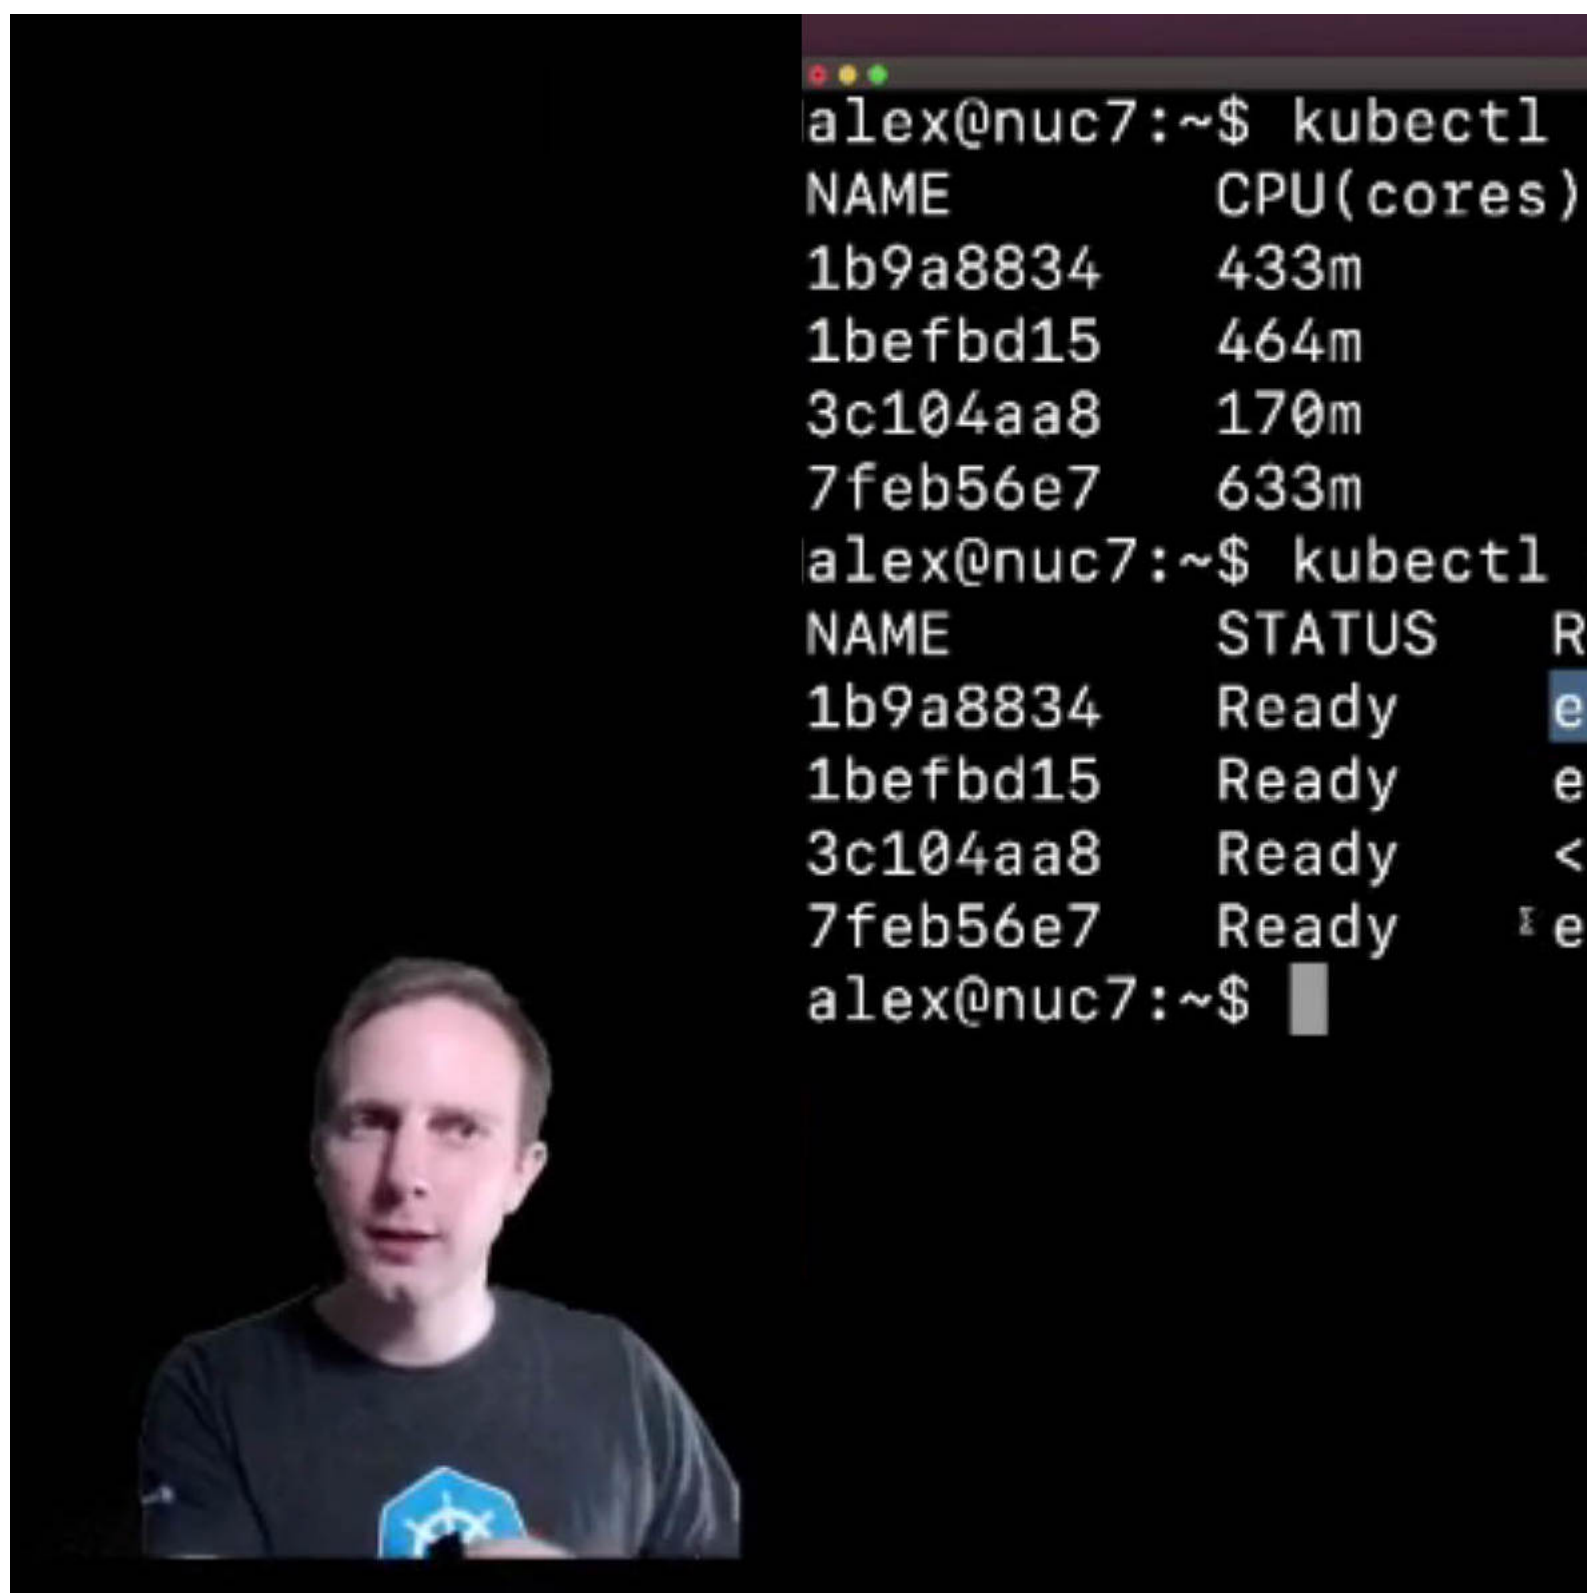


# How to Troubleshoot Applications on Kubernetes

This article was fetched from an [rss feed](#)(01/06/2022)



This is my guide to troubleshooting applications deployed on a Kubernetes cluster. It's based upon my experience with the technology since I switched over from Docker in early 2017.

There are many third-party tools and techniques available, but I want to focus in on things that you'll find on almost every computer, or CLIs that you can download quickly for MacOS, Windows or Linux.

What if you're a Kubernetes expert?

If you're already supporting clusters for customers, or developing software, you may know some of these. So perhaps you may find this content useful to send over to your own users, and colleagues.

In any case, stay with me and see whether there's a new flag or technique you may pick up. You may also like some of the more advanced supplementary material that I'm linking to throughout the post.

Got a comment, question or suggestion? [Let's talk on Twitter](#)

## Is it there?

When your application isn't working, you may want to check that all its resources have been created.

The first command you learn to probably going to be "kubectl get pods"

But remember, that Kubernetes supports various namespaces in order to segregate and organise workloads. What you've installed may not be in the default namespace.

Here's how can change the namespace to look into the "kube-system" or "openfaas-fn" namespace for instance:

```
kubectl get pods --namespace kube-system
kubectl get pods -n openfaas
```

You can query all of the namespaces available with:

```
kubectl get pods --all-namespaces
kubectl get pods -A
```

The -A flag was added to kubectl in the 2-3 years, and means you can save on typing.

Now of course, Pods are just one of the things we care about. The above commands can also take other objects like Service, Deployment, Ingress and more.

## Why isn't it working?

When you ran "kubectl get", you may have seen your resource showing as 0/1, or even as 1/1 but in a crashing or errored state.

How do we find out what's going wrong?

You may be tempted to reach for "kubectl logs", but this only shows logs from applications that have started, if your pod didn't start, then you need to find out what's preventing that.

You're probably running into one of the following:

- The image can't be pulled
- There's a missing volume or secret
- No space in the cluster for the workload
- Taints or affinity rules preventing the pod from being scheduled

Now kubectl get events on its own isn't very useful, because all the rows come out in what appears to be a random order. The fix is something you'll have to get tattooed somewhere prominent, because there's no shorthand for this yet.

```
kubectl get events \
--sort-by=.metadata.creationTimestamp
```

This will print out events in the default namespace, but it's very likely that you're working in a specific namespace, so make sure to include the --namespace or --all-namespaces flag:

```
kubectl get events \
--sort-by=.metadata.creationTimestamp \
-A
```

Events are not just useful for finding out why something's not working, they also show you how pods are pulled, scheduled and started on a cluster.

Add "--watch" or "-w" to the command to watch an OpenFaaS function being created for instance:

```
kubectl get events \
--sort-by=.metadata.creationTimestamp \
-n openfaas-fn
```

And with the watch added on:

```
kubectl get events \
--sort-by=.metadata.creationTimestamp \
-n openfaas-fn \
--watch
```

LAST SEEN	TYPE	REASON	OBJECT	MESSAGE
0s	Normal	Synced	function/figlet	Function synced successfully
0s	Normal	ScalingReplicaSet	deployment/figlet	Scaled up replica set figlet-5485896b55 to 1
1s	Normal	SuccessfulCreate	replicaset/figlet-5485896b55	Created pod: figlet-5485896b55-j9mbd
0s	Normal	Scheduled	pod/figlet-5485896b55-j9mbd	Successfully assigned openfaas-fn/figlet-5485896b55-j9mbd to k3s-pi
0s	Normal	Pulling	pod/figlet-5485896b55-j9mbd	Pulling image "ghcr.io/openfaas/figlet:latest"
0s	Normal	Pulled	pod/figlet-5485896b55-j9mbd	Successfully pulled image "ghcr.io/openfaas/figlet:latest" in 632.059147ms
0s	Normal	Created	pod/figlet-5485896b55-j9mbd	Created container figlet
0s	Normal	Started	pod/figlet-5485896b55-j9mbd	Started container figlet

There's actually quite a lot that's going on in the above events. You can then run something like kubectl scale -n openfaas-fn deploy/figlet --replicas=0 to scale it down and watch even more events getting generated as the pod is removed.

Now, there is a new command called kubectl events which you may also want to look into, however my kubectl version was too old, and it's only an alpha feature at present.

You can upgrade kubectl using [arkade](#) whether you're on Windows, MacOS or Linux:

```
arkade get kubectl@v1.24.0
$HOME/.arkade/bin/kubectl alpha events -n openfaas-fn
```

Now as I understand it, this new command does order the events, to keep an eye on how it progresses to see if the Kubernetes community promote it to generally available (GA) status or not.

## It starts, but doesn't work

So the application starts up with 1/1 pods, or starts up then keeps crashing, so you're seeing a lot of restarts when you type in:

```
kubectl get pod
NAME                                READY   STATUS    RESTARTS   AGE
ingress-nginx-controller-54d8b558d4-591j2  1/1     Running   5          114d
```

This is probably where the age old favourite "kubectl logs" comes in.

Most people do not work with Pods directly, but create a Deployment, which in turn creates a number of Pods depending on the replicas field. That's true of the way I've installed ingress-nginx, which you can see has restarted 5 times. It's now been running for 114 days or nearly 4 months.

```
kubectl logs ingress-nginx-controller-54d8b558d4-591j2|wc -l
```



openfaas	gateway-7c96d7cbc4-d47wh	12m	24Mi
openfaas	inlets-portal-client-5d64668c8d-8f85d	1m	5Mi
openfaas	nats-675f8bbc59-cndw8	2m	12Mi
openfaas	pro-builder-6ff7bd4985-fxswm	1m	117Mi
openfaas	prometheus-56b84ccf6c-x4vr2	10m	28Mi
openfaas	queue-worker-6d4756d8d9-km8g2	1m	2Mi
openfaas	queue-worker-6d4756d8d9-xgl8w	1m	2Mi
openfaas-fn	bcrypt-7d69d458b7-7zr94	12m	16Mi
openfaas-fn	chaos-fn-c7b647c99-f9wz7	2m	5Mi
openfaas-fn	cows-594d9df8bc-z15rr	2m	12Mi
openfaas-fn	shasum-5c6cc9c56c-x5v2c	1m	3Mi

Now we can see how well the pods are balanced across machines:

```
kubect1 top node
NAME          CPU(cores)   CPU%   MEMORY(bytes)  MEMORY%
k3s-agent-1   236m         5%     908Mi          23%
k3s-pi        528m         13%    1120Mi         14%
```

So we actually have plenty of headroom to deploy some more workloads.

Now you cannot add any kind of "--watch" or "--follow" flag to this command, so if you want to watch it whilst you scale some functions or deploy a bunch of new pods, you need to use a bash utility like "watch".

Try this for example:

```
# Terminal 1
watch "kubect1 top pod -A"

# Terminal 2
watch "kubect1 top node"
```

The metrics-server is an optional add-on, which you can install with arkade or helm:

```
arkade install metrics-server
```

## Have you turned it off and on again?

It's tragic, but true, that turnings things off and on again can fix many errors that we run into on a daily basis.

Restarting a deployment in Kubernetes may fix issues due to forcing your code to reconnect to services, pull down an updated image, or just release memory and database connections.

The command you're looking for is `kubect1 rollout restart`.

Let's restart the 2x OpenFaaS queue-workers, whilst watching the logs with `stern`, and the events from the namespace in another window.

```
# Terminal 1
kubect1 get events \
  -n openfaas
  --sort-by=.metadata.creationTimestamp \
  --watch

# Terminal 2
stern -n openfaas queue-worker.* --since 5s

# Terminal 3
kubect1 rollout restart \
  -n openfaas deploy/queue-worker
```

Note the syntax for `stern` is a regular expression, so it'll match on anything that starts with "queue-worker" as a prefix. The `--since 5s` is similar to what we used with `kubect1 logs`, to keep what we're looking at recent.

So we see our two new pods show up in `stern`, something that `kubect1 logs` would not be able to do for us:

```
queue-worker-6d4756d8d9-km8g2 queue-worker 2022/06/01 10:47:27 Connect: nats://nats.openfaas.svc.cluster.local:4222
queue-worker-6d4756d8d9-xgl8w queue-worker 2022/06/01 10:47:30 Connect: nats://nats.openfaas.svc.cluster.local:4222
```

And the events for the queue-worker show the new pods being created and the older ones being removed from the cluster.

LAST SEEN	TYPE	REASON	OBJECT	MESSAGE
69s	Normal	Pulling	pod/queue-worker-6d4756d8d9-xgl8w	Pulling image "ghcr.io/openfaasltd/queue-worker:0.1.5"
68s	Normal	Pulled	pod/queue-worker-6d4756d8d9-xgl8w	Successfully pulled image "ghcr.io/openfaasltd/queue-worker:0.1.5" in 597
68s	Normal	Created	pod/queue-worker-6d4756d8d9-xgl8w	Created container queue-worker
68s	Normal	Started	pod/queue-worker-6d4756d8d9-xgl8w	Started container queue-worker
67s	Normal	ScalingReplicaSet	deployment/queue-worker	Scaled down replica set queue-worker-755598f7fb to 0
67s	Normal	SuccessfulDelete	replicaset/queue-worker-755598f7fb	Deleted pod: queue-worker-755598f7fb-h2cfx
67s	Normal	Killing	pod/queue-worker-755598f7fb-h2cfx	Stopping container queue-worker

## What we didn't talk about

I need to check on applications for two reasons. The first is that I actually write software for Kubernetes such as [openfaas](#) and [inlets-operator](#). During development, I need most of the commands above to check when things go wrong, or to see output from changes I've made. The second reason is that my company supports OpenFaaS and inlets users in production. Supporting users remotely can be challenging, especially if they are not used to troubleshooting applications on Kubernetes.

There are so many things that need to be checked in a distributed system, so for OpenFaaS users, we wrote them all down in a [Troubleshooting guide](#) and you'll see some of what we talked about today covered there.

In my experience, application-level metrics are essential to being able to evaluate how well a service is performing. With very little work, you can record Rate, Error and Duration (RED) for your service, so that when the errors are more subtle, you can start to understand what may be going on.

Metrics are beyond the scope of this article, however if you'd like to get some experience, in [my eBook Everyday Go](#) I show you how to add metrics to a Go HTTP server and start monitoring it with Prometheus.





### [Adding Prometheus metrics to a HTTP server in Everyday Go](#)

You can learn more about how we use Prometheus in OpenFaaS in this talk from KubeCon: [How and Why We Rebuilt Auto-scaling in OpenFaaS with Prometheus](#)

I see Kubernetes as a journey and not a destination. I started my walk in 2017, and it was very difficult for me at the time. If, like me, you now have several years under your belt, try to have some empathy for those who are only starting to use the technology now.

You may also like these two free courses that I wrote for the Linux Foundation/CNCF:

- [Kubernetes at the Edge with K3s](#)
- [Introduction to Serverless on Kubernetes](#)

Kubernetes, K3s, OpenFaaS, Docker, Prometheus, Grafana and many other Cloud Native tools are written in Go, learn what you need to know with [my eBook Everyday Golang](#)

Got a comment, question or suggestion? [Let's talk on Twitter](#)