Project Number:	**IST-2001-32603**
Project Title:	**6NET**
CEC Deliverable Number:	**32603/UOS/DS/2.5.1/A1**
Contractual Date of Delivery to the CEC:	30th September 2002
Actual Date of Delivery to the CEC:	6th December 2002
Title of Deliverable:	D2.5.1: Issues for IPv6 Deployment (missing pieces for IPv6 deployment and IPv6-only operation)
Work package contributing to Deliverable:	WP2
Type of Deliverable*:	R
Deliverable Security Class**:	PU
Editors:	Tim Chown (University of Southampton)
Contributors:	Feico Dillema (Invenia) and WP2 participants

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

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

Abstract:

In this report we present an initial scoping of issues that remain to be resolved for IPv6 deployment to be realised in both site and ISP networks. The report focuses on IPv6 components and "missing pieces" that still require attention in either standardisation processes or within appropriate bodies essential to successful deployment. The scope has changed from the original aim of reporting on IPv6-only networking issues; the broader coverage of general IPv6 deployment issues is more appropriate to this initial scoping report. We also present a case study of an IPv6-only site deployment made at a university in Norway.

Keywords:

IPv6 site transition, IPv6-only networking, IPv6 integration, IPv6 deployment issues

Executive Summary

There has been much effort in the last 2-3 years in standardisation bodies (principally the IETF) and from vendors (operating system and router) towards the availability and deployment of IPv6. However, while IPv6 deployment is beginning to happen in some parts of the world, there are a number of issues that remain to be resolved or improved before that adoption can become more widespread.

In this report we highlight a range of these issues, as they have arisen within the scope of the 6NET project. We do not include in detail issues of transition from IPv4 to IPv6 per se; such interworking issues will be reported in the annual site and ISP transition cookbooks that the 6NET project will produce within Work Package 2.

This report is aimed at scoping the problems. These, and progress on them (by external bodies to the project) will be reported in a later 6NET project deliverable.

These IPv6-specific issues include (in no particular order of importance):

- Use of site-local IPv6 addresses
- Operation of IPv6 Multicast
- IPv6 Firewalls
- Service discovery mechanisms
- IPv6 DNS
- Application preference for use of IPv4 or IPv6 connectivity
- IPv6 prefix delegation
- IPv6 privacy extensions
- Routing stability in “production” IPv6 networks
- Porting code to run using IPv4 or IPv6 stacks

We also present a brief case study of a 6NET site that is already running IPv6-only.

IPv6 is generally ready for deployment, but the issues raised in this document need some consideration for widescale deployment to be successful.

Table of Contents

1. Introduction

IPv6 is beginning to be deployed in sites and networks around the world. The 6NET project has already deployed an international IPv6 backbone spanning some 15 National Research and Education Networks (NRENs). Some of the 6NET partner sites are running quite extensive IPv6 networks, though almost invariably dual-stack IPv4/IPv6. One 6NET site, in Norway, has an IPv6-only network spanning a whole department. Other 6NET sites have IPv6-only network testbeds, for the purposes of learning which IPv6 components (or standards) are still missing before a full IPv6-only deployment could be considered. Note that in this case, “IPv6 only” may mean a device with a hybrid IPv4/IPv6 stack, but with only IPv6 networking configured.

In this document we present a list of problems that still exist for IPv6 deployment. While these may not be critical for individual sites, they represent issues that need to be addressed before a global deployment of IPv6 can be successful.

Note that this deliverable is scoping the issues, and the available systems platforms. More detailed analysis will follow in Deliverables D2.5.2 (due by July 2003) and D2.5.3 (due by the end of December 2004).

2. Vendor and Standards Support

A key basic requirement for IPv6 deployment is IPv6 support in host and router stacks. In the last two years, the state of implementations has improved significantly, to the status where many vendors have commercial support for IPv6 functionality.

2.1. Host OS support

Host operating system support includes the following vendors:

- Apple, Jaguar OS: http://www.apple.com/macosx/jaguar/morefeatures.html
- Compaq (Tru64): http://www.compaq.com/ipv6/
- FreeBSD: http://www.freebsd.org/
- KAME stack (*BSD): http://www.kame.net/
- HP-UX: http://www.hp.com/products1/unix/operating/internet/ipv.html
- IBM AIX: http://www.ibm.com/software/ipv6/
- Linux (USAGI): http://www.linux-ipv6.org/
- Linux (general): http://www.bieringer.de/linux/IPv6/IPv6-HOWTO/IPv6-HOWTO.html
- Microsoft Windows (XP and .NET): http://www.microsoft.com/ipv6/
- NetBSD: http://www.netbsd.org/
- Sun Solaris (8 and 9): http://wwws.sun.com/software/solaris/ipv6/

The level of support in the various operating systems is mixed. The BSD variants, thanks to the KAME stack integration, have probably the best IPv6 support “out of the box”. Many Linux flavours are catching up fast, helped by the USAGI project. Basic support in SuSE and RedHat distributions is good, with IPv6 present at the very least as a loadable module. IPv6 support (as an Experimental option) has to be configured in the stable 2.4.x Linux kernel (it is not usually a default option when building a downloaded kernel source). Solaris support has been good since Solaris 8, carried into Solaris 9. Perhaps most importantly, at least from a user base perspective, is Windows XP and .NET support (as well as support in Windows CE .NET); in XP IPv6 can be turned on by running “ipv6 install” at a command window prompt. Windows XP ships with ISATAP and 6to4 support, and also actively uses RFC3041 privacy extensions. Windows 2000 is still prevalent, but the IPv6 support from MS for Windows 2000 is “experimental”; it is not commercially supported, and applying the IPv6 hotfix may undo some Windows 2000 security patches.

Within the scope of the 6NET project we plan to report both on incidental experiences with IPv6 support in operating systems, but also to utilise IPv6 conformance and interoperability test suites where available and practical, e.g. the TAHI project [tahi] and the Ixia test suite [anvl], and we will also track reports from ETSI and University of New Hampshire IPv6 “plug fests”.

2.2. Router support

Router support includes the following vendors:

- 6WIND: http://www.6wind.com/
- Cisco IOS: http://www.cisco.com/warp/public/732/Tech/ipv6/
- Ericsson Telebit: http://www.ericssontelebit.dk/
- FreeBSD: http://www.freebsd.org/
- Hitachi: http://www.internetworking.hitachi.com/products_ipv6.shtml
- ipinfusion (ZebOS): http://www.ipinfusion.com/
- Juniper (JUNOS): http://www.juniper.net/products/ipv6_overview.html
- MRT: http://sourceforge.net/projects/mrt
- NetBSD: http://www.netbsd.org/
- Nortel Networks: http://www.nortelnetworks.com/corporate/technology/ipv6/
- Zebra: http://www.zebra.org/

Router support also varies quite considerably in terms of functionality, e.g. for features varying from IPv6 in IPv6 tunnels, through IPv6 PIM-SM support, for various transition methods (NAT-PT, ISATAP, 6to4), for OSPFv3 or IS-IS, and, most importantly, for hardware forwarding performance, which is critical in ISP backbone networks.

In the updated version of this deliverable, D2.5.2, we will report in detail on the functionality present in some of the above router products, where used by 6NET project partners.

2.3. IPv6 IETF Standards Working Groups

The major IPv6-related IETF Working Groups are listed below. In recent months the ipng WG has been renamed ipv6, and the ngtrans WG is being phased out and replaced by v6ops. These changes are intended to signal the general readiness of IPv6 for deployment.

- IPv6 [ipv6]
- IPv6 Operations [v6ops] (replacing Next Generation Transition [ngtrans])
- Dynamic Host Configuration [dhc]
- DNS Extensions [dnsext]
- Mobile Ad-hoc Networks [manet]
- IP Routing for Wireless/Mobile Hosts [mobileip]
- Site Multihoming in IPv6 [multi6]
- Network Mobility [nemo]
- Securing Neighbor Discovery [send]

The manet and nemo WGs are still in a relatively "research"-oriented stage; manet is not doing much consideration of IPv6 yet. The zeroconf and proposed zerouter WGs are also set to take on IPv6 items.

It is important that IPv6 standards are taken on board by all IETF WGs, especially at present those in the applications area.

2.4. IPv6-only operation

It is important to recognise that at present IPv6-only operation is not common, and systems built with no IPv4 capability at all are exceptionally rare (except for some micro/embedded systems in Japan). However, it is possible to use BSD and Linux systems in IPv6-only mode "out of the box" (with IPv4 present in the hybrid OS but not enabled), because IPv6 basic applications exist to support the networking protocols, and IPv6 transport DNS is available to support name resolution over IPv6. In router platforms, while IPv6-only routing can be done in many cases, other reasons may exist to retain IPv4, e.g. access to IPv6 SNMP MIB data over IPv4 SNMP connections.

3. Standards-Related Issues Affecting IPv6 Deployment

In the course of discussions between the 6NET project participants a number of issues have arisen that affect the readiness of IPv6 for successful global deployment. We outline these issues in the following four sections. They are generally, though not universally, standards-related issues, although in some cases issues of policy also arise. In a number of cases, the issues have been written up as IETF Internet Drafts by project participants as a result of the discussions.

In addition some issues are reported in an IETF Internet Draft by Hagino and Jinmei [issues00].

There is no particular order of importance to the topics discussed in this document. In the following four sections we categorise the issues somewhat arbitrarily into the areas of

- network robustness and performance (which must exist for users to be encouraged to use IPv6 instead of IPv4)
- network management (including supporting services)
- application and IPv6-specific features (of concern to application developers and end users)

- security considerations

It should also be noted that there are many IETF Internet Drafts that are important to progress to at least Draft Standard so that implementors may be encouraged to put more effort into developing and releasing code. Two notable examples are Mobile IPv6 [mipv6] which is at Draft v19, and DHCPv6 [dhcpv6] at Draft v28. However, both are expected to go for Last Call on or around the end of 2002.

4. Network Robustness and Performance Issues

4.1. General IPv6 Routing Stability

While routing of IPv6 traffic is generally good in a local region (and is within 6NET for example), international IPv6 connectivity is generally poor, due in large part to the combination of long (multiple IPv4 AS-hop) IPv6-in-IPv4 tunnels, of sites/ISPs with high numbers of peers (sites wish to establish direct tunnels to bypass the otherwise poor routing, leading to a peering "arms race"), and inappropriate sites/ISPs giving free transit.

These problems are discussed in a new "6bone mess" Internet Draft [6mess01].

6NET is working with Euro6IX, Abilene (the US research network) and WIDE (in Japan) to establish good routing policy to build a predictable, efficent and well-performing IPv6 infrastructure for research, such that IPv6 can more readily be used for day-to-day activity and work.

4.2. Preferring IPv4 or IPv6 Connections

In attempting to promote the use of IPv6, many applications on receiving both A (IPv4) and AAAA (IPv6) records back from DNS queries will attempt to connect to the IPv6 service in preference to using IPv4, only falling back to IPv4 if IPv6 fails (after some timeout delay).

However, given the observed routing issues described above – e.g. at the recent IETF conference the author experienced 100ms RTT delays for IPv4 connections back to the UK against 400ms RTT delays over IPv6 – preference for IPv6 over IPv4 will more often than not give a worse experience for the user.

Methods to allow per-host or per-application selection of IP version preference may thus be desirable.

4.3. Multihoming with IPv6

For IPv6 to be adopted by sites and ISPs already accustomed to multihoming in IPv4, an IPv6 multihoming solution is required. There have been over 30 Internet Drafts proposed that have been related to IPv6 multihoming in the last 4-5 years, but these have almost universally failed to reach RFC status with the exception of SCTP [rfc2960].

The IETF multi6 WG has been static for some time, having failed to agree a requirements draft text. As a result, an "alternative" WG has been set up to discuss multihoming; this IPv6MH group [ipv6mh] met several times in the November 2002 IETF meeting in Atlanta, and looks set to make some progress.

The basic problem is that punching out /48 prefixes into the IPv6 DFZ will not scale. But running "classic" IPv6 multihoming where every host in a site can inherit an IPv6 address tied to each site connectivity provider has also not yet been shown to work. As a result, a number of host and router-based solutions are being proposed, including methods that split the identifier and routing space.

A 6NET Project deliverable has recently been produced that overviews IPv6 Multihoming solutions [d451].

5. Network Management Issues

5.1. IPv6 and DNS

Although the AAAA vs A6 issue has been resolved, with A6 and DNAME moving to Experimental status [rfc3363] (and thus in effect Historic), a number of DNS-related issues remain:

- Not all operating systems offer IPv6 transport lookups. While BSD and Linux can talk to DNS servers over IPv6, this is not yet possible in Windows XP without using a roll-your-own local forwarding process/proxy.
- There is no standard method for statelessly autoconfiguring IPv6 hosts to discover a DNS server address; instead DHCP(v6) or manual configuration needs to be used. There is a draft proposal to use well-known site-local addresses [sitedns], but nothing as yet widely implemented. Solving this bootstrapping problem is an important issue for IPv6 "plug and play" deployment.
- There are no public IPv6-enabled root DNS servers. Such a pilot server does exist however, which requires a disclaimer to be signed before it can be used.
- You cannot register a new domain with IPv6 DNS entries in any common registrar. This will be important should anyone wish to operate an IPv6-only service.
- The ip6.int to ip6.arpa reverse delegation transition process has been slow, especially for 6bone address space under 3ffe::/16.

As with IPv6 Prefix Delegation – see below – there is a lot of pressure to resolve DNS issues because they are critical to successful widespread deployment.

5.2. IPv6 and SNMP

Although IPv6 MIBs exist (many of which are being reworked to blend more cleanly with IPv4 MIBs), SNMP operations are most commonly run over IPv4 connections. IPv6 transport SNMP is still in its infancy; some vendors support it, but it is not in wide use yet.

Thus while the 6NET backbone has only IPv6 sessions running between the backbone routers, IPv4 is run over the links from the national access PoPs to allow SNMP to those backbone routers. IPv6 SNMP transport is required to remove that dependency on IPv4.

5.3. Service location methods

There are many methods proposed for service discovery in IPv6 environments. Different protocols have different preferred methods for discovering services or particular devices (e.g. routers or relays). It may be desirable to have some consistency in methods. It may be that some methods (e.g. well-known site local addresses – see above) have specific problems cited against them.

The discovery methods include:

- Use of IPv4 or IPv6 Anycast addresses (e.g. IPv4 Anycast for 6to4 relay router discovery)
- Link or site scope IPv6 Multicast (e.g. in Neighbor Discovery)
- Well-known site local addresses, e.g. use of fec0:000:0000:ffff::1,2,3 for DNS server discovery as specified in [sitedns]
- Service Location Protocol [rfc2165]
- Well-known DNS name (e.g. "isatap" as used in the current ISATAP Internet Draft)
- Advertising services in Router Advertisement "piggyback" messages
- Use of DHCP(v6)
- Linklocal Multicast Name Resolution (LLMNR), aka. mDNS [llmnr13]

It would be interesting to survey IPv6 RFCs and Internet Drafts for usage of the above set of methods. This is not a barrier to deployment per se, but it should be noted that not all these methods are available now (e.g. general implementations of DHCPv6).

5.4. IPv6 Multicast Deployment

IPv6 multicast deployment issues have been described in a recent Internet Draft [mcast01]. A key problem for IPv6 Multicast is the lack of an inter-domain communication method for PIM-SM, e.g. there is no MSDP for IPv6. However, some methods have been proposed, e.g. embedding the Rendezvous Point location in the multicast address. There may be a reduced requirement for PIM-SM where single source (source specific) multicast (PIM-SSM) is used instead.

MLD [rfc2710] and its successor MLDv2 [mldv2] are both important for IPv6 Multicast deployment. MLDv2 obsoletes MLD, and includes support for listening for specific sources, and thus for PIM-SSM.

There is some discussion on how best to handle multicast traffic at Layer 2, i.e. whether MLD and MLDv2 snooping is required, or whether a specific new protocol may help solve the problem of multicast traffic "swamping" links on switched Layer 2 networks.

5.5. IPv6 Prefix Delegation

ISPs need a method to assign IPv6 prefixes to customer equipment. This has become a requirement for commercial ISPs in Japan that are beginning to offer native IPv6 DSL services. The most favoured method is currently a DHCPv6 option [v6pd].

5.6. Wireless LAN Access Point Management

Management of WLAN access points is only available over IPv4. Thus an access point used in an IPv6 only WLAN either has to be configured prior to deployment, or via an out of band method (e.g. a serial interface where present), unless only the air interface is run IPv6-only and the wired link from the access point back to the upstream router is dual-stack.

Similar comments could be made about other network hardware, e.g. networked printers.

5.7. IPv6 NTP service

IPv6 support for NTPv4 has recently been added to the NTP development project site [ntp]. The reference ID issue has been resolved as the first 32 bits of the MD5 hash of the IPv6 address.This code is being tested within the 6NET Project.

6. Application and IPv6-specific feature issues

6.1. Use of IPv6 site local addresses

The IPv6 Scoped Address Architecture specification [scope04] defines the implications of having multiple scopes for IPv6 addresses, including link local, site local and global scopes, as originally defined in the IPv6 Addressing Architecture [rfc2373]. The usage of link and global scope addresses has become well understood. However, there is considerable discussion within the IETF (many hundreds of IETF email list messages, and a whole dedicated session at the November 2002 IETF meeting in Atlanta) on best practice for use of site local addressing.

The two basic problems cited for site locals are ambiguity and leakage (routability). Site local addresses are ambiguous because any site using such addressing may choose its own site local prefix from within the fec0::/10 prefix. As a result, there will inevitably be clashes in addressing, and ambiguities where applications cross site boundaries. It is also likely to be inevitable that site local addresses (including source IPv6 addresses) will leak from sites, just as RFC1918 addresses do in IPv4 space now. However, site local addressing should not be seen as a security measure, nor should it promote in any way use of as yet non-existant IPv6 NAT.

It seems clear that some kind of restricted use of site locals is desirable (e.g. only in disconnected networks), but intermittently connected networks, or those reconnecting with varying prefices, need some form of stable addressing. The IETF ipv6 WG is discussing possible creation of Globally Unique Site Local (GUSL) addresses that would at least attempt to solve the ambiguity issue. At present, we cannot create Globally Unique Provider Indepedent (GUPI) addressing because advertising all GUPI /48 prefixes to the IPv6 DFZ would not scale.

6.2. IPv6 Code Porting

The 6NET Project is undertaking a considerable amount of porting work, within the Applications (WP5) and Network Management (WP6) Work Packages. For example, both the Globus Toolkit and the Vocal VoIP packages are being ported. The methodology of such porting is to

a) use the standard advanced APIs [rfc2133][rfc2292] and RFC2292bis [sockapi], and also RFC2553 [rfc2553] and 2553bis [2553bis], and the Single UNIX Specifiation, Version 3 [single3].

b) make the code AF (IP) independent

c) ensure the porting feeds directly back to the main code tree, to avoid having a patch tied to a specific release version

There are examples of existing suggestions for porting best practice, including:

- A KAME Newsletter [kport]
- The Sun Porting Guide [sport]
- The LONG Project Porting Guide [lport]

Other porting references can be found on the IPv6 Forum web site [ipv6forum].

The 6NET Project will seek to report on its porting experiences.

Further discussion of these issues can be found in a recent Application Aspects of IPv6 Transition draft [appasp].

6.3. Missing applications and protocols

There are still many services missing for IPv6 that are standard in IPv4. Two examples are network file sharing and database access.

Linux has no TI-RPC support, so it cannot do RPC over IPv6. Hence NFS, NIS etc. are missing. This is implemented on *BSD and Solaris though. Windows XP users would also need to be able to run SMB over IPv6; there is an IPv6 patch for Samba (see http://v6web.litech.org/samba/) but this is not integrated into the main Samba code base due to lack of IPv6-enabled SMB clients from Microsoft. Having AFS available over IPv6 would also be useful (this is listed as future work on the OpenAFS web site).

Another problem is lack of IPv6 support in SQL databases like mySQL and postgreSQL. We are not aware of any SQL databases supporting IPv6. There is a big problem that the large ISVs do not pay much attention to IPv6 yet, e.g. Oracle, CA etc.

Note that LDAP does have IPv6 support, included as standard in OpenLDAP, and also in Sun ONE (formerly known as iPlanet).

6.4. Use of the IPv6 Flow Label

The IPv6 Flow Label usage [flow03] is still not agreed, beyond a basic Internet Draft specifying how to handle the Flow Label until such usage is determined. This recommends that where used the flow label is set by the sender to be unique for the (source, destination) flow, that it is immutable in transit, and that no semantics should be read into the value. Where not used, the field should be zero value.

The original idea for the flow label tied it to IntServ QoS, which is no longer widely used (for scaling reasons).

6.5. Use of IPv6 Privacy Extensions (RFC3041)

RFC3041 [rfc3041] was devised to prevent IPv6 devices being trackable through the host part of the address being constant where EUI-64 (stateless autoconfiguration) addresses are used to connect a device to different networks in different locations. These privacy extensions allow a host to generate a "random" host part of its address in stateless autoconfiguration, thus preventing correlation of host parts of IPv6 addresses where observed in traffic logs.

Windows XP implements RFC3041.

In some cases, as described in [harm01], RC3041 may be detrimental. One case is where distributed denial of service attacks are run which regularly change the source address that they use (which may be misconstrued as RFC3041 behaviour). While trusting a specific source IP or set of IPs is not necessarily optimal authentication, RFC3041 makes this less easy to do. There is also an argument that use of RFC3041 can be detected (the host does not have an EUI-64 host address), thus use of privacy is itself observable.

Users may still be tracked by cookies and by persistent IPv6 network address prefixes, but many people see RFC3041 as important.

7. Security-related Issues

7.1. Implementation and use of IPv6 IPsec

The base IPv6 specification [rfc2460] states that a "full implementation" of IPv6 must include implementation of the AH and ESP headers. This means that fully IPv6-compliant stacks must support the ability for the application (or user) to use IPsec. However, the use of IPsec itself is not mandated.

At present, few host IPv6 stacks support IPsec. This is expected to change with time. The lack of widely available PKI solutions also hampers the deployment of IPsec except in manually configured or keyed environments. IPsec is also important for routers, e.g. OSPFv3 can use AH for route data authentication (although how exactly is another issue).

The IETF send WG is studying methods that could be applied to secure Neighbor Discovery. At present the fully open "plug and play" IPv6 stateless autoconfiguration has no security methods; a rogue IPv6 router can attach to a network and advertise a bogus prefix and route, and likewise unwanted hosts can join an open network.

7.2. IPv6 Firewalls

There are no commercially available IPv6 firewall products at the time of writing, although some vendors have hinted at future availability (e.g. Checkpoint).

An Internet Draft exists on the subject of firewall issues for IPv6 [fire00]. Problems include:

- Handling extension header chains, and unknown options.

- How to handle end-to-end IPsec sessions
- How to handle peer-to-peer applications
- Lack of availability of fully-featured stateful firewalls

At present the best (only) stateful firewalls exist for Linux and BSD filtering tools (e.g. ipfw).

7.3. Security implications of IPv6 transition mechanisms

There are security implications in every IETF RFC or Internet Draft, as described in the security considerations sections of these documents.

There may be particular risks or threats associated with specific types of protocols. In particular, transition tools may be generally liable to spoofing and denial of service attacks where tunneling methods are used, e.g. in the case of 6to4 relays [6to4sec00].

There is a general security risk in handling two concurrent versions of IP, as both need to be secured individually (and the security policy may not necessarily be the same for both), and in addition interrelationships of IPv4 and IPv6 may pose additional risks.

This area is being studied in 6NET in Work Package 2.

8. Case Study: An IPv6-only Site Deployment

The University of Tromsø and Invenia Innovation AS aimed at running experimental IPv6-only networks from as early as the start of 1999. The main motivation and aim was to live on the bleeding edge of technology and force network users, mainly computer science students, to investigate and learn about network configuration, management and programming. As IPv6 technology was still far from the comfortable world that IPv4 technology provided, switching to IPv6 without the comfort of a backup IPv4 service would force users to get involved again in the underlying technology that provided their connectivity and force them to investigate and fix problems along the way. This worked well over the past few years in that currently most problems in our basic IPv6-only infrastructure have been ironed out, and we are left with a more-or-less plug-and-play IPv6-only world.

This does not mean that every service, protocol and tool of the IPv4 world is also available in our IPv6-only world. But most of those that are still unavailable are of little interest to our users and administrators. In the following sections we will describe the main services, protocols and tools in active use as and on our IPv6-only network, but also those not used due to open problems or limitations. In addition, we will list some technologies we have not tested yet but would like to deploy on our networks.

8.1. Network Topology

Besides introducing IPv6 connectivity to the IPv4 infrastructure of the Department of Computer Science of the University of Tromsø and of Invenia Innovation, an IPv6-only research lab environment and MAN has been set up and is in active operation. The Lab, called Pasta-labben, provides the daily work environment for about 5 masters students and staff members, but its services are used by a far larger user community. Hence it needs to provide a fair level of production quality services. In addition, an experimental IPv6-only Lab has been setup that is used mainly for educational purposes as a tool in the University's undergraduate networking course.

The MAN is a WLAN based network connecting the home networks (residential area networks or RAN) of graduate students and staff members to the IPv6-only Lab and on to 6NET. The routers in this MAN only route IPv6 packets and only connect upstream to IPv6-only infrastructure. Users of the MAN are free to decide whether their RAN runs IPv6-only or consists of a dual-stack setup with IPv4 to IPv6 translators or proxies at its border router. TCP connectivity to the IPv4 world is provided using a TCP connection relay (an in-house variant of faithd) combined with the in-house developed DNS-proxy totd, running on the border router between the IPv6-only infrastructure and the dual-stack infrastructure. No connection or packet translation to the IPv4 world for other protocols, like UDP, is provided.

8.2. IPv6-only MAN core network

The Tromsø IPv6 MAN core network consists of a single 12-interface site-border router that provides upstreams IPv6 connectivity to 6NET, several routers deployed in the field that each interconnect a number of WLAN networks, and a larger number (currently about 15) of residential gateways. Every user of the MAN is required to setup their own residential gateway as part of the MANs operational policy. Apart from simplifying network management of the core network, it is an important factor in educating the user population on network technology and management in general and on IPv6 technology in particular. Finally, a special `protocol boundary gateway' provides the bridge between the IPv6-only network and the dual-stack departmental network. This gateway is an IPv6-only router in that it routes only IPv6 packets, but it has upstream IPv4 connectivity in addition which is used to provide a TCP translation service from IPv6 to IPv4.

The main operating systems in use on the routers are NetBSD and Cisco IOS. The main site-border router is a Cisco, and while small SOHO Cisco routers were used initially as connecting routers, most have been replaced by NetBSD-based PC-routers. The small Cisco routers are now used in more experimental settings and in the educational network laboratory. Due to its mature and early integrated IPv6 stack (from the Kame project), NetBSD has been the operating system of choice also for many residential gateways. However, the occasional FreeBSD and Linux-based system is also found as a residential gateway. Default linux installations proved less than ideal IPv6 router platforms with considerable robustness and standards conformance problems, such that Linux-based routers were less popular in the early days of the MAN. Nowadays, a Linux-based router especially one incorporating the work/patches of the USAGI project is a viable choice as residential gateway that exhibit few if any problems.

The main facility missing from the core network is proper support for remote network management and configuration. Considering the limited size of our network this poses no serious problems, but for a larger production style network this may be a major issue. Currently lacking is:

1. IPv6 transport support for SNMP (the Simple Network Management Protocol) which lies at the heart of most remote network monitoring and management applications. We currently have no alternative or acceptable work-around (other than tunneling IPv4 SNMP packets through a IPv6 SSH connection, for example) that is more attractive than living without SNMP in our network. The NET-SNMP project [netsnmp] has recently built IPv6 transport support, which may be explored in the near future.

2. In the scope of Wireless LANs, IPv6 transport support in the firmware of the WLAN radios for SNMP and/or HTTP to support remote monitoring and configuration. DHCPv6 server support to configure IP addresses to WLAN radios (needed once IPv6 transport is supported by them). Various tricks can be used to work around this limitation as the WLAN radios on our network are all directly connected to a PC router. A feasible approach in some cases is to let the router proxy for the radio boxes using a HTTP proxy. The router in this case communicates with the radios over IPv4 (on the local link only, so no IPv4 routing needed) or over a serial line.

3. NTP or similar time synchronization protocol support. Recently, initial IPv6 support in the xntp protocol has become available. However, not all protocol issues have been resolved yet and the support is not yet ubiquitous.

4. An open workable router configuration protocol of some kind. The router renumbering protocol [rfc2894] relies on IPSEC [rfc2401] for authentication, authorization and integrity and thereby on a trust model that does not easily apply to our `loose' administrative organization where each RAN constitutes its own administrative domain and new RANs typically get connected with little to no centralized control. Hence, currently routers are configured manually which is feasible regarding the current size of the network and the way its administrative control is distributed over its users. However, while for production networks significantly larger than ours this may pose a problem in theory, in practice such networks are typically managed and secured in a much more centralized way making this a non-issue. Still, interesting open issues lie in this area, especially when extremely large networks (which the IPv6 address space permits!) make centralized management and control infeasible.

5. Firewalling tools that avoid hard-coding IPv6 addresses to ease renumbering. Currently we use ipfilter on *BSD and Cisco access control lists to enforce network security policy, but currently we have no tools to rewrite such access control lists in case of a renumbering event. Other firewall features are missing, e.g. state support for UDP.

8.3. IPv6 Service Availability

A wide range of network management tools commonly used in IPv4 networks is available for IPv6 also. We can name available tools like ping6, traceroute6, tcpdump, ethereal, netperf, trafshow, and telnet amongst others but we will not describe them here nor try to list them exclusively. Instead, we list the range of application services provided on our IPv6-only network that are in daily operation and use without other problems than those of their IPv4 counterparts:

Service	**Application**
DNS	bind9
SMTP	sendmail

(A)POP	qpopper
IMAP	imap-uw
HTTP	apache, thhtpd
FTP	ftpd (NetBSD)
SSH	openssh, freessh
NFS	nfsd (NetBSD)
CDDB	freedb
CVS	cvs
RSYNC	rsync
SYSLOG	syslogd (NetBSD)
PRINTER	lpd (NetBSD)

Several of the above mentioned services run on a server host with one `leg' (network interface) in the IPv6-only world and one in the IPv4-only world. For some services (SMTP and PRINTER) this facilitates that the service also can provide proxy service for IPv6-only hosts to communicate with IPv4-only hosts. In addition to these `piggy-backed' proxy services, a pure DNS proxy or Application Level Gateway (ALG) service is provided that runs on a dual-stack host. This DNS proxy is the in-house developed `totd' that has the DNS additional task of rewriting DNS records to assist the transport relay transition mechanism deployed at the edge of the network.

Proxy service	**Applications**
DNS	Totd
SMTP	Sendmail
PRINTER	lpd (NetBSD)

Various other applications and services have been deployed and tested at various times on the IPv6-only network, most notably NNTP (news), IRC, Xpilot, Quake6. Most notable absent services are those we have available and in operation on our IPv4 network but are not available/accessible on the IPv6-only network. These are:

Missing Service
NTP
CVSUPD
SNMP
DHCP

LDAP, RTP, UUCP, KERBEROS, SMB, and various P2P protocols are amongst the protocols that we have not tested nor even tried to deploy on IPv6 as we did (do) not use these on IPv4 either and/or there is little interest in them among our user community so far.

8.4. Service Management Conflicts

The deployment of an IPv6-only network often requires gateway/proxy services at the dual-stack edge where it interconnects with the rest of the (IPv4-only) world. When (services deployed on) the IPv6-only infrastructure of an organization is under different administrative control than the IPv4-only infrastructure, problems may arise in enforcing policies for the dual-stack gateway services. Typically, the dual-stack service gateway will be under control of the IPv6 management staff while it should enforce policies formulated by IPv4 management.

An example we encountered is SMTP (email) relaying between the IPv6 and IPv4 world. Our goal is to transport email over IPv6 if possible, i.e. if the destination is reachable over IPv6). At the same time, the IPv4 management tries to enforce the policy that all email (transported over IPv4) leave the network via one of their main mail servers (for centralized implementation of email filtering policies, for example). It is common and easy to configure a mail server to forward all email to a central server.

But our mailserver software could not (easily) be configured to first try delivery over IPv6, and only when that was not possible to deliver to a central server. In our relatively open university environment we could `solve' this problem by requesting exemption from the central policies for the IPv6-only user community (when communicating with the IPv4 world), but this may not be an option in more restrictive corporate environments.

8.5. Residential Area Networks

The residential area networks vary greatly in number of machines (from as little as 2 to more than 10) and users, type of operating systems in use, technical skills of their user(s) and number of network applications in active use.

Although some of the residential area networks are IPv6-only (in addition to their IPv6-only upstream connectivity), many of them need to accomodate for IPv4-only users and/or applications on their home network. This requires them to setup IPv4 to IPv6 translation or proxy services on their residential gateway.

An alternative strategy is to tunnel the IPv4 traffic through the IPv6-only network via a dual-stack machine that does have upstream IPv4 connectivity. However, such a strategy is considered `cheating' and is actively discouraged and occasionally tolerated at best (although it is worth noting this is a principle also used by DSTM). As tunneling IPv4 in IPv6 is little different from tunneling IPv6 in IPv4, which is extensively discussed in other deliverables, we will omit a discussion of tunneling here.

As no globally-routable IPv4 addresses are assigned to the residential area networks and no IPv4 NAPT (Network Address and Port Translation) service is provided, translation techniques like NAPT-PT (NAPT with Protocol Translation) on the network layer or TRT (Transport Relay Translator) are not part of the solution space. As no implementations of such techniques for IPv4 to IPv6 translation are readily available, this is not a serious limitation.

The remaining possibility of deploying dual-stack proxy or application level gateway (ALG) services on the residential gateway is currently in use on most of these for a limited, but typically most heavily used by home users, set of application protocols. The table below lists these protocols and the applications that implement proxy service for them on most residential gateways.

Proxy service	**Applications**
HTTP, FTP	wwwoffle, www6to4, squid
SMTP	fetchmail, qpopper, sendmail
DNS	bind9, totd

8.6. Next Steps

There is still more work that could be done. For example, to improve the network infrastructure:

- Add smaller cells downtown to allow roaming laptop and PDA users and to facilitate useful deployment or trials of Mobile IPv6.
- Add redudant WLAN links in the MAN to experiment with failure scenarios and to experiment with host-based multihoming solutions.

We could also improve the service/protocol infrastructure:

- Deploy multicast-based application services like local radio streaming. AN example might be the Trondheim Underground Radio (TUR).
- IP-telephony/teleconferencing. This might include the Vocal VoIP package being ported elsewhere in 6NET.
- Deploy multicast-based infrastructure services like service location protocols.

The RANs could also be enhanced:

- Collect and learn from experiences from RAN users migrating to IPv6-ready OSes like Windows XP for the first time.
- Develop and/or deploy new proxy services for RTP, for example.

9. Conclusions

In this report we have identifed a number of issues that may hinder the successful deployment of IPv6 in a large scale. Some of the concerns are more pressing than others, and some may be relatively minor issues. However, for the “big picture” of IPv6 deployment we should address them all within the 6NET scope.

The IPv6-only deployment in Norway shows that in a small scale it is possible to use IPv6 only in a site network.

We will report more fully on these issues in 6NET Deliverable 2.5.2, due by July 2003.

10. References

Note that Internet Drafts expire (sometimes to new versions or RFCs), and are not stable references. Such drafts should be considered as works in progress.

[2553bis] “Basic Socket Interface Extensions for IPv6”, R.E. Gilligan et al, IETF Internet Draft, February 2002, draft-ietf-ipngwg-rfc2553bis-05.txt

[6mess01] “Moving from 6bone to IPv6 Internet”, P. Savola, IETF Internet Draft, November 2002, draft-savola-v6ops-6bone-mess-01.txt

[6to4sec00] “Security Considerations for 6to4”, P. Savola, IETF Internet Draft, October 2002, draft-savola-v6ops-6to4-security-00.txt

[anvl] Ixia ANVL IPv6 conformance test suite, http://www.ixiacom.com/products/caa/anvl_IPv6.php

[appasp] “Application Aspects of IPv6 Transition”, M. Shin et al, IETF Internet Draft, October 2002, draft-shin-v6ops-application-transition-00.txt

[d451] “Report on IETF Multihoming Solutions”, 6NET Project Deliverable, M. Dunmore editor, October 2002, http://www.6net.org/publications/deliverables/D4.5.1.pdf

[dhc] IETF Dynamic Host Configuration WG, http://www.ietf.org/html.charters/dhc-charter.html

[dhcpv6] “Dynamic Host Configuration Protocol for IPv6 (DHCPv6)”, R. Droms et al, IETF Internet Draft, November 2002, draft-ietf-dhc-dhcpv6-28.txt

[dnsext] IETF DNS Extensions WG, http://www.ietf.org/html.charters/dnsext-charter.html

[fire00] “Firewalling Considerations for IPv6”, P. Savola, IETF Internet Draft, September 2002, draft-savola-v6ops-firewalling-00.txt

[flow03] “IPv6 Flow Label Specification”, J. Rajahalme et al, IETF Internet Draft, September 2002, draft-ietf-ipv6-flow-label-03.txt

[harm01] “RFC 3041 Considered Harmful”, F. Dupont and P. Savola, IETF Internet Draft, June 2002, draft-dupont-ipv6-rfc3041harmful-01.txt

[ipv6] IETF IPv6 WG, http://www.ietf.org/html.charters/ipv6-charter.html

[ipv6forum] The IPv6 Forum, http://www.ipv6forum.org/

[ipv6mh] IPv6 Multihoming Group, http://arneill-py.sacramento.ca.us/ipv6mh/

[issues00] “Unidentified issues in IPv6 deployment operation”, J. Hagino, T. Jinmei, IETF Internet Draft, June 2002, draft-itojun-jinmei-ipv6-issues-00.txt

[kport] “Implementing AF-independent application”, Jun-ichiro itojun Itoh, 1998-2002, KAME Newsletter, http://www.kame.net/newsletter/19980604/

[llmnr13] “Linklocal Multicast Name Resolution (LLMNR), L. Esibov et al, IETF Internet Draft, November 2002, draft-ietf-dnsext-mdns-13.txt

[lport] “Guidelines for migration of collaborative work applications”, LONG Project Deliverable D3.2, revised, Eva Castro et al, June 2002, http://long.ccaba.upc.es/long/040Deliverables/083dwnLONG-D32A.pdf

[manet] IETF Mobile Ad-hoc Networks WG, http://www.ietf.org/html.charters/manet-charter.html

[mipv6] “Mobility Support in IPv6”, D. Johnson et al, IETF Internet Draft, October 2002, draft-ietf-mobileip-ipv6-19.txt

[mldv2] “Multicast Listener Discovery Version 2 (MLDv2) for IPv6”, R. Vida et al, IETF Internet Draft, October 2002, draft-vida-mld-v2-05.txt

[mobileip] IETF IP Routing for Wireless/Mobile Hosts WG, http://www.ietf.org/html.charters/mobileip-charter.html

[multi01] “IPv6 Multicast Deployment Issues”, P. Savola, IETF Internet Draft, November 2002, draft-savola-v6ops-multicast-issues-01.txt

[multi6] IETF Site Multihoming in IPv6 WG, http://www.ietf.org/html.charters/multi6-charter.html

[ndissues] “IPv6 Neighbor Discovery trust models and threats”, P. Nikander, IETF Internet Draft, October 2002, draft-ietf-send-psreq-00.txt

[nemo] IETF Network Mobility WG, http://www.ietf.org/html.charters/nemo-charter.html

[netsnmp] NET-SNMP Project, http://net-snmp.sourceforge.net/

[ngtrans] IETF Next Generation Transition WG, http://www.ietf.org/html.charters/ngtrans-charter.html

[ntp] The NTP Information Site, http://www.ntp.org/

[rfc2133] “Basic Socket Interface Extensions for IPv6”, R. Gilligan et al, IETF RFC2133, April 1997, http://www.ietf.org/rfc/rfc2133.txt

[rfc2165] “Service Location Protocol”, J. Veizades et al, IETF RFC2165, June 1997, http://www.ietf.org/rfc/rfc2165.txt

[rfc2292] “Advanced Sockets API for IPv6”, R. Stevens and M. Thomas, IETF RFC2292, February 1998, http://www.ietf.org/rfc/rfc2292.txt

[rfc2373] “IP Version 6 Addressing Architecture”, R. Hinden and S. Deering, IETF RFC2373, July 1998, http://www.ietf.org/rfc/rfc2373.txt

[rfc2401] “Security Architecture for the Internet Protocol”, S. Kent, R. Atkinson, IETF RFC2401, November 1998, http://www.ietf.org/rfc/rfc2401.txt

[rfc2460] “Internet Protocol, Version 6 (IPv6) Specification”, S. Deering and R. Hinden, IETF RFC2460, December 1998, http://www.ietf.org/rfc/rfc2460.txt

[tfc2553] “Basic Socket Interface Extensions for IPv6”, R, Gilligan et al, IETF RFC2553, March 1999, http://www.ietf.org/rfc/rfc2553.txt

[rfc2710] “Multicast Listener Discovery (MLD) for IPv6”, S. Deering et al, IETF RFC2710, October 1999, http://www.ietf.org/rfc/rfc2710.txt

[rfc2894] “Router Renumbering for IPv6”, R. Crawford, IETF RFC2894, August 2000, http://www.ietf.org/rfc/rfc2894.txt

[rfc2960] “Stream Control Transmission Protocol”, R. Stewart et al, IETF RFC2960, October 2000, http://www.ietf.org/rfc/rfc2960.txt

[rfc3041] “Privacy Extensions for Stateless Address Autoconfiguration in IPv6”, T. Narten and R. Draves, IETF RFC3041, January 2001, http://www.ietf.org/rfc/rfc3041.txt

[rfc3363] “Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System (DNS)”, R. Bush et al, IETF RFC3363, August 2002, http://www.ietf.org/rfc/rfc3363.txt

[scope04] “IPv6 Scoped Address Architecture”, S.Deering et al, IETF Internet Draft, June 2002, draft-ietf-ipngwg-scoping-arch-04.txt

[single3] The Single UNIX Specification, Version 3, http://www.unix.org/version3/

[sitedns] “Well known site local unicast addresses to communicate with recursive DNS servers”, A. Durand et al, IETF Internet Draft, draft-ietf-ipv6-dns-discovery-07.txt

[sockapi] “Advanced Sockets API for IPv6”, W. Stevens and E. Nordmark, IETF Internet Draft, October 2002, draft-ietf-ipngwg-rfc2292bis-08.txt

[send] IETF Securing Neighbor Discovery WG, http://www.ietf.org/html.charters/send-charter.html

[sport] “Porting Networking Applications to the IPv6 APIs”, Sun Microsystems, 1999, http://wwws.sun.com/software/solaris/ipv6/porting_guide_ipv6.pdf

[tahi] The TAHI Project, http://www.tahi.org/

[v6ops] IETF IPv6 Operations WG, http://www.ietf.org/html.charters/v6ops-charter.html

[v6pd] “IPv6 Prefix Options for DHCPv6”, O. Troan and R. Droms, IETF Internet Draft, November 2002, draft-ietf-dhc-dhcpv6-opt-prefix-delegation-01.txt