

V60PS
Internet-Draft
Intended status:
Informational
Expires: October 15,
2010

B. Carpenter
Univ. of Auckland
S. Jiang
Huawei Technologies
Co., Ltd
April 13, 2010

[TOC](#)

Emerging Service Provider Scenarios for IPv6 Deployment draft-carpenter-v6ops-isp-scenarios-02

Abstract

This document describes practices and plans that are emerging among Internet Service Providers for the deployment of IPv6. They are based on practical experience so far, as well as current plans and requirements, reported in a survey carried out in early 2010.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on October 15, 2010.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Survey of ISP experience, plans and requirements](#)
- [3. Lessons from experience and planning](#)
- [4. Gap analysis](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [7. Acknowledgements](#)
- [8. Change log](#)
- [9. Informative References](#)
- [Appendix A. Summary of replies](#)
- [Appendix B. Questionnaire](#)
- [§ Authors' Addresses](#)

1. Introduction

[TOC](#)

As is well known, the approaching exhaustion of IPv4 address space will bring about a situation in which Internet Service Providers (ISPs) are faced with a choice between one or more of three major alternatives:

1. Squeeze the use of IPv4 addresses even harder than today, using smaller and smaller address blocks per customer, and possibly trading address blocks with other ISPs.
2. Install multiple layers of network address translation [\[I-D.nishitani-cgn\]](#) (Yamagata, I., Nishitani, T., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common requirements for IP address sharing schemes," March 2010.), or share IPv4 addresses by other methods such as address-plus-port mapping [\[I-D.ymbk-aplusp\]](#) (Bush, R., "The A+P Approach to the IPv4 Address Shortage," October 2009.), [\[I-D.boucadair-port-range\]](#) (Boucadair, M., Levis, P., Bajko, G., and T. Savolainen, "IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion: Port Range based IP Architecture," July 2009.).
3. Deploy IPv6, and operate IPv4-IPv6 coexistence and interworking mechanisms.

This document focuses on alternative (3), while recognizing that many ISPs may be obliged by circumstances to prolong the life of IPv4 by using (1) or (2) while preparing for (3).

The document describes IPv6 deployment scenarios already adopted or currently planned by a set of ISPs who responded to a technical questionnaire. Thus, it is a factual record of the responses from those ISPs. It makes no recommendations; the best choice of scenarios will depend on the circumstances of individual ISPs.

We consider various aspects of IPv6 deployment: addressing, routing, DNS, management and of course IPv4-IPv6 coexistence and interworking. We do not consider application services in detail, but we do cover general aspects.

The reader is assumed to be familiar with IPv6. The IETF's view of core IPv6 requirements is to be found in [\[RFC4294\] \(Loughney, J., "IPv6 Node Requirements," April 2006.\)](#) (currently being updated as [\[I-D.ietf-6man-node-req-bis\] \(Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements RFC 4294-bis," March 2010.\)](#)). However, this does not give a complete view of mechanisms an ISP may need to deploy, since it considers the requirements for an individual node, not for a network as a whole.

[\[RFC4029\] \(Lind, M., Ksinant, V., Park, S., Baudot, A., and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks," March 2005.\)](#) discusses scenarios for introducing IPv6 into ISP networks, as the problem was viewed some years ago. The end goal described in RFC 4029 is simply a dual-stack ISP backbone. Today's view is that this is insufficient, as it does not allow for interworking between IPv6-only and legacy (IPv4-only) hosts. Indeed, the end goal today might be an IPv6-only ISP backbone, with some form of legacy IPv4 support.

[\[RFC4779\] \(Asadullah, S., Ahmed, A., Popoviciu, C., Savola, P., and J. Palet, "ISP IPv6 Deployment Scenarios in Broadband Access Networks," January 2007.\)](#) discusses deployment in broadband access networks such as CATV, ADSL and wireless. [\[RFC5181\] \(Shin, M-K., Han, Y-H., Kim, S-E., and D. Premec, "IPv6 Deployment Scenarios in 802.16 Networks," May 2008.\)](#), [\[RFC5121\] \(Patil, B., Xia, F., Sarikaya, B., Choi, JH., and S. Madanapalli, "Transmission of IPv6 via the IPv6 Convergence Sublayer over IEEE 802.16 Networks," February 2008.\)](#) and [\[I-D.ietf-16ng-ip-over-ethernet-over-802-dot-16\] \(Jeon, H., Riegel, M., and S. Jeong, "Transmission of IP over Ethernet over IEEE 802.16 Networks," September 2009.\)](#) deal specifically with IEEE 802.16 access networks. MPLS-based ISPs may use the 6PE [\[RFC4798\] \(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.\)](#) mechanism.

[\[RFC4942\] \(Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/Co-existence Security Considerations," September 2007.\)](#) covers IPv6 security issues, especially those that are specific to transition and coexistence scenarios. Also related to security, [\[RFC4864\] \(Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6," May 2007.\)](#) discusses "Local Network Protection", i.e., how the internal structure of an IPv6 site network can be protected. Although not directly part of ISP operations, this topic does affect the issue of how well an ISP's customers are protected after they deploy IPv6.

Although the basic IPv6 standards have long been stable, it should be noted that considerable work continues in the IETF, particularly to resolve the issue of highly scalable multihoming support for IPv6

sites, and to resolve the problem of IP layer interworking between IPv6-only and IPv4-only hosts. IPv6/IPv4 interworking at the application layers is handled within the original dual-stack model of IPv6 deployment: either one end of an application session will have dual-stack connectivity, or a dual-stack intermediary such as an HTTP proxy or SMTP server will interface to both IPv4-only and IPv6-only hosts.

[\[RFC5211\] \(Curran, J., "An Internet Transition Plan," July 2008.\)](#)

describes an independent view of a possible sequence of events for IPv6 adoption in the Internet as a whole, with direct implications for ISPs. Its main point, perhaps, is that by 2012 it will be necessary to regard IPv4 networks as the legacy solution.

2. Survey of ISP experience, plans and requirements

[TOC](#)

To obtain a view of the IPv6 experience, plans and requirements of ISPs, a questionnaire was issued by the authors in December 2009 and announced widely to the operational community. We promised to keep the replies strictly confidential and to publish only combined results, without identifying information about individual ISPs in any published results. The raw technical questions are shown in [Appendix B \(Questionnaire\)](#), and a detailed summary of the replies is in [Appendix A \(Summary of replies\)](#). Note that although the questionnaire was widely announced, those who chose to reply were self-selected and we can make no claim of statistical significance or freedom from bias in the results. In particular, we assume that ISPs with a pre-existing interest in IPv6 are more likely to have replied than others. The results should therefore be interpreted with some care.

Thirty ISPs replied before the cutoff date for this analysis. 66% of responses were from European ISPs but large operators in North America and Asia/Pacific regions are included. Commercial ISPs operating nationally predominate, with a vast range of size (from 30 customers up to 40 million). Note that some providers chose not to answer about the exact number of customers. Nevertheless, it can be stated that 6 providers each had millions of customers, the next 9 each had more than 10,000 customers, and the remaining 15 each had fewer than 10,000 customers.

80% of the respondents offer both transit and origin-only service; 27% offer IP multicast service; 80% have multihomed customers. A very wide variety of access technologies is used: xDSL, DOCSIS, leased line (X.25, TDM/PDH, SDH), frame relay, dialup, microwave, FTTP, CDMA, UMTS, LTE, WiMAX, BWA, WiFi, Ethernet (100M-10G), MetroEthernet/MPLS. Most ISPs provide CPE to some or all of their customers, but these CPE are often unable to support IPv6.

Estimates of when ISPs will run out of public IPv4 address space for internal use vary widely, between "now" and "never". Public IPv4 address space for customers is mainly expected to run out between 2010

and 2015, but three or four ISPs suggested it will never happen. About 20% of ISPs offer RFC 1918 space to customers, and about 40% use such addresses internally.

60% of ISPs report that some big customers are requesting IPv6.

Predictions for when 10% of customers will require IPv6 range from 2010 to 2017, and for 50% from 2011 to 2020. These ISPs require IPv6 to be a standard service by 2010 to 2015; the most common target date is 2011. 40% already offer IPv6 as a regular service, although in general it is used by fewer than 1% of customers. Another 47% of ISPs have IPv6 deployment in progress or planned. These all plan at least beta-test service in 2010. Planned dates for regular service are between 2010 and 2013 (except for one ISP who replied that it depends on the marketing department). When asked when IPv6 will reach 50% of total traffic, the most common answer is 2015.

Turning to technology choices, the overwhelming choice of approach (93%) is a dual stack routing backbone, and the reason given is simplicity and cost. 40% run, or plan to run, a 6to4 relay as well, and 17% run or plan a Teredo server. However, 77% of ISPs don't have or plan any devices dedicated to IPv6. A different 77% don't see IPv6 as an opportunity to restructure their network topology.

When asked which types of equipment are unable to support IPv6, the most common answer was CPE (9 mentions). Also mentioned: handsets; DSLAMs; routers (including several specific models); traffic management boxes; load balancers; VPN boxes; management interfaces & systems; firewalls; billing systems. When asked if such devices can be field-upgraded, the answers were gloomy: 5 yes, 4 partially, 10 no, and numerous "don't know" or "hopefully".

83% support or plan DNS AAAA queries over IPv6, and all but one of these include reverse DNS lookup for IPv6.

The ISPs have prefixes ranging from /19 to /48, and have a variety of policies for customer prefixes. Fifteen ISPs offer more than one of /48, /52, /56, /60 or /64. Two offer /56 only, seven offer /48 only. Two wired operators offer /64 only. Mobile operators offer /64 in accordance with 3GPP, but at least one would like to be allowed to offer /128 or /126. Also, as many as 30% of the operators already have IPv6 customers preferring a PI prefix. The methods chosen for prefix delegation to CPEs are manual, DHCPv6[-PD], SLAAC, RADIUS, and PPOE. About 50% of ISPs already operate or plan dual-stack SMTP, POP3, IMAP and HTTP services. In terms of internal services, it seems that firewalls, intrusion detection, address management, monitoring, and network management tools are also around the 50% mark. However, accounting and billing software is only ready at 23% of ISPs.

Considering IPv4-IPv6 interworking, 57% of ISPs don't expect to have IPv6-only customers (but mobile operators are certain they will have millions). Five ISPs report customers who explicitly refused to consider IPv6. When asked how long customers will run IPv4-only applications, the most frequent answer is "more than ten years". Only three ISPs state that IPv6-IPv4 interworking at the IP layer is not needed. On the other hand, only three (a different three!) run or plan

to run NAT-PT. At least 30% plan to run some kind of translator (presumably NAT64/DNS64), but only one felt that a multicast translator was essential. Among those who do not plan a translator, when asked how they plan to connect IPv6-only customers to IPv4-only services, seven rely on dual stack and three have no plan (one said, paraphrasing, "it's their problem").

Asked about plans for Mobile IPv6 (or Nemo mobile networks), 20% said yes, and 70% said no, with the others uncertain. A detailed analysis shows that of the six ISPs who responded positively, three are large mobile operators (and a fourth mobile operator said "No"). The other three who were positive were broadband ISPs ranging from small to very large.

We examined the data to see whether any other differences emerge between the very large (millions of customers), medium (at least 10,000), and small (fewer than 10,000) operators. However, the range of answers seems to be broadly similar in all cases. The major exception is that among the six very large operators, one plans to use 6PE and DS-lite, and another to use 6VPE, instead of a simple dual stack. The other large operators and all the medium and small operators prefer a simple dual stack.

3. Lessons from experience and planning

[TOC](#)

This section is assembled and paraphrased from general comments made in the various questionnaire responses. Any inconsistencies or contradictions are present in the original data. Comments related to missing features and products have been included in [Section 4 \(Gap analysis\)](#).

As noted in the summary above, the ISPs that responded overwhelmingly prefer a native dual stack deployment. Numerous comments mentioned the simplicity of this model and the complexity and sub-optimal routing of tunnel-based solutions. Topology redesign is not generally considered desirable, because congruent IPv4 and IPv6 topology simplifies maintenance and engineering. Nevertheless, 6to4 is considered convenient for cable modem (DOCSIS) users and 6RD is considered an attractive model by some. Various operators, but by no means all, also see a need for Teredo. One very large MPLS-based operator prefers 6PE because it avoids an IPv6 IGP and reduces operational costs. This operator also sees IPv4 scarcity addressed by DS-lite [\[I-D.ietf-softwire-dual-stack-lite\]](#) (Durand, A., Droms, R., Haberman, B., Woodyatt, J., Lee, Y., and R. Bush, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion," March 2010.) and Carrier Grade NAT, also acting as a catalyst for IPv6. Another very large operator with an existing NAT44 infrastructure plans to use 6VPE [\[RFC4659\]](#) (De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN," September 2006.)

and believes that NAT64 will be similar to NAT44 in support requirements.

Several ISPs observe that IPv6 deployment is technically not hard. The most eloquent statement: "Just do it, bit by bit. It is very much an 'eating the elephant' problem, but at one mouthful at a time, it appears to be surprisingly easy." Other comments paraphrased from the replies are:

- *Despite the known gaps, the tools and toolkits are fairly mature at this point.

- *Managerial commitment and a systematic approach to analysing requirements and readiness are essential.

- *Evangelization remains a must, as it seems that many ISP and IT managers are still unaware of the need for an IPv6 plan, and are inclined to dismiss IPv4 depletion as a false alarm, and also seem unaware that NATs create expensive support requirements.

- *Without customers wanting IPv6, getting business backing is very hard, and IPv6 security and scale was not a focus for vendors until very recently.

- *Operators lack real experience with customer usage of IPv6, and the resulting lack of confidence causes delay.

Three further quotations are of interest:

"We are planning to move all our management addressing from IPv4 to IPv6 to free up IPv4 addresses."

"I am actively pushing our larger customers to start testing IPv6 as our development has pretty much stopped as we need some real world testing to be done."

"Customer support needs to be aware that IPv6 is being started in your network, or servers. We experienced many IPv6 blocking applications, applications that do not fall back to IPv4, etc. The most difficult part may be to get engineers, sales, customer support personnel to like IPv6."

4. Gap analysis

[TOC](#)

The survey has shown a certain number of desirable features to be missing, either in relevant specifications, or in many products. This section summarizes those gaps.

As noted above, numerous models of various types of product still do not support IPv6:

CPE

handsets

DSLAMs

routers

traffic management boxes

load balancers

VPN boxes

other appliances

management interfaces and systems

firewalls

intrusion detection systems

accounting and billing systems

It is not the purpose of this document to name and shame vendors, but it is today becoming urgent for all such products to avoid becoming part of the IPv4 legacy. ISPs stated that they want consistent feature-equal support for IPv4 and IPv6 in all equipment and software at reasonable or no extra cost. The problems can be quite subtle: for example, one CPE with "full" IPv6 support does not support IPv6 traffic shaping, so the ISP cannot cap IPv6 traffic to contractual limits. Other needs and issues mentioned:

*A specific CPE need is an intelligent prefix sub-delegation mechanism (ease of use issue).

*Full RA support on LAN side of ADSL CPE.

*PPPoE and RADIUS support. Specifically, global unicast address assignment for L2TP/PPPoE. Currently the PPPoE client will be assigned a link-local address, requiring a second step (DHCPv6 or SLAAC).

*Simple automatic distribution of DNS server details is essential; a DNS server option in RA [\[RFC5006\] \(Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Option for DNS Configuration," September 2007.\)](#) is wanted.

*ISPs need fully featured DHCPv6, especially since SLAAC does not allow settings to be pushed out (except for RFC 5006).

*Firewall systems should not use separate IPv4 and IPv6 rules.

- *Customer side firewalls/routers which can do 25-100 Mbit/s.
- *Gaps in infrastructure security for IPv6 management.
- *Gaps in routers' treatment of header options.
- *RA-Guard in switches.
- *VRRP support.
- *PE-CE routing protocols (OSPF and IS-IS).
- *Problems using a single BGP session to exchange routes for both IPv4 and IPv6.
- *IPv6 support in all the best open source tools.

Several ISPs also noted that much commercial applications software does not correctly support IPv6 and that this will cause customer problems. Also, some operating systems are still shipped with IPv6 disabled by default, or with features such as IPv4-mapped addresses disabled by default.

Numerous ISPs want a scaleable NAT64/DNS64 product. Other protocol-related needs are:

- *"Getting it right" so that a dual stack client doesn't end up trying to use the wrong transport to reach a site.
- *User-side port allocation mechanisms. UPnP IGD and NAT-PMP can be used for IPv4, but nothing exists for IPv6 (yet). [Editor's comment: even though port mapping is in principle unnecessary for IPv6, a method of opening ports through firewalls on demand seems necessary.]

Global IPv6 connectivity still has issues; for example ISPs in the Pacific region may have to obtain IPv6 transit via the USA (a situation faced by IPv4 operators in Europe about twenty years ago). Also, there are persistent PMTUD issues in various places on the net, and a lack of multicast peering.

Finally, there was a comment that real deployment case studies must be shown to operators, along with multihoming and traffic engineering best practices.

5. Security Considerations

[TOC](#)

As a report on a survey, this document creates no security issues in itself. ISPs did not register any general concerns about IPv6 security. However, we note that about half of all firewall and intrusion

detection products are still reported not to support IPv6. Some ISPs expressed concern about firewall performance and about the need for separate firewall configurations for IPv4 and IPv6.

6. IANA Considerations

[TOC](#)

This document makes no request of the IANA.

7. Acknowledgements

[TOC](#)

We are grateful to all those who answered the questionnaire: Pete Barnwell, Cameron Byrne, Gareth Campling, David Freedman, Wesley George, Steinar Haug, Paul Hoogsteder, Mario Iseli, Christian Jacquenet, Kurt Jaeger, Seiichi Kawamura, Adrian Kennard, Simon Leinen, Riccardo Loselli, Janos Mohacsi, Jon Morby, Michael Newbery, Barry O'Donovan, Al Pooley, Antonio Querubin, Anthony Ryan, Marc Schaeffer, Valeriu Vraciu, Bill Walker and those who preferred to remain anonymous.

The ISPs willing to be named were: AISP, Alphanet, Breedband Delft, Claranet, E4A, Fidonet, Finecom, France Telecom, Hungarnet, Imagine, LavaNet, NEC BIGLOBE, Nepustilnet, Net North West, RoEduNet, SWITCH, Snap, Sprint, Star Technology, T-Mobile USA, Ventelo, and Whole Ltd. Useful comments and contributions were also made by Mohamed Boucadair, and others.

This document was produced using the xml2rfc tool [\[RFC2629\]](#) (Rose, M., "Writing I-Ds and RFCs using XML," June 1999.).

8. Change log

[TOC](#)

draft-carpenter-v6ops-isp-scenarios-02: updated summary and discussion of questionnaire results, deleted material to repurpose document as survey results only, 2010-04-13

draft-carpenter-v6ops-isp-scenarios-01: added summary and discussion of questionnaire results, 2010-02-23

draft-carpenter-v6ops-isp-scenarios-00: original version, 2009-10-13

9. Informative References

[TOC](#)

[I-D.boucadair-port-range] Boucadair, M., Levis, P., Bajko, G., and T. Savolainen, "[IPv4 Connectivity Access in the Context of IPv4 Address Exhaustion: Port Range based IP Architecture](#)," draft-boucadair-port-range-02 (work in progress), July 2009 ([TXT](#)).

[I-D.ietf-16ng-ip-over-ethernet-over-802-dot-16] Jeon, H., Riegel, M., and S. Jeong, "[Transmission of IP over Ethernet over IEEE 802.16 Networks](#)," draft-ietf-16ng-ip-over-ethernet-over-802-dot-16-12 (work in progress), September 2009 ([TXT](#)).

[I-D.ietf-6man-node-req-bis] Jankiewicz, E., Loughney, J., and T. Narten, "[IPv6 Node Requirements RFC 4294-bis](#)," draft-ietf-6man-node-req-bis-04 (work in progress), March 2010 ([TXT](#)).

[I-D.ietf-softwire-dual-stack-lite] Durand, A., Droms, R., Haberman, B., Woodyatt, J., Lee, Y., and R. Bush, "[Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion](#)," draft-ietf-softwire-dual-stack-lite-04 (work in progress), March 2010 ([TXT](#)).

[I-D.ietf-v6ops-cpe-simple-security] Woodyatt, J., "[Recommended Simple Security Capabilities in Customer Premises Equipment for Providing Residential IPv6 Internet Service](#)," draft-ietf-v6ops-cpe-simple-security-11 (work in progress), April 2010 ([TXT](#)).

[I-D.ietf-v6ops-ipv6-cpe-router] Singh, H., Beebee, W., Donley, C., Stark, B., and O. Troan, "[Basic Requirements for IPv6 Customer Edge Routers](#)," draft-ietf-v6ops-ipv6-cpe-router-04 (work in progress), January 2010 ([TXT](#)).

[I-D.nishitani-cgn] Yamagata, I., Nishitani, T., Miyakawa, S., Nakagawa, A., and H. Ashida, "[Common requirements for IP address sharing schemes](#)," draft-nishitani-cgn-04 (work in progress), March 2010 ([TXT](#)).

[I-D.ymbk-aplusp] Bush, R., "[The A+P Approach to the IPv4 Address Shortage](#)," draft-ymbk-aplusp-05 (work in progress), October 2009 ([TXT](#)).

[RFC2629] [Rose, M.](#), "[Writing I-Ds and RFCs using XML](#)," RFC 2629, June 1999 ([TXT](#), [HTML](#), [XML](#)).

[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](#)).

[RFC4294] Loughney, J., "[IPv6 Node Requirements](#)," RFC 4294, April 2006 ([TXT](#)).

[RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "[BGP-MPLS IP Virtual Private Network \(VPN\) Extension for IPv6 VPN](#)," RFC 4659, September 2006 ([TXT](#)).

[RFC4779] Asadullah, S., Ahmed, A., Popoviciu, C., Savola, P., and J. Palet, "[ISP IPv6 Deployment Scenarios in Broadband Access Networks](#)," RFC 4779, January 2007 ([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](#)).
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "[Local Network Protection for IPv6,](#)" RFC 4864, May 2007 ([TXT](#)).
- [RFC4942] Davies, E., Krishnan, S., and P. Savola, "[IPv6 Transition/Co-existence Security Considerations,](#)" RFC 4942, September 2007 ([TXT](#)).
- [RFC5006] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "[IPv6 Router Advertisement Option for DNS Configuration,](#)" RFC 5006, September 2007 ([TXT](#)).
- [RFC5121] Patil, B., Xia, F., Sarikaya, B., Choi, JH., and S. Madanapalli, "[Transmission of IPv6 via the IPv6 Convergence Sublayer over IEEE 802.16 Networks,](#)" RFC 5121, February 2008 ([TXT](#)).
- [RFC5181] Shin, M-K., Han, Y-H., Kim, S-E., and D. Premec, "[IPv6 Deployment Scenarios in 802.16 Networks,](#)" RFC 5181, May 2008 ([TXT](#)).
- [RFC5211] Curran, J., "[An Internet Transition Plan,](#)" RFC 5211, July 2008 ([TXT](#)).

Appendix A. Summary of replies

[TOC](#)

This summarises the 30 replies received by February 16, 2010. Note that the answers to some questions do not total to 30, due to missing answers or to multiple choices in some cases. The ISPs are distributed as follows:

Europe: 20

North America: 6

Asia/Pacific: 4

They can additionally be classified as:

Commercial: 26

Academic/research: 4

Operating internationally: 5

Operating nationally: 25

Basic data about IP service offerings:

*Offering both origin-only and transit service: 24

*Offering no transit: 6

*Number of private/small office customers (one IPv4 address): a few with zero, then from 30 units up to 40 million

*Number of corporate customers (block of IPv4 addresses): a few with zero, then from 40 units up to 40000

*IP multicast service? 8 yes, 22 no

*Do any customers require multihoming to multiple ISPs? 24 yes, 6 no

*Access technologies used: xDSL, DOCSIS, leased line (X.25, TDM/PDH, SDH), frame relay, dialup, microwave, FTTP, CDMA, UMTS, LTE, WiMAX, BWA, WiFi, Ethernet (100M-10G), Ether/MPLS, IPv6-in-IPv4 tunnels

Customer Premises Equipment:

*Do customers use CPE that ISP supplies? 26 yes (20% to 100% of customers), 4 no

*Does the CPE provided support native IPv6? 16 yes (some), 10 no

*(Note that these numbers include answers that apply to tens of millions of mobile handsets.)

IPv4 Address Space:

*When do you expect to run out of public IPv4 address space inside your own network? Estimates run from "now" to 2020, but 4 ISPs say "never" or "unforeseeable".

*Do you run RFC1918 addresses and NAT within your network? 9 yes; 17 no; 3 others say yes, but only for equipment management or L3VPN.

*What % of IPv4 space is needed for ISP use (not for customers)? 1% to 30% (and 100% for NRENs with PI customers).

*When do you expect to run out of public IPv4 address space for customers? Answers range from 2010 to 2015; 4 ISPs say it depends on their registry. 3 or 4 give answers suggesting it will never happen.

*Do you offer RFC1918 addresses to customers? 6 yes, 23 no

IPv6 Requirements:

*Are some big customers requesting IPv6? 18 yes, 12 no

*When do you predict 10% and 50% of customers to require IPv6 service? Ignoring unclear answers, answers for 10% range from 2010 to 2017, and for 50% from 2011 to 2020.

*When do you require IPv6 to be a standard service available to all customers? Answers range from 2010 to 2015; the most common answer is 2011.

*When do you predict IPv6 traffic to reach 50% of total traffic? Answers range from 2011 to 2020; the most common answer is 2015.

IPv6 Status and Plans:

*Do you currently offer IPv6 as a regular service? 12 yes, 5 partial, 12 no

*What % of customers currently use IPv6? <1% to 30%; mostly 0 or <1%

*When do you plan to start IPv6 deployment? 11 complete, 12 in progress, 3 in plan, 4 have no plan

*When do you plan to offer IPv6 as a special or beta-test service? For all those in progress or with a plan, the answer was either "now" or during 2010.

*When do you plan to offer IPv6 as a regular service to all customers? For all those in progress or with a plan, the answer was between "now" and 2013 (except for one ISP who replied that it depends on the marketing department).

IPv6 Technologies. Note the answers refer to actual deployment or to plans, as the case may be:

*Which basic IPv6 access method(s) apply?

-Dual stack routing backbone: 28 yes, 1 maybe

-Separate IPv4 and IPv6 backbones: 2 yes, 1 maybe

-6to4 relay: 12 yes

-Teredo server: 5 yes

-Tunnel broker: Unfortunately this question was unclear and obviously misunderstood by most respondents. Several respondents

mentioned that they are getting their own transit connectivity via static tunnels.

-Something else: Answers were 6VPE + NAT64/DNS64; PNAT; "considering 6RD"

*Please briefly explain the main reasons/issues behind your choice: The best summary of the answers is "dual stack is simplest, why do anything else?".

*Which types of equipment in your network are unable to support IPv6? The most common answer was CPE (9 mentions). Also mentioned: handsets; DSLAMs; routers (including several specific models); traffic management boxes; load balancers; VPN boxes; management interfaces & systems; firewalls; billing systems.

*Can they be field-upgraded? 5 yes, 4 partially, 10 no and numerous "don't know" or "hopefully"

*Is any equipment 100% dedicated to IPv6? 7 yes, 23 no

*Is IPv6 an opportunity to restructure your whole topology? 2 yes, 5 partial, 23 no

*Do you include support for DNS AAAA queries over IPv6? 20 yes, 5 plan, 4 no

*Do you include support for reverse DNS for IPv6 addresses? 21 yes, 3 plan, 5 no

*What length(s) of IPv6 prefix do you have or need from the registry? 1 /19, 1 /21 (plus several /32s), 1 /22 1 /30, 3 multiple /32s, 20 /32, 3 /48

*What length(s) of IPv6 prefix are offered to customers? 15 ISPs offer more than one of /48, /52, /56, /60 or /64. 2 offer /56 only, 7 offer /48 only. Two wired operators offer /64 only. Mobile operators offer /64 in accordance with 3GPP, but at least one would like to be allowed to offer /128 or /126.

*Do any customers share their IPv6 prefix among multiple hosts? Unfortunately this question was unclear and obviously misunderstood by most respondents.

*Do any of your customers prefer to use PI IPv6 prefixes? 9 yes, 17 no

*How are IPv6 prefixes delegated to CPEs? 15 manual, 10 DHCPv6[-PD], 8 SLAAC, 8 RADIUS, 2 PPOE

*Are your SMTP, POP3 and IMAP services dual-stack? 10 yes, 6 plan, 12 no

*Are your HTTP services, including caching and webmail, dual-stack? 9 yes, 1 partial, 4 plan, 14 no

*Are any other services dual-stack? 11 yes, 2 plan, 10 no

*Is each of the following dual-stack?

-Firewalls: 12 yes, 3 partial, 3 plan, 8 no

-Intrusion detection: 10 yes, 2 plan, 12 no

-Address management software: 15 yes, 1 plan, 12 no

-Accounting software: 7 yes, 20 no

-Monitoring software: 16 yes, 2 partial, 2 plan, 10 no

-Network management tools: 13 yes, 4 partial, 1 plan, 10 no

*Do you or will you have IPv6-only customers? 13 yes (or maybe), 17 no (or not likely)

*Do you have customers who have explicitly refused to consider IPv6? 5 yes, 22 no

*Interworking issues:

-How many years do you expect customers to run any IPv4-only applications? Answers range from 2 years to infinity, most frequent answer is >10 years.

-Is IPv6-IPv4 interworking at the the IP layer needed? 15 yes, 9 uncertain, 3 no

-Do you include a NAT-PT IPv6/IPv4 translator? 2 yes, 1 plan, 25 no

-If yes, does that include DNS translation? 1 plan, 2 no

-If not, do you plan to operate an IPv6/IPv4 translator? 10 plan (NAT64), 7 no, others uncertain

-If not, how do you plan to connect IPv6-only customers to IPv4-only services? 7 rely on dual stack; 3 have no plan (one said "their problem")

-If you offer IP multicast, will that need to be translated too?
1 yes, 2 uncertain, 5 no

*Any plans for Mobile IPv6 (or Nemo mobile networks)? 6 yes, 2
uncertain, 21 no

Appendix B. Questionnaire

[TOC](#)

This appendix reproduces the technical body of the questionnaire that was made available for ISPs to express their requirements, plans and experience.

I. General questions about IP service

1. Do you offer origin-only (stub, end-user) IP service, transit IP service, or both?

2. Approximate number of private/small office customers (one IPv4 address)

3. Approximate number of corporate customers (block of IPv4 addresses, not included in Q2)

4. Do you offer IP multicast service?

5. Do any of your customers require multihoming to multiple ISPs?

6. Access technologies used (ADSL, etc.)

7. Do your customers use CPE that you supply?

7.1. What % of customers?

7.2. Does the CPE that you provide support native IPv6?

8. When do you expect to run out of public IPv4 address space inside your own network?

8.1. Do you run private (RFC1918) addresses and NAT within your network (i.e., a second layer of NAT in the

case of customers with their own NAT)?

8.2. What % of your IPv4 space is needed for your own use (not for customers)?

9. When do you expect to run out of public IPv4 address space for customers?

9.1. Do you offer private (RFC1918) addresses to your customers?

II. Questions about requirements for IPv6 service

10. Are some big customers requesting IPv6?

11. When do you predict 10% and 50% of your customers to require IPv6 service?

12. When do you require IPv6 to be a standard service available to all customers?

13. When do you predict IPv6 traffic to reach 50% of total traffic?

III. Questions about status and plans for IPv6 service

14. Do you currently offer IPv6 as a regular service?

14.1. What % of your customers currently use IPv6?

14.2. When do you plan to start IPv6 deployment?

14.3. When do you plan to offer IPv6 as a special or beta-test service to customers?

15. When do you plan to offer IPv6 as a regular service to all customers?

IV. Questions about IPv6 technologies

16. Which basic IPv6 access method(s) apply:

16.1. dual stack routing backbone?

16.2. separate IPv4 and IPv6 backbones?

16.3. 6to4 relay?

16.4. Teredo server?

16.5. tunnel broker? If so, which one?

16.6. Something else? Please briefly describe your method:

16.7. If possible, please briefly explain the main reasons/issues behind your choice.

17. Which types of equipment in your network are unable to support IPv6?

17.1. Can they be field-upgraded to support IPv6?

- 17.2. Is any equipment 100% dedicated to IPv6?
18. Is IPv6 an opportunity to restructure your whole topology?
19. Do you include support for DNS AAAA queries over IPv6?
20. Do you include support for reverse DNS for IPv6 addresses?
21. What length(s) of IPv6 prefix do you have or need from the registry?
22. What length(s) of IPv6 prefix are offered to customers?
 - 22.1. Do any customers share their IPv6 prefix among multiple hosts?
23. Do any of your customers prefer to use PI IPv6 prefixes instead of a prefix from you?
24. How are IPv6 prefixes delegated to CPEs? (Manual, PPPoE, RADIUS, DHCPv6, stateless autoconfiguration/RA, etc...)
25. Are your SMTP, POP3 and IMAP services dual-stack?
26. Are your HTTP services, including caching and webmail, dual-stack?
27. Are any other services dual-stack?
28. Is each of the following dual-stack?
 - 28.1. Firewalls
 - 28.2. Intrusion detection
 - 28.3. Address management software
 - 28.4. Accounting software
 - 28.5. Monitoring software
 - 28.6. Network management tools
29. Do you or will you have IPv6-only customers?
30. Do you have customers who have explicitly refused to consider IPv6?
31. How many years do you expect customers to run any IPv4-only applications?
32. Is IPv6-IPv4 interworking at the IP layer needed?

33. Do you include a NAT-PT IPv6/IPv4 translator?

33.1. If yes, does that include DNS translation?

33.2. If not, do you plan to operate an IPv6/IPv4 translator?

33.3. If not, how do you plan to connect IPv6-only customers to IPv4-only services?

33.4. If you offer IP multicast, will that need to be translated too?

34. Any plans for Mobile IPv6 (or Nemo mobile networks)?

35. What features and tools are missing today for IPv6 deployment and operations?

36. Any other comments about your IPv6 experience or plans? What went well, what was difficult, etc.

Authors' Addresses

[TOC](#)

Brian Carpenter
Department of Computer Science
University of Auckland
PB 92019
Auckland, 1142
New Zealand

Email: brian.e.carpenter@gmail.com

Sheng Jiang
Huawei Technologies Co., Ltd
KuiKe Building, No.9 Xinxu Rd.,
Shang-Di Information Industry Base, Hai-Dian District, Beijing
P.R. China

Email: shengjiang@huawei.com