

Arvados - Feature #14922

[crunch-dispatch-cloud] Run multiple containers concurrently on a single VM

03/05/2019 10:02 PM - Tom Clegg

Status:	New	Start date:	
Priority:	Normal	Due date:	
Assigned To:		% Done:	0%
Category:	Crunch	Estimated time:	0.00 hour
Target version:	Arvados Future Sprints		
Description			
Run a new container on an already-occupied VM (instead of using an idle VM or creating a new one) when the following conditions apply:			
<ul style="list-style-type: none">• the occupied VM is the same price as the instance type that would normally be chosen for the new container• the occupied VM has enough unallocated RAM, scratch, and VCPUs to accommodate the new container• either all containers on the VM allocate >0 VCPUs, or the instance will still have non-zero unallocated space after adding the new container (this ensures that an N-VCPU container does not share an N-VCPU instance with anything, even 0-VCPU containers)• the occupied VM has IdleBehavior=run (not hold or drain)			
If multiple occupied VMs satisfy these criteria, choose the one that has the most containers already running, or the most RAM already occupied. This will tend to drain the pool of shared VMs rather than keeping many underutilized VMs alive after a busy period subsides to a less-busy period.			
Typically, "same price as a dedicated node, but has spare capacity" will only happen with the cheapest instance type, but it might also apply to larger sizes if the menu has big size steps. Either way, this rule avoids the risk of wasting money by scheduling small long-running containers onto big nodes. In future, this rule may be configurable (to accommodate workloads that benefit more by sharing than they lose by underusing nodes).			
Typically, the smallest node type has 1 VCPU, so this feature is useful only if container requests and containers can either			
<ul style="list-style-type: none">• specify a minimum of zero CPUs, or• specify a fractional number of CPUs.			
...so ensure at least one of those is possible.			
Related issues:			
Related to Arvados - Feature #15370: [arvados-dispatch-cloud] loopback driver			New

History

#1 - 03/06/2019 06:00 PM - Tom Clegg

- Description updated

#2 - 03/06/2019 06:43 PM - Tom Morris

- Story points set to 2.0

#3 - 03/06/2019 06:43 PM - Tom Morris

- Target version changed from To Be Groomed to Arvados Future Sprints

#4 - 06/14/2019 09:02 PM - Tom Clegg

- Related to Feature #15370: [arvados-dispatch-cloud] loopback driver added

#5 - 10/02/2019 03:27 PM - Tom Clegg

- Subject changed from [crunch-dispatch-cloud] Run multiple containers on a single VM to [crunch-dispatch-cloud] Run multiple containers concurrently on a single VM