

Bug 459756 - DNS resolver library doesn't seem to be working reliably[Last Comment](#) - [Format For Printing](#) - [XML](#) - [Clone This Bug](#)**Status:** CLOSED RAWHIDE**Reported:** 2008-08-21 18:51 EDT by Tom Horsley**Aliases:** None ([edit](#))**Modified:** 2009-04-16 15:24 EDT ([History](#))**Product:** Fedora**Fixed In Version:****Component:** glibc ([Show Fedora/glibc bugs](#))**Release Notes:****Version:** 10**Clone Of:****Platform:** All Linux**Priority:** medium **Severity:** high ([vote](#))**Target Milestone:** ---**Assigned To:** Jakub Jelinek**QA Contact:** Fedora Extras Quality Assurance**URL:****Whiteboard:** Workaround: [http://www.fedorafaq.org/...](http://www.fedorafaq.org/)**Keywords:** Reopened**Depends on:****Blocks:** [477145](#)Show dependency [tree](#)**Restrict Group Visibility:** Only users in any of the selected groups can view this bug:

(Unchecking all boxes makes this a more public bug.)

Current Groups: ([edit](#))

None Set

Attachments	
strace output from failed wget (20.91 KB, text/plain) 2008-08-28 17:54 EDT, Tom Horsley	Details
strace output from wget which worked (23.13 KB, text/plain) 2008-08-28 17:57 EDT, Tom Horsley	Details
Here's the tcpdump from when wget fails (944 bytes, text/plain) 2008-08-29 12:27 EDT, Tom Horsley	Details
strace of that failing wget (with -tt this time) (24.84 KB, text/plain) 2008-08-29 12:28 EDT, Tom Horsley	Details
tcpdump from when wget works (875 bytes, text/plain) 2008-08-29 12:29 EDT, Tom Horsley	Details
strace of the working wget (with -tt) (27.20 KB, text/plain) 2008-08-29 12:30 EDT, Tom Horsley	Details
tcpdump from nslookup (663 bytes, text/plain) 2008-08-29 12:31 EDT, Tom Horsley	Details
old-Debian-patch (1.78 KB, patch) 2008-09-02 10:27 EDT, Denys Vlasenko	Details Diff
Proof-of-concept patch (11.98 KB, patch) 2008-09-04 07:01 EDT, Denys Vlasenko	Details Diff
Test program (3.37 KB, text/plain) 2008-09-04 07:02 EDT, Denys Vlasenko	Details
Interface configuration on my F10 laptop (806 bytes, text/plain) 2008-11-29 03:44 EDT, David Miller	Details
tcpdump trace of DNS query on my laptop F10 system (289 bytes, text/plain) 2008-11-29 03:46 EDT, David Miller	Details
Interface list on FC9 desktop (655 bytes, text/plain) 2008-11-29 04:10 EDT, David Miller	Details

FC9 DNS query, does not timeout (560 bytes, text/plain) 2008-11-29 04:12 EDT, David Miller	Details
strace -o wget.log -tt wget (25.49 KB, application/octet-stream) 2008-11-29 18:38 EDT, Richard Zuber	Details
Wireshark results of just DNS traffic during a 'wget -O /dev/null google.com' (12.34 KB, text/plain) 2008-12-30 00:12 EDT, Kai Meyer	Details
Add an attachment (proposed patch, testcase, etc.)	Hide Obsolete (1)

External Bugs

Description From Tom Horsley 2008-08-21 18:51:37 EDT

Description of problem:

[Bug 459672](#) was the first appearance of this problem, cropping up in yum frequently, but it doesn't appear to be isolated to yum, I've now seen the same thing in wget:

```
[root@zooty ~]# wget
http://mirrors.fedoraproject.org/mirrorlist?repo=rawhide&arch=x86_64
--2008-08-21 18:40:03--
http://mirrors.fedoraproject.org/mirrorlist?repo=rawhide&arch=x86_64
Resolving mirrors.fedoraproject.org... failed: Name or service not known.
wget: unable to resolve host address `mirrors.fedoraproject.org'
```

There is a very long pause at the ... before it reports the error. If I repeat the process it sometimes works (but the long pause is still there).

On the other hand, if I run nslookup (which I believe doesn't use the glibc resolver code, but has its own code), this always works perfectly and is near instantaneous:

```
Server: 68.87.74.162
Address: 68.87.74.162#53
```

```
Non-authoritative answer:
mirrors.fedoraproject.org canonical name = wildcard.fedoraproject.org.
Name: wildcard.fedoraproject.org
Address: 209.132.176.120
```

Version-Release number of selected component (if applicable):
glibc-2.8.90-11.x86_64

How reproducible:

random, but frequent

Steps to Reproduce:

1. so far, I've seen this in both yum and wget
- 2.
- 3.

Actual results:

name lookup error

Expected results:

just works

Additional info:

The exact same hardware with the exact same resolv.conf file running both fedora 8 and fedora 9 does not display this problem, but I'm still just guessing that the problem is in glibc, could be network I suppose (but the total reliability of nslookup makes me think it isn't the network).

Comment #1 From Denys Vlasenko 2008-08-28 05:07:59 EDT

Can you do

```
strace -o wget.log -tt wget ...
```

and attach strace output to this bug?

Comment #2 From Tom Horsley 2008-08-28 08:29:22 EDT

I'll get the strace when I get home today, but one possibly relevant item may be IPv6 - I unchecked the "enable IPv6" box when I installed, mainly because I have no idea what the devil to type in the fields I need to fill in for IPv6 addresses, and my router doesn't support v6 anyway.

Comment #3 From Tom Horsley 2008-08-28 17:54:21 EDT

[Created an attachment \(id=315304\) \[details\]](#)

strace output from failed wget

I tried this a couple of times as normal user, and it worked, when I switched to root I got this failure. I don't know if that is random behaviour or actually significant (but I'll always be root when running yum, so it needs to work for root :-).

I just tried it 4 times in a row as root, and it failed the 1st 3 and worked on the 4th try. I'll attach the strace that worked after this attachment (might be interesting to compare).

Comment #4 From Tom Horsley 2008-08-28 17:57:04 EDT

[Created an attachment \(id=315305\) \[details\]](#)

strace output from wget which worked

Here's the one that worked.

I should also mention that I have selinux disabled, so I wouldn't think root should wind up being special.

Comment #5 From Denys Vlasenko 2008-08-29 06:18:26 EDT

You forgot to add -tt option to strace.

Failing wget took 5.05 seconds to execute. The failure happened after a DNS reply packet has arrived. I think your DNS server has 5 second timeout, if it can't resolve a host in 5 seconds, it replies to the client with "Name or service not known".

Successful wget took a little bit over 1 second to execute.

Comment #6 From Tom Horsley 2008-08-29 08:02:11 EDT

Sorry about the -tt, my pore old eyes just missed that :-). I'll see if I can rerun the strace soon.

The 5 second timeout may be believable, but I never get these timeouts on F8 or F9 booted on the same hardware, or even with nslookup on F10. It is almost like something really nasty is going on like a compiler error or uninitialized variable that results in it asking the wrong question, which triggers the timeout. (Yes, F8 and F9 have the same DNS servers configured).

Comment #7 From Denys Vlasenko 2008-08-29 08:42:36 EDT

Can you simultaneously run "tcpdump -nliethN -s0 udp port 53 -vvv" and capture DNS traffic? Just to test the theory that DNS requests are garbled.

Comment #8 From Tom Horsley 2008-08-29 12:27:48 EDT

[Created an attachment \(id=315374\) \[details\]](#)

Here's the tcpdump from when wget fails

OK, got a chance to boot f10, so here's a slew of new dumps of various things coming. The bad udp checksum comments bother me, but I see them on fedora 8 as well, so I guess I just don't understand what is being dumped.

Comment #9 From Tom Horsley 2008-08-29 12:28:51 EDT

[Created an attachment \(id=315375\) \[details\]](#)

strace of that failing wget (with -tt this time)

Comment #10 From Tom Horsley 2008-08-29 12:29:28 EDT

[Created an attachment \(id=315376\) \[details\]](#)

tcpdump from when wget works

Comment #11 From Tom Horsley 2008-08-29 12:30:35 EDT

[Created an attachment \(id=315377\) \[details\]](#)

strace of the working wget (with -tt)

Comment #12 From Tom Horsley 2008-08-29 12:31:56 EDT

[Created an attachment \(id=315378\) \[details\]](#)

tcpdump from nslookup

Just as a comparison, here's a tcpdump from doing nslookup on mirrors.fedoraproject.org (which always seems to work with no problems or time delays).

Comment #13 From Denys Vlasenko 2008-09-01 10:58:21 EDT

Look at tcpdumps.

"Failing wget" tcpdump shows that your machine asked for both IPv4 and IPv6 IP address of mirrors.fedoraproject.org, and DNS server replied with IPv6 address (only).

At the first glance, this reply says "I don't know the [IPv6] address, but here is some information which may be useful":

```
AAAA? mirrors.fedoraproject.org. 1/1/0 mirrors.fedoraproject.org. CNAME
wildcard.fedoraproject.org. ns: fedoraproject.org. SOA fedoraproject.org.
hostmaster.fedoraproject.org. 2008082802 28800 7200 2419200 86400 (113)
```

Apparently this info isn't useful, so wget wats for IPv4 answer for 5 seconds, but it does not come.

"working wget" - your machine asked for both IPv4 and IPv6 IP address of mirrors.fedoraproject.org again, but this time DNS server replied with *IPv4* address (only), and reply does contain the IPv4 address. So wget can proceed.

nslookup asks only "IPv4 question", so it always succeeds.

Unless you are really connected to IPv6 backbone, you probably need to reconfigure your DNS server to not give back only IPv6 address, but try to find IPv4 address too.

Which DNS server is it (vendor, version, etc)?

Alternatively, you may disable IPv6 on your machine (unload ipv6 module, etc...)

Comment #14 From Tom Horsley 2008-09-01 11:16:38 EDT

OK, that probably explains everything. The real problem is that I did disable IPv6 when I installed, but that apparently didn't "take", so maybe this is really an anaconda/network config problem.

I didn't think I was supposed to need to go as far as disabling the IPv6 module merely to make it not use IPv6, but I guess I could try that and see if things start working.

The DNS servers aren't under my control, they are merely the ones comcast tells my router to use when the router gets the DHCP lease.

Thanks for explaining those dumps, they are mostly just gibberish to me.

I wonder what component I should redirect this bug to next :-).

Comment #15 From Tom Horsley 2008-09-01 11:32:45 EDT

Maybe it is really the library. I added the /etc/modprobe.d/noipv6 file containing the line "install ipv6 /bin/true" and booted into F10.

I can lsmod and see that no ipv6 module is loaded, but the exact same random failures still happen in yum and wget. If I run system-config-network I see the IPv6 checkbox is NOT checked for eth0. As near as I can tell, I have IPv6 as disabled as I can possibly get it, so if the resolver is still making Ipv6 DNS requests, it must be getting carried away.

Comment #16 From Denys Vlasenko 2008-09-02 04:31:28 EDT

There seems to be no way to instruct glibc to not perform AAAA resolution by editing of resolv.conf etc :(

The thing which seems to work for most people is to replace "alias net-pf-10 ipv6" by "alias net-pf-10 off" in /etc/modprobe.d/aliases. Can you try this and let me know whether this helped? (You will need to reboot)

Comment #17 From Denys Vlasenko 2008-09-02 10:27:01 EDT

It seems that Debian people were trying to improve this situation:

<http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=435646>

The idea is: "if there is no non-link-local IPv6 addresses, we are probably not connected to 'big' IPv6 network and resolving hostnames into IPv6 addresses is pointless"

This caused a regression:

<http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=441857>

Basically, getaddrinfo("ip6-localhost", "119", &hints, &result) fails under above conditions. I assume that ip6-localhost was set to ::1 in /etc/hosts.

IOW: the above heuristic should be applied only when we try to do real DNS resolution; resolving a name from /etc/hosts to IPv6 address should be ok.

Unfortunately, it seems that Debian patch cannot be easily adapted to do it. I will attach it now for reference anyway.

Comment #18 From Denys Vlasenko 2008-09-02 10:27:45 EDT

~~Created an attachment (id=315555) [details]~~
old Debian patch

Comment #19 From Denys Vlasenko 2008-09-02 10:50:16 EDT

Here is how it can be done.

getaddrinfo() already has code which finds out whether IPv6 addrs exist (bool seen_ipv6 variable). This is how it percolates down to actual DNS resolution

```
getaddrinfo ->  
gai_inet -> gethosts (this is a macro, it does dynamic NSS call) ->
```

```
_nss_dns_gethostbyname3_r ->
__libc_res_nsearch ->
__libc_res_nquery[domain] ->
res_nmkquery...
```

In one of these functions, we need to check `seen_ipv6`, or a new analogous variable `ok_to_emit_AAAA_requests` (or whatever), and act accordingly.

Currently `seen_ipv6` is not passed down. Can we use a bit in `_res.options` for this?

Comment #20 From Denys Vlasenko 2008-09-02 10:57:01 EDT

Oh, and btw, `sysdeps/unix/sysv/linux/check_pf.c` seems to have a weaker form of above mentioned bug:

```
if (ifam->ifa_family == AF_INET)
{
    if (*(const in_addr_t *) address
        != htonl (INADDR_LOOPBACK))
        *seen_ipv4 = true;
}
else
{
    if (!IN6_IS_ADDR_LOOPBACK (address))
        *seen_ipv6 = true;
}
```

So, if system has only loopback addresses, `seen_ipvN` would not be set, and this will suppress resolution of names even from `/etc/hosts`. (Did not test whether this is really happening...)

Comment #21 From Tom Horsley 2008-09-02 14:54:31 EDT

Potentially silly question here:

If the DNS library wants to return both IPv4 and IPv6 addresses if they are available.

And if (as seems to be the case with comcast's DNS servers) the server feels like it has satisfied the request by simply returning one of an IPv4 or IPv6 result at random.

Then shouldn't the library be written to do two separate requests, one IPv4 only and the other IPv6 only?

Seems like that algorithm would work correctly no matter what, then it could be modified to avoid wasting time on the IPv6 request if IPv6 isn't configured on the system.

Comment #22 From Denys Vlasenko 2008-09-03 04:32:23 EDT

> Then shouldn't the library be written to do two separate requests,
> one IPv4 only and the other IPv6 only?

It does this already:

```
tcpdump: listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
```

```
12:07:35.453095 IP (tos 0x0, ttl 64, id 61002, offset 0, flags [DF], proto UDP (17), length 71) 192.168.1.106.38681 > 68.87.74.162.domain: [bad udp cksum 2337!] 15402+ A? mirrors.fedoraproject.org. (43)
```

```
12:07:35.454233 IP (tos 0x0, ttl 64, id 61003, offset 0, flags [DF], proto UDP (17), length 71) 192.168.1.106.38681 > 68.87.74.162.domain: [bad udp cksum a0af!] 43180+ AAAA? mirrors.fedoraproject.org. (43)
```

These two packets are two separate requests.

The thing we are trying to solve here is: DNS servers have bugs in IPv6 handling, therefore we should avoid using AAAA queries when we know we can't use the result anyway (because IPv6 routing is not set up). Sending AAAA

queries only increases network traffic and triggers bugs in DNS servers in this case.

You are not the first person to report DNS+IPv6 problem. I googled for it - it's quite common. This proves that this is a real world problem and we'd better fix it.

Comment #23 From Denys Vlasenko 2008-09-03 05:43:23 EDT

What we have now:

If hints.ai_flags includes the AI_ADDRCONFIG flag, then IPv4 addresses are returned in the list pointed to by result only if the local system has at least one IPv4 address configured, and IPv6 addresses are only returned if the local system has at least one IPv6 address configured.

Implementation interprets this as "if the local system has at least one NON-LOOPBACK address configured" (this "non-loopback" check happens inside __check_pf):

```
__check_pf (&seen_ipv4, &seen_ipv6, &in6ai, &in6ailen);

if (hints->ai_flags & AI_ADDRCONFIG)
{
    /* Now make a decision on what we return, if anything. */
    if (hints->ai_family == PF_UNSPEC && (seen_ipv4 || seen_ipv6))
    {
        /* If we haven't seen both IPv4 and IPv6 interfaces we can
           narrow down the search. */
        if (!seen_ipv4 || !seen_ipv6)
        {
            local_hints = *hints;
            local_hints.ai_family = seen_ipv4 ? PF_INET : PF_INET6;
            hints = &local_hints;
        }
    }
    else if ((hints->ai_family == PF_INET && !seen_ipv4)
            || (hints->ai_family == PF_INET6 && !seen_ipv6))
    {
        /* We cannot possibly return a valid answer. */
        free (in6ai);
        return EAI_NONAME;
    }
}
}
```

Is it a bug that if hints->ai_family == PF_UNSPEC, addresses are still returned even if seen_ipv4 == seen_ipv6 == false?

I verified it. With network disabled:

```
# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet6 ::1/128 scope host
...

```

```
and hints.ai_flags = AI_ADDRCONFIG; hints.ai_family = AF_INET6; -
```

```
# LD_LIBRARY_PATH=. /root/srcdevel/glibc/fix_rhbug457506/a.out
getaddrinfo <=> debug output from glibc
seen_ipv4 = seen_ipv6 = false <=> debug output from glibc
E: Failed to get addrinfo: Name or service not known
```

```
with hints.ai_flags = AI_ADDRCONFIG; hints.ai_family = AF_UNSPEC; -
```

```
# LD_LIBRARY_PATH=. /root/srcdevel/glibc/fix_rhbug457506/a.out
getaddrinfo
seen_ipv4 = seen_ipv6 = false
getaddrinfo 2
getaddrinfo 3
getaddrinfo 4
getaddrinfo 5
```

```
Addrinfo for 0x2270370
```

```

Flags:          32
Family:        2
Socket Type:   1
Protocol:      6 (tcp)
Canonical name: (null)
Socket Address (len=16):
  Port:        119
  IPv4 Address: 127.0.0.1
...

```

Questions:

- (1) do we need to fix AF_UNSPEC to also fail here?
- (2) is "local system has at least one NON-LOOPBACK address configured" interpretation of AI_ADDRCONFIG flag correct?
- (3) if yes, should it also include "...and if IPv6 addresses are not link-local?"
- (4) I think it's impractical to expect that people will use AI_ADDRCONFIG as often as needed, I think we still need to avoid A/AAAA DNS queries if IPv4/IPv6 routing is not configured. What others (esp. Ulrich as maintainer) think?

Comment #24 From Ulrich Drepper 2008-09-03 13:25:31 EDT

First: fix wget. getaddrinfo should always be called with AI_ADDRCONFIG. I don't know why not all programs are already fixed.

Second: Debian's patch is of course completely wrong. The standards demand the current behavior and programs can depend on it and break.

Third: Comcast is known to have broken DNS servers. Some will not reply at all to IPv6 replies. This is something you must bring up with Comcast. There is nothing the resolver can do.

Comment #25 From Denys Vlasenko 2008-09-04 04:18:33 EDT

> First: fix wget

This won't work. If only loopback is configured:

```

# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST> mtu 1500 qdisc pfifo_fast state DOWN qlen 1000
    link/ether 00:1e:37:d0:50:06 brd ff:ff:ff:ff:ff:ff

```

then this fails:

```

void gaifail(const char *msg, int code)
{
    fprintf(stderr, "E: %s: %s\n", msg, gai_strerror(code));
    exit(EXIT_FAILURE);
}

int main(int argc, char **argv)
{
    int status;
    struct addrinfo hints;
    hints.ai_flags = AI_ADDRCONFIG;
    hints.ai_family = AF_INET;
    hints.ai_socktype = 0;
    hints.ai_protocol = 0;
    struct addrinfo *result = NULL;

    if((status = getaddrinfo("localhost", "119", &hints, &result)) != 0)
        gaifail("Failed to get addrinfo", status);
    ...
}

# gcc addrtest.c
# ./a.out

```


E: Failed to get addrinfo: Name or service not known

This is clearly wrong.

BTW, with AF_UNSPEC instead of AF_INET it works, which I noted in question #1 - we seem to have a discrepancy here, in which directions do we need to fix it?

Comment #26 From Denys Vlasenko 2008-09-04 06:51:30 EDT

I made a patch, very rough. This is what happens on the wire when test program is resolving "google.com":

```
sh-3.2# tcpdump -nlieth0 -s0 udp port 53
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
12:46:03.088423 IP 10.34.33.233.38677 > 10.34.32.125.domain: 38234+ A?
google.com. (28)
12:46:03.088622 IP 10.34.32.125.domain > 10.34.33.233.38677: 38234 3/4/4 A
72.14.207.99, A 64.233.167.99, A 64.233.187.99 (212)
```

Patch contains instrumentation which shows how it decides that there is no routable IPv6 address on our interfaces, and therefore sending AAAA requests is not done:

```
getaddrinfo
seen_ipv4 = seen_ipv6 = 0
seen_ipv4 = 1
seen_ipv4 = 3
seen_ipv6 = 1
seen_ipv6 = 2
seen_ipv6:2
__libc_res_nquery: __vda_seen_ipv6:2
__vda_seen_ipv6 < SEEN_IPVx_ROUTABLE in __libc_res_nquery
__libc_res_nquery: __vda_seen_ipv6:2
__vda_seen_ipv6 < SEEN_IPVx_ROUTABLE in __libc_res_nquery
__libc_res_nquery: __vda_seen_ipv6:2
__vda_seen_ipv6 < SEEN_IPVx_ROUTABLE in __libc_res_nquery
__libc_res_nquery: __vda_seen_ipv6:2
__vda_seen_ipv6 < SEEN_IPVx_ROUTABLE in __libc_res_nquery
__libc_res_nquery: __vda_seen_ipv6:2
```

But addresses from /etc/hosts (localhost, localhost6 etc) would be resolved just fine, patch only suppresses AAAA requests on the wire, not IPv6 resolution in general.

For comparison, tcpdump with unpatched glibc:

```
sh-3.2# tcpdump -nlieth0 -s0 udp port 53
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
12:46:17.545478 IP 10.34.33.233.35403 > 10.34.32.125.domain: 40766+ AAAA?
google.com. (28)
12:46:17.545699 IP 10.34.32.125.domain > 10.34.33.233.35403: 40766 0/1/0 (78)
12:46:17.546089 IP 10.34.33.233.39325 > 10.34.32.125.domain: 27454+ AAAA?
google.com.englab.brq.redhat.com. (50)
12:46:17.546341 IP 10.34.32.125.domain > 10.34.33.233.39325: 27454 NXDomain*
0/1/0 (112)
12:46:17.546582 IP 10.34.33.233.37369 > 10.34.32.125.domain: 18916+ AAAA?
google.com.brq.redhat.com. (43)
12:46:17.546812 IP 10.34.32.125.domain > 10.34.33.233.37369: 18916 NXDomain
0/1/0 (98)
12:46:17.547061 IP 10.34.33.233.34683 > 10.34.32.125.domain: 27036+ AAAA?
google.com.redhat.com. (39)
12:46:17.547280 IP 10.34.32.125.domain > 10.34.33.233.34683: 27036 NXDomain
0/1/0 (87)
12:46:17.547511 IP 10.34.33.233.41772 > 10.34.32.125.domain: 3927+ A?
google.com. (28)
12:46:17.547786 IP 10.34.32.125.domain > 10.34.33.233.41772: 3927 3/4/4 A
64.233.187.99, A 72.14.207.99, A 64.233.167.99 (212)
```

Will attach patch and test program now

Comment #27 From Denys Vlasenko 2008-09-04 07:01:34 EDT

[Created an attachment \(id=315735\) \[details\]](#)

Proof-of-concept patch

This patch only demonstrates the basic idea, do not use in production

Comment #28 From Denys Vlasenko 2008-09-04 07:02:10 EDT

[Created an attachment \(id=315736\) \[details\]](#)

Test program

Comment #29 From Denys Vlasenko 2008-09-04 07:06:05 EDT

Build patched glibc 2.8, build test program

```
# gcc -o addrtest addrtest.c
```

Then run it

```
# LD_LIBRARY_PATH=./nss:./resolv ./addrtest
```

and observe tcpdump and program output. Edit this part:

```
// if((status = getaddrinfo("ip6-localhost", "119", &hints, &result)) != 0)
// if((status = getaddrinfo("localhost", "119", &hints, &result)) != 0)
// if((status = getaddrinfo("localhost6", "119", &hints, &result)) != 0)
  if((status = getaddrinfo("google.com", "119", &hints, &result)) != 0)
    gaifail("Failed to get addrinfo", status);
```

to test resolution of localhost[6] names. Play with networking on/off (in Gnome, the easiest way is to right-click on networking icon and switch off "[x] Enable networking")

Comment #30 From Tom Horsley 2008-09-05 20:21:06 EDT

Just to note a potential solution in case anyone else runs into this:

I worked around this by configuring bind to run as a caching nameserver and editing the /etc/sysconfig/named file to add the -4 startup option.

Now I can point resolv.conf at localhost, and nobody does IP46 lookups on the comcast servers anymore :-).

Comment #31 From Mads Kiilerich 2008-11-13 14:24:00 EDT

I have created [Bug 471450](#) which might be related

Comment #32 From Tom Horsley 2008-11-13 16:39:10 EDT

And I'll add that as of the latest Fedora 10 Preview release, the resolver lib still had this problem (and I still work around it by running a caching nameserver that only does IPv4 lookups).

Comment #33 From Bug Zapper 2008-11-25 21:50:28 EDT

This bug appears to have been reported against 'rawhide' during the Fedora 10 development cycle.
Changing version to '10'.

More information and reason for this action is here:
<http://fedoraproject.org/wiki/BugZappers/HouseKeeping>

Comment #34 From Jonathan Steffan 2008-11-26 13:13:08 EDT

This is still an issue with Fedora 10 x86_64 Final.

Comment #35 From Tom Horsley 2008-11-26 14:00:03 EDT

Yep. Definitely still fails in released fedora 10, and I also reported [bug-473073](#) which may be related to the reorganization of the host lookup code (this new bug is caused by NIS).

Comment #36 From Ulrich Drepper 2008-11-26 20:00:13 EDT

In all this endless discussion I've seen one valid point: if no interface is configured at all and AI_ADDRCONFIG is used, then we shouldn't try to perform any lookup. I've changed that.

If the only configured IPv4 address is the localhost address the system is regarded as not having any IPv4 addresses. Similarly for IPv6. Therefore some of the comments above are wrong.

Anyway, none of this is likely to be the reason for the problem/delay. The original report was about a system with network interface. And I see nothing wrong with that (except buggy programs not using AI_ADDRCONFIG). If the name server is broken and doesn't reply then the delay is what you get. Complain to the ISV.

If there any other `_real_` issues open a new bug. This one is already too overloaded with comments and references which might or might not have anything to do with each other. I'll close the bug when we have the first rawhide build for F11 which will have the change I checked in.

Comment #37 From Tyler K 2008-11-26 21:17:15 EDT

I was having this same issue. I all the items listed above for disabling ipv6, with no luck. Using wireshark filtering for DNS, I was getting 'UDP checksum errors'. DNS queries seemed to work fine when using a local DNS resolver, and only 10-20% of the time while using external DNS servers.

I disabled UDP checksum offloading, and now my DNS works fine, in all cases.

```
'ethtool -K eth0 rx on tx off'
```

Comment #38 From Tom Horsley 2008-11-26 21:37:19 EDT

But as a practical matter, ISPs aren't going to give a ****, and if the resolver code also doesn't give a ****, then the end result is great breakage everywhere.

Seems like it ought to be possible do something like spiff up the nsswitch syntax to support "dnsv4" and "dnsv6" in addition to the simple "dns" keyword so there would be a relatively simple fix available by just changing all instances of "dns" in /etc/nsswitch.conf to "dnsv4".

Comment #39 From David Miller 2008-11-27 05:06:28 EDT

FWIW I'm hitting this too, and it took me a while to find this bugzilla. But be prepared for a deluge of reports on this as F10 starts to get deployed out there.

Nothing is going to make Comcast fix this any time soon.

The net effect of not adding a workaround is that all users using Comcast as their ISP are going think Linux is slow and sucks.

Comment #40 From Ulrich Drepper 2008-11-27 13:38:16 EDT

I think people are majorly confused.

Nothing whatsoever changed wrt to deciding whether to make V6 lookups or not. The same rules are in place for many years. If you think F1 to F9 are OK, so is F10.

To the contrary, F10 will be slightly faster because we perform v4 and v6 lookups in parallel now instead of sequentially. There were a number of problems with that change and there is perhaps one more left. But that's it.

Nothing else changes.

The other possible changes are in the system configuration (perhaps more systems have IPv6 loaded?) and in applications which got converted to use `getaddrinfo` without using `AI_ADDRCONFIG` as they should.

The whole details of the `getaddrinfo` implementation allow to make the decision about using v6 in exactly one place (the setup of network interfaces) instead of duplicating it in many places. Configuration options are the worst possible way. Just get the setup and the applications fixed and all is fine.

Comment #41 From Tom Horsley 2008-11-27 16:56:14 EDT

>Nothing whatsoever changed wrt to deciding whether to make V6 lookups or not.
>The same rules are in place for many years.

Something changed somewhere. No power on earth can reproduce this problem in fedora 8 or 9 with same hardware and same comcast DNS servers. With fedora 10 (and test versions leading up to 10), the problem always exists - at least 20% of the time the I get the unable to lookup name problem. I usually need to try "yum install system-config-bind" about 5 times before it finally gets it downloaded so I can setup a local named. This is not "OK" :-).

Maybe you are doing the lookups in parallel, then rejecting them both on the first error, and I get lots of errors from the v6 lookups? Something is clearly wrong, and I'm not confused about that.

Comment #42 From David Miller 2008-11-29 03:32:11 EDT

What happens is that the IPV4 part of the response comes back from the DNS server, and then the resolver sits there and waits for the IPV6 part to come but that never happens and we timeout instead.

Anyways, I wonder what you mean by "ipv6 address configured" because just bringing an interface up with an ipv4 address gives it an ipv6 link local address automatically. So I hope your test is a little bit more sophisticated than it sounds.

Every single interface gets a link-local IPV6 address merely as a side effect of being brought up in any way.

I just checked the current glibc code after your changes and it's not going to fix this situation at all.

The `__check_pf()` code marks "seen_ipv6" as true if any non-loopback address is seen. This means the automatic link-local address will cause `seen_ipv6` to be set to true. So ipv6 DNS queries will be done on pretty much every system out there regardless of whether real global scope IPV6 addresses are configured on the interface.

Comment #43 From David Miller 2008-11-29 03:44:29 EDT

[Created an attachment \(id=325068\) \[details\]](#)

Interface configuration on my F10 laptop

Notice the automatic IPV6 link-local address assigned to my wireless interface, eth1

Comment #44 From David Miller 2008-11-29 03:46:03 EDT

[Created an attachment \(id=325069\) \[details\]](#)

tcpdump trace of DNS query on my laptop F10 system

Note both AAAA and A record request sent.

A record response arrives, at this point the resolver hangs waiting for the AAAA response that never arrives.

Comment #45 From David Miller 2008-11-29 03:54:33 EDT

As a comparison, Windows Vista's algorithm is that if only link-local or Teredo IPv6 addresses are assigned to the interface, AAAA lookups will not be performed.

Doing some more research online suggests that it is extremely common for AAAA DNS requests to be silently dropped by firewalls and other intermediate devices.

Therefore the conservative choice to only ask for AAAA records when we have something more than a link-local IPv6 address assigned to some interface makes a lot of sense and will fix this Comcast issue completely.

Comment #46 From Jakub Jelinek 2008-11-29 03:54:50 EDT

AFAIK F9 and earlier glibc was doing both AAAA and A requests as well in such case, only it wasn't sending both AAAA and A requests together, but instead AAAA request first and when it arrived (or timed out) the A request. So if your DNS never responds to AAAA queries, F9 and earlier should time out the same way...

Comment #47 From David Miller 2008-11-29 03:59:01 EDT

FC9 and before never had the DNS timeout behavior on any of my systems, with all updates applied.

Comment #48 From David Miller 2008-11-29 04:10:58 EDT

[Created an attachment \(id=325070\) \[details\]](#)
Interface list on FC9 desktop

List of interfaces on my FC9 desktop, to be used in analyzing the tcpdump trace I'm about to post.

Comment #49 From David Miller 2008-11-29 04:12:21 EDT

[Created an attachment \(id=325071\) \[details\]](#)
FC9 DNS query, does not timeout

This is a DNS query from an FC9 system with all updates applied. This is behind the same ISP, Comcast, as my laptops from which the FC10 traces were recorded.

Comment #50 From Richard Zuber 2008-11-29 18:38:05 EDT

[Created an attachment \(id=325104\) \[details\]](#)
strace -o wget.log -tt wget

```
[root@risiko-laptop ~]# strace -o wget.log -tt wget
http://mirrors.fedoraproject.org/mirrorlist?repo=fedora-10&arch=i386
[1] 4724
[root@risiko-laptop ~]# --2008-11-29 23:29:28--
http://mirrors.fedoraproject.org/mirrorlist?repo=fedora-10
Resolving mirrors.fedoraproject.org... failed: Name or service not known.
wget: unable to resolve host address `mirrors.fedoraproject.org'
```

```
[1]+ Done strace -o wget.log -tt wget
http://mirrors.fedoraproject.org/mirrorlist?repo=fedora-10
[root@risiko-laptop ~]# less wget.log
```

Comment #51 From Andre Robatino 2008-11-29 19:49:32 EDT

My father also seems to have this issue. For him, the "host" command seems to work reliably, but nothing else (we haven't tried doing lookups dozens of times to see if it would finally work). Does the host command also use its own resolver code (like nslookup)?

Comment #52 From Ulrich Drepper 2008-11-29 20:39:27 EDT

I never said the current code is correct. Read [comment #40](#). There is one known bug left (see [bug 471450](#)). This bug has the consequence that sometimes replies get dropped and therefore it appears as if there is a timeout.

Comment #53 From David Miller 2008-11-30 00:26:44 EDT

Actually your comments in #40 have a lot of falsehoods in it.

Something did change, and it was not in people's configuration and it was not applications being converted to be ipv6 aware.

The applications effected have had ipv6 support for years.

Interfaces have been obtaining a link-local ipv6 address merely as a result of being brought up (even with just an ipv4 address), for years. In fact the very first IPV6 stack in Linux did this.

It was, in fact, glibc's behavioral change that broke things, nothing else.

Now that we have that established, could you please at least entertain the idea of adding a link-local address check to `__check_pf()`, as I suggested?

That would solve all of these problems permanently, without having to be concerned with what AAAA handling peculiarities might exist in some large ISPs DNS implementation.

To Andre Robatino, wrt. [comment #51](#), the host command uses it's own resolver code and doesn't use glibc's stuff. That's why it works without timeouts, and application DNS lookups have the problem.

Comment #54 From Andre Robatino 2008-11-30 00:44:45 EDT

I should also have mentioned that my father is using i386. In fact, we are using identical PCs, and each using DSL with dynamic IP, except that I'm using x86_64 F10 without seeing this problem at all, and he's using i386 F10. So if it is the same bug, it's definitely not limited to 64-bit.

Comment #55 From Andre Robatino 2008-11-30 13:14:50 EDT

We were able to work around the problem by entering the secondary DNS address 208.67.222.222 under the DNS tab for eth0 in system-config-network. The primary DNS is the same as that for the normally working F9 machine, namely the LAN address for the DSL router. I got the idea from the fact that one of the few differences between my father's setup and mine is that we are using different ISPs with different DNS servers.

Comment #56 From David Miller 2008-12-01 01:35:46 EDT

Looking at these traces with Herbert Xu, we have a theory that the problem is probably exactly sending the A and AAAA request out at the same time.

We believe that if different ports were used, or the requests were sent in sequence (only sending the next after the first has been replied to, as FC9 does), the AAAA response would be sent back by the DNS server.

We think this behavior is meant as a countermeasure to the DNS server DoS vulnerabilities from earlier this year.

Comment #57 From Richard Zuber 2008-12-01 03:49:26 EDT

I just arrived at work with my F10 laptop and everything with yum works fine here. looks like the connection at home has some name resolution issue... but its outside my house

Comment #58 From Jonathan Steffan 2008-12-01 20:42:28 EDT

(In reply to [comment #56](#))

Thanks for the update. Please let us know when we have something you would like us all to test (koji build, something.) This issue has become a daily issue for myself and colleagues and I can't imagine is pleasant for anyone else.

Comment #59 From Max Kanat-Alexander 2008-12-01 21:20:36 EDT

*** [Bug 473863](#) has been marked as a duplicate of this bug. ***

Comment #60 From Max Kanat-Alexander 2008-12-01 21:45:03 EDT

I have entirely disabled ipv6 on my system--"lsmod | grep ipv6" shows no results--and applications are still making AAAA requests, according to Wireshark.

Changing the checksum offloading does not help.

Comment #61 From Max Kanat-Alexander 2008-12-01 21:48:14 EDT

Another interesting fact is that "host" (which I assume uses bind-libs instead of glibc) also always makes AAAA requests (even though I have ipv6 turned off), and yet they never fail or cause problems, and I always get a response from Comcast with no delay.

Comment #62 From Max Kanat-Alexander 2008-12-01 21:50:08 EDT

(In reply to [comment #56](#))

I'm sorry for the comment spam, but I see and suspect the same thing as David. Serial requests work, parallel requests often fail.

Comment #63 From Ulrich Drepper 2008-12-01 22:13:31 EDT

(In reply to [comment #60](#))

> I have entirely disabled ipv6 on my system--"lsmod | grep ipv6" shows no
> results--and applications are still making AAAA requests, according to
> Wireshark.

Read the thread. I've already said multiple times that buggy programs which don't use AI_ADDRCONFIG will perform AAAA lookups. File bugs for the programs which cause the lookups.

Comment #64 From Max Kanat-Alexander 2008-12-02 02:09:13 EDT

If every single program should call getaddrinfo with AI_ADDRCONFIG, then why isn't it the default behavior when not specified? (I'm sure there's some good explanation, but I'm sure that most of us on this bug don't know it.) Is it just that they're already specifying some other flags explicitly?

Note that it appears that wget actually specifically *removed* AI_ADDRCONFIG support:

<http://osdir.com/ml/web.wget.patches/2005-06/msg00030.html>

As far as I can tell, every application I'm using is affected. (Including Firefox, xchat-gnome, pidgin, ssh, gwibber, and claws-mail, at the least.)

In any case, I've filed a Core: Networking bug at bugzilla.mozilla.org, for Firefox.

Comment #65 From Denys Vlasenko 2008-12-02 08:00:31 EDT

With AI_ADDRCONFIG, on a machine with no network connectivity whatsoever, but with loopback IP address configured on interface "lo", and "localhost" entered in /etc/hosts,

wget <http://localhost/a/page.html>

would fail. I would find that very wrong.

Comment #66 From Andre Costa 2008-12-02 19:27:11 EDT

Well, I'm both relieved and disappointed. Relieved because now I know I'm not alone. Disappointed because, needless to say, I am also suffering with terrible network performance due to slow and/or failed DNS resolution on my brand new F10 installation. This, along with problems with NetworkManager and static IP, and s-c-n messing with configuration parameters, gives F10 the crown of "worst-network-setup-ever" on Linux as far as I can tell.

For the first time I can remember, Windows XP is running more efficiently on my machine than Linux. That sucks =/

... sorry for the rant, it's just really frustrating. I know there's (competent) people working to fix these issues, I just hope it happens soon.

Comment #67 From Need Real Name 2008-12-02 22:58:45 EDT

For those of you who end up here, searching for your problem as I did, here is what solves it.

With the help of the people in #fedora on freenode, we figured it out.

It is obvious that it's the IPv4 / IPv6 issue, that has already been confirmed. Here is what we did to solve it for now until there is a fix.

As root:

```
Disable NetworkManager ( service NetworkManager stop; chkconfig NetworkManager off )
```

```
Enable network ( service network start; chkconfig network on )
```

```
Have the DVD media available, and install bind, if not already installed ( yum --localinstall --disable-repo=* /media/FC10.../Packages/bind...rpm ) The disable-repo=* kept it from just hitting our bug again when it tried to pull down mirrors :)
```

```
Enable named ( service named start; chkconfig named on )
```

```
Modify /etc/resolv.conf to have  
nameserver 127.0.0.1
```

```
Also modify /etc/sysconfig/network-scripts/ifcfg-eth# to set ONBOOT=yes and PEERDNS=no
```

If I am remember everything I did, that did it!!

The main thing was that we just needed to disable NetworkManager and install bind so that the machine looked at itself for DNS and life was groovy...

I am going to hold off updating the rest of my machines to 10 until this is resolved, so I don't have to do it again, but hopefully this helps someone out!

Adam

Comment #68 From Jonathan Steffan 2008-12-03 12:53:25 EDT

There is no reason to disable NetworkManager for this workaround. I am against having this information on the bug report but the change that is needed to not have your /etc/resolv.conf clobbered is the PEERDNS=no. You can still use NetworkManager with a local DNS server (named, dnsmasq, etc.) Additionally, you don't need the DVD, you could just use yum... assuming you can get *any* DNS resolution.

Still waiting for a fix to glibc to test. Thanks.

Comment #69 From Jonathan Steffan 2008-12-03 21:28:44 EDT

Annd.. sorry:


```
# No nameservers found; try putting DNS servers into your
# ifcfg files in /etc/sysconfig/network-scripts like so:
#
# DNS1=xxx.xxx.xxx.xxx
# DNS2=xxx.xxx.xxx.xxx
# DOMAIN=lab.foo.com bar.foo.com
```

So you would need to define DNS1=127.0.0.1, etc.

Comment #70 From Andre Costa 2008-12-04 07:31:05 EDT

Ok, now things really don't make any sense to me.

I tried to enable named, but its config file is too complicated; since I just want DNS caching, I went for dnsmasq. The standard F10 RPM has IPv6 support enabled, so I grabbed the SRPM and generated a custom RPM with COPTS=-DNO_IPV6, which apparently worked just fine:

```
Dec  4 09:53:12 localhost dnsmasq[5677]: compile time options: no-IPv6
GNU-getop
t no-ISC-leasefile DBus no-I18N TFTP
```

However, if I start dnsmasq and point /etc/resolv.conf to 127.0.0.1, queries are received by dnsmasq and are forwarded to "real" DNS servers -- but replies are never received (I checked with Wireshark). If I replace 127.0.0.1 by the real DNS servers (eg. OpenDNS), queries go out and replies come in as expected.

Could this be a firewall issue? Or is it a dnsmasq issue? I enabled access to port 53 no iptables, here's the /etc/sysconfig/iptables file generated by s-c-f:

```
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
-A INPUT -m state --state NEW -m tcp -p tcp --dport 53 -j ACCEPT
-A INPUT -m state --state NEW -m udp -p udp --dport 53 -j ACCEPT
-A INPUT -m state --state NEW -m tcp -p tcp --dport 22 -j ACCEPT
-A INPUT -j REJECT --reject-with icmp-host-prohibited
-A FORWARD -j REJECT --reject-with icmp-host-prohibited
COMMIT
```

Some additional info:

```
- I disabled ipv6 adding these to modprobe.conf:
install ipv6 /bin/true
alias net-pf-10 off
```

Any help would be much appreciated, this "DNS Hell" is a huge pain in the a**
=/

Comment #71 From Chris Tirpak 2008-12-04 10:03:51 EDT

I pulled my hair out trying to find this. IMHO, this bug is not a 'medium' priority - it makes F10 useless. I disagree that users should have to go and try and find downstream software that uses glibc when it is glibc that changed (and I read every entry in this thread). F6,7,8,9 were all fine on the exact same hardware.

This should be critical priority not medium. Forcing a user to install BIND or DNSMASQ as a work around is utter nonsense. I've gone back to F9 until this is resolved.

Comment #72 From Chris Tirpak 2008-12-04 10:46:31 EDT

I pulled my hair out trying to find this. IMHO, this bug is not a 'medium' priority - it makes F10 useless. I disagree that users should have to go and try and find downstream software that uses glibc when it is glibc that changed (and I read every entry in this thread). F6,7,8,9 were all fine on the exact

same hardware.

This should be critical priority not medium. Forcing a user to install BIND or DNSMASQ as a work around is utter nonsense. I've gone back to F9 until this is resolved.

Comment #73 From Andre Costa 2008-12-04 12:41:43 EDT

Agreed, it's making F10 literally unusable, I'll have to boot on Windows XP to work. This is a *huge* step back for Fedora -- and for Linux. Changing glibc and saying the solution is to wait for everone else to adapt their apps is complete nonsense, if this was indeed necessary then F10 should have come out only when everything was ready and working as it should (and always did on previous versions).

I can barely upgrade my system -- have to get past dozens of "<urlopen error (-2, 'Name or service not known')>". It's even mentioned on latest Fedora News issue as "Strange Resolution Problems" [http://fedoraproject.org/wiki/FWN/LatestIssue#Strange_Resolution_Problems], so it's affecting even Fedora developers.

If there had been a warning like "*** ATTENTION *** F10 will come with changes on glibc that might severely affect DNS resolution and make your computer go back to dial-up era" I would have skipped F10 and waited for F11... =/

Comment #74 From Max Kanat-Alexander 2008-12-04 14:19:00 EDT

I've posted very clear workaround instructions at:

<http://www.fedorafaq.org/f10/#dns-slow>

Those instructions are working for me.

Comment #75 From Andre Costa 2008-12-04 14:36:53 EDT

That's good, I hope it helps others. However, neither dnsmasq (my preferred choice) nor named are working for me as caching name servers (see post #70).

Comment #76 From Chris Tirpak 2008-12-04 15:53:12 EDT

(In reply to [comment #74](#))

> I've posted very clear workaround instructions at:

>

> <http://www.fedorafaq.org/f10/#dns-slow>

>

> Those instructions are working for me.

I saw these too late but they look like they might be a good workaround to this bug. As I said previously, this bug makes F10 unusable for me out of the box and the workaround should not be necessary in the first place.

I've fallen back to F9 which does not have this bug. I've done exactly what I suspect countless others have silently done.

Also, lets face it, most users will never make it to a point where they find your instructions, especially a noob. And many a noob will take one look at them and give up on F10 or Linux entirely and say to their friends "yeah, I installed Linux/Fedora once and all I had was a bunch of networking issues".

I hope those capable of fixing this will rethink the situation and fix this.

Comment #77 From Andre Costa 2008-12-04 20:09:37 EDT

Newsflash: things were so strange (see post #70) that I decided to start all over again. I reinstalled F10 from scratch, and now things are looking better (that is: now I am able to workaround the most critical issues):

- NetworkManager still denies me access to eth0 configuration with static IP, but system-config-network-tui allowed me to workaround it, as others have suggested

- I enabled dnsmasq as my caching nameserver (the same way I was doing before), and now it is working as it should

I am following a cautious approach now: I'll make changes incrementally, so that I can see if/when things break. I still haven't applied any updates, I did not try to disable IPv6 in any way, and I did not mess with firewall settings.

Next step: apply all latest updates.

Comment #78 From Behdad Esfahbod 2008-12-05 00:22:35 EDT

Interesting. I've been seeing this on my two F10 machines too. I thought something's messed up with my home network, but it seems unlikely.

For me, ssh is affected too. When doing "svn commit" to gnome.org servers, sometimes (say, 10% of the times) I get a DNS error. Retrying immediately works. Also affects yum for me. I assume my firefox getting stalled at name resolution may be related too. When it happens, multiple tabs get stalled. Then after some 10-20 seconds they all work.

Comment #79 From Jonathan Steffan 2008-12-05 00:45:24 EDT

(In reply to [comment #78](#))

> I assume my firefox getting stalled at name resolution may be related too.

Yes. A workaround (until this issue is properly fixed) for firefox is to set "network.dns.disableIPv6" to true via about:config

Comment #80 From Tobias Ringstrom 2008-12-05 01:53:06 EDT

I'm getting this occasionally (but often enough to be annoying) with the Subversion client resolving a local server name in our company LAN. The DNS server is running Windows Server 2003. It happens on all machines upgraded to F10, but it never happened with F9 and earlier.

I've traced it using Wireshark, and what I see is that both an A and an AAAA query are sent simultaneously, and that the server sometimes replies to only one of them. It's as if it would consider the second request to be a duplicate, but that is only speculation, of course.

Comment #81 From Andre Robatino 2008-12-05 02:59:43 EDT

This bug is not limited to x86_64. Can the Platform be changed to "All"?

Comment #82 From Christopher Stone 2008-12-05 12:44:04 EDT

Hi, I am running into this bug also with a php script running apache. Note, the php script works perfectly fine when run as a normal user, but fails 100% of the time when run through a web server as the apache user.

```
<?php
$fp = fsockopen ("redhat.com", 80, $errno, $errstr, 30);
if (!$fp) echo "$errstr ($errno)<br>\n";
else { echo "worked"; fclose($fp); }
?>
```

The steps outlined in [comment #74](#) work as a suitable work-around.

Note, I do not use comcast, my ISP is AT&T/Yahoo.

Comment #83 From Kevin Fenzi 2008-12-05 12:59:01 EDT

In reply to [comment #63](#):

>Read the thread. I've already said multiple times that buggy programs which
>don't use AI_ADDRCONFIG will perform AAAA lookups. File bugs for the programs
>which cause the lookups.

It's sounding like there are a ton of such programs. ;(

Can we possibly get a workaround in glibc now, and then in rawhide look at finding all these broken programs and trying to fix them?
We don't want to try and do this in a stable release, IMHO.

Comment #84 From Andre Robatino 2008-12-05 16:02:30 EDT

Why was the platform changed back to x86_64 only? Is it really a separate bug with the same symptoms for i386?

Comment #85 From Stephen Adler 2008-12-06 08:58:16 EDT

OMG... The comments are endless on this!!!! Feroda team this is a totally urgent bug to fix!!!!

Here's my story. As everyone else here, my general access to the internet was slow and often fatal. (i.e. firefox would tell me it could not resolve names like google.com and yahoo.com)

I then went off the following path to fix it. Bascially its to setup a caching name server. My setup is different from most of you. I have 3 servers, 2 of them run red hat enterprise linux 5.2, the 3rd is my fedora 10 desktop. My main server is server01 which run nis, dhcpd etc. So now I setup named to run off of it. I had to install the caching-nameserver package which gave me a template named.conf to work from. Once setup, and with my fedora 10 using it as the domain nameserver, my access to the internet is blazing fast! I mean totally night and day.

As I was debugging the configuration of the caching name server, I did notice the A and AAAA lookups. What I did was run named in the foreground with all error messages going to stderr. (i.e. named -f -g) When fedora 10 desktop was doing queries, it would hit it with A and AAAA lookups. It may be a bit more complicated. When I first set up the caching name server I didn't pay much attention to access field so that it would only allow 127.0.0.1 to do its name resolution. Thus my fedora 10 desktop would try then fail, and try again, then fail and with each of these tries, it would do first an A lookup then an AAAA lookup. So it may be that for some reason the A lookup is failing first followed by a AAAA lookup. Anyway, once I allowed named to let others on its subnet to query from it, my fedora 10 desktop was then properly resolving names and now its so blazing fast it seems unreal. When I surf to yahoo.com, the whole page pops up almost instantaneously. Before it would take a while (10 to 20 secs) to draw the whole front page of yahoo.com. I supposed the reason for the slow drawing of the yahoo.com web page was the multiple dns lookups it had to perform.

Comment #86 From Christopher Stone 2008-12-06 11:20:10 EDT

Just curious, what kind of bugs would be considered more urgent than this for Fedora 10?

Are there any bugs for Fedora 10 which are a higher priority than this one? I am amazed at the bug triaging going on here...

Comment #87 From Christopher Stone 2008-12-06 11:28:40 EDT

I have added an entry to the Common Bugs page:
https://fedoraproject.org/wiki/Common_F10_bugs#DNS_Resolver_not_Reliable

Feel free to add more detailed information to the wiki text.

Comment #88 From Tom Horsley 2008-12-06 11:31:51 EDT

Yea, if it is hard to fix, why not just retrieve the DNS resolver code for the old glibc, call the result glibc-2.9-3 and do 2.9-4 someday when you can actually make it work correctly (or even just stop trying to improve things that work perfectly fine).

Comment #89 From David Miller 2008-12-07 04:09:20 EDT

I definitely think that would be the appropriate way to handle

this, revert to the old querying behavior until a better workaround is figured out.

Comment #90 From Andre Costa 2008-12-07 16:37:38 EDT

Continuing my report started with [comment #77](#): after a full reinstall, dnsmasq is finally working as a local DNS cache (don't know exactly what I did before to prevent this from happening).

However, something is still seriously broken with DNS on F10, since some queries are repeated over and over again, as if they simply didn't stick on dnsmasq's cache (for example if I *repeatedly* run 'dig www.mozilla.com', it never returns immediately). It specially hurts web browsing, but also affects ssh, mail clients etc. (pretty much everything... =/)

Even if technically some (most?) apps are not performing AAAA lookups the right way, I think it is a very bad decision to simply make glibc "right", break everything else and let the rest of the world play catch up. It only hurts user experience and keeps F10 from really shining.

Comment #91 From Jakub Jelinek 2008-12-08 09:54:21 EDT

Please try <http://kojipkgs.fedoraproject.org/packages/glibc/2.9/3/> which has temporarily the simultaneous IPv{4,6} query disabled to give us time to see how can we keep the lookups fast while still not timing out on buggy DNS servers.

Comment #92 From Ben Williams 2008-12-08 11:03:44 EDT

As soon as this issue can be resolved Fedora Unity will do a F10 re-spin to help the community

Comment #93 From Jonathan Steffan 2008-12-08 12:55:28 EDT

(In reply to [comment #91](#))

> Please try <http://kojipkgs.fedoraproject.org/packages/glibc/2.9/3/>

After brief testing, it looks like these packages fix the issue. Please submit to Bodhi (and updates-testing) for a wider testing sample.

Comment #94 From Jonathan Steffan 2008-12-08 13:06:06 EDT

I don't have access to submit this to Bodhi. Please do so promptly.

Comment #95 From Scott Glaser 2008-12-08 13:35:37 EDT

After some local testing on my laptop, it looks like these packages fixes the issue. Please submit to Bodhi (and updates-testing).

Comment #96 From Fedora Update System 2008-12-08 15:34:57 EDT

glibc-2.9-3 has been submitted as an update for Fedora 10.
<http://admin.fedoraproject.org/updates/glibc-2.9-3>

Comment #97 From Ulrich Drepper 2008-12-08 17:19:05 EDT

(In reply to [comment #95](#))

> After some local testing on my laptop, it looks like these packages fixes the
> issue.

No, the problem is work-around. It's not fixed.

And that's the critical point here. We will re-enable the code soon again, with some changes to handle broken DNS servers a bit differently. These will have to be tested. All those people complaining here had the chance to test rawhide for weeks and months and apparently didn't do it. There was one single report which didn't point exactly at the problem. As far as this is concerned, rawhide failed miserable.

If you don't want to get stuck with badly working DNS again use the test release as soon as we have one. This bug should remain open until the problem is actually fixed.

Comment #98 From Ulrich Drepper 2008-12-08 17:23:28 EDT

*** [Bug 471450](#) has been marked as a duplicate of this bug. ***

Comment #99 From Tom Horsley 2008-12-08 17:35:57 EDT

>All those people complaining here had the chance to test
>rawhide for weeks and months and apparently didn't do it.

Um, the original report was 2008-08-21 18:51 EDT by Tom Horsley, more than 2 months before the f10 final release, followed up only a few days later by several straces and tcpdumps. I don't know how the problem could have been more obviously pointed out before the release.

Comment #100 From Andre Costa 2008-12-08 19:11:19 EDT

Good news, patched glibc version indeed stopped sending AAAA queries. Thks for the workaround, hope you guys find out the definitive solution soon.

Comment #101 From Ulrich Drepper 2008-12-08 19:25:58 EDT

(In reply to [comment #99](#))

> Um, the original report was 2008-08-21 18:51 EDT by Tom Horsley, more
> than 2 months before the f10 final release,

There was one single person (or two) and it was at no point clear that this is a) a problem with the DNS server (could as well be network issues) and b) that this is a wide-spread problem.

Realize that disabling this code is a correctness issue and a performance issue for those people who are not using braindead DNS servers (which certainly makes 95% of the people or more).

Comment #102 From Ulrich Drepper 2008-12-08 19:28:35 EDT

(In reply to [comment #100](#))

> Good news, patched glibc version indeed stopped sending AAAA queries.

What? Nothing should have changed in this regard. The only thing that changed is that the requests are not at the same time and hence broken servers can handle one requests after the other. if you see anything else, that's not expected.

Comment #103 From Andre Costa 2008-12-08 20:10:35 EDT

(In reply to [comment #102](#))

> (In reply to [comment #100](#))

> > Good news, patched glibc version indeed stopped sending AAAA queries.

>

> What? Nothing should have changed in this regard. The only thing that changed
> is that the requests are not at the same time and hence broken servers can
> handle one requests after the other. if you see anything else, that's not
> expected.

Sorry, I can't really say what has changed. It's just that I recall seeing AAAA queries/replies while monitoring traffic with wireshark during my first attempts to make F10 work (before I reinstalled), and now I am not seeing them anymore with patched glibc.

But, forget what I said about AAAA queries: I just wanted to say it seems to be working better now.

Comment #104 From Jan Teichmann 2008-12-09 10:03:54 EDT

I also have big problems with the DNS on a VPN connection using OpenVPN and the NetworkManager. It's a new installed, virgin F10 on a Thinkpad R61. If I'm connected to my university using OpenVPN I couldn't access most of the internet sites because the servers are unknown. The big problem is that the license server of my university is also unknown :(

Comment #105 From Mike 2008-12-09 10:29:39 EDT

(In reply to [comment #13](#))

```
> At the first glance, this reply says "I don't know the [IPv6] address, but here
> is some information which may be useful":
```

```
>
> AAAA? mirrors.fedoraproject.org. 1/1/0 mirrors.fedoraproject.org. CNAME
> wildcard.fedoraproject.org. ns: fedoraproject.org. SOA fedoraproject.org.
> hostmaster.fedoraproject.org. 2008082802 28800 7200 2419200 86400 (113)
>
> Apparently this info isn't useful, so wget wats for IPv4 answer for 5 seconds,
> but it does not come.
```

I had some issues with ipv6 lookup failures flooding my log files for f10 and following advice from an expert in networking it turned out that I could cure my problems by adding the line:

```
OPTIONS="-4"
to the file /etc/sysconfig/named
and then:
# service named restart
```

This prevents any ipv6 lookups from dns and the network lookups work well.

I wonder if this might be relevant here also?

Comment #106 From Fedora Update System 2008-12-09 23:36:58 EDT

glibc-2.9-3 has been pushed to the Fedora 10 stable repository. If problems still persist, please make note of it in this bug report.

Comment #107 From Ulrich Drepper 2008-12-10 00:04:21 EDT

No, we don't close it. This is a work-around.

Comment #108 From Phil Oester 2008-12-10 09:41:32 EDT

Just a further datapoint on this, since I too spent a few days scratching my head on it. It looks like what changed in F10 is that both the AAAA and A requests are sent using the SAME SOURCE PORT, while pre-F10 used different source ports for the two requests.

For me, that change spelled trouble in the form of a race for my loadbalancer. I saw this:

- 1) receive A request, creating session table entry with NAT'd reply IP
- 2) receive AAAA request on port x, reusing session table entry from #1
- 3) respond to AAAA request on port x and remove session table entry
- 4) loadbalancer receives response from DNS server for A request, but since session table entry (with VIP response IP) is gone, it simply forwards the traffic, so client receives a reply from a different IP (the IP of the server itself, NOT the vip) and ignores it

So for me, the simple solution to this is to go back to the old behaviour of having the A and AAAA requests use unique source ports. Wouldn't that be more secure anyway? Seems like a step backward to reuse the port.

Comment #109 From Maurizio Paolini 2008-12-11 02:55:27 EDT

I would like to confirm [comment #108](#) From Phil Oester. I tried a tcpdump on both the client that makes the A and AAAA requests *and* on the server where the DNS server is running (my ISP is in between). The

DNS server receives both requests and replies to both requests with two packets, but only one of those packets arrive at the client. Since my ISP performs IP masquerading (private IP) I guess that [comment #108](#) is a perfect explanation.

Comment #110 From Andre Robatino 2008-12-11 09:19:29 EDT

> [glibc-2.9-3](#) has been pushed to the [Fedora 10 stable repository](#).

No, it hasn't (or any of the other updates from the 10th as indicated by [fedora-package-announce](#)). The ones from the 11th, however, showed up immediately.

Comment #111 From Tom 2008-12-12 13:12:06 EDT

I just installed the [glibc-2.9-3](#) series packages from koji, and they seem to cure the problem. I'll keep my fingers crossed.

Comment #112 From Jonathan Steffan 2008-12-12 13:28:51 EDT

(In reply to [comment #108](#))

This is also what we saw. DNS that was routed directly to a server would work, but any DNS requests that went over a firewall or load balancer did not.

Comment #113 From Maurizio Paolini 2008-12-12 15:24:35 EDT

(In reply to [comment #111](#))

> I just installed the [glibc-2.9-3](#) series packages from koji, and they seem to
> cure the problem. I'll keep my fingers crossed.

Let me ask: how is the problem solved in [glibc-2.9-3](#)? Using different source ports for the two queries (A and AAAA)? Or making the two queries sequentially instead of in parallel?

It seems that [comment #108](#) clearly explains the origin of the problem...

Comment #114 From Andre Robatino 2008-12-14 01:13:56 EDT

[glibc-2.9-3](#) has finally been pushed (for real, this time) to the [updates-released mirrors](#).

Comment #115 From Harshad RJ 2008-12-14 09:19:36 EDT

+1 for [comment #113](#)

(how is the problem solved in [glibc-2.9-3](#)? Using different source ports for the two queries (A and AAAA)? Or making the two queries sequentially instead of in parallel?)

Comment #116 From D. Hugh Redelmeier 2008-12-14 16:53:39 EDT

Re: [Comment #108](#) From Phil Oester

I assume that these queries are using UDP (DNS can use TCP).

- 1) NAT is an evil fudge. But a handy and widely deployed one.
- 2) the UDP protocol does not have a notion of a session.
- 3) NAT software typically "invents" some kind session notion for UDP. These inventions all have weaknesses that break certain legitimate uses of the UDP protocol.
- 3) your NAT software imposes a particular notion of a session. This causes legitimate use of UDP to fail. In particular, the NAT software is misbehaving when it tears down the session (which blocks the second reply).

So: the bug is in your load balancer.

I don't know your world so I don't know the best way for you to avoid the problem. Perhaps you should run a local (caching only?) DNS server. Local to your site, bypassing the load balancer.

Forcing queries to use TCP instead of DNS could help because TCP does have a notion of session. On the other hand, not all DNS servers are willing to talk TCP and it does cause some overhead.

Comment #117 From Need Real Name 2008-12-15 14:45:34 EDT

All I know is I did a yum update today, and all the computers I had 'fixed' this bug on, can no longer get online.... again... :-p

Comment #118 From Need Real Name 2008-12-15 14:56:34 EDT

yup, 2.9-3 seems to have addressed my original issue. After I un-did all the changes I had to do originally to make it work, it is back online! Everyone that worked on it, thanks!

Comment #119 From William F. Acker 2008-12-15 17:49:43 EDT

Now that we have the updated version of glibc, I can no longer resolve addresses via IPv6. if i put, for example, "nameserver 207.224.49.209", everything works. But if I use 2001:470:80ee:0:207:e9ff:fe09:c032, I can only resolve addresses using host or nslookup.

Comment #120 From Mars Magic 2008-12-22 12:26:24 EDT

I updated to glibc 2.9-3 last night along with a lot of other updates. I have been having no resolver issues. I have bind running and therefore have no /etc/resolv.conf. Until I created a /etc/resolv.conf with nameserver 127.0.0.1, I had absolutely no IPv4 DNS resolution in yum, firefox, etc. Only dig and host were able to resolve addresses.

Comment #121 From Robert Scheck 2008-12-27 16:26:45 EDT

Jakub, can we investigate for Rawhide now as the problem is worked around for Fedora 10? It seems worse to me, that IPv6 nameservers are no longer usable as it is mentioned in [comment #119](#). IPv6 is the future, IPv4 has anyway to die...

Comment #122 From Ulrich Drepper 2008-12-28 01:03:24 EDT

Don't overload bugs. If you think you found a new Problem open a new bug and don't mention it in some unrelated BZ.

Comment #123 From David Miller 2008-12-28 02:41:19 EDT

How is it unrelated? This ipv6 issue people are reporting now was introduced by the workaround for this specific bug.

Why in the world wouldn't we want that important information logged here? It's a regression added by the fix for this bug, so it's just as much a part of this bug as any new one people would open.

Opening a new bug is just more red tape, this is one issue.

Comment #124 From Ulrich Drepper 2008-12-28 12:26:43 EDT

(In reply to [comment #123](#))
> How is it unrelated?

It cannot possibly be related. Just because you cannot see it doesn't change that. Not opening a new bug means the details are hidden between all the irrelevant other information.

Comment #125 From Ulrich Drepper 2008-12-29 15:19:54 EDT

I fixed the handling of installations with just IPv6 name servers upstream. Will be in the next build.

Comment #126 From Kai Meyer 2008-12-30 00:12:50 EDT

[Created an attachment \(id=327953\) \[details\]](#)
Wireshark results of just DNS traffic during a 'wget -O /dev/null google.com'

Comment #127 From Kai Meyer 2008-12-30 00:16:45 EDT

Sorry, new to bugzilla. To give more information about my last post, I meant to also say that this problem is not resolved for me with the newest version of glibc (2.9-3). I'm on Fedora 10 x86_64.

When this problem first cropped up for me in Rawhide (about a month before release) the best solution I could find was to add:

```
install ipv6 /bin/true  
in /etc/modprobe.conf
```

This caused another error to come up all the time "E: socket-client.c: socket(): Address family not supported by protocol" which also caused a 5-10 second delay in whatever action I was taking, whether it was a flash video in firefox, an mp3 in amarok, or a video with vlc. I would prefer the delay with the ipv6 module not even loaded to the delay caused by the dns resolving failure.

Comment #128 From Ulrich Drepper 2008-12-30 11:15:00 EDT

Again, 2.9-3 cannot possibly cause the same problem for people with broken servers and firewalls as the previous version. In fact, if F9 and earlier worked for you, this versions DNS lookup will work.

There have been a few other problems which are now fixed and will be in 2.9-4 but those have nothing to do with this specific problem.

Comment #129 From Kai Meyer 2008-12-30 13:23:07 EDT

Ulrich, I got your reply, and was going to provide some supporting evidence of other computers on my network that do not have this problem (A i386 Fedora 10 laptop, and an Ubuntu 8.04 laptop) to rule out firewall or ISP issues. Today, for whatever reason, I am not getting the same delays in AAAA lookups that I was getting last night. I haven't rebooted or reloaded the ipv6 kernel module. I'll cross my fingers and hope that the issue is in fact resolved, and just pretend that what I saw last night was just a misaligned star or something.

Thanks

Comment #130 From Marc Rechte 2009-01-01 11:40:20 EDT

Hello,
I did a fresh F10 install on a machine which was running F9 before. This is a machine configured with named as a caching / forwarder. I did an update of the system yesterday and since then I have no more name resolution. Host doamin or dig domain both work fine but if I try ping, elinks or firefox, they say this same domain does not exist ! My F10 machine has become unusable since this update !!!

Comment #131 From Marc Rechte 2009-01-01 12:46:09 EDT

Sorry my previous post is wrong. I had no 'nameserver 127.0.0.1' sentence in /etc/resolv.conf (in fact no nameserver at all). Since I explicitly set it up, it is now OK. Sorry for the disturbance on this thread.

Comment #132 From Andre Costa 2009-01-22 20:15:20 EDT

Any update on a final solution for this? AFACIS this is still a workaround, right?

Here on my box DNS still works significantly better on Windows XP: name resolution happens at 1-2s, while on Fedora 10 it takes usually around 5-8s, which leads to lots of timeouts and makes internet experience painful. Using dnsmasq as local DNS cache improves things but still isn't good enough.

Even though my ISP DNS setup probably has some issues, the facts are that Windows "just works" and that F10 made things much worse on Linux.

I would love to provide any additional data that could help, so, please, let me know if I can help.

Comment #133 From Ulrich Drepper 2009-01-22 21:08:13 EDT

(In reply to [comment #132](#))

> Here on my box DNS still works significantly better on Windows XP: name
> resolution happens at 1-2s, while on Fedora 10 it takes usually around 5-8s,
> which leads to lots of timeouts and makes internet experience painful.

This has nothing to do with this bug.

What you have is a DNS server which doesn't serve IPv6 replies and a setup which has IPv6 addresses configured on applications which don't use `getaddrinfo` correctly (as I explained many times already). The resolution of the bug will not change this at all.

Set up your machines correctly (disable IPv6) and/or files bugs to get the applications fixed.

> Even though my ISP DNS setup probably has some issues, the facts are that
> Windows "just works" and that F10 made things much worse on Linux.

You most probably compare apples and oranges. Set the machines up identically.

Comment #134 From Tobias Ringstrom 2009-01-23 03:28:59 EDT

Ulrich,

I don't think it's likely that ordinary users will figure out that they need to disable IPv6 to get good name resolution performance. They will just get worse performance than they get with other OS'es, and it makes Fedora look bad. Even if they do figure out that IPv6 should be disabled, it's not very easy to do. Even if "Enable IPv6 configuration for this interface" is unchecked in `system-config-network`, the interface still gets an IPv6 link-local address.

What are your thoughts on David Miller's comments in this issue that every IPv4 interface automatically get a link local IPv6 address ([comment #42](#)), and that a certain other OS does not issue AAAA requests if only link local IPv6 are configured ([comment #45](#))? Would it be a good idea to do this in Linux too? It would likely take care of the bulk of the issues people are having.

People in IPv6 supported networks probably have working AAAA name resolution, and it seems common for networks without IPv6 support to break AAAA lookups in various ways.

Comment #135 From Andre Costa 2009-01-23 07:25:35 EDT

(In reply to [comment #133](#))

> (In reply to [comment #132](#))

> > Here on my box DNS still works significantly better on Windows XP: name
> > resolution happens at 1-2s, while on Fedora 10 it takes usually around 5-8s,
> > which leads to lots of timeouts and makes internet experience painful.

```
>
> This has nothing to do with this bug.
>
> What you have is a DNS server which doesn't serve IPv6 replies and a setup
> which has IPv6 addresses configured or applications which don't use getaddrinfo
> correctly (as I explained many times already). The resolution of the bug will
> not change this at all.
>
> Set up your machines correctly (disable IPv6) and/or files bugs to get the
> applications fixed.
```

I already tried disabling IPv6, and it didn't improve things much. Maybe I did it the wrong way, though. I disabled IPv6 for eth0, and tried aliasing ipv6 and net-pf kernel modules to 'off', and I also tried the "install xxx /bin/true" approach, but neither seemed to have improved things much. I'll give it another try.

BTW: which one is the right approach to completely and correctly disable IPv6?

```
> > Even though my ISP DNS setup probably has some issues, the facts are that
> > Windows "just works" and that F10 made things much worse on Linux.
>
> You most probably compare apples and oranges. Set the machines up identically.
```

Yes, I know I am comparing apples and oranges, but I wanted to show that both systems on the same machine, with the same router and ISP settings perform completely different (granted, network settings are probably different, since XP probably ignores IPv6 completely). And I'd love to setup the machines identically, but you're telling me that in order to do so I need to figure out which components are broken and go after them. This is not as simple as configuring network settings.

The bottomline is: even if it's not directly related to this bug anymore, Fedora 10 as a whole provided a bad internet experience out-of-the-box, and this holds true almost 5 months after this bug has been reported. This sucks big time.

I will gladly file bug reports, but relying on users to do it for different apps when they don't even know which ones are broken isn't the right thing to do IMHO. This should have been done internally _before_ F10 was released, and now that F10 is in the open, Fedora developers should be leading this effort.

It's like selling a state-of-the-art car (ok, giving it for free ;-)) that doesn't work as it should, and in response to complaints say "ok, some parts need to be fixed, you figure out which ones they are and go bother the respective manufacturers".

If you already know which apps are broken, please at least provide a list, and bug reports will start coming in.

Comment #136 From Andre Costa 2009-01-23 07:32:52 EDT

(In reply to [comment #134](#))

```
> Ulrich,
>
> I don't think it's likely that ordinary users will figure out that they need to
> disable IPv6 to get good name resolution performance. They will just get worse
> performance than they get with other OS'es, and it makes Fedora look bad. Even
> if they do figure out that IPv6 should be disabled, it's not very easy to do.
> Even if "Enable IPv6 configuration for this interface" is unchecked in
> system-config-network, the interface still gets an IPv6 link-local address.
```

Thks, Tobias, these are my thoughts exactly (see [comment #135](#)). All these "DNS+IPv6+[whatever]" issues do make Fedora look bad, specially because there are no precise instructions on how to workaroud it (eg. disable IPv6).

Comment #137 From Mike 2009-01-23 16:26:05 EDT

I must admit that adding the line:

```
OPTIONS="-4"
```

to the bottom of the file /etc/sysconfig/named after installing bind-chroot and starting the "named" dns service fixed the issues for me.

However it seems that maybe others have additional dns issues. I am now

confused as to whether or not those who remain with dns problems attempted the above work-around or not?

Comment #138 From Andre Costa 2009-01-25 17:27:12 EDT

I tried using named, but I find dnsmasq easier to configure. However, it has no explicit support to disabling IPv6, so I probably should give named a try again. I've read somewhere that by simply installing named and pointing /etc/resolv.conf to 127.0.0.1 would give me local DNS caching (no extra configuration needed, aside from OPTIONS=-4; is that right?). Any pointers to a good HOWTO for Fedora?

I also created a /etc/modprobe.d/ipv6-off file with

```
alias ipv6 off
alias net-pf-10 off
```

It seems to work, since no ipv6 modules are being loaded by the kernel. I also disabled ip6tables service.

It doesn't seem to have improved things much, though.

I've also been monitoring traffic to/from port 53 with Wireshark, and here's the list I've compiled so far of apps which make AAAA queries: ssh, whois and (guess what?), yum. Seeing yum on this list is IMHO a clear example that Fedora developers should be leading this "DNS sanitization" effort, since it's a tool essential to the system, used primarily (uniquely?) by Fedora and maintained by its own personnel.

(I found something weird, maybe dnsmasq's fault: AAAA queries are tried first for the correct FQDN, and later for the FQDN + ".localdomain" -- which is clearly bogus and a useless query. Any idea on how to fix this?)

From all these apps, I have only been able to "fix" ssh by adding "AddressFamily inet" to /etc/ssh/ssh_config. I did not find any simple way to configure the other apps to avoid IPv6, so I guess I'll need to file bug reports for them. As I said on previous posts, no problem, I will gladly do this if it helps with this DNS/IPv6 hell.

However, what should I say exactly on the bug report?

Comment #139 From Andre Costa 2009-01-25 17:33:42 EDT

... on a second thought, there's nothing really "clear" for me regarding all this mess, so I can't really say it's yum's fault that AAAA queries are being sent when it queries the servers for updates. Please apologize if I made wrong assumptions.

Comment #140 From Andre Costa 2009-01-29 12:14:56 EDT

Some news:

- wget also tries to make IPv6 queries. Passing "-4" on the command line or adding "inet4_only = on" to /etc/wgetrc fixes this

- I found some references to yum's (actually Python's) "IPv6 obsession": <http://lists.baseurl.org/pipermail/yum/2006-November/020463.html> (Nov/2006 = (). The closest thing reported on Bugzilla seems to be [bug 171664](#) (but the reporter seems to be more concerned with the bad encoding than with the fact that it shouldn't be doing IPv6 queries on a IPv4-only system)

- I tried using named, but for some reason its queries to the DNS root servers never get any answer (could be my ISP's fault?), so I fell back to dnsmasq.

Comment #141 From Maurizio Paolini 2009-01-29 12:44:20 EDT

(In reply to [comment #140](#))

> Some news:

>

> - wget also tries to make IPv6 queries. Passing "-4" on the command line or
> adding "inet4_only = on" to /etc/wgetrc fixes this

The point is that "modern" applications should resolve names using "getaddrinfo(3)" (in place of gethostbyname). There is a flag that *can* be specified by the application in order to require ipv4-only query; however the default is to query for both ipv4 *and* ipv6.

So, in my opinion, applications like wget should not be blamed in this case.

Probably there should be a way to globally require that "getaddrinfo" makes ipv4 only requests when called with default settings by applications, either by means of some global configuration (perhaps something in /proc/sys/...) or by some kind of heuristics (which I don't like very much).

Comment #142 From Andre Costa 2009-01-29 13:07:20 EDT

100% agreed, I just mentioned wget's settings for future reference (since I have already mentioned settings for ssh).

Comment #143 From Andre Costa 2009-01-31 10:28:37 EDT

Ok, final comment: to make a long story short, I finally solved my problem, and it was not Fedora's fault after all.

I've been trying all I could to pinpoint the cause of this slow DNS resolution problem, and I finally managed to borrow a HUAWEI E226 3G USB modem from a friend. To my surprise, I experienced no delays at all regarding DNS resolution.

This meant either the problem was on my ISP or on my DI-624 router. Checking this was simple: I bypassed the router and connected the ethernet cable directly to the cable modem, and -- bingo! -- DNS was working as it should with my ISP.

After going through all router's settings and founding nothing suspicious (as far as I could see), and before throwing it away for good, I decided I should try one last measure: I reset the router to factory defaults, and reconfigured only the essentials. That did the trick, and now all is working as it should (with the router). I even turned dnsmasq off.

Even though this proves me wrong regarding my previous complaints, I am actually glad to realize that Fedora was not to blame. My faith in Fedora has been restored =) Also, I've learned my lesson, next time I will investigate further before making any conclusions about matters I'm not familiar with (such as networking low-level details).

So, please apologize for all the noise, I hope this thread at least helps others.

PS: I do believe that bugs should be filled about those apps (yum/Python, whois etc.) that insist on making AAAA queries even when IPv6 is disabled system-wide. But, now this is just a minor issue...

Comment #144 From Need Real Name 2009-02-07 14:19:09 EDT

The current rawhide seems to be quite thoroughly hosed. I installed F11 Alpha and updated to rawhide. Now I have

```
glibc-2.9.90-3.i686
kernel-2.6.29-0.93.rc3.git10.fc11.i586
```

I disabled the installation of the ipv6 module. DNS resolution seems to mostly work with 'ping' and 'host', but not with yum, ssh or Firefox, and also the system could not connect to a LDAP server until I added the server in /etc/hosts. I can ping a machine by its name, but a second later ssh says it cannot resolve the address of the same machine. Firefox does not find any addresses at all. The DNS server on the local network is bind-9.3.4-6.0.3.P1.e15_2 running on CentOS 5.2.

Comment #145 From Mike 2009-02-19 17:41:17 EDT

By an odd coincidence this evening I booted two machines on my LAN to the new

F10 kernel and noticed that ssh from one to the other was taking around 5 seconds to connect - previous to rebooting networking was fine. I spent a considerable time checking all dns and network related settings and found nothing wrong. Out of desperation I powered down my Linksys WAG54G2 wireless router into which the ethernet from both machines was connected, and immediately networking speed was back to normal (the machines had detected the loss of the connection and re-established the connection without restarting anything on either machine!)

This sounds remarkably similar to the experience in [comment #143](#) - and was something I would not have expected to make a difference.

Maybe there is an explanation but it is not something I understand.

Comment #146 From Ulrich Drepper 2009-04-08 17:37:58 EDT

The rawhide build

<http://koji.fedoraproject.org/koji/buildinfo?buildID=97098>

of glibc contains changes to the DNS lookup. The problematic behavior is reenabled. But we are now handling the situation where only one reply is received differently. In that situation we are switching (permanently for that process) to a mode where the second request is sent only when the first answer has been received. I.e., we should transparently fall back to a slower mode for broken DNS servers.

This will mean, though, that people with these broken DNS servers will experience delays. There are ways around it, though:

- use nscd. Should be done anyway. This way only one delay per system start applies
- adding single-request to the options in /etc/resolv.conf. This will only try the mode for broken DNS servers

I decided to go this route and not fall back on the slow method because the number of people affected is relatively small and there cannot be a justification for the self-inflicted problems of the few to cripple the rest of the world.

Those with broken hardware are asked to test this glibc version. Please test it with and without nscd, with and without the /etc/resolv.conf option.

Comment #147 From Andre Costa 2009-04-09 17:21:49 EDT

Hi Ulrich,

as one of the people affected, I would like to give this a try to see how this goes. Is there a F10 version?

Comment #148 From Ulrich Drepper 2009-04-10 14:38:15 EDT

(In reply to [comment #147](#))

> as one of the people affected, I would like to give this a try to see how this
> goes. Is there a F10 version?

I don't think we have an F10 version. Jakub's in charge of all this.

But you can take the F11 binary, extract the libnss_dns.so.2 and libresolv.so.2 files, put them in a new directory, kill nscd, and then use LD_LIBRARY_PATH to point to the new directory with the files.

Comment #149 From Ulrich Drepper 2009-04-15 15:48:26 EDT

I have not seen any feedback at all so far on this. The code is active in rawhide and nobody complaint so far but the history of this bug (i.e., rawhide for F10) showed this doesn't say much. Not that many people run rawhide.

Anyway, we're not far away from F11. The code with my latest changes will be

activated unless I hear about problems. I don't expect any problems, but it's still necessary to verify.

Comment #150 From Tom Horsley 2009-04-15 17:55:59 EDT

If it actually made it to the mirrors, I probably downloaded it as an update on my f11 beta system - same hardware where the original bug was filed in an earlier release (still using comcast), and I haven't noticed any name lookup problems.

Poking around in yum.log on that partition, it looks like the last glibc I got was glibc-2.9.90-15.x86_64, so that probably does have the new code.

Comment #151 From Need Real Name 2009-04-16 07:58:56 EDT

I have

glibc-2.9.90-16.i686

on an up-to-date F11 i686 system, and it seems to work without problems. The ipv6 module is loaded and ipv6 is enabled in e.g. Firefox, and it works.

Comment #152 From Ulrich Drepper 2009-04-16 15:24:15 EDT

Thanks for the feedback. I think we can close this now. F10 could potentially get the change backported but it's not urgent. I leave this up to Jakub.

Note

You need to [log in](#) before you can comment on or make changes to this bug.

[Format For Printing](#) - [XML](#) - [Clone This Bug](#) - [Top of page](#)