V6OPS
Internet-Draft
Intended status: Informational
Expires: August 26, 2012

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

February 23, 2012

IPv6 Guidance for Internet Content and Application Service Providers
draft-carpenter-v6ops-icp-guidance-03

Abstract

This document provides guidance and suggestions for Internet Content Providers and Application Service Providers who wish to offer their service to both IPv6 and IPv4 customers. Many of the points will also apply to any enterprise network preparing for IPv6 users.

Table of Contents

1. Introduction

The deployment of IPv6 [RFC2460] is now in progress, and users with no IPv4 access are likely to appear in increasing numbers in the coming years. Any provider of content or application services over the Internet will need to arrange for IPv6 access or else risk losing large numbers of potential customers. The time for action is now, while the number of such customers is small, so that appropriate skills, software and equipment can be acquired in good time to scale up the IPv6 service as demand increases. An additional advantage of early support for IPv6 customers is that it will reduce the number of customers connecting later via IPv4 "extension" solutions such as double NAT, which will otherwise degrade the user experience.

Nevertheless, it is important that the introduction of IPv6 service should not make service for IPv4 customers worse. In some circumstances, technologies intended to assist in the transition from IPv4 to IPv6 are known to have negative effects on the user experience. A deployment strategy for IPv6 must avoid these effects as much as possible.

The purpose of this document is to provide guidance and suggestions for Internet Content Providers (ICPs) and Application Service Providers (ASPs) who wish to offer their services to both IPv6 and IPv4 customers. For simplicity, the term ICP is mainly used in the body of this document, but the guidance also applies to ASPs. Many of the points in this document will also apply to enterprise networks that do not classify themselves as ICPs. Any enterprise or department that runs at least one externally accessible server, such as an HTTP server, may also be concerned. Although specific managerial and technical approaches are described, this is not a rule book; each operator will need to make its own plan, tailored to its own services and customers.

2. General Strategy

The most importance advice here is to actually have a general strategy. Adding support for a second network layer protocol is a new departure for most modern organisations, and it cannot be done casually on a day-by-day basis. Even if it is impossible to write a precisely dated plan, the intended steps in the process need to be defined well in advance. There is no single blueprint for this. The rest of this document is meant to provide a set of topics to be taken into account in defining the strategy.

In determining the urgency of this strategy, it should be noted that the central IPv4 registry (IANA) ran out of spare blocks of IPv4

addresses in February 2011 and the various regional registries are expected to exhaust their reserves over the next one to two years. After this, Internet Service Providers (ISPs) will run out at dates determined by their own customer base. No precise date can be given for when IPv6-only customers will appear in commercially significant numbers, but - particularly in the case of mobile users - it may be quite soon. Complacency about this is therefore not an option for any ICP that wishes to grow its customer base over the coming years.

The most common strategy for an ICP is to provide dual stack services - both IPv4 and IPv6 on an equal basis - to cover both existing and future customers. This is the recommended strategy in [RFC6180] for straightforward situations. Some ICPs who already have satisfactory operational experience with IPv6 might consider an IPv6-only strategy, with IPv4 clients being supported by translation or proxy at their ISP border. However, the present document is addressed to ICPs without IPv6 experience, who are likely to prefer the dual stack model to build on their existing IPv4 service.

Within the dual stack model, two approaches could be adopted, sometimes referred to as "outside in" and "inside out":

- Outside in: start by providing external users with an IPv6 public access to your services, for example by running a reverse proxy that handles IPv6 customers (see Section 7 for details). Progressively enable IPv6 internally.
- Inside out: start by enabling internal networking infrastructure, hosts, and applications to support IPv6. Progressively reveal IPv6 access to external customers.

Which of these approaches to adopt depends on the precise circumstances of the ICP concerned. "Outside in" has the benefit of giving interested customers IPv6 access at an early stage, and thereby gaining precious operational experience, before meticulously updating every piece of equipment and software. For example, if some back-office system, that is never exposed to users, only supports IPv4, it will not cause delay. "Inside out" has the benefit of completing the implementation of IPv6 as a single project. Any ICP could choose this approach, but it might be most appropriate for a small ICP without complex back-end systems.

A point that must be considered in the strategy is that some customers will remain IPv4-only for many years, others will have both IPv4 and IPv6 access, and yet others will have only IPv6. Additionally, mobile customers may find themselves switching between IPv4 and IPv6 access as they travel, even within a single session. Services and applications must be able to deal with this, just as easily as they deal today with a user whose IPv4 address changes (see

the discussion of cookies in Section 8.2).

Neverthless, the end goal is to have a network that does not need major changes when at some point in the future it becomes possible to transition to IPv6-only, even if only for some parts of the network. That is, the IPv6 deployment should be designed in such a way as to more or less assume that IPv4 is absent, so the network will function seamlessly when it is indeed no longer there.

An important first step in every strategy is to determine from every hardware and software supplier details of their planned dates for providing full IPv6 support, with performance equivalent to IPv4, in their products and services.

3. Education and Skills

Some older staff may have experience of running multiprotocol networks, which were common twenty years ago before the dominance of IPv4. However, IPv6 will be new to them, and also to younger staff brought up on TCP/IP. It is not enough to have one "IPv6 expert" in a team. On the contrary, everybody who knows about IPv4 needs to know about IPv6, from network architect to help desk responder. Therefore, an early and essential part of the strategy must be education, including practical training, so that all staff acquire a general understanding of IPv6, how it affects basic features such as the DNS, and the relevant practical skills. To take a trivial example, any staff used to dotted-decimal IPv4 addresses need to become familiar with the colon-hexadecimal format used for IPv6.

There is an anecdote of one IPv6 deployment in which prefixes including the letters A to F were avoided by design, to avoid confusing sysadmins unfamiliar with hexadecimal notation. This is not a desirable result. There is another anecdote of a help desk responder telling a customer to "disable one-Pv6" in order to solve a problem. It should be a goal to avoid having untrained staff who don't understand hexadecimal or who can't even spell "IPv6".

It is very useful to have a small laboratory network available for training and self-training in IPv6, where staff may experiment and make mistakes without disturbing the operational IPv4 service. This lab should run both IPv4 and IPv6, to gain experience with a dual-stack environment and new features such as having multiple addresses per interface.

A final remark about training is that it should not be given too soon, or it will be forgotten. Training has a definite need to be done "just in time" in order to properly "stick." Training, lab

experience, and actual deployment should therefore follow each other immediately. If possible, training should even be combined with actual operational experience.

4. Arranging IPv6 Connectivity

There are, in theory, two ways to obtain IPv6 connectivity to the Internet.

- Native. In this case the ISP simply provides IPv6 on exactly the same basis as IPv4 - it will appear at the ICP's border router(s), which must then be configured in dual-stack mode to forward IPv6 packets in both directions. This is by far the better method. An ICP should contact all its ISPs to verify when they will provide native IPv6 support, whether this has any financial implications, and whether the same service level agreement will apply as for IPv4. Any ISP that has no definite plan to offer native IPv6 service should be avoided.
- Tunnel. It is possible to configure an IPv6-in-IPv4 tunnel to a remote ISP that offers such a service. A dual-stack router in the ICP's network will act as a tunnel end-point, or this function could be included in the ICP's border router.

  A tunnel is a reasonable way to obtain IPv6 connectivity for initial testing and skills acquisition. However, it introduces an inevitable extra latency compared to native IPv6, giving users a noticeably worse response time for complex web pages. It is also likely to limit the IPv6 MTU size. In normal circumstances, native IPv6 will provide an MTU size of at least 1500 bytes, but it will almost inevitably be less for a tunnel, possibly as low as 1280 bytes (the minimum MTU allowed for IPv6). Apart from the resulting loss of efficiency, there are cases in which Path MTU Discovery fails, therefore IPv6 fragmentation fails, and in this case the lower tunnel MTU will actually cause connectivity failures for customers.

  For these reasons, ICPs are strongly recommended to obtain native IPv6 service before attempting to offer a production-quality service to their users.

5. IPv6 Infrastructure

5.1. Address and subnet assignment

An ICP must first decide whether to apply for its own Provider Independent (PI) address prefix for IPv6. The default is to obtain a

Provider Aggregated (PA) prefix from each of its ISPs, and operate them in parallel. Both solutions are viable in IPv6. However, scaling properties of the wide area routing system (BGP4) limit the routing of PI prefixes, so only large content providers can justify the bother and expense of obtaining a PI prefix and convincing their ISPs to route it. Millions of enterprise networks, including smaller content providers, will use PA prefixes. In this case, a change of ISP would necessitate a change of the corresponding PA prefix, using the procedure outlined in [RFC4192].

An ICP that has multiple connections via multiple ISPs will have multiple PA prefixes. This results in multiple PA-based addresses for the servers, or for load balancers if they are in use.

An ICP may also choose to operate a Unique Local Address prefix [RFC4193] for internal traffic only, as described in [RFC4864].

Depending on its projected future size, an ICP might choose to obtain /48 PI or PA prefixes (allowing 16 bits of subnet address) or longer PA prefixes, e.g. /56 (allowing 8 bits of subnet address). Clearly the choice of /48 is more future-proof. Advice on the numbering of subnets may be found in [RFC5375].

Since IPv6 provides for operating multiple prefixes simultaneously, it is important to check that all relevant tools, such as address management packages, can deal with this. In particular, the need to allow for multiple PA prefixes with IPv6, and the possible need to renumber, means that using manually assigned static addresses for servers is problematic [I-D.carpenter-6renum-static-problem].

Theoretically, it would be possible to operate an ICP's IPv6 network using only Stateless Address Autoconfiguration [RFC4862]. In practice, an ICP of reasonable size will probably choose to operate DHCPv6 [RFC3315] and use it to support stateful and/or on-demand address assignment.

5.2. Routing

In a dual stack network, IPv4 and IPv6 routing protocols operate quite independently and in parallel. The common routing protocols all exist in IPv6 versions, such as OSPFv3 [RFC5340], IS-IS [RFC5308], and even RIPng [RFC2080] [RFC2081]. For trained staff, there should be no particular difficulty in deploying IPv6 routing without disturbance to IPv4 services.

The performance impact of dual stack routing needs to be evaluated. In particular, what performance does the router vendor claim for IPv6? If the performance is significantly inferior compared to IPv4,

will this be an operational problem? To answer this question, the ICP will need a projected model for the amount of IPv6 traffic expected initially, and its likely rate of increase. [[Note: further input from the WG is needed on this point.]]

If a site operates multiple PA prefixes as mentioned in Section 5.1, complexities may appear in routing configuration. In particular, source-based routing rules may be needed to ensure that outgoing packets are routed to the appropriate border router and ISP link. Normally, a packet sourced from an address assigned by ISP X should not be sent via ISP Y, to avoid ingress filtering by Y [RFC2827] [RFC3704]. Additional considerations may be found in [I-D.ietf-v6ops-ipv6-multihoming-without-ipv6nat].

Each IPv6 subnet normally has a /64 prefix, leaving another 64 bits for the interface identifiers of individual hosts. In contrast, a typical IPv4 subnet will have no more than 8 bits for the host identifier, thus limiting the subnet to 256 or fewer hosts. A dual stack design will typically use the same subnet topology for IPv4 and IPv6, and therefore the same router topology. This means that the limited subnet size of IPv4 will be imposed on IPv6. It would be theoretically possible to avoid this limitation by implementing a different subnet and router topology for IPv6, for example by ingenious use of VLANs. This is not advisable, as it would result in extremely complex fault diagnosis when something went wrong.

5.3. DNS

This is largely a case of "just do it." Each externally visible host (or virtual host) that has an A record for its IPv4 address needs an AAAA record [RFC3596] for its IPv6 address, and a reverse entry if applicable. One important detail is that some clients (especially Windows XP) can only resolve DNS names via IPv4, even if they can use IPv6 for application traffic. It is therefore advisable for all DNS servers to respond to queries via both IPv4 and IPv6.

6. Load Balancers

It is to be expected that IPv6 traffic will initially be low, i.e. a small percentage of IPv4 traffic. For this reason, updating load balancers to fully support IPv6 can perhaps be delayed; however, such an update needs to be planned in anticipation of significant growth over a period of several years. The same would apply to TLS or HTTP proxies used for load balancing purposes. It is important to obtain appropriate assurances from vendors about their IPv6 support, including performance aspects (as discussed for routers in Section 5.2).

7. Proxies

An HTTP proxy [RFC2616] can readily be configured to handle incoming connections over IPv6 and to proxy them to a server over IPv4. Therefore, a single proxy can be used as the first step in an outside-in strategy, as shown in the following diagram:

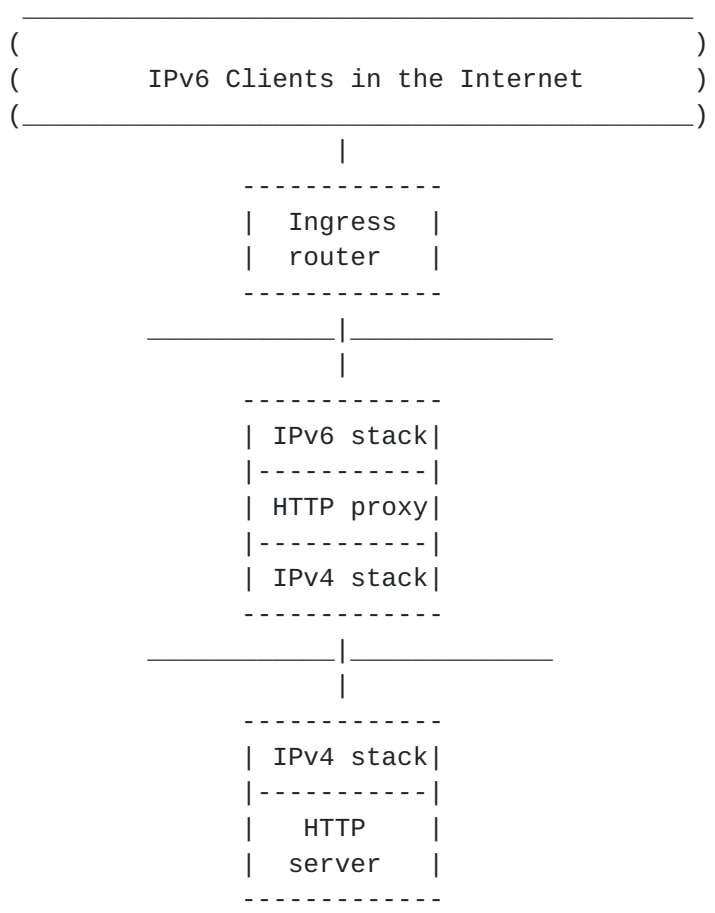

In this case, the AAAA record for the service would provide the IPv6 address of the proxy. This approach will work for any HTTP or HTTPS applications that operate successfully via a proxy, as long as IPv6 load remains low.

8. Servers

8.1. Network Stack

The TCP/IP network stacks in popular operating systems have supported IPv6 for many years. In most cases, it is sufficient to enable IPv6 and possibly DHCPv6; the rest will follow. Servers inside an ICP

network will not need to support any transition technologies beyond a simple dual stack, with a possible exception for 6to4 mitigation noted below in Section 9.

8.2. Application Layer

Basic HTTP servers have been able to handle an IPv6-enabled network stack for some years, so at the most it will be necessary to update to a more recent software version. The same is true of generic applications such as email protocols. No general statement can be made about other applications, especially proprietary ones, so each ASP will need to make its own determination.

One important recommendation here is that all applications should use domain names, which are IP-version-independent, rather than IP addresses. Applications based on middlware platforms which have uniform support for IPv4 and IPv6, for example Java, may be able to support both IPv4 and IPv6 naturally without additional work.

A specific issue for HTTP-based services is that IP address-based cookie authentication schemes will need to deal with dual-stack clients. Servers might create a cookie for an IPv4 connection or an IPv6 connection, depending on the setup at the client site and on the whims of the client operating system. There is no guarentee that a given client will consistently use the same address family, especially when accessing a collection of sites rather than a single site. If the client is using privacy addresses [RFC4941], the IPv6 address (but not its /64 prefix) might change quite frequently. Any cookie mechanism based on 32-bit IPv4 addresses will need significant remodelling.

Generic considerations on application transition are discussed in [RFC4038], but many of them will not apply to the dual-stack ICP scenario. An ICP that creates and maintains its own applications will need to review them for any dependency on IPv4.

8.3. Geolocation

As time goes on, it is to be assumed that geolocation methods and databases will be updated to fully support IPv6 prefixes. There is no reason they will be more or less accurate in the long term than those available for IPv4. However, we can expect many more clients to be mobile as time goes on, so geolocation based on IP addresses alone may become problematic. Initially, at least, ICPs may observe some weakness in geolocation for IPv6 clients.

9. Coping with Transition Technologies

As mentioned above, an ICP should obtain native IPv6 connectivity from its ISPs. In this way, the ICP can avoid most of the complexities of the numerous IPv4-to-IPv6 transition technologies that have been developed; they are all second-best solutions. However, some clients are sure to be using such technologies. An ICP needs to be aware of the operational issues this may cause and how to deal with them.

In some cases outside the ICP's control, clients might reach a content server via a network-layer translator from IPv6 to IPv4. ICPs who are offering a dual stack service and providing both A and AAAA records, as recommended in this document, should not normally receive traffic from NAT64 translators [RFC6146]. Exceptionally, however, such traffic could arrive via IPv4 from an IPv6-only client whose DNS resolver failed to receive the ICP's AAAA record for some reason. Such traffic would be indistinguishable from regular IPv4-via-NAT traffic.

Alternatively, ICPs who are offering a dual stack service might exceptionally receive IPv6 traffic translated from an IPv4-only client that somehow failed to receive the ICP's A record. An ICP could also receive IPv6 traffic with translated prefixes [RFC6296]. These two cases would only be an issue if the ICP was offering any service that depends on the assumption of end-to-end IPv6 address transparency.

In other cases, also outside the ICP's control, IPv6 clients may reach the IPv6 Internet via some form of IPv6-in-IPv4 tunnel. In this case a variety of problems can arise, the most acute of which affect clients connected using the Anycast 6to4 solution [RFC3068]. Advice on how ICPs may mitigate these 6to4 problems is given in Section 4.5. of [RFC6343]. For the benefit of all tunnelled clients, it is essential to verify that Path MTU Discovery works correctly (i.e., the relevant ICMPv6 packets are not blocked) and that the server-side TCP implementation correctly supports the Maximum Segment Size (MSS) negotiation mechanism [RFC2923] for IPv6 traffic.

Some ICPs have implemented an interim solution to mitigate transition problems by limiting the visibility of their AAAA records to users with validated IPv6 connectivity [I-D.ietf-v6ops-v6-aaaa-whitelisting-implications].

Another approach taken by some ICPs is to offer IPv6-only support via a specific DNS name, e.g., ipv6.example.com, if the primary service is www.example.com. In this case ipv6.example.com would have an AAAA record only. This has some value for testing purposes, but is

otherwise only of interest to hobbyist users willing to type in special URLs.

There is little an ICP can do to deal with client-side or remote ISP deficiencies in IPv6 support, but it is hoped that the "happy eyeballs" [I-D.ietf-v6ops-happy-eyeballs] approach will improve the ability for clients to deal with such problems.

10. Content Delivery Networks

DNS-based techniques for diverting users to Content Delivery Network (CDN) points of presence (POPs) will work for IPv6, if AAAA records are provided as well as A records. In general the CDN should follow the recommendations of this document, especially by operating a full dual stack service at each POP. Additionally, each POP will need to handle IPv6 routing exactly like IPv4, for example running BGP4+ [RFC4760] if appropriate.

Note that if an ICP supports IPv6 but its CDN does not, its clients will continue to use IPv4 and any IPv6-only clients will have to use a transition solution of some kind. This is not a desirable situation, since the ICP's work to support IPv6 will be wasted. The converse is not true: if the CDN supports IPv6 but the ICP does not, dual-stack and IPv6-only clients will obtain IPv6 access.

An ICP might face a complex situation, if its CDN provider supports IPv6 at some POPs but not at others. IPv6-only clients could only be diverted to a POP supporting IPv6. There are also scenarios where a dual-stack client would be diverted to a mixture of IPv4 and IPv6 POPs for different URLs, according to the A and AAAA records provided and the availability of optimisations such as "happy eyeballs." These complications do not affect the viability of relying on a dual-stack CDN, however.

The CDN itself faces related complexity: "As IPv6 rolls out, it's going to roll out in pockets, and that's going to make the routing around congestion points that much more important but also that much harder," stated John Summers of Akamai in 2010.

11. Business Partners

As noted earlier, it is in an ICP's or ASP's best interests that their users have direct IPv6 connectivity, rather than indirect IPv4 connectivity via double NAT. If the ICP or ASP has a direct business relationship with some of their clients, or with the networks that connect them to their clients, they are advised to coordinate with

those partners to ensure that they have a plan to enable IPv6. They should also verify and test that there is first-class IPv6 connectivity end-to-end between the networks concerned. This is especially true for implementations that require IPv6 support in specialized programs or systems in order for the IPv6 support on the ICP/ASP side to be useful.

12. Operations and Management

There is no doubt that, initially, IPv6 deployment will have operational impact, as well as requiring education and training as mentioned in Section 3. Staff will have to update network elements such as routers, update configurations, provide information to end users, and diagnose new problems. However, for an enterprise network, there is plenty of experience, e.g. on numerous university campuses, showing that dual stack operation is no harder than IPv4-only in the steady state.

Whatever management, monitoring and logging is performed for IPv4 is also needed for IPv6. Therefore, all products and tools used for these purposes must be updated to fully support IPv6. Note that since an IPv6 network may operate with more than one IPv6 prefix and therefore more than one address per host, the tools must deal with this as a normal situation. This includes any address management tool in use (see Section 5.1) as well as tools used for creating DHCP and DNS configurations. There is significant overlap here with the tools involved in site renumbering [I-D.jiang-6renum-enterprise].

As far as possible, however, mutual dependency between IPv4 and IPv6 operations should be avoided. A failure of one should not cause a failure of the other. One precaution to avoid this would be for back-end systems such as network management databases to be dual stacked as soon as convenient. It should also be possible to use IPv4 connectivity to repair IPv6 configurations, and vice versa.

Dual stack, while necessary, does have management scaling and overhead considerations. As noted earlier, the long term goal is to move to single-stack IPv6, when the network and its customers can support this. This is an additional reason why mutual dependency between the address families should be avoided in the management system in particular; a hidden dependency on IPv4 that had been forgotten for many years would be highly inconvenient.

13. Security Considerations

Essentially every threat that exists for IPv4 exists or will exist

for IPv6. Therefore, it is essential to update firewalls, intrusion detection systems, denial of service precautions, and security auditing technology to fully support IPv6. Otherwise, IPv6 will become an attractive target for attackers.

When multiple PA prefixes are in use as mentioned in Section 5.1, firewall rules must allow for all valid prefixes, and must be set up to work as intended even if packets are sent via one ISP but return packets arrive via another.

Performance aspects of dual stack firewalls must be considered (as discussed for routers in Section 5.2).

In a dual stack operation, there may be a risk of cross-contamination between the two protocols. For example, a successful IPv4-based denial of service attack might also deplete resources needed by the IPv6 service, or vice versa. This risk strengthens the argument that IPv6 security must be up to the same level as IPv4.

A general overview of techniques to protect an IPv6 network against external attack is given in [RFC4864]. Assuming an ICP has native IPv6 connectivity, it is advisable to block incoming IPv6-in-IPv4 tunnel traffic using IPv4 protocol type 41. Outgoing traffic of this kind should be blocked except for the case noted in Section 4.5 of [RFC6343]. ICMPv6 traffic should only be blocked in accordance with [RFC4890]; in particular, Packet Too Big messages, which are essential for PMTU discovery, must not be blocked.

Scanning attacks to discover the existence of hosts are much less likely to succeed for IPv6 than for IPv4 [RFC5157]. However, this is only true if IPv6 hosts are configured with interface identifiers that are hard to guess; for example, it is not advisable to manually configure servers with static interface identifiers starting from "1".

Transport Layer Security version 1.2 [RFC5246] and its predecessors work correctly with TCP over IPv6, meaning that HTTPS-based security solutions are immediately applicable. The same should apply to any other transport-layer or application-layer security techniques.

If an ASP uses IPsec [RFC4301] and IKE [RFC5996] in any way to secure connections with clients, these too are fully applicable to IPv6, but only if the software stack at each end has been appropriately updated.

14. IANA Considerations

This document requests no action by IANA.

15. Acknowledgements

Valuable contributions were made by Erik Kline. Useful comments were received from Tassos Chatzithomaoglou, Wesley George, Victor Kuarsingh, Bing Liu, John Mann, and other participants in the V6OPS working group.

This document was produced using the xml2rfc tool [RFC2629].

16. Change log [RFC Editor: Please remove]

draft-carpenter-v6ops-icp-guidance-03: additional WG comments, 2012-02-23.

draft-carpenter-v6ops-icp-guidance-02: additional WG comments, 2012-01-07.

draft-carpenter-v6ops-icp-guidance-01: multiple clarifications after WG comments, 2011-12-06.

draft-carpenter-v6ops-icp-guidance-00: original version, 2011-10-22.

17. References

17.1. Normative References

[RFC2080] Malkin, G. and R. Minnear, "RIPng for IPv6", RFC 2080, January 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

[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.

[RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, October 2003.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, October 2008.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

[RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.

17.2. Informative References

[I-D.carpenter-6renum-static-problem]
Carpenter, B. and S. Jiang, "Problem Statement for Renumbering IPv6 Hosts with Static Addresses", draft-carpenter-6renum-static-problem-01 (work in progress), December 2011.

[I-D.ietf-v6ops-happy-eyeballs]
Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", draft-ietf-v6ops-happy-eyeballs-07 (work in progress), December 2011.

[I-D.ietf-v6ops-ipv6-multihoming-without-ipv6nat]
Troan, O., Miles, D., Matsushima, S., Okimoto, T., and D. Wing, "IPv6 Multihoming without Network Address Translation", draft-ietf-v6ops-ipv6-multihoming-without-ipv6nat-04 (work in progress), February 2012.

[I-D.ietf-v6ops-v6-aaaa-whitelisting-implications]
Livingood, J., "Considerations for Transitioning Content to IPv6", draft-ietf-v6ops-v6-aaaa-whitelisting-implications-09 (work in progress), February 2012.

[I-D.jiang-6renum-enterprise]
Jiang, S., Liu, B., and B. Carpenter, "IPv6 Enterprise Network Renumbering Scenarios and Guidelines", draft-jiang-6renum-enterprise-02 (work in progress), December 2011.

[RFC2081] Malkin, G., "RIPng Protocol Applicability Statement", RFC 2081, January 1997.

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

[RFC2923] Lahey, K., "TCP Problems with Path MTU Discovery", RFC 2923, September 2000.

[RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, June 2001.

[RFC4038] Shin, M-K., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "Application Aspects of IPv6 Transition", RFC 4038, March 2005.

[RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.

[RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.

[RFC4890] Davies, E. and J. Mohacsi, "Recommendations for Filtering ICMPv6 Messages in Firewalls", RFC 4890, May 2007.

[RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in

IPv6", RFC 4941, September 2007.

[RFC5157] Chown, T., "IPv6 Implications for Network Scanning", RFC 5157, March 2008.

[RFC5375] Van de Velde, G., Popoviciu, C., Chown, T., Bonness, O., and C. Hahn, "IPv6 Unicast Address Assignment Considerations", RFC 5375, December 2008.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

[RFC6180] Arkko, J. and F. Baker, "Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment", RFC 6180, May 2011.

[RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011.

[RFC6343] Carpenter, B., "Advisory Guidelines for 6to4 Deployment", RFC 6343, August 2011.

Authors' Addresses

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
Q14, Huawei Campus
No.156 Beiqing Road
Hai-Dian District, Beijing 100095
P.R. China

Email: jiangsheng@huawei.com