Project Number:	**IST-2001-32603**
Project Title:	**6NET**
CEC Deliverable Number:	**32603/UOS/DS/2.5.3/A1**
Contractual Date of Delivery to the CEC:	31st May 2005
Actual Date of Delivery to the CEC:	16th June 2005
Title of Deliverable:	D2.5.3: Final Report on IPv6 Deployment Issues (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)
Reviewer:	David Mills (University of Southampton)
Contributors:	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 a discussion and summary 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. This document updates and is in effect a final version of Deliverables D2.5.1 and D2.5.2.

Keywords:

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

Executive Summary

There has been much effort made in the first half of the decade in standardisation bodies (principally the IETF) and from vendors (operating system, router and to a lesser extent application) 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, and we see IPv6 used on a regular basis. While the issues may be shrinking in number, it is important that they be documented and addressed.

In this report we highlight the issues we have identified, as they have arisen within the scope of the 6NET project. This means issues identified since 2002, some of which are ongoing, many of which have been addressed, either completely or to some extent.

We do not include in detail issues of transition from IPv4 to IPv6 per se; such interworking issues are reported in the final versions of the site and ISP transition cookbooks that the 6NET project has produced within Work Package 2, see [d224] and [d234v2].

The initial report D2.5.1 [d251] was aimed at scoping the open issues. These, and progress on them (by the project or by external bodies) was updated in D2.5.2 [d252], and this document represents the final version of the summary of the statuses.

These IPv6-specific issues are being addressed, and are in various stages of resolution. We do not consider the political or social issues in IPv6 deployment in this report. Instead we focus on technical and standards-related issues.

The political and social issues are being addressed by the European IPv6 Task Force [ipv6tf], by National IPv6 Task Forces and by the IPv6 Forum [ipv6forum].

In the initial report D2.5.1 we presented a brief case study of a 6NET site that is already running IPv6-only. The focus for actual deployment in 6Net enterprises (campuses) has been dual-stack, but this is done with the operation of an IPv6-only node in that environment in mind, i.e. the services are deployed such that an IPv6-only node could operate successfully.

IPv6 is generally ready for deployment, but the issues raised in this document need some consideration for wide-scale adoption (where IPv6 is used on a day-to-day basis for everyday applications and services) to occur.

It should be noted that a number of the issues that were raised in [d251] have been advanced within the IETF or elsewhere, and largely resolved. The conclusions section highlights this progress.

Table of Contents

1. Introduction

IPv6 is beginning to be deployed in production ISP and site networks around the world. The 6NET project deployed an international IPv6-only backbone spanning some 15 National Research and Education Networks (NRENs) in 2002. With 6NET completing, IPv6 connectivity between NRENs is available through the pan-European research network backbone GÉANT [geant], has now also deployed IPv6 natively, by the dual-stack approach. Through the work of DANTE and the NRENs in the 6NET project, the testing and deployment of the GÉANT infrastructure was significantly accelerated.

Most of the 6NET partners are running quite extensive national-scope IPv6 networks, though invariably dual-stack IPv4/IPv6 (described in the 6NET ISP transition cookbooks, concluding in [d224], rather than IPv6-only. The partner universities have early IPv6 deployments in place, with one exception (Tromso, described later in this text) also dual-stack. Southampton (UK) is an example of a production dual-stack deployment (spanning over 1,000 hosts and 1,500 users), documented in [d234v2].

The deployment of a dual-stack infrastructure spanning GÉANT and many NRENs allows the efficient connection of end sites (universities), such that the performance and behaviour of the IPv6 routing and transport closely matches that of IPv4. The challenge now is validation of applications and services (undertaken in Work Packages 3, 4 and 5 in 6NET), of management tools (Work Package 6) and identification and promotion of early IPv6 deployment scenarios (which we discuss in Section 8 of this document).

In Deliverable D2.4.2 we cover IPv6-specific issues for deployment of IPv6-enabled (IPv6-only or dual-stack) Wireless LANs. Many of the issues reported there apply generically to site IPv6 deployments (wired or wireless).

As mentioned above, one 6NET site, Tromso in Norway, has an IPv6-only network spanning a whole department; this was described in the previous instance of this Deliverable, D2.5.1 [d251] and updated notes are presented later in this text. 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. These may complement dual-stack 'production' deployments. Note that in this case "IPv6 only" may mean a device with a hybrid IPv4/IPv6 stack, but with only the IPv6 networking and routing information configured.

In this document we present a list of ongoing 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 production-quality adoption of IPv6 can be successful. Many are in the process of being tackled, and that progress is captured in the text. We focus on technical and standards-oriented issues, rather than the political and social issues that are addressed by the European and national IPv6 Task Forces [ipv6tf] and the IPv6 Forum [ipv6forum], to which many 6NET partners contribute.

Note that this Deliverable is the final version of a deliverable that updates Deliverable D2.5.1 [d251] and D2.5.2 [d252].

2. Vendor and Standards Support

A key basic requirement for IPv6 deployment is IPv6 support in host and router stacks. The state of implementations has improved significantly since the 6NET project started, to the state where almost all vendors now have commercial support for IPv6 functionality.

2.1. Host OS support

Host operating system support includes commercial and open source solutions, and well as those for both embedded and PDA systems (which have some overlap).

2.1.1. Commerical operating systems

Host operating system support includes the following vendors:

- Apple, Mac OS/X: http://www.apple.com/
- Compaq (Tru64): http://www.compaq.com/ipv6/ (discontinued)
- HP-UX 11i IPv6: http://www.hp.com/
- IBM AIX: http://www.ibm.com/software/ipv6/
- Microsoft Windows (XP, .NET, and 2003 Server): http://www.microsoft.com/ipv6/
- SGI Irix (from 6.5.19): http://www.sgi.com/
- Sun Solaris (8, 9 and 10): http://www.sun.com/

2.1.2. Open source operating systems

Host operating system support includes the following open source systems:

- FreeBSD: http://www.freebsd.org/
- KAME stack (*BSD): http://www.kame.net/
- Linux (USAGI): http://www.linux-ipv6.org/
- Linux (general): http://www.bieringer.de/linux/IPv6/IPv6-HOWTO/IPv6-HOWTO.html
- NetBSD: http://www.netbsd.org/

2.1.3. Functionality

The level of support in the various operating systems is mixed. The BSD variants, thanks to the KAME stack integration, have very good IPv6 support “out of the box”.

Many Linux flavours (e.g. SuSE, Mandrake, RedHat, Debian, Gentoo and so forth) now have good support, helped in part by the USAGI project. IPv6 support (as an Experimental option) had to be configured in the stable 2.4.x Linux kernel, but as of the 2.6 kernel series is included by default.

Solaris support has been good since Solaris 8, carried into Solaris 9 and 10. The IPv6 support can be enabled during the OS installation, or at a later date by an administrator.

Perhaps most important, at least from a user base perspective, is support in Microsoft Windows platforms; in XP IPv6 can be turned on by running "ipv6 install" at a command window prompt. The next version of Windows (Longhorn) is expected to enable IPv6 by default. Windows XP ships with ISATAP and 6to4 transition tool support, and also actively uses RFC3041 privacy extensions [rfc3041]. Windows 2000 is still prevalent in everyday home and office deployment, 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. Older Windows versions have no official support.

2.1.4. PDA devices

IPv6 support for PDA devices includes:

- Familiar Linux (on a Compaq iPAQ for example), http://familiar.handhelds.org/
- OpenZaurus (for the Sharp Zaurus), http://www.openzaurus.org/
- Windows Pocket PC 2003, http://www.microsoft.com/

The status of Microsoft support for Windows-based PDAs is particularly important.

2.1.5. Embedded systems

There have been a number of releases of embedded operating systems with IPv6 support. These include:

- Elmic Systems, http://www.elmic.com/
- Embedded Linux, http://www.embedded-linux.org/
- Symbian OS (as used on the Sony Ericsson P800/900 and Nokia 9500 Communicator handsets), http://www.symbian.com/
- Windriver , http://www.windriver.com/

The Nokia 9500 Communicator is an interesting device – it supports IPv6 and WLAN operation, with the included web browser supporting IPv6 by default. A community-modified version of the Putty ssh client allows IPv6 secure shell access on a 80x24 display on the handset.

2.1.6. Comments

It is worth highlighting also the TAHI project [tahi] and the Ixia test suite [anvl] that can be used to test IPv6 conformance on a host system. ETSI and the University of New Hampshire organise various IPv6 "plug fest" events (though specific results are not rteported, participation is generally open), and the IPv6 Forum now has an "IPv6 Ready" programme [v6ready], entering its second phase at the time of writing.

2.2. Router support

IPv6 router support comes from both commercial and open source solutions.

2.2.1. Commercial router support

Router support includes the following vendor platforms:

- 6WIND: http://www.6wind.com/ (available now only as software, not a hardware product)

- Cisco IOS: http://www.cisco.com/
- Ericsson Telebit: http://www.ericssontelebit.dk/ (now obsolete)
- Hitachi: http://www.internetworking.hitachi.com/
- ipinfusion (ZebOS): http://www.ipinfusion.com/
- Juniper (JUNOS): http://www.juniper.net/
- Nortel Networks: http://www.nortelnetworks.com/

2.2.2. Open source router support

Router support includes the following open source solutions:

- FreeBSD: http://www.freebsd.org/
- MRT (Multi-threaded Routing Toolkit): http://www.mrtd.net/
- NetBSD: http://www.netbsd.org/
- Xorp: http://www.xorp.org/ (a research platform)
- Zebra: http://www.zebra.org/

2.2.3. Functionality

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.

2.2.4. Enterprise router-switch platforms

There are a number of platforms that blend Layer 2 and Layer 3 functions for the enterprise network. Vendors of such equipment with IPv6 product include:

- Cisco: http://www.cisco.com/ (e.g. the 6509 and 3750 products)
- Foundry: http://www.foundrynet.com/technologies/ipv6/
- Extreme: http://www.extremenetworks.com/ (e.g. the Black Diamond 10K)
- Procket: http://www.procket.com/

General requirements in the enterprise network will include IPv6 routing, VLAN support for IPv6, IPv6 packet filtering and multicast support (e.g. PIM-SM routing, embedded RP [embedrp] and MLD snooping). We are not aware of any system yet offering a full set of enterprise features (the Cisco product is the most advanced, but lacks MLD snooping, which is just entering the EFT stage at the time of writing).

2.2.5. IPv6 Multicast support

A number of vendors offer multicast routing capabilities, including Cisco, Juniper, Hitachi and 6WIND, in addition to open source solutions including BSD and Xorp.

Features implemented vary, e.g. for PIM-SM, SSM, and for MLD snooping. There is some support for handling of multicast routes (e.g. Cisco and Juniper offer MBGP). Cisco and Juniper have implemented the embedded RP method for inter-domain PIM-SM [rfc3956].

There are not any IPv6 multicast MIBs available at the time of writing.

Hardware support for multicast is also important, both for routing and snooping. The status will almost certainly vary with specific equipment ranges from vendors.

2.3. IPv6 Support in Programming Languages

The major programming environments supporting IPv6 are currently:

- C - uses the new APIs and data structures for TCP/IP sockets as defined in the Basic Socket Extensions for IPv6 [rfc3493] and the Advanced Socket API for IPv6 [rfc3542]
- Java - the Java Development Kit (JDK) as of version 1.4.0 supports basic IPv6 functionality for Linux and Solaris platforms [jdkv6]. MS Windows support was made available in JDK1.5.
- Perl - there is support in Perl 5 where the underlying system supports IPv6. If you want to access IPv6 specific or AF neutral functions you should use the IO-INET6 module[1] and/or Socket6 module[2].
- Python - the standard Python interpreter since python 2.3.4 natively supports IPv6, but unfortunately the binaries for Microsoft Windows available on the python.org website do not have compiled-in IPv6 support. Starting from the upcoming 2.4 release of the Python interpreter, the binaries for Microsoft Windows will be built with IPv6 support. To access the advanced socket API for IPv6 (RFC 3542) you can use the Python Socket Module Extension[3].
- Ruby - Ruby 1.6 and Ruby 1.8 support IPv6 natively if the host operating system supports it. The Unix version compiles with IPv6 support, while there is also a Windows version available[4].

Programming issues are described in the Applications section below.

2.4. IPv6 IETF Standards Working Groups

The major IPv6-related IETF Working Groups are listed below. In late 2002 the ipng WG was renamed the ipv6 WG, and the ngtrans WG was phased out and replaced by v6ops. These changes were 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]

[1] http://www.cpan.org/modules/by-module/IO/INET6-2.01.tar.gz

[2] http://www.cpan.org/modules/by-module/Socket6/Socket6-0.18.tar.gz

[3] http://www.pps.jussieu.fr/~ylg/PyXAPI/

[4] http://win6.jp/index.html

- Mobile Ad-hoc Networks [manet]
- IP Routing for Wireless/Mobile Hosts [mobileip] (superceded by mip4 and mip6)
- MBONE Deployment [mboned] (for IPv6 Multicast deployment)
- Mobility for IPv6 [mip6]
- Mobile IPv6 Signalling and Handoff Optimisation (mipshop)
- Site Multihoming in IPv6 [multi6] (superceded by shim6)
- Network Mobility [nemo]
- Securing Neighbor Discovery [send] (completed)
- Shim 6 [shim6]

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

The Mobile IP [mobileip] IETF WG split into the Mobility for IPv4 WG, the Mobility for IPv6 WG [mip6] and the Mobile IPv6 Signalling and Handoff Optimisation WG in 2003.

In 2004, the [multi6] WG completed its work and spun off the [shim6] WG.

The [send] WG completed its charter in late 2004 – producing [rfc3971] and [rfc3972] - and has since been disbanded.

An IPv6 Tunnel Configuration (v6tc) BoF was held at IETF 62, but no WG has officially been formed from this BoF yet.

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

The 6LINK project [6link] has produced IETF IPv6 Standardisation reports three times each year since 2002, which have been published on the IST IPv6 Cluster web site [ist-ipv6]. 6LINK completed in March 2005, so further updates may not be available.

2.5. 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 (which probably means with IPv4 present in the hybrid OS but not enabled). In most cases the basic IPv6-only operation works out of the box. That said, depending on the precise flavour/install, the most common problem that surfaces is use of and access to the IPv4 loopback device, which may need to be left enabled, even though IPv4 communication and configuration is otherwise disabled.

Solutions offering IPv6 transport for DNS are becoming available on some host platforms to support name resolution over IPv6, e.g. in BSD, but not in Windows XP.

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 that is only available over IPv4 SNMP connections, or for export of network flow information.

In a later section we describe potential early IPv6 deployment scenarios, which may include IPv6-only operation. While the 6NET backbone links (between backbone routers) are IPv6-only, most site deployments are heavily dual-stack, with some exceptions including the Tromso network, as described in Deliverable D2.5.1 [d251] and commented on later in this text.

We should recognise that in the migration to an "end game" of IPv6 networking, there will be a long coexistence period with IPv4. Topics such as automatic tunnelling of IPv4 data over IPv6 networks will only likely come to the mainstream in the medium to long term.

In the next section, we look at the status of technical issues identified within the 6NET project that affect the deployment and adoption of IPv6.

3. Technical Issues Affecting IPv6 Deployment

In the course of attending a variety of events and meetings, and from 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, technical, standards-related issues, although in some cases issues of policy also arise.

In many cases, the specific issues have been captured and presented in IETF Internet Drafts by the project participants. Examples include 6to4 security [rfc3964], global routing stability [6mess], firewalling [fire], multicast [multicast], transition architectures [transarch] and transition security [transsec].

These drafts have helped focus discussion and progression on the issues. Some of these drafts have expired as the issues are resolved, others have made RFC status as informational documents, e.g. [rfc3964]. Each has been a useful discussion point of focus for the issue concerned.

There are also other publications of IPv6 deployment issues, e.g. an 'IPv6 Barriers' publication from the EC IPv6 TF [ipv6tf] and an IETF Internet Draft by Hagino and Jinmei [issues]. A new Japanese initiative also tracks deployment issues [ipv6fix]. It is worth noting that the 'issues' list is a shifting target.

The 6NET deliverable on IPv6 Wireless LAN deployment D2.4.2 [d242] includes many IPv6-specific deployment considerations (including ways of thinking for managers used to operating IPv4 networks, rather than just 'missing pieces').

3.1. Categories of deployment issues

There is no particular order of importance to the topics discussed in this document. In the following sections we categorise the issues (somewhat arbitrarily, but categorising the issues aids their presentation) into general areas.

Note that the 'issues' may not all be actual barriers to deployment as such, but are cited by as issues in discussions in various forums.

- Network robustness and performance (which must be of good quality for users to be encouraged to use IPv6 instead of IPv4, whether they are aware they are using it or not)

  - General IPv6 routing stability
  - Preference for use of IPv4 or IPv6 connectivity
  - Interworking between IPv4 and IPv6 systems
  - IPv6 multihoming
- IPv6 and DNS
  - IPv6 transport issues
  - Registry-related issues
  - Gaps in standards
- Network management
  - IPv6 and SNMP
  - Management of devices over IPv6 transport
  - Service discovery methods
  - Transition management
  - Performance measurement
  - IPv6 Prefix Delegation (assignment to, for use by, routers)
  - IPv6 network renumbering
  - NTP services
- Application and IPv6-specific features (of concern to application developers and end users)
  - Site-local addressing requirements
  - IPv6 code porting and development
  - IPv6 multicast support
  - IPv6 deployment over ADSL (SOHO environments)
  - Other service support
  - Available and missing applications
  - Use of the IPv6 Flow Label
  - IPv6 Privacy Extensions issues
- Security considerations
  - IPv6 IPsec
  - IPv6 firewalls
  - Network Address Translation (NAT)
  - Stateless autoconfiguration
  - Transition security issues

IPv6 is quite advanced in terms of standardisation in the IETF. The core protocols have been set since around 2000. However, two of the core specifications are undergoing revision as a result of deployment experience:

- RFC2461-bis [rfc2461bis]
- RFC2462-bis [rfc2462bis]

Both revisions are now pretty much completed at the time of writing.

The production and availability of stable IPv6 standardisation documents is important to give implementers the confidence to put more effort into developing and releasing code.

The standardisation of IPv6 has taken significant steps forward with the publication of 'advanced' protocol standards. For example, in 2004 we saw the publication of DHCPv6 [rfc3315], Staless DHCPv6 [rfc3736] and Mobile IPv6 [rfc3775] and [rfc3776].

As mentioned above, the state of IPv6 standardisation has been captured in Standardisation Reports issued every four months by the 6LINK project [6link]; the last report of the project was released in March 2005 [ist-ipv6].

4. Network Robustness and Performance Issues

4.1. General IPv6 Routing Stability

While routing of IPv6 traffic is generally good in local regions (and is within the NREN networks that interconnect through GÉANT, for example), international IPv6 connectivity, particularly outside the scope of the academic research networks, is generally not as reliable or predictable, 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 while the key commercial transit providers themselves are being slow to offer that transit.

The routing problems are discussed in [6mess]. Partly as a result of this draft, a phase-out plan was drawn up for the 6bone [rfc3701] whereby no new 6bone pTLA allocations have been made since January 1st 2004, and no 6bone prefixes should be routed after June 6th 2006 (06/06/06). The use of 6bone prefixes is already significantly reduced since this plan was published.

Since late 2001 6NET has cooperated with Euro6IX, Abilene (the US research network) and WIDE (in Japan) through discussions at IETF meetings to establish good routing policy to build a predictable, efficient and well-performing IPv6 infrastructure for research, such that IPv6 can more readily be used for day-to-day activity and work. These discussions have proven useful in also helping to determine appropriate policy for GÉANT.

There is now good IPv6 native (dual-stack) connectivity between the European NRENs and GÉANT, and from GÉANT to the Abilene, ESnet and Canarie international networks. With native IPv6 deployed across Abilene, and many US regional networks also deploying, it is now

possible for IPv6 routing from European universities to follow a similar (if not the same) path to IPv6-connected US universities, with very similar traffic properties (e.g. round trip time).

Stable, efficient routing outside of these domains requires a number of measures to be adopted, not least native transit provision by the bigger commercial players, and removal of the problematic long distance (including transatlantic) "open" tunnels. This is slowly happening.

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 a timeout/delay that is heavily OS/application specific, but anecdotally seems to range from instant to 40-60 seconds, in some cases not falling back at all).

However, given the observed routing issues described above, preference for use of IPv6 rather than use of IPv4 will more often than not give a worse experience for the user.

The preference and fallback issue has been studied in an Internet Draft [v6on] that showed how the timeout and fallback behaviour varied across various stacks and operating systems.

It is also the case that assuming IPv6 connectivity by the presence of a DNS AAAA record is not a reliable indicator. The target network may not be globally well connected. Even if some host services on a host are IPv6-enabled, they may not all be, and thus an IPv6 connection may be attempted to an IPv4-only service on the target dual-stack host.

Methods to allow per-host or per-application selection of IP version preference may be viewed by some to be desirable. Discussions about this issue have shown that, though being able to choose a preference (IPv4 over IPv6 or similar) might be desirable in terms of performance, it is generally not a good idea to leave these low level choices to the user. Most users lack the skills to make such choices and in a lot of cases it is probably not productive to overload options of applications with these settings.

Given that multiple versions of IP were not present when the initial timeout behaviours for TCP were defined, some reconsideration of the TCP behaviour may be desirable. However, the situation should improve as IPv6 internetworking becomes more reliable and more (all) host services can be offered over both protocols.

4.3. Interworking between IPv4 and IPv6 systems

IPv6 deployment in 6NET networks (NRENs) and sites (universities) has been almost universally dual-stack; the only exceptions are the parallel German IPv6 backbone (6WiN) and the site deployment at Tromso.

As a result, IPv4-IPv6 interworking is simplified, since all systems can speak both protocols. Where both protocols are available, the applications or services can select their IP version preference.

In cases where IPv4-only and IPv6-only systems need to communicate, some form of translation is required, as discussed in [d234v2]. The translation can be done at the network layer (e.g. NAT-PT), the transport layer (e.g. TRT) or application layer (an ALG). Consensus in the transition studies

was to use ALGs wherever possible. NAT-PT is a method of last resort, and indeed is proposed for deprecation in the IETF [natdepr].

One may argue that 6NET has somewhat sidestepped the IPv4-only to IPv6-only communication issue by sites deploying dual-stack. However, these deployments have been made on the basis of analysis of the options, and availability of features from open source and commercial platforms. A dual-stack approach is the method that we see being dominant for the short to medium term, even where the IPv4 nodes of the deployment use IPv4 NAT (in conjunction with global IPv6 addresses).

In summary, a dual-stack approach – for new and updated deployments - alleviates the concern. For legacy IPv4-only interoperation with IPv6-only services, an ALG approach would be preferred over NAT-PT.

4.4. Multihoming and Provider Independence with IPv6

For IPv6 to be adopted by sites and ISPs already accustomed to being multihomed with IPv4, an IPv6 multihoming solution is required. Currently, an IPv4 site would typically have its own Provider Independent (PI) address space, and have that address space advertised by two (or more) upstream ISPs. Alternatively, a secondary ISP may advertise a more-specific block from the primary ISP's address space. Both approaches in effect 'bloat' the Internet default free zone (DFZ) routing table size, because additional prefixes are being advertised that ordinarily would not be.

There have been over 40 Internet Drafts proposed that have been related to potential IPv6 multihoming solutions in the last five years or so, but these have almost universally failed to reach RFC status with the exception of SCTP [rfc2960]. A 6NET Project deliverable has been produced that overviews the status of IPv6 multihoming solutions as of February 2005 [d453].

A number of people consider the lack of an IPv6 multihoming solution a significant factor that is slowing IPv6 deployment for large enterprises for who resilience (and freedom from being tied to the address space of a single provider) is seen as important.

The basic problem is that punching out /48-sized IPv6 prefixes into the IPv6 DFZ, akin to current IPv4 multihoming practice, will not scale. Running "classic" IPv6 multihoming by multiaddressing, where every host in a site can inherit an IPv6 address tied to each site connectivity provider, also has complexities (e.g. ISP ingress filters on source addresses). As a result, a number of host and router-based solutions have been proposed, including methods that split the identifier and locator (routing) space.

The IETF multi6 WG was revived in 2003/04, and has progressed to a conclusion with recommendations for future work that will be undertaken by the SHIM6 [shim6] WG, where the locator-identifier approach looks likely to be adopted as a longer-term solution. This does not solve the immediate problem for medium to large enterprise sites however.

Recently, a proposal was made to ARIN for sites to acquire IPv6 PI address space, on the basis that they could qualify for an ASN from the Regional Registry. This proposal is still open, and does not seem to have progressed. It is unlikely to help small to medium size sites, even if adopted.

Thus at present IPv6 sites must use their ISP's PA address space, unless they themselves can obtain LIR status and obtain their own /32 size IPv6 (essentially PI) prefix. IPv6 multihoming solutions from the IETF are still on the long-term radar.

5. IPv6 and DNS

The Domain Name System/Service (DNS) is a critical piece of Internet infrastructure technology. Successful deployment of IPv6 requires IPv6 considerations of DNS to be identified and where necessary resolved.

We can loosely view the IPv6 issues as falling into three categories: standards, transport and registry issues.

Early in the IPv6 standardisation process, an option existed to represent IPv6 prefixes with special A6 DNS records. This became a bit of a 'religious' discussion point, with many arguing that A6 added unnecessary complexity and scope for administrators to 'shoot themselves in the foot'. Although the AAAA vs A6 issue has been resolved, with A6 and the associated DNAME moving to Experimental status [rfc3363] (and thus in effect Historic), a number of DNS-related issues remain.

The DNS issues are captured and discussed in an IETF Internet Draft [dnsissues].

Another IETF Draft exists for transport guidelines in IPv6 DNS [rrfc3901].

5.1. IPv6 standards issues

The following areas need further work in the IETF:

- 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 [dnsdisc], implemented by Microsoft, but this is now in effect deprecated due to the IETF's deprecation in turn of site local addressing. Solving this bootstrapping problem is an important issue for IPv6 "plug and play" deployment. There is currently a split in the IETF between using DHCPv6 or a modification to Router Advertisements. The specification for 'stateless DHCPv6' [rfc3736] does not require the maintenance of IP lease information; it is designed to respond to requests only for other configuration option data, like DNS search path or resolver IP addresses. The RA method has the advantage of multicasting the information. Currently, the DHCPv6 approach seems preferred.
- Within the DNS protocol (RFC1034) there is a stipulation of a 512-byte limit for payloads. IPv6's longer addresses may thus result in truncated responses, in particular where A6 records (now deprecated) or DNSSEC are used. This issue is discussed in RFC3226 [rfc3226], where the EDNS0 solution is recommended (RFC2671). This requires changes to installed IPv6 client systems.
- There is no method yet for populating the reverse DNS data for a network, particularly for statelessly autoconfiguring hosts. Such reverse lookups are used as a weak authentication is some instances, e.g. by sendmail to accept SMTP from local hosts. One IETF Draft has recently been produced in this area [dnsrev].
- There is also a requirement for name registration of communication endpoints, and for updates of that information to be available. Dynamic DNS exists as a solution in this space, as does ENUM [rfc3761] and the use of Mobile IPv6. This is not uniquely an IPv6 issue.

5.2. IPv6 transport for DNS

The following areas are related to IPv6 transport:

- Not all operating systems offer IPv6 transport lookups. At present at least BSD and Linux systems can talk to DNS servers over IPv6, while Windows XP requires a port proxy to be installed. Windows Server 2003's DNS client and server supports it, to enable it you just have to execute 'dnscmd /Config /EnableIPv6 1'.
- There is no widespread availability of IPv6-enabled root DNS servers. Some servers have been enabled as a first step.
- The address blocks allocated for use by root DNS servers are generally likely to be /48 sized, as per delegations to IX nodes. It is thus important that these prefixes are not filtered by the general '48 catchall' filters that exist on the IPv6 Internet today. ISPs should list and offer exceptions for these special prefixes.

5.3. Registry issues for DNS

The following areas are registry and registrar related issues:

- You cannot generally 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. This situation is beginning to change however, e.g. Nominet allows such registrations in the UK.
- Perhaps more importantly, users generally do not have a way today to add AAAA records for their own domains where those domains are managed by a 3rd party registrar service.
- The ip6.int to ip6.arpa reverse delegation transition process has been slow, especially for 6bone address space under 3ffe::/16. There is now a firm recommendation to deprecate ip6.int [intdepr].
- There is no agreed process for reverse DNS lookup delegation under 2002::/16, the 6to4 transition address space. However, at least one proposal for a solution has been made, see [rev6to4].

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. The main issue is the local discovery of an IPv6 DNS resolver; many feel that deploying DHCPv6 to provide the information is heavyweight, but as yet there is no real alternative on the table. Other issues are being progressed; the deprecation of ip6.int seems to be taking longer than it should.

6. Network Management Issues

In this section we describe issues for network management.

6.1. SNMP and management of devices over IPv6 transport

Although some IPv6 MIBs exist, many are being reworked to blend more cleanly with IPv4 MIBs. This process is ongoing, but is reaching a conclusion with the publication of new IPv6-aware MIBs over recent months. These remain to be implemented by many vendors however. IPv6 multicast MIBs are still in the draft stage.

SNMP operations are most commonly run over IPv4 connections. This is, among other factors, due to the poor support for IPv6 transport SNMP in most commercial management solutions. The NET SNMP project includes IPv6 support [netsnmp].

Thus while the 6NET backbone had only IPv6 traffic running between the backbone routers, IPv4 was 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.

There are many types of devices on a network that will need to be managed over IPv6 connections eventually, including printers, switches and wireless access points.

Management of WLAN access points is also only available over IPv4 at the time of writing. 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. Additionally, access points that are aware of IPv4 traffic types need to be operated in bridged mode in almost all cases to function properly. Thus, features present for controlling IPv4 traffic are often not usable for IPv6 traffic.

Deliverable D2.4.2 [d242] discusses support for IPv6 in other aspects of wireless LANs, e.g. support for protocols for access control and to enable roaming.

6.2. 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, so that administrators do not need to support all the methods (for a variety of protocols that use 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 [dnsdisc]. However the deprecation of site-local unicast addressing in the IETF means that any existing use of the site local prefix needs to be reworked for the new ULA replacements [ula, ula-central].
- 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) [rfc3315] or “Stateless DHCPv6” [rfc3736]. Stateless DHCPv6 records no state (e.g. IP leases) and is thus a lighter service to support.

- Linklocal Multicast Name Resolution (LLMNR), aka. mDNS [llmnr]

It would be interesting to survey IPv6 RFCs and Internet Drafts for usage of the above set of methods. This is not a deployment issue per se, but it should be noted that not all these methods are available now (e.g. general implementations of DHCPv6, even though the standard has been finalised for nearly a year at the time of writing).

6.3. IPv6 Transition Management

It would be desirable to have tools to manage transition, including handling of routing, addressing, protocol translation and security aspects.

A new IETF Draft discusses the architectures for transition [transarch], while another considers security issues in transition tools [transsec] (e.g. requirements for secure end-to-end communication through firewalls, mapped address handling, tunnels and 6to4 relay issues).

There is some need to compare and highlight relative advantages of manual and automatic tunnel connectivity methods (e.g. Teredeo or 6to4 compared to the tunnel broker).

This is not a deployment barrier; however, implementers of IPv6 network management systems should consider the management of transition tools.

6.4. Tools for performance measurement

There is a requirement to have tools to measure performance of network paths. Companies such as Ixia, Agilent and Spirent have IPv6 network products in this area. There are also many open source tools, as reported in 6NET Work Package 6.

Some new features are required for IPv6 testing, e.g. for new IPv6-specific features and headers, as well as for IPv6 transition tools (which could be tested implicitly rather than explicitly).

We expect such products to become more widely available in the near future.

6.5. Tools for network monitoring

These are covered in 6NET Work Package 6; tools in use in 6NET are listed in [d624].

6.6. 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 [rfc3633]. We expect DHCP Prefix Delegation (DHCP-PD) to be widely used for prefix delegation.

6.7. IPv6 network renumbering

Because Provider Independent (PI) address space is not available to most enterprise networks (only those qualifying as an LIR may obtain it), an enterprise wishing to change ISP must renumber its network. Studies of procedures for IPv6 network renumbering are reported in 6NET Work Package 3, including enterprise renumbering analysis and experiments [d361][d362].

6.8. IPv6 NTP service

IPv6 support for NTPv4 has been added to the open source NTP development project [ntp]. The reference ID issue has been resolved as the first 32 bits of the MD5 hash of the IPv6 address. This code has been tested within the 6NET Project, as has a dedicated IPv6 NTP appliance supplied by Meinberg.

A small number of IPv6 NTP servers are now available. Further products can be expected to become available.

7. Application and IPv6-specific feature issues

Here we describe application-oriented issues related to IPv6 deployment.

7.1. Deprecation of IPv6 site local addresses

The IPv6 Scoped Address Architecture specification [rfc4007] 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 of [rfc3513].

The usage of link and global scope addresses has become well understood. However, there has been considerable discussion within the IETF (many hundreds of IETF email list messages, and generally a whole dedicated session at three IETF meetings) on best practice for use of site local addressing. The result of these discussions is that site local scope unicast addressing has been deprecated [rfc3879]. Their reference has been removed from some documents already [rfc4007]. There is also a new version of the IETF IPv6 Addressing Architecture RFC [rfc3513], but this still includes site-local references (section 2.5.6), which will need to be removed.

The two basic problems cited for site locals were ambiguity and leakage (routability). Site local addresses were ambiguous because any site using such addressing may have chosen its own site local prefix from within the fec0::/10 prefix. As a result, there would inevitably be clashes in addressing, and ambiguities where applications cross site boundaries. It is also likely that in many cases site local addresses (including source IPv6 addresses) would leak from sites, just as RFC1918 addresses do in IPv4 space now (whether directly or in application data payloads).

The requirements for a site-local replacement were discussed in an IETF Draft [lsareqts].

The IETF ipv6 WG has proposed two solutions from this requirement set, namely Unique Local Addresses (ULAs) [ula] and centrally assigned ULAs [ula-central]. The former are probabilistically unique, the latter unique by assignment/allocation. The solutions solve the ambiguity issue, though not necessarily the leakage issue. But this has been enough to appease most objectors to site locals as they were originally defined. The new schemes use a new prefix, with fc00::/7 reserved, with the 8th bit set to 1 for locally assigned (probabilistically unique) prefixes, and to 0 for centrally assigned ULAs.

Private addressing and NAT would be less in demand if effective IPv6 renumbering were available. The router renumbering protocol [rfc2894] is only one piece of that requirement (see Section 6 above), although that protocol itself requires stable local addressing to operate.

It appears that the 'site local' issue can be considered resolved, for now.

7.2. IPv6 Code Porting

The 6NET Project is undertaking a considerable amount of porting work, within the Applications (WP5) and Network Management (WP6) Work Packages.

7.2.1. Porting or development methodology

The methodology of such porting is to

a) Use the new APIs and data structures for TCP/IP sockets (as used in the C programming language) as defined in the Basic Socket Extensions for IPv6 [rfc3493] (which obsoletes RFC2553) and the Advanced Socket API for IPv6 [rfc3542] (which obsoletes RFC2292). Also relevant is the Single UNIX Specifiation, Version 3 [single3].

b) Make the code AF (IP) independent. The above standards specify the socket address structures, address conversion functions, socket options and name resolution functions. The definitions include IP-independent functions, as well as those for IPv6-only applications. In the current state of IPv6 deployment, IP-independent applications are preferred, such that they can operate in the presence of either or both protocols (without recompilation).

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 still currently some subtleties in behaviour between platforms, e.g. in binding to IPv4 and IPv6 simultaneously, due to different *bind()* call implementations.

The Java Development Kit (JDK) as of version 1.4.0 supports basic IPv6 functionality for Linux and Solaris platforms. MS Windows support was made available in JDK1.5. The JDK includes network preferences for IPv6 (i.e. *java.net.preferIPv4Stack, java.net.preferIPv6Addresses*) [jdkv6].

There is as yet no definition within Java for advanced API functions, e.g. writing a Flow Label field from a Java application. There needs to be action within the Java community to investigate and specify advanced API functionalities where required, including handling of IPv4-mapped addresses.

IPv6 API issues have been documented by 6NET participants within the Global Grid Forum (GGF) as dcoument GFD.40 [ggfspec], produced by the GGF IPv6 Working Group.

7.2.2. IP version independent specifications and code

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]
- A Global Grid Forum draft [ggfspec]

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

7.2.3. Software tools to aid porting

A number of tools exist to aid porting. These generally identify the code sections to be altered, rather than making any automatic changes. These include:

- The "socket scrubber" from Sun
- The Microsoft Checkv4.exe utility
- A porting tool under development as an M.Sc. project at Lancaster University.

Further discussion of these issues can be found in the IETF's Application Aspects of IPv6 Transition guide, RFC4038 [rfc4038], which draws heavily on the LONG guide mentioned above.

7.3. IPv6 Multicast Deployment

IPv6 multicast deployment issues have been described in an Internet Draft [multicast]. A key problem for IPv6 Multicast is that the model for Any-Source Multicast with PIM-SM requires a method for PIM-SM rendezvous points (RPs) to exchange information about sources. In IPv4 this can be done with MSDP, but there is no MSDP defined for IPv6, on the grounds that it is not a scalable solution. However, a solution for this has been proposed as a result of 6NET work, i.e. embedding the Rendezvous Point location in the multicast address, and this technique is now described in [rfc3956] and is implemented by Cisco and Juniper. This is an example of a technology that is simply not possible with IPv4.

There may be a reduced requirement for PIM-SM if source specific multicast (SSM) can be used instead. MLD [rfc2710] and its successor MLDv2 [rfc3810] are both important for IPv6 Multicast deployment. MLDv2 obsoletes MLD, and includes support for listening for specific sources, and thus for 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. Vendors are just beginning to implement MLD snooping in Layer 2 equipment (e.g. Cisco has an EFT for this, tested successfully in 6NET, at the time of writing). It is not known if Cisco will implement CGMP for IPv6.

When considering snooping, one should bear in mind that existing switches will likely be in place for many years, and not be upgraded.

There is a debate in the IETF MBONED WG [mboned] on whether ASM should be promoted with IPv6, using the embedded RP for inter-domain communication, or whether IPv6 is an opportunity to migrate all multicast to SSM (with the applications being modified to adapt). The 6NET project has studied both areas in Work Package 3, and as a result is highlighting many issues. Both technologies are being used successfully (over the same infrastructure).

Thus there seem to be no significant IPv6 multicast deployment issues, although the option to run MLD snooping in enterprise equipment is desirable (as is done for IGMP today for IPv4) but not yet widely implemented.

7.4. IPv6 Deployment over ADSL (SOHO environments)

There has been a lot of discussion on the IPv6 deployment scenarios that can most benefit end users. Many agree that home networking (SOHO networks) is a very strong scenario to show IPv6's added value – most home networks today lie behind IPv4 NATs, inhibiting communications into and between homes.

IPv6 transition methods exist for advanced home users, e.g. Tunnel Brokers or 6to4. However, for this scenario to take off, native ADSL IPv6 deployment is required.

This is generally not a technical issue; the protocol building blocks exist, e.g. IPv6 over PPP [rfc2472]. Some ISPs in Japan and Asia are deploying native ADSL services. It is one of cost and willingness (visible return on investment (RoI)) for European ISPs, in the absence of 'obvious' 'killer' applications. This is an issue for IPv6 Task Forces to address, not the 6NET project directly, although 6NET has validated much of the required technology. There are limited services available in Europe, often provisioned via tunnelled methods.

An example of IPv6 ADSL deployment is described in [d514].

7.5. Available applications

Basic application support is good (ping, telnet, ftp, etc), and a growing number of general applications have become or are becoming available with IPv6 support.

7.5.1. General applications

These include but are not limited to:

- Sendmail, postfix.
- OpenLDAP
- Apache2
- OpenSSH
- BIND9 (DNS)
- OpenH.323
- UW-imap
- kPhone

A good description and listing of Linux applications can be found the DeepSpace6 site [deepspace6].

A discussion of IPv6 support in peer-to-peer applications was carried out at a recent TERENA TF-NGN meeting[5].

7.5.2. Proxy services

Proxy services are also available, e.g. squid as a Web proxy, totd for DNS, or fetchmail for retrieving emails.

7.6. 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.

[5] http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn15/20040930_jm_p2p_ipv6.pdf

Linux appears to still be lacking on TI-RPC support, so it cannot do RPC over IPv6. Hence NFS, NIS etc. are missing, but there is some work on fixing this. 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.

Another problem is lack of IPv6 support in SQL databases like mySQL and postgreSQL (though a patches exist for both, we don't believe they have been integrated into the official build yet). We are not aware of any SQL databases supporting IPv6 at this time. 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).

IPv6 support is included in at least two RADIUS implementations – Radiator and FreeRADIUS.

The general issue here is one of IPv6 functionality not being present in many commercial products, whereas most open source products do include IPv6 support. This situation will change, but may take time.

7.7. Use of the IPv6 Flow Label

The IPv6 Flow Label usage [flow] 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 the Integrated Services mode for QoS, which is no longer widely used. Interestingly, there is now some discussion of Flow Label usage for MPLS labels [6lsa].

This does not represent a deployment barrier; implementations should set the Flow Label value to zero by default.

7.8. Use of IPv6 Privacy Extensions (RFC3041)

RFC3041 [rfc3041] was devised to prevent IPv6 devices being trackable in cases where stateless address autoconfiguration [rfc2462] is being used and a host attaches to various IPv6 networks, using a common 64-bit host part of the address (the address being formed with the host's unique MAC address included).

The privacy extensions allow a host to generate a "random" host part of its address in the stateless autoconfiguration process, thus preventing correlation of common host elements of IPv6 addresses where observed in various traffic or server logs.

Windows XP and USAGI Linux implement RFC3041. It may be used on static nodes too, e.g. a Windows XP desktop system will create a new privacy address to use in outbound connections on a daily basis.

In some cases, as described in [harm], RC3041 may be detrimental. One case is where distributed denial of service attacks (which regularly change the source address that they use) may be confused with RFC3041 behaviour. Also, while trusting a specific source IP or set of IPs is not necessarily

optimal authentication, RFC3041 makes this less easy to do. And there is 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. Finally, in 6NET it has become clear that multiply addressed hosts present problems for network management and monitoring tools – how do these know which IPv6 addresses belong to the same nodes?

It is desirable to have RFC3041 usage controllable per application, e.g. so it could be turned off for ssh to allow (weak) IP-based authentication.

RFC3041 is only one privacy tool; users may still be tracked by cookies (for example) and by persistent IPv6 network address prefixes.

Further operational experience of IPv6 Privacy Extensions is required.

8. Security-related Issues

8.1. 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. in principle OSPFv3 can use AH for route data authentication (although operational examples may be limited at the time of writing).

IPv6 IPsec implementations to date include:

- Net/FreeBSD, which use KAME
- OpenBSD
- KAME
- HP/UX
- 6WIND
- FreeS/WAN
- The 2.6 Linux kernel

It is expected that the US DoD's announcement that it has a requirement IPv6 support in procurements will lead to an acceleration in IPsec implementations.

8.2. IPv6 Firewalls and Intrusion Detection Systems (IDS)

Commercial IPv6 firewall products are emerging at the time of writing, including CheckPoint Firewall-1, the Nokia IP range and NetScreen.

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

- Handling extension header chains, and unknown options.
- Handling header options related to MIPv6 usage
- How to handle end-to-end IPsec sessions
- How to handle peer-to-peer applications
- Lack of availability of fully-featured stateful firewalls

Stateful firewalls exist for Linux and BSD filtering tools (e.g. ipfw). New commercial IPv6 firewall products can be expected to mature in the relative near term.

In contrast, there is as yet no evidence of IPv6 features in commercial IDS products that are capable of detecting IPv6 header manipulation (e.g. repeated or bad IPv6 header options). There is as yet also no IPv6 support in the most popular open source IDS package, snort.

8.3. Network Address Translation (NAT)

The use of IPv4 NAT as a 'security' measure has been engrained into the minds of network administrators for a considerable time.

Using NAT for IPv6 would destroy a large part of the benefit of deploying IPv6. The IETF has recongised this fact and documented a philosophy for Network Architecture Protection (NAP) [nap] that shows how a NAT's 'advantages' can be realised and improved upon with IPv6 deployment.

8.4. Securing Stateless Autoconfiguration

The IETF send WG [send] has studied methods that could be applied to secure Neighbor Discovery [rfc2461]. The motivation for this work was that the "plug and play" IPv6 stateless autoconfiguration [rfc2462] by default 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.

Trust in Neighbor Discovery is discussed in [rfc3756].

The send WG has recently defined a cryptographically generated address (CGA) approach [rfc3972] to solve this problem within the SEcure Neighbour Discovery (SEND) protocol [rfc3971]. The send WG has been disbanded on completion of this work.

8.5. 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 [rfc3964] or transport relay translators (TRT).

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.

A summary of transition security is presented in [transsec].

This area is being studied in 6NET in Work Package 2, see [d224] and [d234v2].

9. IPv6 Early Deployment Scenarios

In this section we briefly overview some potential early IPv6 deployment and application scenarios. These scenarios are presented with an academic, research network viewpoint. They may be realised through dual stack or IPv6-only. We are beginning to see these scenarios realised in 6NET campus deployments already.

An example of the operation of an IPv6-only site was described in [d251].

9.1. Wireless Campus

Here we deploy a wireless infrastructure in support of teaching and research in a university setting.

The components described in Deliverable D2.4.2 [d242] would be be required, including WLAN access and roaming support, and deployment of Mobile IPv6 for persistent connectivity for services including VoIP, video streaming and remote access sessions (e.g. using ssh).

Access to learning material for various laptop and PDA devices may require adaptation. Location awareness and presence notification will be important building blocks for services. Multicast may be advantageous for content distribution.

Visiting researchers will seek to collaborate with colleagues in the local network environment.

New 'phone' devices will more commonly feature 802.11 WLAN interfaces and IPv6 capability (e.g. Symbian OS on the Nokia Communicator 9500) in increasing numbers, and be used by students on campuses.

9.2. Home entertainment and networking

Currently the possibilities for home networking are restricted by the general presence of IPv4 NATs in such home networks. With IPv6 the application view becomes simplified, at the expense of a greater reliance on end-to-end security.

A typical example may be videoconferencing or collaborative working between a tutor and student.

New consumer devices are set to emerge with IP(v6) functionality.

IPv6 may have advantages for gaming applications, where multiparty games could be operated without the need for central servers.

Native IPv6 access is also required in the long run, while intermediate transition tools are used (such as the Teredo solution used by Microsoft's ThreeDegrees application 'experiment').

9.3. Peer-to-peer applications and messaging

Here we foresee peer-to-peer applications in the form of file sharing, VoIP and similar examples. The "server" systems may be static or mobile. Issues of naming and addressing, and in the user-network interface, will be important.

Use of Dynamic DNS, ENUM [rfc3761] or Mobile IPv6 [rfc3775] may be required for a consistent name space view from the applications.

9.4. Multicast video distribution

With IPv4, multicast has failed to gain widespread adoption. While there is no single reason for that, the complexity of PIM-SM and its associated inter-domain RP source discovery exchange protocol (i.e. MSDP) is one barrier. With IPv6, we have an opportunity to adopt SSM from the outset, adapting applications from the any-source multicast (ASM) to source-specific multicast (SSM) mode as necessary, or to use PIM-SM with Embedded RP [rfc3956].

Application examples may include high-quality video such as DVTS.

10. Conclusions

In this report we have identified a number of issues that may hinder the successful deployment of IPv6 on a large scale. Some of the concerns are more pressing than others, and some may be relatively minor issues.

Many of the issues first reported in [d251] – the first instance of this report - have been resolved by the time of this final version. This is highlighted by the number of IPv6-related Internet Drafts that have progressed to RFC status in the course of the 6NET project, and the number of existing standards nearing completion of their update based on deployment experience, the references at the end of this document being a (subset) of the full list.

The following table highlights the more significant issues:

	Issue	Solutions	Timeframe
1	Lack of widespread IPv6 transit between networks (in particular commercial networks)	Wider adoption of IPv6 by the commercial carrier and transit network operators.	Medium
2	Preference for use of IPv6 may give worse performance than use of IPv4	Better IPv6 connectivity (see 1) Ensuring host services are all dual-stack capable	Medium Medium
3	Lack of IPv6 multihoming capability or Provider Independent (PI) address space	IETF solutions to be taken forward in the shim6 WG. Availability of PI address space from ARIN (proposed, but not adopted)	Long Short
4	Lack of single agreed method for local IPv6 DNS resolver discovery	Use of DHCPv6 (full or ‘stateless’) Modification of RA messages	Available Medium

5	Potential for DNS UDP message size to be exceeded, in particular for DNSSEC deployments	Use of EDNS0 (RFC2671)	Available, but not widely implemented
6	Lack of ability to pre-populate reverse DNS zone for a prefix	Potential use of wildcards? Some argue reverse DNS not needed	Unclear
7	Implementation of new IPv6-aware MIBs by vendors and IPv6 transport for SNMP	Requires vendor action	Short to Medium
8	No advanced IPv6 API for Java, e.g. to write IPv6 Flow Label values	Action required within Java community	Medium?
9	General lack of IPv6 ADSL services to SOHO networks (home users) as a highly desirable deployment scenario.	ISPs need to see a return on investment to be attracted to deploy IPv6 natively for SOHO users	Medium to Long?
10	Lack of IPv6 support in many commercial software products (applications)	Requires vendor action	Short to Medium?
11	Pros and cons of IPv6 Privacy Extensions need to be evaluated in production networks	Gain operational experience to determine best practice	Short to Medium
12	No intrusion detection systems capable of checking IPv6 headers/transport	Vendor and open source community action required	Medium?

The IPv6-only deployment in Norway, reported in D.2.5.1, shows that in a small scale it is possible to use IPv6 only in a site network.

In the meantime dual-stack operation is the most realistic and practical stepping stone towards the IPv6 end game, and much of the work in this report and in 6NET is general is focused there.

11. References

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

[6link] The 6LINK project, http://www.6link.org and http://www.ist-ipv6.org

[6lsa] "IPv6 Label Switching Architecture", S. Chakravorty, IETF Internet Draft (work in progress), February 2005, http://www.ietf.org/internet-drafts/draft-chakravorty-6lsa-01.txt

[6mess] "Moving from 6bone to IPv6 Internet", P. Savola, IETF Internet Draft, November 2002 (expired), http://www.watersprings.org/pub/id/draft-savola-v6ops-6bone-mess-01.txt

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

[d224] "Final IPv4 to IPv6 Transition Cookbook for Organisational/ISP (NREN) and Backbone Networks", 6NET Project Deliverable, T. Chown editor, February 2005, http://www.6net.org/publications/deliverables/D2.2.4.pdf

[d234v2] "Final IPv4 to IPv6 Transition Cookbook for end-site networks/universities", 6NET Project Deliverable, C. Schild editor, June 2005, http://www.6net.org/publications/deliverables/D2.3.4v2.pdf

[d242] "Final report on IPv6-specific implications for Wireless LAN/MAN transition to IPv6", 6NET Project Deliverable, T. Chown editor, October 2003, http://www.6net.org/publications/deliverables/D2.4.2.pdf

[d251] "Issues for IPv6 Deployment (missing pieces for IPv6 deployment and IPv6-only operation)", 6NET Project Deliverable, T. Chown editor, December 2002, http://www.6net.org/publications/deliverables/D2.5.1.pdf

[d252] "Updated IPv6 Deployment Issues (missing pieces for IPv6 deployment and IPv6-only operation)", 6NET Project Deliverable, T. Chown editor, September 2003, http://www.6net.org/publications/deliverables/D2.5.2.pdf

[d361] "IPv6 network renumbering guide for SOHO and Backbone networks", 6NET Project Deliverable, T. Chown editor, June 2005, http://www.6net.org/publications/deliverables/D3.6.1.pdf

[d362] "IPv6 network renumbering guide for ISP and Enterprise networks", 6NET Project Deliverable, T. Chown editor, June 2005, http://www.6net.org/publications/deliverables/D3.6.2.pdf

[d453] "Evaluation of Multihoming Solutions", 6NET Project Deliverable, M. Dunmore editor, February 2005, http://www.6net.org/publications/deliverables/D4.5.3.pdf

[d514] "Cookbook on Deploying IPv6 in School Networks", 6NET Project Deliverable, D. Kalogeras, C. Friacas, J. Ferreira editors, June 2005, http://www.6net.org/publications/deliverables/D5.14.pdf

[d624] "Final Report on IPv6 management tools, developments and tests", 6NET Project Deliverable, B. Gajda and W. Procyk editors, September 2004, http://www.6net.org/publications/deliverables/D6.2.4.pdf

[deepspace6] The DeepSpace6 Linux Portal, http://www.deepspace6.net/

[defrouter] "Default Router Preferences, More Specific Routes and Load Sharing", R. Draves, D. Thaler, IETF Internet Draft (work in progress), October 2004, http://www.ietf.org/internet-drafts/draft-ietf-ipv6-router-selection-06.txt

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

[dhcplinux] DHCPv6 for Linux, a Sourceforge project, http://dhcpv6.sourceforge.net/

[dnsdisc] "Well known site local unicast addresses to communicate with recursive DNS servers", A. Durand et al,, IETF Internet Draft (expired), October 2002, http://www.watersprings.org/pub/id/draft-ietf-ipv6-dns-discovery-07.txt

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

[dnsissues] "Operational Considerations and Issues with IPv6 DNS", A. Durand, J. Ihren, P. Savola, IETF Internet Draft (work in progress), October 2004, http://www.ietf.org/internet-drafts/draft-ietf-dnsop-ipv6-dns-issues-10.txt

[dnsrev] "Dynamic reverse DNS for IPv6", A. Durand, IETF Internet Draft (expired), February 2003, http://www.watersprings.org/pub/id/draft-durand-dnsop-dynreverse-00.txt

[fire] "Firewalling Considerations for IPv6", P. Savola, IETF Internet Draft (expired), March 2003, http://www.watersprings.org/pub/id/draft-savola-v6ops-firewalling-01.txt

[geant] The GÉANT network, http://www.dante.net/geant/

[ggfspec] "Guidelines for IP version Independence in GGF Specifications", T. Chown, J. Bound, S. Xiang, P. O'Hanlon, GGF Draft, September 2003, http://www.ggf.org/

[harm] "RFC 3041 Considered Harmful", F. Dupont, P. Savola, IETF Internet Draft (expired), January 2003, http://www.watersprings.org/pub/id/draft-dupont-ipv6-rfc3041harmful-02.txt

[intdepr] "Deprecation of 'ip6.int'", G. Huston, IETF Internet Draft (work in progress), May 2005, http://www.ietf.org/internet-drafts/draft-huston-ip6-int-03.txt

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

[ipv6fix] The WIDE IPv6 Fix WG, http://www.wide.ad.jp/project/wg/v6fix.html

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

[ipv6tf] The European IPv6 Task Force Portal, http://www.ipv6tf.org/

[issues] "Unidentified issues in IPv6 deployment operation", J. Hagino, T. Jinmei, IETF Internet Draft (expired), June 2002, http://www.watersprings.org/pub/id/draft-itojun-jinmei-ipv6-issues-00.txt

[ist-ipv6] IST IPv6 Cluster web site, http://www.ist-ipv6.org/

[jdkv6] Java Development Kit 1.4.1, IPv6 Guide, http://java.sun.com/j2se/1.4.1/docs/guide/net/ipv6_guide

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

[llmnr] "Linklocal Multicast Name Resolution (LLMNR), L. Esibov, B. Aboba, D. Thaler, IETF Internet Draft (work in progress), May 2005, http://www.ietf.org/internet-drafts/draft-ietf-dnsext-mdns-40.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

[lsareqts] "Requirements for Limited-scope Unicast Addressing in IPv6", F. Templin, IETF Internet Draft (expired), June 2003, http://www.watersprings.org/pub/id/draft-templin-lsareqts-00.txt

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

[mboned] IETF Mbone Deployment (MBONED) WG, http://www.ietf.org/html.charters/mboned-charter.html

[mip6] Mobility for IPv6, IETF WG, http://www.ietf.org/html.charters/mip6-charter.html

[mipl] Mobile IPv6 for Linux, http://www.mipl.mediapoli.com

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

[multicast] "IPv6 Multicast Deployment Issues", P. Savola, IETF Internet Draft (expired), November 2002, http://www.watersprings.org/pub/id/draft-savola-v6ops-multicast-issues-01.txt

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

[nap] "IPv6 Network Architecture Protection", G. Van de Velde at al, IETF Internet Draft (work in progress), June 2005, http://www.ietf.org/internet-drafts/draft-ietf-v6ops-nap-01.txt

[natdepr] "Reasons to Move NAT-PT to Experimental", C. Aoun, E. Davies, IETF Internet Draft (work in progress), January 2005, http://www.ietf.org/internet-drafts/draft-ietf-v6ops-natpt-to-exprmntl-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/

[pana] IETF Protocol for Carrying Authentication for Network Access WG, http://www.ietf.org/html.charters/pana-charter.html

[rev6to4] "6to4 Reverse DNS Delegation", G. Huston, IETF Internet Draft (expired), October 2004, http://www.watersprings.org/pub/id/draft-huston-6to4-reverse-dns-03.txt

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

[rfc2461] "Neighbor Discovery for IPv6", T. Narten, E. Nordmark, W. Simpson, IETF RFC2461 (under update), December 1998, http://www.ietf.org/rfc/rfc2461.txt

[rfc2461bis] "Neighbor Discovery for IP version 6 (IPv6)", T. Narten et al, IETF Internet Draft (work in progress), May 2005, http://www.ietf.org/internet-drafts/draft-ietf-ipv6-2461bis-03.txt

[rfc2462] "IPv6 Stateless Autoconfiguration", S. Thomson, T. Narten, IETF RFC2462 (under update), December 1998, http://www.ietf.org/rfc/rfc2462.txt

[rfc2462bis] "IPv6 Stateless Address Autoconfiguration", S. Thomson, T. Narten, T. Jinmei, IETF Internet Draft (work in progress), May 2005, http://www.ietf.org/internet-drafts/draft-ietf-ipv6-rfc2462bis-08.txt

[rfc2472] "IP Version 6 over PPP", D. Haskin and E. Allen, IETF RFC 2472 (under update), December 1998, http://www.ietf.org/rfc/rfc2472,txt

[rfc2472bis] "IP Version 6 over PPP", S. Varada and E. Allen, IETF Internet Draft, June 2005, http://www.ietf.org/internet-drafts/draft-ietf-ipv6-over-ppp-v2-02.txt

[rfc2553] "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 (under update), January 2001, http://www.ietf.org/rfc/rfc3041.txt

[rfc3041bis] "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", T. Narten and R. Draves, S. Krishnan, IETF Internet Draft (work in progress), May 2005, http://www.ietf.org/internet-drafts/draft-ietf-ipv6-privacy-addrs-v2-04.txt

[rfc3162] "RADIUS and IPv6", B. Aboba et al, IETF RFC3162, August 2001, http://www.ietf.org/rfc/rfc3162.txt

[rfc3226] "DNSSEC and IPv6 A6 aware server/resolver message size requirements", O. Gudmundsson, IETF RFC3226, December 2001, http://www.ietf.org/rfc/rfc3226.txt

[rfc3315] "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", R. Droms et al, IETF RFC3315, July 2003, http://www.ietf.org/rfc/rfc3315.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

[rfc3484] "Default Address Selection for IPv6", R. Draves, IETF RFC3484, February 2003, http://www.ietf.org/rfc/rfc3484.txt

[rfc3493] "Basic Socket Interface Extensions for IPv6", R. Gilligan et al., IETF RFC3493, February 2003, http://www.ietf.org/rfc/rfc3493.txt

[rfc3513] "IP Version 6 Addressing Architecture ", S. Deering, R. Hinden, , IETF RFC3513, April 2003, http://www.ietf.org/rfc/rfc3513.txt

[rfc3513bis] "IP Version 6 Addressing Architecture", R. Hinden, S. Deering, IETF Internet Draft (work in progress), May 2005, http://www.ietf.org/internet-drafts/draft-ietf-ipv6-addr-arch-v4-04.txt

[rfc3542] "Advanced Sockets Applications Program Interface (API) for IPv6", R. Stevens et al., IETF RFC3542, May 2003, http://www.ietf.org/rfc/rfc3542.txt

[rfc3633] "IPv6 Prefix Options for DHCPv6", O. Troan and R. Droms, IETF RFC 3633, December 2003, http://www.ietf.org/rfc/rfc3633.txt

[rfc3697] "IPv6 Flow Label Specification", J. Rajahalme et al, IETF RFC 3697, March 2004, http://www.ietf.org/rfc/rfc3697.txt

[rfc3701] "6bone (IPv6 Testing Address Allocation) Phaseout", R. Fink, R. Hinden, IETF RFC 3701, March 2004, http://www.ietf.org/rfc/rfc3701.txt

[rfc3736] "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6, R. Droms, IETF RFC 3736, April 2004, http://www.ietf.org/rfc/rfc3736

[rfc3756] "IPv6 Neighbor Discovery trust models and threats", P. Nikander editor, IETF RFC 3756, May 2004, http://www.ietf.org/rfc/rfc3756.txt

[rfc3761] "The E.164 to URI DDDS Application (ENUM)", P. Faltstrom, M. Mealling, IETF RFC3761 (updating RFC2916), April 2004, http://www.ietf.org/rfc/rfc3761

[rfc3775] "Mobility Support in IPv6", D. Johnson, C. Perkins, J. Arkko, IETF RFC 3775, June 2004, http://www.ietf.org/rfc/rfc3775.txt

[rfc3776] "Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents", J. Arkko, V. Devarapalli, F. Dupont, IETF RFC 3776, June 2004, http://www.ietf.org/rfc/rfc3776.txt

[rfc3810] "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", R. Vida et al, IETF RFC 3810, June 2004, http://www.ietf.org/rfc/rfc3810.txt

[rfc3879] "Deprecating Site Local Addresses", C. Huitema, B. Carpenter, IETF RFC 3879, September 2004, http://www.ietf.org/rfc/rfc3879.txt

[rfc3901] "DNS IPv6 transport operational guidelines", A. Durand, J. Ihren, IETF RFC 3901, September 2004, http://www.ietf.org/rfc/rfc3901.txt

[rfc3956] "Embedding the Address of RP in IPv6 Multicast Address", P. Savola, B. Haberman, IETF RFC3956, November 2004, http://www.ietf.org/rfc/rfc3956.txt

[rfc3964] "Security Considerations for 6to4", P. Savola, C. Patel, IETF RFC3964, December 2004, http://www.ietf.org/rfc/rfc3964.txt

[rfc3971] "SEcuring Neighbour Discovery (SEND)", J. Arjjo editor, IETF RFC3971, March 2005, http://www.ietf.org/rfc/rfc3971.txt

[rfc3972] "Cryptographically Generated Addresses (CGAs)", T. Aura, IETF RFC3972, March 2005, http://www.ietf.org/rfc/rfc3972.txt

[rfc4007] "IPv6 Scoped Address Architecture", S.Deering et al, IETF RFC 4007, March 2005, http://www.ietf.org/rfc/rfc4007.txt

[rfc4038] "Application Aspects of IPv6 Transition", M. Shin, P. Savola et al, IETF RFC4038, March 2005, http://www.ietf.org/rfc/rfc4038.txt

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

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

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

[shim6] IETF Shim6 WG Draft Charter, http://www3.ietf.org/proceedings/05mar/shim6.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/

[transarch] "A view on IPv6 transition architecture", P. Savola, IETF Internet Draft (expired), January 2004, http://www.watersprings.org/pub/id/draft-savola-v6ops-transarch-03.txt

[transsec] "IPv6 Transition/Co-existence Security Considerations", E. Davies, S. Krishnan, P. Savola, IETF Internet Draft (work in progress), May 2005, http://www.ietf.org/internet-drafts/draft-ietf-v6ops-security-overview-00.txt

[ula] "Unique Local IPv6 Unicast Addresses", R. Hinden, B. Haberman, IETF Internet Draft (work in progress), January 2005, http://www.ietf.org/internet-drafts/draft-ietf-ipv6-unique-local-addr-09.txt

[ula-central] "Centrally Assigned Unique Local IPv6 Unicast Addresses", R. Hinden, B. Haberman, IETF Internet Draft (work in progress), February 2005, http://www.ietf.org/internet-drafts/draft-ietf-ipv6-ula-central-01.txt

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

[v6on] "Dual stack IPv6 on by Default", S. Roy et al, IETF Internet Draft (expired), June 2003, http://www.watersprings.org/pub/id/draft-roy-v6ops-v6onbydefault-01.txt

[v6ready] The IPv6 Ready Programme, http://www.ipv6ready.org/