

Kepler - Bug #2895

Distributed Execution Tracking Bug

07/23/2007 10:31 AM - Chad Berkley

Status:	New	Start date:	07/23/2007
Priority:	Immediate	Due date:	
Assignee:	Chad Berkley	% Done:	0%
Category:	core	Estimated time:	0.00 hour
Target version:	3.X.Y	Spent time:	0.00 hour
Bugzilla-Id:	2895		

Description

Chad and Lucas are developing the distributed execution system for kepler. The system is currently working in a very simplified way. This bug is a consolidation of bug 1891 and bug 1899.

The following items need to be added:

- Make sure that the JNI libraries can be accessed via the slave and that the ENM actors will work on the slave
- we might have to solve the problem that kepler has where you can't run multiple instances of the application with the same user account. The problem is that the cache uses an embedded database which only allows one connection at a time. the db is stored in the .kepler directory so if you try to run kepler twice at the same time, you'll get an error on the 2nd one that the db is already in use. If we have a cluster where the slave is distributed via a single home directory, this will be a problem
- Matt came up with the idea of using the ecogrid registry as a way of doing node discovery.
- Get this to run on the NCEAS ROCKS cluster.
- we need to deal with transferring support files to the slave(s). This includes doing the indirect transfers between slaves (instead of transferring results back to the master then to the next slave, the slaves should be able to transfer data between each other).

History

#1 - 03/27/2013 02:21 PM - Redmine Admin

Original Bugzilla ID was 2895