

**"glibc" package**[Log in / Register](#)[Overview](#)[Code](#)**Bugs**[Blueprints](#)[Translations](#)[Answers](#) **Bug #313218:**Bug #313218 reported by [sojourner](#) on 2009-01-02 ([Activity log](#))

IPV6 causes slow internet access

[This report is public](#) [Mark as duplicate](#) (This bug has 3 duplicates)[Convert to a question](#)

Affects	Status	Importance	Assigned to	Milestone
glibc (Fedora)	Fix Released	Unknown	redhat-bugs #459756	
glibc (Ubuntu)	Fix Released	High	Colin Watson	
Jaunty	Fix Released	High	Colin Watson	

[Also affects project](#)[Also affects distribution](#)[Nominate for release](#)

starting with kernel 2.6.28-4 ipv6 is built in and cannot be disabled . this causes very slow internet access for people who have routers or isp's that are non compliant . Prior to 2.6.28-4 ipv6 was loaded as a module and could be disabled by setting " alias net-pf-10 ipv6 off " in /etc/modprobe.d/aliases and also adding blacklist ipv6 to /etc/modprobe.d/blacklist . please provide a method of disabling IPV6 for those who need it .

ProblemType: Bug

Architecture: amd64

DistroRelease: Ubuntu 9.04

Lsusb:

Bus 002 Device 001: ID 1d6b:0002 Linux Foundation 2.0 root hub

Bus 001 Device 001: ID 1d6b:0001 Linux Foundation 1.1 root hub

NonfreeKernelModules: nvidia

Package: linux-image-2.6.28-4-generic 2.6.28-4.5

ProcCmdLine: root=UUID=1f603d8b-0fe4-4846-92f9-54b572ff1924 ro quiet splash

ProcEnviron:

SHELL=/bin/bash

LANG=en_US.UTF-8

ProcVersionSignature: Ubuntu 2.6.28-4.5-generic

SourcePackage: linux

Duplicates of this bug

[Bug #311941](#) [Bug #312104](#) [Bug #326718](#)Tags: [apport-bug](#) [qa-jaunty-foundations](#) [regression-potential](#)[Update description / tags](#) (see [original description](#))











Related branches

[~cjwatson/ubuntu-toolchain/bug313218 \(Development\)](#) Fix Available[Link another branch](#)[Link to CVE](#)











sojourner wrote on 2009-01-02: ([permalink](#))

[Subscribe someone else](#)



Subscribers

-  [BootDmesg.txt](#) (43.2 KiB, text/plain; charset="utf-8")
-  [CurrentDmesg.txt](#) (2.1 KiB, text/plain; charset="utf-8")
-  [Dependencies](#) (1018 bytes, text/plain; charset="utf-8")
-  [HalComputerInfo.txt](#) (3.1 KiB, text/plain; charset="utf-8")
-  [Lspci.txt](#) (14.3 KiB, text/plain; charset="utf-8")
-  [ProcCpuinfo.txt](#) (1.4 KiB, text/plain; charset="utf-8")
-  [ProcInterrupts.txt](#) (1.2 KiB, text/plain; charset="utf-8")
-  [ProcModules.txt](#) (2.3 KiB, text/plain; charset="utf-8")
-  Colin Watson
-  Daniel Stovannov

Neal McBurnett wrote on 2009-01-02: ([permalink](#))





-  umizer
-  François Letendre
-  Thank you for reporting this issue.
-  geckon
-  Gonzo
-  Please provide more details on why it seems slow. Where do you see the delays? How is your system and network configured? How can this be reproduced for testing?
-  Gonçalo Marrafa
-  Jack Wasey
-  Jason Harman
-  JoeKlein

Neal McBurnett on 2009-01-02

-  Changed in linux:
-  **status:** New → Incomplete

-  Matteo Gazzoni
-  Mike

sojourner wrote on 2009-01-02: ([permalink](#))

-  pashok99
-  network is configured by network manager as open (no encryption) network is wired with a belkin f5d7230-4 router and mororola cable modem , ISP is comcast (us) , system is jaunty ,currently instaled with daily alternate and64 destop dated dec 31 08. I can tell ipv6 is the problem because prior to kernel 2.6.28-4 with ipv6 disabled, website access was fast and loading was fast and smooth, with ipv6 enabled there was a long delay >10sec
-  Tactics Ting, any web page and loading was slow and choppy (burst of data pause burst of data pause finally loads)
-  this was evident in all browsers FF 3.05,3.1 (PPA),3.2 (PPA) , opera 64bit 9.62 and 10 beta (builds 4103 and 4116) , epiphany , dillo, seamonkey and galeon . connecting to a different network (wireless) with a different router the problem dissapears , connecting with the original network the problem dissapears , connecting directly to the cable router the problem dissapears , these results are duplicated with another system running ubuntu 8.04 (lts) upto date . In my case it appears to be the router that is the bottleneck but there are reports in the forums and elsewhere about the same thing with noncompliant ISPs . You will not be able to duplicate this problem unless you have a noncompliant router or ISP , please feel free to request any tests I can perform for you .

Also notified

-  Adam Conrad
-  Andrew

Jonathan Ernst wrote on 2009-01-02: ([permalink](#))

I think my problem is related. Since updating to jaunty I cannot use my default dns server (my router) and I have to change my resolv.conf to use my isp's resolv.conf (my other computers running intrepid and other OSs have no problem with my router's dns handling). See [Bug #312104](#)

Jonathan Ernst wrote on 2009-01-02: ([permalink](#))

> use my isp's resolv.conf

I meant use my isp's dns server

Chad Waters wrote on 2009-01-02: ([permalink](#))

Regardless of the router slowness: the issue is that the choice to disable it *apparently* is no longer there.

I currently have daemons listening on a non-existent IPv6 network. Likewise if I had IPv6, the choice to disable IPv4 would be appreciated.

Thanks

JoeKlein wrote on 2009-01-02: **Re: [Bug 313218] Re: IPV6 causes slow internet access** ([permalink](#))

All,

I am unsure that it would be considered a bug, instead, a set of decisions, based on your environment (routers, hosts, ISP, security) and a set of configurations. I am working on a decision tree and documentation to help fix this problem. Please give me a few weeks.

Joe Klein

On Fri, Jan 2, 2009 at 2:31 PM, Chad Waters <[email address hidden](#)> wrote:

[\[...\]](#)

sojourner wrote on 2009-01-02: ([permalink](#))

in my opinion it is a bug and also a regression in that it causes problems for people with some very common hardware . atleast here in the US (the only country for which I can speak) belkin is a very well known and widely available

brand and my model is their "low end" and low price model . It has functioned very satisfactorily up to the point when IPV6 was built into the kernel, and by failing to resolve this problem we will essentially be locking out many potential users . the same applies to non compliant ISPs since here in the US ,broadband is often only available from your monopoly phone company or monopoly cable company .

Neal McBurnett wrote on 2009-01-03: ([permalink](#))

Sounds like we need to get to the bottom of what is causing the timeouts. Hopefully it can be resolved in a way that is not slow, and also allows both ipv4 and ipv6 without requiring manual configuring or disabling by the user.

I expect that it would help to compare tcpdump traces on your host accessing something that demonstrates the problem Both with and without the delays, but with all else as similar as possible. But perhaps the procedure Joe is working on incorporates that, or is even nicer.

sojourner wrote on 2009-01-03: ([permalink](#))

how do I do a tcpdump trace ? I can easily configure my system to have or not have the delay , simply unplug a cat5 cable and enable a wireless network or if you want it done on the same box/router/isp it will have to be done under an earlier kernel because in 2.6.28-4 ipv6 is not loaded as a module and cannot be disabled .

marmuta wrote on 2009-01-03: ([permalink](#))

Replace the ip with the one of your router.

This dumps dns traffic only, probably what you need:
`sudo tcpdump -nls128 host 192.168.1.1 and port 53 | tee tcpdump.txt`

This dumps everything going from/to your router:
`sudo tcpdump -nls128 host 192.168.1.1 | tee tcpdump.txt`

marmuta wrote on 2009-01-03: ([permalink](#))

There are other problems too with forcing ipv6 enabled for everyone.

Some apps aren't prepared well and totally drop ip4 connectivity.

In the short time since the change to build-in ipv6 I had already vino-server ([Bug #196675](#)) and gobby ([Bug #313393](#)) fail me. I'm sure there are more waiting in the repos.

It's a good thing to bring these issues to light now, but I wouldn't hold my breath for having all of them fixed until release day. There is a real need for being able to turn ipv6 off completely and so far there appears to no other reliable way besides having it build as a module.

Shirish Agarwal wrote on 2009-01-03: ([permalink](#))

Hi Marmuta,

What do I do, I have two ethernet ports, one (eth0) not used, the other (eth1) being used

```
~$ ifconfig
```

```
eth0 Link encap:Ethernet HWaddr 00:07:95:44:10:db
      UP BROADCAST MULTICAST MTU:1500 Metric:1
      RX packets:0 errors:0 dropped:0 overruns:0
```

```
frame:0
```

```
      TX packets:0 errors:0 dropped:0 overruns:0
```

```
carrier:0
```

```
      collisions:0 txqueuelen:1000
```

```
      RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
```

```
      Interrupt:18 Base address:0xc800
```

```
eth1 Link encap:Ethernet HWaddr 00:08:a1:92:56:33
      inet addr:192.168.1.3 Bcast:192.168.1.255
```

```
Mask:255.255.255.0
```

```
      inet6 addr: fe80::208:a1ff:fe92:5633/64
```

```
Scope:Link
```

```
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
```

```
      RX packets:6801947 errors:0 dropped:0 overruns:0
```

```
frame:0
```

```
      TX packets:7413867 errors:0 dropped:0 overruns:0
```

```
carrier:0
```

```
      collisions:0 txqueuelen:1000
```

```
      RX bytes:4229253367 (4.2 GB) TX bytes:2328395793
```

```
(2.3 GB)
```

```
      Interrupt:22 Base address:0xcc00
```

```
lo Link encap:Local Loopback
```

```
      inet addr:127.0.0.1 Mask:255.0.0.0
```

```
      inet6 addr: ::1/128 Scope:Host
```

```
      UP LOOPBACK RUNNING MTU:16436 Metric:1
```

```
      RX packets:5403 errors:0 dropped:0 overruns:0
```

```
frame:0
```

```
      TX packets:5403 errors:0 dropped:0 overruns:0
```

```
carrier:0
```

```
      collisions:0 txqueuelen:0
```

```
      RX bytes:1102931 (1.1 MB) TX bytes:1102931 (1.1
```

```
MB)
```

Can you tell how should I give my tcpdump?

Using the above I get this :-

```
$ sudo tcpdump -nls128 host 192.168.1.1 | tee tcpdump.txt
[sudo] password for shirish:
tcpdump: WARNING: eth0: no IPv4 address assigned
tcpdump: verbose output suppressed, use -v or -vv for full
protocol decode
```

```
listening on eth0, link-type EN10MB (Ethernet), capture
size 128 bytes
^C0 packets captured
0 packets received by filter
0 packets dropped by kernel
```

This is though on Intrepid atm but still would love to know how.

--

Regards,
Shirish Agarwal

This email is licensed under <http://creativecommons.org/licenses/by-nc/3.0/>

<http://flossexperiences.wordpress.com>


065C 6D79 A68C E7EA 52B3 8D70 950D 53FB 729A 8B17

[Dean Loros](#) on 2009-01-03

Changed in linux:


status: Incomplete → Confirmed

[sojourner](#) wrote on 2009-01-03: ([permalink](#))

 [tcpdump2.txt](#) (5.2 KiB, text/plain)

attached are tcpdumps with ipv6 disabled (tcpdump1.txt) and enabled and loaded (module) . Note these were taken with my hardy box (ubuntu 8.04) since that was easier to do , it shows exactly the same symptoms with ipv6 enabled . hardy kernel is 2.6.24-22-generic . I can do tcpdumps with my jaunty box and kernel 2.6.28-4 later if you need them , just a a little more work .

[sojourner](#) wrote on 2009-01-03: ([permalink](#))

 [tcpdump1.txt](#) (3.1 KiB, text/plain)

opps dump1 didn't attach here it is

[marmuta](#) wrote on 2009-01-03: ([permalink](#))

@Shirish Agarwal

Try this:

```
sudo tcpdump -nls128 -i eth1 | tee tcpdump.txt
```

[marmuta](#) wrote on 2009-01-03: ([permalink](#))

sojourner, did your router reboot during tcpdump2.txt?
It didn't respond for >30s beginning from 06:28:31.

sojourner wrote on 2009-01-03: ([permalink](#))

no that was a period where nothing was loading , after my initial pages loaded I paused to take a peek at what had been written so far and then loaded a couple of more pages

.

JoeKlein wrote on 2009-01-03: **Re: [Bug 313218] [NEW] IPv6 causes slow internet access** ([permalink](#))

Problem Space:

Below is a detailed definition of the problem space I have encountered when implementing IPv6 over the last 6 years across many different system and network environments.

A. **Network Stack:

As many of you know, the IPv6 RFC's defines a precedence order for network communications. Here is the order when IPv6 is enabled:

1. **Use Native IPv6

a. Generate link local address

b. Obtain global addresses

i. Host file address

ii. Check for stateless autoconfiguration (Router/Neighbor Discovery)

iii. Check for stateful autoconfiguration (DHCPv6)

2. **Use Tunneled IPv6 (If available)

a. Connect to tunnel endpoint (requires a DNS lookup in many cases)

b. Generate Global address

3. **Use Native IPv4

a. Host file

b. DHCPv4

c. Generate Link Local

From a booting perspective, generating link local address and obtaining global address via host file is the fast. After that, each step requires a delay and implements a timeout to fall though the logic.

B. **Name Resolution

DNS name order precedence also exists with IPv6 and it is:

1. Cache
2. AAAA over IPv6
3. A over IPv6
4. AAAA over IPv4
5. A over IPv4

For simply sake, I have not included name order precedence when LLDP, UPnP, Zero Configuration, WS-Discovery, NFS, SMB or other name management techniques are in use. Also, the DNS request should be send over IPv6 and IPv4 transport in parallel, but in many implementations, there is a timeout period. This is also under the assumption that the DNS server hard coded or obtained via DHCP, supports both IPv4 and IPv6 transport.

There are 81 variations of IPv4 only, IPv6 only, IPv4 and IPv6, when applied against the host, DNS, Internal Network Segment support, support for IPv6 at the network edge and ISP support.

Here are the following cases that I have seen in the 'wild' and some of the issues:

- *DNS Server only supports IPv4 transport, and A records*

- o Host has IPv4 only enabled - Resource is only on IPv6 (i.e. ipv6.google.com)

- § Host requests A records, over IPv4 transport and times out

- § Result: Delay and Application may fail

- o Host has IPv4 and IPv6 enabled - Resource is only on IPv6 (i.e. ipv6.google.com)

- § Host requests AAAA records, over IPv6 transport and times out

- § Host requests AAAA records, over IPv4 transport and times out

- § Result: Delay and Application mostly fails

- o Host has IPv4 and IPv6 enabled - Resource has both IPv4 and IPv6 DNS records.

- § Host requests AAAA records, over IPv6 transport and times out

- § Host requests A records, over IPv6 transport and times out

§ Host requests AAAA records, over IPv4 transport and times out

§ Host requests A records, over IPv4 transport and receives a response

§ Result: Delay and if Application has built in timers, may fail

o Host only has IPv6 - Resource is only IPv6

§ Host requests AAAA records, over Ipv6 transport and times out

§ Application timeout or locks

§ Result: Delay and Fail

- *DNS server only supports IPv4 transport, and supports A/AAAA records*


o Host has IPv4 only enab...

[Read more...](#)

[sojourner](#) wrote on 2009-01-03: ([permalink](#))

I have already purchased a new router to replace the belkin that does not play nice with IPV6 , however I do not think that this is a way that casual users or first time tryers will find acceptable . the "quick and dirty" way to handle this atleast for the time being would be to either return to the "module" method of loading IPV6 or to provide a way to switch it off if it is to remain "built in" , if it is possible to do so . In the long run IPV6 is necessary but the quick and dirty way will buy time in which to develop a more elegant solution . I will for awhile maintain my system as is with the belkin router to be used as a "test bed" for any ideas you have and would like to try .

[Shirish Agarwal](#) wrote on 2009-01-04: ([permalink](#))

 [tcpdump.txt](#) (153.6 KiB, text/plain)

[Shirish Agarwal](#) wrote on 2009-01-04: ([permalink](#))

As one can see the tcpdump.txt seems to be going fine. But what I wanna know is :-

- a. Is my router ipv6 ready or no? Is there anyway to test that?
- b. I know quite a few dns nameservers that my ISP has. Is there anyway to test if they are Ipv6 ready?

marmuta wrote on 2009-01-04: ([permalink](#))

Your router seems to do fine with ipv6 name resolution over ipv4.

Can't tell what happens with an ipv6-based network. Tried pinging inside your local network with ping6?

Your nameserver already responds to ipv6 queries:

```
dig @218.248.240.181 ipv6.google.com AAAA +short
ipv6.l.google.com.
2001:4860:c003::68
```

Mine too, yay!

```
dig @192.168.2.1 ipv6.google.com AAAA +short
ipv6.l.google.com.
2001:4860:0:1001::68
```

Jonathan Ernst wrote on 2009-01-04: ([permalink](#))

I can resolve ipv6 addresses too (but cannot access ipv6 websites) but I'm still affected by this bug.

Disabling ipv6 in firefox ([about:config](#), set network.dns.disableIPv6) fixes the issue (in firefox) for me (pages load visibly faster). Can others here test it?

Shirish Agarwal wrote on 2009-01-04: **Re: [Bug 313218] Re: IPV6 causes slow internet access** ([permalink](#))

Is there a way to know which are ipv6 websites?

--

Regards,
Shirish Agarwal

This email is licensed under <http://creativecommons.org/licenses/by-nc/3.0/>
<http://flossexperiences.wordpress.com>
065C 6D79 A68C E7EA 52B3 8D70 950D 53FB 729A 8B17

Shirish Agarwal wrote on 2009-01-04: ([permalink](#))

Ignore that. This might be a starting point.

<http://www.ipv6.org/v6-www.html>

--

Regards,
Shirish Agarwal

This email is licensed under <http://creativecommons.org/licenses/by-nc/3.0/>
<http://flossexperiences.wordpress.com>
065C 6D79 A68C E7EA 52B3 8D70 950D 53FB 729A 8B17

JoeKlein wrote on 2009-01-04: ([permalink](#))

Shirish,

If you use Firefox as a webbrowser, you can install the ShowIP plugin. It will show the IPv6/IPv4 address of the current page in the status bar.

If you are not using Firefox, my I suggest one of the following sites to "check" if you are getting both DNS resolution (IPv4 or IPv6 transit) and endpoint connection (IPv6):

Logo Animation Sites:

<http://ipv6.google.com> - The Google logo Jumps when you arrive on the site via IPv6

<http://www.kame.net/> - Dancing Turtle

Shows your IPv6 and IPv4 address

<http://www.whatismyipv6.net/>

<http://whatismyv6.com/>

<http://www.runningipv6.net/what-is-my-ipv6-address.php>

On Sun, Jan 4, 2009 at 7:28 AM, Shirish Agarwal <email address hidden>wrote:

[...]

Shirish Agarwal wrote on 2009-01-05: ([permalink](#))

JoeKlein,

I have installed show IP, it gives me Ipv4 addresses all the time, not ipv6.


Today morning, the link died unexpectadly, not on the router but for some unexplainable people not able to hook up/listen to the net.

I tried repeatedly on 2.6.28.4 as well as 2.6.28.3 but to no avail.

I had to go back down to 2.6.28.2-ub-generic to get a working link. In this one I have disabled ipv6.

Giving another tcpdump. Perhaps this would tell people something.

Shirish Agarwal wrote on 2009-01-05: ([permalink](#))

 [tcpdump.txt](#) (2.7 KiB, text/plain)

Shirish Agarwal wrote on 2009-01-05: ([permalink](#))

For a long time it was not giving any output hence in another window gave a

```
$ ping -ac 5 192.168.1.1
```

after which one starts seeing some action on that tcpdump.txt

Lemme know if anything more is needed.

[BUGabundo](#) on 2009-01-07

description: updated

[Colin Watson](#) wrote on 2009-01-13: ([permalink](#))

Although I have not investigated the network traces in this bug in any detail, people investigating it (particularly JoeKlein, who touches on similar issues) may like to bear in mind that we used to have a problem along these lines in Ubuntu, but adjusted glibc's resolver code in order to avoid it. As a result you may find that old documentation or documentation based on other systems is not as applicable as you might think at first glance. The original change was:

```
glibc (2.5-0ubuntu13) feisty; urgency=low
```

```
* debian/patches/any/local-ipv6-sanity.diff: Only do AAAA lookups if we have an interface with better than link-local addresses available.
```

```
-- Tollef Fog Heen <email address hidden> Tue, 3 Apr 2007 14:11:26 +0200
```

In the current glibc source package, you'll find that change in debian/patches/any/local-ipv6-lookup.diff.

On general troubleshooting principles, if you are encountering a problem with AAAA lookups happening when true IPv6 connectivity is not available (as should be the case with non-IPv6-capable routers or ISPs), then I would strongly advise starting by figuring out if the behaviour described in this patch has regressed.

[Colin Watson](#) wrote on 2009-01-13: ([permalink](#))

With regard to the specifics of this bug report, we would prefer, if possible, to make any necessary further fixes to glibc and other similar resolver code rather than returning ipv6 to its previous modular state in the kernel. Remember that this has only happened in a development cycle, which is a good time to try to iron out these kinds of bugs that


may affect several packages.

sojourner (or anyone else who is actually experiencing a *real* IPv6 problem; Shirish does not appear to be experiencing such a problem and as far as I can see is simply asking for help configuring IPv6), could you please post the output of 'ifconfig'? What I'm interested in is whether there is a global-scope address on any of your network interfaces.

[sojourner](#) wrote on 2009-01-14: ([permalink](#))


attached is the output of ifconfig . eth0 is the connection to the router that doesn't play nice with ipv6 . wlan0 is connected through a different router and isp and does work with ipv6 , if you wish I can shut down wlan0 if it is confusing the situation.

[sojourner](#) wrote on 2009-01-14: ([permalink](#))

 [ifconfig](#) (1.8 KiB, text/plain)

oops forgot to attach the file

[Matteo Gazzoni](#) wrote on 2009-01-17: ([permalink](#))

 [ifconfig](#) (908 bytes, text/plain)

Since Ubuntu 9.04 the dns lookups are very slow, and it's indeed ipv6 fault (proven using network.dns.disableIPv6=true in Firefox).

[Colin Watson](#) on 2009-01-30

Changed in glibc:
assignee: nobody → kamion

[Colin Watson](#) wrote on 2009-01-30: ([permalink](#))

OK, so neither of you have global-scope addresses. That points to either a regression in glibc, or some other application implementing its own resolver that doesn't play by our rules.

[Colin Watson](#) wrote on 2009-01-30: ([permalink](#))

Curiously, I can't reproduce this using any of the

applications I've tried. While I do have some level of IPv6 networking, I've deliberately turned it off for this test, and in any case I'm testing at the tcpdump level; I can clearly see the A and AAAA queries going out in parallel, rather than what sojourner was seeing which was the AAAA query going out and then a long timeout.

That means I need to figure out what's different about our environments. One useful place to start would be an strace of a fairly simple client such as w3m. Could you please:

```
* install the strace and w3m packages
* run: strace -f -s 4096 -o w3m.trace w3m http://
some/url/of/your/choice
* quit w3m once the page has loaded (shift-Q)
* attach w3m.trace to this bug
```

This may not be sufficient for me to work out what's wrong, but it will give me a bit more to go on.

Colin Watson wrote on 2009-01-30: ([permalink](#))

I wonder if this is http://sources.redhat.com/bugzilla/show_bug.cgi?id=7060 ? Using a certain amount of guesswork, I traced the problem as far as gethostbyname4_r myself, then found that bug.

That said, the comments on the patch near the end say that Firefox is *not* affected.

Matteo Gazzoni wrote on 2009-01-31: ([permalink](#))

 [w3m.trace](#) (600.0 KiB, text/html)

If Firefox is not affected then is the wrong bug, AFAIK every component in the system is affected, even Synaptic is slow to begin to download the packages.

Anyway, here's the output of strace.

Matteo Gazzoni wrote on 2009-01-31: ([permalink](#))

I think this bug (that seems related to the one Colin's posted) is worth a look: https://bugzilla.redhat.com/show_bug.cgi?id=459756

Note: I've used Fedora 9 and I've had the same issues reported here, then Fedora 10, same issues but after one month or so from release, suddenly, the dns slowness has disappeared (before I used to disable the kernel module).

Colin Watson on 2009-01-31


Changed in glibc:

importance: Undecided → Critical

[Colin Watson](#) wrote on 2009-01-31: ([permalink](#))

Looks like it, thanks. I'm travelling this weekend but will look at pulling out the relevant patches from Fedora next week; however I'm not sure I'll be able to construct an IPv6 testbed next week to make sure that we don't regress that, so an upload might have to wait until the week after.

[sojourner](#) wrote on 2009-01-31: ([permalink](#))

 [w3m.trace](#) (325.7 KiB, text/html)

here is the w3m strace from my sys , this is with a cabled connection to the router in question , no wireless .

[Bug Watch Updater](#) on 2009-02-01

Changed in glibc:

status: Unknown → In Progress

[Jordan Wilberding](#) wrote on 2009-02-10: ([permalink](#))

I am also having this problem.

I am almost positive it is not a router issue but an ISP issue. My reasoning is this. When I use my ISP's dns wget will not work. However if do wget -4(which forces ipv4) it does work. Further, if I switch the DNS server in /etc/resolv.conf to a ipv6 capable DNS server, then wget works fine without needing the -4 switch.

Therefore, it appears very likely to not be a router issue, but to be specific DNS servers not supporting the IPv6 requests coming from Ubuntu.


I would just like to add my support for the call that there is a simple way to disable ipv6 for those of us who have no need for it.

[sojourner](#) wrote on 2009-02-10: ([permalink](#))

it is probably both a problem with some routers and some dns servers , on my system bypassing the router and plugging directly into my cable modem eliminates the problem for me . the problem is that some routers and some servers are not ipv6 capable at this time so we need to be

able to disable it .

Daniel Stoyanov wrote on 2009-02-11: ([permalink](#))

 [ifconfig](#) (1.7 KiB, text/plain)

I'm also experiencing this issue. I just upgraded to Jaunty and tested the following applications :

- apt, Failed to fetch <http://archive.ubuntu.com/ubuntu/dists/jaunty/Release.gpg> Could not resolve 'archive.ubuntu.com'
- pidgin, connects to GTalk (gmail.com), but refuses to connect to ICQ (login.messaging.aol.com)
- firefox very slow, almost unusable until I set network.dns.disableIPv6 to true
- rhythmbox listening radio works without any problem (<http://mp3.live.tv-radio.com/franceinter/all/franceinterhautdebit.mp3>)

I think this bug is very annoying because some network services (like messenger for example) are very important for most users. I have attached ifconfig, if you need further info, don't hesitate.

Regards,

Colin Watson on 2009-02-12

Changed in glibc:
assignee: nobody → kamion

clerum wrote on 2009-02-15: ([permalink](#))

I think that this may be related to <https://bugs.launchpad.net/ubuntu/+source/linux/+bug/254622>

What I'm seeing on the wireshark captures is that even between two systems on the same LAN connectivity between two IPv6 addresses will be much slower due to the packets being sent with sizes much larger than the interface MTU. In cases where the switch MTU is smaller than the size sent the packets will be dropped and thus connections suffer.

Colin Watson wrote on 2009-02-19: ([permalink](#))

clerum: I honestly think this is unlikely to be related to that bug. It will probably help us if multiple maybe-similar-but-possibly-not bug reports aren't conflated into one issue; when that happens it is very difficult to know when we're making progress!

Could people affected by the original problem reported at

the top of this bug report please install the libc6 packages from here:

```
deb http://ppa.launchpad.net/cjwatson/ppa/ubuntu jaunty
main
```

You can find the necessary public key and instructions here:

<https://launchpad.net/~cjwatson/+archive/ppa>

This version of glibc incorporates a patch from Fedora which is said to address a similar-sounding issue, and I'd like to confirm whether it also fixes the problems people are encountering here.

It would be nice if anyone with real IPv6 connectivity could also test this to ensure that I haven't broken anything for them. I've had a working IPv6 setup myself in the past, but it's broken at the moment and I thought I'd get this out for testing before spending too long trying to fix it.

Thanks in advance!

[Loïc Minier](#) wrote on 2009-02-20: ([permalink](#))

I can confirm that the libc in cjwatson's PPA doesn't regress on a natively connected IPv6 host.

[Evan Dandrea](#) wrote on 2009-02-20: ([permalink](#))

I can confirm that the libc in cjwatson's PPA fixes the problem I was having on my network where Intrepid was having normal interaction with the nameserver, but jaunty was unable to reach Ubuntu domains. A strace of a failing w3m process can be found here:
<http://people.ubuntu.com/~evand/tmp/w3m.strace.txt>

[Jonathan Ernst](#) wrote on 2009-02-20: ([permalink](#))

The libc6 package from Colin Watson's PPA both fixes non-working connection in GNOME apps ([Bug #312104](#)) and makes the name resolution fast again. Thanks

[Jordan Wilberding](#) wrote on 2009-02-20: ([permalink](#))

The libc6 package from Colin Watson's PPA fixed my problem of not being able to resolve DNS entries on a DNS server that did not support IPv6.

There appears to be no ill side effects.

[sojourner](#) wrote on 2009-02-21: ([permalink](#))

I am unable to install the packages from the ppa , I had already updated to a later version by the time I saw them and the do not showup in synaptic . any suggestions ?

[Colin Watson](#) wrote on 2009-02-21: ([permalink](#))

Yes, Loic uploaded a newer version to Jaunty proper. You could install them with 'sudo apt-get install libc6=2.9-0ubuntu10ppa1', or I could upload a newer version to my PPA - but, given the confirmations above, I think I'm just going to upload 2.9-0ubuntu12 to Jaunty with this fix.

[Colin Watson](#) wrote on 2009-02-21: ([permalink](#))

Thanks to everyone who tested!

[Launchpad Janitor](#) wrote on 2009-02-21: ([permalink](#))

This bug was fixed in the package glibc - 2.9-0ubuntu12

```
-----
glibc (2.9-0ubuntu12) jaunty; urgency=low

 * debian/patches/all/fedora-nss_dns-gethostbyname4-
  disable.diff: Patch
   from Fedora 2.9-3 to temporarily disable _nss_dns_
  gethostbyname4_r,
   which caused problems for systems with broken IPv6
  connectivity
  (LP: #313218, https://bugzilla.redhat.com/show\_bug.cgi?id=459756).

 -- Colin Watson <email address hidden> Sat, 21 Feb 2009
 07:40:16 +0000
```

Changed in glibc:
status: Confirmed → Fix Released

[sojourner](#) wrote on 2009-02-21: ([permalink](#))

I have updated to libc6 2.9-0ubuntu12 from the repos , while this may give some small improvement (too small to be sure) it has not fixed my problem , which as I noted is with my router , not my ISP . at least NM now allows me to switch between eth0 (the problem router) and wlan0 through a different router without having to unplug the cable .

[Jason Waters](#) wrote on 2009-02-24: ([permalink](#))

I'm wondering if I'm having a similar issue. I'm having slow internet access when I set a static IP. If I use DHCP it works great. I have this Ubuntu 8.04 Server directly connected to a comcast router. It has 5 static IPs. I use these all the time so I just set the static ip in /etc/network/interfaces and it normally works. I get 1.5mb down. Right now if I set the static IP I get about 300k down. Then if I change my interfaces file to iface eth0 inet dhcp and reload networking I get my 1.5mb down. I have IPV6 disabled. I use the same DNS with each test. If I plug my laptop running windows directly into the router I can get 1.5mb down with DHCP or static so I don't think it is a ISP hardware problem. Maybe this is a new bug? Any thoughts? Thanks.

[dmizer](#) wrote on 2009-02-25: ([permalink](#))

@Jason Waters:
This bug is for Jaunty.

You probably do not have a bug at all. For more information, please post in the ubuntuforums, or log into #ubuntu and give a detailed explanation of your problem.

[Colin Watson](#) wrote on 2009-02-27: ([permalink](#))


sojourner: Could you repeat the previous strace command, but this time add the -tt option to strace to produce timestamps? Sorry I didn't ask for this earlier. Could you also please confirm that your ifconfig output is the same as before (I'm specifically interested in the "inet6 addr:" lines)?

I'm going to reopen this bug, since sojourner was the original reporter, but at a reduced severity since it is no longer affecting so many people. I would GREATLY appreciate it if anyone else who believes that they have a similar problem would please file a separate bug; this one has already got very confusing.

[Colin Watson](#) on 2009-02-27


Changed in glibc:
importance: Critical → High
status: Fix Released → Triaged

[sojourner](#) wrote on 2009-02-28: ([permalink](#))

 [ifconfig2](#) (1.1 KiB, text/plain)

here is the strace and ifconfig , at the top of the copy of ifconfig is the cmd line I used to run strace , note that the response is can't the url .

[sojourner](#) wrote on 2009-02-28: ([permalink](#))

 [w3m.trace2](#) (275.2 KiB, text/plain)

nuts it only attached 1 file , here is the strace .

[Jason Harman](#) wrote on 2009-02-28: ([permalink](#))

I am also having what I am guessing is this problem. It happened around the time of Jaunty Alpha 5's release as my Internet was unaffected earlier. The disconnect is between my computer and my router, which can be verified by transfers on the LAN and WAN being interrupted. Firefox and all other apps are affected.

The router is fairly new (DLink DIR 655 Wireless N Series) so I'm unsure why it would not be able to switch between IPV6 and IPV4 as the other commenter noted regarding his Belkin router.

If there is any date (ifconfig, etc.) I can add let me know.

[Colin Watson](#) wrote on 2009-03-01: ([permalink](#))

Jason, as I noted in a previous comment:

"I would GREATLY appreciate it if anyone else who believes that they have a similar problem would please file a separate bug; this one has already got very confusing."

At this point I'm afraid I am going to disregard all reports in comments in this bug that are not from the original reporter, just so that I can have some hope of keeping track.

[Colin Watson](#) wrote on 2009-03-01: ([permalink](#))


sojourner: Thanks for the data. However, it's unsurprising that you would be unable to use `ipv6.google.com`, given that you do not have IPv6 connectivity (according to your ifconfig output); so this is not a very good test.

Could you please try the following two URLs instead:

<http://www.google.com/>
<http://www.ietf.org/>


Google has only an A record (IPv4); the IETF has both an A record (IPv4) and an AAAA record (IPv6). This should exercise the two interesting cases rather than the uninteresting one.

[sojourner](#) wrote on 2009-03-01: ([permalink](#))

 [w3m.tracegoogle](#) (261.0 KiB, text/html)

I went ahead and did strace on both those url's here they are .

[sojourner](#) wrote on 2009-03-01: ([permalink](#))

 [w3m.traceietf](#) (264.3 KiB, text/html)


and the other , actually both of those loaded quickly .

[Colin Watson](#) wrote on 2009-03-02: ([permalink](#))

Thanks for those. Hmm. Indeed, neither of those shows any sign of a problem at all.

Maybe I should have phrased my question differently. Can you give me an example of a URL other than `ipv6.google.com` (i.e. a URL that isn't expected to be IPv6-only) with which you're experiencing a problem, or otherwise describe your *current* problem in more detail?

[sojourner](#) wrote on 2009-03-02: ([permalink](#))

 [w3m.tracelp](#) (1.4 MiB, text/html)

here is a trace to LP which shows a little delay but overall things seem to be loading faster today , maybe all of the other updates that have come with alpha 5 had some effect ?

[Colin Watson](#) wrote on 2009-03-02: ([permalink](#))

Thanks. What seems to be happening now is that your system does a perfectly ordinary IPv4 A record query for `bugs.launchpad.net` to `192.168.2.1` (the first nameserver in `/etc/resolv.conf`, which I'm guessing is your router), gets no reply for five seconds, and then retries successfully to `68.87.74.162`.

In other words, there is still a problem, but it no longer appears to be anything to do with IPv6. More to the point,

I don't think it's anything that glibc can do anything about either; it's just doing what it's told to do. At this point I feel comfortable in saying that you should fix this by way of a configuration change. For example, if 192.168.2.1 is simply a broken DNS server and you don't really want to use it, but your router always sends its address out via DHCP, you could edit /etc/dhcp3/dhclient.conf and use the directives described in the "OPTION MODIFIERS" section of the dhclient.conf(5) manual page to override it: for example 'supersede domain-name-servers 68.87.74.162 68.87.68.162 68.87.73.242;' if those addresses are stable.

Changed in glibc:
status: Triaged → Fix Released

[Jason Waters](#) wrote on 2009-03-02: ([permalink](#))

Colin,

I think you were right that I didn't have a bug at all. I actually posted to ubuntuforums and then figured it out on my own.

<http://ubuntuforums.org/showthread.php?p=6798431>

It was an ACPI problem. Thanks for your input though. Take care.

Jason

[sojourner](#) wrote on 2009-03-02: ([permalink](#))

Thankyou for all your efforts in resolving this . I do have a few observations though , this did not occur with any earlier version of Ubuntu or with jaunty prior to kernel 2.6.28-4 where ipv6 was built in to the kernel , as long as I disabled ipv6 in /etc/modprobe.d/aliases and blacklisted ipv6 . My solution will be to install the new router I bought as soon as this came up , I have just been keeping the old one online to test fixes .

[Jason Harman](#) wrote on 2009-03-03: ([permalink](#))

I've listed a new bug report for a bug I was experiencing that I thought *might* have been related to this one. If anyone is interested in commenting or relating please see here:

<https://bugs.launchpad.net/ubuntu/+bug/337488>

[Bug Watch Updater](#) on 2009-04-17

Changed in glibc (Fedora):
status: In Progress → Fix Released

[Gonzo](#) wrote on 2009-04-25: [\(permalink\)](#)

I'm sorry to say that all of you have gone too far in the debris of protocols putting aside the possible workaround...

I'm from Ukraine and I confirm that connection became slow with this new 9.04 release.

As 'sojourner' already told there was an excellent fix for that in the past (blacklisting the 'ipv6' in /etc/modprobe.d/aliases).

Can anyone tell the new fix for that? I mean can I still blacklist ipv6 because its unsupported in my country at all.

[geckon](#) wrote on 2009-04-25: [\(permalink\)](#)

What about compiling kernel with IPv6 support compiled as a module and then blacklist it? This should work, shouldn't it?

To post a comment you must [log in](#).

What next?

- » [Report another bug](#) about glibc in fedora
- » [List open bugs](#) for glibc in fedora

[What is Launchpad?](#) | [Get help with Launchpad](#)

© 2004-2009 Canonical Ltd. | [Terms of use](#)