

## Arvados - Story #8543

### [NodeManager] Don't use Futures when not expecting a reply

02/24/2016 09:35 PM - Peter Amstutz

|   |                                  |
|---|----------------------------------|
| <b>Status:</b> Resolved   | <b>Start date:</b> 03/04/2016    |
| <b>Priority:</b> Normal   | <b>Due date:</b>                 |
| <b>Assigned To:</b> Peter Amstutz   | <b>% Done:</b> 100%              |
| <b>Category:</b> Node Manager   | <b>Estimated time:</b> 0.00 hour |
| <b>Target version:</b> 2016-03-16 sprint  |                                  |
| <b>Description</b><br>Quoting <a href="#">#8437</a><br><br>on_failure is only called when there is no future associated with the message. As it turns out, all calls that use ActorProxy have an associated Future object, and all messaging between actors in node manager uses ActorProxy. This means unhandled exceptions are stored in a Future object to be returned to the caller. However, if the caller never calls get() on the Future object (because it never stored it), this means the exception is silently ignored.<br><br>These lingering future objects may also be creating circular references that is causing the memory leak.<br><br>Anywhere a message is sent between Actors where a response is not required, such as subscriptions, use tell() on the ActorRef object instead of using ActorProxy (which uses ask()). (May want to write a helper similar to ActorProxy, such as TellActorProxy).<br><br>Anticipated benefits: <ul style="list-style-type: none"><li>• Unhandled exceptions will be logged, and the on_failure method of the Actor be called as intended</li><li>• Possibly uncover previously unknown bugs (previously hidden due to exceptions being lost)</li><li>• May mitigate longstanding memory leak bug by avoiding creating Future objects that are not needed (and which may be actively harmful)</li></ul> |                                  |
| <b>Subtasks:</b><br>Task # 8604: Use TellActorProxy when future is not going to be used <b>Resolved</b><br>Task # 8602: Review 8543-nodemanager-fewer-futures <b>Resolved</b><br>Task # 8605: Fix tests <b>Resolved</b><br>Task # 8601: Implement TellActorProxy <b>Resolved</b>  |                                  |
| <b>Related issues:</b><br>Related to Arvados - Bug #8541: [NodeManager] Use sys.exc_clear() to release ... <b>Closed</b><br>Related to Arvados - Bug #7026: [Node Manager] Mishandles stop signals <b>Closed</b> 08/19/2015   |                                  |

#### Associated revisions

##### Revision 2dbbaef - 03/08/2016 08:08 PM - Peter Amstutz

Merge branch '8543-nodemanager-fewer-futures' closes #8543

##### Revision d1de3281 - 03/10/2016 08:54 PM - Peter Amstutz

Fix node manager to send ActorProxy instead of TellActorProxy to ComputeNodeStateChangeBase subscribers. refs #8543

#### History

##### #1 - 02/24/2016 09:42 PM - Peter Amstutz

- Description updated

##### #2 - 02/25/2016 02:29 AM - Peter Amstutz

- Description updated

- Category set to Node Manager

##### #3 - 02/29/2016 03:43 PM - Brett Smith

- Target version set to Arvados Future Sprints

**#4 - 03/01/2016 07:02 PM - Brett Smith**

- Target version changed from Arvados Future Sprints to 2016-03-16 sprint

**#5 - 03/01/2016 07:37 PM - Brett Smith**

- Story points set to 1.0

**#6 - 03/02/2016 08:21 PM - Peter Amstutz**

- Assigned To set to Peter Amstutz

**#7 - 03/07/2016 08:36 PM - Tom Clegg**

Looking at 8543-nodemanager-fewer-futures @ [bf49d05](#)...

Could you squash this branch and synthesize all the commit messages? I'd like to have less noise in the git history, like commit logs that just say "tests pass", and these changes that got made and then reverted...

- ```
self.subscribers.clear()
```
- ```
if key in self.sizes_booting_shutdown:
    del self.sizes_booting_shutdown[key]
```

This error message doesn't scan. I assume the "has no attribute" part is a copy/paste error?

```
raise AttributeError('%s not a callable on %s has no attribute' % (name, self))
```

BaseNodeMangerActor should probably be BaseNodeManagerActor

I think it would be worth providing a little more documentation beyond "Like ActorProxy, except only permits asynchronous stuff". Even coming fresh from this story description I find the workings of the "Tell" stuff a bit mysterious. Perhaps it would help to stick a docstring on tell\_proxy() saying when one should use that instead of proxy()?

This is a complaint about existing code, not the new code, but I can't help thinking it should be easy/worthwhile to log exception\_type and/or exception\_value here...

- ```
"Unhandled exception is a fatal error, killing Node Manager"
```

**#8 - 03/08/2016 04:34 PM - Peter Amstutz**

Tom Clegg wrote:

Looking at 8543-nodemanager-fewer-futures @ [bf49d05](#)...

Could you squash this branch and synthesize all the commit messages? I'd like to have less noise in the git history, like commit logs that just say "tests pass", and these changes that got made and then reverted...

- [...]
- [...]

Rebased and squashed.

This error message doesn't scan. I assume the "has no attribute" part is a copy/paste error?

[...]

Yes I think so.

BaseNodeMangerActor should probably be BaseNodeManagerActor

Thanks for catching that.

I think it would be worth providing a little more documentation beyond "Like ActorProxy, except only permits asynchronous stuff". Even coming fresh from this story description I find the workings of the "Tell" stuff a bit mysterious. Perhaps it would help to stick a docstring on tell\_proxy() saying when one should use that instead of proxy()?

I recall now being interrupted in the middle of starting to write that comment and obviously I forgot to come back to it. Expanded docstrings for those classes.

This is a complaint about existing code, not the new code, but I can't help thinking it should be easy/worthwhile to log `exception_type` and/or `exception_value` here...

- [...]

The exception already gets logged by pykka. It doesn't show up in the tests because the exception gets logged as "ERROR" but the tests by default only log "CRITICAL". Try `ANMTEST_LOGLEVEL=ERROR python setup.py test --test-suite=tests.test_failure` to see for yourself. In production we currently run at log level INFO or DEBUG.

**#9 - 03/08/2016 05:41 PM - Tom Clegg**

LGTM @ [054f950](#), thanks

**#10 - 03/08/2016 08:15 PM - Peter Amstutz**

- *Status changed from New to Resolved*

Applied in changeset `arvados|commit:2dbbaefc6a4a46a7f17b9e7799fc455cd722113`.