

poplib 100% cpu usage

Source: <http://coding.derkeiler.com/Archive/Python/comp.lang.python/2008-07/msg01714.html>

- *From:* Oli Schacher <python@xxxxxxxxxxxxxx>
 - *Date:* Wed, 16 Jul 2008 17:01:08 +0200
-

Hi all

I wrote a multithreaded script that polls mails from several pop/imap accounts. To fetch the messages I'm using the getmail classes (<http://pyropus.ca/software/getmail/>), those classes use the poplib for the real pop transaction.

When I run my script for a few hours cpu usage goes up to 100%, sometimes even 104% according to 'top' :-)
This made our test machine freeze once. First I thought I maybe didn't stop my threads correctly after polling an account but I attached a remote debugger and it showed that threads are stopped ok and that the cpu gets eaten in poplib in the function "_getline" which states in the description:

```
---snip---  
# Internal: return one line from the server, stripping CRLF.  
# This is where all the CPU time of this module is consumed.  
# Raise error_proto('-ERR EOF') if the connection is closed.  
  
def _getline(self):  
---snip---
```

So for testing purposes I changed this function and added:
time.sleep(0.0001)
(googling about similar problems with cpu usage yields this time.sleep() trick)

It now looks ok, cpu usage is at about 30% with a few spikes to 80-90%.

Of course I don't feel cozy about changing a standard library as the changes will be overwritten by python upgrades.

Did someone else from the list hit a similar problem and maybe has a better solution?

Thanks for your hints.

Best regards,
Oli Schacher

.