

# Re: The reliability of python threads

---

*Source:* <http://coding.derkeiler.com/Archive/Python/comp.lang.python/2007-01/msg03444.html>

---

- *From:* Steve Holden <[steve@xxxxxxxxxxxxxx](mailto:steve@xxxxxxxxxxxxxx)>
  - *Date:* Tue, 30 Jan 2007 05:14:33 +0000
- 

Carl J. Van Arsdall wrote:

Aahz wrote:

[snip]

My response is that you're asking the wrong questions here. Our database server locked up hard Sunday morning, and we still have no idea why (the machine itself, not just the database app). I think it's more important to focus on whether you have done all that is reasonable to make your application reliable — and then put your efforts into making your app recoverable.

Well, I assume that I have done all I can to make it reliable. This list is usually my last resort, or a place where I come hoping to find ideas that aren't coming to me naturally. The only other thing I thought to come up with was that there might be network errors. But i've gone back and forth on that, because TCP should handle that for me and I shouldn't have to deal with it directly in pyro, although I've added (and continue to add) checks in places that appear appropriate (and in some cases, checks because I prefer to be paranoid about errors).

I'm particularly making this comment in the context of your later point about the bug showing up only every three or four months.

Side note: without knowing what error messages you're getting, there's not much anybody can say about your programs or the reliability of threads for your application.

Right, I wasn't coming here to get someone to debug my app, I'm just looking for ideas. I constantly am trying to find new ways to improve my software and new ways to reduce bugs, and when i get really stuck, new ways to track bugs down. The exception won't mean much, but I can say that the error appears to me as bad data. I do checks prior to performing actions on any data, if the data doesn't look like what it should look like, then the system flags an exception.

The problem I'm having is determining how the data went bad. In tracking down the problem a couple guys mentioned that problems like that usually are a race condition. From here I examined my code, checked out all the locking stuff, made sure it was good, and wasn't able

## Re: The reliability of python threads

to find anything. Being that there's one lock and the critical sections are well defined, I'm having difficulty. One idea I have to try and get a better understanding might be to check data before its stored. Again, I still don't know how it would get messed up nor can I reproduce the error on my own.

Do any of you think that would be a good practice for trying to track this down? (Check the data after reading it, check the data before saving it)

Are you using memory with built-in error detection and correction?

regards

Steve

--

Steve Holden +44 150 684 7255 +1 800 494 3119

Holden Web LLC/Ltd <http://www.holdenweb.com>

Skype: holdenweb <http://del.icio.us/steve.holden>

Blog of Note: <http://holdenweb.blogspot.com>

See you at PyCon? <http://us.pycon.org/TX2007>

.