

Customer problems that could occur

If you know of any issues that customers might encounter during the transition to IPv6, please add a subheading to this page and explain both the problem and how to address it.

Broken users unable to access dual-stacked content

A small number of users have some kind of misconfiguration or bug in their internet connection that makes them unable to properly access dual-stacked web sites. More often than not, these users have no problems accessing IPv4-only sites, which causes them to perceive the dual-stacked site as the one having problems. A web site operator can detect these users using JavaScript and potentially warn them about the problem, see [Warning broken users with JavaScript](#).

The existence of these users is unfortunately causing content providers to put off enabling IPv6. For more information, see these presentations by Yahoo ([https://sites.google.com/site/ipv6implementors/2010/agenda/07 Fesler Y!atGIPv6ImpConf.pdf?attredirects=0](https://sites.google.com/site/ipv6implementors/2010/agenda/07_Fesler_Y%21atGIPv6ImpConf.pdf?attredirects=0)), [Google](#), and [Redpill Linpro](#). There's also an [article about the problem](#) in Wikipedia.

This section attempts to document the most common causes as to why this happens, and how end users can solve it. It is based on real operational experience from running dual-stacked web sites.

Use of transitional IPv6 connectivity

Transitional IPv6 connectivity (6to4 or Teredo), is usually less reliable than IPv4. For this reason, the end user's web browser should prefer to use IPv4 when attempting to access dual-stacked site. This is well documented in [I-D.vandevelde-v6ops-harmful-tunnels](#), as well as by researchers [E mile Aben](#) and [Geoff Huston](#).

6to4 router functionality is implemented and often enabled by default in a large number of home gateway products made by several different vendors. This is unfortunately in accordance with Microsoft's [requirements for Windows-compatible home routers](#). Due to the fact that not all operating systems and applications de-prioritize transitional IPv6 connectivity, it is recommended to disable it (especially 6to4) whenever possible.

Due to the operational problems with 6to4, the IETF has recently moved towards deprecating it and discouraging any further use of it, by adopting [I-D.ietf-v6ops-6to4-to-historic](#) as a IETF v6ops Working Group document.

Android

Earlier versions of Android preferred transitional IPv6 connectivity above IPv4, as its system resolver did not implement RFC 3484.

Solution: Upgrade Android to version 2.2 Froyo or later.

Apple Mac OS X

Mac OS X, in versions earlier than 10.6.5, do not de-prioritize 6to4 compared to IPv4.

Solution: Upgrade Mac OS X to the latest available version.

Versions predating 10.6.x «Snow Leopard»

At the time of writing, the patch that de-prioritizes 6to4 is only available for the «Snow Leopard» series (10.6.x). Avoiding 6to4 entirely is the best solution, but might not always be possible or feasible when another device in the network is announcing itself as a 6to4 router. Users on old PowerPC systems have no upgrade path from 10.5.

Work-around: Use the latest version of [Google Chrome](#), which will mask the underlying problem.

Work-around: Disable IPv6 completely in the operating system: *System Preferences -> Network -> Advanced -> TCP/IP -> Configure IPv6 -> Off*.

GNU/Linux

The GNU C Library will de-prioritize transitional IPv6 connectivity if the local IPv4 address is not a private one (RFC 1918), due to a strict interpretation of RFC 3484 (see [SW#11438](#)).

Solution: Upgrade the Linux distribution to [Debian Squeeze](#), [Fedora 13 Goddard](#), [Ubuntu 10.04 Lucid Lynx](#), [Mandriva 2010.1 Spring](#), [openSUSE 11.3](#), [Gentoo 2010-04-25](#), or any later versions.

Alternative solution: Add the following lines to the file `/etc/gai.conf` (create it if it doesn't exist):

```
scopev4 ::ffff:169.254.0.0/112 2
scopev4 ::ffff:127.0.0.0/104 2
scopev4 ::ffff:0.0.0.0/96 14
```

Microsoft Windows

MS Windows automatically enables Teredo and 6to4 whenever possible. The system resolver library will de-prioritize their use, but not all applications use the system resolver library and will therefore end up preferring the less reliable transitional IPv6 connectivity.

Solution: Disable 6to4 and Teredo, by entering the following commands in an Administrator shell:

```
netsh interface ipv6 6to4 set state disabled
netsh interface teredo set state disabled
```

Alternatively, visit [Microsoft KB#929852](#), and choose the **Disable IPv6 tunnel interfaces** fix (solution 50443).

Internet Connection Sharing

If Internet Connection Sharing is enabled on **any** interface (it doesn't even have to be connected), the Windows hosts will announce itself as an IPv6 router to the local network. This will in turn cause problems for other operating systems that doesn't de-prioritize transitional IPv6 connectivity, notably Mac OS X (see above).

Solution: Disable 6to4 and Teredo as described above, and also ensure Internet Connection Sharing is disabled whenever unused.

Opera web browser

The Opera web browser uses its own resolver library, and would in versions earlier than 10.50 (on Microsoft Windows) and 10.63 (on Mac OS X and Linux) prefer transitional IPv6 connectivity above IPv4.

Solution: Upgrade Opera to the latest available version.

Home gateways and broadband routers

Apple

The AirPort and Time Capsule wireless base stations will advertise the prefix `::/64` on their LAN interfaces if they are set up with IPv6 mode **tunnel** while having a private (RFC 1918) address assigned to their WAN interfaces. As the resulting auto-configured addresses are invalid, hosts on the LAN will suffer dual-stack breakage. The bug was last reported in firmware version 7.4.2, and is confirmed to be fixed in firmware version 7.5.2).

Solution: Upgrade the firmware of the AirPort/Time Capsule to the latest available version ([instructions](#)).

AVM

Certain AVM FRITZ!Box models (at least 7270 International v2) is known to have ULA (RFC 4193) functionality based on recommendations from IETF [-D.ietf-v6ops-ipv6-cpe-router](#) revision 07 or earlier (the advertised router lifetime is greater than 0). When IPv6 is enabled on the device, it will announce a ULA prefix, which is withdrawn as soon as a global prefix is available, including a 6to4-derived one. Like ULA, 6to4 is enabled by default by the overall IPv6 configuration setting.

If the device is used in a IPv4-only environment with CGNAT/NAT444 where the WAN interface is numbered using RFC 1918 addresses, 6to4 is never activated and the ULA addresses will continue to be announced. Because ULAs are globally scoped addresses that are not special-cased in RFC 3484, this will cause hosts on the LAN to attempt to use them when connecting to dual-stacked destination, which cannot work.

Solution: Ensure ULA functionality is disabled. This can be done from the router's web interface - instructions follows:

1. Open <http://fritz.box/> with your web browser.
2. Enter your FRITZ!Box password.
3. Click *Settings* (towards the top of the window).
4. Click *Internet*.
5. On the left, click *Account Information*.
6. Select the *IPv6* tab on the right side, towards the top of the window, under the heading *Account information*.
7. You are now on the IPv6 settings page. The last section of this page is called *Unique Local Addresses*. In this section, select **do not assign unique local addresses (ULA) (not recommended)**.
8. Click *Apply*.

Cisco Linksys

E2000, E3000, and E4200

These models (and possibly others which share the same code) offers by default DHCPv6 service that assigns addresses to the LAN hosts from the documentation prefix 2001:db8::/32. However, no ICMPv6 Router Advertisements are being sent in this case, which causes the LAN hosts to not have a default route, thus causing the host operating stack to generate fast internal failures when IPv6 is attempted.

They also enable 6to4 router functionality by default. If 6to4 is active (i.e. if the WAN interface is numbered using a public IPv4 address), ICMPv6 Router Advertisements will be transmitted to the LAN segment, containing a Prefix Information Option for the 6to4 prefix.

Luckily, the networking stack in Microsoft Windows will immediately remove the DHCPv6-assigned bogon address from 2001:db8::/32 upon receiving an ICMPv6 Router Advertisement. As Microsoft Windows is the only major operating system that supports DHCPv6 (and uses it even when there are no ICMPv6 RAs), end users are unlikely to experience brokenness due to this issue.

The recent firmware revisions disable both 6to4 and the DHCPv6 service with the bogon addresses by default, solving all known problems.

Solution: Upgrade the router firmware to the firmware version indicated below (or newer):

- E2000: [1.0.04 \(build 07\)](#)
- E3000: [1.0.04 \(build 06\)](#)
- E4200: [1.0.04 \(build 07\)](#)

Select the *Downloads* tab on the pages linked to above in order to acquire the latest firmware image and its corresponding release notes.

As of 10/8/2011, the E2000 link above is 404 and the E4200 link leads to firmware 1.0.03 (released 9/28/11).

Firmware 1.0.03 of E4200 suffers from the same problem as the Fritz in that it provides ULA addresses which Windows 7 uses as source addresses. There does not appear to be a way to disable ULA on the E4200 router in 1.0.03.

On a related note: Firmware 1.0.00, which performs tunneling by default, but does not provide ULA, works out of the box. Do not update to 1.0.03 if your configuration is working.

IPv6/6to4 settings can be configured from (assuming 192.168.1.1 is the IPv4 address of the device):

- E4200: <http://192.168.1.1/SystemConfig.asp>
- E3000: <http://192.168.1.1/System.asp> (the on/off toggle is labeled *Vista Premium*)
- E2000: Unknown - probably one of the above.

WRVS4400N

This model (and possibly others) uses the bogon range 2005:123:456:789::/64 as its default DHCPv6 pool. These addresses cannot be used to communicate with the IPv6 internet, and their presence on end-user hosts will cause end-user brokenness.

Solution: Ensure the router is operating in IPv4-only mode. This setting is found under *Setup -> IP Versions* in its web interface.

D-Link

Several D-Link models from the *DSL* series (at least DSL-584T, DSL-G604T, DSL-G624T, DSL-G664T, and DSL-G684T), do not correctly forward DNS responses for hostnames with both A and AAAA records published. What it does is to stuff the first 32 bits of the AAAA record into the A record that's being returned to the end user's computer. In other words, *getipv6.info* will incorrectly resolve to 32.1.5.0 (the 2001:0500: part of the IPv6 address). If the operating system or web browser prefers to use IPv4, it will be unable to connect to the destination.

Italian ISP Wind/Infostrada is reported to have distributed the DSL-G624T to its customer base over a period of several years.

It doesn't happen all the time - it appears to be timing-dependent. Older Mozilla Firefox browsers are hit particularly bad, due to the fact that they will request AAAA lookups even if the local host does not have an IPv6 address.

Work-around: Disable DNS forwarding support in the router. This will cause the D-Link to advertise the ISP's upstream DNS resolvers (instead of itself) in DHCPv4, and the hosts on the LAN will query them directly. Instructions follows:

Log in to your D-Link box as user 'admin'.
On the Home page, select DNS from the left hand menu.
[This is not shown in the documentation, and may not be in the same menu in all devices. If you cannot find it, there may be no easy solution for this model.]
On the DNS page, select DNS Relay Selection/Disable DNS Relay
Then click Apply and Save.

To make the change permanent, go back to the Home page, select Tools/System/Save & Reboot. Your D-Link will then take a minute or so to restart itself.

Bugs in operating system TCP/IP stacks

Apple iOS

No fallback from IPv6 to IPv4

Apple iOS is not able to fall back from IPv6 to IPv4, if the initial IPv6 connection attempt fails due to blackholing. An error message is displayed after a timeout of about 60 seconds instead. If errors are generated by the network, on the other hand, it will successfully fail over; either instantly (TCP RSTs), or after four seconds (ICMPv6 unreachable).

Short of ensuring that any IPv6 connectivity present work perfectly, there is no known workaround. IPv6 cannot be disabled in iOS.

This bug is reported to Apple in #8702877. The latest iOS version the bug has been confirmed to be present in is 4.3.

Solution: Upgrade iOS to the latest available version - it is reported to be fixed in iOS 5, which includes an «Happy Eyeballs» implementation.

Apple Mac OS X

Use of IPv6 with only link/site-local addresses

When attempting to connect to a dual-stacked destination when the system has only link-local IPv6 addresses (but at the same time a default IPv6 route), a 75 second timeout is incurred per AAAA record. IPv6 is preferred over IPv4 in this case due to lack of support for RFC 3484 in Mac OS X's system resolver. A system will be vulnerable to this bug if it has received an ICMPv6 Router Advertisement message that does not contain any Prefix Information Options, for example.

Portugese ISP SAPO is known to distribute a CPE to its customers (Pirelli A1000G) that emits such RAs by default, and the D-Link DIR series of routers will do so as well (at least DIR-635, DIR-655, DIR-825, and DIR-855).

Hosts running Microsoft Windows (at least Vista) with Internet Connection Sharing enabled can also emit such prefix-less RAs, in certain situations (possibly limited to when two network interfaces are bridged).

This will especially affect users of Mozilla Firefox older than version 4.0, as it will request AAAA records even though the system does not have any global IPv6 addresses (see [bug #614526](#)), and will attempt to connect to the IPv6 addresses in preference to falling back to IPv4. Safari also suffered from this problem up until Mac OS X 10.6.4. It will also affect users of the virtualisation software Parallels Desktop (regardless of the browser used), because its virtual network interfaces have site-local IPv6 addresses assigned, in turn making the `AI_ADDRCONFIG` flag to `getaddrinfo()` ineffective.

Similarly, the "ssh" client application (accessible in a MacOS X shell) initiates connections to global scope IPv6 addresses even if the OS only has link-local addresses in its interfaces. Bug with ID 8820407 is open with Apple.

Solution: Upgrade to Mac OS X 10.6.8. The bug is fixed there by sorting IPv6 addresses last in `getaddrinfo()`'s result set when there's only link-local IPv6 addresses present on the host.

Versions predating 10.6.x «Snow Leopard»

There is no fix this bug available for older versions of Mac OS X like «Tiger» and «Leopard». While it is preferred to upgrade to «Snow Leopard»

to get the bug fixed properly, this might be undesired due to it not being a free upgrade, or entirely impossible if the hardware used is PowerPC-based. One of the following work-arounds might help for users of old Mac OS X versions:

Work-around: If the user is using Mozilla Firefox, ensure it is upgraded to [version 4.0](#) or newer. This *might* avoid the problem by suppressing AAAA lookups.

Work-around: Use the latest version of [Google Chrome](#), which will mask the underlying problem.

Work-around: Disable IPv6 completely in the operating system: *System Preferences -> Network -> Advanced -> TCP/IP -> Configure IPv6 -> Off*

Handling of Router Advertisements with a lifetime of 0

When receiving an Router Advertisement packet with a lifetime of 0, Mac OS X-based hosts will incorrectly install a default route. It appears that a Prefix Information Option must be present in the RA packet for the bug to take effect.

The technique of using a lifetime of 0 to announce an IPv6 router without global connectivity is used by devices conforming to RFC 6204 in the case where ULA addressing are used within the residential site. In this situation the Mac OS X host will attempt to use ULA IPv6 addresses for global connectivity, leading to timeouts towards dual-stacked content.

Apple bug ID: 8705091.

Solution: Upgrade to Mac OS X 10.6.8, where this bug is fixed.

Versions predating 10.6.x «Snow Leopard»

There is no fix this bug available for older versions of Mac OS X like «Tiger» and «Leopard». While it is preferred to upgrade to «Snow Leopard» to get the bug fixed properly, this might be undesired due to it not being a free upgrade, or entirely impossible if the hardware used is PowerPC-based. One of the following work-arounds might help for users of old Mac OS X versions:

Work-around: Use the latest version of [Google Chrome](#), which will mask the underlying problem.

Work-around: Disable IPv6 completely in the operating system: *System Preferences -> Network -> Advanced -> TCP/IP -> Configure IPv6 -> Off*

Handling of ICMPv6 Destination Unreachable

When a router in the network responds to a TCP SYN packet with an ICMPv6 Destination Unreachable, Mac OS X will re-transmit the TCP SYN packet five times with intervals of one second before giving up, even though the ICMPv6 code indicates that the situation is permanent and further attempts are pointless (e.g. Code 2 - Beyond scope of source address). In other words, an avoidable four-second timeout is incurred for every outgoing TCP connection.

Work-around: Use the latest version of [Google Chrome](#), which will mask the underlying problem.

Work-around: Upgrade to Mac OS X 10.7 «Lion», which will mask the underlying problem (only works if Safari is used)

Microsoft Windows

No «un-deprecation» of SLAAC-assigned IPv6 addresses

If a Windows Vista/7 machine has a SLAAC-assigned IPv6 address that becomes deprecated (due to its Valid Lifetime countdown timer reaching zero), it fails remove the **deprecated** flag after receiving an RA with a Valid Lifetime for the same prefix of >0 (the timers themselves will however be refreshed). This could happen for example if the user's CPE router intentionally deprecated the prefix in response to a WAN link failure event (and later attempted to un-deprecate it after the WAN link came back up and was assigned the same prefix), or if the computer was suspended/hibernated for a duration longer than the remaining Valid Lifetime and then woken up again (while remaining connected to the same network as before).

In the typical case, the problem will cause IPv4 to be used in preference to IPv6 in situations where it should have been the other way around. It will therefore not cause a «dual-stack brokenness» problem for content providers, however it might lead to a performance impact in the case where the IPv4 connectivity is inferior to IPv6 (for example if there's a CGN in the IPv4 path).

Solution: Install update KB2563894. This is available through Windows Update as an important update, and may also be downloaded from [security bulletin MS11-064](#).

Handling of ICMPv6 Destination Unreachable

Microsoft Windows appears to ignore received ICMPv6 Destination Unreachable completely, instead falling back on the overall connection timeout mechanism. This causes a completely avoidable 21 second long timeout to occur for every outbound TCP connection.

Bugs in web browsers

Mozilla Firefox

Mozilla Firefox does not set the AI_ADDRCONFIG flag when looking up names using the system `getaddrinfo()` library function, which causes it to solicit AAAA records even though the system has no IPv6 addresses (non-loopback and non-linklocal). This can trigger other problems, such as the D-Link AAAA mangling bug and the Mac OS X bug regarding the use of link-local IPv6 addresses when connecting to global destinations. The bug is fixed as of Firefox 4.0, see the [bug report](#)

Solution/work-around: Upgrade to Firefox [version 4.0](#).

Opera

When Opera, in versions older than 11.10, is used on a Mac OS X machine that has Parallels Desktop installed, it is unable to connect to dual-stacked web sites (except if the machine in question also has connectivity to the IPv6 internet). This is Opera Software bug DSK-326913.

Solution: Upgrade Opera to [version 11.10 or later](#).

Problematic network deployments

NTT NGN (Japan)

The NTT Next Generation Network includes the physical last-mile infrastructure (the local loop) to the customer's premises. Other ISPs lease these lines in order to provide internet service, similar to a LLU arrangement.

However, the NGN infrastructure also includes a walled-garden IPv6 deployment, which is used to deliver IPTV services, at least. While this IPv6 connectivity cannot communicate with the Internet at large, it uses globally scoped IPv6 addresses that looks just like ordinary IPv6 internet connectivity to the devices on the residential LAN. To compensate for this, the NGN centrally spoofs TCP RST packets for all connections that attempt to cross the walled garden boundary and get out to the global Internet. While this limits the impact on end users, it still is known to cause at least 1-second connection timeouts and failed image loading on (older) MSIE browsers. There's also a concern that the central TCP RST generators won't be able to keep up with the load, once a significant number of popular destinations on the Internet deploy IPv6.

Work-around: Use the latest version of [Google Chrome](#), which will mask the underlying problem.

Work-around: Install a modified RFC 3484 policy table that specifically de-prefers the use of the NGN IPv6 prefixes for communication with global destinations, see <http://www.attn.jp/maz/p/i/policy-table/> for more information.

Alternate work-around: Disable IPv6 in the operating system.

Firewall Config Issues

Especially if you have users on Vista. It does this IPv6 tunnelling thing that on the surface appears really cool. When you try and talk IPv6 to something other than link-local: (in order)

- If you have a non-RFC1918 (ie. 'public') address, it fires up 6to4.
- If you have an RFC1918 address, it fires up Teredo.
Seems cool in theory, and you'd think that it would really help global IPv6 deployment - I'm sure that's how it was intended, and I applaud MS for taking a first step. But in practice, however, this has essentially halted any IPv6 /content/ deployment that people want to do, as user experience is destroyed.

You can help, though - here's the problem:

6to4 uses protocol 41 over IP. This doesn't go through NAT, or stateful firewalls (generally). Much like GRE.

Because of this, if you're a enterprise-esque network operator who runs non-RFC1918 addresses internally and do NAT, or you do stateful firewalling, PLEASE, run a 6to4 relay on 192.88.99.1 internally, but return ICMPv6 unreachable/admin denied/whatever to anything that tries to send data out through it. Better yet, tell your firewall vendor to allow you to

inspect the contents of 6to4 packets, and optionally run your own 6to4 relay, so outgoing traffic is fast.

Even if you don't want to deploy IPv6 for some time, do this at the very least RIGHT NOW, or you're preventing those of us who want to deploy AAAA records alongside our A records from doing so. If you need configs for <vendor/OS B/C/J/L>, post a message to the [NANOG list](#) and I'll write some templates.

I see this sort of IPv4 network quite commonly at universities, where students take their personal laptops and throw them on the campus 802.11 network. While disabling the various IPv6 things in Vista at an enterprise policy level might work for some networks, it doesn't for a university with many external machines visiting. So, if you're a university with a network like this (ie. most universities here in NZ, for example), please spend a day or two to fix this problem in your network - or better yet, do a full IPv6 deployment.

--Nathan Ward

Allow ICMPv6 through firewalls

Ensure that your firewalls allow through ICMPv6 types 1-4 and 128/129 as per NIST recommendation. In particular if you inadvertently block type 2 (Packet too big), you may find that users behind tunnels can connect to a web site but not get any content back.

--John Gibbins

Increased Latency to your IPv6 Content

If you do deploy an IPv6 network for your content, [set up a Teredo relay](#), and point 2001::/32 at it. Your viewers/users will automatically use this relay when accessing your content, and their traffic to you will be over IPv4, all they way from their PC to your network - so, equivalent performance as IPv4. Note that I say relay here, not server.

Mozilla.org are doing this for example. Cue Matthew Zeier.

--Nathan Ward

Check out [Enabling IPv6 on a Mail Server](#).

Unfortunately, not all Domain Registrars are providing IPv6 Glue yet

It may be tough to get IPv6 AAAA records for your nameservers into the DNS Glue records, depending on your registrar.

Check out [DNS Registrars IPv6 Support Status](#), let's build a list of who does and does not.

Troubleshooting Steps

The following steps should be followed in order. After each step, try testing again to see if your problem has been resolved. In general, keep your operating system and software up to date.

Update your Operating System

Windows

Go to <http://update.microsoft.com> for the latest version.

Mac OS X

Go to the Apple menu, choose Software Update, and update all available updates.

Update your browser or application

Internet Explorer

Go to <http://update.microsoft.com> and follow the instructions.

Mozilla Firefox

Firefox is normally updated automatically, but you can easily check. In Firefox, click Help, then Check for Updates. Get the New Version, then Restart Firefox.

Chrome

If a wrench icon appears on the browser toolbar, click it. Select "Update Google Chrome" then Restart.

Opera

Go to <http://www.opera.com/download/> to get the latest version, and follow the prompts.

Safari

Should have been updated when you updated your operating system: Go to the Apple menu, choose Software Update, and update all available updates.

Update your home gateway

Check the website of your gateway vendor to learn how to update it.

Disable transition technologies

Windows

Go to <http://support.microsoft.com/kb/929852> and select the "Disable IPv6 tunnel interfaces" fix.

Or, from a Windows command prompt: Click Start > Run > cmd (enter).

Then type (without the quotation marks) "netsh interface ipv6 6to4 set state disabled" and press enter.

Finally, type (without the quotation marks) "netsh interface teredo set state disabled" and press enter.

Mac OS X

Teredo is not commonly used by Mac OS X. Updating to the latest version of OS X should resolve any problems with 6to4, but if you cannot upgrade, you may have to disable IPv6 altogether. Continue troubleshooting, until you have no other choice.

Configure your firewall

Ensure that your firewalls allow through ICMPv6 types 1-4 and 128/129 as per NIST recommendation. In particular if you inadvertently block type 2 (Packet too big), you may find that users behind tunnels can connect to a web site but not get any content back.

If you are using tunnels intentionally, permit protocol 41 traffic.

Special cases

If you are using an AVM FRITZ!Box, follow the directions above to disable ULA.

If you are using a Linksys WRVS4400N, follow the directions above to disable IPv6 on the gateway.

If the problem still isn't fixed

Use Mac OS X 10.7 «Lion» (Safari only)

Mac OS X 10.7 «Lion» has high-level APIs that implement Happy Eyeballs-ish parallel connections, which efficiently mask any connection problems related to IPv6 (or IPv4 for that matter). More information about this [here](#). Currently, Safari is the only major browser to utilise these APIs.

Use Google Chrome

The latest version of [Google Chrome](#) has built-in robustness (Happy Eyeballs) when it comes to dual-stack brokenness, which will in the worst case result in a 300ms delay once per host, instead of long timeouts for every single connection. While it doesn't cure the underlying problem, it will mask the symptoms in a very efficient manner.

Disable IPv6

If no other workaround or attempted fix has remedied the problem, the last thing to try is to disable IPv6 outright. While this will prevent the user from successfully using IPv6 once the underlying problem has been fixed, it is likely to help improve the user experience in the short term.

Windows

Go to <http://support.microsoft.com/kb/929852> and select the "Disable IPv6 except for loopback" fix.

Mac OS X

Go to the Apple menu, select System Preferences, then Network (make sure Built-In Ethernet is selected) and choose Advanced, TCP/IP, then Configure IPv6 and select "off." Then choose OK and Apply.