
Subject: Timeout handling

Posted by [oxygen](#) on Wed, 31 Aug 2005 16:27:35 GMT

[View Forum Message](#) <> [Reply to Message](#)

Looks like there's a timeout problem with NST & Thunderbird

I have a slow dialup and when I download 100 - 200 kb email it may take from 30 to 60 seconds. Everything worked fine, until I found that I have 1 message on the server, thunderbird downloads it, but there's no email in folder.

Looking at logs revealed this:

Aug 31, 2005, 17:03:17 Session 1: RETR 1

Aug 31, 2005, 17:03:17 Session 1: +OK 88940 bytes will follow

Aug 31, 2005, 17:03:18 Session 1: (Attachment Filter) From: "OKHA"

Aug 31, 2005, 17:03:18 Session 1: (Attachment Filter) To:

Aug 31, 2005, 17:03:53 Session 1: Received end of data, mail size 87kB

Aug 31, 2005, 17:03:53 Session 1: (Attachment Filter) Filter result is accept/deliver

Aug 31, 2005, 17:03:53 Session 1: (SpamAssassin) Executing: sa\spamassassin.exe -x -c "sa\ruleset" -e 255

Aug 31, 2005, 17:04:02 Session 1: Connection closed by client

Aug 31, 2005, 17:04:02 Session 1: Connection from 127.0.0.1 closed

Aug 31, 2005, 17:04:03 Statistics load: 1 rows 14 columns

Aug 31, 2005, 17:04:03 Statistics save: 1 rows 14 columns

And something like this (shorter sample):

Aug 31, 2005, 17:03:18 Session 1: (Attachment Filter) From: "OKHA"

Aug 31, 2005, 17:03:18 Session 1: (Attachment Filter) To:

Aug 31, 2005, 17:04:02 Session 1: Connection closed by client

Aug 31, 2005, 17:04:02 Session 1: Connection from 127.0.0.1 closed

Looked like thunderbird just timed out and closed connection (without saying a word to me!). Downloading through Outlook Express went fine. So i tweaked Thunderbird a bit and set "mail.pop3_response_timeout" param from 45s to 5min. Everything went fine then.

So. It's all nice and great, but. I can get messages which are more than 5 min to download -- meaning I will have to tinker timeout again. I think NST should handle this problem.

You can send every 30s-60s a blank line to email client, or send some headers to prevent timing out, while downloading big message or scanning it with spamassassin. POPFile uses "X-POPFile-TimeoutPrevention" header so some message headers start like this:

X-POPFile-TimeoutPrevention: 0

X-POPFile-TimeoutPrevention: 1

X-POPFile-TimeoutPrevention: 2

I hope you'll address this problem in NST, because other program are aware of it and handle it fine.

P.S. NST! V2.3.1.2

Post Edited (08-31-05 18:28)

Subject: Re: Timeout handling
Posted by [support](#) on Mon, 05 Sep 2005 12:54:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

Sending blank lines wouldn't work, because it would break the MIME format. Also,emitting X-headers as timeout prevention is likely to (mildly) confuse filters such as SpamAssassin (all the Received headers are supposed to come first in a MIME message).

But you are right, given the fact that popular mail clients have rather short timeouts (T-Bird 45s, Outlook 60s), it is better for most NoSpamToday! users to get bogus headers instead of having to modify the default timeouts depending on the speed of the download connection to the POP3 server.

I added your suggestion to our list of requested features, to be included in one of the next releases.

Subject: Re: Timeout handling
Posted by [oxygen](#) on Mon, 05 Sep 2005 14:34:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks for your response!

As for headers:

1) Just looked at the first recieved message header:

"From hosting@rbc.ru Mon Sep 05 16:24:26 2005

Return-path:

Received:"

So we have Return-path pror to Recieved, looks like normal practice here

2) As for bayes, I read somewhere 3 years ago that it filters out many default headers from analyzing, including headers which start with "X-Spam-..."

Use header "X-Spam-Timeout-Prevention"?

And, by the way, I found out that my POP3 server sometimes timeouts too :)
Maybe you should consider blank lines (or anything) for POP3 server to prevent timeout?

This reveals the next problem: sometimes spamassassin is SLOW. I don't know why, but for some messages it's a snap, for some it can be a minute or two. It doesn't depend on the message size (all messages ~10-30 kb). I'm not even sure it depends on the particular message or on the time of a day or on demographic situation in Africa, but still, sometimes I have to wait for spamassassin LONG to process the message. And I don't turn RBL on (plus dns config is set to no).

I have Centrino 1.6 Ghz, 512 Mb ram. I don't think hardware is an issue.

Post Edited (09-05-05 20:44)

Subject: Re: Timeout handling
Posted by [oxygen](#) on Mon, 05 Sep 2005 18:44:13 GMT
[View Forum Message](#) <> [Reply to Message](#)

Just a hint: maybe you can add additional log detail option to show all spamassassin debug output (-D) option, so if spamassassin processes mail message for too long, I can see if it is a particular problem (it stumbled upon some action) or it's a general unresponsivness or something else by looking at debug output.

Post Edited (09-05-05 20:44)

Subject: Re: Timeout handling
Posted by [oxygen](#) on Mon, 05 Sep 2005 20:08:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

Ok, here's some investigation concerning speed issues. I found out the problem -- RBL checks. But it's not as obvious as you may think.

With rbl checks ON (skip_rbl_checks 0) spamassassin does these tests:

```
-- 1 (reverse DNS) -----  
debug: is Net::DNS::Resolver available? yes  
debug: Net::DNS version: 0.49  
debug: looking up PTR record for '62.118.249.63'  
debug: PTR for '62.118.249.63': 'mox.rbc.ru'  
debug: received-header: parsed as [ ip=62.118.249.63 rdns=mox.rbc.ru helo= by=mx17.mail.ru  
ident= envfrom= intl=0 id=1EAiLr-000G7D-00 auth= ]  
-----
```

```
-- 2 (URIDNSBL check) -----
debug: Running tests for priority: 500
debug: URIDNSBL: queries completed: 0 started: 0
debug: URIDNSBL: queries active: DNSBL=1 NS=1 at Mon Sep 5 23:33:00 2005
debug: URIDNSBL: queries completed: 0 started: 0
..... 19 tests more .....
debug: URIDNSBL: queries active: DNSBL=1 NS=1 at Mon Sep 5 23:33:20 2005
debug: URIDNSBL: escaping: must have lost requests
debug: URIDNSBL: aborting remaining lookups
-----
```

```
-- 3 (RBL check) -----
debug: RBL: success for 0 of 17 queries
debug: DNS: timeout for njabl-notfirsthop,njabl after 30 seconds
debug: DNS: timeout for njabl after 30 seconds
debug: DNS: timeout for sorbs after 30 seconds
debug: DNS: timeout for sblxbl after 30 seconds
debug: DNS: timeout for sblxbl-notfirsthop,sblxbl after 30 seconds
debug: DNS: timeout for ahbl after 30 seconds
debug: DNS: timeout for NO_DNS_FOR_FROM after 30 seconds
debug: DNS: timeout for sorbs after 30 seconds
debug: DNS: timeout for rfc1_envfrom after 30 seconds
debug: DNS: timeout for NO_DNS_FOR_FROM after 30 seconds
debug: DNS: timeout for bsp-untrusted after 30 seconds
debug: DNS: timeout for rsl after 30 seconds
debug: DNS: timeout for dsbl-notfirsthop after 30 seconds
debug: DNS: timeout for rsl after 30 seconds
debug: DNS: timeout for spamcop after 30 seconds
debug: DNS: timeout for spamcop after 30 seconds
debug: DNS: timeout for bsp-firsttrusted after 30 seconds
-----
```

With rbl checks OFF (skip_rbl_checks 1) spamassassin does these tests:

```
-- 1 (reverse DNS) -----
debug: is Net::DNS::Resolver available? yes
debug: Net::DNS version: 0.49
debug: looking up PTR record for '217.132.229.79'
debug: PTR for '217.132.229.79': 'CBL217-132-229-79.bb.netvision.net.il'
debug: received-header: parsed as [ ip=217.132.229.79
rdns=CBL217-132-229-79.bb.netvision.net.il helo= by=mox.smtp.ru ident= envfrom= intl=0
id=40648324 auth= ]
debug: looking up A records for 'mox.smtp.ru'
debug: A records for 'mox.smtp.ru': 81.211.64.121
debug: looking up A records for 'mox.smtp.ru'
debug: A records for 'mox.smtp.ru': 81.211.64.121
-----
```

```
-- 2 (URIDNSBL check) -----
debug: Running tests for priority: 500
```

```
debug: URIDNSBL: queries completed: 0 started: 0
debug: URIDNSBL: queries active: DNSBL=1 NS=1 at Mon Sep 5 23:47:54 2005
debug: URIDNSBL: queries completed: 0 started: 0
..... 19 tests more .....
debug: URIDNSBL: queries active: DNSBL=1 NS=1 at Mon Sep 5 23:48:14 2005
debug: URIDNSBL: escaping: must have lost requests
debug: URIDNSBL: aborting remaining lookups
-----
-- 3 (RBL check??) -----
debug: RBL: success for 0 of 2 queries
-----
```

So it's URIDNSBL that eats much time when RBL is OFF! Here it looks like it eats only 20 seconds (from URIDNSBL log), but aborting lookups and other network stuff add up much more. So it's not unusual to have 1 min email check (even 3-5 min for connection problems). Then I set `dns_available no` and `spamassassin` flies! I got messages checked in 3-6 seconds!

The next question. What's about my DNS? I have this line in NST log:
SpamAssassin DNS server(s): 192.168.1.1 192.168.1.2 213.59.64.67

First two servers are my work's local DNS and the last is my home dialup DNS (I'm testing it under dialup right now). It looks like it tries to reach first servers first and then 30s RBL timeout hits. Can I reorder it somehow?

P.S. About methodology. I've downloaded `redir` program that can redirect `stderr` to a file and set up following command for NST:

```
CommandLine=redir -ea spamassassin.err sa\spamassassin.exe -D -x -c sa\ruleset -e 255
```

and when `spamassassin` checks the message, I can see the debug output almost in real time. I have changed `dns_available` and `skip_rbl_checks` variables in `local.cf` file.

P.S.S. What a long post! But I hope my findings can help you and other people which experience unpredictable "slowness"

Subject: Re: Timeout handling
Posted by [oxygen](#) on Mon, 05 Sep 2005 22:31:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

Ok, the last post.

With the help of command line I found the following:

- * If I set `SET RES_NAMESERVERS=192.168.1.1 192.168.1.2 213.59.64.67` (like NST detects it) SA checks messages for 40-60sec and all DNS (rbl & uridnsbl) checks fail
- * If I set `SET RES_NAMESERVERS=213.59.64.67` (my home dialup ISP's DNS) SA checks messages for 4-6sec and all DNS checks succeed

Conclusion: with proper configured SA's DNS settings there's almost no time penalty. Getting the right server order is vital.

Strange that even with the first config, SA is able to reverse lookup (PTR records), just like in the post above, but not able to use rbl

Subject: Re: Timeout handling
Posted by [support](#) on Tue, 06 Sep 2005 10:31:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

There is one other thing I want to add, for completeness' sake:

> So. It's all nice and great, but. I can get messages which are more than 5 min to
> download -- meaning I will have to tinker timeout again. I think NST should
> handle this problem.

If a large message download time exceeds 5minutes, it will probably be larger than 256kB, and exceed the default skip size setting. Once this happens, the message will be forwarded to the POP3 client directly.

So it will be enough if the timeout is long enough to accommodate a 256kB download. You need a slow modem connection to run into problems.

Another reason for a timeout can be Bayes database auto-expire. This has been discussed in other threads and in the FAQ.

Subject: Re: Timeout handling
Posted by [support](#) on Tue, 06 Sep 2005 10:50:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

> So it's URIDNSBL that eats much time when RBL is OFF! Here it
> looks like it eats only 20 seconds (from URIDNSBL log), but
> aborting lookups and other network stuff add up much more. So
> it's not unusual to have 1 min email check (even 3-5 min for
> connection problems). Then I set dns_available no and
> spamassassin flies! I got messages checked in 3-6 seconds!

Checking "skip RBL checks" in the config wizard automatically sets dns_available to "no", unchecking it sets it back to "yes" again. So you should be fine when using the config wizard to modify your local.cf file.

> The next question. What's about my DNS? I have this line in NST
> log:

> SpamAssassin DNS server(s): 192.168.1.1 192.168.1.2
> 213.59.64.67
>
> First two servers are my work's local DNS and the last is my
> home dialup DNS (I'm testing it under dialup right now). It
> looks like it tries to reach first servers first and then 30s
> RBL timeout hits. Can I reorder it somehow?

The SpamAssassin DNS servers are auto-detected from your system's network settings. Usually, these settings are fine, if they aren't, you should experience delays in other programs as well, such as your Internet browser.

You can statically set the DNS server(s) you want to use, overriding auto-detection. For this, add the following line to the [WorkstationSessionManager] section of the nospamtoday.cf config file:

```
DNSServer="213.59.64.67"
```

>
> P.S. About methodology. I've downloaded redir program that can
> redirect stderr to a file and set up following command for NST:
>
> CommandLine=redir -ea spamassassin.err sa\spa massassin.exe -D
> -x -c sa\ruleset -e 255
> and when spamassassin checks the message, I can see the debug
> output almost in real time. I have changed dns_available and
> skip_rbl_checks variables in local.cf file.

You can also make the stderr output appear in the log, by checking the "log filter error output" option in the config wizard, though not in real-time.

> P.S.S. What a long post! But I hope my findings can help you
> and other people which experience unpredictable "slowness"

Thank you for the time you devoted to this problem!