

## Arvados - Feature #17756

### Initial implementation of LSF dispatcher

06/04/2021 01:58 PM - Peter Amstutz

<b>Status:</b> Resolved	<b>Start date:</b> 07/14/2021
<b>Priority:</b> Normal	<b>Due date:</b>
<b>Assigned To:</b> Tom Clegg	<b>% Done:</b> 100%
<b>Category:</b> Crunch	<b>Estimated time:</b> 0.00 hour
<b>Target version:</b> 2021-08-04 sprint	
<b>Description</b> Similar to crunch-dispatch-slurm. Should be packaged as a subcommand of arvados-server, though.  Use Arvados SDK "Dispatcher" class to interact with API server.  Start with basic dispatch to queue (bsub) and monitoring.  Stub out commands to support unit testing (as done with slurm).  Test on 9tee4	
<b>Subtasks:</b> Task # 17795: Review 17756-dispatch-lsf <span style="float: right;"><b>Resolved</b></span>	
<b>Related issues:</b> Related to Arvados Epics - Story #16304: LSF support <span style="float: right;"><b>In Progress</b> 04/01/2021 09/30/2021</span>	

#### Associated revisions

##### Revision 5d56a1af - 07/29/2021 01:57 PM - Tom Clegg

Merge branch '17756-dispatch-lsf' into main

closes #17756

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <[tom@curii.com](mailto:tom@curii.com)>

#### History

##### #1 - 06/04/2021 01:58 PM - Peter Amstutz

- Category set to Crunch

##### #2 - 06/08/2021 05:35 PM - Tom Clegg

- Description updated

##### #3 - 06/08/2021 05:41 PM - Peter Amstutz

- Description updated

##### #4 - 06/08/2021 05:42 PM - Peter Amstutz

- Description updated

##### #5 - 06/09/2021 04:09 PM - Peter Amstutz

- Assigned To set to Tom Clegg

##### #7 - 06/22/2021 05:28 PM - Peter Amstutz

- Target version changed from 2021-06-23 sprint to 2021-07-07 sprint

##### #8 - 07/02/2021 05:39 PM - Tom Clegg

- Status changed from New to In Progress

testing [c1ca6ec2e](#) on 9tee4:

- add DispatchLSF: InternalURLs: "http://0.0.0:9009": true to Services section in /etc/arvados/config.yml
- add Containers: RuntimeEngine: singularity in /etc/arvados/config.yml
- systemctl stop crunch-dispatch-slurm
- install ...path/to/arvados-server /usr/local/bin/arvados-dispatch-lsf
- arvados-dispatch-lsf
- submit a container request

```

{"Listen":["[:]:9009"],"PID":29877,"Service":"arvados-dispatch-lsf","URL":"http://0.0.0:9009/","level":"info",
,"msg":"listening","time":"2021-07-02T17:17:50.230755112Z"}
{"PID":29877,"level":"warning","msg":"FIXME: checkLsfQueueForOrphans","time":"2021-07-02T17:17:50.230759598Z"}
{"PID":29877,"level":"info","msg":"Submitting container 9tee4-dz642-unjdidp7zb08qdm
to LSF","time":"2021-07-02T17:17:56.390746354Z"}
{"PID":29877,"level":"info","msg":"bsub command [\"sudo\" \"-E\" \"-u\" \"crunch\" \"bsub\" \"-J\" \"
9tee4-dz642-unjdidp7zb08qdm
\" \"-R\" \"rusage[mem=757MB:tmp=640MB] affinity[core(1)]\"] script \"#!/bin/sh\\nexec 'crunch-run' '--runtime
-engine=singularity' '-cgroup-parent-subsystem=memory' '9tee4-dz642-unjdidp7zb08qdm
'\\n\"","time":"2021-07-02T17:17:56.390943355Z"}
{"PID":29877,"level":"info","msg":"bsub finished","stdout":"Job \u003c190\u003e is submitted to default queue
\u003cnormal\u003e.\n","time":"2021-07-02T17:17:56.426583960Z"}
{"PID":29877,"level":"info","msg":"Start monitoring container 9tee4-dz642-unjdidp7zb08qdm
in state \"Locked\"","time":"2021-07-02T17:17:56.426686121Z"}
{"PID":29877,"level":"info","msg":"container 9tee4-dz642-unjdidp7zb08qdm
changed state from Locked to Running","time":"2021-07-02T17:19:27.781201740Z"}
{"PID":29877,"level":"info","msg":"container 9tee4-dz642-unjdidp7zb08qdm
is done","time":"2021-07-02T17:19:28.828679722Z"}
{"PID":29877,"level":"info","msg":"Bkill(190) ","time":"2021-07-02T17:19:29.389316198Z"}
{"PID":29877,"level":"info","msg":"container 9tee4-dz642-unjdidp7zb08qdm
job disappeared from LSF queue","time":"2021-07-02T17:19:30.389644004Z"}
{"PID":29877,"level":"info","msg":"Done monitoring container 9tee4-dz642-unjdidp7zb08qdm
","time":"2021-07-02T17:19:33.836542040Z"}

```

todo/todo:

- dispatching to docker doesn't work, even though docker-via-slurm works on the same compute nodes
  - 2021-07-02T15:15:27.602617730Z error in Run: While loading container image: While loading container image into Docker: Got permission denied while trying to connect to the Docker daemon socket at unix:///var/run/docker.sock: Post "http://%2Fvar%2Frun%2Fdocker.sock/v1.21/images/load?quiet=1": dial unix /var/run/docker.sock: connect: permission denied"
  - my guess is bsub sets up the user's primary group (or the groups that user has on the submitting node), which wouldn't include docker in this case. I tried adding groups with "bsub -G docker" but that group doesn't exist on the submitting node, so: "Bad user group name. Job not submitted."
  - is this a blocker, or should we just document for now that lsf+docker isn't yet supported?
- implement checkLsfQueueForOrphans (see checkQueueForOrphans)
- propagate arvados container priority to lsf job priority (strategy might depend on lsf config)
- add doc page
- add deb/rpm package

## #9 - 07/05/2021 08:58 PM - Tom Clegg

17756-dispatch-lsf @ [8145f8e8e6ab99184fcd41dea042ede63e9ff0d5](#) -- <https://ci.arvados.org/view/Developer/job/developer-run-tests/2566/>

todo:

- update docs: add DispatchLSF: InternalURLs: "http://0.0.0:9009": true to Services section in /etc/arvados/config.yml
- priority?

## #10 - 07/07/2021 03:17 PM - Tom Clegg

- Target version changed from 2021-07-07 sprint to 2021-07-21 sprint

## #11 - 07/16/2021 07:03 PM - Nico César

review @ [3b63632698de9868a501191e8989f14c23e4e743](#)

I think documentation for BsubArgumentsList should take the form of array [] instead of a list. This is correct in config.yml documentation. An even better option for sysadmins will be a string as in "-C 0.." since is likely they will copy and paste from some working integrations. Internally can be transformed to an array.

How does a non existant BsubSudoUser in compute node error comes up? I'm thinking of a newly added misconfigured compute node in an already working cluster. also will be good if there is a good way to check that a compute node is properly configured (maybe arvados-client diagnostics should have this chec?)

func (disp \*dispatcher) runContainer() seems slightly convoluted the first block with a bunch of nested if's and else's and error handling is unclear to me. I like the x, err := foo ; if (err != nil) { return... } style because it gives a "line of sight" in the same column number all errors to be returned instead the 7 tabs/spaces while nested.

in runContainer, what will be the case that ok false? shouldn't we log this?

```
case updated, ok := <-status:
  if !ok {
    done = true
    break
  }
```

Also I noticed that we check if Priority 0 {} else {} ; can Priority be negative at any point? I remember sometime we had issues with slurm and some negative priority cases, but that could all be something from the past.

**#12 - 07/20/2021 05:19 PM - Peter Amstutz**

- Related to Story #16304: LSF support added

**#13 - 07/21/2021 03:10 PM - Peter Amstutz**

- Target version changed from 2021-07-21 sprint to 2021-08-04 sprint

**#14 - 07/22/2021 03:20 PM - Tom Clegg**

17756-dispatch-lsf @ [efc1846a758929bdb57b87bdbb3f757f8907c69b](https://ci.arvados.org/view/Developer/job/developer-run-tests/2602/) -- <https://ci.arvados.org/view/Developer/job/developer-run-tests/2602/>

- Change example yaml formatting from list to array: BsubArgumentsList: ["-C", "0"]
- Nested if's and else's addressed by moving the two copies of the "no suitable instance type" reporting code from c-d-slurm/lfsf into sdk/go/dispatch.
- Check Priority<1 instead of ==0 (just in case it happens, although it shouldn't).

I tried "sudo -E -u postgres bsub echo ok" on a test cluster and LSF seemed to just leave that job in PEND state (maybe waiting for a compute node to appear that has that user?). I don't see any complaints about it in the LSF logs. Not sure how hard we should try to help troubleshoot this, but at least I added a note to the config file comment & install doc page.

**#15 - 07/22/2021 06:46 PM - Nico César**

review @ [efc1846a758929bdb57b87bdbb3f757f8907c69b](https://ci.arvados.org/view/Developer/job/developer-run-tests/2602/)

LGTM ready to merge

**#16 - 07/29/2021 02:07 PM - Tom Clegg**

- % Done changed from 0 to 100

- Status changed from In Progress to Resolved

Applied in changeset [arvados|5d56a1af42f64df57ef7a1bcef6d016ff2310900](https://ci.arvados.org/view/Developer/job/developer-run-tests/2602/).