

# Re: Problems with named default configuration in 6-STABLE

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2007-07/msg00213.html>

---

- *From:* Volker <[volker@xxxxxxxxxxx](mailto:volker@xxxxxxxxxxx)>
  - *Date:* Tue, 17 Jul 2007 20:39:58 +0200
- 

Doug,

On 07/17/07 18:14, Doug Barton wrote:

Volker,

I'm sorry to say that you've provided a great deal of incorrect information in this thread.

sorry? really? I couldn't find one!

Volker wrote:

Remember, AXFR requires a TCP transfer and not every firewall will happily let it pass.

This is true, although since to the firewall an AXFR looks just like any other stateful TCP connection out to the wide world, it's actually more likely (percentage wise) that this will succeed than it is that the DNS requests (using UDP) will. Obviously, for those that cannot transfer the zone, the hints mechanism is still in the comments.

I've seen setups (not mine but I also tend to suppress 53/tcp) to deny 53/tcp traffic. To see this as a common or uncommon setup is more likely a philosophical discussion.

I (partially) agree to the speedup effects you mentioned

I think you should read the paper that David posted, <http://www.imconf.net/imc-2004/papers/p15-malone.pdf> before you comment further.

## Re: Problems with named default configuration in 6-STABLE

I never said, using hint or slave will have no impact on speed. Where did I write that? Read: "I agree to the speedup effects."

It's also worth nothing that even if the only benefits were greater reliability vs. a root DDoS attack (which is sadly no longer a theoretical issue) and the substantial improvements to local NXDOMAIN answers, it would be worth it. Add the benefits of at worst a wash with overall root traffic for the root zone, and the extra benefits of also having local copies of ARPA and IN-ADDR.ARPA (which are much smaller, and usually more frequently queried than the root zone) and it's a clear win.

Ok, a DDoS attack against the root servers has happend a few times in the past but there has been no single moment in time when not a single root server could be reached. So yes, a DDoS attack is possible, slaving the root zone will make you still being able to resolve DNS addresses but this is a theoretic scenario to have none of 13 root DNS servers worldwide being unreachable.

I should add for the sake of completeness that not every DNS professional has reached the same conclusions I have

Sorry, but this reads as "nobody else has reached my wisdom". And it's comments like these which will get me away from using FreeBSD sometime in the future (while FreeBSD still being a great OS).

, however the main objection that is usually raised is not operational from the root server operators, rather it's that DNS admins who are not paying attention might miss a change that would prevent their local resolver from slaving the zone at some time in the future. Given that the IP addresses of the root servers hardly ever change, and given that we have 5 servers to choose from (and we only need one good transfer to make it work), and given that I (and as this thread points out, others) actually do pay attention, I don't think this is going to be a problem for us.

The point is, when relying on something which is not guaranteed to work ("SHOULD NOT reply to zone transfer requests" – RFC2870, 2.7) and you're using this in a default OS configuration, you'll have trouble *if* the remaining 5 root DNS servers refuse to answer zone transfer requests. Even if one after another is changed to refuse, you'll most likely notice after the last remaining root server will refuse to answer. In that situation, the average user will not be able to reach

Re: Problems with named default configuration in 6-STABLE

the internet at all as his DNS resolver will not work anymore.

We're not talking (at least I'm not) about the experienced admin. I am talking about the average user which does not know a single bit about DNS. Just for the case, the remaining 5 root DNS server refuse to AXFR: You want him (the average user) to figure out himself what's wrong with DNS if he can't reach the internet?

but if just 5  
out of 13 root servers support AXFR, your bind will sit for a while to  
find a root server responding to it's AXFR requests.

I'm sorry, but that comment means that either you haven't read the new named.conf, or you don't understand what you've read. Either way, you should seriously consider whether or not it's a good idea for you to continue offering DNS advice. The following:

You're wrong. I've read named.conf and understood. In it, there're only 5 (of the total 13) root DNS servers configured as the root zone's master servers. Where have I been wrong? What did I don't understand?

There are a total of 13 root DNS servers (not physical hosts) worldwide. You have configured these 5, which still respond to AXFR requests.

Yes, a single root server serving AXFR requests is enough to get the job done, but fact is, you're relying on only 5 out of 13 servers. Again, where have I been wrong?

I know you've got the wisdom... :(

BTW, so far we've only talked about the savings for an individual resolver. If you are responsible for a network of resolvers you can slave the zones from the roots on one server then slave them out to your network from your local master for an overwhelming savings of overall traffic to the roots. If you decide to do that, you should take a look at the ixfr-from-differences option to save yourself even more local traffic.

I never disagreed to the performance improvements.

Now I'm waiting for you to tell me one by one where I've "provided a great deal of incorrect information".

Volker

Re: Problems with named default configuration in 6-STABLE

Re: Problems with named default configuration in 6-STABLE

---

freebsd-stable@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-stable>

To unsubscribe, send any mail to "freebsd-stable-unsubscribe@xxxxxxxxxxx"