


IPv6

TECHNICAL GUIDE

The background of the cover is a dark red color. It features a network diagram composed of several circular nodes connected by lines. Two nodes are highlighted in a bright orange color, while the others are a darker shade of red. The lines connecting the nodes are also in a light red color, creating a complex, interconnected pattern.

Tim Chown
School of Electronics
and Computer Science,
University of Southampton

JANET(UK) Technical Guides

JANET(UK) Technical Guides are a series of user guides available to JANET customers. Technical Guides contain detailed technical information and are intended for technical support staff at JANET sites, network specialists, or those with a particular interest in the specialist area.

If you have any queries or comments about the Guides or would like to obtain copies, please contact:

JANET Service Desk

Lumen House, Library Avenue

Harwell Oxford

Didcot, Oxon OX11 OSG

Tel: 0300 300 2212

Fax: 0300 300 2213

E-mail: service@ja.net

Further details of the documents in this series are available at:

<http://www.ja.net/services/publications/technical-guides.html>

Table of Contents

1	Introduction	7
2	IPv6 Overview.....	8
2.1	The Emergence of IPv6.....	8
2.2	The JANET Context.....	9
3	Arguments for IPv6 Deployment	11
3.1	Global Address Space	11
3.2	Supporting Teaching and Research.....	11
3.3	Early Experience/Insight into IPv6	12
3.4	Internet Architecture and End-to-End Transparency.....	12
3.5	Security Viewpoint	13
4	Basic Operation and Features of IPv6	14
4.1	IPv6 Core Protocol and Features.....	14
4.2	Address Format.....	14
4.2.1	IPv6 Literal Addresses	14
4.2.2	IPv6 Address Architecture	15
4.3	IPv6 Header Format	16
4.4	Fragmentation	18
4.5	ND (Neighbour Discovery) Protocol.....	18
4.6	SLAAC (Stateless Address Autoconfiguration)	18
4.6.1	IPv6 Privacy Addresses.....	19
4.6.2	Securing SLAAC and Neighbour Discovery	19
4.7	DHCPv6.....	20
4.8	Address Management.....	20
5	Main Differences to IPv4	22
5.1	Addressing	22
5.2	Operation	22
5.3	Multicast.....	23
5.4	IPv6 and NAT	23
6	IPv6 Address Allocation and Management.....	24
6.1	Hierarchical Addressing.....	24
6.2	Allocations to JANET.....	24
6.3	Allocations to End Sites (Universities, Colleges)	24
6.4	Allocations to Regional Network Operators.....	25
6.5	Network Addressing Plans.....	25
7	IPv6 Transition and Coexistence Tools.....	26
7.1	Dual-stack	26
7.2	Tunnelling	26
7.2.1	Manual Tunnels	27
7.2.2	Tunnel broker.....	28
7.2.3	6to4.....	29
7.2.4	6rd	31
7.2.5	Teredo.....	31
7.3	Translation	32

7.4	IPv4 Mapped Addresses	33
7.5	Transition Tools on JANET	33
8	JANET and IPv6	35
8.1	History of IPv6 on JANET	35
8.1.1	Production IPv6 networking	35
8.1.2	International peerings	35
8.2	Connecting to JANET with IPv6.....	36
8.3	IPv6-enabled JANET Services	36
8.3.1	Web site.....	36
8.3.2	DNS	36
8.3.3	Other services	36
9	Deploying IPv6 in academic enterprise networks	37
9.1	Multiphase deployment	37
9.1.1	Dual-stack or IPv6-only	37
9.2	Phase 1: Advance planning	37
9.2.1	Assessing IPv6 capability	38
9.2.2	Address planning.....	39
9.3	Phase 2: Testbed/Trial deployment	40
9.3.1	Testbed considerations.....	40
9.4	Phase 3: Production deployment	42
9.4.1	Preparing the IPv6 Network.....	42
9.4.2	DNS	43
9.4.3	E-Mail	44
9.4.4	Web servers	44
9.5	Phase 4: Ongoing operation	45
10	Deployment of IPv6 on Regional Networks	46
10.1	Deployment Methodology	46
10.2	Choice of Routing Protocols.....	46
10.3	Address Space and DNS	47
10.3.1	Point-to-point links	47
10.4	Other Services.....	47
11	IPv6 Security	48
11.1	IPv4 Security Equivalence.....	48
11.2	New IPv6 Security Issues	48
11.3	IPv6 firewalls	49
11.4	Network port scanning	49
11.5	IP Security (IPsec), CGAs and SeND.....	50
11.6	Summary.....	50
12	Advanced IPv6 Network Services	51
12.1	IPv6 Multicast	51
12.1.1	Embedded RP	51
12.1.2	IPv6 SSM.....	52
12.1.3	MLD	52
12.1.4	IPv6 Multicast on JANET.....	52
12.2	IPv6 QoS	52

12.2.1	DiffServ	53
12.2.2	IntServ	53
12.3	IPv6 Multihoming	53
12.4	Mobile IPv6	53
13	Case Study: Enterprise IPv6 Deployment at Southampton	56
13.1	Scenario	56
13.2	Dual-stack vs IPv6-only	56
13.3	Capability review.....	56
13.4	Address Management.....	57
13.5	Network Elements	57
13.6	IPv6 Services	58
13.7	Monitoring and service examples	59
13.7.1	IPv6 external web traffic	59
13.7.2	IPv6 email.....	60
13.7.3	Switch/router monitoring – NAV	60
13.7.4	Network flows	61
13.8	Reflections.....	62
13.9	Future steps.....	63
14	IPv6 Roadmap	64
14.1	Recommendations for UK Academic Deployment	64
14.1.1	When to Deploy?	64
14.1.2	How to Deploy?	65
14.2	Feedback.....	65
14.3	Acknowledgements	65
15	References.....	66
	Glossary.....	69
	Appendix A: Sample Cisco® IOS Configuration File.....	71

List of Figures

Figure 4.1: Differences in IPv4 and IPv6 header formats	16
Figure 4.2: IPv6 header format.....	17
Figure 7.1: IPv6 in IPv4 tunnel	27
Figure 7.2: A tunnel to the JANET IPv6 Experimental Service.....	28
Figure 7.3: IPv6 tunnel broker.....	28
Figure 7.4: Basic 6to4 usage between two 6to4 routers	30
Figure 7.5: 6to4 with a 6to4 relay.....	30
Figure 7.6: 6rd architecture	31
Figure 7.7: Teredo overview	32
Figure 7.8: Dual-stack MX relay as an ALG	36
Figure 9.1: One possible testbed topology	41
Figure 9.2: A potential topology for a dual-stack DMZ pilot.....	41
Figure 9.3: A dual-stack site topology.....	43
Figure 12.1: Mobile IPv6	54
Figure 12.2: Mobile IPv6 with Route Optimisation.....	54
Figure 13.1: Dual-stack with separate firewalls	58
Figure 13.2: External web visits, 2005-2010	60
Figure 13.3: NAV searching for host by MAC	61
Figure 13.4: NAV, dual-stack VLAN	61
Figure 13.5: nfsen showing DNS Netflow data.....	62
Figure 13.6: nfsen showing specific SMTP flows	62

1. Introduction

This technical guide is intended to assist site and network administrators in the UK academic community to deploy IPv6 services, ranging from a small experimental testbed through to a fuller production campus deployment. This version of the document updates the original version released in 2006.

The guide covers the basics of IPv6 and how it evolved, highlighting some of the differences between IPv6 and the existing IPv4 protocol. It describes, with a case study example, how IPv6 can be deployed on sites (i.e. campus, whether HE or FE) and regional networks, including discussion of the popular IPv6 transition and coexistence techniques. IPv6 address allocations are explained in the context of the JANET network, along with notes on deployment of basic IPv6 enabled services.

Since the original version of this guide was written, IPv6 on the JANET backbone has progressed from an experimental service on SuperJANET4 to a production service on SuperJANET5. The SLA that was introduced under SuperJANET5 with the RNOs requires them to provide native IPv6 on request for their customers. While the experimental service may remain in place for a while as a fallback tunnelling solution, there should now be no barrier to sites connecting natively.

A further sign of the evolution of IPv6 since the original guide was written is the number of IETF Internet Drafts that have progressed to RFC status in the intervening years. IPv6 is now a mature technology with very broad support in hosts, routers and applications. IPv6 is present, and enabled by default, in Windows 7, Linux and MacOS X, and widely supported on common router platforms. At the time of writing, both Google and Facebook have IPv6 content available, Google via a test programme¹ and Facebook via a specific IP6-only domain.² By the time you read this, both may well be offering IPv6 connectivity to their production domains.

The most important event since the last version of the guide was written is the exhaustion of the unallocated IPv4 global address pool in February 2011. The last unused /8 IPv4 blocks have now been given one to each Regional Internet Registry (RIR), so from 2011 onwards sites and ISPs will have to look to increased leasing and trading of address blocks to get the address space they need for future growth.

Within JANET, most academic sites already have enough IPv4 address space to meet their immediate needs, many of them having received generous allocations up to 20 years ago, but they will need to understand when it is timely to consider introducing IPv6 and how they should go about planning that process.

The author of this guide welcomes any feedback that may help improve future versions. Such feedback should be sent to the author, Tim Chown (tjc@ecs.soton.ac.uk).

1 <http://www.google.com/intl/en/ipv6/>

2 <http://www.v6.facebook.com>

2. IPv6 Overview

IPv6 is the new version of IP, the common protocol underpinning all Internet communications, and as such it will ultimately succeed the current version, IPv4, to become the dominant version of IP used on the Internet. The gradual transition from IPv4 towards majority IPv6 deployment will take many years and IPv4 itself may never be phased out completely. In the meantime the two protocols can co-exist and be used together in various ways, as described in Section 7.

2.1 The Emergence of IPv6

In the 1990s, the main driver for the research and development that led to the standardisation of the base IPv6 protocol was the foreseen exhaustion of the IPv4 address space. When IPv4 was first deployed in the 1970s the 32-bit address format, offering a theoretical four billion addresses, was not seen to be a limiting factor, but the global uptake of the Internet, which is still in its infancy in some areas, has proven otherwise. While the web was the primary driver for growth, the future Internet looks set to embrace whole new application domains of networked devices, e.g. smart metering devices in the home.

The unallocated IPv4 address pool became exhausted in February 2011 when the last five blocks of IPv4 address space were allocated to the RIRs (Regional Internet Registries). While increased address space leasing and trading may provide some relief for some organisations, in general IPv4 address space will become an ever more difficult to acquire resource. The question then posed is how much more difficult it becomes to operate IPv4 networks when no new global address space is available to draw upon, and how IPv6 might be introduced to help alleviate such problems? At the time of writing JANET(UK) does not have a position on acquiring additional IPv4 address space (be it by leasing or trading) for its connected organisations.

When the potential exhaustion of IPv4 address space was recognised in the early 1990s, work began on a successor to IPv4 as well as on technologies and policies to help extend the lifetime of IPv4. Two key events in that area were the deployment of CIDR (Classless Inter-Domain Routing) [RFC1519] in 1993 (later updated to [RFC4632]) and NAT (Network Address Translation) [RFC1631] in 1994 (later updated to RFC3022). The former allowed routed subnets to be of any size and thus sites and ISPs (Internet Service Providers) to be *allocated* network address blocks of any size, rather than being limited to Class A, B or C blocks (which have 24, 16 and 8 bits of host address space respectively³). One major goal of this addressing plan was to allocate Internet address space in such a manner as to allow aggregation of routing. This made address space utilisation far more efficient, while also helping to minimise the size of routing tables seen on backbone Internet routers. Thus CIDR has allowed IPv4 addresses to be allocated with less wastage, and is generally seen as a good thing.

The introduction of NAT allowed multiple IP devices behind a Network Address Translator, typically a router, to use private address [RFC1918] internally and share a limited number, or even just one, global address. NAT is almost de facto in home ADSL networks. However, while NAT is widely deployed, and it clearly ‘works’ from an end-user perspective, it has a number of significant and potentially harmful architectural implications, as documented in [RFC2993]. These concerns have increased the desire for a new version of IP that would allow all hosts to be uniquely globally addressed again.

The discussions which led to IPv6’s development were held within the Internet Engineering Task Force (IETF), the standards body that has defined most of the TCP/IP layer protocols in documents known as ‘Request for Comments’ or RFCs. In principle, the IETF works by focused working groups on specific topics, organised into general areas, with ‘consensus and running code’ being important factors in its processes. The IETF’s mission statement is most recently documented in [RFC3935].

The evolution of the core IPv6 specification within the IETF took some time, having begun back in the early 1990s with proposals such as TUBA (TCP and UDP with Bigger Addresses)

³ A Class A IPv4 network has 8 network prefix bits, with 24 bits available for internal use; this may also be written /8, so for example the Class A network prefix 24.0.0.0 is written 24.0.0.0/8.

[RFC1347], SIPP (Simple Internet Protocol Plus) [RFC1710] and CATNIP (Common Architecture for the Internet) [RFC1707]. Eventually a 'winner' emerged in 1998 with the publication of the IPv6 core specification [RFC2460], which combined some of the features of alternative proposals, as described in [RFC1752].

Where IPv4 is limited to four billion unique addresses, IPv6 uses 128-bit addressing, allowing for 340 undecillion (3.4×10^{38}) unique addresses. Since the original RFC was released, hundreds of additional RFCs have been published further defining IPv6, as well as updating many former IPv4-only protocols to support IPv6.

General information about IPv6 including deployment cookbooks can be found at the 6NET web site [6NET], and the IETF (standardisation) [IETF] and IPv6 Cluster (news) [IST-IPV6] web sites. The 6DEPLOY [6DEPLOY] project web site includes general information as well as training-oriented material. Slides and material used in JANET(UK)'s first IPv6 hands-on workshop, held at Southampton in September 2005, are still very relevant and are also available online.⁴

2.2 The JANET Context

JANET(UK), has played an active role in undertaking technical studies and feeding back its experience into the standardisation path of IPv6. These began back in May 1997 when JANET first attached to the 6bone testbed network and offered IPv6 tunnels to customers that were interested in experimenting. In 1999/2000 the Bermuda and Bermuda 2 projects [BERMUDA], involving the universities of Southampton, UCL and Lancaster, undertook more detailed IPv6 testing and analysis. Then still known as UKERNA, JANET(UK) then participated along with those three universities in the 6NET project [6NET] from 2002-2005, which validated IPv6 technology for deployment on the upgraded JANET backbone.

At present JANET consists of a core network, a set of connecting regional networks (managed by the RNOs, which are generally comprised of consortiums of academic institutions) and the end-site HE and FE campuses attaching to those regional networks. This means that any IPv6 deployment not only has to consider the site and JANET core technology but also that of the intermediate regional networks. In the case of the authors' site, Southampton is connected to the LeNSE regional network, which in turn is connected to one of the JANET C-PoPs (core points of presence). Since the rollout of the current backbone, IPv6 has been a production service and, with LeNSE also supporting IPv6, Southampton has a native IPv6 path to the JANET core and beyond.

JANET also has interconnections with other commercial and academic networks. These include a native IPv6 peering with GÉANT, the pan-European backbone network interconnecting the European NRENs (National Research and Education Networks).

From the point of view of IPv6 deployment, there are thus the cases of the core, regional and end site networks to consider. Each of these cases are discussed in more detail in their own sections later in this guide.

Internationally, best current practice is to run dual-stack backbone networks such that IPv4 and IPv6 coexist on the same routing equipment and network links (as described in Section 7). GÉANT and all of the NRENs served by it have been dual-stack since around 2005, assisted by research and deployment work undertaken in the 6NET project. The US Internet2 network is dual-stack, as is the link between it and GÉANT. Thus native IPv6 connectivity is available throughout these backbone networks, alongside IPv4, without any tunnelling mechanisms being required.

JANET has also followed this dual-stack path and the challenge now is to encourage and assist deployment in the end sites (campuses). Since the deployment of the current backbone, the regional networks have been required to make IPv6 service available to end sites on demand, and thus IPv6 penetration into the RNOs has increased. The small number of HE and FE sites on JANET that have deployed some level of IPv6 all currently do so with dual-stack networking. This allows systems at those sites to use IPv4 to communicate with external IPv4-only sites,

⁴ See <http://www.webarchive.ja.net/services/events/calendar/2005/ipv6-2005/programme.html>

and IPv6 to communicate with external IPv6-only sites (as and when these become more commonplace). Where two dual-stack sites communicate, the choice of protocol to use is left to the application to decide – many applications (like MS Internet Explorer) prefer IPv6 by default.

At present, IPv6-only networks are very rare, and the only academic network to adopt this strategy is the Chinese CERNET2 backbone. Part of the reason for doing this was to encourage sites to adopt IPv6 so they could use the new, higher capacity, less congested backbone network. It also helped generate some focus on developing solutions for IPv6-only systems accessing IPv4-only systems, and tunnelling IPv4 in IPv6.

However, it is unlikely that major European academic networks will deploy anything other than dual-stack in the foreseeable future, or that the upcoming refresh of the JANET backbone will be anything other than dual-stack for its lifetime. However, IPv6-only networks are inevitably going to be more widely deployed over time, and thus interoperability with 'legacy' IPv4 systems will be an important issue.

3. Arguments for IPv6 Deployment

Given that IPv6-only backbones may be many years away in European academic networks, one might ask why any site or network should deploy IPv6 at all at this stage. This section discusses some of the answers to that question (in no particular order of importance).

3.1 Global Address Space

Sites will want to have global address space for a variety of reasons. The principal requirement is likely to be the capability to provide public-facing Internet services that can be reached reliably by other sites and users on the Internet. At the moment, global IPv4 addresses are required for this purpose. A site needs enough IPv4 address space to uniquely number the systems providing those public-facing services.

It is desirable to have enough global addresses also to uniquely number systems within any given site. Universities who received address space before CIDR (prior to around 1993) probably have Class B (/16) IPv4 allocations, giving them 65,000 or so globally unique IPv4 addresses. These sites are likely to have enough global address space for the immediate future, although some of these may be using NAT for student residence (halls) networking.

It is more likely that FE and 'smaller' sites have less global IPv4 address space and may use private addressing and NAT more routinely on their network, often because the provision of the local networking is outsourced. In such networks global IPv4 addresses may largely be used for externally facing services, with NAT used for client systems to access Internet resources and with internal-only communications using private addressing.

Many Higher and Further Education institutions may be split over more than one location and need to access services and servers located on their other campuses. Without using VPNs between sites (to encapsulate private addresses within public addresses routed on JANET), unique addressing is required to perform this.

Given the very low levels of IPv6 deployment on the Internet today (in a later section the stats at Southampton suggest that 2% or less of its external traffic is IPv6, and some of that may be from sites that also run IPv4 anyway) there is, currently, no compelling argument to deploy IPv6 to communicate with other IPv6-enabled sites. However, as IPv6-only networks begin to emerge, making your public-facing services available via IPv6 as well as IPv4 is prudent, as it removes the dependency on translation services (like [NAT64], as discussed in Section 7.3) being required by the other sites and networks communicating with you.

Thus in the short term sites should secure sufficient IPv4 address space to enable their public-facing services to be reachable. In the medium term sites may deploy dual-stack networking to enable access to and from their systems and services via either protocol. In the long term sites may be able to run IPv6-only, assuming technology like NAT64 matures to enable access to IPv4-only content from those IPv6-only systems, or that enough momentum grows with IPv6 deployment that it becomes the de-facto protocol of choice for the Internet.

3.2 Supporting Teaching and Research

Teaching and research are a university's prime roles. There is a good case for universities who may teach networking topics to deploy IPv6 at least in parts of their teaching and research laboratories, to expose undergraduate and postgraduate/research students to the protocol. All graduates will be now emerging from university into a commercial world where no new IPv4 address space is available; thus exposure to IPv6 while studying should be considered a significant benefit to them.

Sites that have made some early IPv6 deployment have observed that the availability of IPv6 has encouraged students to develop their own innovative applications. For example, at Southampton students have created IPv6 radio streaming services [SURGE] as well as the

ECS-TV multicast service [ECSTV], and written a variety of IPv6 file-sharing and similar applications.

Some sites may wish to be more adventurous and experiment with deployment of protocols such as Mobile IPv6 [RFC3775][3775BIS] in support of mobile wireless users on campus. However, it is not expected that there will be significant production usage of MIPv6 in the short term due to a lack of implementations in commonly used devices.

3.3 Early Experience/Insight into IPv6

By deploying IPv6, if only in a small part of a campus or FE site, a considerable amount of experience and insight into the protocol can be gained. In parallel with appropriate training, this allows a site to be better informed when defining its IPv6 requirements for tenders, thus ensuring appropriate IPv6 capability is secured in all procurements, even if that capability is not immediately used. By only sourcing IPv6-capable solutions, the cost of future IPv6 service provision should be minimised because IPv6 protocol support has been included in technology refreshes.

By exposing IT staff to IPv6, confidence and expertise can be built. This can help inform future technology decisions, and ensure all projects – not just the obvious network-specific projects – can consider IPv6 implications.

Identification of ‘problem’ systems (in terms of IPv6 readiness) can ensure early engagement with suppliers/ vendors to establish timely solutions to those problems. If solutions cannot be found, alternative suppliers or products that are IPv6 compliant can be explored.

3.4 Internet Architecture and End-to-End Transparency

One of the original goals of the Internet architecture was end-to-end transparency, made possible by all hosts having a globally unique IPv4 address space. While IPv4 addresses were in abundant supply (MIT received a Class A network allocation for example, big enough for over 65,000 networks of over 250 hosts each), and firewalls and proxies were rarely deployed, the majority of communications enjoyed direct communications, without any device on the path altering packet headers or filtering traffic. As a result, any error in communications could be handled by one of the end systems; they were thus said to be able to participate in fate-sharing their communications.

In today’s Internet, the presence of firewalls, NATs, proxies, dynamically allocated IPv4 addresses etc. means that very often two systems communicating do not fate share – the communication relies on and may be affected by other ‘middleboxes’ [RFC3234] on the communication path. In particular, where NAT is deployed, a new communication into a node behind a NAT may be difficult (requiring special NAT configuration) or impossible. By changing IP headers, NAT also makes classic fate-sharing impractical. Of course, to many people NAT provides a level of security by ‘protecting’ internal nodes but deployment of NAT comes at an architectural price.

One can see a good discussion of the importance of Internet transparency in [RFC2775] and of the architectural implications of NAT in [RFC2993]. The IETF text on Local Network Protection for IPv6 [RFC4864] discusses how IPv6 can provide the perceived advantages of NAT while still using globally unique IPv6 addresses. These documents provide reasonably neutral views of the issues involved.

One could thus argue that with IPv6 applications can be developed without a need to consider and ‘waste’ effort on NAT traversal code and support. As the number and type of IP enabled devices (PDAs, cameras, home appliances, etc.) grows, IPv6 provides support for that growth in a way that IPv4 cannot. In practice any application being written today needs to include consideration of NAT traversal, as that is the environment (typically SOHO ADSL networks) in which most users lie. Until IPv6 is widely deployed, developers will quite understandably write applications to work as best as they can with IPv4 and NAT.

In dual-stack networks, deploying IPv6 alongside IPv4+NAT is a perfectly viable solution, using IPv4+NAT for the conventional mail and web applications, and IPv6 with global addresses for both existing and new applications that leverage predictable end-to-end connectivity, e.g. conferencing, messaging or peer-to-peer style applications. Such applications should be simpler and quicker to develop, and give a more consistent user experience.

With the exhaustion of IPv4 address space now a reality, it may not be long before some ISPs begin to use 'carrier grade' NAT (CGN), i.e. the ISP uses address sharing and NAT within its own network, and thus their customers see two layers of NAT, one internally in their own network and one within their ISP. Such practice will only further complicate application behaviour.

3.5 Security Viewpoint

IPv6 support now ships as standard in all common host and router platforms, and is usually enabled by default, as is the case with Windows 7, MacOS X and most Linux distributions. There is thus a good case to be made for a site administrator to control the usage of IPv6 on their network before more 'adventurous' or malicious users do so themselves (inadvertently or with intent). IPv6 is in your network now, whether you formally support it or not.

Thus, even in a supposedly IPv4-only network, administrators should understand how to ensure IPv6 is not used to create 'back doors' in their network infrastructure, and be aware of potential network problems caused by features of IPv6 that are not being managed or monitored (e.g. rogue IPv6 Router Advertisements [RFC6104]). Ensuring that management and monitoring tools, even in a supposed IPv4-only network, can detect and report IPv6 traffic is important.

4. Basic Operation and Features of IPv6

In this section we discuss the core features of the new IPv6 protocol, as originally defined in [RFC2460] and subsequently followed by further standardisation. We cover core features, the address and packet header formats, packet fragmentation, the Neighbour Discovery protocol and address allocation and management.

Being a new version of IP, IPv6 sits below the transport layer (TCP, UDP) and above the link layer (Ethernet). Thus existing applications and services can continue to work if they are adapted to use new APIs that include the new (for example) IP version-independent socket functions and they run over link layers adapted to support IPv6.⁵ For example [RFC2464] defines transmission of IPv6 packets over Ethernet (the IPv6 Ethernet packet type being **0x86dd**.)

4.1 IPv6 Core Protocol and Features

RFC 2460 defines the nature of the packet format, the packet header(s) and the way in which packet headers may be created and processed.

The key advantage of IPv6 is its 128-bit address space, compared to the 32 bits used for IPv4.

4.2 Address Format

An IPv6 address is written in the form of eight sets of four hexadecimal characters, i.e. xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx, for example

2001:0db8:0001:0035:0bad:beef:0000:cafe.

Leading zeros within each group of four hex digits may be omitted, so one can instead write

2001:db8:1:35:bad:beef:0:cafe.

Further, one can use '::' once in an address to abbreviate a sequence of zeros, e.g.

2001:db8:d0:0:0:0:1 can be written **2001:db8:d0::1**.

The IETF has published recommendations for how IPv6 addresses should commonly be written [RFC5952]. This text also describes the impact of the new IPv6 address format on organisational IT systems, e.g. at a very simple level the higher number of characters used should be considered when designing user interfaces. The impact spreads to other areas of an organisation's IT systems, including logging, auditing and searching functions.

The notation used for specifying the lengths of network prefixes is the same in IPv6 as it is in IPv4. Thus one might write **192.0.2.0/24** to represent an IPv4 network with 24 bits specifying the network prefix length, and thus 8 bits available for host addressing within that subnet. In a similar way, an IPv6 prefix of **2001:db8:2::/64** is a prefix with 64 bits specifying the network prefix length. As we'll see later, IPv6 host subnets are by default /64 in size, due to the way IPv6 SLAAC (Stateless Address Autoconfiguration) works.

4.2.1 IPv6 Literal Addresses

It is worth noting that because of the clash between ':' as a separator in IPv6 addresses and its use as a delimiter in URLs to indicate a port number, a new format for using IPv6 literal addresses in URLs has been defined by [RFC2732]. Rather than writing the ambiguous

http://2001:db8:1::1:8080/

one would instead write the address as

http://[2001:db8:1::1]:8080/

⁵ Existing applications may also use IPv6 automatically if the API does not need updating, e.g. as is the case with Java, if `java.net.preferIPv6Addresses` is set to true.

using '[' delimiters.

4.2.2 IPv6 Address Architecture

The IPv6 address architecture is described in [RFC4291], which defines the various types of addresses and address scopes available in IPv6. The address types are unicast, multicast and anycast. IPv6 does not include the concept of broadcast addresses.

With the addition of ULAs (Unique Local IPv6 Unicast Addresses) as defined in [RFC4193] the available unicast address scopes include:

- *Link-local unicast addresses*, under **fe80::/10**, used on a link (subnet) and unique on the link, e.g. **fe80::230:48ff:fe23:58df**. Nodes can use link-local addresses where no on-link router exists, e.g. on disconnected networks. Packets sent to link-local destinations are not forwarded by routers.
- *Unique Local IPv6 Unicast Addresses (ULAs)*, a unique **/48** prefix available for use in a network, similar to private IPv4 addresses under **10.0.0.0/8** or **192.168.0.0/16**. ULAs have the advantage of being probabilistically unique globally (if chosen randomly as per RFC 4193) and thus should not clash when networks using them merge. Use of ULAs does not imply use of NAT or IPv6. Rather, a site may choose to deploy ULA and global prefixes together, such that hosts are multi-addressed, and may use ULAs for persistent internal connectivity. However, at the time of writing, the author knows of no academic site using ULAs internally.
- *Global unicast addresses*, e.g. **2001:db8:d0:1::2**. There are three specific 'special' prefixes that may be seen.
 - **3ffe::/16** was address space used on the (old) 6bone test network for (tunnelled) IPv6 connectivity in the very early days of deployment from 1996; this has now been phased out [RFC3701] so packets with this prefix should not be seen any more.
 - **2002::/16** is reserved for use by the 6to4 transition method (see Section 7 below).
 - **2001:db8::/32** is reserved for IPv6 documentation examples [RFC3849], as used throughout this text.

The special loopback address also exists in IPv6 and is expressed as **::1**, which is the equivalent of **127.0.0.1** in IPv4. There is now no place like **::1**.

Multicast addresses all fall under **ff00::/8**. Bits 13-16 indicate the scope of IPv6 multicast packets, so for example, **ff02::/16** is link-local scope, **ff05::/16** is site-local scope and **ff0e::/16** is global scope. IPv6 multicast is discussed in much more detail in a separate JANET Technical Guide.⁶

One should assume that it is common for IPv6 nodes to be multi-addressed, at least with a link-local address (which all hosts have) and a global scope address. IPv6 Default Address Selection [RFC3484] rules are used to allow hosts to pick appropriate address pairs when communicating between multi-addressed hosts. For example, if the destination is global scope, a global scope source address will be picked, not a link-local one (which could not be replied to). RFC3484 also prefers longest matching prefixes as a 'tie breaker' for address selection between addresses of the same scope.

As hinted at above, IPv6 has no broadcasts and instead uses (link-local scope) multicast on subnets where IPv4 uses broadcasts. An IPv6 node will join a number of multicast groups, including the 'all nodes' link local scope multicast group (**ff02::1**). One can ping **ff02::1** on a link to discover all hosts on the link.

Another special multicast address is the solicited-node multicast address **ff02::1:ffXX:XXXX** where **XX:XXXX** are the last 24 bits of the node's unicast address. For each unicast address a node has, it responds to link-local multicast traffic directed to the corresponding solicited-node multicast address. This scheme reduces the 'noise' that a node needs to listen for on a link,

⁶ <http://www.ja.net/development/network-engineering/multicast/>

since only it, or hosts with the same last 24 bits in their unicast address, will need to process the multicast packet.

In some cases an IPv6 address output may be seen in the form **2001:db8:1:2::1%dc0**. The % here indicates a scope identifier that indicates a specific interface on the node. To force traffic through a certain interface where the destination link may otherwise be ambiguous, the scope identifier can be used, e.g. **ff02::1%eth0** is the 'all nodes' multicast link-local target on interface eth0. On a BSD system, %dc0 may typically be seen. On Windows 7 it will be an index number like %4 (this can be checked with the command-line ipconfig command).

4.3 IPv6 Header Format

Aside from the size of the addresses, the most obvious difference in the two IP protocols is in the format of the packet header. IPv6 has 'streamlined' the header, reducing the number of fields and making the base header a fixed size, which in principle should help packet processing by routers. The differences between the headers are highlighted in Figure 4.1 below, where the IPv4 header is shown with fields not present in the IPv6 header highlighted with a red (or dark, for viewers watching in black and white) background.

Version	IHL	Service Type	Total Length	
Identifier		Flags	Fragment Offset	
TTL	Protocol	Header Checksum		
32-bit Source Address				
32-bit Destination Address				
Options				

Figure 4.1: Differences in IPv4 and IPv6 header formats

The simplification of the IPv6 header is achieved by the 'Next Header' system, whereby IPv6 headers can be chained, as described later in this section; essentially a set of special headers can be used between the main header and the payload data. The IPv4 fragmentation field has been replaced by a new IPv6 fragmentation header while IPv6 has no header checksum.

The format of the main IPv6 header is illustrated in Figure 4.2 below.

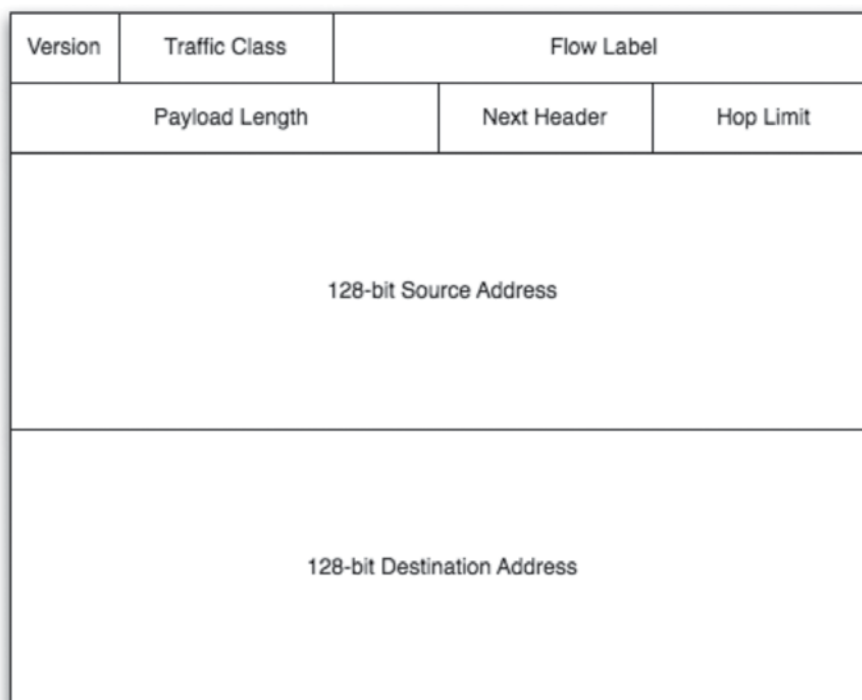


Figure 4.2: IPv6 header format

The IPv6 'Next Header' construct offers potential future expansion of IPv6 by definition of new headers, rather than carving out specific bits of the header for functions as happened with IPv4. Each header contains a field indicating the next header, or the payload data (TCP, UDP, ICMPv6) if it is the last header.

The original (as per RFC 2460) six basic header types are (with Next Header value in brackets):

- *Hop-by-hop*, which is processed by all nodes on a path (0)
- *Routing*, which allows a form of source routing for IPv6 (43)
- *Fragment*, which allows packet reassembly (44)
- *Authentication header* (51)
- *Encapsulated security payload (ESP)* (50)
- *Destination options*, processed only by the final node on a path (60).

In principle the IPv6 header format makes packet processing more efficient and new header types can easily be added. Since the base specification was released the following IPv6 header types have been added:

- *Mobility* (Mobile IPv6) (135)
- *HIP* (Host Identity Protocol) (139)
- *Shim6*, for multihoming (140)
- *WESP* (Wrapped Encapsulating Security Payload) (141).

While new headers can be defined, there is usually some delay before these have support for processing in hardware added, or support in firewall products (for example). Given concerns about adding new features that may require 'slow path' processing, it remains to be seen how readily the IETF will accept further new proposals for new header types.

4.4 Fragmentation

In IPv6 only end nodes may fragment packets. Routers must not fragment packets. Nodes should thus undertake PMTU (Path Maximum Transmission Unit) discovery to facilitate this. For PMTU to be successful it is important that routers and firewalls treat ICMPv6 (Internet Control Message Protocol) packets appropriately, as described in [RFC4890].

All IPv6 links must have a MTU (Maximum Transmission Unit) of at least 1280 bytes.

4.5 ND (Neighbour Discovery) Protocol

The IPv6 ND (Neighbour Discovery) protocol [RFC4861] replaces the function of ARP (Address Resolution Protocol) in IPv4 and makes use of ICMPv6 and multicast.

A node may send a NS (Neighbour Solicitation) to discover the link layer address of a node. It may also send a NA (Neighbour Advertisement) in response to a solicitation message or to announce a link layer address change. Nodes keep ND caches for IPv6 just as they keep ARP caches for IPv4.

When one IPv6 node wishes to find the link layer address of another IPv6 node, it sends an NS message to the solicited-node multicast address of the target address, including its link layer address. The target responds with an NA message containing its link layer address. Using ICMPv6 for this is more media independent than ARP and also allows for use of IP security mechanisms.

4.6 SLAAC (Stateless Address Autoconfiguration)

A new feature for IPv6 is the concept of SLAAC (Stateless Address Autoconfiguration), as described in [RFC4862]. The goal of SLAAC is to increase the 'plug and play' nature of IPv6.

Nodes can autoconfigure the following settings through SLAAC:

- a link-local unicast address as well as (usually) a global unicast address
- the address of the default gateway (router) for off-link traffic
- a list of valid on-link network prefixes (usually just one, /64 size)
- a valid and preferred lifetime for any network prefix in use
- an indication of whether a managed address allocation service (DHCPv6 (Dynamic Host Configuration Protocol)) is available to the node, via a pair of 'Managed' and 'Other' bits
- the link Maximum Transmission Unit size.

The basic mode of operation is that a router on an IPv6 link sends out a periodic link-local scope multicast RA (Router Advertisement). If a node sees the managed bit set in the RA it may choose to change to using DHCPv6 rather than SLAAC, otherwise it will use autoconfiguration.

The node forms an address automatically by taking the observed network prefix in the RA and appending its EUI-64 address, which is constructed from the node's 48-bit Layer 2 MAC address and 16 bits of 'padding' (**ff:fe**)⁷.

For example, if a node has Ethernet address

00:30:48:23:58:df

and the network prefix in the RA is

2001:db8:1:cafe::/64

then the SLAAC address is

2001:db8:1:cafe:0230:48ff:fe23:58df

⁷ If and when devices use 64-bit EUI addresses for Ethernet, the 64-bit identifier may be used directly.

(The change in the top byte of the address from '00' to '02' comes from the global bit being set in the translation from IEEE MAC-48 to EUI-64.)

Rather than waiting for a RA on link, which the router administrator may have configured to be sent only every few minutes (a typical default is every 10 minutes), a node may send a multicast RS (Router Solicitation) message, to which any on-link routers will respond with a unicast or (more commonly) multicast RA. A node will typically send a RS message on startup, or if its interface is brought down and back up, for example.

Note that DHCPv6 does not at the time of writing include a Default Gateway option like IPv4, and thus all links need RAs for hosts to learn of both their default router as well as valid on-link prefixes. Use of SLAAC for DNS resolver discovery is discussed later in this section.

The prefix preferred lifetime is useful for network renumbering. Here, the prefix being deprecated can have its preferred lifetime set to zero, such that while both it and the new prefix are in use together, new connections will prefer and use the valid (new) address on the node. A procedure for how this can be used in IPv6 renumbering can be found in [RFC4192].

An additional feature of IPv6 is DAD (Duplicate Address Detection), as described in [RFC4862]. In principle, IPv6 interfaces will never proceed beyond 'tentative' address binding where clashes are present. When using DAD, a node will join the all nodes multicast group **ff02::1**, then join the solicited node multicast address of the tentative address and send an NS on this address from the unspecified address '::', because it has no as-yet known to be non-clashing address to use. If no NA response on the multicast all nodes group **ff02::1** is seen, the node can assume that the tentative address is safe to use.

4.6.1 IPv6 Privacy Addresses

The fact that a node's Ethernet address may be used to form part of its globally unique IPv6 address led to some privacy concerns. A node using SLAAC in two different wireless hotspot networks, for example, could be correlated by observation of the lower 64 bits of its address. While cookies and other tokens may also be (ab)used in this way, the IETF defined IPv6 Privacy Extensions [RFC4941] to address these concerns.

RFC 4941 essentially allows a node to pick a 'random' host part for its source address on startup, and, if it is configured to do so, to pick a new privacy address periodically. IPv6 privacy addresses are designed to be used as source addresses when initiating communications; a separate static global IPv6 address may still be configured and DNS-advertised for the node to be contacted by other hosts.

A side effect of privacy addresses is that the changing IP addresses of hosts will cause far more unique IP addresses to show up in, for example, web server log files, giving an apparent large increase in unique visitors. There may also be some ambiguity in how network monitoring software behaves – unless multiple addresses can be identified as belonging to the same host over time the software may report 'phantom' systems as being on the network.

4.6.2 Securing SLAAC and Neighbour Discovery

The SEND (SEcure Neighbour Discovery) protocol, as described in [RFC3971], provides a method to secure ND and thus SLAAC. The protocol uses CGAs (Cryptographically Generated Addresses) [RFC3972] that offer assurance of the ownership of an IPv6 address by including a cryptographic hash of the host's public key in its own address – a method only possible in IPv6 due to the size of the address.

While SEND offers an improved security model, there are few public implementations available at this time. These are expected to become available in the near future; support is appearing in JunOS and IOS. Whether academic sites deploy SEND, given they (for example) rarely use Authenticated DHCPv4, is another question.

4.7 DHCPv6

IPv6 address assignment can also be performed by DHCP, for which there are two variants for IPv6.

- The full (stateful) version of DHCPv6 [RFC3315] offers both address assignment (with leases, as per DHCPv4) as well as other configuration information, e.g. DNS resolver address(es).
- The Stateless DHCPv6 [RFC3736] variant only offers other configuration information. This version is intended for (optional) use by hosts that are using SLAAC for address autoconfiguration.

As mentioned previously, neither version of DHCPv6 currently has a Default Gateway IP address option; all hosts must learn the address of their default gateway from observing the source address of RAs seen on their link.

Administrators should also be aware that DHCPv6 uses DHCP Unique Identifiers (DUIDs) rather than MAC addresses in message exchanges. Because DUIDs are generally not known a priori of a host operating system being installed, this raises an interesting problem for administrators pre-provisioning systems (in contrast a host's MAC address is typically available on its case). Recognising this issue, the IETF is defining a new DUID-UUID option for DHCPv6 [UUID], which would allow some persistent firmware-based identifier to be used. It is expected that this will be made available in implementations from later in 2011.

4.8 Address Management

Because of the way SLAAC works, all IPv6 links (subnets) use a **/64** network prefix and thus use 64 bits of host addressing, even where host addresses are manually assigned or assigned by DHCPv6. This may seem rather wasteful of the 128 bits in use for IPv6, but the big benefit for the administrator is that the days of resizing subnets to maximise efficiency in use of IPv4 addresses are gone as far as IPv6 is concerned. As many or as few nodes as the administrator likes can be put on a given IPv6 link.

Network administrators can choose to use SLAAC, DHCPv6, or a combination of both for address assignments and passing other configuration information to end hosts.

In addition, an option to convey DNS configuration information in RAs has been defined [RFC5006]. While few implementations of this currently exist, it offers a method for hosts only using SLAAC to configure DNS information.

Thus the choices open to administrators for IPv6 address management is to use one of

- full stateful DHCPv6 for address and other configuration information
- SLAAC for configuring address and default gateway, and Stateless DHCPv6 for other configuration information (DNS, NTP, etc).
- SLAAC for configuring address and default gateway, and DNS configuration from RAs. This approach assumes no other configuration information is required.

The most common method used today is SLAAC. However, as DHCPv6 implementations mature, and the DUID-UUID option is progressed, it is reasonable to expect that DHCPv6 will get more traction, especially with administrators comfortable with the DHCPv4 accountability and provisioning model (and who have existing tools/ scripts to support that model).

In dual-stack networks, SLAAC can be used for IPv6 address/ gateway configuration information and other configuration can use IPv4 services (e.g. IPv4 DNS resolvers can be used and happily return IPv4 or IPv6 DNS records for queried domains).

On server systems it is quite reasonable to use manually configured IPv6 addresses, just the same as manually configured IPv4 nodes. IPv6 can also use 'tricks' like numbering a node by its function, e.g. **2001:db8:1:1::53** for a DNS server or **2001:db8:1:2::25** for an SMTP relay. Using

SLAAC on servers runs the risk of DNS-advertised addresses changing should a host's Ethernet address change (e.g. with a motherboard or NIC swap-out).

It is also worth noting that DHCPv6 PD (Prefix Delegation) can be used to assign prefixes to downstream routers in a network. In commercial networks the requesting router is usually a CPE router and the delegating router is an ISP's aggregating device/router.

5. Main Differences to IPv4

While there are some notable differences between IPv4 and IPv6, it is important to remember that IPv6 is still 'just IP' so it can run alongside the existing IPv4 protocol and utilise the same protocols above and below IP.

There are for example IPv6 versions of most of the basic command line tools that administrators are used to using in running and troubleshooting IPv4 networks.

However, there are some fundamental differences, some of which have been discussed above in the previous section, and these are highlighted here. While this is not an exhaustive list, it should help capture the flavour of the main differences.

It is useful to keep these differences in mind both when designing IPv6 networks and when troubleshooting them.

5.1 Addressing

The following differences relate to IPv6 addressing:

- IPv6 was designed with hierarchical addressing in mind from the outset. While PI (Provider Independent) address space can be obtained, sites will generally take PA (Provider Aggregatable) address space from the provider, which for UK academic sites is JANET(UK). All UK sites will have a site prefix allocated from within JANET's own allocation (**2001:630::/32**).
- IPv6 nodes will generally be multi-addressed, with at least link-local and global unicast addresses. Dual-stack nodes will of course also have an IPv4 address.
- IPv6 Privacy Addresses mean that IPv6 nodes may change their preferred source addresses over time, even if deployed in a fixed, static network. You should be aware of this for network monitoring, logging and related purposes. Hosts using Privacy Addresses would normally also have a static global address in the DNS to which connections from other hosts can be initiated.
- IPv6 includes Duplicate Address Detection (DAD) by default. You should thus not see address clashes on an IPv6 subnet under normal conditions.
- IPv6's 'infinitely large' subnets make network resizing exercises redundant and deter classic port-scanning attacks. However, while your network is dual-stack, you may still need to resize your IPv4 subnets.
- Point-to-point links may use 127-bit IPv6 prefixes as described in [P2PLINKS], but you may also reasonably use a /64 for a point-to-point link if you are comfortable with or have methods to mitigate the concerns expressed in [P2PLINKS], particularly concerning possible denial-of-service attacks against ND caches.

5.2 Operation

The following differences apply to general IPv6 protocol operation:

- IPv6 has no IP header checksum; checks are assumed to happen at other layers.
- IPv6 has no broadcast traffic, instead it uses link layer multicast.
- There is no IPv6 packet fragmentation done by routers. This means that IPv6 fragmentation must be performed by the end nodes. Hence PMTU discovery is required for IPv6, and firewalls and filters should be configured to allow PMTUD to work.
- The IPv4 ARP functionality is replaced by Neighbour Discovery (ND) in IPv6, which uses ICMPv6 and multicast. A host will maintain a ND cache containing IP to link-layer mappings.

- In dual-stack networks, an application may need to fall back between protocols should one not be available for some reason. It is thus useful for any filtering devices (firewalls) to return unreachable messages rather than just silently dropping traffic, to allow the host to fall back to the other protocol rapidly.

5.3 Multicast

We discuss IPv6 multicast in a later section but the key differences to IPv4 multicast are:

- Multicast Listener Discovery (MLD) [RFC3810] replaces the function of IGMP for IPv4. MLDv2 supports SSM operation. Where you may currently procure switches with IGMP snooping capability, you will now want MLD snooping in addition.
- Scoping of group addresses is explicit in the group address itself, through a 4-bit scope field. The common scopes are 2 (link), 5 (site), 8 (organisation) and e (global), e.g. ff02::1 is all hosts on the local link. This makes configuring scope boundaries on routers much simpler.
- Embedded-RP [RFC3956] removes the need for Multicast Source Discovery Protocol (MSDP) by defining how an IPv6 RP address can be embedded/included in an IPv6 multicast group address, if the RP address is restricted to a certain format. This allows routers supporting Embedded-RP to know exactly where the RP for the given group is, and no longer need to have that location configured or passed via MSDP. End-to-end Embedded-RP support is needed though between a sender and any potential receiver.
- It is easy to obtain a global scope multicast address based on your site's unicast IPv6 allocation, as defined in [RFC3306]. But with no MSDP there can only be one RP per multicast group.

5.4 IPv6 and NAT

It is not expected that NAT will be used for IPv6, given that doing so would defeat one of the primary design objectives of IPv6.

However, there is an argument for defining a NAT66 stateless mapping service between internal and external IPv6 network prefixes to allow a site to avoid the need to renumber its internal network when changing ISP. NAT66 is currently being progressed as a personal Internet Draft as described in [NAT66].

The address independence offered by NAT66 is now arguably less compelling since the RIRs now have policies to assign /48 IPv6 PI space to end sites (subject to certain qualifications).

UK educational sites should be able to get IPv6 address space from JANET(UK), and should not generally need their own independent address space given that JANET(UK) is their default connectivity provider.

There are some sites which today use IPv4 NAT for additional perceived benefits beyond address independence. These properties are achievable in IPv6 networks, as described in [RFC4864], Local Network Protection for IPv6.

6. IPv6 Address Allocation and Management

This section describes how IPv6 addresses are allocated to the NRENs – in the case of the UK, to JANET(UK) – and from there to their constituent academic sites and networks.

6.1 Hierarchical Addressing

While the larger address space is IPv6's most important new feature, there are other benefits as well. One of these is the preferred use of hierarchical address allocations for all ISPs and end sites. ISPs in Europe get their IPv6 address allocations from the RIPE NCC, which manages allocations as the RIR for the European region. End sites get their address space from the ISPs. This means the address space allocated to networks and end sites can be aggregated within the ISP providing service, reducing the number of specific network routes that need to be held in the backbone routers of the Internet (in the so-called Default Free Zone). With IPv4 there are over 360,000 such routes seen as of February 2011 but only 4,500 IPv6 routes.⁸

If one takes the example of the UK universities, in IPv4 many universities have distinct prefixes, often 16-bit network allocations. This is because, prior to 1993, JANET site/network prefixes were allocated in a non-hierarchical fashion directly from the top-level authority. Only after then, with CIDR introduced, did JANET begin to allocate prefixes from contiguous blocks received from RIPE NCC. Thus a backbone router outside JANET may need around 200 routes for UK universities, while internally on JANET around two thousand routes may be seen. With IPv6, all UK universities take address space from JANET's allocated /32 prefix and thus the route for only that one IPv6 prefix needs to be advertised outside JANET. This level of organisation may also help with tracing origins of traffic more easily, when sites outside JANET observe JANET-originated datagrams.

In principle, using hierarchical, or Provider Aggregatable (PA), addressing should reduce the size of backbone routing tables. However in practice end sites will still have the same desire for Provider Independent (PI) address space as they do in IPv4 – for traffic engineering or to avoid renumbering when changing provider. For this reason PI addressing policies for IPv6 have been agreed by the RIRs, making PI IPv6 address space available. While PA address space remains the default for IPv6, the IPv6 backbone routing tables are likely to grow as IPv4-style practices begin to be adopted for IPv6. Again though, HE/FE sites should not need to use anything other than the /48 PA site prefix allocated to them by JANET(UK).

6.2 Allocations to JANET

In Europe, production IPv6 and IPv4 address space is managed and allocated by the RIPE NCC [RIPE]. JANET obtained its production IPv6 prefix via RIPE NCC back in October 1999 and has been using it for IPv6 services in UK academic networks ever since, starting with the Bermuda 2 project [BERMUDA] pilots in 1999/2000.

Prior to 1999, JANET used the 6bone prefix **3ffe:2100::/24**, but JANET's use of 6bone address space was phased out with the introduction of the production prefix.

The recommended default prefix to allocate to an ISP is a /32, as agreed by a common policy between the Regional Internet Registries. JANET(UK) has been allocated **2001:630::/32** by the RIPE NCC.

6.3 Allocations to End Sites (Universities, Colleges)

The recommendations for sizes of prefixes used in allocations to IPv6-enabled end sites are discussed in [RFC3177]. Until recently, the default recommendation for all end sites is that they receive a /48 size allocation.

⁸ See <http://www.potaroo.net>.

An update to RFC3177 [ALLOC56] is however suggesting removing /48 as a 'one size fits all' allocation, which would more readily allow use of /56 size prefixes for SOHO (Small Office/ Home Office) size sites. The aim of this recommendation is to reduce the size of the top level prefix needed by a large ISP offering IPv6 services to millions of SOHO users.

Regardless of the above refinements to RFC3177, a campus site can assume it will receive a /48 size prefix from the JANET Service Desk for use on its network. In IPv4 terms, that is similar to an IPv4 Class A (or /8) prefix, in that one can deploy over 65,000 subnets containing hosts (253 hosts in IPv4's case, and an effectively unlimited number of hosts for IPv6).

Full details of how IPv6 address space can be obtained for use on JANET-connected sites and networks from the JANET Service Desk.⁹

Once a /48 is allocated to a campus site, the registration details are held by the JANET Service Desk and also passed on to RIPE NCC. Should a campus site wish to receive more than a /48 allocation, this can be considered on a case-by-case basis where justification is provided. It is not expected that campus sites will require more than a /48 allocation unless they provide significant IPv6 services to third party sites.

6.4 Allocations to Regional Network Operators

At the present time, the regional networks attached to JANET do not provide address space directly to their connected campus sites. It is expected that for IPv6 networking, the sites will obtain address space directly from the JANET Service Desk as described above, and that the regional networks will facilitate the routing of that address block through their network between the JANET core network and the end site.

A regional network will typically be able to obtain its own /48 prefix for numbering its own IPv6 equipment from JANET(UK). The procedure is as described above for end sites.

There have not yet been cases where regional networks have required more than a /48 allocation; such scenarios should be discussed with JANET(UK) as and when they arise.

6.5 Network Addressing Plans

A general strategy for an IPv6 internal networking (subnetting) plan would be, where IPv6 is deployed dual-stack, to have congruent IPv4 and IPv6 subnets. Given that many sites make their subnetting plans based on either physical or 'political' (policy) grounds, those factors are unlikely to be different for IPv4 and IPv6. This is the approach taken at Southampton as reported in the case study in a later section.

More detailed comments on address planning are included in Section 9.2 below.

⁹ <http://www.ja.net/services/connections/ip-address-application.html>

7. IPv6 Transition and Coexistence Tools

IPv6 transition includes the integration of, co-existence of and inter-operation between IPv4 and IPv6 networks and devices. It is a very broad subject; witness that the reports written on backbone/ISP and site transition methods in the 6NET project back in 2005 took up a few hundred pages between them. For details, including configuration examples, the reader is referred to [D224] and [D234], much of which is still valid today. This section briefly overviews the most common tools.

The term 'IPv6 transition' is a little fuzzy. In principle it means a transition or migration from IPv4-only operation to IPv6-only operation. In practice, IPv6 deployment will be phased in such a way that IPv4 and IPv6 co-existence is the medium-term strategy. IPv4 is likely to remain around for a long time, even as pockets of IPv6-only deployment grow, so it is preferable to talk of 'coexistence' rather than 'transition'.

Transition tools need to handle a variety of different scenarios of IPv4 and IPv6 interworking. There are broadly three classes of transition and coexistence tools:

- *Dual-stack* – here systems and networks run both protocols and can communicate with each other using either protocol.
- *Tunnels* – these generally involve encapsulating IPv6 over IPv4 links, to allow IPv6 'islands' to communicate over IPv4-only paths, such that the IPv6 packet is the payload of the IPv4 packet. This requires open 'holes' in firewalls to achieve, i.e. packets whose Protocol field is '41'. As IPv6-only networking begins to emerge, IPv4-in-IPv6 tunnels are likely to be required.
- *Translation* – for communicating between IPv6-only and IPv4-only systems; translation can occur at the IP layer (e.g. using [NAT64] with [DNS64]), the TCP layer (e.g. through dual-stack TCP relays) or the application layer (e.g. using dual-stack application layer gateways).

Each of these approaches is discussed below.

7.1 Dual-stack

Dual-stack is currently the preferred solution, due to the overheads and complexity of using tunnelling or translation methods.

Most of the European and international NRENs have deployed IPv6 dual-stack on their backbone networks, including JANET. The case study later in this guide describes IPv6 deployed dual-stack at the University of Southampton (Electronics and Computer Science). Deploying dual-stack assumes enough IPv4 addresses are available, but it is also perfectly possible to deploy IPv6 with global addresses alongside IPv4 and NAT.

Dual-stack networking places the 'complexity' in management in the network as a whole. All network elements, and their associated management and monitoring systems, need to support dual-stack operation. Ideally, these systems should be manageable consistently as one network, rather than feeling like two independent networks. In contrast, translation methods such as NAT64 place the complexity at the edge; NAT64 also implicitly assumes all the necessary internal elements can operate in IPv6-only mode.

An important consideration for dual-stack is security, i.e. securing both protocols consistently and appropriately. A summary of IPv6 transition security issues can be found in [RFC4942].

7.2 Tunnelling

Tunnelling may be used host-to-host, host-to-router or router-to-router. Tunnels may be set up manually, on demand by the end site/user, e.g. using a tunnel broker [RFC3053], or automatically, e.g. as described for 6to4 [RFC3056] below.

7.2.1 Manual Tunnels

The most common examples of manually configured tunnels are generally router-to-router tunnels between IPv6-capable networks where the intervening network does not support IPv6 (i.e. is IPv4-only).

Figure 7.1 shows an example of router-to-router tunnelling. The original IPv6 packet from A has source address **2001:db8:1::1** and destination **2001:db8:2::1**. It is encapsulated at the dual-stack site exit router in an IPv4 packet as Protocol 41, with source address **152.78.1.1** (the local tunnel endpoint) and destination **152.78.70.12** (the remote tunnel endpoint).




Figure 7.1: IPv6 in IPv4 tunnel

Manually configured tunnels are still used by a small number of JANET sites today for connecting to the JANET Experimental Service, where the tunnel is required to traverse an IPv4-only regional network between the campus/site and the JANET core network, or where a site is not yet in a position to enable dual-stack IPv4/IPv6 on its access router. Currently JANET supports a tunnel server on the core network for the experimental IPv6 service, even though IPv6 has become part of the RNO's standard SLA with JANET. However, now that many RNOs offer IPv6 natively the need for JANET to support a tunnel service should decline.

If a site has only one IPv6 uplink to its ISP (e.g. JANET) as a tunnel, setting up a manual entry is not a big overhead. However, the provider has to maintain a set of endpoint relationships, which can become a significant task as the number of tunnelled customers grows, especially if the rate of change of tunnel endpoints is high.

The next section looks at the tunnel broker, which can be used to allow end sites or, much more commonly, individual users to get IPv6 connectivity without provider intervention.

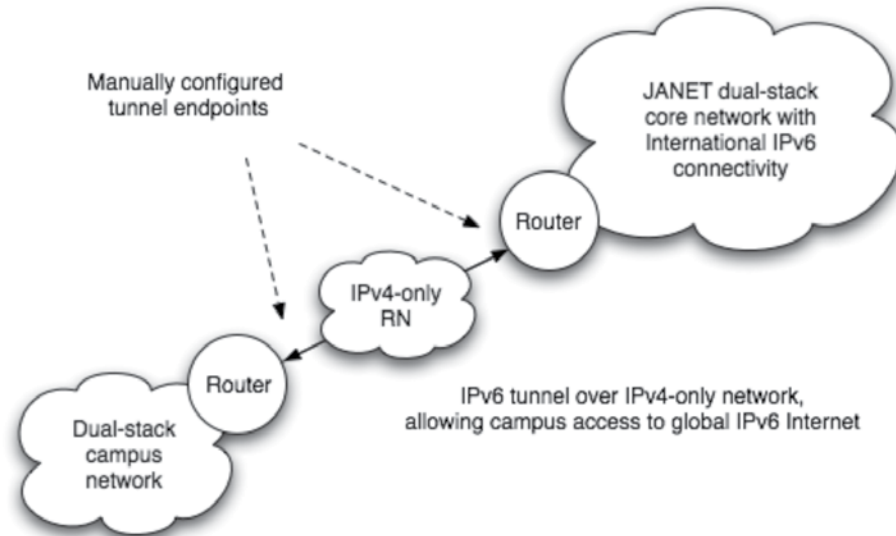


Figure 7.2: A tunnel to the JANET IPv6 Experimental Service

7.2.2 Tunnel broker

The IPv6 tunnel broker, described in [RFC3053], allows an IPv6-in-IPv4 tunnel to be established by a user with an IPv6 capable system (typically a dual-stack system) in an IPv4-only network, by tunnelling from the system to the tunnel broker server. It is primarily aimed at end-users wanting connectivity, but could be used to connect sites (though the JANET Experimental Service uses manually configured tunnels). The architecture is illustrated in Figure 7.3 below.

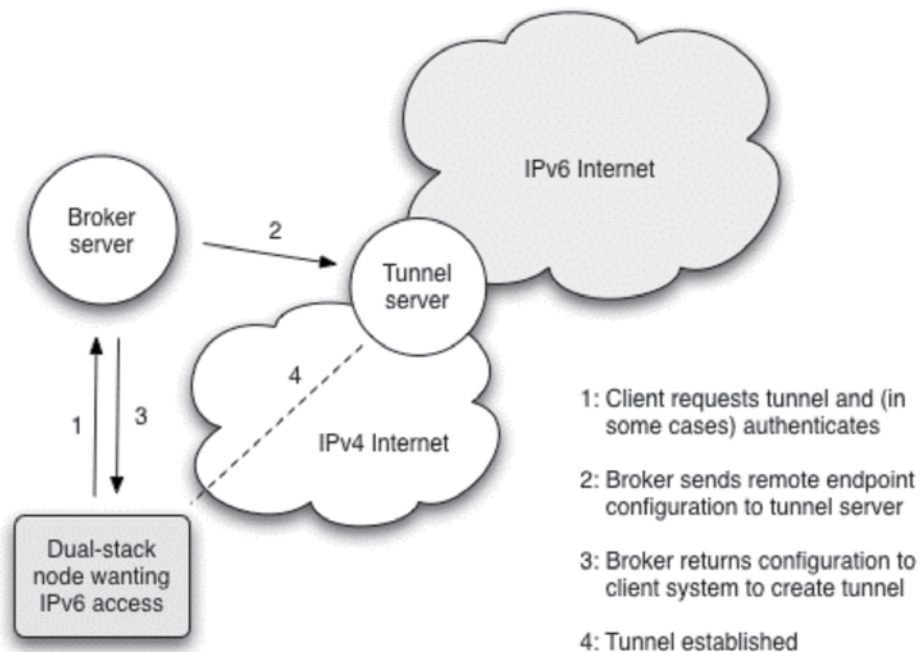


Figure 7.3: IPv6 tunnel broker

The broker service is generally presented to the end user as a web site to visit to request a tunnel. The user connects to the broker's web interface and will typically need to register to then be able to request a tunnel. When the user does this, the broker returns a script or data that the client can use to set up their end of the tunnel, while it also communicates parameters to the broker tunnel

server to set up the remote end of the tunnel. When the client executes the script, they gain tunnelled IPv6 connectivity.

A tunnel broker is a very convenient tool for a user who wants to do some early IPv6 testing or wants IPv6 connectivity off campus, e.g. in their home ADSL network. The broker can provide a single endpoint address, e.g. for a user on a laptop at a conference, or an IPv6 prefix (to connect a small site).

TSP (the Tunnel Setup Protocol) [RFC5572] may also be used to assist the broker in setting up connection parameters and handling elements such as authentication. TSP also supports broker client operation behind an IPv4 NAT by using UDP tunnelling in conjunction with a dedicated piece of client software. This allows users in typical home ADSL networks to use a TSP-based broker for IPv6 connectivity.

JANET(UK) deployed a broker pilot [BROKER] in 2006 which has been open for use by anyone in the UK academic community (that is, anyone with a working .ac.uk e-mail address). The pilot broker also supports TSP. The broker is under review at the time of writing.

A key consideration for tunnel setup is to use a topologically close tunnel endpoint, else the first hop to any IPv6 destination may be a long one. The most well-known public IPv6 brokers available with a UK presence are the Hurricane Electric broker¹⁰ and the SixXS broker.¹¹ Both are free to use and give good performance to UK academic users. Both brokers are quite 'friendly' in supporting the user's IPv4 endpoint address changing periodically. The HE broker also includes some 'tutorials' by which users adding additional capabilities to their connected IPv6 hosts or networks score 'points' – there is no reward as such but it is an interesting idea to encourage users to experiment further with IPv6.

7.2.3 6to4

The 6to4 protocol, defined in [RFC3056], offers automatic IPv6-in-IPv4 tunnelling between IPv6 site networks. Its primary goal is to facilitate automatic tunnelling of IPv6-in-IPv4 between 6to4 site routers deployed on the IPv4 Internet.

The architecture of 6to4 in its basic form is shown below in Figure 7.4.

There is a special IPv6 prefix allocated to 6to4, namely **2002::/16**. The next 32 bits of the 6to4 site **/48** prefix are then set to the IPv4 address of the tunnel end-point device, which is usually a site border router, though 6to4 can be used by a host. When a 6to4-aware router sees this special prefix in the IPv6 destination field of a packet it needs to route, it knows it should tunnel that IPv6 traffic in IPv4 towards the destination IPv4 address indicated in bits 17 to 48 of the 6to4 site prefix. For example, the 6to4 prefix **2002:c000:0101::/48** indicates a site network behind a 6to4 router that can receive traffic tunnelled to **192.0.1.1**. The 6to4 site router will typically be dual-stack on the internal-facing interface and IPv4-only on the external-facing interface.

10 <http://tunnelbroker.net/>

11 <http://www.sixxs.net/>




Figure 7.4: Basic 6to4 usage between two 6to4 routers

This method can be very effective for connectivity between 6to4 sites (routers), because it uses the optimal IPv4 routing paths.

However, the catch with 6to4 is that it is not very good at communicating between 6to4 and the ‘real’ IPv6 Internet. Imagine in Figure 7.4 that one of the hosts in the 6to4 site networks wished to speak to an IPv6 device on the address **2001:db8:10::80**. It would not have connectivity since it only knows how to tunnel to other 6to4 networks under the prefix **2002::/16**. In this scenario the 6to4 router needs to tunnel towards a special 6to4 router that is known to have connectivity to the IPv6 Internet. This special 6to4 router is known as a 6to4 relay.

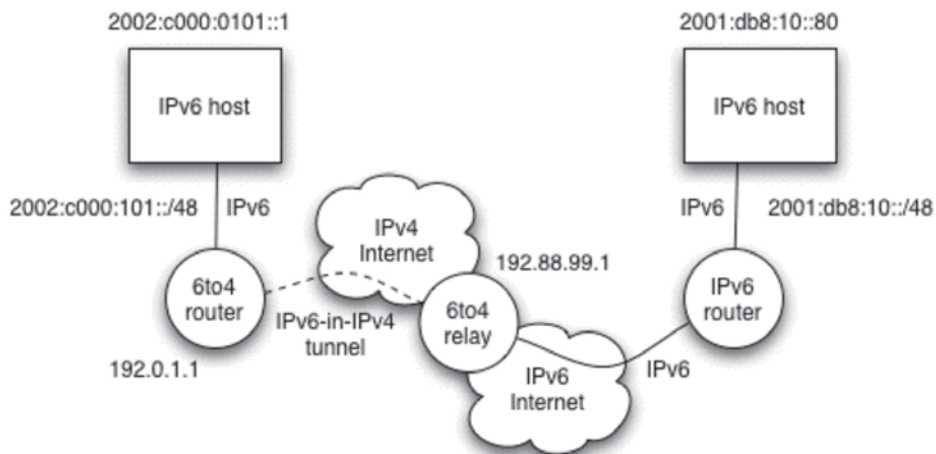


Figure 7.5: 6to4 with a 6to4 relay

The relay has a 6to4 interface on the IPv4 Internet and an IPv6 interface on the IPv6 Internet. It advertises **2002::/16** towards connected IPv6 networks, offering them the ability to route to 6to4 networks, and it advertises a host route to the IPv4 address of the relay on the IPv4 Internet. The advertised IPv4 relay address is a special anycast address (see [RFC3068]) and is set to **192.88.99.1** in the case of JANET’s 6to4 relay. The 6to4 architecture including use of a 6to4 relay is shown in Figure 7.5.

However, because such a relay could be abused, routes for it are not advertised outside JANET, so non-JANET users, e.g. those on home ADSL networks, won’t be able to use it. A JANET-connected site should be able to see the JANET 6to4 router and can test this by pinging or tracerouting to the 192.88.99.1 anycast address.

The common operational issue with 6to4 is that any 6to4 routers external to JANET may have trouble finding an appropriate relay, if their (IPv4) ISP does not offer such a service. That is to say, a remote site using 6to4 can only communicate with JANET sites using ‘real’ IPv6 if the remote site has access to a working 6to4 relay.

In some scenarios basic 6to4 can be very useful, e.g. connections between two student SOHO ADSL networks, but given the issues commonly observed when RFC3068 is used as well, a tunnel broker service is recommended instead for general personal IPv6 connectivity.

7.2.4 6rd

The 6rd protocol [RFC5569] was developed for use in a large ISP in France, Free.fr. The main advantage of 6rd, as shown in Figure 7.6, is that as an automatic tunnelling solution it uses the IPv6 prefix of the ISP supporting 6rd, rather than a dedicated prefix like 6to4. As a result, all traffic for 6rd users in the Free.fr network is routed normally to the ISP and all tunnelling is constrained within the ISP.

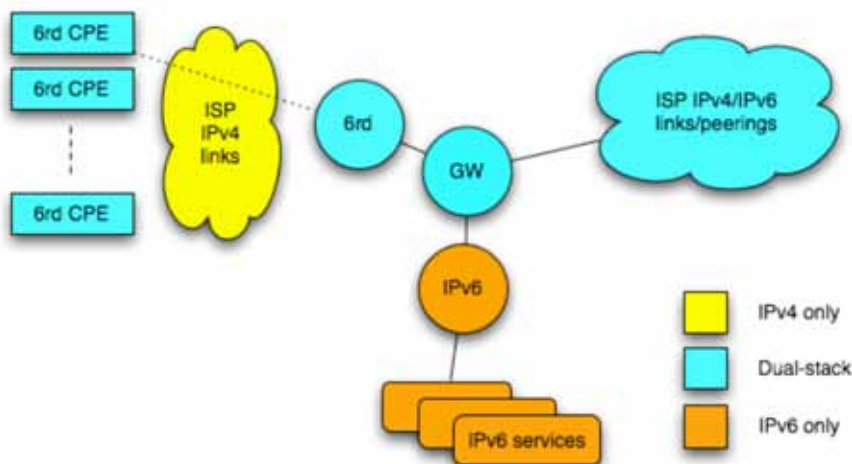


Figure 7.6: 6rd architecture

The tunnelling is automated between the 6rd router and participating customer CPEs, and thus occurs totally within the ISP's own address space; to external sites, the customers appear to have native IPv6 connectivity. Because 6rd has less bits to embed the IPv4 tunnel end point in the site IPv6 prefix, the protocol can recover extra bits to use by omitting the fixed IPv4 network prefix bits belonging to the Free.fr ISP.

Free.fr has approximately 1,500,000 customers. Each user can turn on 6rd on their CPE as and when they wish to try using IPv6. The ISP can also choose to introduce new IPv6-only services in their network for those users as shown in Figure 7.6.

7.2.5 Teredo

Teredo is a tunnelling protocol [RFC4380] designed to be used as a 'method of last resort' when a host wishes to use IPv6 to communicate with a remote IPv6-only system (if the remote system is dual-stack, RFC3484 address selection rules should mean the host prefers native IPv4 ahead of Teredo). It uses UDP encapsulation to allow operation through most types of IPv4 NAT.

Figure 7.7 shows an overview of Teredo. The client initially connects to a Teredo server to get its configuration information, the default server being the Microsoft-provided ones running on `teredo.ipv6.microsoft.com`, port 3544. The initial messaging is designed to punch a NAT hole to allow the client to talk to remote IPv6 systems through a Teredo relay. Thus, akin to 6to4, ISPs need to deploy local 6to4 Teredo relays to improve/ensure proper operation.

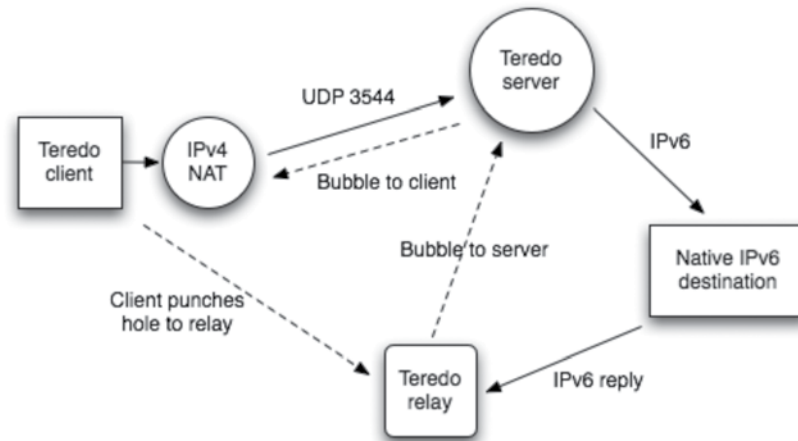


Figure 7.7: Teredo overview

The Teredo service uses the reserved prefix **2001:0::/32**, which as a side effect allows easy identification of peers using Teredo. Site administrators can also look for signs of clients in their site trying to use Teredo by looking for outbound UDP traffic on port 3544.

For a more detailed description of Teredo see the MS TechNet page.¹²

For a Linux implementation of Teredo, see Miredo.¹³

7.3 Translation

Translation methods are designed to allow IPv6-only systems to communicate with IPv4-only systems. Translation can happen at various layers:

- at the IP layer, for example using [NAT64] with [DNS64]
- at the TCP layer, for example using a TCP relay, e.g. TRT as described in [RFC3142]
- at the application layer, using an application layer gateway, e.g. a web proxy, IRC server, SMTP gateway, or H.323 proxy such as OpenMCU.

The recommended way to translate between IPv4 and IPv6 only systems is to use an application layer gateway; these are often in use for IPv4 today, for various reasons including caching and NAT traversal. Many applications or services, like MX relays or DNS resolvers, support proxying or caching naturally.

As an example, a dual-stack mail (MX) relay can receive emails sent from internal IPv6-only mail clients/servers, and forward these to IPv4-only or IPv6-only external mail relays, as shown in Figure 7.8.

¹² <http://technet.microsoft.com/en-us/library/bb457011.aspx>

¹³ <http://www.remlab.net/miredo/>

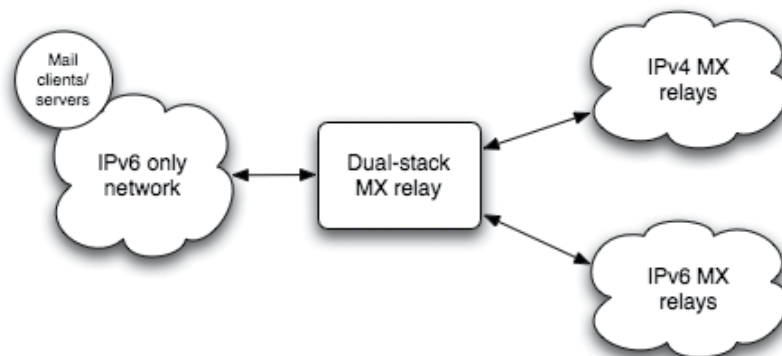


Figure 7.8: Dual-stack MX relay as an ALG

The TRT approach has been deployed at a small Norwegian university site as part of the 6NET project, but has not been used on a large scale.

IP layer translation was initially proposed using NAT-PT [RFC2766] but a number of deficiencies have been documented in this protocol and it has since been deprecated by the IETF, as described in [RFC4966], and thus is not currently recommended.

By focusing on the problem of IPv6 clients talking to ‘legacy’ IPv4 servers, a more specialised IPv6-to-IPv4 translation and DNS mapping service is being defined, as per [NAT64] and [DNS64]. Early implementations¹⁴ of NAT64 are encouraging but have not at the time of writing been tested on a large scale.

7.4 IPv4 Mapped Addresses

There is a special IPv6 address format used to represent an IPv4 address where a socket API (Application Programming Interface) may receive an IPv4 datagram on an IPv6 socket, known as an IPv4-mapped address.

The format is `::ffff:<IPv4 address>`, e.g. `::ffff:152.78.64.1`.

One socket is used for both address families. Such addresses should not be seen on the wire, i.e. not as source or destination addresses. They may surface in certain applications, e.g. in web log files, depending on the software used and its configuration.

7.5 Transition Tools on JANET

The following are examples of the use of transition tools on the JANET network, supported by the JANET NOC (Network Operations Centre) or JANET(UK). Some have already been mentioned above.

- The JANET core network backbone is dual-stack.
- JANET currently offers a manually configured tunnel service for sites wishing to get IPv6 connectivity, e.g. where their regional network does not yet support IPv6 [JANETEXP]. This is the recommended JANET-connected site connection method, for a testbed or fuller deployment where it isn’t possible to enable dual-stack on a site’s JANET access router.
- JANET(UK) offers a pilot IPv6 Tunnel Broker service [BROKER]. This is the recommended host/user or small testbed connectivity method. Alternative free brokers with UK servers are offered by Hurricane Electric and SixXS.

¹⁴ <http://ecdysis.viagenie.ca/>

- A 6to4 relay is operated by JANET, using the IPv4 anycast address 192.88.99.1, as defined in [RFC3068]. Use of 6to4 to connect sites is not recommended but provision of a relay on JANET improves the probability that native IPv6 sites on JANET can successfully reply to traffic received from 6to4-based sources.

The JANET(UK) Tunnel Broker pilot service uses its own prefix from the JANET allocation. Thus if a site wishes to use its JANET Service Desk allocated production IPv6 /48 prefix from the start of deployment, and for some reason cannot establish native IPv6 connectivity via its regional network, it should instead use the Experimental Service and a manually configured tunnel to the JANET NOSC router. Most RNOs now have IPv6 deployed, so this should become increasingly less required.

It should be noted that there is no single 'best method' for IPv6 transition; each tool has a place and a use, and the important consideration is deploying them in a complementary way.

8. JANET and IPv6

The JANET network currently supports IPv4/IPv6 dual-stack in production, and many of JANET(UK)'s services are either dual-stack now or will be in their next refresh.

8.1 History of IPv6 on JANET

JANET(UK) has been actively involved in IPv6 deployment since being one of the first participants in the 6bone experimental network in 1997. The 6bone, using the specially assigned and reserved prefix **3ffe::/16**, was an overlay network of IPv6-in-IPv4 tunnels, largely using routing equipment separate to the production IPv4 networks. While connectivity was somewhat hit and miss at the time, and the Gordian knot of tunnels and (sometimes) lacking routing policies meant 6bone users experienced rather mixed results, the lessons learnt even at that early stage were very valuable.

8.1.1 Production IPv6 networking

In 1999 JANET requested production address space from RIPE NCC and was allocated the prefix **2001:630::/35**, which was later expanded to **2001:630::/32** when RIPE's assignment policy changed. This prefix has been the basis for all production deployment on JANET. The 6bone 'experiment' was finally phased out in 2006; there should no longer be any use of **3ffe::/16** prefixes on the Internet today.

While the early piloting work, including the Bermuda [BERMUDA] project in 2000, gave good insight to IPv6 deployment issues, it was the 6NET project [6NET] from 2002-2005 that proved invaluable to JANET and to the UK universities involved, namely Southampton, Lancaster and UCL.

6NET has produced a whole raft of cookbooks on IPv6 deployment best practice and the reader is referred to these for details. The project allowed JANET(UK) to gain shared experience in IPv6 with other European NRENs, and this assisted in the dual-stack enabling of the core JANET network, at the same time as other European networks and GÉANT were also deploying IPv6. As a result, by 2005, where end sites/universities and their regional networks supported it, end-to-end native IPv6 paths were available across JANET on a production basis.

During the lifetime of the present backbone, the RNOs were expected to provide IPv6 services on demand of connected sites as part of the standard SLA with JANET. As a result, since the introduction of the present backbone, IPv6 deployment has extended into many more regional networks.

8.1.2 International peerings

JANET has IPv6 interconnections with other commercial and academic networks, such as the dual-stack peering with GÉANT, the pan-European backbone network that interconnects JANET with research and education networks worldwide. Most other R&E networks also run a dual-stack network; CERNET, the Chinese R&E network, being a notable exception where their next generation network is IPv6 only. IPv6 is also a requirement for JANET's providers of transit to the global Internet, and where possible all peerings between JANET and other service providers are dual-stack. One notable dual-stack peering is with Google, which has permitted JANET sites to participate in Google's IPv6 trial¹⁵ during 2010/11.

As a result, JANET is fully connected to the IPv6-enabled Internet. In the same way as a site can expect IPv4 connectivity to any other part of the Internet, if it enables an IPv6 connection to JANET it can expect access to any other IPv6 network connected to the Internet.

¹⁵ <http://www.google.com/intl/en/ipv6/>

8.2 Connecting to JANET with IPv6

JANET sites should seek to establish IPv6 connectivity natively. The RNOs are required by their SLA to provision IPv6 natively on demand. The JANET IPv6 Experimental Service is now only a fallback position to facilitate a tunnelled connection for sites that for some reason cannot go dual-stack on their access router, or where there is some other extenuating circumstances with the RNO. The connecting site need only complete a brief application form on the JANET(UK) web site.¹⁶

The default address allocation to requesting organisations is a **/48** prefix.

JANET(UK) had already allocated the 100th IPv6 prefix to connected organisations before the last IPv4 /8 blocks was allocated to the RIRs. While at the time of writing only a small number of these are in use in significant deployments, the volume of allocations is an indication of interest in IPv6 in the JANET community.

The service also extends to IPv6 Multicast connectivity, although this uses a different experimental router and may thus be subject to a differing level of support (as an experimental service).

8.3 IPv6-enabled JANET Services

JANET(UK) has enabled IPv6 on many of its production services. Other services will follow as and when technology refreshes happen.

8.3.1 Web site

The JANET web site at <http://www.ja.net> is available over IPv6.

```
$ ping6 www.ja.net
PING www.ja.net(2001:630:1:107:1::65) 56 data bytes
64 bytes from 2001:630:1:107:1::65: icmp_seq=0 ttl=56 time=6.69 ms
```

8.3.2 DNS

The ja.net DNS servers have responded over IPv6 transport since 1994, as illustrated by querying for information on ns0.ja.net:

```
$ dig -t any ns0.ja.net
;; ANSWER SECTION:
ns0.ja.net. 49783 IN AAAA 2001:630:0:8::14
ns0.ja.net. 49783 IN AAAA 2001:630:0:9::14
ns0.ja.net. 33206 IN A 193.63.94.20
ns0.ja.net. 33206 IN A 128.86.1.20
```

As IPv6 has grown into a mature technology, the rollout of the protocol on JANET has been progressive. Support for IPv6 transport DNS on JANET's primary nameservers indicates the level of confidence in the technology, as does the deployment of dual-stack networking services on the backbone without adverse impact on the IPv4 service.

8.3.3 Other services

Other services that are IPv6-enabled include NTP and the MX servers for ja.net and nosc.ja.net. Sites running eduroam may opt to run their RADIUS peerings to the NRPS via IPv6 transport.

¹⁶ <http://www.ja.net/services/connections/ip-address-application.html>

9. Deploying IPv6 in academic enterprise networks

Deploying an IPv6 capable network infrastructure may not yet be a top priority for everyone. As with other networking requirements, a large university undertaking research has different requirements from a school or college whose focus is centred entirely on education.

That said, a prudent step for all is to ensure that all new IT equipment and services being procured are able to support IPv6 – even if the technology is not deployed immediately.

It is important to distinguish between deploying an IPv6 network and deploying only new or upgraded equipment and services that are IPv6 capable. The former could be considered desirable in the JANET environment; the latter should now be considered essential.

9.1 Multiphase deployment

An IPv6 deployment is likely to be different for each site considering it, depending on their rationale for deployment, their prior knowledge and experience, and where on campus IPv6 services and capability are to be enabled.

The general principles for deployment however probably share a good deal in common. Whatever project methodology you use, the steps are likely to be similar, even if the technical solutions may differ.

In this section we do not talk of specific project frameworks (e.g. PRINCE2); instead we suggest a phased deployment model that should be of use to any generic site.

The pace of deployment between the phases is for the site to decide. In principle the steps could be followed rapidly but the advance planning and piloting can be, and usually are, done some time in advance of production deployment, and production deployment itself can be staged, e.g. enabling external IPv6 web access before turning on any internal IPv6 capability.

9.1.1 Dual-stack or IPv6-only

At present it is recommended that any campus deploying IPv6 does so as a dual-stack deployment, i.e. running IPv6 alongside IPv4, whether as global IPv4 with global IPv6 addressing, or IPv4 NAT with global IPv6 addressing.

The alternative is to run IPv6-only. However to do so a site would need to remove all existing IPv4 capability, to ensure that all systems and services in its network can use IPv6, and then deploy large-scale NAT64 (and DNS64) translation support at the campus edge for access to legacy external IPv4 resources. As yet, NAT64 is not proven on a large scale, but it is quite probable that it will be within 2-3 years, at which point smaller sites may consider IPv6-only networking. Some experimental results of small-scale IPv6-only networking have been published, e.g. [V6ONLY] and at RIPE61.¹⁷

Just as the planning for IPv6 can be phased, so can the deployment of services across a site. Provided the underlying base network services are IPv6 enabled, further IPv6 capability can be turned on as and when a site is ready. The discussion of the JANET IPv6 service enabling in the previous section illustrates this, as do the notes on the Southampton case study later in this text.

9.2 Phase 1: Advance planning

The first phase for deployment should begin immediately for any site that has not yet considered IPv6 deployment. It is a preparation step that includes the following actions:

1. Establish IPv6 training for both management and operational staff. This is necessary to understand both the strategic aspects and implications of IPv6, to be able to write appropriate project initiation documents/cases, and to gain technical (configuration)

¹⁷ http://ripe61.ripe.net/presentations/140-ripe_rome_jari.pdf

and standards (protocol) familiarity. JANET currently offers an IPv6 Fundamentals course¹⁸ with a more advanced hands-on course being planned during 2011.

2. Include IPv6 requirements in all future tenders. This is an important task but not necessarily easy to do without appropriate templates or protocol knowledge, e.g. understanding that layer 2 switches should support MLDv1 and MLDv2 snooping for multicast. Some sample procurement guidance has been gathered into [RIPE501].
3. Survey existing software and systems for IPv6 capability. All common current OS and router platforms now have good IPv6 support – the ‘problems’ are more likely to lie in commercial applications, services and any bespoke software your site may be running. The survey should include all network services including management and monitoring tools, which should also be examined to see that where they do support IPv6 they preferably provide integrated management with IPv4.
4. Some of your tools/ software may not yet support IPv6 so you should assess the effort to port software where required. You may also view this as an opportunity to change to different software that has better or more cleanly integrated IPv6 support. Note that not all software may need to be IPv6 capable from Day One of deployment.
5. Review IPv6 security issues. IPv6 is enabled by default on many host platforms; this should be considered when enforcing security policies on systems and networks. You should be confident that security, management and accountability functions will be available as required for IPv6 systems.
6. Speak to the JANET(UK) Service Desk to acquire IPv6 address space (a /48 prefix) for your site. This will be needed at some point so the space may as well be acquired sooner rather than later. This will include delegation of IPv6 forward and reverse DNS for your site.

We discuss some of these issues in the next couple of subsections. There is also a good set of material about preparing for IPv6 deployment in RIPE NCC’s IPv6 Act Now web site.¹⁹

9.2.1 Assessing IPv6 capability

Modern OSes – for example Windows 7, Linux, Mac OS X v10.6, Solaris, BSD – all have good IPv6 support and most have IPv6 enabled by default. Windows XP includes IPv6 support but it is not as feature-rich as Windows 7 and is not enabled by default.

The major networking equipment vendors now support IPv6 in their equipment, although some require an extra licence to be purchased to enable its use. Support in firewalls and other security appliances has lagged behind a little, although most now provide IPv6 support in recent software releases.

The IETF has provided updated specifications for IPv6 support in common routing protocols, including RIPng [RFC2080], OSPFv3 [RFC2740], IS-IS [RFC5308] and BGP4+ [RFC2858].

The level of support in particular releases or by different vendors may vary, so it is prudent to check carefully that the software you intend to use supports IPv6 in all required protocols, be that tunnelling or running BGP.

As an example of configuring an IPv6-in-IPv4 tunnel end-point on a Cisco® router:

```
interface Tunnel100
  ipv6 address 2001:630:d0:8000::2/112
  tunnel source 152.78.63.241
  tunnel destination 152.78.108.2
  tunnel mode ipv6ip
```

The tunnel can then be viewed with ‘show ipv6 interfaces’.

An example basic IPv6 BGP (Border Gateway Protocol) peering in IOS might look like this:

¹⁸ <http://www.ja.net/services/training/courses/ipv6.html>
¹⁹ <http://www.ipv6actnow.org/>


```
router bgp 65001
  no bgp default ipv4-unicast
  neighbor 2001:630:d0:8000::1 remote-as 65010
  address-family ipv6 unicast
  neighbor 2001:630:d0::1 activate
```

and the peering can be inspected by 'show bgp ipv6 unicast summary'.

Ethernet switches 'support' IPv6 by default in as much as the switch only deals in Ethernet frames, but you should be aware of other desirable layer 2 features such as MLD or RA snooping (RA Guard), and the ability to manage such devices over IPv6 transport.

Gaps in IPv6 support in networking equipment tend to lie in areas such as load-balancers or resilience/failover protocols; support for these should be checked carefully.

Support for IPv6 tends to be less common at the application level; while many network infrastructure software implementations are IPv6-capable (for example DNS and e-mail servers) other software such as videoconferencing and virtual learning environments are less likely to support IPv6 at the present time. Where an essential part of IT infrastructure does not support IPv6, there are methods that can be used to translate between the IPv4 and IPv6 worlds (see the transition section).

9.2.2 Address planning

As with IPv4, some form of addressing plan is necessary. For some smaller sites this might be as simple as selecting a small number of /64 prefixes to use from their IPv6 address space assignment. For large sites that already have a large-scale routed IPv4 network, a little more consideration is needed.

The most obvious strategy is to create an addressing plan based on the existing IPv4 network, with the only difference being that the equivalent IPv6 prefix length will be /64, regardless of the size of the IPv4 subnet.

For some sites, the opportunity to start from scratch with IPv6 may offer the opportunity to make a more logical addressing structure, particularly given the size of address space available their IPv6 /48. For example, remote campuses, buildings or departments could be assigned contiguous blocks of /64 prefixes over and above the current requirement.

This means that the location of any IPv6 address is easily visible, as it is within the block for (for example) the building it is in. In many cases this is not so easy for IPv4, as addressing plans have evolved in a relatively anarchic fashion, where adjacent subnets may well be on the other side of campus, city or even region.

In practice though, most sites will likely build their IPv6 address plan around their existing IPv4 plan – the reasons for subnet boundaries are typically geographic or administrative, and those reasons will remain the same.

Currently there are very few, if any, campus sites that are using IPv6 ULAs (Unique Local Addresses) [RFC4193] – their theoretical benefit is stable addressing through a renumbering or similar event, but in practice campuses have stable long-term address allocations from JANET(UK).

If a campus is planning to multi-home through one or more local RNOs/ISPs, it may be worth the campus exploring the availability of PI address space, as available through the new RIPE policy and request form.²⁰ However, for the vast majority of sites, resilience is now factored into the JANET backbone and RNO networks, so this is unlikely to be necessary and thus JANET-assigned address space should be sufficient.

There is a good overview of IPv6 address planning considerations in [RFC5375], and SURFnet has published an address planning guide²¹ in conjunction with RIPE NCC.

²⁰ <http://www.ripe.net/ripe/docs/ripe-467>

²¹ http://www.ripe.net/training/material/IPv6-for-LIRs-Training-Course/IPv6_addr_plan4.pdf

9.3 Phase 2: Testbed/Trial deployment

In this stage a site is validating IPv6 for deployment and will want to build an appropriate testbed to evaluate the new technology. Actions to take at this stage include the following:

1. Establish IPv6 connectivity to JANET. Sites can use the tunnelled JANET Experimental Service for a testbed, or contact their RNO to check for any native IPv6 offering. The current RNO SLAs require that the regional networks supply IPv6 on demand, but a tunnel through a site edge router may be preferable during the testbed phase, as this would not require IPv6 to be enabled on the edge router.
2. Assign and deploy an IPv6-capable router for your testbed, and (strongly recommended) a firewall or filtering device, e.g. BSD with pf, or Cisco® IOS IPv6 ACLs.
3. Connect testbed host systems on the internal router interface, using at least one subnet prefix (a /64) from the site's allocated IPv6 /48 prefix. The connected subnet(s) should be dual-stack if that is the intended production deployment mode, and could if preferred be a dual-stack DMZ that has one interface in the site IPv4 DMZ and one in the site IPv6 testbed DMZ, as discussed below.
4. Enable IPv6 on the host systems, and test IPv6 functions on services and applications (e.g. BIND for DNS, Apache for Web, sendmail or exim for mail transport).
5. Gain experience in network operation, management, monitoring and applications / service porting.
6. Discuss IPv6 issues arising – technical, policy or otherwise.

Once the site is satisfied in the testbed deployment, a production deployment can be considered, enabling IPv6 for appropriate services.

9.3.1 Testbed considerations

Establishing an IPv6 test facility gives staff the opportunity to become familiar with IPv6 before embarking upon the rollout of the technology across an organisational network. The facility might also be used to provide familiarity to interested parties across the organisation, e.g. researchers in a computer science department.

The test environment does not necessarily need to connect to JANET immediately; it might be that the network is put together and enabled internally first, building up to connecting to the outside world. The address space assigned by JANET(UK) to the site can be used for the testbed, whether it is connected externally or not.

There are many ways a testbed could be set up on a campus. If a specific research group is interested in IPv6, a tunnel directly from that group's labs to the JANET Experimental Service tunnel server may be sufficient. However, for a more formal testbed activity by the campus computing service department, a topology like Figure 9.1 might be appropriate. Here, a dual-stack testbed is being deployed, including a dual-stack router from which an IPv6-in-IPv4 tunnel is created to the JANET(UK) tunnel server, and internally a dual-stack firewall protects the testbed systems.

It must be stressed that once a site is connected by IPv6 to the outside world IPv6 provides a new potential attack route into the network, so IPv6 security policies should be implemented in advance of any deployment that may connect critical systems.

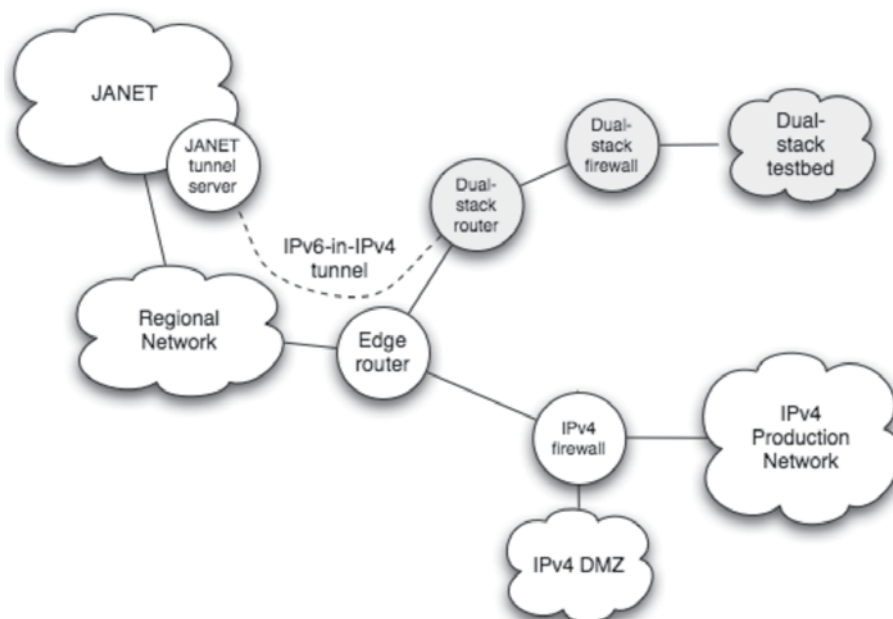


Figure 9.1: One possible testbed topology

An alternative testbed topology is shown in Figure 9.2. This is more adventurous as it exposes IPv6 routing and RAs to IPv4 systems in a common dual-stack DMZ. This may be similar to the topology for production dual-stack DMZ services, and requires IPv6 enabled on the edge router. In this scenario the site also has native IPv6 connectivity through its RNO to JANET, so is routing IPv6 through the IPv6 firewall to the dual-stack DMZ. By using a separate IPv6 firewall, filtering only IPv6 traffic, testing of IPv6 filtering can be done without impacting the production IPv4 firewall(s), but any potential compromise in the DMZ is still mitigated by the IPv4 firewall protecting the internal IPv4 network. However, a dual-stack firewall may be used instead if desired.

If possible, the test facility should be constructed with sufficient equipment to provide a realistic environment to test all features that will need to be deployed. For example, a single router is not enough equipment to simulate the configuration and operation of any routing protocols that will be necessary on the live network.

It is recommended that the pilot deployment is constrained to specific subnets and physical infrastructure, both for security purposes and to avoid accidental exposure of IPv6 to unintended systems.

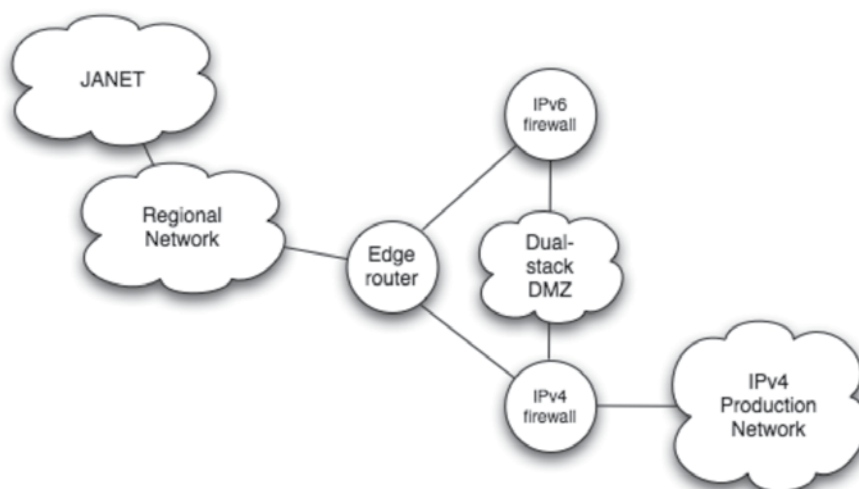


Figure 9.2: A potential topology for a dual-stack DMZ pilot

The topology above, possibly with a dual-stack firewall, would allow a site to expose some services externally via IPv6 before any internal services need to be enabled. This for many sites may be a useful stepping stone to fuller campus deployment.

9.4 Phase 3: Production deployment

A production deployment includes the following action points:

1. Plan which parts of the network and which existing subnets will be IPv6-enabled (made dual-stack). The latter may be certain server subnets, a DMZ, or a WLAN network, for example.
2. Determine how your production IPv6 connectivity will be handled; it can ideally be via a single dual-stack entry point, or through separate IPv4 and IPv6 links.
3. Put IPv6 capability in the network first. Enable IPv6 on the wire in routing protocols and appropriate security components (firewalls, IDS, etc); ensure IPv4 equivalent security is in place (common policy implementation). Do not turn on RAs on user subnets yet.
4. Add IPv6 addresses to your DNS systems and configure them to respond over IPv4 or IPv6 transport.
5. Deploy IPv6 support in appropriate management and monitoring tools.
6. Enable IPv6 in selected production services and applications (e.g. Apache or IIS for web servers), and turn on RAs on those server subnets (servers should have manually configured addresses, but need RAs to learn their default router).
7. Enable IPv6 in selected client subnets (turn on RAs) – enabling a small number of ‘friendly’ subnets first, e.g. computer science research locations.
8. Monitor resulting IPv6 usage and performance, adding more client subnets as you go.
9. Include IPv6 transport in all ongoing security audit/penetration tests.

The depth of the IPv6 deployment may vary from site to site. There is no requirement to enable all services or client subnets from Day One; the pace and scope of the deployment is for the site to decide.

9.4.1 Preparing the IPv6 Network

If the connection to JANET was not enabled as part of the pre-deployment work, now is also the time to configure and secure that connection. As previously discussed, this connection may be a dual-stack direct connection, or if that is not possible a tunnelled connection to a router elsewhere on JANET.

The process above recommends not enabling end-user subnets until the network and its services are ready and tested. Enabling the use of IPv6 can be performed on a network-by-network basis, perhaps starting with the network used by ICT support staff, or by ‘friendly’ and interested computer science department users. This should make any problems with basic IPv6 connectivity, either internal or to external IPv6 available services, evident only to the ICT staff and therefore not affect the network service or application services delivered to end-users.

An appropriate topology might then become as shown in Figure 9.3.

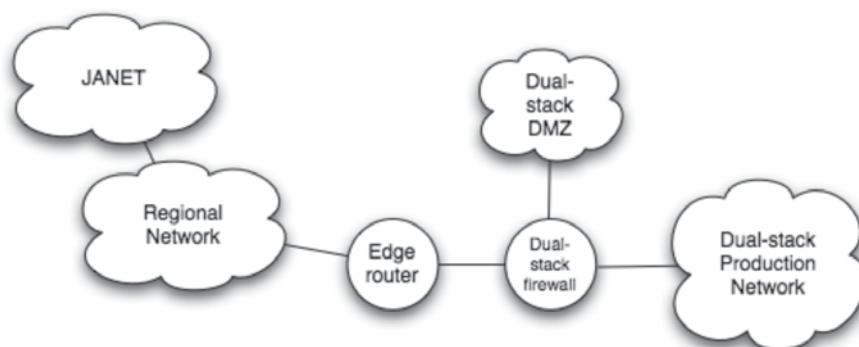


Figure 9.3: A dual-stack site topology

A full deployment would involve making all key production services dual-stack where possible – for example, e-mail, web servers and DNS servers. Bear in mind that enabling a service for IPv6 may involve separate configuration of any host or network based firewalls or other security device, e.g. where the service resides in a DMZ network.

If all services currently available via IPv4 are included in the IPv6 deployment, this theoretically allows the use of IPv6-only devices on the organisational network. However, without the use of additional transition mechanisms such as NAT64 or ALGs any such devices would not be able to communicate with IPv4-only services elsewhere.

We now discuss some aspects of typical services that might be made dual-stack and exposed externally.

9.4.2 DNS

Once a service has been configured (and tested) to accept IPv6 connections, the DNS entry for the server/service needs to be adjusted to include a record for its IPv6 address. The DNS uses AAAA (or ‘quad-A’) records to denote IPv6 addresses, in exactly the same way as A records are used for IPv4 addresses. A host should not be DNS-advertised for IPv6 until all services on it reachable by that domain name are IPv6 ready.

A popular DNS platform is BIND, which only needs a minor change to its configuration file to listen on IPv6:

```
listen-on-v6 { any; }
```

In addition you may need to add ‘transfer-source-v6 *’ to specify IPv6 source address for transfers and ‘query-source-v6 address * port *’ to specify IPv6 source for queries, and remember that IPv6 addresses can be used in ACLs, e.g. to restrict zone transfers

Adding IPv6 records to the DNS is then simple; where A records are added for IPv4, AAAA records are added for IPv6., e.g.

```
; Dual-stack host
websites1 IN A 152.78.189.43
          IN AAAA 2001:630:d0:f104::80e
```

The reverse DNS for IPv6 uses records under ip6.arpa (formerly ip6.int) and works using PTR records just like IPv4. Sites should get both forward and reverse DNS delegations from the JANET Service Desk, e.g. a delegation may look like:

```
0.63.78.152.in-addr.arpa IN NS ns0.ecs.soton.ac.uk.
7.0.d.0.0.0.3.6.0.1.0.0.2.ip6.arpa IN NS ns0.ecs.soton.ac.uk.
```

A DNS forward entry of

```
ipv6lab.ecs.soton.ac.uk AAAA 2001:630:d0:7000::9:2
```

has the corresponding reverse entry

9.5 Phase 4: Ongoing operation

Once the initial IPv6 services are running (dual-stack) in your production environment, the process of monitoring and reviewing the ongoing deployment begins.

It is likely that the initial deployment does not provide pervasive dual-stack services on all systems and to all subnets. Thus a running plan should be maintained for how and when those deployment 'gaps' will be filled. Some may depend on releases from router or firewall vendors for example. Others may be more political in nature. The service delivery elements will however in general be the same as they were for IPv4, and IPv6 should not be treated differently in that respect – it is just another service.

The other aspect to ongoing operation is seeking to exploit new opportunities presented by IPv6 being deployed, e.g.

- offering new services to users who now have global IPv6 addresses where previously they had only IPv4 NAT
- IPv6 multicast – perhaps adding new streaming services and ad-hoc streamed user content using Embedded-RP
- mobile IPv6 – support for IP mobility on and between campus(es) – though note that currently implementations of MIPv6 are not common and operational experience is limited
- IPv6-only device applications – ensuring these can operate within the campus (e.g. that IPv6 transport DNS is available) and considering options to allow them to access external IPv4-only content, e.g. running a NAT64 trial, or configured ALGs for specific applications
- testing new applications that benefit from IPv6, e.g. MS DirectAccess
- following IPv6 developments from Google, Facebook, etc.

Such opportunities should be tracked, and ongoing plans adjusted accordingly.

10. Deployment of IPv6 on Regional Networks

Currently, the deployment of IPv6 on a regional network should be a dual-stack solution, such that the RNO offers IPv4 and IPv6 to an equivalent standard for any of its connected sites. Since the rollout of the present backbone, any RNO is expected to provide IPv6 to any site requesting it.

Any UK HE/FE site may still obtain IPv6 connectivity using the JANET Experimental Service [JANETEXP], but this will be a tunnelled connection through a regional network's IPv4 infrastructure and should be considered a last resort connectivity mechanism, or limited to testbed use.

Ideally a site's IPv6 connectivity through a regional network should be native; to support this the regional network should deploy IPv6 natively itself.

10.1 Deployment Methodology

There is a growing set of resources available for RNOs to consider when deploying IPv6. For example, the 6NET reports on IPv6 transition [D224] for NRENs and the IETF's ISP transition analysis document [RFC4029] are both relevant reading.

Ideally, one would look for the regional networks to deploy IPv6 natively in dual-stack mode, just as has been done on the JANET core network. However, if the regional network already has MPLS (Multi-Protocol Label Switching) deployed then 6PE [RFC4798] is a viable interim solution, as adopted on the LeNSE regional network. However, MPLS adds some extra complexity to a network and thus its deployment *purely* as an IPv6 enabler is not recommended.

10.2 Choice of Routing Protocols

The choice of interior routing protocol is relatively open. RIPng is unlikely to be suitable for use on regional networks, so one would expect either of the link state OSPFv3 [RFC2740] or IS-IS for IPv6 [RFC5308] protocols to be used.

As described in [RFC4029] a RNO generally has the choice of running either:

- Separate routing processes (often referred to as 'ships in the night'):
 - OSPFv2 for IPv4, IS-IS for IPv6 (only)
 - OSPFv2 for IPv4, OSPFv3 for IPv6, or
 - IS-IS for IPv4, OSPFv3 for IPv6.
- One routing process:
 - IS-IS for both IPv4 and IPv6.

It is widely considered that single topology IS-IS (a common IS-IS instance for both IPv4 and IPv6) is the preferred solution for running dual-stack networking. While running separate processes may mitigate some risk, the single process approach is considered the more efficient. It is usually not possible to have separate IS-IS instances for v4 and v6, but Cisco® also has multi-topology IS-IS that allows this.

This implies that networks running OSPFv2 should migrate to IS-IS for IPv4 as well – a decision that is entirely up to the regional network itself. Some networks have converted to IS-IS to have a single routing process, which uses less resource and is easier to manage. This is also a reason why the JANET core switched from OSPF to IS-IS. However, one could use OSPFv2 for IPv4 and OSPFv3 for IPv6 if necessary.

When considering the network topologies for IPv4 and IPv6, the best approach is to seek to have the same topologies for both IPv4 and IPv6. It is particularly important when using the

IS-IS routing protocol for both IPv4 and IPv6 to have the topologies 'convex', as described in Appendix A of [RFC4029].

A regional network will use BGP for external routing and would thus use BGP4+ for IPv4 and IPv6 [RFC2858].

10.3 Address Space and DNS

In terms of IPv6 address space, a regional network can expect to obtain a network prefix from the JANET Service Desk just as any end site would do, and thus receive a /48 size prefix for numbering its own network devices and systems. Campus sites served by the regional network would be expected to get their address space directly from JANET.

The regional network should also ideally handle DNS for the allocated prefix, both forward and reverse, as the regional networks do today for IPv4. The DNS server(s) used should support IPv6 transport. This makes addressing the network elements by name easier, but also helps end sites by traceroute reports displaying meaningful names on the hop reports.

10.3.1 Point-to-point links

Regarding prefix lengths to be used on a link, it is quite possible to use a /127, as described in [P2PLINKS], despite concerns over Subnet-Router anycast addresses [RFC4291].

Many ISPs are using /64 prefixes for point-to-point links, which is also legitimate, but in doing so you should be aware of the concerns in [P2PLINKS], in particular regarding the potential for DoS attacks on ND caches.

The author is aware of ISPs who also use /112 and /126 for point-to-point links.

10.4 Other Services

It would not be expected that a regional network would run a 6to4 relay or tunnel broker, since the JANET services should – at least at the current load levels of transition – be ample for users/sites in all connected regional networks. Adding a local 6to4 relay would improve connectivity for local native sites to remote 6to4 sites.

11. IPv6 Security

IPv6 is often perceived as being more secure than IPv4, possibly as a result of some ‘bold’ claims in the early days of IPv6 development.

In reality, network security is a function of the security policies in effect for a network and the way they are applied. Firewalls, packet filtering and other basic security devices and configuration are required for IPv6 just as they are for IPv4. A site will need to determine appropriate policies and ensure they are applied consistently for both protocols.

Some features of IPv6 may make a network harder to be attacked; however the security of an IPv6 network should be treated in exactly the same way as an existing IPv4 network. This is important to bear in mind when deploying IPv6, particularly in a dual-stack fashion – a drop in the security practice for IPv6 might permit unauthorised access to the IPv4 network via dual-stacked host.

Thus, where possible, a site should determine and apply security policies equally for IPv4 and IPv6. But it is also important to understand new IPv6 security threats – while many have IPv4 equivalents, e.g. ND spoofing rather than ARP spoofing, other threats are new, like Rogue RAs [RFC6104].

All common modern operating systems ship with IPv6 support and typically with IPv6 enabled by default. Thus if you don’t manage IPv6 – even if you just disable it as an interim measure – you may find your users abusing it, unintentionally or deliberately.

11.1 IPv4 Security Equivalence

Security considerations should be standard in all your activities. You should seek to implement common measures for whichever area is required, e.g.

- firewall/router ACLs – including host-based firewalls
- intrusion detection – whether the traffic is IPv4 or IPv6, it should be possible to check payloads (likely to be common to both protocols, e.g. http attacks) and protocol-specific issues (e.g. IPv6 Routing Header DoS attacks).
- anti-spam/virus – IPv4 or IPv6 transport emails, with IPv6 RBLs when they become available
- network authentication – 802.1X is a fundamental part of eduroam²² but also excellent for dual-stack environments as it is fundamentally IP version agnostic.

When security policies are documented, consideration for how they can be implemented consistently for IPv4 and IPv6 needs to follow. This becomes simpler if the network and system components allow objects to be managed and configured consistently (and intuitively) for both protocols.

11.2 New IPv6 Security Issues

IPv6 introduces many new security concerns. These include:

- General security considerations of various IPv6 transition tools (see [RFC4942]); these relate to the more obvious potential for IP tunnels to bypass security measures, through to potential DoS attacks for certain transition methods. In general a native, dual-stack deployment mitigates these concerns by avoiding support for tunnelling tools.
- There may be more subtle use of IPv6 encapsulation in your network, e.g. possible use of IPv6-in-UDP tunnelling using Teredo [RFC4380]. This can be spotted if Teredo is using the default UDP port of 3544 for hosts contacting Teredo servers.

²² <http://www.eduroam.org>

- Ensuring users can be accounted to hosts where hosts may be multiply addressed and/or using IPv6 Privacy Addresses. Many network monitoring packages may assume such addresses belong to different hosts.
- Accidental or malicious rogue Router Advertisements (see [RFC6104]).
- Possible DoS attacks on ND caches, where the (effectively) unlimited IPv6 subnets facilitate the possible exhaustion of ND cache capacity if a large-scale host scan is performed on a given subnet. This is implementation dependent, e.g. routers should not expire active neighbours in favour of incomplete ND processing.

There are some hacker toolkits publicly available for new vulnerabilities, e.g. Duplicate Address Detection (DAD) denial of service. The most well known package is the THC Toolkit,²³ which you can download and test against your own network.

Some 'defensive' packages have been released. For example, NDPmon²⁴ allows detection of many of the attacks included in the THC toolkit. The ramond²⁵ package allows detection of and deprecation of rogue RAs, though a more proper solution to rogue RAs that switch vendors should be implementing soon is defined in [RFC6105].

It is important to gain familiarity with these types of new threat, e.g. by deploying a small IPv6 security testbed at an early stage.

11.3 IPv6 firewalls

The general principles for firewall operation are the same for IPv4 and IPv6. Ideally the firewall configuration interface will allow handling of dual-stack objects in an efficient, intuitive way.

An IPv6 firewall may need to route IPv6 multicast, just as it might handle IPv4 multicast, though it is not uncommon for sites to route multicast around firewalls. Where a DMZ is used, the IPv6 firewall should be able to issue RAs on the DMZ interface.

Host-based firewalls such as iptables for Linux and pf for BSD either integrate or have parallel tools to provide IPv6 security features to match those provided for IPv4. Windows 7 supports IPv6 filtering and MacOS X has the ip6fw utility.

IPv6 support is included in many firewall platforms. For example, Cisco® IOS has access lists that use the same principle as the IPv4 filters, e.g.

```
ipv6 access-list-name permit tcp 2001:0db8:0300:0201::/64 eq 22
```

The traffic-filter command can also be used then to apply named rules inbound or outbound, e.g.

```
interface ethernet 0
ipv6 traffic-filter access-list-name in
```

The BSD pf tool can also be used, e.g.

```
pass out quick on $if proto tcp from any to 2001:db8::22 port ssh keep state
```

Linux ip6tables has identical syntax as iptables for IPv4, e.g.

```
ip6tables -A FORWARD -d 2001:db8::22 -p tcp --dport 22 -i eth0 -j ACCEPT
```

11.4 Network port scanning

As a typical IPv6 network will have a /64 prefix, traditional port-scanning attacks should be far less practical as they are more time consuming. While the targeted address space may be whittled down via various strategies (e.g. assuming the middle bytes are fffe from SLAAC), a full scan of 2⁶⁴ addresses at just one node per second would take over 500 billion years to

²³ <http://www.thc.org/thc-ipv6>

²⁴ <http://ndpmon.sourceforge.net>

²⁵ <http://ramond.sourceforge.net>

complete. At 2000 probes per second a complete scan would still take 250 million years (see [RFC5157]).

The experience at the author's site is that some port scanning is seen to certain ports on hosts that are DNS-advertised, such as web or MX servers, but there is very little sign of any classic port sweep scanning.

11.5 IP Security (IPsec), CGAs and SeND

The original IPv6 specification mandated support for authentication (via the Authentication Header) and encryption (via the ESP header). This did not mean that IPsec must be used; rather it meant that a full implementation of IPv6 must support the use of the Authentication Header and ESP IPv6 headers.

This requirement was echoed in Section 8.2 of [RFC4294] which describes IPv6 node requirements. However, the requirement has subsequently been relaxed to a 'SHOULD be supported' in Section 11.1 of the updated version of the Node Requirements RFC that is close to publication at the time of writing [NODE-BIS].

In practice only some IPv6 implementations currently support Authentication and ESP headers, and of those some support only the 'null' ESP header.

IPv6, through the size of an address, enables some new security methods. In general, ND is vulnerable to attack and addresses may be spoofed. However, by using Cryptographically Generated Addresses (CGAs), as defined in [RFC3972], the ownership of an address can be assured by use of a hash of a node's public key being used within the 128-bit address.

The SEcure Neighbour Discovery (SEND) protocol (RFC 3971) uses CGAs to secure ND messaging in a network, without the need for IPsec.

There are at the time of writing very few public implementations of SEND, but in principle it offers improved security where it is deemed required (though one could argue Authenticated DHCP has been available for many years, but the compromise for DHCP security is just to use DHCP snooping on switch ports).

11.6 Summary

IPv6 security must be considered because all hosts ship with IPv6 support and usually with IPv6 enabled by default.

Security policies should be applied equally for both protocols – don't leave gaps.

Be aware of new security issues, e.g. in transition tools or with rogue RAs.

Track new IPv6 security measures, e.g. CGAs and SEND.

12. Advanced IPv6 Network Services

It is very likely that sites and RNOs at some point will want to implement IPv6 support for both QoS and Multicast, if they already do so for IPv4.

Mobile IPv6 is a new IPv6-specific protocol for supporting host mobility in IPv6 networks. Implementations are not as yet common; support by sites for MIPv6 would require provision of Home Agent(s) and bootstrapping methods (e.g. for assigning IPv6 Home Addresses and host to Home Agent IPsec relationships), as well as ensuring appropriate capability exists in firewalls to allow appropriate MIPv6 traffic to pass.

12.1 IPv6 Multicast

A separate JANET guide is available that documents IPv6 Multicast,²⁶ but we describe some of the basics of the protocol and differences to IPv4 here.

The primary difference between IPv4 and IPv6 Multicast is that IPv6 Multicast does not have an equivalent of MSDP. Instead, the Embedded RP [RFC3956] protocol provides a mechanism for multicast routers to deduce implicitly the address of the RP for the destination group.

It is expected that inter-site multicast in IPv6 will either use Embedded-RP or SSM.

Another important difference is that IPv6 Multicast has explicit scope; the scope identifier (5 bits) is included in the group address used, from 0 through 15, with 5 being reserved for site scope, 8 for organisation and e for global. This means scope boundaries can be easily defined (no need to use TTLs, or rely on specific group address ranges being used as is the case with IPv4).

Further, IPv6 multicast group addresses can be generated based on the unicast prefix for a site [RFC3306], so it is very easy to obtain globally unique IPv6 group addresses to use, even if that group is limited to having only one RP globally.

12.1.1 Embedded RP

The traditional multicast service model allows for applications to join a group (*,G) in order to receive any data sent to the group. A source simply sends IP datagrams to the multicast IP address of G. An application does not need to know which sources are available. This model is called ASM (Any-Source Multicast).

The most popular multicast routing protocol is PIM-SM. When a source starts sending to the group, the packets are sent from a router on the source's local network. Initially the multicast packets are actually encapsulated into unicast packets and sent as PIM register messages from the router next to the source to the Rendezvous Point. At the same time PIM JOIN messages reach the RP from the receiver, and thus the RP knows which interfaces to use to reach interested receivers.

In this way the packets are forwarded from the source to the receivers via the Rendezvous Point. However, once the multicast packets reach the last-hop router (the last router before a receiving host), the last-hop router can shortcut this by building a shortest-path tree towards the source. The last-hop router learns the source address by looking at the source address of the multicast packets it receives.

One problem is all the routers in a network knowing where the Rendezvous Point is for a group. This can be manually configured. Another solution is to use BSR for Rendezvous Point-to-group mappings, but this is very hard to do throughout the Internet. For IPv6, Embedded-RP allows the IP address of the RP to be embedded in the multicast group address, which can be achieved given a suitably constrained RP address. With Embedded-RP, a multicast router knows the RP to use by inspecting the group address.

²⁶ <http://www.ja.net/development/network-engineering/multicast/>

In IPv4, RPs in different multicast domains can exchange information about sources using MSDP and Source Active (SA) messages. In IPv6, there is no MSDP, so any ASM group must use Embedded-RP, and only one RP per group. The attraction of this method is that small, ad-hoc multicast group networks can be set up given all the multicast routers on the path between them support Embedded-RP. The RP may typically be close to the multicast source, but it need not be. Embedded-RP is supported in most academic networks; the author has tested it successfully between his own university and universities in the US on Internet 2.

The downside of Embedded-RP is that there is only one RP per group; the operation of the group depends on that RP being available.

12.1.2 IPv6 SSM

In contrast, the SSM service model is based on applications joining so-called channels rather than groups. A channel is a pair (S,G) consisting of a unicast source address S and a multicast group address G. For each source S and group G from which the application is to receive data, it must join the channel (S,G). This means that the application explicitly specifies the source and hence it must know the source address in advance.

PIM-SM can also provide an SSM service. When a host joins a channel (S,G) the last-hop router immediately knows the source address S, so it can immediately start building the shortest-path tree towards the source without waiting for multicast packets. In order to provide SSM service only, no RPs are necessary. This basically removes all the complexities. There is no need for multicast protocols such as BSR (Bootstrap Router) and MSDP to be used, and no shared trees, no PIM register with encapsulation and decapsulation. This means that an SSM-only network is much easier to manage.

In IPv6 the reserved address range for SSM is **ff3x::/32**.

12.1.3 MLD

When using PIM-SM, SSM and ASM can co-exist in the same network. Operation of PIM-SM and PIM-SSM requires that hosts support MLD (Multicast Listener Discovery) and MLDv2 respectively. For IPv4 and IPv6, there are protocols IGMP and MLD respectively that are used between routers and hosts for signalling which groups the host wants to receive data for. The latest versions of these protocols, IGMPv3 and MLDv2 [RFC3810] also let the host specify specific sources. In addition to SSM, they also allow a host to join a group as in ASM, but blocking certain sources.

For just sending SSM multicast the above protocols are not needed. A host does not in any way need to join and receive multicast in order to send.

At the time of writing, Mac OS X does not yet support MLDv2 for SSM.

12.1.4 IPv6 Multicast on JANET

Ideally, IPv6 Multicast would be supported natively between sites. The JANET core supports IPv6 Multicast, but for end-to-end native multicast the regional networks will ultimately need to support it natively on their networks.

As with unicast, the JANET Experimental Service includes a facility for a site to set up a (tunnelled) IPv6 Multicast connection.

12.2 IPv6 QoS

There are two broad flavours of Quality of Service (QoS) – DiffServ (Differentiated Services) and IntServ (Integrated Services).

12.2.1 DiffServ

Both IPv4 and IPv6 have a common 8-bit field in their header for the Type of Service or Traffic Class, and can use the same 6-bit DSCP (Differentiated Service Code Point) values to indicate how marked traffic should be handled. While these bits are in different locations in the header, the treatment of IPv4 and IPv6 traffic can in principle be the same for DiffServ.

Thus it should be possible for DiffServ handling to be consistent for IPv4 and IPv6 traffic in dual-stack networks.

12.2.2 IntServ

In support of IntServ, IPv6 introduces the 20-bit Flow Label field (which was 24 bits in the very first IPv6 specification). The usage of the Flow Label for IPv6 is described in [RFC3697].

At present, the default behaviour for nodes not using the Flow Label is to set it to zero. Otherwise a unique 'random' value should be used to mark all common flows. There is no widespread usage of the Flow Label field at the time of writing.

12.3 IPv6 Multihoming

RIPE-NCC has a policy by which Provider Independent (PI) address blocks (**/48**) can be obtained by end sites under certain conditions. This is described in Section 8 of RIPE-512.²⁷

In principle such PI blocks could be used by JANET sites to multihome, but in practice the JANET network now has resilience designed in, so sites can use JANET Provider Aggregatable (PA) address space allocated to them under the JANET prefix **2001:630::/32**.

The original intention with IPv6 was that sites using multihoming would use two PA blocks, and advertise both internally such that hosts become multi-addressed from both prefixes. In practice, few if any sites are using this approach and instead IPv6 PI is the likely method for site multihoming.

An alternative for multihoming is stateless prefix-based NAT [NAT66], which is supported on Cisco IOS. However, this approach is as yet somewhat untested, and many of the usual NAT limitations would still apply.

Other multihoming mechanisms that are designed for IPv6 or have IPv6 support include shim6 [RFC5533], HIP [RFC4423] and the Locator/Identifier Separation Protocol [LISP].

12.4 Mobile IPv6

Mobile IPv6 (MIPv6), as defined in [RFC3775] and currently under minor update in [3775BIS], is designed with two general goals in mind:

- to allow a node to be contacted via a static IPv6 (Home) address at all times
- to allow an IPv6 node that is moving between IPv6 links/subnets to maintain IP connectivity while doing so, allowing persistent applications to survive the underlying change of network.

The general principle of MIPv6 is that the Mobile Node registers with a Home Agent in its home network, and while in that network it behaves like any normal IPv6 node. As part of the initial bootstrapping the Home Agent assigns the node a Home Address, which it is reachable by at all times. An IPsec relationship is also typically established at this time between the Home Agent and the node.

When the Mobile Node moves to a different IPv6 network, it configures itself with a Care-of Address on that new network, and then sends a message, a Binding Update, back to its Home

²⁷ <http://www.ripe.net/ripe/docs/ripe-512>

Agent informing the Home Agent of its new Care-of Address. The Home Agent is then able to act as a proxy node on the home network for the Mobile Node, forwarding (tunnelling) traffic targeted to it from a remote Correspondent Node to the Mobile Node's Care-of address.

This process is illustrated in Figure 12.1.

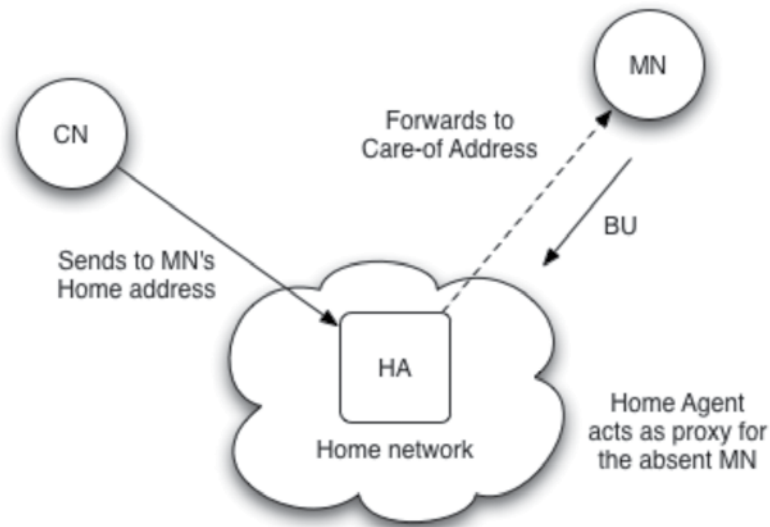


Figure 12.1: Mobile IPv6

Mobile IPv6 also supports a route optimisation mode whereby the Mobile Node may also send a Binding Update directly to the Correspondent Node, so that future communication need not take the 'triangular' route via the Home Agent. The latter property is very useful where two nodes are both mobile and in the same location away from their home networks.

This mode is shown in Figure 12.2.

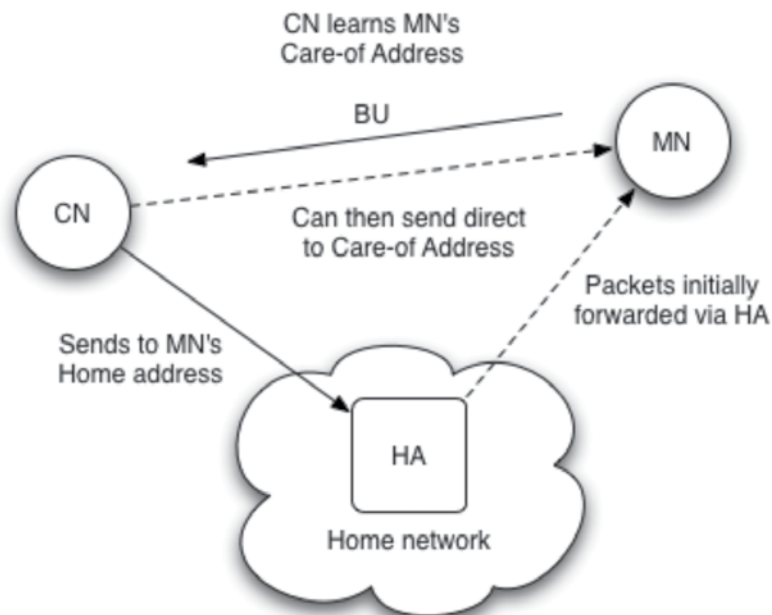


Figure 12.2: Mobile IPv6 with Route Optimisation

One of the major concerns with MIPv6 is security of the session, in particular that the Binding Updates can be trusted to be from the genuine Mobile Node.

As hinted above, IPSec may be used between the Mobile Node and its Home Agent.

The problem of securing the Binding Update from the Mobile Node to the Correspondent Node has been solved – after much discussion in the IETF – by the addition of the Return Routability test to the MIPv6 specification. This test essentially relies on secrets in messages sent both via the Home Agent and directly, to ensure a malicious node is not trying to spoof a Binding Update to hijack a session. Full details can be found in [RFC3775].

Another important consideration that arises from MIPv6 is how firewalls are configured to pass appropriate MIPv6 traffic, as described in [RFC4487]. An example is that firewalls need to understand the IPv6 Destinations Options header to support carrying of the Home Address when route optimisation is used (because the Mobile Node sends using its Care-of address in that mode to avoid packets being dropped by RPF checks at the visited site).

At the time of writing the main inhibitor to Mobile IPv6 deployment is the lack of implementations in commercial off-the-shelf devices.

13. Case Study: Enterprise IPv6 Deployment at Southampton

There are three case studies of IPv6 site deployment in 6NET deliverable [D234]: Tromsø, Lancaster and Southampton.

The Lancaster study has a heavy focus on address planning. The Tromsø case study is unusual in that it covers an IPv6-only deployment (not currently recommended for general deployments).

The deployment at Southampton has developed since the time of 6NET, and this section reflects that, but the general description and discussion in the 6NET text should still be relevant and of interest.

13.1 Scenario

The scenario is a large departmental network (1,500+ users with around 1,000 wired hosts and a greater number of wireless devices) that wishes to introduce IPv6 into the network so that IPv4 and IPv6 systems, services and applications can coexist. The network is that of the School of Electronics and Computer Science at the University of Southampton, where the School has deployed IPv6 in advance of a fuller campus service.

13.2 Dual-stack vs IPv6-only

It was decided from the outset that the deployment would be incremental dual-stack; i.e. at this time using IPv6-only networking was not practical. A dual-stack infrastructure can support IPv6-only devices internally, and offer those devices external Internet access via dual-stack proxies or a NAT64/DNS64 service. Such experience may be valuable towards a longer-term objective of IPv6-only operation.

13.3 Capability review

We reviewed all our systems and network components, which are cited in detail in [D234], and also an (expired) IETF Internet Draft on the topic of campus transition [CAMPUS]. This the School did in full for all its applications, services, network elements, etc. before embarking on the deployment procedure as described in Section 9 above.

The deployment at Southampton has spanned over ten years, but for the sake of this case study we consider the point from which we deployed new Cisco equipment in 2005. At the start of 2005 we wrote our requirements for a procurement process to acquire IPv6-capable router and switch equipment, as a result of which we deployed a Cisco solution involving 6509, 7206 and 3750 series components. Had it existed at the time, [RIPE501] would have been an excellent reference for writing our requirements. From 2005 onwards, all our procurements have had mandatory IPv6 requirements included.

As a computer science department we largely use open source applications and services, so for example ISC BIND and DHCP rather than commercial options, and Apache rather than other web server products. This meant the IPv6 support was good, though since the time of our deployment support in commercial platforms, particularly those from Microsoft, has improved significantly. Most of our administrative systems are delivered via web interfaces, rather than requiring custom client/server packages, which helped make IPv6 support relatively straightforward.

Current client host operating systems all have good IPv6 support, especially Windows 7 and Linux. While Mac OS X has IPv6 enabled by default, it does not as yet support some important IPv6 capabilities, including DHCPv6 client, Default Address Selection [RFC3484] and Source Specific Multicast (SSM), though MacOS X Lion may include the first two of these. At the time of

writing support in mobile devices is rapidly improving, e.g. Android 2.2 includes IPv6, as does the iPhone on IOS4 on its wireless LAN interface.

Our security review looked at various areas but identified quite quickly that our border firewall platform did not provide appropriate IPv6 capability, both in terms of general filtering but also in terms of support as an IPv6 router issuing IPv6 RAs on DMZ interface(s) and as a PIM-SM IPv6 router (with Embedded-RP support). As a result we decided that we would maintain independent IPv4 and IPv6 border firewalls. An open source solution was chosen, which had the added advantage of allowing research-oriented customisations to be made, e.g. to add support for passing certain MIPv6-related traffic.

13.4 Address Management

The department has the equivalent of approximately 15-20 (non-contiguous) /24 IPv4 prefixes allocated to it by the campus computing services department (iSolutions).

For IPv6, the University has been allocated **2001:630:d0::/48** by JANET. After discussion, the department has been allocated a /52 size prefix, **2001:630:d0:f000:/52**, by the University. Though a /56 allocation would have probably sufficed for production use, some additional address space was required for internal testbed purposes.

In the initial deployment, we decided that IPv4 and IPv6 subnets would be congruent and thus share and run over the same VLANs (virtual LANs). This is because the reasons for subnet division tend to be administrative or geographical, and thus IP version independent.

The processes for IPv4 address management, including for example those for pre-provisioning desktop systems, assume a DHCP service is available and can be linked to device MAC addresses. For IPv6, the School plans to use DHCPv6 when support is mature on core client platforms (support on Mac OS X is missing at the time of writing). Until then, the deployment uses SLAAC for clients and manually configured server addresses.

13.5 Network Elements

It was decided that enabling IPv6 'on the wire' throughout the department was the first objective, as far as the router element serving each link. Once that process was completed, and the necessary security components in place (including a border IPv6 firewall), subnets could then be activated on demand by turning on RAs on the desired links.

We also decided to put network monitoring systems in place in advance of turning on IPv6 on client subnets or IPv6 on application servers. Our goal was to be able to monitor and control IPv6 traffic on the network, then add the desired client and server capabilities.

The focus is thus to provide increasing IPv6 functionality in a dual-stack environment, with the implied goal of allowing IPv6-only devices to be introduced and to operate successfully using IPv6 transport. (That does not mean they have to interoperate with all 'legacy' services, but they should be able to use DNS, NTP, SMTP and similar basic services.)

The School's deployment notes for 6NET were drawn up while it had Alcatel network infrastructure. In the summer of 2005 it upgraded to new Cisco® equipment. Prior to that upgrade we ran a parallel IPv6 routing infrastructure, injecting IPv6 RAs into IPv4 VLANs from a hierarchy of open source BSD routers, as described in [RFC4554]. Such methods should no longer be required due to more mature router support, and thus the notes in this section refer only to deployment since we deployed the Cisco platforms.

Both IPv4 and IPv6 connectivity for the department come through JANET and the regional network (LeNSE) into the campus. IPv6 is delivered natively via JANET (dual-stack) and LeNSE (through use of 6PE) to the campus border router.

As discussed above, we run separate IPv4 and IPv6 firewall platforms. Thus from our campus border IPv6 traffic is carried via a dedicated VLAN to the School's IPv6 firewall (a Linux system

running iptables), while the IPv4 traffic is carried to the IPv4 firewall (a Checkpoint commercial solution).

The split handling of firewalling is not ideal – keeping the policies consistent requires care – but until an appropriate dual-stack firewall can be procured this interim solution is sufficient, as shown in Figure 13.1.

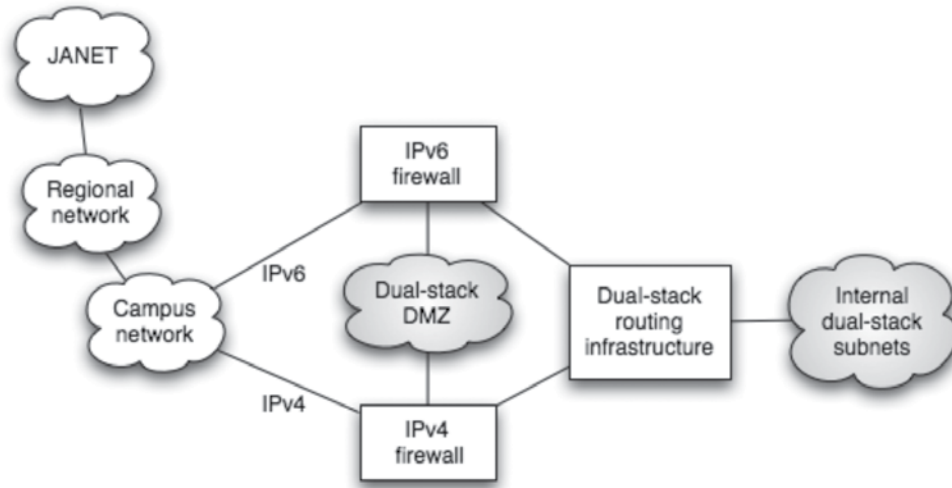


Figure 13.1: Dual-stack with separate firewalls

The order of deployment would thus be:

1. Enabling IPv6 in routing platforms;
2. Deploy and configure IPv6 security components (e.g. firewall, internal ACLs);
3. Enable IPv6 on appropriate server subnets;
4. Enable core IPv6 network services, including network monitoring and DNS transport;
5. Enable appropriate application servers, adding IPv6 DNS entries;
6. Enable client systems by turning on RAs on desired subnets.

It is important to stress that the last two points can certainly be done incrementally – there is no need to have a ‘flag day’ approach to dual-stack rollout.

13.6 IPv6 Services

The School’s three primary DNS servers running BIND9 have been enabled for IPv6 transport. This includes reverse delegation of the School’s prefix under **0.d.0.0.3.6.0.1.0.0.2.ip6.arpa**. The DNS server information is currently as follows:

```

ns0.ecs.soton.ac.uk.      1800   IN     A      152.78.68.137
ns0.ecs.soton.ac.uk.      1800   IN     AAAA   2001:630:d0:f110::53a
ns1.ecs.soton.ac.uk.      1800   IN     A      152.78.71.14
ns1.ecs.soton.ac.uk.      1800   IN     AAAA   2001:630:d0:f110::53b
ns2.ecs.soton.ac.uk.      1800   IN     A      152.78.71.210
ns2.ecs.soton.ac.uk.      1800   IN     AAAA   2001:630:d0:f110::53c
  
```

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

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

The School's SMTP and MX servers now exchange external and internal e-mail over IPv6. IPv6 DNS records were added for the hosts that provide these services. If the sending or receiving node the School is communicating with supports IPv6, IPv6 transport for e-mail is usually preferred.

Almost all of the School's web servers/sites have been made available using Apache 2 – for example the main School web site at <http://www.ecs.soton.ac.uk>, and many of the 200 or so hosted domains that the School runs.

The School's WLAN (over 30 access points) is IPv6 enabled. Some Mobile IPv6 has been deployed and tested between the WLAN and the local community wireless network (SOWN), using the MIPL code which is now integrated into the mainstream Linux kernel. The more advanced WLAN network the School has deployed uses eduroam with 802.1x based access control which is IP version neutral and thus can be used to secure the IPv4 and IPv6 WLAN access (in contrast to commercial web-redirection gateways that currently generally only support IPv4).

Monitoring is achieved by use of a variety of systems. These include Netflow v9, Nagios, the School's RIPE TTM server and a package called NAV²⁸ that was developed by UNINETT.

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

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

An H.323 IPv6 conferencing system has been tested (GnomeMeeting for Linux), and the School hosts a dual-stack Open H.323 MCU server for videoconferencing, which can interact with IPv4 or IPv6 end stations.

13.7 Monitoring and service examples

In this section we look at some example statistics and monitoring tool views in our network.

13.7.1 IPv6 external web traffic

We have seen only slow growth in our IPv6 web traffic. Figure 13.2 shows a very slow growth up until 2010. More recent stats show a peak at 1.8% of external visits coming over IPv6 transport, but the average remains closer to 1%.

²⁸ <http://metanav.uninett.no>

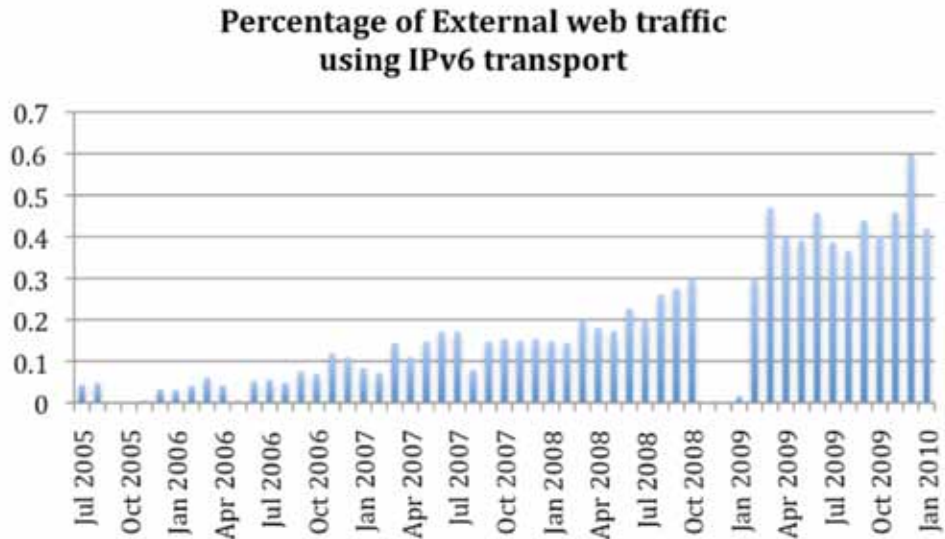


Figure 13.2: External web visits, 2005-2010

Of these visits, less than 1% of IPv6 accesses are via 6to4.

13.7.2 IPv6 email

We configured our MX DNS records as per recommendations in [RFC3974]. We receive 250,000 or so external mails per day over IPv4, as opposed to some 1,000 or so over IPv6. The spam level on IPv4 is currently about 88%, while it is 25% over IPv6. The spammers are clearly yet to discover IPv6.

13.7.3 Switch/router monitoring – NAV

The NAV package written by UNINETT does a good job of handling IPv4 and IPv6 information consistently.

NAV Network Administration Visualized

Home > Machine Tracker Preferences Toolbox Useradmin Userinfo Logout admin

Machine Tracker

IP Search **MAC Search** Switch Search

MAC dns days

MAC Search results
1 hits

Switch	Module	Interface	Start time	End time	MAC
b59-i2-cal1	2	FastEthernet2/0/48 (details)	2009-11-26 08:26	Still active	00:1a:a0:16:2a:9e

1 hits

Interface details

Please note that the MAC search results are historic data, while the information found at the interface details link is the current data on the interface. They may have no other connection than being related to the same interface at different points in time.

IP Search results
4 hits

DNS	IP	MAC	Start time	End time
crow.ecs.soton.ac.uk	152.78.71.14	<input checked="" type="checkbox"/> 00:1a:a0:16:2a:9e	2011-03-08 03:27	Still active
crow.ecs.soton.ac.uk	2001:630:d0:f110::25b	<input checked="" type="checkbox"/> 00:1a:a0:16:2a:9e	2011-03-08 02:57	Still active
ns1.ecs.soton.ac.uk	2001:630:d0:f110::53b	<input checked="" type="checkbox"/> 00:1a:a0:16:2a:9e	2011-03-08 02:27	Still active
--	2001:630:d0:f110:21a:a0ff:fe16:2a9e	<input checked="" type="checkbox"/> 00:1a:a0:16:2a:9e	2011-03-15 08:57	Still active

4 hits

Figure 13.3: NAV searching for host by MAC

In Figure 13.3, NAV uses MAC correlations to know that the same server has one global IPv4 address and two global IPv6 addresses.

NAV also understands dual-stack VLANs, as shown in Figure 13.4.

NAV Network Administration Visualized

Home > Report > Router port prefixes Preferences Toolbox Useradmin Userinfo Logout admin

Report

Advanced Search

Router port prefixes
Result generated 01:41:56
2 hits

Router	I/index	Router port	module	Mbps	hsrp	IP address	prefix	vlan	nettype	netident	Description
s50-i1-c6500	83	Vlan11	S	1000.0		152.78.85.254	152.78.84.0/23	11	lan		
s59-i1-c6500	83	Vlan11	S	1000.0		2001:630:d0:f110::2	2001:630:d0:f110::/64	11	lan		

2 hits

Figure 13.4: NAV, dual-stack VLAN

NAV is at the time of writing under active development.

13.7.4 Network flows

Netflow v9 includes IPv6 support. We configured our Cisco routers to export flow data to an open source collector running nfsen.²⁹ The examples in Figures 13.5 and 13.6 show queries for DNS flows (port 53) and individual port 25 SMTP connections.

²⁹ <http://nfsen.sourceforge.net>

The more complex issues included:

- ensuring we had the right management and monitoring tools to manage a dual-stack environment consistently (to feel like one network rather than two)
- IP address accountability without DHCP (which the administrators were familiar with from IPv4 practices)
- living with multi-addressed hosts, e.g. those using IPv6 Privacy Extensions
- support in some Microsoft services, though these have improved a lot in the last couple of years, e.g. MS Exchange now support IPv6
- some LAN security issues, in particular rogue RAs.

But none of those issues are critical. DHCPv6 support is maturing, with Mac OS X Lion having signs of DHCPv6 client support being included at the time of writing. There is now RA Guard defined [RFC6105] to help counter rogue RAs, and management and monitoring tools are becoming more aware of the trend for IPv6 systems to commonly be multi-addressed.

The key lesson we feel we have learnt from deployment is that IPv4 and IPv6 should, as far as possible, be considered equivalent, and the tools used to operate a dual-stack environment should be chosen to allow administrators to feel they are managing one consistent environment and not two separate networks. Examples include how IPv6 firewalls can be configured for devices with IPv4 and IPv6 addresses, or how DHCP management is provisioned. These are important questions to ask suppliers.

13.9 Future steps

A key point to emphasise is that in making the transition to support IPv6 dual-stack pervasively in its network, the School has not seen any significant adverse affect on IPv4 services. Making its DNS, MX and web servers all able to send/receive data over IPv4 or IPv6 has not impacted its network's robustness or reliability.

Much has been achieved but there is also still work to be done. A pilot deployment of DHCPv6 is highly desirable – we can then better assess the pros and cons of stateless versus stateful address configuration in our enterprise network. We also plan to test IPv6-only devices in our network, including provision of proxies and/or NAT64 at the edge for access to external IPv4 resources for those devices.

We also expect to make more use of IPv6 Multicast, which is currently being used for some lectures and for a local TV service.

We are also seeing both researchers and students writing new applications for IPv6, some of which are quite innovative, and the presence of IPv6 has encouraged that development. New equipment arriving in our labs includes IPv6 by default, e.g. new sensor nodes. Our student wireless society (SOWN) is also deploying IPv6 on its open community network.

14. IPv6 Roadmap

At present there are very few commercial IPv6 services in the UK. A number of ISPs have initial trial networks but only a handful of (relatively) niche ADSL suppliers including Andrews and Arnold and Goscomb are offering IPv6 to their customers.

The early IPv6 site (enterprise) deployments are being made in academic networks, though despite 100 UK academic sites having received their IPv6 /48 allocation as of 2010, very few have any significant production deployment yet. Loughborough, Southampton and Lancaster are amongst the early adopters.

Bodies and projects such as the IPv6 Forum,³⁰ 6DEPLOY [6DEPLOY] and the IPv6 Portal³¹ are involved in promoting the technology and presenting roadmaps for deployment. In the UK, a new group called 6UK³² is trying to promote IPv6 primarily to industrial and commercial organisations. More information can be found on their web sites.

Nominet, which manages the .uk namespace, now supports some IPv6 services, e.g. IPv6 transport for the .uk DNS, and also the ability to register an IPv6 DNS server address for a domain.

The question as to when IPv6 deployment will take off is an open one. When queried, many ISPs will say that their procurements of equipment include an eye on IPv6 readiness and that some even have IPv6 available in their core networks (but not the customer-facing edge).

The two most commonly cited reasons by UK ISPs for lack of more widespread deployment are the lack of consumer priced IPv6-ready DSL devices and a lack of IPv6 support in Operations Support System (OSS) tools for network management. While Cisco® 800 series ADSL routers support IPv6 and cheap Linksys routers can be modified to do so, adoption will still be limited until cheap, commodity devices support native IPv6 by default.

For the immediate future, IPv6 home networking is thus limited to niche ADSL providers using non-commodity DSL routers, or to users who wish to use a tunnelling service, such as those from SixXS or Hurricane Electric.

14.1 Recommendations for UK Academic Deployment

The key questions for HE/FE sites are when and how to begin IPv6 deployment.

14.1.1 When to Deploy?

JANET includes dual-stack IPv6 networking as a production service.

JANET's approach is in line with the many other NRENs and the GÉANT pan-European network who have deployed IPv6 in a similar way. The core academic backbones across Europe can all carry IPv6 natively.

The RNO deployments on JANET should continue to follow as and when their connected sites request IPv6 service; the SLA for JANET requires RNOs to commit to that deployment.

So the question is now one for sites/campuses – when to deploy? In principle, if they have enough IPv4 address space they can continue to operate as they have done in the past. But to ignore the exhaustion of the unused IPv4 address pool would be somewhat short sighted. While a site may continue to use IPv4 internally, and have enough global IPv4 address space for its externally facing services, it should consider the position of other sites it may wish to communicate with in the future. It is likely that many emerging networks, especially in countries with little IPv4 address space, will turn to IPv6, and thus supporting IPv6 in your own site would be the best approach to ensure you can reach and be reached by other IPv6 sites as smoothly as possible.

³⁰ <http://www.ipv6forum.com>

³¹ <http://www.ipv6tf.org>

³² <http://www.6uk.org.uk>

While there are no significant IPv6-only deployments at the time of writing, it is only a matter of time before these emerge. The Chinese CERNET2 backbone is IPv6-only; its connected sites are thus encouraged to use IPv6, or to face tunnelling IPv4-in-IPv6. For IPv6-only sites to flourish, they need to consider how they access IPv4-only content; the likely solution is NAT64 as discussed in the transition section in this guide. NAT64 is in its early stages of production testing but is likely to provide a workable solution for many scenarios.

If your site is concerned with talking to IPv6 sites, running dual-stack is the most appropriate medium-term solution; you can then use IPv6 to talk to other dual-stack or IPv6-only sites and IPv6-only sites can reach you without needing to use NAT64.

In parallel many large content providers and distributors are testing IPv6 – Google and Facebook both have versions of their content available over IPv6. For Google you need to apply (at the time of writing) to be DNS white-listed so you get AAAA records returned to you for Google content (including YouTube), while Facebook is available via <http://www.v6.facebook.com>. Akamai is working with IPv6 content distribution.

Thus it makes sense for all sites on JANET at the very least to include IPv6 capabilities in their procurements now, even if they do not plan to turn IPv6 on yet.

The next step is when to turn on IPv6. One approach is to make some of your public-facing services available over IPv6 and measure activity on those. Another is to find ‘tame’ groups in your site who may benefit from IPv6, e.g. a Computer Science department for teaching and research use, or the Computing Services department.

There are already some UK HE sites running IPv6 in production, so the answer to the original question is ‘whenever you are ready’: but at least take measures now to ensure you are ready, even if you do not deploy campus-wide for some time.

14.1.2 How to Deploy?

The phased approach described in a previous section should be able to be tailored to any HE/FE site. A site may choose to focus on enabling certain internal services first, or instead add IPv6 capability to externally facing services. Either way, the same approach should apply.

The best place to seek IPv6 help is on the IPv6 users JISCmail list, available at:

<http://www.jiscmail.ac.uk/archives/ipv6-users.html>

Any IPv6 queries can be asked on this list. Any IPv6 news related to JANET will also be sent to this list, so it is useful to be on the list to stay abreast of developments.

JANET’s general IPv6 web information is here:

<http://www.ja.net/development/network-engineering/ipv6/>

14.2 Feedback

If you have any comments on this guide, please contact the editor Tim Chown (tjc@ecs.soton.ac.uk). Your comments may be incorporated into or help shape future versions of this guide.

14.3 Acknowledgements

The author would like to thank numerous people who gave feedback and comments for the production of the original version of this guide, including Duncan Rogerson, Rob Evans, Rina Samani, Stig Venaas and John Linn.

For the revised version, the assistance of Martin Dunmore and Rob Evans was much appreciated in checking the accuracy and details of the updates.

15. References

IETF RFCs can be found at <http://www.ietf.org/rfc/rfcNNNN.txt>, where NNNN is the number of the RFC document.

- [3775BIS] Mobility Support in IPv6, IETF Internet Draft (work in progress), draft-mext-rfc3775bis-12, Feb 2011.
- [6DEPLOY] 6DEPLOY and 6DEPLOY2 web site:
<http://www.6deploy.org>
- [6NET] 6NET web site:
<http://www.6net.org>
- [ALLOC56] IPv6 Address Assignment to End Sites, IETF Internet Draft (work in progress), draft-ietf-v6ops-3177bis-end-sites-01, Jan 2011.
- [BERMUDA] The Bermuda 2 project:
<http://www.ipv6.ac.uk/bermuda2/>
- [BROKER] JANET IPv6 Tunnel Broker pilot:
<http://www.broker.ipv6.ac.uk>
- [CAMPUS] IPv6 Campus Transition Scenario Description and Analysis, IETF Internet Draft (expired), draft-ietf-v6ops-campus-transition-01, Mar 2007.
- [D224] IPv4/IPv6 Transition Cookbook for ISP/NREN networks, 6NET Deliverable D2.4.4, Feb 2005:
<http://www.6net.org/publications/deliverables/D2.2.4.pdf>
- [D234] IPv4/IPv6 Transition Cookbook for end site networks/universities, 6NET Deliverable D2.3.4v2, Jun 2005:
<http://www.6net.org/publications/deliverables/D2.3.4v2.pdf>
- [DNS64] DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers, IETF Internet Draft (work in progress), draft-ietf-behave-dns64-11, Oct 2010.
- [ECSTV] ECS-TV:
<http://www.ecstv.ecs.soton.ac.uk/>
- [IETF] The Internet Engineering Task Force:
<http://www.ietf.org>
- [IST-IPV6] The IPv6 Cluster and Portal:
<http://www.ist-ipv6.org>
- [JANETEXP] JANET IPv6 Experimental Service:
<http://www.ja.net/development/network-engineering/ipv6/ipv6-experimental-service.html> (under update at time of writing)
- [LISP] Locator/ID Separation Protocol (LISP), IETF Internet Draft (work in progress), draft-ietf-lisp-10, Mar 2011.
- [NAT64] Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers, IETF Internet Draft (work in progress), draft-ietf-behave-v6v4-xlate-stateful-12, Jul 2010.
- [NAT66] IPv6-to-IPv6 Network Prefix Translation, IETF Internet Draft (work in progress), draft-mrw-nat66-09, Mar 2011.
- [NODE-BIS] IPv6 Node Requirements RFC 4294-bis, IETF Internet Draft (work in progress), draft-ietf-6man-node-req-bis-08, Mar 2011.
- [P2PLINKS] Using 127-bit IPv6 Prefixes on Inter-Router Links, IETF Internet Draft (work in progress), draft-koho-ipv6-prefixlen-p2p-3, Oct 2010.
- [RFC1347] TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing, IETF RFC 1347, Jun 1992.
- [RFC1519] Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy, IETF RFC 1519, Sep 1993.
- [RFC1631] The IP Network Address Translator (NAT), IETF RFC 1631, May 1994.
- [RFC1707] CATNIP: Common Architecture for the Internet, IETF RFC 1707, Oct 1994.

-
- [RFC1710] Simple Internet Protocol Plus White Paper, IETF RFC 1710, Oct 1994.
- [RFC1752] The Recommendation for the IP Next Generation Protocol, IETF RFC 1752, Jan 1995.
- [RFC1918] Address Allocation for Private Internets, IETF RFC 1918, Feb 1996.
- [RFC2080] RIPng for IPv6, IETF RFC 2080, Jan 1997.
- [RFC2460] Internet Protocol, Version 6 (IPv6) Specification, IETF RFC 2460, Dec 1998.
- [RFC2464] Transmission of IPv6 Packets over Ethernet Networks, IETF RFC 2464, Dec 1998.
- [RFC2732] Format for Literal IPv6 Addresses in URL's, IETF RFC 2732, Dec 1999.
- [RFC2740] OSPF for IPv6, IETF RFC 2740, Dec 1999.
- [RFC2766] Network Address Translation – Protocol Translation (NAT-PT), IETF RFC 2766, Feb 2000.
- [RFC2775] Internet Transparency, IETF RFC 2775, Feb 2000.
- [RFC2858] Multiprotocol Extensions for BGP-4, IETF RFC 2858, Jun 2000.
- [RFC2993] Architectural Implications of NAT, IETF RFC 2993, Nov 2000.
- [RFC3022] Traditional IP Network Address Translator (Traditional NAT), IETF RFC 3022, Jan 2001.
- [RFC3053] IPv6 Tunnel Broker, IETF RFC 3053, Jan 2001.
- [RFC3056] Connection of IPv6 Domains via IPv4 Clouds, IETF RFC 3056, Feb 2001.
- [RFC3068] AN Anycast Prefix for 6to4 Relay Routers, IETF RFC 3068, Jun 2001.
- [RFC3142] An IPv6-to-IPv4 Transport Relay Translator, IETF RFC 3142, Jun 2001.
- [RFC3177] IAB/IESG Recommendations on IPv6 Address Allocations to Sites, IETF RFC 3177, Sep 2001 (under revision by draft-narten-ipv6-3177bis-48boundary-00).
- [RFC3234] Middleboxes: Taxonomy and Issues, IETF RFC 3234, Feb 2002.
- [RFC3306] Unicast-Prefix-based IPv6 Multicast Addresses, IETF RFC 3306, Aug 2002.
- [RFC3315] Dynamic Host Configuration Protocol for IPv6 (DHCPv6), IETF RFC 3315, Jul 2003.
- [RFC3484] Default Address Selection for Internet Protocol Version 6 (IPv6), IETF RFC 3484, Feb 2003.
- [RFC3493] Basic Socket Interface Extensions for IPv6, IETF RFC 3493, Feb 2003.
- [RFC3542] Advanced Sockets Application Program Interface (API) for IPv6, IETF RFC 3542, May 2003.
- [RFC3627] Use of /127 Prefix Length Between Routers Considered Harmful, IETF RFC 3627, Sep 2003.
- [RFC3697] The IPv6 Flow Label Specification, IETF RFC 3697, Mar 2004.
- [RFC3701] 6bone (IPv6 Testing Address Allocation) Phaseout, IETF RFC 3701, Mar 2004.
- [RFC3736] Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6, IETF RFC 3736, Apr 2004.
- [RFC3775] Mobility Support in IPv6, IETF RFC 3775, Jun 2004.
- [RFC3810] Multicast Listener Discovery version 2 (MLDv2) for IPv6, IETF RFC 3810, Jun 2004.
- [RFC3849] IPv6 Address Prefix Reserved for Documentation, IETF RFC 3849, Jul 2004.
- [RFC3935] A Mission Statement for the IETF, IETF RFC 3935, Oct 2004.
- [RFC3956] Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address, IETF RFC 3956, Nov 2004.
- [RFC3971] Secure Neighbour Discovery (SEND), IETF RFC 3971, Mar 2005.
- [RFC3972] Cryptographically Generated Addresses (CGAs), IETF RFC 3972, Mar 2005.
- [RFC3974] SMTP Operational Experience in Mixed IPv4/v6 Environments, IETF RFC 3974, Jan 2005.
- [RFC4029] Scenarios and Analysis for Introducing IPv6 into ISP Networks, IETF RFC 4029, Mar 2005.
- [RFC4038] Application Aspects of IPv6 Transition, IETF RFC 4038, Mar 2005.

- [RFC4192] Procedures for Renumbering an IPv6 Network without a Flag Day, IETF RFC 4192, Sep 2005.
- [RFC4193] Unique Local IPv6 Unicast Addresses, IETF RFC 4193, Oct 2005.
- [RFC4291] IP Version 6 Addressing Architecture, IETF RFC 4291, Feb 2006.
- [RFC4294] IPv6 Node Requirements, IETF RFC 4294, Apr 2006.
- [RFC4380] Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs), IETF RFC 4380, Feb 2006.
- [RFC4423] Host Identity Protocol (HIP) Architecture, IETF RFC 4423, May 2006.
- [RFC4487] Mobile IPv6 and Firewalls: Problem Statement, IETF RFC 4487, May 2006.
- [RFC4554] Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks, IETF RFC 4554, Jun 2006.
- [RFC4632] Classless Inter-domain Routing, IETF RFC 4632, Aug 2006.
- [RFC4798] Connecting IPv6 Islands over IPv4 MPLS using IPv6 Provider Edge Routers (6PE), IETF RFC 4798, Feb 2007.
- [RFC4861] Neighbour Discovery for IP version 6 (IPv6), IETF RFC 4861, Sep 2007.
- [RFC4862] IPv6 Stateless Address Autoconfiguration, IETF RFC4862, Sep 2007.
- [RFC4864] Local Network Protection for IPv6, IETF RFC 4864, May 2007.
- [RFC4890] Recommendations for Filtering ICMPv6 Messages in Firewalls, IETF RFC 4890, May 2007.
- [RFC4941] Privacy Extensions for Stateless Address Autoconfiguration in IPv6, IETF RFC 4941, Sep 2007.
- [RFC4942] IPv6 Transition/Co-existence Security Considerations, IETF RFC 4942, Sep 2007.
- [RFC4966] Reasons to move NAT-PT to Historical Status, IETF RFC 4966, Jul 2007.
- [RFC5006] IPv6 Router Advertisement Option for DNS Configuration, IETF RFC 5006, Sep 2007.
- [RFC5157] IPv6 Implications for Network Scanning, IETF RFC 5157, Mar 2008.
- [RFC5308] Routing IPv6 with IS-IS, IETF RFC 5308, Oct 2008.
- [RFC5375] IPv6 Unicast Address Assignment Considerations, IETF RFC 5375, Dec 2008.
- [RFC5533] Shim6: Level 3 Multihoming Shim protocol for IPv6, IETF RFC 5533, Jun 2009.
- [RFC5569] IPv6 Rapid Deployment on IPv4 Infrastructures, IETF RFC 5569, Jan 2010.
- [RFC5572] IPv6 Tunnel Broker with the Tunnel Setup Protocol (TSP), IETF RFC 5572, Feb 2010.
- [RFC5375] IPv6 Unicast Address Assignment Considerations, RFC 5375, Dec 2008.
- [RFC5952] A Recommendation for IPv6 Address Text Representation, IETF RFC 5952, Aug 2010.
- [RFC6104] Rogue IPv6 Router Advertisement Problem Statement, RFC 6104, Feb 2011.
- [RFC6105] IPv6 Router Advertisement Guard, RFC 6105, Feb 2011.
- [RIPE501] Requirements for IPv6 in ICT Equipment, RIPE-501, RIPE NCC, Nov 2010, <http://www.ripe.net/ripe/docs/ripe-501>.
- [SURGE] IPv6 Surge Radio:
<http://www.surgeradio.co.uk>
- [UUID] Definition of the UUID-based DHCPv6 Unique Identifier (DUID-UUID), IETF Internet Draft (work in progress), draft-ietf-dhc-duid-uuid-03, Feb 2011.
- [V6ONLY] Experiences from an IPv6-only Network, IETF Internet Draft (work in progress), draft-arkko-ipv6-only-experience-02, Oct 2010.

Glossary

Term	Explanation
ACL	Access Control List.
AH	Authentication Header.
ALG	Application Layer Gateway.
ASM	Any Source Multicast.
BSR	Bootstrap Router
BU	Binding Update.
CGA	Cryptographically Generated Address (see RFC 3972).
CIDR	Classless Inter Domain Routing (see RFC1519).
DAD	Duplicate Address Detection (see RFC2462).
DSCP	Differentiated Service Code Point.
ESP	Encapsulated Security Payload.
HA	Home Agent.
IETF	Internet Engineering Task Force.
IPv4	Internet Protocol version 4. The current widely deployed version of the Internet Protocol.
IPv6	Internet Protocol version 6. The successor to the existing version of IP, IPv4.
JSD	JANET Service Desk
MIB	Management Information Base.
MIPv6	Mobile IPv6 (see RFC3775).
MN	Mobile Node.
MSDP	Multicast Source Discovery Protocol.
MTA	Mail Transfer/Transport Agent.
NA	Neighbour Advertisement, part of ND (see RFC2461).
NAT	Network Address Translation (see RFC1631).
NAT64	Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers
NAT66	Stateless IPv6 Prefix Translation
NAT-PT	Network Address Translation – Protocol Translation.
ND	Neighbour Discovery (see RFC2461).
NS	Neighbour Solicitation (see RFC2461).
NREN	National Research and Education Network.
PIM-SM	Protocol Independent Multicast – Sparse Mode
PMTU	Path Maximum Transmission Unit.
RFC	'Request For Comments', an IETF standardisation/informational document.
RA	Router Advertisement (see RFC2462).
RIPE	Réseaux IP Européens.
RIR	Regional Internet Registry.
RS	Router Solicitation (see RFC2462).
SLAAC	Stateless Address Autoconfiguration (see RFC2462).
SEND	Secure Neighbour Discovery (see RFC3971).
SNMP	Simple Network Management Protocol.

SSM Source Specific Multicast.
TUBA TCP and UDP with Bigger Addresses (see RFC1347).

Appendix A: Sample Cisco® IOS Configuration File

In this Appendix the authors cite as an example the IOS configuration file for the Cisco® 7206 router used to provide the ECS IPv6 connectivity during the 6NET project. This has an externally facing dual-stack interface and an internally facing IPv6-only interface (forwarding traffic to a BSD-based IPv6 firewall). The internal ECS prefix is a /52. External unicast IPv6 connectivity via the regional regional network (LeNSE) is native, but IPv6 multicast is tunnelled to JANET's multicast IPv6 router. Some other external tunnels are in use from the router.

Some interesting points to note:

- the DNS server the router uses can be over IPv6 transport
- the 'general-prefix' construct is useful for renumbering
- IPv6 NetFlow is configured and running
- the external uplink (to LeNSE) has a ULA filter on it to prevent leakage of 'private' IPv6 address space
- embedded-RP for multicast is in use on the site.

NB. This configuration is for example purposes only, and has been edited for (some) brevity, so is thus not complete. It is intended as an example, to highlight IPv6-specific elements.

!

```

version 12.4
!
hostname ford
!
flow exporter ecs-nfsen-export
  destination 152.78.68.183
!
flow monitor sown-nfsen
  description Clone of ECS nflow to prove it works
  record netflow ipv6 original-input
  exporter sown-nfsen-export
!
flow monitor ecs-nfsen
  description ECS NetFlow monitor (mrtg.ecs.soton.ac.uk/nfsen.php)
  record netflow ipv6 original-input
  exporter ecs-nfsen-export
!
no ip source-route
ip cef
!
ip domain name ecs.soton.ac.uk
ip name-server 2001:630:D0:F116::53
ip name-server 152.78.68.1
ip name-server 152.78.70.1
ipv6 general-prefix UOS 2001:630:D0::/48
ipv6 general-prefix ECS 2001:630:D0:F000::/52
ipv6 general-prefix SOWN 2001:630:D0:F600::/55

```

```
ipv6 general-prefix TBROKER48 2001:630:C2::/48
ipv6 general-prefix TBROKER56 2001:630:D0:FC00::/56
ipv6 host sown-dc 2001:630:D0:F000::8:2
ipv6 host tbroker-dc 2001:630:D0:F000::3:2
ipv6 host ecs-dc 2001:630:D0:F000::2
ipv6 host ha-dc 2001:630:D0:F000::3:3
ipv6 unicast-routing
ipv6 cef
ipv6 multicast-routing
!
interface Tunnel5
  description IPv6 multicast tunnel to Janet
  no ip address
  ipv6 address 2001:630:0:1::6A/126
  ipv6 mtu 1450
  ipv6 multicast boundary scope 8
  tunnel source 152.78.108.2
  tunnel destination 194.82.173.253
  tunnel mode ipv6ip
  tunnel checksum
!
interface GigabitEthernet0/1
  description Downlink to ECS v6 firewall (zaphod)
  no ip address
  duplex auto
  speed auto
  media-type rj45
  no negotiation auto
  ipv6 address ECS ::/112 anycast
  ipv6 address ECS ::1/112
  ipv6 nd ra suppress
  ipv6 verify unicast reverse-path
  ipv6 rip internal enable
  ipv6 rip internal default-information originate
  ipv6 multicast boundary scope 5
  ipv6 flow monitor ecs-nfsen input
  ipv6 flow monitor ecs-nfsen output
!
interface GigabitEthernet0/2
  description Downlink to hexago tunnel broker
  no ip address
  duplex auto
  speed auto
  media-type rj45
  no negotiation auto
  ipv6 address ECS ::0.3.0.0/112 anycast
```

```
ipv6 address ECS ::0.3.0.1/112
ipv6 nd ra suppress
ipv6 verify unicast reverse-path
ipv6 multicast boundary scope 5
ipv6 flow monitor ecs-nfsen input
ipv6 flow monitor ecs-nfsen output
!
interface GigabitEthernet0/3
description Downlink to SOWN
no ip address
duplex auto
speed auto
media-type rj45
no negotiation auto
ipv6 address ECS ::0.8.0.0/112 anycast
ipv6 address ECS ::0.8.0.1/112
ipv6 nd ra suppress
ipv6 verify unicast reverse-path
ipv6 multicast boundary scope 5
ipv6 flow monitor sown-nfsen input
ipv6 flow monitor ecs-nfsen input
ipv6 flow monitor ecs-nfsen output
!
interface GigabitEthernet1/0
description Uplink to LENSE
ip address 152.78.108.2 255.255.255.248
negotiation auto
ipv6 address 2001:630:C1:100::2/64
ipv6 enable
ipv6 traffic-filter ulaleak out
ipv6 nd ra suppress
ipv6 multicast boundary scope 8
ipv6 flow monitor ecs-nfsen input
ipv6 flow monitor ecs-nfsen output
!
router bgp 8933
bgp router-id 152.78.108.2
no bgp default ipv4-unicast
bgp log-neighbor-changes
neighbor 2001:630:0:1::69 remote-as 786
neighbor 2001:630:C1:100::1 remote-as 786
!
address-family ipv6
network 2001:630:D0::/48
exit-address-family
!
```

```
address-family ipv6 multicast
neighbor 2001:630:0:1::69 activate
neighbor 2001:630:0:1::69 send-community both
network 2001:630:D0::/48
exit-address-family
!
ip forward-protocol nd
ip route 0.0.0.0 0.0.0.0 152.78.108.6
no ip http server
no ip http secure-server
!
ipv6 route 2001:630:C2::/48 2001:630:D0:F000::3:2
ipv6 route 2001:630:D0:F000::/55 2001:630:D0:F000::2
ipv6 route 2001:630:D0:F400::/55 2001:630:D0:F000::3:3
ipv6 route 2001:630:D0:F600::/55 2001:630:D0:F000::8:2
ipv6 route 2001:630:D0:FC00::/56 2001:630:D0:F000::3:2
ipv6 route 2001:630:D0::/48 Null0
ipv6 route ::/0 2001:630:C1:100::1 unicast
ipv6 route ::/0 2001:630:0:1::69 multicast
ipv6 router rip internal
!
ipv6 pim rp-address 2001:630:D0:F000::1 rpuos-org
ipv6 pim rp-address 2001:660:3007:300:1:: rpm6bone
!
!
route-map router-to-internal permit 10
match ip address internal-traffic
set ip next-hop 10.78.64.60
!
route-map router-to-internal permit 20
set ip next-hop 152.78.108.6
!
ipv6 access-list ulaleak
deny ipv6 FC00::/7 any
permit ipv6 any any
!
ipv6 access-list rpuos-org
permit ipv6 any FF08::/16
permit ipv6 any FF18::/16
sequence 90 permit ipv6 any FF78::/16
permit ipv6 any FF7E::/16
!
```

```
ipv6 access-list rpm6bone
  permit ipv6 any FF0E::/16
  permit ipv6 any FF1E::/16
!
```

```
end
```

JANET(UK) manages the networking programme on behalf of the higher and further education and research community in the United Kingdom. JANET, the United Kingdom's academic and research network, is funded by the Joint Information Systems Committee (JISC).

For further information please contact:

JANET Service Desk
Lumen House, Library Avenue
Harwell Oxford
Didcot, Oxon, OX11 0QS

Tel: 0300 300 2212
Fax: 0300 300 2213
E-mail: service@ja.net

Copyright:

This document is copyright The JNT Association trading as JANET(UK). Parts of it, as appropriate, may be freely copied and incorporated unaltered into another document unless produced for commercial gain, subject to the source being appropriately acknowledged and the copyright preserved. The reproduction of logos without permission is expressly forbidden. Permission should be sought from JANET Service Desk.

Trademarks:

JANET® is a registered trademark of the Higher Education Funding Councils for England, Scotland and Wales. The JNT Association is the registered user of this trademark. JANET(UK) is a registered trademark of the JNT Association.

Cisco® is a registered trademark of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries.

The term Linux® is a registered trademark of Linus Torvalds, the original author of the Linux kernel.

Microsoft® and Windows® are registered trademarks of Microsoft Corporation in the United States and/or other countries.

Disclaimer:

The information contained herein is believed to be correct at the time of issue, but no liability can be accepted for any inaccuracies.

The reader is reminded that changes may have taken place since issue, particularly in rapidly changing areas such as internet addressing, and consequently URLs and e-mail addresses should be used with caution.

The JNT Association cannot accept any responsibility for any loss or damage resulting from the use of the material contained herein.

Availability:

Further copies of this document may be obtained from JANET Service Desk at the above address.

This document is also available electronically from: <http://www.ja.net/services/publications/technical-guides.html>



