

Debugging K8s pod throughput with Calico CNI

Thomas Hartland

User report: “Cluster is slow to create pods”

Me:




Clusterloader2

- Kubernetes benchmark tool ¹.
- In this case, using it to
 - Create many pods in the cluster at once (N pods per node).
 - Measure the pod startup latencies.

¹<https://github.com/kubernetes/perf-tests/tree/master/clusterloader2>

Pod startup latency



Measuring against

- Kubernetes pod startup latency SLI:

Startup latency of schedulable stateless pods, excluding time to pull images and run init containers, measured from pod creation timestamp to when all its containers are reported as started and observed via watch, measured as 99th percentile over last 5 minutes. ²

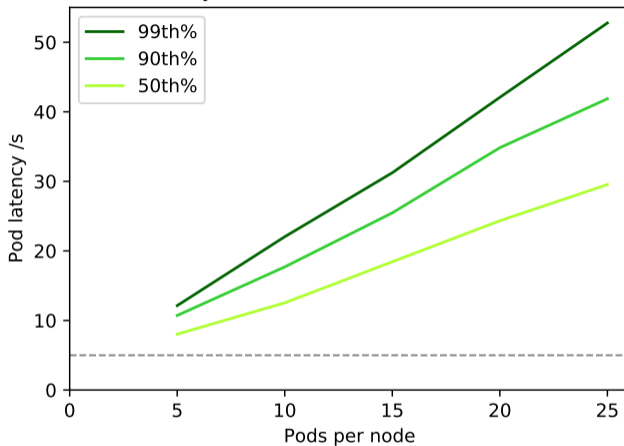
- The target: < 5 seconds

²https://github.com/kubernetes/community/blob/master/sig-scalability/slos/pod_startup_latency.md

Measuring this myself

- Clusterloader2 measures the 50th, 90th and 99th percentiles of pod startup latency.
- I created a cluster with 20 nodes, to match the user's cluster.
- I started testing with clusterloader2, first with 5 pods per node and then increasing.

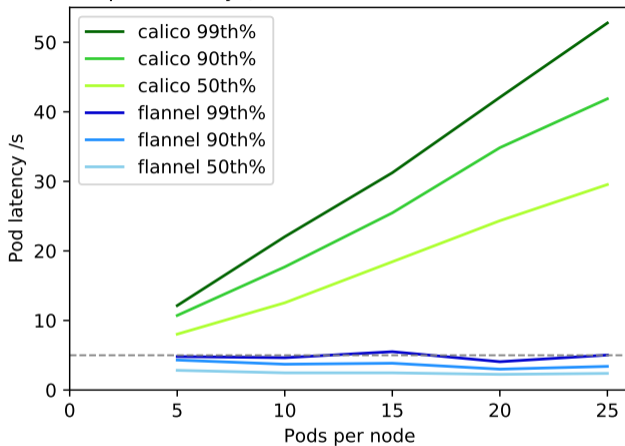
Pod latency (medium master/worker, 20 nodes)



- I looked in the kubelet logs for one node to check for anything unusual.
- I found a lot of errors like this:

```
[ERROR][1278786] customresource.go 136: Error updating resource
Key=IPAMBlock(10-100-145-192-26)
error=Operation cannot be fulfilled on ipamblocks.crd.projectcalico.org:
the object has been modified; please apply your
changes to the latest version and try again
```


CNI pod latency (medium master/worker, 20 nodes)



Is the API server ok?

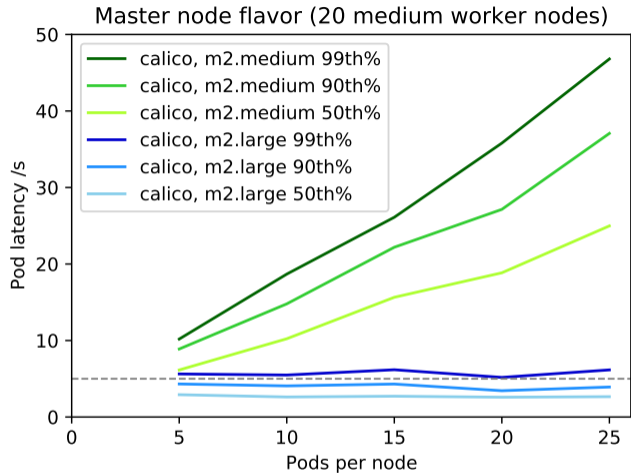
- No.

Is the API server ok?

- No.
- On the master node, the API server and etcd were fully using the available CPU.


Is the API server ok?

- No.
- On the master node, the API server and etcd were fully using the available CPU.
- I didn't change any defaults when creating the cluster, so it was only a 2-core m2.medium master node.



Bigger is better

- Changing the master node to be a 4-core m2.large flavor brought things mostly under control.
- But, the 99th percentile is still not meeting the target.
- Let's go bigger! 16 cores!



So we're done?

- Yes, just make you sure choose a large enough master node to handle the expected amount of load in your cluster.

So we're done?

- Yes, just make you sure choose a large enough master node to handle the expected amount of load in your cluster.
- But.....

So we're done?

- Yes, just make you sure choose a large enough master node to handle the expected amount of load in your cluster.
- But.....
- I'm not satisfied yet!

What's an IPAMBlock?

```
apiVersion: crd.projectcalico.org/v1
kind: IPAMBlock
metadata:
  name: 10-100-32-128-26
  resourceVersion: "1500"
spec:
  affinity: host:tgh-clusterloader-118-dz7jvsoaiqcc-node-0
  cidr: 10.100.32.128/26
  allocations:
    - 0
    - 1
    - 2
    - 3
```

In the code

- When the calico-ipam CNI plugin is assigning an IP for a pod on a node, there is a section where it
 - Gets the IPAMBlock for the node
 - Picks an unused IP from the block
 - Creates an IPAMHandle
 - Tries to update the IPAMBlock to mark the IP as in use

Pod 1

GET

Pod 2

GET

Pod 3

GET

Pod 4

GET

Pod 1



Pod 2



Pod 3



Pod 4



Pod 1



Pod 2



Pod 3




Pod 4




Pod 1




Pod 2



Pod 3



Pod 4



Pod 1




Pod 2



Pod 3




Pod 4




Pod 1




Pod 2



Pod 3



Pod 4



Feedback loop

- The load on the API server causes calico to make conflicting updates to IPAMBlocks.
- Calico retries the updates which failed.
- Putting more load on the API server!

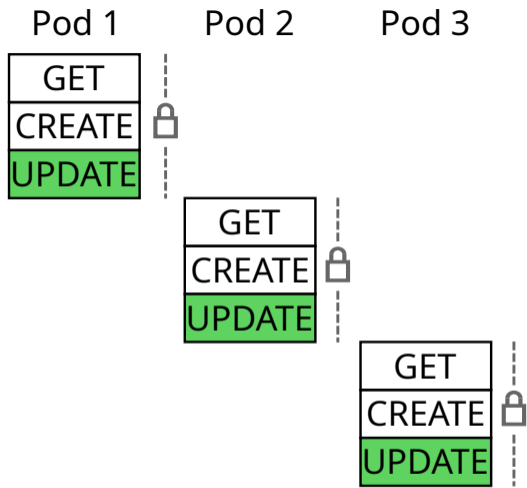
So what's the solution?

- The solution is still to use a larger master node.
- But at this point I went to the Calico #kubernetes Slack channel to speak to the developers.

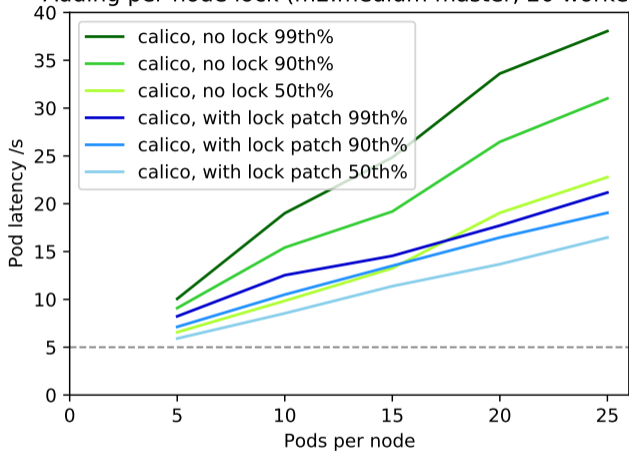
So what's the solution?

- There are improvements that can be made in the worst case.
- Using a shared (file based) lock on each node to guard the IPAMBlock API, contention can be eliminated³.

³<https://github.com/projectcalico/cni-plugin/pull/985>







Adding per-node lock (m2.medium master, 20 workers)



Checking the API server

- The API server logs contain request latencies, the time from receiving a request to returning a response.
- We can look at those, and the raw request rate, to see how this change impacts the API.





The result

- In total, number of requests from the CNI is reduced from 20k to 10k over the duration of the test.
- The peak API response latency is cut in half as well.

And on larger master nodes

- Possible 60ms \rightarrow 30ms peak response time on m2.large master node with the patch applied.
- m2.2xlarge master node maintains \sim 3ms response time even without patch.

Finally

- And in the end...

Finally

- And in the end...
- The problem in the user report was unrelated.

Finally

- And in the end...
- The problem in the user report was unrelated.
- (Spinning disks on the nodes).

End

Extra slides

