

Arvados - Bug #4751

[Node Manager] Can erroneously pair cloud nodes with stale Arvados node records

12/09/2014 06:07 PM - Brett Smith

Status:	Resolved	Start date:	03/02/2015
Priority:	Normal	Due date:	
Assigned To:	Brett Smith	% Done:	100%
Category:	Node Manager	Estimated time:	0.00 hour
Target version:	2015-03-11 sprint		
Description			
Node Manager pairs cloud nodes with Arvados node records based solely on an IP address match. See <code>arvnodemanager.computenode.dispatch.ComputeNodeMonitorActor.offer_arvados_pair</code> .			
It can happen that a cloud node comes up with an IP address that happens to match a stale Arvados node record. Make the testing stricter so there's no pairing in this case.			
Subtasks:			
Task # 5351: Review 4751-node-manager-stricter-node-pairing-wip			Resolved
Related issues:			
Related to Arvados - Support #5251: [Support] Fix bugs and write tests (first...		Resolved	
Has duplicate Arvados - Bug #4891: [Node Manager] Should not associate node w...		Closed	
Has duplicate Arvados - Bug #5292: [Node Manager] Failed to recognize busy no...		Closed	02/23/2015
Copied from Arvados - Story #4293: [Node Manager] Write off cloud nodes that ...		Resolved	10/27/2014

Associated revisions

Revision 6be95f5c - 03/02/2015 07:08 PM - Brett Smith

Merge branch '4751-node-manager-stricter-node-pairing-wip'

Closes #4751, #5351.

History

#1 - 12/09/2014 06:11 PM - Brett Smith

- Description updated

#2 - 12/09/2014 06:22 PM - Tom Clegg

- Story points set to 0.5

#3 - 12/09/2014 06:22 PM - Tom Clegg

- Target version changed from Bug Triage to Arvados Future Sprints

#4 - 01/23/2015 04:17 PM - Brett Smith

I think there are basically two possible approaches:

- EC2 compute nodes, at least, put their EC2 id in the Arvados node record's info. If we check against that, we can't go wrong—but it has the downside of meaning we have to reimplement this check for every cloud driver.
- Check that the Arvados node's `first_ping_at` is greater than the cloud node's boot time before accepting a pairing. This is completely generic, and very safe, although it could still go wrong if the total garbage data is getting into the node records.

I think I prefer #2, but I wanted to note the alternatives at least.

#5 - 03/02/2015 03:46 PM - Brett Smith

- Target version changed from Arvados Future Sprints to 2015-03-11 sprint

Moving to current sprint because it came up again during science support, and it's likely to become more pressing now that we've increased our `max_nodes` setting.

#6 - 03/02/2015 06:24 PM - Peter Amstutz

I feel like this came up in an earlier code review (discussing the pitfalls of reusing computed node records generally) so it's good to tighten up the check.

4751-node-manager-stricter-node-pairing-wip LGTM

#7 - 03/02/2015 07:15 PM - Brett Smith

- *Status changed from New to Resolved*

- *% Done changed from 0 to 100*

Applied in changeset arvados|commit:6be95f5c3a2fcbe6321bba52c20393060e33e637.

#8 - 03/03/2015 04:40 PM - Brett Smith

- *Assigned To set to Brett Smith*