cisco 7206 router. The testbed is then divided into Virtual Local Area Networks (VLANs) with each subnet (generally a ::/60 or a ::/64) comprising its own VLAN. The Extreme Summit Ethernet switch is capable of assigning VLAN tags

according to IPv6 prefixes and switching them accordingly (thus, it acts as a virtual IPv6 router). All of the IPv6 VLANs are terminated on the Cisco 7206.

The 802.11 wireless LAN is configured so that each access router (i.e. HA) of the wireless LAN is analogous to one ESSID. Only one Access Point per ESSID is illustrated in the figure, although wireless coverage for each ESSID can be extended with more Access Points. Of course, configuring multiple Mobile IPv6 networks in this manner will eat up the available 802.11b channel space rather quickly. The figure above illustrates how multiple Mobile IPv6 HAs can have their associated Access Points arranged so that a maximum of 3 overlapping channels is seen in one cell footprint. Obviously, this is rather straightforward to accomplish in a lab environment where our main focus is to investigate the interactions between multiple Mobile IPv6 networks. In a real (semi)production environment the close geographical proximity of different distinct Mobile IPv6 networks is much less likely.

14.3.2 Components

The various components in our Mobile IPv6 testbed are as follows.

Home Agents

The testbed employs four different Home Agents. One using the Microsoft MIPv6 technical preview implementation for Windows XP, another using MIPL v1.1, another using the Cisco Ohanami EFT and finally one Home Agent using KAME. The Microsoft, Linux and KAME Home Agents are all PC based routers each comprising AMD Athlon XP 2.6Ghz CPU 512MB RAM. The Cisco Home Agent is a 2611 XM with 128MB on board RAM.

Mobile Nodes

The Mobile Nodes in the testbed are dual-boot (Windows XP and Linux 2.4.26) machines and can therefore use either MIPv6 implementation of the two available operating systems. The Windows XP operating system has Service Pack 1 installed along with Microsoft's Advanced Networking Pack update. The MIPv6 implementation on the Windows XP operating system

Correspondent Nodes

The two PC-based Correspondent Nodes comprise AMD Athlon XP 1.6Ghz CPUs and 256MB RAM. One of the PCs is a dual-boot Windows XP and Redhat Linux system, the other is a FreeBSD system running FreeBSD 4.10-RELEASE. Note that the Mobile Nodes may also perform the role of a Correspondent Node.

The make-up of the various components in our Mobile IPv6 testbed are summarised in Table 14-2.

Table 14-2 Lancaster MIPv6 Testbed Components

	Hardware	System	MIPv6 Implementation
Cisco HA	2611 XM, 128MB RAM	Ohanami IOS EFT	Included in IOS
Microsoft HA	AMD Athlon XP 2.6Ghz, 512MB RAM	Windows XP SP1 + advanced networking update	Draft v24 compliant Technical Preview
KAME HA	AMD Athlon XP	FreeBSD 4.10-RELEASE	SNAP /kame/snap/kame-

	Hardware	System	MIPv6 Implementation
	2.6Ghz, 512MB RAM		20050103-freebsd410-snap.tgz
Linux HA	AMD Athlon XP 2.6Ghz, 512MB RAM	Redhat Linux 2.4.26	MIPL v1.1
Mobile Node 1	Sony Vaio Intel Pentium III Mobile 1Ghz, 256MB	Dual boot Windows XP SP1 with advanced networking update and Redhat Linux 2.4.26	Draft v24 compliant Technical Preview MIPL v1.1
Mobile Node 2	Dell Latitude C610 Pentium III Mobile 1Ghz, 256MB	Dual boot Windows XP SP1 with advanced networking update and Redhat Linux 2.4.26	Draft v24 compliant Technical Preview MIPL v1.1
Correspondent Node 1	AMD Athlon XP 1.6Ghz, 256MB	Free-BSD 4.10 – RELEASE	SNAP /kame/snap/kame-20050103-freebsd410-snap.tgz
Correspondent Node 2	AMD Athlon XP 1.6Ghz, 256MB	Dual boot Windows XP SP1 with advanced networking update and Redhat Linux 2.4.26	MIPL v1.1

14.3.3 Addressing and Subnetting

The rationale behind the addressing and subnet allocations is as follows. The prefix allocated to Lancaster University by JANET is 2001:0630:0080::/48.

Other than some special addresses reserved for e.g. router interface numbering and DNS, Lancaster University IPv6 addresses observe the following format:

<48 UNI> <1 Res> <3 Site> <12 Subnetting> <64 Host>

where:

<48 UNI> - 48 bit University prefix 2001:630:80::/48

<1 Res> - 1 reserved bit for future aggregation

<3 Site> - 3 bit code identifying the site of the network

<12 Subnetting> - 12 bits for site subnetting

<64 Host> - 64 bit host identifier

All of the sites, with the exception of site 7 (Research and Development), use their 12 bits for subnetting as follows:

<8 Building> <4 Network>

where:

<8 Building> - 8 bit code identifying a building within the site

<4 Network> - 4 bits to allocate subnets within each building

However, for our experimental Mobile IPv6 testbed we assigned it to site 7, research and development. We felt it was wise to have production and experimental research networks attributed to different sites and our Mobile IPv6 testbed is very much non-production traffic. The prefix for research and development networks is: 2001:630:80:7000::/52

Site 7, research and development uses its 12 bits for subnetting as follows:

<2 Reserved> <6 Networks> <4 Subnets>

where:

<2 Reserved> - 2 bits reserved for future aggregation

<6 Networks> - 6 bits to allocate to R&D networks

<4 Subnets> - 4 bits to allocate Subnets

Hence the full address format for addresses on the research and development network is:

<48 UNI> <1 Res> <3 Site> <2 Reserved> <6 Networks> <4 Sub-Networks> <64 Host>

There are 6 bits to allocate research and development networks (26 = 64 networks) and they are allocated in a flat manner.

For example:

2001:630:80:700::/60 - Network 0

2001:630:80:701::/60 - Network 1

2001:630:80:702::/60 - Network 2

2001:630:80:703::/60 - Network 3

…

2001:630:80:73F::/60 - Network 63

Further subnets can be defined arbitrarily based on the 4 bits to allocate for subnets.

So each Mobile IPv6 network is assigned a ::/60 prefix as illustrated in Figure 14-2. Currently each Mobile IPv6 network comprises one Home Agent. Each HA will advertise a ::/64 prefix via Router Advertisements to hosts on its link. Currently there is no stateful address configuration (e.g. using DHCPv6) on our MIPv6 networks and all hosts use stateless address autoconfiguration (remember that not every host on the HA's link needs to be a MN). A MN will know that it is on its home network when it sees the RA from its HA. It is possible to have multiple HAs per network but we chose this topology as it facilitates easier debugging during testing.

14.3.4 Testing

We have tested manually configured IPSec security associations between Microsoft MN and KAME HA. This seems to work well in that the Microsoft MN is able to register with the KAME HA and binding updates are performed successfully. When testing a Linux MN or HA we have only used non-IPSec authenticated communication between the MN and HA. The same is true when using a Cisco HA. In general we have found that all the implementations in our testbed are interoperable in a basic form.

For obvious reasons we do not employ any AAA or access control mechanisms when performing handover latency tests. This is especially important when using any streaming applications as the delay incurred when waiting for network access to be granted can result in a considerable number of lost packets. We have found that using IPerf [IPERF] between the CN and MN is a very useful tool for discovering the extent of packet loss during handovers when comparing different network and different tunings of the RA intervals.

Since each Home Agent also acts as the default router on its respective link as shown in Figure 14-2 we configure each Home Agent to announce unsolicited Router Advertisements at more frequent intervals than the default. In general, we use a Router Advertisement interval of 1 second which is sufficient for most adaptable TCP applications such as email, web, ftp etc. However, we have noticed considerable (i.e. unpleasant) disruption when using a 1 second interval when streaming video and/or audio. Reducing the interval to around 300ms seems to be fair benchmark to set for supporting streaming media. Although it is possible to reduce the interval even further, in most tests we did not observe enough improvement in the stream disruption to justify the extra link bandwidth and Home Agent CPU cycles being used.

Initially, the most simple way of testing the operation of the MIPv6 testbed is by testing basic connectivity between the Mobile Node and Correspondent Node with a stream of ping requests from the Correspondent Node to the Mobile Mode while the Mobile Node moves between its home network and a foreign network. In a correctly configured Mobile IPv6 testbed, the ping requests should continue to be answered as the Mobile Node moves. Sometimes we observe one ping timing out as the packet is lost. This can occur if the ping request is sent just as the Mobile Node becomes unreachable in its original location but has not yet completed the binding update procedure to be reachable at its new location. However, on a properly configured testbed we rarely observe more than one ping timeout.

For testing that TCP connections survive movement of the Mobile Node, a simple test is to run telent or similar (using the client and server in either combination of Mobile Node and Correspondent Node) and observe that the session remains active after movement. You should not observe any noticeable effects (save some possible interruption in character echo on the terminal) if all is configured correctly.

Beyond these basic tests some users may want to experiment with the TAHI compliance test suite described at [TAHI].

14.3.4.1 Handover Latency Tests

Figure 14-3 shows the small MIPv6 testbed used for performing the handover tests. A MN running MIPL v1.1 is away from home (the HA also running MIPL v1.1) and can attach to one of two networks represented by the SSIDs ‘roam1’ and ‘roam2’. We decided to conduct handover tests from one foreign network to another (e.g. rather than from home network to foreign network) as the nature of mobility implies that when one is mobile, one is very rarely located at the home network.

Handovers from one network to another were forced by turning off one of the APs so that the MN would immediately associate with the other AP and thus receive different RAs than on the previously connected network.

The handover times were measured from the point at which the AP is switched off (link down notification) to when the Binding Acknowledgement is received from the Correspondent Node with each event being timestamped in the relative logs. Note that the Return Routability protocol was enabled for Route Optimisation.

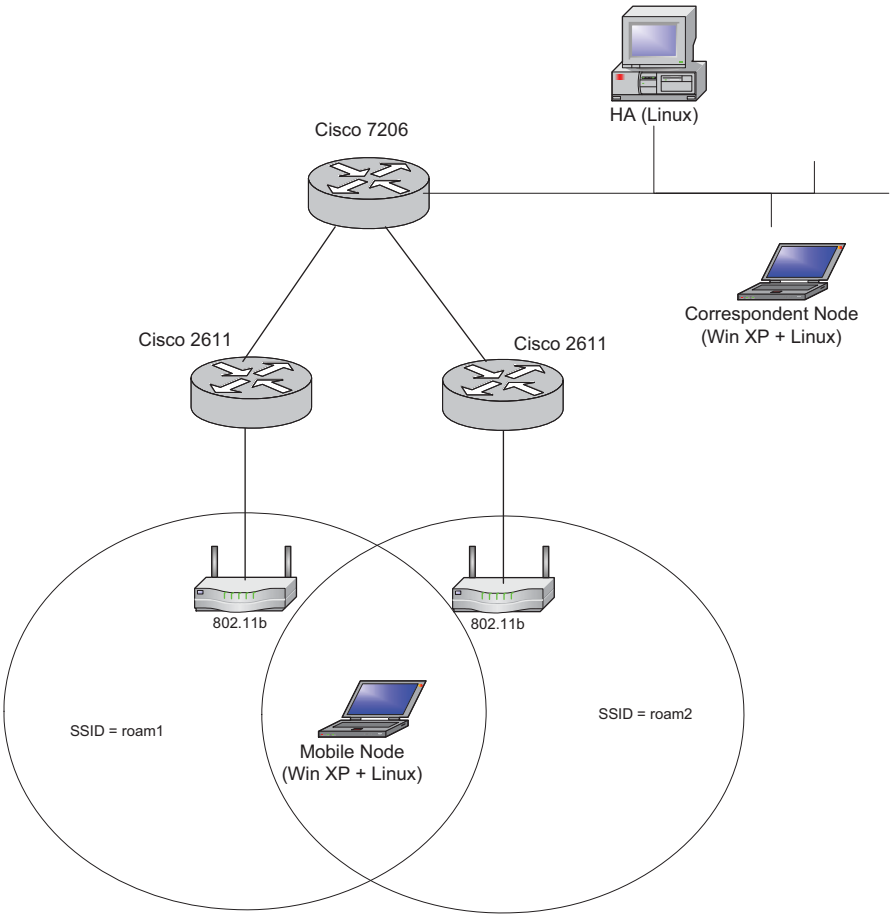

Figure 14-3 Simple MIPv6 Handover Testbed

Reducing the Router Advertisement Intervals

In order to demonstrate the effects of reducing the RA interval we performed handover tests with various configurations of RA intervals on the Cisco 2611 access routers. As described earlier, upon detecting movement, the MN will issue a RS assuming it hasn't received a new RA already. As can be seen from Figure 14-4, the time it takes for the MN to receive a solicited RA is fairly random within a given (configurable) time window.

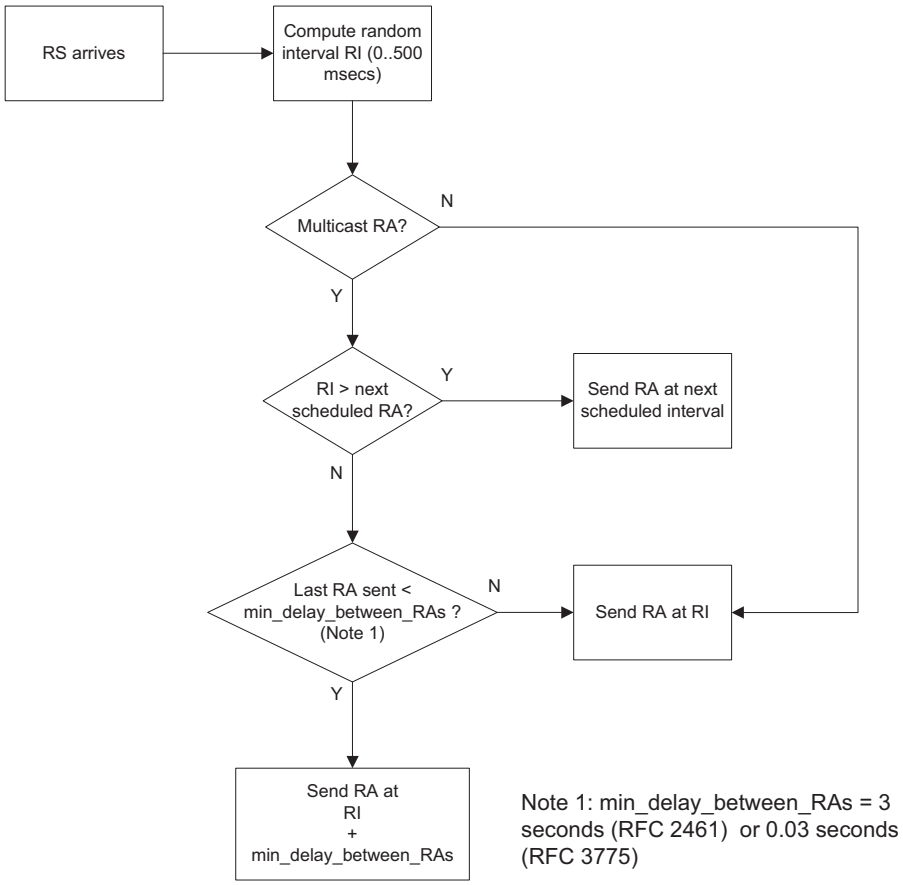

Figure 14-4 Processing Router Solicitations

The test experiment was as follows. The IPv6 capable VoD (Video on Demand) server was located at the Correspondent Node and streamed 1.5Mbps MPEG1 video clip of 30 seconds duration was streamed to the Mobile Node. At 10 seconds into the video clip the AP to which the MN was associated with was switched off, forcing a handover to the other network. At the end of the clip the handover latency and packet loss (reported by the VoD client) were noted. This was repeated 10 times for each value of RA interval configured on the Cisco routers. These RA intervals were 300ms (the RFC 3775 minimum), 1000ms and 3000ms.

Table 14-3 Results of RA Interval Tests

	Avg Latency (seconds)	Avg # Packets Lost
300 ms	1.917	245.376
1000 ms	2.448	313.344
3000 ms	3.013	385.664

It can be seen that changing the RA interval does not have as much effect on reducing the overall latency as we would like. This can be explained in that the rest of the handover procedure after receiving a RA, i.e. CoA configuration, DAD and CoA registration with the HA and CN is completely

unaffected by reducing the RA interval. One can also see the number of packets lost in the video stream. On the client playback the stream would recover itself after handover but the break in the video and audio seemed about 1 or 2 seconds longer than the handover latency reported in the logs. It is easy to conclude that even tuning the RA interval to the lowest possible value will not suffice for real-time voice and video applications in a mobile environment.

Unicasting Solicited RAs

Another possible trick is to change the default behaviour of neighbour discovery so that a Router Solicitation is answered with a unicast RAs rather than the default multicast RA. In the standard algorithm depicted in Figure 14-4 it can be seen how a unicast RA only incurs the random delay interval and is not affected by the configured RA interval parameter (since this only applies to multicast RAs). To see what effect this would have on handover latency we had to replace a 2611 router with a linux PC based equivalent (since we were unable to configure the IOS accordingly).

Table 14-4 Using Unicast RAs

	Avg Latency (seconds)	Avg # Packets Lost
Unicast RA	2.072	265.216

From the table we can see that the results are slightly worse than the best we can get from configuring the RA interval. However, since the random interval is between 0 and 500 ms (the MAX_RA_DELAY_TIME constant in [RFC2461], we are unable to reduce this parameter further.

Eliminating DAD / Optimistic DAD

By removing the DAD procedure altogether we can reduce the handover latency even further (potentially by a second or so). However, removing this check altogether is not a realistic option both in terms of ratification by the IETF or by tuning an implementation's configuration. Thus, we are left with the option of fine tuning the DAD procedure in some way that reduces the time it takes for a MN to be able to use its CoA. A procedure called 'Optimistic DAD' which modifies [RFC2461] and [RFC2462] is proposed in [Moo05], which essentially allows a CoA to be used before it has completed DAD. The CoA is marked as 'optimistic' as opposed to 'tentative' before completing DAD and is marked as 'preferred' once DAD is complete.

Unfortunately, we have not been able to source a suitable implementation of Optimistic DAD with which to test. An implementation will soon be made available from Monash University, but this will appear too late for the lifetime of the 6NET project.

Conclusions

Our tests have demonstrated that even with fine tuning the parameters of routers for optimum MIPv6 handover performance, we still do not approach anywhere near good enough handover times for real-time voice/video applications.

We must therefore conclude that MIPv6, in its current form is not by itself sufficient to be the de-facto mobility management model in the mobile IPv6 Internet. Further optimisations relating to handover performance must be made in order to support interactive and real-time IPv6 applications in a mobile context.

We have examined the fast handover protocol for MIPv6, FMIPv6 [RFC4068]. This aims to improve handover latency by eliminating IPv6 configuration latency and also prevents packet loss by the use of a bi-directional tunnel while physical movement and MIPv6 CoA registration are taking place.

To the best of our knowledge no implementation has yet been developed for us to perform handover tests. Yet this does not prevent us from reasoning that FMIPv6 will indeed reduce handover latency in almost all cases. In some cases, e.g. predicted handovers and relatively infrequent movement FMIPv6 promises to be sufficient for the real-time applications, most notably the killer mobile application VoIP. However, without being able to perform real tests it would be rather hasty to take this for granted.

14.4 University of Oulu

The University of Oulu has carried out experiments with Mobile IPv6 for Linux in heterogeneous wireless networks. In particular, the handover performance has been studied.

14.4.1 Testbed

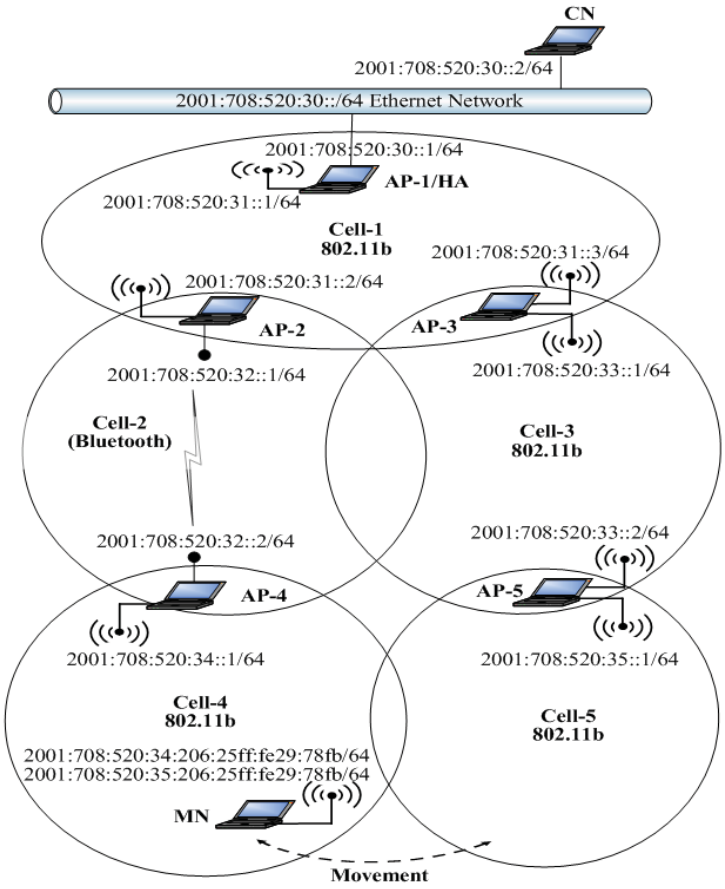

Figure 14-5 University of Oulu Heterogeneous Wireless MIPv6 Testbed

14.4.2 Handover Performance

This series of experiments investigates the effect of handover duration on the TCP disruption time. We will describe handoffs between AP-4 and AP-5 (i.e. between the fast path and the slow path) as shown in Figure 14-5. AP-4 and AP-5 generate one router advertisements every 1.5 seconds. We define the handover time as the amount of time starting from when the mobile node can not get any packets from its old access point until the time when it receive binding acknowledgement from its home agent via the new access point. The first handoff experiment is from AP-5 to AP-4 (i.e. from the fast path to the

slow path). The MN starts sending TCP traffic to the CN using AP-5 and moves towards AP-4. Figure 14-6 shows the time sequence number plot of the TCP connection including the handover time. We found that it takes around 4.9 seconds to handoff from AP-5 to AP-4. The TCP disruption time was found to be 10.5 seconds. As illustrated in the figure, the TCP goes through 4 consecutive timeouts and retransmissions. Because of TCP exponential backoff mechanism, TCP doubles the size of its timeout interval each consecutive timeouts. Following the handover, the TCP waits for the last timeout to elapse before starts sending data using the new care-of-address. This result demonstrate that the handover duration contribute approximately 40% of the TCP disruption time. The remaining time delay is due to TCP congestion control mechanism. The average handover duration and TCP disruption time are shown in Table 14-5.

Table 14-5 Handover Duration and TCP Disruption Time

	Average handover duration [seconds]	Average TCP disruption time [seconds]
From fast to slow path	3.94	8.40
From slow to fast path	1.49	3.91

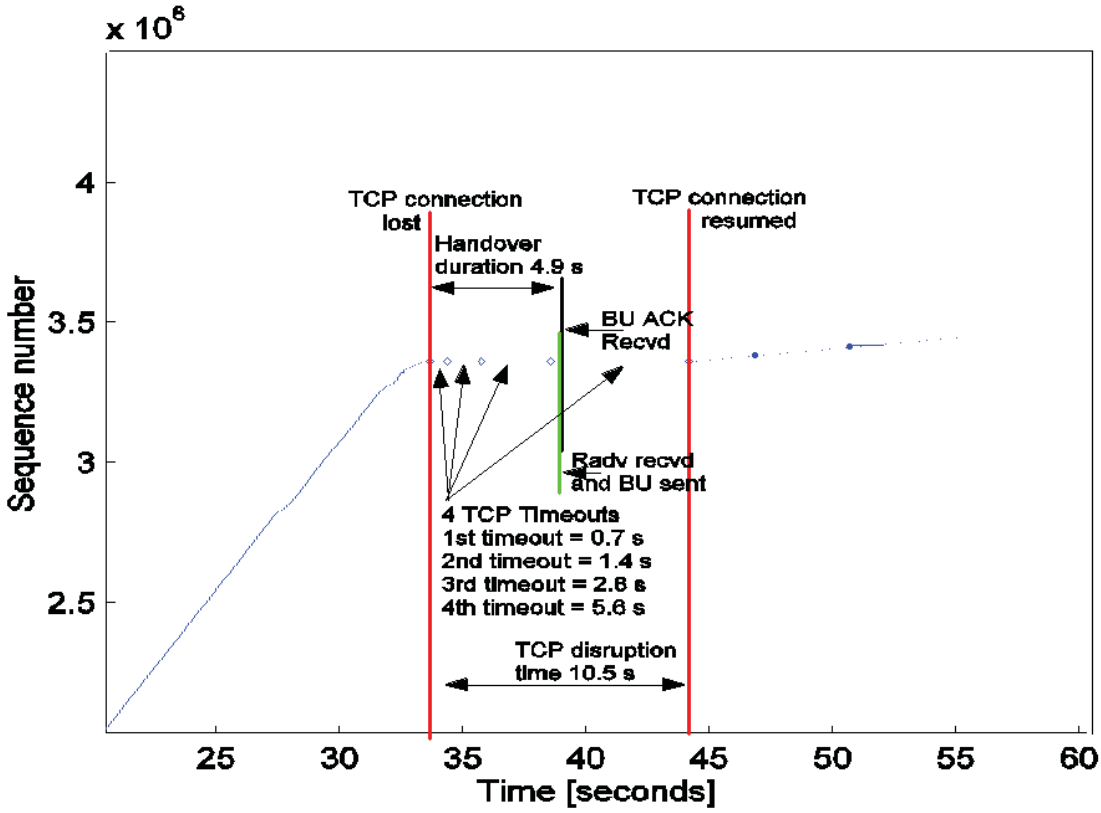

Figure 14-6 TCP Packet Trace During Handover from AP-5 to AP-4

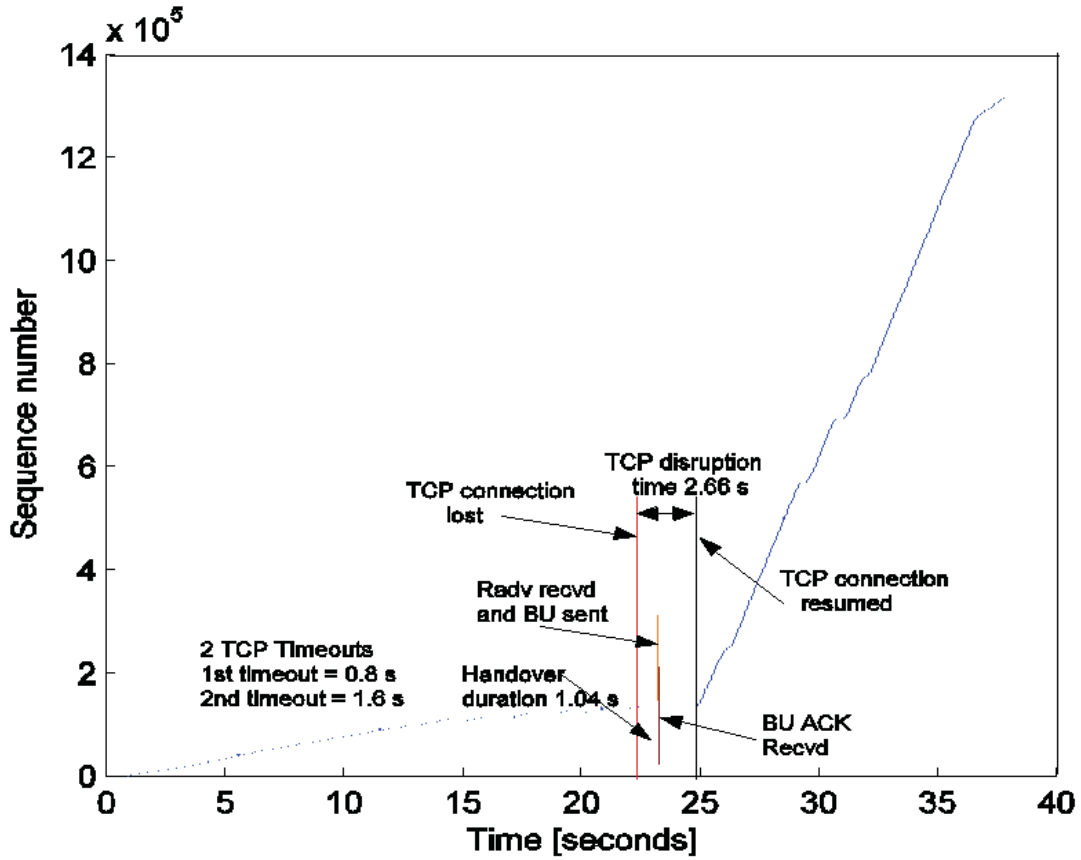

Figure 14-7 TCP Packet Trace During Handover from AP-4 to AP-5

The same behaviour was noticed when the handover was performed from AP-4 to AP-5 (i.e. from the slow path to the fast path). However, handover duration significantly decreased to about 1.04 seconds. The TCP disruption time takes around 2.66 seconds. TCP Packet Trace During Handover from AP-4 to AP-5 shows the TCP behaviour during a handover from AP-4 to AP-5. It can be seen that fewer timeouts and retransmissions have occurred in this experiment. The overall average handover duration and TCP disruption times are shown in Table 14-5.

These results indicate that handover between AP4 to AP5 is faster than from AP5 to AP4. This is because related signalling packets travel faster in the fast path. It also shows that TCP performs well when handover is from a low data rate to a high data rate network. However, due to the TCP slow start mechanism, the recovery is slow when the handover is from the high data rate to the low data rate path. After the handover, the MN starts sending packets at slow rate (sets the congestion window maximum segment size to 1) and increases its sending rate exponentially fast. Several methods have been introduced to optimize the Mobile IPv6 handover process, including hierarchical MIPv6 and fast handover. The fast Handover protocol has been proposed as away to minimize the interruption in service experienced by a Mobile IPv6 node as it changes its point of attachment to the Internet. Using the fast handover mechanism would allow the MN to send and receive packets from the time that it disconnects from one access point to the time it registers a new care-of address from the new access point.