

IPv6 Operations T. Chown
Internet-Draft University of
 Southampton
Intended status: March 28, 2007
Informational
Expires: September 29,
2007

[TOC](#)

IPv6 Campus Transition Scenario Description and Analysis draft-ietf-v6ops-campus-transition-01

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 29, 2007.

Abstract

In this document we consider and analyse the specific scenario of IPv6 transition and deployment in a large department of a university campus network. The department is large enough to operate its own instances of all the conventional university services including (for example) web, DNS, email, filestore, interactive logins, and remote and wireless access. The scenario is a dual-stack one, i.e. transition to IPv6 means deploying IPv6 in the first instance (and probably for some time) alongside IPv4. This analysis identifies the available components for IPv6 transition, while validating the applicability of the IPv6 Enterprise Network Scenarios informational text. It focuses on the network and associated service elements of the transition, rather than the application elements.

Table of Contents

- [1.](#) Introduction
 - [1.1.](#) Site Philosophy
- [2.](#) Discussion of Scenarios Network Infrastructure Components
 - [2.1.](#) Component 1: Enterprise Provider Requirements
 - [2.2.](#) Component 2: Enterprise Application Requirements
 - [2.3.](#) Component 3: Enterprise IT Department Requirements
 - [2.4.](#) Component 4: Enterprise Network Management System
 - [2.5.](#) Component 5: Enterprise Network Interoperation and Coexistence
- [3.](#) Discussion of Network Infrastructure Component Requirements
 - [3.1.](#) DNS
 - [3.2.](#) Routing
 - [3.3.](#) Configuration of Hosts
 - [3.4.](#) Security
 - [3.5.](#) Applications
 - [3.6.](#) Network Management
 - [3.7.](#) Address Planning
 - [3.8.](#) Multicast
 - [3.9.](#) Multihoming
- [4.](#) Specific Scenario Component Review
 - [4.1.](#) Network Components
 - [4.1.1.](#) Physical connectivity (Layer 2)
 - [4.1.2.](#) Routing and Logical subnets (Layer 3)
 - [4.1.3.](#) Firewall
 - [4.1.4.](#) Intrusion Detection System
 - [4.1.5.](#) Management
 - [4.1.6.](#) Monitoring
 - [4.1.7.](#) Remote access
 - [4.1.8.](#) IPv6 External Access
 - [4.2.](#) Address Allocation Components
 - [4.2.1.](#) IPv6 network prefix allocation
 - [4.2.2.](#) IPv6 Address allocation
 - [4.3.](#) Core Services
 - [4.3.1.](#) Email
 - [4.3.2.](#) Web Hosting
 - [4.3.3.](#) Directory Services
 - [4.3.4.](#) DNS
 - [4.3.5.](#) NTP
 - [4.3.6.](#) Multicast
 - [4.4.](#) Hard-coded address points
- [5.](#) IPv6 Enterprise Deployment Procedure
 - [5.1.](#) Advanced Planning
 - [5.2.](#) Testbed/Trial Deployment
 - [5.3.](#) Production Deployment
- [6.](#) Analysis: Dual-Stack Deployment - Transition toolbox
- [7.](#) Analysis: Considerations beyond the Scenarios Document
- [8.](#) Summary

- [9. Acknowledgements](#)
- [10. IANA Considerations](#)
- [11. Security Considerations](#)
- [12. Informative References](#)
- [§ Author's Address](#)
- [§ Intellectual Property and Copyright Statements](#)

1. Introduction

[TOC](#)

The scope of the enterprise network transition scenarios being considered by the IETF is very large, much more so than that of the other three IPv6 transition areas that have been studied ([ISP \(Lind, M., Ksinant, V., Park, S., Baudot, A., and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks," March 2005.\)](#) [RFC4029], [unmanaged \(Huitema, C., Austein, R., Satapati, S., and R. van der Pol, "Evaluation of IPv6 Transition Mechanisms for Unmanaged Networks," September 2004.\)](#) [RFC3904] and [3GPP \(Wiljakka, J., "Analysis on IPv6 Transition in Third Generation Partnership Project \(3GPP\) Networks," October 2005.\)](#) [RFC4215]). However, an [IPv6 Enterprise Network Scenarios \(Bound, J., "IPv6 Enterprise Network Scenarios," June 2005.\)](#) [RFC4057] description has been produced. In this case study document we present our experience in a specific area for IPv6 transition, namely a large department (1,500 staff and students, over 1,000 hosts) in an academic campus network. The purpose of this document is to both define and analyse the IPv6 transition of such a network, but also to test and validate the applicability of the IPv6 Enterprise Network Scenarios document to a specific example. This document describes the transition focusing on the network elements. Our campus study falls under Scenario 1 of the [IPv6 Enterprise Network Scenarios \(Bound, J., "IPv6 Enterprise Network Scenarios," June 2005.\)](#) [RFC4057] document, i.e. the campus network is an existing IPv4 network, where IPv6 is to be deployed in conjunction with the IPv4 network.

Scenario 1 has the assumption that the IPv4 network infrastructure used has an equivalent capability in IPv6. This document analyses that assumption. The Scenario also has requirements, i.e. that the existing IPv4 network infrastructure is not disrupted, and that IPv6 should be equivalent or better than the network infrastructure in IPv4. The Scenario also notes that it may also not be feasible to deploy IPv6 on all parts of the network immediately.

These assumptions and requirements will be discussed later in this text. An incremental deployment strategy may, for example, be a desirable property.

It should also be noted why Scenarios 2 and 3 did not apply to this campus transition scenario. Scenario 2 talks of specific applications, but in the campus case we wish to deploy IPv6 pervasively, in wired and

wireless networks, as an enabler for education and research, to encourage new application development. Scenario 3 focuses on using IPv6 as the basis for most network communication, but in the campus we already have a significant IPv4 deployment that will be utilised for the foreseeable future (Scenario 3 would perhaps be more appropriate for a green field deployment).

1.1. Site Philosophy

[TOC](#)

The site which is the subject of this study is a large departmental network on a campus. That network (prior to transition) is an IPv4 network with around 20 subnets, using a core network infrastructure that combines switch-router functionality in central devices, with switches at the network edge. The main switching equipment is all VLAN (IEEE 802.1q) capable. There are around 1,000 networked nodes and 1,500 users, not including transient (mainly wireless) visitors.

The site wishes to deploy IPv6 dual-stack to support its own users along with its teaching and research needs. The goal is to IPv6 enable the network (on the wire) and services (DNS, SMTP, etc) such that the whole operation is dual-stack. This in due course would allow IPv6-only devices to be deployed within the fully IPv6-capable environment. Some network links may become IPv6-only in a subsequent phase in the future. This text has evolved over time. When we began writing, the department did not have IPv6 capability on its existing IPv4 routing equipment, thus an interim deployment method was required until the next router procurement. We discuss that interim solution within this document, and present the discussion from an initial point of an interim parallel IPv6 deployment prior to unifying the IPv4 and IPv6 routing on a single platform. Our initial deployment plan used a separate IPv6 path into the department with a parallel routing infrastructure for IPv6. In practice this meant that our initial deployment used a parallel IPv6 routing infrastructure, using BSD routers, for over three years, prior to deployment of a unified solution on a commercial platform.

2. Discussion of Scenarios Network Infrastructure Components

[TOC](#)

In this section, we look at the issues raised by following step by step the questions and considerations in the Scenarios Network Infrastructure Components of the [IPv6 Enterprise Network Scenarios \(Bound, J., "IPv6 Enterprise Network Scenarios," June 2005.\)](#) [RFC4057] document, section 3.2. This section is written from the perspective of pre-transition planning, although we are writing this document having undergone transition.

[TOC](#)

2.1. Component 1: Enterprise Provider Requirements

The answers to the questions posed in this section of the IPv6 Enterprise Network Scenarios document are as follows:

- *There is external access to/from the campus network, regional MAN and National Research Network beyond.
- *There are needs for access by remote staff, student and researchers.
- *It is a single site, with four geographically close buildings.
- *There are no leased lines or wide-area VPNs between remote buildings.
- *The department has 12 IPv4 Class C's, the campus has a Class B, independent from its provider (assigned prior to use of [CIDR \(Fuller, V. and T. Li, "Classless Inter-domain Routing \(CIDR\): The Internet Address Assignment and Aggregation Plan," August 2006.\)](#) [RFC4632]).
- *The IPv4 and IPv6 provider is the National Research and Education Network (JANET in the UK). JANET provides a /48 IPv6 site prefix for the university. The university offers a /52 prefix for the department.
- *The university and department make their own prefix allocations for subnets.
- *There is no multihoming, and thus no multihomed clients. The regional academic MAN supports network resilience measures.
- *The provider (JANET) offers an IPv6 [Tunnel Broker \(Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker," January 2001.\)](#) [RFC3053] service and a [6to4 \(Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds," February 2001.\)](#) [RFC3056] relay, though the campus is offered native IPv6 connectivity via its regional MAN.
- *There is no external IPv6 routing protocol needed due to the use of static route configuration between the campus and the regional network.
- *There is no external data centre.
- *IPv6 will run over the same access links to campus as IPv4 does (the JANET backbone uses true dual stack, the regional MAN uses [6PE \(De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge](#)

[Routers \(6PE\)," February 2007.\)](#) [RFC4798]). On campus, the IPv4 traffic to the department is received through a commercial firewall solution, while the IPv6 traffic will initially be received through a BSD firewall. Thus the access links into the department for IPv4 and IPv6 are different, though the goal in the longer term is to make them the same.

2.2. Component 2: Enterprise Application Requirements

[TOC](#)

The focus of this document is network transition and services, but the answers to the next IPv6 Enterprise Network Scenarios section on application aspects are as follows:

- *The application inventory is out of scope of this text.
- *We expect the first applications to be IPv6-enabled will be the support services, including DNS. The first applications should be the common IPv4 applications, e.g. web, remote login and email, such that IPv6 offers as least an equivalent service to IPv4 for the core service applications.
- *The academic environment has a good mix of open source and commercial software, predominantly either Microsoft or Linux, but with a growing number of Mac OS/X users. An exhaustive list of desktop, laptop and PDA platforms is out of scope of this text. Most open source applications have been upgraded to allow IPv6 operation out of the box; others can be upgraded given time.
- *The general goal is for applications to support both IPv4 or IPv6 operation, i.e. to be IP agnostic.
- *There is no use of NAT in the department's network. Home users, or users access into the network remotely from certain locations, may experience NAT at their client side.
- *NAT issues are relevant from the end-to-end perspective, for establishment of end-to-end security where desired, and in relation to IPv6 transition (remote access) methods that may need to be run through NATs.
- *There is a mix of internal and external applications. Where limitations occur, it is mainly by policy not technology, with that policy typically implemented through firewall restrictions.

[TOC](#)

2.3. Component 3: Enterprise IT Department Requirements

Here we list responses to the next IPv6 Enterprise Network Scenarios section on IT Department Requirements. Again, in this section we write our comments from a pre-transition perspective.

*Network and system ownership and support is all in-house.

*Remote VPNs are supported.

*No inter-site networking is required.

*No IP mobility support is required at this point, though we expect to use Mobile IPv6 between the department network and a local community wireless network, on our wireless LAN deployment as it grows in size, and to pilot its use inter-campus.

*The IPv6 address plan for the department requires a /52 prefix.

*There is no detailed asset database, though one exists providing a host inventory (for DNS and DHCP use).

*There are no (significantly) geographically separate sites.

*The internal IPv4 address assignment mechanism is DHCP for clients, with manual configuration for servers. We thus expect to use DHCPv6 for at least some, if not all, IPv6 clients. This will depend on availability of DHCP client and server software.

*Internal IPv4 routing is static or uses RIP. We thus expect to use RIPng internally.

*We expect our IPv6 network management policy to be very similar to that for IPv4. Having coherent policies, and a consistent means to configure them, should make network operation simpler.

*There is no QoS provision at present, largely due to the ample campus bandwidth (1Gbit/s uplink).

*Security is applied through many technologies implementing our policies, e.g. firewall, email scanning, IDS and wireless LAN access controls. We expect similar policies for IPv6, but need to analyse potential differences (e.g. considering use of RFC3041 privacy [addresses \(Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," January 2001.\)](#) [RFC3041]).

*Training will be done in-house.

*The impacted software service components are discussed in the next main section. Not all functions are upgradeable to IPv6; those that are not are discussed in the analysis sections. Some are, e.g. use of OpenLDAP (IPv6 capable) as an interim step in place of MS Active Directory (not IPv6 capable at the time of the analysis). Our view is that if components cannot be given immediate IPv6 equivalents, this functionality will come in due course, and IPv4 transport can be used in the interim. But the ultimate goal is to facilitate IPv6 capability.

*The impacted hardware components are discussed in the next main section. Not all hardware is upgradeable, e.g. network printers. There are no load balancing systems in use. There are wireless LAN hosts in the network that are mobile, but currently the wireless network is a single flat IPv4 subnet. There may be nodes moving to external wireless networks (i.e. the local community wireless network).

2.4. Component 4: Enterprise Network Management System

[TOC](#)

The responses to the next IPv6 Enterprise Network Scenarios section are as follows:

*No performance management is required. Systems are monitored for loading for purposes for future capacity planning.

*There are a number of network management and monitoring tools in use, which will need to be used in a dual stack or IPv6 mode, e.g. the nocol availability monitoring tools, and SNMP-based management.

*The configuration management may include use of tools to configure services including DNS and email. In-house DNS management tools are used.

*No policy management and enforcement tools are required.

*No detailed security management is required, though we expect to manage the implementations including firewalls and intrusion detection, and here a consistent management interface for both protocols is desirable..

*We may need to manage any specific deployed transition tools and mechanisms.

*We need to analyse the considerations IPv6 creates for network management, e.g. use (or not) of IPv6 privacy addresses. The need

for user privacy is recognised, but the need for simplified management is also a strong consideration.

2.5. Component 5: Enterprise Network Interoperation and Coexistence

[TOC](#)

Answers to the final IPv6 Enterprise Network Scenarios section on Coexistence are as follows:

- *An exhaustive list of platforms that are required to be IPv6 capable is out of scope of this text.
- *There is only one network ingress and egress point to the site that needs to be IPv6 capable; this is a Gigabit Ethernet interface.
- *The required transition mechanisms are discussed in the analysis section. In the initial phase of deployment, with the existing IPv4 switch-router equipment not supporting IPv6 routing, we expect to mainly use the [VLAN \(Chown, T., "Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks," June 2006.\)](#) [RFC4554] mechanism for internal IPv6 transport, with a parallel IPv6 routing infrastructure based on BSD routers, until the core infrastructure is able to support IPv6 (via upgrade or a new procurement).
- *The transition to IPv6 will be enabled on the wire first, enabling clients, with a phased introduction of service capability, as discussed below in the analysis section.
- *The preferred mechanism for interoperation with legacy nodes is to use dual-stack and thus IPv4 to communicate with IPv4 nodes and IPv6 to communicate to IPv6 nodes. We have not identified any in-house, non-upgradeable legacy software applications (most in-house applications are presented to users as web applications).

3. Discussion of Network Infrastructure Component Requirements

[TOC](#)

In this section, we discuss the network infrastructure component requirements raised in the [IPv6 Enterprise Network Scenarios \(Bound, J., "IPv6 Enterprise Network Scenarios," June 2005.\)](#) [RFC4057] document, in section 4. We document current IPv4 practices, and how we see these being facilitated when IPv6 is deployed and enabled.

3.1. DNS

[TOC](#)

The open source package BIND (version 9) is used for our three internal name servers. The servers will be made dual stack, to be available for IPv6 transport for local dual-stack or IPv6-only nodes. The three servers will each be listed with AAAA records, and AAAA glue added.

3.2. Routing

[TOC](#)

Internal unicast routing is either statically configured or uses RIP. We thus expect to use RIPng for internal IPv6 routing. The external routing is statically configured for IPv4, and thus is likely to be statically configured for IPv6.

3.3. Configuration of Hosts

[TOC](#)

IPv4 clients use DHCP for address and other configuration options. We expect to use [Dynamic Host Configuration Protocol for IPv6 \(DHCPv6\) \(Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 \(DHCPv6\)," July 2003.\) \[RFC3315\]](#) for IPv6 clients. This will require analysis of the [IPv4 and IPv6 Dual-Stack Issues for DHCPv6 \(Chown, T., Venaas, S., and C. Strauf, "Dynamic Host Configuration Protocol \(DHCP\): IPv4 and IPv6 Dual-Stack Issues," May 2006.\) \[RFC4477\]](#). We expect some clients, perhaps those in wireless LANs, to use [IPv6 Stateless Autoconfiguration \(Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration," December 1998.\) \[RFC2462\]](#), and these nodes will need support for [Stateless Dynamic Host Configuration Protocol \(DHCP\) Service for IPv6 \(Droms, R., "Stateless Dynamic Host Configuration Protocol \(DHCP\) Service for IPv6," April 2004.\) \[RFC3736\]](#) for other configuration options, including the IPv6 address of a local DNS resolver. Although IPv6 offers Stateless Autoconfiguration, it is expected that the managed environment will continue, driven from the asset database, for some time. The site administrators are comfortable with the use of DHCP for IPv4, and wish to use it for IPv6, for global address and potentially IPv6 Privacy Address assignment. Thus DHCPv6 is required. Use of Stateless Autoconfiguration implies a requirement for dynamic DNS updates for such nodes. It is not yet decided how to apply or enforce that plan; it may certainly be flexible with time.

[TOC](#)

3.4. Security

We need to identify new IPv6 related security considerations, and those associated with [transition mechanisms \(Davies, E., "IPv6 Transition/Co-existence Security Considerations," October 2006.\)](#)

[I-D.ietf-v6ops-security-overview]. Site policies may need to be updated as a result.

3.5. Applications

[TOC](#)

Discussion of applications is out of scope of this document. However, the [Application Aspects of IPv6 Transition \(Shin, M-K., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "Application Aspects of IPv6 Transition," March 2005.\)](#) [RFC4038] document describes best porting practice for applications. A new Advanced Sockets API for [IPv6 \(Stevens, W., Thomas, M., Nordmark, E., and T. Jinmei, "Advanced Sockets Application Program Interface \(API\) for IPv6," May 2003.\)](#) [RFC3542] defines the IP version independent API that is now widely supported. Recent versions of Java support IPv4 and IPv6 operation. There should also be consideration for making any required application proxies dual-stack.

3.6. Network Management

[TOC](#)

The network management and monitoring systems will need to support IPv6, and the management and monitoring of any transition mechanisms used to deploy IPv6. Monitoring includes usage tracking (e.g. via open source packages such as MRTG) and availability monitoring (e.g. via the Nagios package).

3.7. Address Planning

[TOC](#)

The department has been allocated 12 Class C prefixes for IPv4 use, and uses only globally routable addresses internally. No IPv4 NAT is used. The IPv4 address space for the campus was obtained prior to CIDR, but the IPv6 address space is allocated from the UK National Research Network (JANET) address space. The university receives a /48 prefix, and the department has a /52 allocation within this block.

Given that global IPv4 addresses are in use throughout our network, we plan to use global IPv6 addresses as well. Since we also do not expect to renumber (our IPv6 provider is expected to be JANET indefinitely) and our connectivity is expected to be stable we do not see any real need to deploy Unique Local [Addresses \(ULAs\) \(Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses," October 2005.\)](#)

[RFC4193]. Doing so would require full support for Default Address Selection (Draves, R., [“Default Address Selection for Internet Protocol version 6 \(IPv6\),” February 2003.](#)) [RFC3484] (so that ULA source addresses are used for ULA destinations, and global source addresses for global destinations) and running a two-faced DNS (with ULAs advertised only internally), which we do not currently do for IPv4. IPv6 address assignment planning for a campus-style enterprise is discussed separately in more detail in the [IPv6 Unicast Address Assignment Considerations \(Velde, G., Popoviciu, C., Chown, T., Bonness, O., and C. Hahn, “IPv6 Unicast Address Assignment Considerations,” September 2008.\)](#) [I-D.ietf-v6ops-addcon] text.

3.8. Multicast

[TOC](#)

IPv4 multicast is used for a number of applications, including AccessGrid multi-party videoconferencing. Connectivity is provided via the local campus and regional network. We expect to use [PIM-SM \(Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, “Protocol Independent Multicast - Sparse Mode \(PIM-SM\): Protocol Specification \(Revised\),” August 2006.\)](#) [RFC4601] for IPv6, initially as Any Source Multicast (ASM). We also plan to make use of Source Specific Multicast (SSM) more heavily in IPv6, bringing IPv6 and SSM together in one deployment cycle.

The use of IPv6 multicast makes it much simpler for our site to generate its own globally unique multicast group addresses than is the case for IPv4, where we need to use GLOP [space \(Albanna, Z., Almeroth, K., Meyer, D., and M. Schipper, “IANA Guidelines for IPv4 Multicast Address Assignments,” August 2001.\)](#) [RFC3171] from an upstream provider. For IPv6, you can generate your own unique multicast group address for regular [groups \(Haberman, B. and D. Thaler, “Unicast-Prefix-based IPv6 Multicast Addresses,” August 2002.\)](#) [RFC3306] or Embedded-RP [groups \(Savola, P. and B. Haberman, “Embedding the Rendezvous Point \(RP\) Address in an IPv6 Multicast Address,” November 2004.\)](#) [RFC3956] based on your unicast prefix (typically /48 or /64).

Since there is no [MSDP \(Fenner, B. and D. Meyer, “Multicast Source Discovery Protocol \(MSDP\),” October 2003.\)](#) [RFC3618] equivalent for IPv6, we only expect to use regular unicast prefix based group [addresses \(Haberman, B. and D. Thaler, “Unicast-Prefix-based IPv6 Multicast Addresses,” August 2002.\)](#) [RFC3306] within our own organisational scope. For wider scope multicast we expect to use Embedded-RP where possible, running our own IPv6 Rendezvous Point(s) to support our own content. In terms of the IPv6 address [architecture \(Hinden, R. and S. Deering, “IP Version 6 Addressing Architecture,” February 2006.\)](#) [RFC4291], we plan to use a site scope (ff05) for our department, with the university having organisational scope (ff08). Locally assigned group IDs would honour the guidelines of [RFC3307](#)

[\(Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses," August 2002.\)](#) [RFC3307].

3.9. Multihoming

[TOC](#)

The site is not multihomed for IPv4, and thus will not be for IPv6. This is typical for UK academic networks, where resilience is provisioned through the regional MAN links.

4. Specific Scenario Component Review

[TOC](#)

Here we describe specific technology in use now in the department. Later in this section we discuss any items not included in the above section, i.e. those not explicitly mentioned in the IPv6 Enterprise Network Scenarios document. Note that not all applications and services have at the time of writing been made IPv6 capable; in general open source packages are IPv6 capable out of the box, but discussion of specific applications is outside the scope of this text (this text aims to be a stable description of the processes and thinking followed during our campus transition).

4.1. Network Components

[TOC](#)

4.1.1. Physical connectivity (Layer 2)

[TOC](#)

*Switched Ethernet

*Gigabit Ethernet

*Wireless networking (802.11a/b/g)

4.1.2. Routing and Logical subnets (Layer 3)

[TOC](#)

The hybrid Layer 2/3 routing equipment has approximately 20 internal IPv4 subnets (in effect, routed VLANs). The only specific internal routing protocol used is [RIP \(Malkin, G., "RIP Version 2," November 1998.\)](#) [RFC2453]. There is a static route via the site firewall to the main upstream provider (academic) running at 1Gbit/s.

We would expect to use [RIPng \(Malkin, G. and R. Minnear, "RIPng for IPv6," January 1997.\)](#) [RFC2080] for IPv6 internally.

4.1.3. Firewall

[TOC](#)

The firewall is currently one running on a commercial hardware platform without IPv6 support. There is one internal facing interface, one external facing interface, and two DMZ interfaces, one for wired hosts and one for the Wireless LAN provision. We expect the topology to remain the same, with the DMZ(s) becoming dual-stack.

4.1.4. Intrusion Detection System

[TOC](#)

The Snort open source package is used locally for IPv4 IDS. Work on IPv6 capability for Snort is ongoing, but needs to consider both similar (e.g. application transport) issues as IPv4 as well as IPv6-specific issues (e.g. excessive use of Hop-by-Hop options).

4.1.5. Management

[TOC](#)

Some network management is performed by SNMP; there is no specific package for this (scripts used are in-house).

4.1.6. Monitoring

[TOC](#)

A number of open source tools are used, to monitor network usage as well as systems availability, e.g. Nagios and MRTG. The network testing tools include iperf, rude and crude.

4.1.7. Remote access

[TOC](#)

*RADIUS servers (our current RADIUS package supports IPv6)

*VPN servers

[TOC](#)

4.1.8. IPv6 External Access

*IPv6 connectivity will come via our regional MAN (which runs 6PE) through trunked (unrouted) VLANs across campus to our departmental network. Because the existing IP firewall pre-transition does not support IPv6, IPv6 will need to be transported into the departmental network via a separate parallel IPv6 capable firewall (e.g. a BSD system using a package such as pf).

4.2. Address Allocation Components

[TOC](#)

The department receives its IPv4 and IPv6 address allocations from the University. For IPv4, the University has a Class B allocation which is not aggregated under the JANET NREN address space post-CIDR. For IPv6, the University receives its allocation from JANET.

4.2.1. IPv6 network prefix allocation

[TOC](#)

For IPv6, JANET currently has a /32 prefix from RIPE-NCC, as the national academic ISP in the UK. The university has been allocated a /48 from this block by JANET. The department IPv6 deployment will be allocated a /52 size prefix from the university allocation. In the initial deployment, we expect that IPv4 and IPv6 subnets will be congruent (and share the same VLANs). This is because the existing subnet divisions are made for geographic or administrative reasons that are not IP version dependent (e.g. by building location or research group membership). One advantage of IPv6 is that subnets will not need to be resized to conserve or efficiently utilise address space as is the case currently for IPv4 (as subnet host counts rise and fall for administrative or research group growth/decline reasons).

4.2.2. IPv6 Address allocation

[TOC](#)

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

*Manually configured addresses (in some servers)

*Stateful DHCPv6 (probably in fixed, wired devices and some servers)

*Stateless address autoconfiguration (probably in wireless and mobile devices)

*RFC3041 privacy addresses (in some client devices)

For devices using stateless or RFC3041 mechanisms, at least a Stateless DHCPv6 [service \(Droms, R., "Stateless Dynamic Host Configuration Protocol \(DHCP\) Service for IPv6," April 2004.\)](#) [RFC3736] will be required for other (non-address) configuration options, e.g. DNS and NTP servers. It is likely that a full DHCPv6 service would provide this function however.

As discussed above, due to current experience with DHCP for IPv4, where all addresses are managed centrally, we expect that use of DHCPv6 for address allocation and management will be preferred (once implementations are mature).

4.3. Core Services

[TOC](#)

4.3.1. Email

[TOC](#)

There are three MX hosts for inbound email, and two main internal mail servers. Sendmail is the MTA. MailScanner is used for anti-spam/anti-virus. This uses external services including various RBLs for part of its spam checking. Successful reverse DNS lookup is required for sendmail to accept internal SMTP connections for delivery. Email access is provided by a variety of open source and commercial client and server applications (including a web front end) the details of which are outside the scope of this document.

We expect to continue to use sendmail for MX and MTA functions, as it supports IPv6 out of the box. Each of our MX servers will be made dual stack, noting the considerations in [RFC3974 \(Nakamura, M. and J. Hagino, "SMTP Operational Experience in Mixed IPv4/v6 Environments," January 2005.\)](#) [RFC3974].

4.3.2. Web Hosting

[TOC](#)

Web content hosting is provided either with Apache 2.x (open source) or in some cases commercial equivalents. Common components used to build systems with are MySQL, PHP and Perl; these enable local tools such as Wikis to be run. Apache 2.x has support for IPv6 included.

[TOC](#)

4.3.3. Directory Services

The following directory services are used:

- *NIS (being phased out)
- *LDAP (OpenLDAP has IPv6 support)
- *Active Directory
- *RADIUS (Our current RADIUS package has IPv6 support)

4.3.4. DNS

[TOC](#)

The three DNS servers are running BIND9. A DNS secondary is held at another UK university site. While we will make our three DNS servers dual-stack, our DNS secondary would remain IPv4-only since it is out of our administrative control.

4.3.5. NTP

[TOC](#)

The JANET NREN offers a stratum 1 NTP server. The department also has a GPS-based NTP server built-in to its own RIPE NCC test traffic server and an NTP device from a commercial provider. Both support IPv6 operation and transport.

4.3.6. Multicast

[TOC](#)

PIM-SM IPv4 multicast is facilitated via a dedicated commercial router, using a Rendezvous Point operated by our regional network. This supports applications including the IPv4 AccessGrid conferencing system. A number of bugs in the existing IPv4 routing equipment prevent heavy use of IPv4 Multicast within the department network (another reason that an IPv6 Multicast solution is desirable). An IPv4 Multicast beacon is used for monitoring Multicast. Our IPv6 multicast deployment plans are discussed in Section 3.8 above.

4.4. Hard-coded address points

[TOC](#)

Usage of IPv4 hard-coded addresses is interesting for at least two reasons. One is that it illustrates where IPv6 hard-coded addresses may appear, and thus secondly it is useful to analyse which hard-coded

addresses may be barriers to smooth IPv6 renumbering. A procedure for renumbering has been described in [Procedures for Renumbering an IPv6 Network without a Flag Day \(Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day," September 2005.\)](#) [RFC4192]. A non-exhaustive list of instances of such addresses includes:

- *Provider based prefix(es)
- *Names resolved to IP addresses in firewall at startup time
- *IP addresses in remote firewalls allowing access to remote services
- *IP-based authentication in remote systems allowing access to online bibliographic resources
- *IP address of both tunnel end points for IPv6 in IPv4 tunnel
- *Hard-coded IP subnet configuration information
- *IP addresses for static route targets
- *Blocked SMTP server IP list (spam sources)
- *Web .htaccess and remote access controls
- *Apache .Listen. directive on given IP address
- *Configured multicast rendezvous point
- *TCP wrapper files
- *Samba configuration files
- *DNS resolv.conf on Unix
- *Nocol monitoring tool
- *NIS/ypbind via the hosts file
- *Some interface configurations
- *Unix portmap security masks
- *NIS security masks

5. IPv6 Enterprise Deployment Procedure

In this section we document (retrospectively) the procedure we went through in deploying IPv6 within our campus site.

The work described in this document has also been fed into the [IPv6 Enterprise Analysis \(Bound, J., "IPv6 Enterprise Network Analysis - IP Layer 3 Focus," December 2006.\)](#) [I-D.ietf-v6ops-ent-analysis]. The reader is referred in particular to Section 4 ("Wide-Scale Dual-Stack Deployment") and Section 7 ("General issues and applicability for all Scenarios") which were directly contributed from the work here.

As described in the [IPv6 Enterprise Analysis \(Bound, J., "IPv6 Enterprise Network Analysis - IP Layer 3 Focus," December 2006.\)](#) [I-D.ietf-v6ops-ent-analysis] document, the scenario here is one of wide-scale dual-stack deployment. The plan for deployment follows the general guidelines of Section 7 of that document, though we have expanded that description here from subsequent experience.

Note that our analysis does not include issues relating to deployment of new IPv6-specific technology, e.g. [MIPv6 \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#) [RFC3775]. The focus of our deployment has been deploying dual-stack pervasively on the wire, with core network oriented services being IPv6 enabled.

5.1. Advanced Planning

[TOC](#)

A first phase for deployment includes the following actions.

- *Include IPv6 requirements in all future tenders. Consult to understand IPv6 specification requirements for tenders; this may prove particularly valuable where new IPv6 specific technology is desirable, e.g. Embedded-RP support for Multicast.
- *Identify your IPv6 ISP. This will most likely be your IPv4 ISP also, but in some cases it may not be.
- *Speak to your IPv6 ISP to acquire IPv6 address space (a /48 prefix) for your site; you will need this at some point, so may as well acquire the space sooner rather than later. This will include delegation of IPv6 forward and reverse DNS for your site. Our campus address space is a /48 prefix allocated by JANET.
- *Establish IPv6 training for operational staff, for administration of host and router platforms.
- *Investigate how to deploy basic IPv6 network services: DNS, routing, host configuration.
- *Encourage operational staff to get some IPv6 familiarity by trying IPv6 through services such as a public or ISP-supported

IPv6 [tunnel broker \(Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker," January 2001.\)](#) [RFC3053].

*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.

Following these action points should allow sites or networks to be ready for a trial or pilot IPv6 deployment, and to have confidence they understand and have control of their current - perhaps unwitting - usage of IPv6 (from systems which have it enabled by default).

5.2. Testbed/Trial Deployment

[TOC](#)

In this stage a site is validating IPv6 for deployment, and will thus take actions including the following:

*Assign and deploy an IPv6-capable router and (we recommend) a firewall or filtering device.

*Establish IPv6 connectivity to the IPv6 ISP. Sites might use a tunnelled service, or check for any native IPv6 offering. In our case, the connectivity is native IPv6 from JANET, via the regional MAN (using [6PE \(De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers \(6PE\)," February 2007.\)](#) [RFC4798]) and the campus (using a VLAN to carry IPv6 natively).

*Connect testbed host systems on the internal router interface, using one subnet prefix (a /64) from the site's allocated IPv6 / 48 prefix. At this stage your trial network may be standalone (disconnected from other internal networks) or, as we did, it may be that you dual-stack your existing IPv4 DMZ(s) for the pilot phase.

*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).

In parallel, other preparation can be undertaken for a production deployment:

*Survey systems, applications and services for IPv6 capability, including management, monitoring and access control devices and systems. The Enterprise Scenarios text as evaluated earlier in this document is a good basis to undertake this task from.

*Formulate an IPv6 address plan for your site/network. Our campus has allocated the department network a /56 prefix that can grow into a /52 prefix, i.e. the department can in theory create up to 256 IPv6 subnets initially. We discuss address planning issues in a separate document on IPv6 [addressing considerations \(Velde, G., Popoviciu, C., Chown, T., Bonness, O., and C. Hahn, "IPv6 Unicast Address Assignment Considerations," September 2008.\)](#) [I-D.ietf-v6ops-addcon].

*Discuss and document IPv6-related policies (e.g. use of Privacy Addresses, and of stateless or stateful address assignment).

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

5.3. Production Deployment

[TOC](#)

A production deployment includes the following action points:

*Plan which parts of the network will be IPv6-enabled first, and which existing subnets will be IPv6-enabled (made dual-stack). This may be certain server subnets, a DMZ, or a Wireless LAN network, for example.

*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.

*Enable IPv6, and IPv6 routing, such that IPv6 is on the wire, prior to host system activation. Ensure filtering and firewall policies are implemented as required.

*Add IPv6 address configuration to your DNS systems, and configure them to respond over IPv4 or IPv6 transport. Do not advertise AAAA records for a node until it is IPv6-reachable. Be aware that multiple services may run on a node, all of which may need to be IPv6-enabled before a AAAA record for the node is published.

*Deploy IPv6 support in appropriate management and monitoring tools.

*Enable IPv6 in selected production services and applications (e.g. Apache or IIS for web servers). In our case, we focused initially on DNS (bind), mail/MX (sendmail) and web services (Apache) in dual-stack mode.

*Include IPv6 transport in all ongoing security audit/penetration tests.

*Support IPv4-IPv6 interworking. As there are not (yet) any IPv6-only links on our site, interworking methods are not required. Should IPv6-only devices be deployed on the dual-stack infrastructure, we anticipate using proxy tools (web cache, SMTP relay, etc) to support their access to legacy IPv4 services, rather than deploying translation-based tools.

*Supporting remote users. These may connect via an IPv4 VPN and then use an IPv6 connection over that VPN, or use the remote IPv6 services of your ISP (e.g. our ISP, JANET, runs a tunnel broker and a 6to4 relay).

The depth of the IPv6 deployment may vary from site to site. By enabling key services you make your site ready for a fuller deployment, and gain confidence and experience in the technology, which is good for your support staff, your students, staff and researchers.

6. Analysis: Dual-Stack Deployment - Transition toolbox

[TOC](#)

Within our network we initially deployed IPv6 such that it was routed in parallel to IPv4, but with data running on the same end-host links, using a VLAN-based method as cited below. This allowed us to pilot IPv6 without risking adverse impact on our existing IPv4 platforms. This method was used for over two years. Towards the end of its use, the BSD platforms used to facilitate this were showing signs of strain under the load, in terms of pure unicast and multicast forwarding requirements under heavier traffic bursts. We have since deployed a unified IPv4 and IPv6 commercial routing platform from a single vendor, which met all our IPv4 and IPv6 procurement requirements for IPv4 and IPv6 unicast and multicast functions, including but not limited to:

*IPv6 unicast routing protocols;

*IPv6 multicast routing protocols (PIM-SM including SSM);

*IPv6 multicast Embedded-RP support;

*IPv6 multicast MLD(v1 and v2) snooping (see [RFC4541 \(Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol \(IGMP\) and Multicast Listener Discovery \(MLD\) Snooping Switches," May 2006.\)](#) [RFC4541]).

We have used the following mechanisms in our department's transition process:

- *[VLANs \(Chown, T., "Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks," June 2006.\)](#) [RFC4554] in an initial phase to distribute IPv6 connectivity over the non-dual-stack capable network infrastructure. This VLAN solution was an interim step until full dual protocol capable equipment was deployed during 2005;
- *A [Tunnel broker \(Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker," January 2001.\)](#) [RFC3053] for remote access, provided by our IPv6 ISP (JANET). We initially deployed our own tunnel broker, but now refer our home and roaming users to the JANET solution. This is only used for remote access to our network, not within our network;
- *A [6to4 \(Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds," February 2001.\)](#) [RFC3056] relay for remote access, provided by our IPv6 ISP. Again, we used to run our own relay, but the relay operated by our IPv6 ISP is perfectly adequate at this time for communicating with 6to4 sites. We do not believe 6to4 is an acceptable solution as a campus connectivity method (because we do not then use our own IPv6 address space as allocated by JANET, and 6to4 itself is prone to failure and [abuse \(Savola, P. and C. Patel, "Security Considerations for 6to4," December 2004.\)](#) [RFC3964]).

We do NOT currently see a requirement for:

- *[NAT-PT \(Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation \(NAT-PT\)," February 2000.\)](#) [RFC2766], because we are dual-stack with no IPv6-only networks (yet), and as we introduce such networks, or IPv6-only nodes in the dual-stack networks, we expect to use application layer gateways and proxies for legacy IPv4 access. Where dual-stack nodes may in future be used on IPv6-only links, the Dual Stack Transition Mechanism (DSTM) may be of value, but preference would be given to use IPv6 transport where possible;
- *[ISATAP \(Templin, F., Gleeson, T., Talwar, M., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol \(ISATAP\)," October 2005.\)](#) [RFC4214], because we prefer to use a structured internal IPv6 deployment, and are doing so in a pervasive fashion (i.e. not as a sparse deployment). ISATAP may be useful for sparse deployment of IPv6 in sites who are happy to IPv6 pilot in a less structured fashion. We do not wish to see arbitrary automatic tunnels being used between links on our network;

*[Teredo \(Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations \(NATs\)," February 2006.\) \[RFC4380\]](#) - we considered deploying servers/relays to support home users behind NATs, but chose not to do so since our ISP's tunnel broker service supports NAT traversal, and we feel it offers better management and monitoring facilities. This decision may be reviewed if we see a rise in demand for Teredo service support.

7. Analysis: Considerations beyond the Scenarios Document

[TOC](#)

Here we mention issues or scenario components that were not explicitly listed in the IPv6 Enterprise Network Scenarios document. Due to the scope, that document could not embrace all details. We mention here components that other sites may also wish to consider:

*Support for WLAN and other access control. Most sites tend to use a web-redirection portal to authenticate users, but these invariably do not detect or support use of IPv6. One solution is to use 802.1x which is IP-agnostic as a Layer 2 port control mechanism.

*Consideration for hard-coded addresses.

The brevity of this list shows that the IPv6 Enterprise Network Scenarios text includes very good coverage of the issues and considerations faced by our enterprise site.

8. Summary

[TOC](#)

In this document we have analysed the specific campus transition scenario for the author's site, and reported the analysis for the benefit of others who may be in a similar scenario, or for whom parts of the scenario may be relevant.

In our case transition does not mean from IPv4 only to IPv6 only, rather from IPv4 only to a dual-stack environment that could support IPv6 only nodes at a later date. We would probably best describe the process as dual stack integration.

We have described how a phased approach to transition can be adopted at a campus site (or part thereof), from a planning stage through a pilot to a fuller deployment. During our transition we initially ran a parallel IPv6 routing infrastructure, then in 2005 unified the routing to a single platform, for unicast and multicast IPv6. We enabled key services for dual-stack operation from the outset (DNS, web and mail/MX) and have enabled other services as and when they have become

available. The VLAN-based interim step was useful for two years until a dual-protocol routing solution could be procured.

We do not discuss detailed availability of IPv6 capability in the services described in Section 4 above in this text, and we leave application support as an issue out of scope (though we observe that open source support for IPv6 is in general very good). For the purposes of our network-oriented transition, we are happy that the path taken and current solution is stable and complete.

The deployment has now been in full dual-stack operation for over two years, with key services enabled (including public-facing DNS, SMTP, web) without any significant adverse effects on the IPv4 service. The author welcomes discussion with other sites that are undergoing or have undergone a similar transition or integration process.

9. Acknowledgements

[TOC](#)

Discussions with fellow participants on the 6NET and Euro6IX projects have been valuable. Input from the IETF IPv6 Operations WG list have also been welcomed.

10. IANA Considerations

[TOC](#)

The document contains no IANA considerations.

11. Security Considerations

[TOC](#)

There are no specific new considerations from this scenario description and analysis.

12. Informative References

[TOC](#)

- [RFC2080] [Malkin, G. and R. Minnear, "RIPng for IPv6,"](#) RFC 2080, January 1997 ([TXT](#)).
- [RFC2453] [Malkin, G., "RIP Version 2,"](#) STD 56, RFC 2453, November 1998 ([TXT](#), [HTML](#), [XML](#)).
- [RFC2462] [Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration,"](#) RFC 2462, December 1998 ([TXT](#), [HTML](#), [XML](#)).
- [RFC2766] [Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation \(NAT-PT\),"](#) RFC 2766, February 2000 ([TXT](#)).
- [RFC3041]

- Narten, T. and R. Draves, "[Privacy Extensions for Stateless Address Autoconfiguration in IPv6](#)," RFC 3041, January 2001 ([TXT](#)).
- [RFC3053] Durand, A., Fasano, P., Guardini, I., and D. Lento, "[IPv6 Tunnel Broker](#)," RFC 3053, January 2001 ([TXT](#)).
- [RFC3056] Carpenter, B. and K. Moore, "[Connection of IPv6 Domains via IPv4 Clouds](#)," RFC 3056, February 2001 ([TXT](#)).
- [RFC3171] Albanna, Z., Almeroth, K., Meyer, D., and M. Schipper, "[IANA Guidelines for IPv4 Multicast Address Assignments](#)," RFC 3171, August 2001 ([TXT](#)).
- [RFC3306] Haberman, B. and D. Thaler, "[Unicast-Prefix-based IPv6 Multicast Addresses](#)," RFC 3306, August 2002 ([TXT](#)).
- [RFC3307] Haberman, B., "[Allocation Guidelines for IPv6 Multicast Addresses](#)," RFC 3307, August 2002 ([TXT](#)).
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "[Dynamic Host Configuration Protocol for IPv6 \(DHCPv6\)](#)," RFC 3315, July 2003 ([TXT](#)).
- [RFC3484] Draves, R., "[Default Address Selection for Internet Protocol version 6 \(IPv6\)](#)," RFC 3484, February 2003 ([TXT](#)).
- [RFC3542] Stevens, W., Thomas, M., Nordmark, E., and T. Jinmei, "[Advanced Sockets Application Program Interface \(API\) for IPv6](#)," RFC 3542, May 2003 ([TXT](#)).
- [RFC3618] Fenner, B. and D. Meyer, "[Multicast Source Discovery Protocol \(MSDP\)](#)," RFC 3618, October 2003 ([TXT](#)).
- [RFC3736] Droms, R., "[Stateless Dynamic Host Configuration Protocol \(DHCP\) Service for IPv6](#)," RFC 3736, April 2004 ([TXT](#)).
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "[Mobility Support in IPv6](#)," RFC 3775, June 2004 ([TXT](#)).
- [RFC3904] Huitema, C., Austein, R., Satapati, S., and R. van der Pol, "[Evaluation of IPv6 Transition Mechanisms for Unmanaged Networks](#)," RFC 3904, September 2004 ([TXT](#)).
- [RFC3956] Savola, P. and B. Haberman, "[Embedding the Rendezvous Point \(RP\) Address in an IPv6 Multicast Address](#)," RFC 3956, November 2004 ([TXT](#)).
- [RFC3964] Savola, P. and C. Patel, "[Security Considerations for 6to4](#)," RFC 3964, December 2004 ([TXT](#)).
- [RFC3974] Nakamura, M. and J. Hagino, "[SMTP Operational Experience in Mixed IPv4/v6 Environments](#)," RFC 3974, January 2005 ([TXT](#)).
- [RFC4029] Lind, M., Ksinant, V., Park, S., Baudot, A., and P. Savola, "[Scenarios and Analysis for Introducing IPv6 into ISP Networks](#)," RFC 4029, March 2005 ([TXT](#)).
- [RFC4038]

- Shin, M-K., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "[Application Aspects of IPv6 Transition](#)," RFC 4038, March 2005 ([TXT](#)).
- [RFC4057] Bound, J., "[IPv6 Enterprise Network Scenarios](#)," RFC 4057, June 2005 ([TXT](#)).
- [RFC4192] Baker, F., Lear, E., and R. Droms, "[Procedures for Renumbering an IPv6 Network without a Flag Day](#)," RFC 4192, September 2005 ([TXT](#)).
- [RFC4193] Hinden, R. and B. Haberman, "[Unique Local IPv6 Unicast Addresses](#)," RFC 4193, October 2005 ([TXT](#)).
- [RFC4214] Templin, F., Gleeson, T., Talwar, M., and D. Thaler, "[Intra-Site Automatic Tunnel Addressing Protocol \(ISATAP\)](#)," RFC 4214, October 2005 ([TXT](#)).
- [RFC4215] Wiljakka, J., "[Analysis on IPv6 Transition in Third Generation Partnership Project \(3GPP\) Networks](#)," RFC 4215, October 2005 ([TXT](#)).
- [RFC4291] Hinden, R. and S. Deering, "[IP Version 6 Addressing Architecture](#)," RFC 4291, February 2006 ([TXT](#)).
- [RFC4380] Huitema, C., "[Teredo: Tunneling IPv6 over UDP through Network Address Translations \(NATs\)](#)," RFC 4380, February 2006 ([TXT](#)).
- [RFC4477] Chown, T., Venaas, S., and C. Strauf, "[Dynamic Host Configuration Protocol \(DHCP\): IPv4 and IPv6 Dual-Stack Issues](#)," RFC 4477, May 2006 ([TXT](#)).
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "[Considerations for Internet Group Management Protocol \(IGMP\) and Multicast Listener Discovery \(MLD\) Snooping Switches](#)," RFC 4541, May 2006 ([TXT](#)).
- [RFC4554] Chown, T., "[Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks](#)," RFC 4554, June 2006 ([TXT](#)).
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "[Protocol Independent Multicast - Sparse Mode \(PIM-SM\): Protocol Specification \(Revised\)](#)," RFC 4601, August 2006 ([TXT](#), [PDF](#)).
- [RFC4632] Fuller, V. and T. Li, "[Classless Inter-domain Routing \(CIDR\): The Internet Address Assignment and Aggregation Plan](#)," BCP 122, RFC 4632, August 2006 ([TXT](#)).
- [RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "[Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers \(6PE\)](#)," RFC 4798, February 2007 ([TXT](#)).
- [I-D.ietf-v6ops-addcon] Velde, G., Popoviciu, C., Chown, T., Bonness, O., and C. Hahn, "[IPv6 Unicast Address Assignment Considerations](#)," draft-ietf-v6ops-addcon-10 (work in progress), September 2008 ([TXT](#)).
- [I-D.ietf-v6ops- Davies, E., "[IPv6 Transition/Co-existence Security Considerations](#)," draft-ietf-v6ops-security-overview-06 (work in progress), October 2006 ([TXT](#)).

security-
overview]
[I-D.ietf-
v6ops-ent-
analysis]

Bound, J., "[IPv6 Enterprise Network Analysis - IP Layer 3 Focus](#)," draft-ietf-v6ops-ent-analysis-07 (work in progress), December 2006 ([TXT](#)).

Author's Address

[TOC](#)

Tim Chown
University of Southampton
School of Electronics and Computer Science
Southampton, Hampshire SO17 1BJ
United Kingdom
Email: tjc@ecs.soton.ac.uk

Full Copyright Statement

[TOC](#)

Copyright © The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this

standard. Please address the information to the IETF at ietf-ipr@ietf.org.