

JobSub - Feature #7305

jobsub_client retry logic

11/10/2014 01:59 PM - Joe Boyd

Status:	New	Start date:	11/10/2014
Priority:	Normal	Due date:	
Assignee:	Dennis Box	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	v1.3.3	Spent time:	0.00 hour
Stakeholders:	FIFE		
Description			
It would be nice if job submission fails that the jobsub_client would retry X times (maybe try every machine in the round robin at least once???). Especially with our HA setup, the first machine tried could be in some bad state and you'd want to then submit the job to the next one.			

History

#1 - 12/19/2014 12:46 PM - Parag Mhashilkar

- Assignee set to Parag Mhashilkar
- Target version set to v1.1.1

#2 - 02/27/2015 10:25 AM - Parag Mhashilkar

- Target version changed from v1.1.1 to v1.2

#3 - 07/07/2015 10:59 AM - Parag Mhashilkar

- Assignee changed from Parag Mhashilkar to Dennis Box
- Stakeholders updated

NOTE: We should brainstorm before implementing this.

We want to be very careful not to resubmit in case the jobs were submitted successfully but for some reason the client did not get the info back from the server.

Here is my proposal and we should look for better alternatives if any

- Submit re-try logic
 - Figure the server to use by our usual method
 - Generate a unique token that is added to the jobs classad when you do the submit
 - If the client gets a failure code from the server, for every jobsubserver/schedd in the system
 - Wait for few sec
 - Check if there were jobs submitted with that unique token in the classad
 - If no jobs were submitted, move to next schedd in the list and submit jobs
 - Add one liner to the client's output clear mentioning the server/schedd that failed and the server/schedd that will be tried next
- Fetchlog retry logic
 - Wait and retry same server after some time. Exit with failure on second failed attempt.
- Other commands re-try logic
 - Unless I am missing something, other commands can be executed from either servers

#4 - 07/22/2015 03:59 PM - Joe Boyd

Running into a failure case right now I just thought I'd add here as food for thought.

We have three schedd's on pre-prod. fermicloud393 is currently broken in some way that I haven't investigated but my submissions there are failing every time with

User authorization has failed: Error authorizing DN='DC=gov/DC=fnal/O=Fermilab/OU=People/CN=Joe B. Boyd/CN=UID:boyd' for AcctGroup='mu2e'

even though I just sent an exact same job to fermicloud391.fnal.gov. Anyway, broken schedd.

The problem is that since I already submitted one job to fermicloud391 and fermicloud383 each, my submit now every time is going to fermicloud393 I assume because of the load balancing stuff that was added. So, we need retry logic if there's a failure, but it can't just keep trying the same schedd

because of the load balancing stuff either because that schedd may be broken. In this case since there is something wrong with auth the job will never submit even with retrys.

#5 - 02/02/2016 06:57 AM - Dennis Box

- Target version changed from v1.2 to v1.2.1

#6 - 03/10/2016 03:43 PM - Dennis Box

- Target version changed from v1.2.1 to v1.2.2

#7 - 07/13/2016 06:26 PM - Dennis Box

- Target version changed from v1.2.2 to v1.2.3

#8 - 09/13/2016 09:26 AM - Dennis Box

- Target version changed from v1.2.3 to v1.2.4

#9 - 05/09/2017 10:00 AM - Dennis Box

- Target version changed from v1.2.4 to v1.2.5

#10 - 10/30/2017 11:02 AM - Dennis Box

- Target version changed from v1.2.5 to v1.2.7

Related to, if not the same issue as [#15257](#)

#11 - 05/08/2018 09:41 AM - Dennis Box

- Target version changed from v1.2.7 to v1.2.8

#12 - 09/18/2018 03:39 PM - Dennis Box

- Target version changed from v1.2.8 to v1.2.9

#13 - 11/20/2018 08:24 AM - Dennis Box

- Target version changed from v1.2.9 to v1.3

#14 - 05/17/2019 01:48 PM - Dennis Box

- Target version changed from v1.3 to v1.3.3