

pitertools - Feature #5306

Reimplementation of queue

05/13/2014 12:58 PM - Joshua A. Drake

Status:	Feedback	Start date:	05/13/2014
Priority:	Normal	Due date:	
Assignee:	Joshua Drake	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Resolution:			
Description			
Reimplementation of PITRTools queue The current queue system possesses the following limitations: 1. No job control - If we are generating xlogs faster than we can ship them, rsync processes will be spawned uncontrollably. 2. A continuous stream of transaction logs must be generated for correct operation of the archiver. If there are no new logs coming in, logs still in the queues will not get shipped. Suggested solution: Queue management daemon - Have cmd_archiver only queue logs, with a separate daemon shipping them. * cmd_queue daemon - write PID to a file - In worker thread] Ship logs in queue with subprocess calls when self to-do list is not empty] If successful, remove paths from self to-do list and log appropriate output] Else, log which queues failed] Wait a few seconds [configurable]] Repeat - In main thread] Wait for data [absolute paths to queued logs] on a pipe or UNIX socket] On data, parse and extract paths, append paths to self to-do list] Repeat * When cmd_archiver is called with a new archive, copy the archive to slave queues and write paths to cmd_queue via pipe/socket			

History

#1 - 05/16/2014 12:10 PM - Joshua A. Drake

Updated, simplified plan:

```
* cmd_queue daemon
- Check if any queues have logs in them
- Send the logs with calls from the subprocess module
  ] If a log ship was successful, delete the link and its source in the global directory. (This can possibly
be done with some fancy rsync flags - not sure yet)
  ] Else, continue. We'll try again on the next cycle.
- Wait a few seconds [configurable]
* When cmd_archiver is called with a new archive, put archive in global queue and hard link to slave queues as
per #5296
```

#2 - 11/12/2014 07:46 AM - Alexander Shulgin

- Status changed from New to In Progress
- Assignee set to Alexander Shulgin

The latest code for this feature is pushed to: <https://github.com/commandprompt/PITRTools/tree/5306>

It would be nice if someone could have a look on it.

#3 - 11/12/2014 07:48 AM - Alexander Shulgin

Alexander Shulgin wrote:

The latest code for this feature is pushed to: <https://github.com/commandprompt/PITRTools/tree/5306>

It would be nice if someone could have a look on it.

Summary of the changes:

- cmd_queue will be started by cmd_archiver (though a user invocation is not prohibited specifically)
- cmd_queue will allow only one instance running per l_archive dir, by checking for the \$l_archive/cmd_queue.pid file then bailing out if it is present **and** there is a process running with that PID¹
- this allows us to run cmd_queue from cmd_archiver for every new WAL file and let it figure the status itself; this also makes sure that manually starting or stopping cmd_queue doesn't interfere with the cmd_archiver plans
- optionally, cmd_queue can check for \$PGDATA/postmaster.pid after each iteration and bail out if it has gone missing; this should make sure cmd_queue lifetime is tied to the postmaster's one, but that's not strictly necessary
- cmd_queue will accept --daemon option to be used from cmd_archiver; this makes sure the archiver won't get stuck when running cmd_queue for the first time; all further invocations are expected to exit immediately due to an already running instance of cmd_queue

It is currently running through slaves in order. I'm looking to see if adding some parallelism is an option.

#4 - 11/12/2014 08:37 AM - Joshua Drake

On 11/12/2014 07:48 AM, pitertools-tickets@lists.commandprompt.com wrote:

- cmd_queue will accept --daemon option to be used from cmd_archiver; this makes sure the archiver won't get stuck when running cmd_queue for the first time; all further invocations are expected to exit immediately due to an already running instance of cmd_queue

It is currently running through slaves in order. I'm looking to see if adding some parallelism is an option.

That would be awesome (the parallelism).

--

Command Prompt, Inc. - <http://www.commandprompt.com/> 503-667-4564
PostgreSQL Support, Training, Professional Services and Development
High Availability, Oracle Conversion, @cmdpromptinc
"If we send our children to Caesar for their education, we should not be surprised when they come back as Romans."

#5 - 11/12/2014 09:30 AM - Alexander Shulgin

- Status changed from In Progress to Feedback

- Assignee changed from Alexander Shulgin to Joshua Drake

Joshua Drake wrote:

That would be awesome (the parallelism).

Here you go (look ma, no threads!): <https://github.com/commandprompt/PITRTools/commit/eff015884a46950900bbf42b8ab365861973151b>

It would be nice though if someone other than me grabbed this latest version and tried running it independently.

--

Alex

#6 - 11/13/2014 10:12 AM - Alexander Shulgin

Just updated the docs and pushed to master.