

Re: Re: Re: Bind9: rndc reload doesn't work for slave servers

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/isp/2006-08/msg00002.html>

- *From:* Edda Hochstrate <eh@xxxxxxxxxx>
 - *Date:* Tue, 01 Aug 2006 15:06:30 +0200
-

Brian Candler wrote:

If you are going to do a 'kill -HUP' every 5 minutes, say, then you are generating a lot of work for your own server and every customer's master server.

Once an hour. Then it takes 10 min. apx. for 3000 zones.

If you have a particular customer who has made a change and explicitly asks you to refresh their zone, then you can 'rndc reload <zone>' as said before.

Yes. But not, if your customer himself has 900 zones and wants your slaves to be as up-to-date as his primary. But maybe we will script it with 'rndc reload <zone>', right.

However, note that your slave server will poll each of the masters at the refresh interval in their SOA record anyway. So if a particular customer wants you to poll their zone more frequently, then they can just reduce the refresh time in their SOA record, and your server will honour their request. That gives you the best of all worlds – frequent polling for those customers who want or need it, and occasional polling for everyone else.

In our opinion as an ISP the refresh time is for the world of resolvers. Authoritative nameservers (especially slaves serving hidden primaries) should be as close as possible to the primary.

Best regards,
Edda

Re: Re: Re: Bind9: rndc reload doesn't work for slave servers

freebsd-isp@xxxxxxxxxxx mailing list

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

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