

RE: How should I connect via DBD::Oracle to efficiently obtain 2000+ simultaneous connections?

RE: How should I connect via DBD::Oracle to efficiently obtain 2000+ simultaneous connections?

Source: <http://coding.derkeiler.com/Archive/Perl/perl.dbi.users/2006-11/msg00111.html>

- *From:* Philip.Garrett@xxxxxxxxxxx (Philip Garrett)
 - *Date:* Sun, 19 Nov 2006 00:22:14 -0500
-

It sounds to me like you have a good case for using some type of pooling.

Here are a few options off the top of my head:

- 1) Pooling via SQLRelay
- 2) Pooling via MTS
- 3) Use fewer mod_perl processes by putting a lightweight reverse proxy in front of them. This doesn't pool connections, just reduces the number of connections required. See Stas Bekman's overview at <http://tinyurl.com/y5k8ka> (perl.apache.org) if you haven't already.

You might also want to consider aggressive caching (memcached perhaps?) to take some load off of your database server.

Regards,
Philip

-----Original Message-----

From: Chris [<mailto:dispo41@xxxxxxxxxxx>]
Sent: Friday, November 17, 2006 9:02 PM
To: dbi-users@xxxxxxxx
Subject: Re: How should I connect via DBD::Oracle to efficiently obtain 2000+ simultaneous connections?

Doing some digging around, I discover that each Oracle connection consumes
4megs of RAM when "doing nothing".

that would add up to 8gigs for 2000 connections... and since I've only got
8gigs in total, this is about 4 times more than I can spare.

It would seem that I will definitely need some kind of connection pooling –
does anyone disagree ?

RE: How should I connect via DBD::Oracle to efficiently obtain 2000+ simultaneous connections? 1

RE: How should I connect via DBD::Oracle to efficiently obtain 2000+ simultaneous connections?

```
# pmap -d 13040 | tail -1  
mapped: 810872K writeable/private: 3392K shared: 722944K
```

.