

IPv6 Operations	T.J. Chown
Internet-Draft	University of Southampton
Intended status: Informational	M. Ford
Expires: December 09, 2011	Internet Society
	S. Venaas
	Cisco Systems
	June 07, 2011

World IPv6 Day Call to Arms
draft-chown-v6ops-call-to-arms-03

Abstract

The Internet Society (ISOC) has declared that June 8th 2011 will be World IPv6 Day, on which some major organisations are going to make their content available over IPv6. With the likes of Google and Facebook providing IPv6 access to their production services and domains, it is very likely we will see more IPv6 traffic flowing across the Internet than has ever been seen before. With this in mind, it seems timely to issue a call to arms for systems and network administrators to review their organisation's IPv6 capabilities in order to mitigate common causes of IPv6 connectivity problems in advance of the day. The increased traffic on World IPv6 Day should also create an excellent opportunity to observe the behaviour and performance of IPv6; it is thus very desirable to have appropriate measurement tools in place in advance. We discuss some appropriate tools from the network and application perspective.

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 December 09, 2011.

Copyright Notice

Copyright (c) 2011 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. [Connectivity Issues](#)
 - *2.1. [Unmanaged Tunnels](#)
 - *2.2. [Tunnel Broker first-hop delays](#)
 - *2.3. [Connection Timeouts](#)
 - *2.4. [PMTU Discovery](#)
 - *2.5. [Rogue Router Advertisements](#)
 - *2.6. [Tunnel performance](#)
 - *2.7. [AAAA record advertised but service not enabled](#)
 - *2.8. [IPv6 Reverse DNS](#)
- *3. [Instrumentation](#)
 - *3.1. [IPv6 traffic levels](#)
 - *3.2. [Network flow records](#)
 - *3.3. [Client Web Access Success Rate](#)
 - *3.4. [Tools to measure IPv6 brokenness](#)
 - *3.5. [IPv4 Performance Comparison](#)
 - *3.6. [User Tickets](#)
 - *3.7. [Security monitoring](#)
- *4. [IPv6-only testing](#)
- *5. [Conclusions](#)
- *6. [Security Considerations](#)
- *7. [IANA Considerations](#)

*8. [Acknowledgments](#)

*9. [References](#)

*[Authors' Addresses](#)

1. Introduction

Despite the recent exhaustion of the available IPv4 address pool, deployment of IPv6 remains limited. To help encourage organisations to trial production deployment, ISOC has declared June 8th 2011 as World IPv6 Day [\[ISOC\]](#). Organisations are encouraged to use this day to test IPv6 in production by making their main, externally-facing websites available over IPv6. Sites planning to turn on IPv6 for access in their network in the interest of World IPv6 Day should ensure this is completed well before the day, and commit to leaving it active after the event, and thus using the method they would choose to do so indefinitely. At the current time, this would generally mean enabling dual-stack networking with IPv4 running alongside IPv6. However, IPv6-only networks are ultimately inevitable, and so some sites may choose to use June 8th to undertake some focused tests on that deployment model.

The purpose of this document is two-fold. One is to discuss common IPv6 connectivity issues that are likely to arise on June 8th, with a focus on dual-stack networking (which is likely to be how the vast majority of sites take part). Most of the issues discussed in this text are those that would affect an end site or enterprise network running IPv6, but may be applicable elsewhere. Highlighting the issues should help raise awareness of those problems and possible mitigations. The other purpose is to encourage organisations to think about how they might get useful instrumentation in place to observe what happens in and to/from their networks on the day, both from the network and application perspective. Such measurement tools are likely to be useful in the longer term, so once deployed they could be left in place beyond June 8th.

For sites providing content, June 8th will be a chance to make some public facing services available over IPv6, most likely web content using their production domain (e.g. [www.example.com](#)) rather than a contrived IPv6 test domain (e.g. [www.ipv6.example.com](#)). Enabling public-facing Internet services is a reasonable first step for any organisation deploying IPv6. For ISPs, supporting IPv6 for their Internet-facing services (web, mail, etc.) and recording the impact of World IPv6 Day on their IPv4-only customers is an appropriate action. For sites enabling clients, doing so initially in their IT department may be appropriate; for educational sites enabling IPv6 on eduroam wireless networks could be appropriate given the underlying 802.1x authentication technology is IP version independent.

It should be emphasised that while World IPv6 Day is in many senses an 'experiment' or 'test flight' for IPv6, organisations should strongly

consider deploying IPv6 in exactly the same robust way that they would do if they were deploying IPv6 and leaving it enabled indefinitely. Similarly, applying measures to improve IPv6 robustness, e.g. improved ICMPv6 filtering practice, should be considered long term benefits. That they 'affect' the experiment is not a problem; indeed all measures that improve the robustness of IPv6 deployment should be seen as worthwhile. There will still be problems found, but these can at least be recognised and work done to make them better. The document also includes a brief section on tools that might be used to test IPv6-only operation. The scope of this document is purely informational to provoke discussion.

2. Connectivity Issues

In this section we review some common causes of IPv6 connectivity issues, oriented towards those that end sites or enterprises may have some ability to influence or mitigate. Some issues, such as transit arrangements, are not included - currently the focus is on end sites (or users) who may take part in the World IPv6 Day. Some IPv6 connectivity test sites are emerging, for example [\[testipv6\]](#). There is no significance to the order in which issues are listed.

2.1. Unmanaged Tunnels

One cause of connectivity problems is the use of unmanaged tunnels, in particular 'automated' methods that are not provisioned by the user's ISP. The most common example is 6to4 [\[RFC3056\]](#), or more specifically the 6to4 relay approach described in [\[RFC3068\]](#). A native IPv6 host communicating with a 6to4 host will require both hosts to have access to an appropriately capable 6to4 relay (which may or may not be the same relay). If a host in a native IPv6 network has no route to 2002::/16 it cannot send traffic to a 6to4 host. Similarly, a 6to4 router that cannot reach the well-known IPv4 anycast relay address cannot send traffic to a native IPv6 network. There are also potential issues with Protocol 41 filtering at site borders close to the client. A presentation by Geoff Huston at IETF80 [\[Huston2011\]](#) highlighted the connection failure rates with 6to4, measured in excess of 15%, as well as the additional latency in 6to4 communications, with 6to4 showing an average additional 1.2s latency per retrieval.

One approach to this problem is to encourage sites/ISPs to run local relays, as discussed in [\[I-D.carpenter-v6ops-6to4-teredo-advisory\]](#). This draft discusses how to make 6to4 more robust in situations where there is a conscious decision to use it. Sites using 6to4 should consider deploying local relays to increase the chance of a good IPv6 experience. The alternative to reduce such problems is simply to move 6to4 to Historic, as proposed in [\[I-D.troan-v6ops-6to4-to-historic\]](#). This would mean 6to4 would not be enabled by default anywhere, and once its usage had reduced enough, relays could be turned off.

There may still be some CPE routers that do enable 6to4 by default; it is likely that devices behind such routers will experience problems on World IPv6 Day.

Connection failures and latency with the Teredo protocol [\[RFC4380\]](#) were also highlighted by Geoff Huston's IETF80 presentation. Teredo connection failure rates were as high as 35%, with 1-3s additional latency. One of the connection issues is reliance on the ICMPv6 probe packet being able to reach the destination host; in practice filters may block these. Thus Teredo should not be considered a reliable means of accessing the IPv6 Internet.

2.2. Tunnel Broker first-hop delays

IPv6 tunnel brokers, such as those provided by SixXS (<http://www.sixxs.net>) and Hurricane Electric (<http://tunnelbroker.net>) provide a more robust, managed approach to IPv6-in-IPv4 tunnelled access than 6to4. Individual users interested in IPv6 access for World IPv6 Day, in the absence of IPv6 support from their ISP, should consider registering to use a free tunnel broker. It would be sensible to register for and test your broker client well in advance of IPv6 Day, and ideally plan to keep it available beyond that date, until your ISP provides IPv6 natively for you. One set of test sites to use would be the list cited on the ISOC World IPv6 Day site [\[ISOCsites\]](#).

When choosing a broker service, it is prudent to pick one with a presence near to you that has a minimal round trip time. Providers such as SixXS and HE have tunnel broker servers in many countries. Beware picking a broker in another continent that may add 150ms+ to your round trip times.

2.3. Connection Timeouts

One of the main drivers for IPv6 Day is identifying and fixing the problems that can lead to connection timeouts. Because unreliable IPv6 connectivity leads to intensely frustrating problems for end-users, it is essential that people motivated to deploy IPv6 connectivity, whether for themselves, or for a larger network, only do so in a well-supported, production-quality fashion.

Where dual-stack systems - or rather the applications running on them - have a choice of IPv4 or IPv6 connectivity, timeouts can occur if there is no connectivity on the preferred protocol. For example, if both A and AAAA DNS records exist for a web server, and IPv6 connectivity is broken, there is likely to be some timeout for the browser before the connection drops back to IPv4.

A bigger problem exists if the application or OS tries IPv6 first and then does not fall back to IPv4. A bug in versions of Opera prior to 10.5 caused such behaviour, which was obviously a big issue for Opera users trying to access dual-stack web sites with broken IPv6 connectivity.

The author has undertaken some informal tests at his own site, which shows how different combinations of browsers and operating systems behave in the event of IPv6 connections failing or when ICMP unreachables are received. On Linux/Firefox, web connections timeout after 20 seconds for 'no response', but immediately for unreachables. In contrast, Windows Vista/IE was 20 seconds regardless of unreachables being received. Any non-trivial delay will cause significant user frustration.

A more complete set of tests was run by Teemu Savolainen and reported at IETF80 [\[Savolainen2011\]](#). Although the tests were only samples, they confirmed the results, also showing experiences across a much broader range of platforms, and that the problems with Vista/IE are repeated with Win 7/IE. It's thus clear that if major content providers enable IPv6 on World IPv6 Day, and end users for some reason try to access the content with broken IPv6 connectivity, they are likely to experience significant timeout issues.

This problem is probably the main reason that Google implemented a AAAA whitelisting system for its test sites. The sites had to demonstrate they had good IPv6 connectivity before being allowed into the test programme. The topic is discussed in [\[I-D.ietf-v6ops-v6-aaaa-whitelisting-implications\]](#). For the sake of World IPv6 Day, it is expected that no such whitelisting is in place - that is, after all, the point of having a day dedicated to testing IPv6 in production. An interesting suggestion to handle the problem is the 'happy eyeballs' approach described in [\[I-D.ietf-v6ops-happy-eyeballs\]](#). This approach is now also being suggested for multiple interface systems, as per [\[I-D.chen-mif-happy-eyeballs-extension\]](#). The happy eyeballs philosophy is to try both IPv4 and IPv6 together, and keep the first working connection up, remembering the result for future connection attempts. It may prefer IPv6 slightly in initial connections rather than trying connections exactly simultaneously. It is an interesting approach, though some people are concerned about the additional connection load, or that this 'workaround' is simply masking underlying problems that should be fixed.

[2.4. PMTU Discovery](#)

IPv6 mandates that fragmentation is only undertaken by the sending node, and thus IPv6 requires working PMTU Discovery [\[RFC1981\]](#). An existing RFC gives Recommendations for Filtering ICMPv6 Messages in Firewalls [\[RFC4890\]](#); if this guidance is not followed, connectivity problems are likely to arise. Blindly filtering all ICMPv6 messages is not good practise. Filtering ICMP is a common practice in some IPv4 networks today. Adopting the same approach to ICMPv6 when deploying IPv6 networks will cause connectivity issues for users of the network filtering ICMPv6 and hosts trying to reach the filtered network. RFC 4890 is therefore an important document for IPv6 deployment engineers to read and it is similarly important to verify that IPv6 firewall deployments support appropriate configurations for ICMPv6 filtering.

The minimum MTU for IPv6 is 1280 bytes. Checking the MTU is an important step when connectivity issues arise. Where PMTUD is not working or not implemented, the using the minimum MTU is likely to resolve the problem, though not give optimal performance (the cause should still be investigated and resolved for longer term benefit). Tunnel broker services such as SixXS and HE set their MTUs to default to 1280, probably due to the varying conditions their customers may be in. However, it is preferable for enterprise networks to configure appropriate ICMPv6 filtering to allow PMTUD to operate and establish the most efficient MTUs for a link.

2.5. Rogue Router Advertisements

Within a site, hosts may use IPv6 Stateless Address Autoconfiguration (SLAAC) [\[RFC4862\]](#). However, it is possible for accidental (or malicious) rogue RAs to cause connectivity issues, as described in the Rogue Router Advertisement Problem Statement [\[RFC6104\]](#).

A typical cause of rogue RAs is Windows ICS, which can present a rogue 6to4 router on its wireless interface. This will cause hosts to potentially autoconfigure two global IPv6 addresses and pick the wrong default router, with unpredictable results. As a (bad) example the author experienced a scenario where he had a rogue 6to4 RA, but because the rogue 6to4 was working he was able to access IPv6 networks outside his own network, but could not access most internal hosts inside his own network because he was unwittingly using 6to4 from outside into his own network, and thus being firewalled from those internal hosts.

In many cases, default address selection [\[RFC3484\]](#) (and [\[I-D.ietf-6man-rfc3484-revise\]](#)) would avoid such cases, because the address selection rules should prefer, or can be configured to prefer, native IPv6 over 6to4. However not all operating systems implement RFC 3484 yet, in particular MacOS X (though support may be appearing in Lion). Where rogue RAs cause broken IPv6 behaviour, the timeout issues discussed above may apply.

Adding ACLs to your switches to block ICMPv6 Type 134 packets on ports that do not have routers connected would also minimise the impact of rogue RAs. A more elegant solution is RA Guard [\[RFC6105\]](#), and another is use of SEcure Neighbour Discovery (SEND) [\[RFC3971\]](#). However neither is widely implemented yet. Indeed, any reported operational experience of SEND in an enterprise network would be very welcome.

Finally, there is a tool called RAMond, available freely from <http://ramond.sourceforge.net>, that can be configured to detect and issue deprecating RAs against observed rogue RAs. This software is based on rafxid.

2.6. Tunnel performance

In scenarios where sites currently have manually configured tunnels to gain IPv6 connectivity, it may be the case that such encapsulation is performed by a router's CPU, in which case unexpected high volumes of

traffic may cause problems. Bear in mind that on World IPv6 Day, you may start using IPv6 by default for some high bandwidth applications that you had not used before, e.g. YouTube from Google. It may be prudent to estimate your load for such applications in advance, and test the capability of your tunnelling solution to handle that load.

2.7. AAAA record advertised but service not enabled

If enabling a service for World IPv6 Day, be aware of other existing services that may be running on the same system. If a server has multiple functions, all services should be IPv6 enabled before a AAAA record is entered into the DNS for services that may use that name. A related consideration is to make sure that firewalls don't just drop IPv6 packets to ports that are not in use. It's better if the firewall or host sends an unreachable indication or a TCP RST to avoid a potential timeout. For example, if you add a AAAA record for your web server that also runs say FTP, where FTP is IPv4 only, either the firewall should have port 21 open or the firewall should be configured to send a TCP RST. There are of course tradeoffs in enabling ICMP unreachables.

2.8. IPv6 Reverse DNS

Presence of IPv6 reverse DNS records is used by many systems as a security method. For example, many mail exchangers will only accept SMTP connections from IP addresses with a reverse DNS entry. It is thus important for such records to exist where, for example, a site is sending mail out over IPv6 transport. It is not necessarily the case that such connections will fall back to IPv4 if reverse records are not present.

3. Instrumentation

In this section we discuss potential instrumentation approaches that may be configured in advance of World IPv6 Day, and then retained longer term after the event. These are particularly useful if your site is turning on AAAA records for its production web presence (for example) and wants to get the best insight into how the systems performed and the nature of the end user experience.

These measurements should complement informal, subjective reports from users at participating sites. It is probably prudent to make at least your organisation's IT staff aware of the 'at risk' day, and actions they should take should they experience problems. It may also be desirable to undertake some form of user survey soon afterwards; whether you inform general users in advance is an issue for each site. The ARIN IPv6 wiki is a good source of such advice [\[ARINwiki\]](#).

3.1. IPv6 traffic levels

It should be possible to measure raw IPv6 traffic levels independently on dual-stack switch/router platforms, given implementations of appropriate MIBs. Sites should take steps to ensure they have the tools in place to be able to view the relative levels of IPv4 and IPv6 traffic over time.

Application level measurement is also desirable, because handling of choice (preference) of protocol used lies with the application if both A and AAAA records are returned. Sites should be aware that due to IPv6 Privacy Extensions [[RFC4941](#)] application logs may show more apparent different clients connecting, due to clients cycling the source IPv6 address they use over time.

The types of information gathered might for example include:

- *IPv6 traffic volume, sources of IPv6 traffic by AS, types of IPv6 traffic (e.g. native, 6to4, Teredo, tunnelled);
- *IPv6 application mix, comparison with IPv4;
- *The number and type of IPv6 client connections.

3.2. Network flow records

Where available, sites should seek to generate and record network flow records for traffic, to maximise opportunities to analyse traffic patterns after the event, or in the case of reports of specific problems. Netflow v9 supports IPv6. Open source IPv6-capable Netflow collectors also exist, e.g. nfsen, from <http://nfsen.sourceforge.net>.

3.3. Client Web Access Success Rate

There have been some recent studies on the capabilities of web clients to access content on dual-stack servers by IPv4 or IPv6 in the presence of both A and AAAA records existing for a web domain.

One good example is that of [[Anderson10](#)], as reported at RIPE-61, where the author set up some application (web server) oriented tests for his newspaper content in Norway. The methodology was to add an invisible IFRAME to his site that would include IMG links randomly to 1x1 images that were served either via an IPv4-only target or a dual-stack target. Variation in the hit rates would imply IPv6 brokenness. By analysing the http metadata information could be gleaned on the cause of the brokenness. Results in Q4'2009 showed 0.2-0.3% brokenness, including the Opera bug mentioned above.

Recent figures published by Google suggest at most a 0.1% level of brokenness, indicating some improvement, but that level is still potentially 1 in 1000 users with a problem.

3.4. Tools to measure IPv6 brokenness

Sites may wish to make their own measurements of IPv6 brokenness rather than relying on third party reports. There are some openly available tools available that work along similar principles to the method proposed by Tore Anderson above.

The APNIC Labs test tool uses a combination of JavaScript and Google Analytics to measure various types of brokenness [\[APNIC\]](#). Eric Vyncke's tool [\[Vyncke\]](#) measures a slightly smaller set of types of brokenness, but also looks very useful, with additional reports on the browser type for each failure. The author is currently using the latter tool, and plans to enable the APNIC measurement system shortly when other Analytics updates are applied locally.

3.5. IPv4 Performance Comparison

Where a dual-stack service is deployed, measuring the relative performance of both protocols is desirable. This may primarily be a measurement of throughput or delay, but may also include availability/uptime measurement. A site may choose to set up its own performance measuring framework, for example using open source bandwidth and throughput test tools. Participants in World IPv6 Day will be monitored from a broad range of locations and measurements will be available to show availability of AAAA records, reachability to http service, latency and availability over time.

3.6. User Tickets

It is possible a higher than usual user ticket rate for connectivity issues may be experienced. being able to categorise these cases for subsequent analysis is desirable.

3.7. Security monitoring

We mentioned RAmond above in the context of watching for rogue RAs. There is another useful package called NDPmon, also available freely from <http://ndpmon.sourceforge.net>, that can be configured to watch for certain types of IPv6 'abuse' on your local network. It may be interesting to run the tool to confirm whether any 'bad' traffic is observed within your network on World IPv6 Day.

4. IPv6-only testing

The long-term IPv6 deployment plan is IPv6-only networking, rather than dual-stack. It is not clear how quickly significant IPv6-only networks will emerge, but testing of approaches to IPv6-only operation is desirable as soon as possible. A draft by Jari Arkko and Ari Keranen describes some such experiences [\[I-D.arkko-ipv6-only-experience\]](#). Some experience of NAT64 [\[RFC6146\]](#) has been described in [\[I-D.tan-v6ops-nat64-experiences\]](#), though this appears to have used only NAT-PT

so far. An implementation of NAT64 is available at <http://ecdysis.viagenie.ca>. Operational experience of IVI is also desirable. An implementation of IVI is available at <http://www.ivi2.org/IVI>.

5. Conclusions

With the ISOC World IPv6 Day event due on June 8th 2011, this document aims to help focus attention on both improving awareness and mitigations of common causes of IPv6 connectivity problems, and encouraging sites and organisations to introduce appropriate instrumentation into their networks so they can observe traffic behaviour appropriately.

This is still an early version of the text, and is thus a little drafty. All comments are very welcome towards a mature version in advance of June.

6. Security Considerations

There are no extra security consideration for this document.

7. IANA Considerations

There are no extra IANA consideration for this document.

8. Acknowledgments

To be added.

9. References

[RFC1981]	McCann, J. , Deering, S. and J. Mogul , " Path MTU Discovery for IP version 6 ", RFC 1981, August 1996.
[RFC3056]	Carpenter, B. and K. Moore, " Connection of IPv6 Domains via IPv4 Clouds ", RFC 3056, February 2001.
[RFC3068]	Huitema, C., " An Anycast Prefix for 6to4 Relay Routers ", RFC 3068, June 2001.
[RFC3484]	Draves, R., " Default Address Selection for Internet Protocol version 6 (IPv6) ", RFC 3484, February 2003.
[RFC3971]	Arkko, J., Kempf, J., Zill, B. and P. Nikander, " SEcure Neighbor Discovery (SEND) ", RFC 3971, March 2005.
[RFC4380]	Huitema, C., " Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs) ", RFC 4380, February 2006.
[RFC4862]	

	Thomson, S., Narten, T. and T. Jinmei, " IPv6 Stateless Address Autoconfiguration ", RFC 4862, September 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.
[RFC6104]	Chown, T. and S. Venaas, " Rogue IPv6 Router Advertisement Problem Statement ", RFC 6104, February 2011.
[RFC6105]	Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C. and J. Mohacsi, " IPv6 Router Advertisement Guard ", RFC 6105, February 2011.
[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.
[I-D.carpenter-v6ops-6to4-teredo-advisory]	Carpenter, B, " Advisory Guidelines for 6to4 Deployment ", Internet-Draft draft-carpenter-v6ops-6to4-teredo-advisory-03, March 2011.
[I-D.ietf-v6ops-happy-eyeballs]	Wing, D and A Yourtchenko, " Happy Eyeballs: Success with Dual-Stack Hosts ", Internet-Draft draft-ietf-v6ops-happy-eyeballs-05, October 2011.
[I-D.tan-v6ops-nat64-experiences]	Tan, J, Lin, J and W Li, " Experience from NAT64 applications ", Internet-Draft draft-tan-v6ops-nat64-experiences-00, March 2011.
[I-D.troan-v6ops-6to4-to-historic]	Troan, O, " Request to move Connection of IPv6 Domains via IPv4 Clouds (6to4) to Historic status ", Internet-Draft draft-troan-v6ops-6to4-to-historic-01, March 2011.
[I-D.ietf-v6ops-v6-aaaa-whitelisting-implications]	Livingood, J, " Considerations for Transitioning Content to IPv6 ", Internet-Draft draft-ietf-v6ops-v6-aaaa-whitelisting-implications-08, November 2011.
[I-D.chen-mif-happy-eyeballs-extension]	Chen, G, Williams, C, Wing, D and A Yourtchenko, " Happy Eyeballs Extension for Multiple Interfaces ", Internet-Draft draft-chen-mif-happy-eyeballs-extension-03, October 2011.

[I-D.ietf-6man-rfc3484-revise]	Matsumoto, A, Kato, J, Fujisaki, T and T Chown, " Update to RFC 3484 Default Address Selection for IPv6 ", Internet-Draft draft-ietf-6man-rfc3484-revise-05, October 2011.
[I-D.arkko-ipv6-only-experience]	Arkko, J and A Keranen, " Experiences from an IPv6-Only Network ", Internet-Draft draft-arkko-ipv6-only-experience-04, October 2011.
[APNIC]	, , "IPv6 Capability Tracker", .
[Vyncke]	Vyncke, E., "Estimation of IPv6 brokenness", .
[ARINwiki]	, , "ARIN IPv6 Wiki", .
[testipv6]	, , "Test IPv6", .
[ISOC]	, , "World IPv6 Day", .
[Huston2011]	Huston, G., "Stacking it Up: Experimental Observations on the operation of Dual Stack Services", 2011.
[Savolainen2011]	Savolainen, T., "Experiences of host behaviour in broken IPv6 networks", 2011.
[ISOCsites]	, , "IPv6 Enabled Websites", .
[Anderson10]	Anderson, T., "Measuring and Combating IPv6 Brokenness", 2010.

[Authors' Addresses](#)

Tim Chown Chown University of Southampton Highfield Southampton ,
Hampshire SO17 1BJ United Kingdom EMail: tjc@ecs.soton.ac.uk

Mat Ford Ford Internet Society Geneva, Switzerland EMail:
ford@isoc.org

Stig Venaas Venaas Cisco Systems
Tasman Drive San Jose, CA 95134 USA EMail: stig@cisco.com