Managing Kubernetes without losing your cool

DDD East Midlands

October 7th 2023

Hi 👋,

I’m **Marcus Noble,** a *platform engineer* at Giant Swarm

I’m found around the web as ✨***AverageMarcus***✨ in most places and **@Marcus@k8s.social** on Mastodon 🐘

6+ years experience running Kubernetes in production environments.

💙

Summary

My 10 tips for working with Kubernetes

#1 → #5

Anyone can start using these today

#6 → #7

Good to know a little old-skool ops first

#8 → #10

Good have some programming knowledge

#0 - Pay someone else to deal with it

#1 - Love your terminal

#1 - Love your terminal

★ Bash? ZSH? Fish? 🤷 - Doesn't matter as long as you're comfortable with it.

★ "rc" files - e.g. `.bashrc`, `.zshrc`
These set configuration for each terminal session you open.

★ `alias` - easily create your own terminal commands

★ Look for "dotfiles" on GitHub - e.g. https://github.com/averagemarcus/dotfiles

#2 - Learn to love `kubectl`

#2 - Learn to love `kubectl`

- ★ Add `alias k='kubectl'` to your .bashrc / .zshrc / .whateverrc

Tip #1 in action

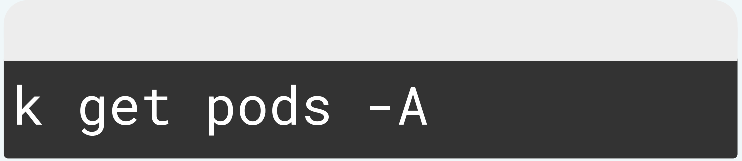

```
k get pods -A
```

- ★ The official docs offer a single page view of all built in commands: kubernetes.io/docs/reference/generated/kubectl/kubectl-commands

- ★ `kubectl explain` is your friend! Find out what any property of any Kubernetes resource is for. ➡

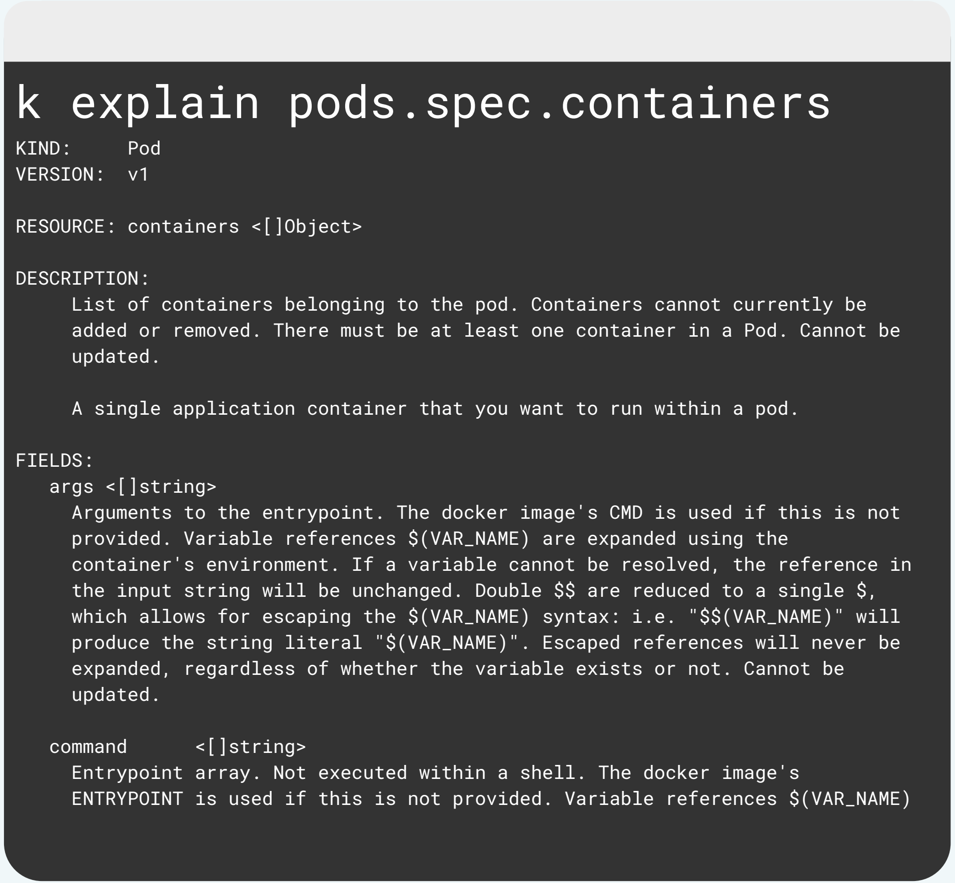

```
k explain pods.spec.containers
KIND:     Pod
VERSION:  v1

RESOURCE: containers <[]Object>

DESCRIPTION:
     List of containers belonging to the pod. Containers cannot currently be
     added or removed. There must be at least one container in a Pod. Cannot be
     updated.

     A single application container that you want to run within a pod.

FIELDS:
   args <[]string>
     Arguments to the entrypoint. The docker image's CMD is used if this is not
     provided. Variable references $(VAR_NAME) are expanded using the
     container's environment. If a variable cannot be resolved, the reference in
     the input string will be unchanged. Double $$ are reduced to a single $,
     which allows for escaping the $(VAR_NAME) syntax: i.e. "$$(VAR_NAME)" will
     produce the string literal "$(VAR_NAME)". Escaped references will never be
     expanded, regardless of whether the variable exists or not. Cannot be
     updated.

   command      <[]string>
     Entrypoint array. Not executed within a shell. The docker image's
     ENTRYPOINT is used if this is not provided. Variable references $(VAR_NAME)
```

#3 - Multiple kubeconfigs

#3 - Multiple kubeconfigs

- Quick switch between different Kubernetes contexts (clusters) and between different namespaces.

- `kubectx` and `kubens`
  https://github.com/ahmetb/kubectx

- `kubie`
  https://github.com/sbstp/kubie

- `kubeswitch`
  https://github.com/danielfoehrKn/kubeswitch

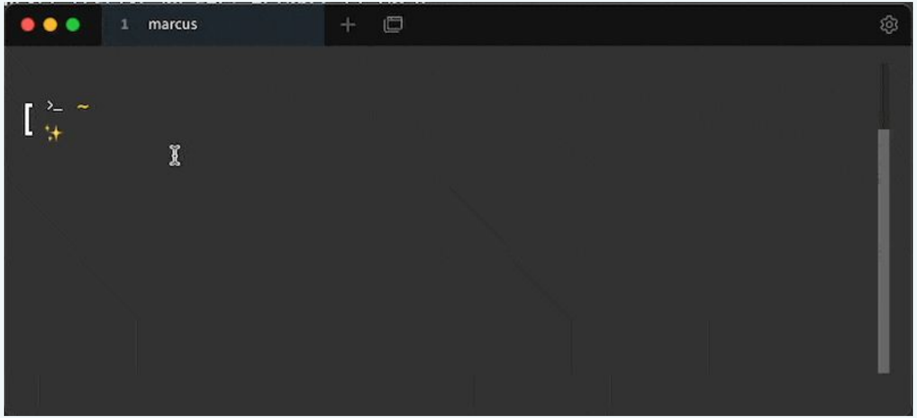

#4 - Interactive UIs

#4 - k9s

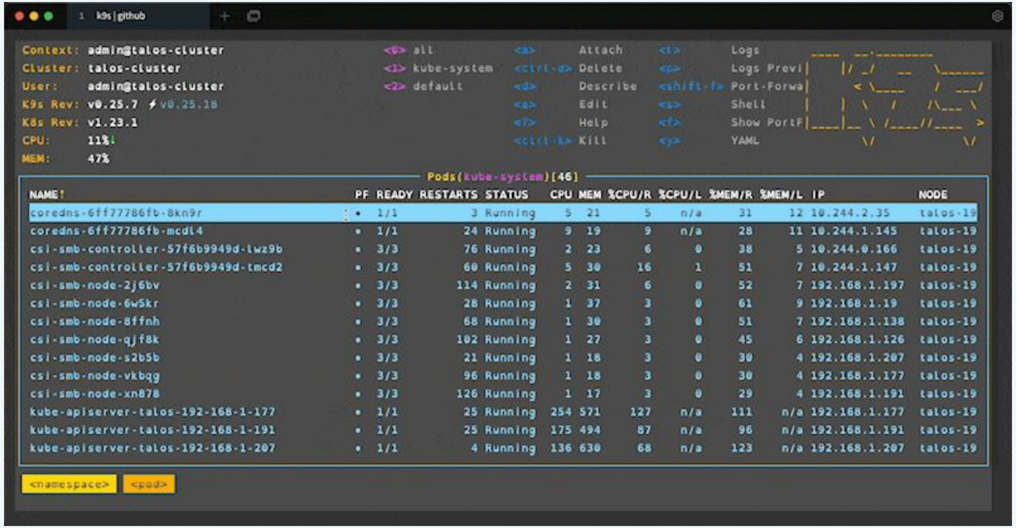

```
Context: admin@talos-cluster                    <0> all            <a>       Attach     <l>       Logs
Cluster: talos-cluster                          <1> kube-system    <ctrl-d>  Delete     <p>       Logs Previ
User:    admin@talos-cluster                    <2> default        <d>       Describe   <shift-f> Port-Forwa
K9s Rev: v0.25.7 ⚡v0.25.18                                          <e>       Edit       <s>       Shell
K8s Rev: v1.23.1                                                   <?>       Help       <f>       Show PortF
CPU:     11%↓                                                      <ctrl-k>  Kill       <y>       YAML
MEM:     47%
                                        Pods(kube-system)[46]
NAME↑                                     PF READY RESTARTS STATUS   CPU MEM %CPU/R %CPU/L %MEM/R %MEM/L IP             NODE
coredns-6ff77786fb-8kn9r                  •  1/1          3 Running    5  21      5    n/a     31     12 10.244.2.35    talos-19
coredns-6ff77786fb-mcdl4                  •  1/1         24 Running    9  19      9    n/a     28     11 10.244.1.145   talos-19
csi-smb-controller-57f6b9949d-lwz9b       •  3/3         76 Running    2  23      6      0     38      5 10.244.0.166   talos-19
csi-smb-controller-57f6b9949d-tmcd2       •  3/3         60 Running    5  30     16      1     51      7 10.244.1.147   talos-19
csi-smb-node-2j6bv                        •  3/3        114 Running    2  31      6      0     52      7 192.168.1.197  talos-19
csi-smb-node-6w5kr                        •  3/3         28 Running    1  37      3      0     61      9 192.168.1.19   talos-19
csi-smb-node-8ffnh                        •  3/3         68 Running    1  30      3      0     51      7 192.168.1.138  talos-19
csi-smb-node-qjf8k                        •  3/3        102 Running    1  27      3      0     45      6 192.168.1.126  talos-19
csi-smb-node-s2b5b                        •  3/3         21 Running    1  18      3      0     30      4 192.168.1.207  talos-19
csi-smb-node-vkbqg                        •  3/3         96 Running    1  18      3      0     30      4 192.168.1.177  talos-19
csi-smb-node-xn878                        •  3/3        126 Running    1  17      3      0     29      4 192.168.1.191  talos-19
kube-apiserver-talos-192-168-1-177        •  1/1         25 Running  254 571    127    n/a    111    n/a 192.168.1.177  talos-19
kube-apiserver-talos-192-168-1-191        •  1/1         25 Running  175 494     87    n/a     96    n/a 192.168.1.191  talos-19
kube-apiserver-talos-192-168-1-207        •  1/1          4 Running  136 630     68    n/a    123    n/a 192.168.1.207  talos-19

<namespace> <pod>
```

github.com/derailed/k9s

#4 - OpenLens

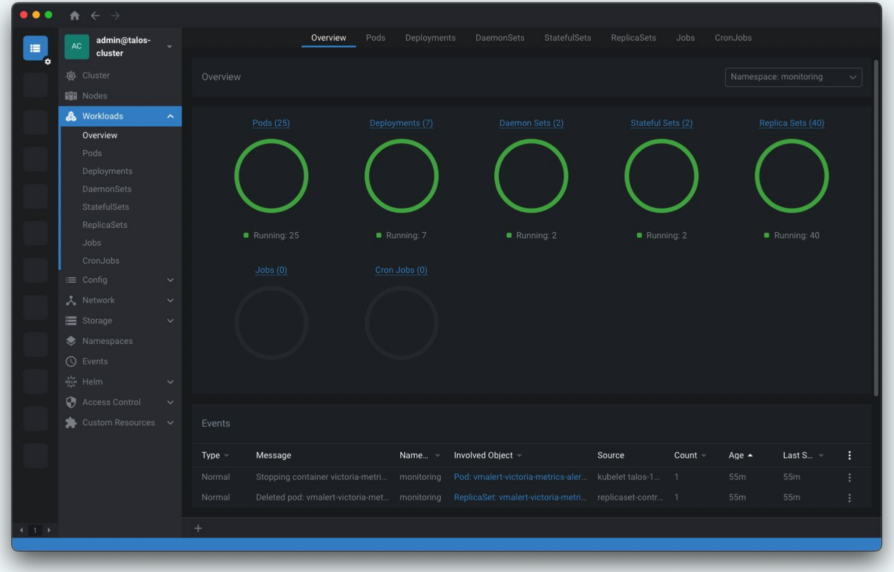

Overview

Namespace: monitoring

Pods (25) — Running: 25

Deployments (7) — Running: 7

Daemon Sets (2) — Running: 2

Stateful Sets (2) — Running: 2

Replica Sets (40) — Running: 40

Jobs (0)

Cron Jobs (0)

Events

Type	Message	Name...	Involved Object	Source	Count	Age	Last S...
Normal	Stopping container victoria-metri...	monitoring	Pod: vmalert-victoria-metrics-aler...	kubelet talos-1...	1	55m	55m
Normal	Deleted pod: vmalert-victoria-met...	monitoring	ReplicaSet: vmalert-victoria-metri...	replicaset-contr...	1	55m	55m

github.com/MuhammedKalkan/OpenLens

#5 - kubectl plugins

#5 - kubectl plugins

- Any command in your `$PATH` that is prefixed with `kubectl-` becomes a kubectl plugin

- Krew - package manager for kubectl plugins
  github.com/kubernetes-sigs/krew
- Install plugins with:
  `kubectl krew install <PLUGIN NAME>`

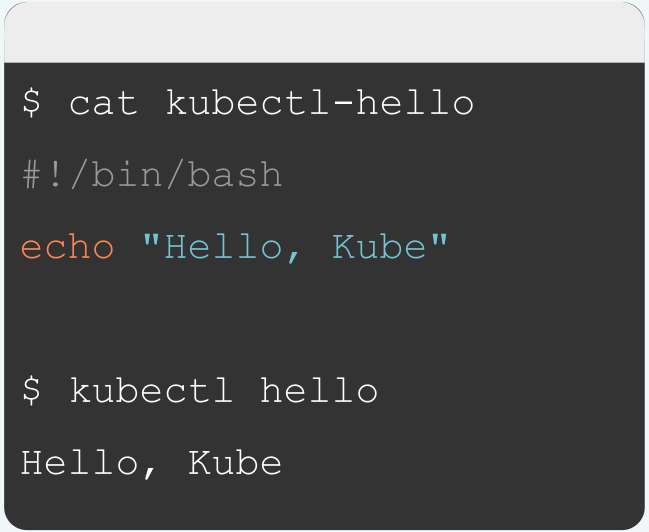

```
$ cat kubectl-hello
#!/bin/bash
echo "Hello, Kube"


$ kubectl hello
Hello, Kube
```

- Some of my fave plugins:
  - **stern** - Multi-pod/container log tailing
  - **tree** - Show hierarchy of resources based on ownerReferences
  - **community-images** - Find images still referencing the **k8s.gcr.io** registry.
  - **gs** - Giant Swarm’s plugin for working with our managed clusters

Summary

My 10 tips for working with Kubernetes

✅ **~~#1 → #5~~**

~~Anyone can start using these today~~

#6 → #7

Good to know a little old-skool ops first

#8 → #10

Good have some programming knowledge

#6 - Pod Debugging

#6 - Pod Debugging: kshell

Launch a temporary pod running a bash shell for cluster debugging

Tip #1 in action, again

```
alias kshell='kubectl run \
    -it                   \
    --image bash          \
    --restart Never       \
    --rm                  \
    shell'
```

Need more tools? Replace this with alpine or ubuntu

#6 - Pod Debugging: kshell

Launch a temporary pod running a bash shell for cluster debugging

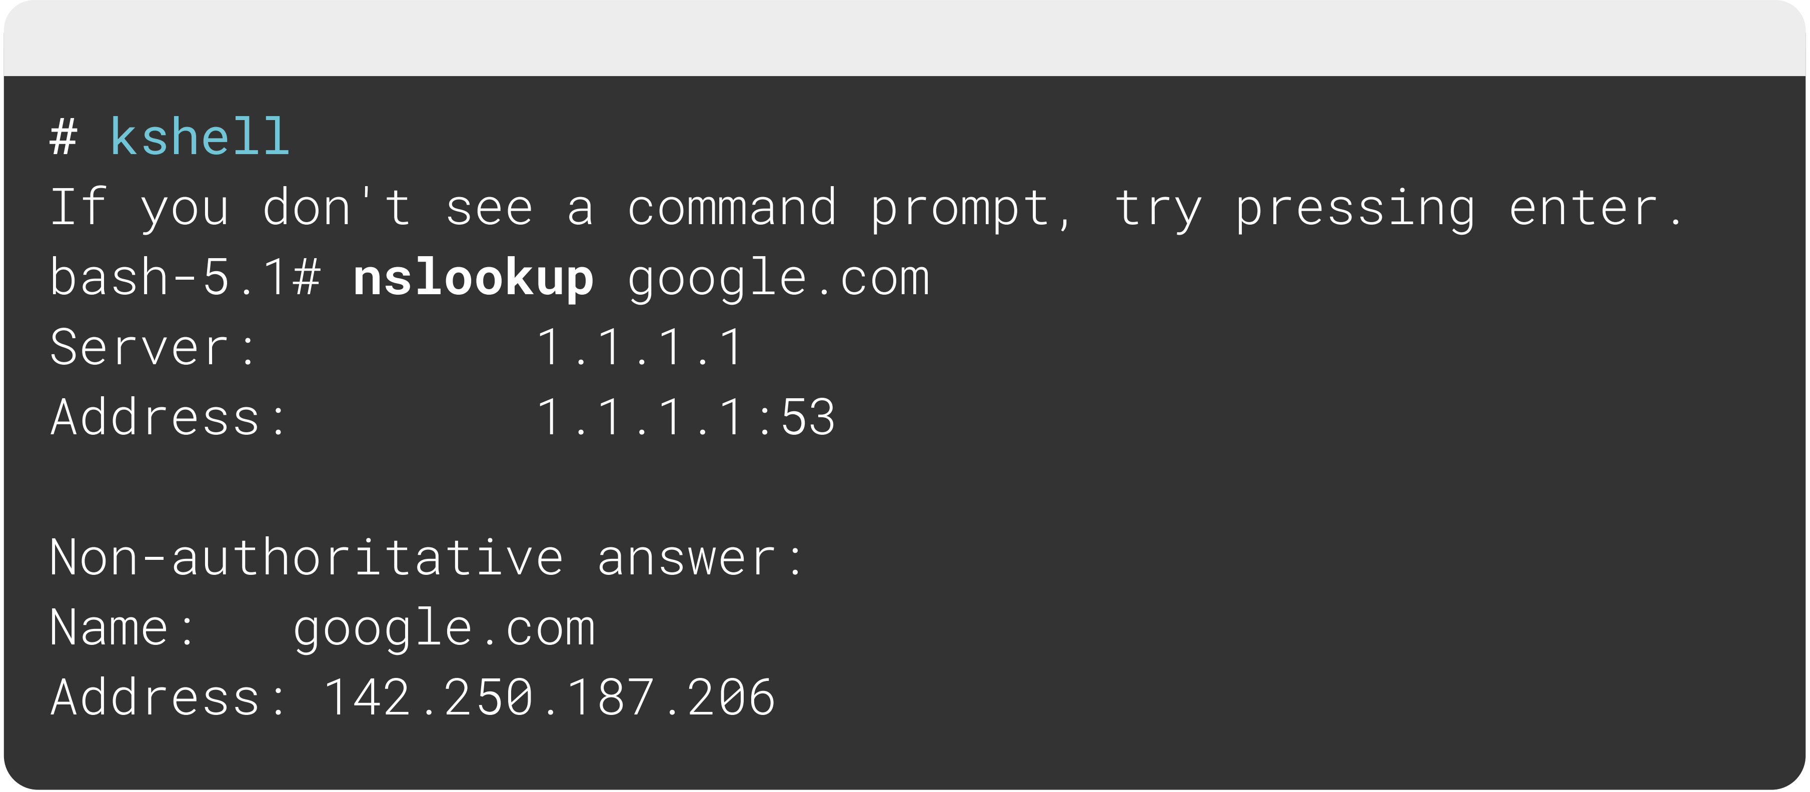

```
# kshell
If you don't see a command prompt, try pressing enter.
bash-5.1# nslookup google.com
Server:         1.1.1.1
Address:        1.1.1.1:53

Non-authoritative answer:
Name:   google.com
Address: 142.250.187.206
```

#6 - Pod Debugging: kubectl exec

Debugging an existing, running pod - `kubectl exec`

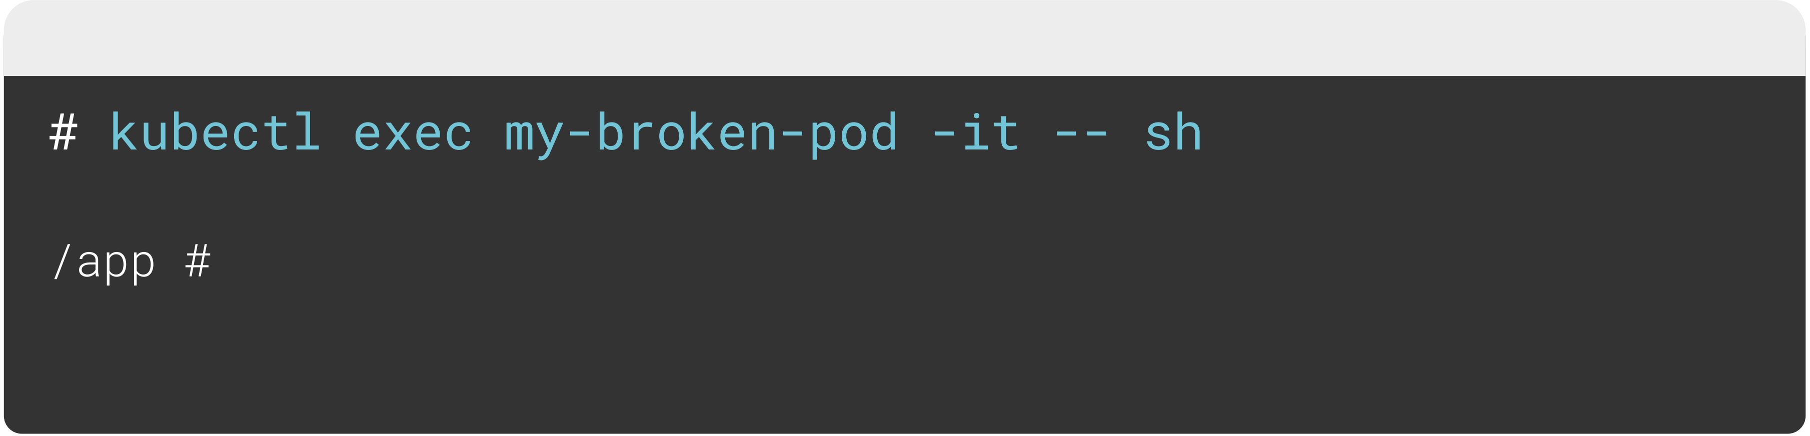

```
# kubectl exec my-broken-pod -it -- sh

/app #
```

Note:

- ★ Needs a shell environment within the container
- ★ Limited to what’s available in the container (or what you can pull in from the ‘net)
- ★ Container needs to be running

#6 - Pod Debugging: kubectl debug

Debugging a running pod - `kubectl exec`

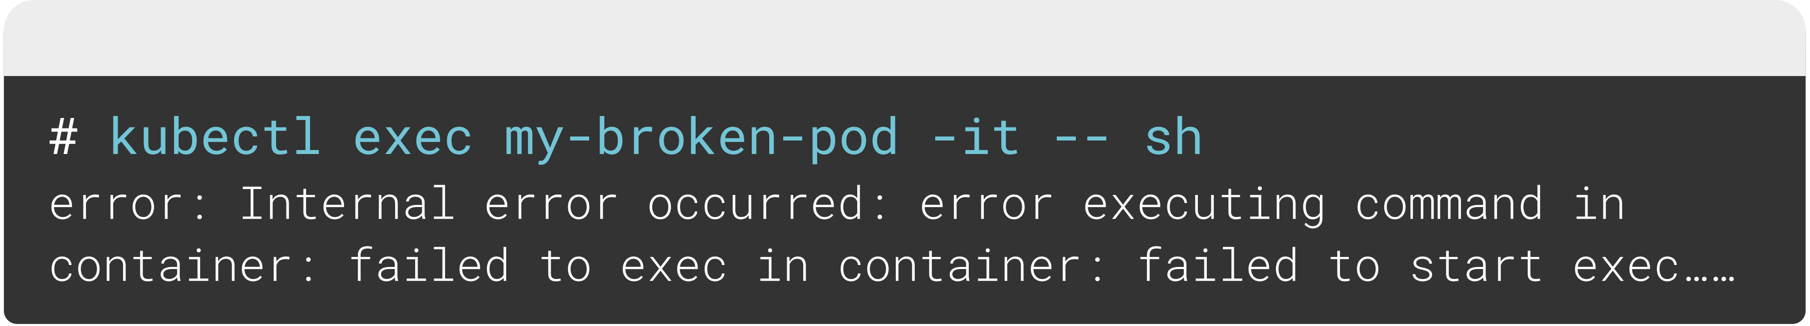

```
# kubectl exec my-broken-pod -it -- sh
error: Internal error occurred: error executing command in
container: failed to exec in container: failed to start exec……
```

Debugging a running pod - `kubectl debug` Requires Kubernetes v1.23+

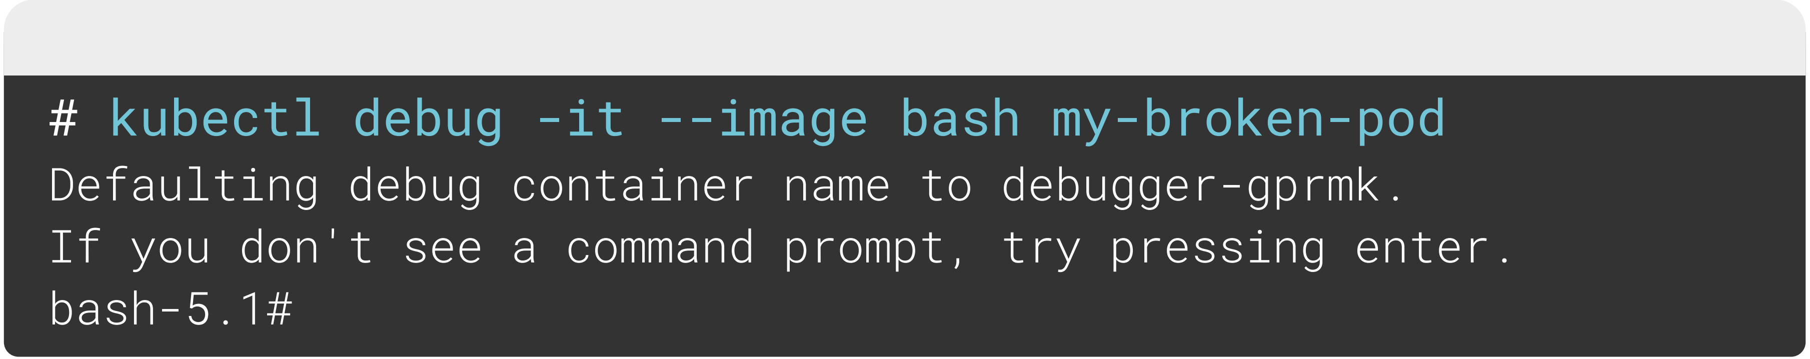

```
# kubectl debug -it --image bash my-broken-pod
Defaulting debug container name to debugger-gprmk.
If you don't see a command prompt, try pressing enter.
bash-5.1#
```

#6 - Pod Debugging: kubectl debug

Example - investigate a CrashLooping pod

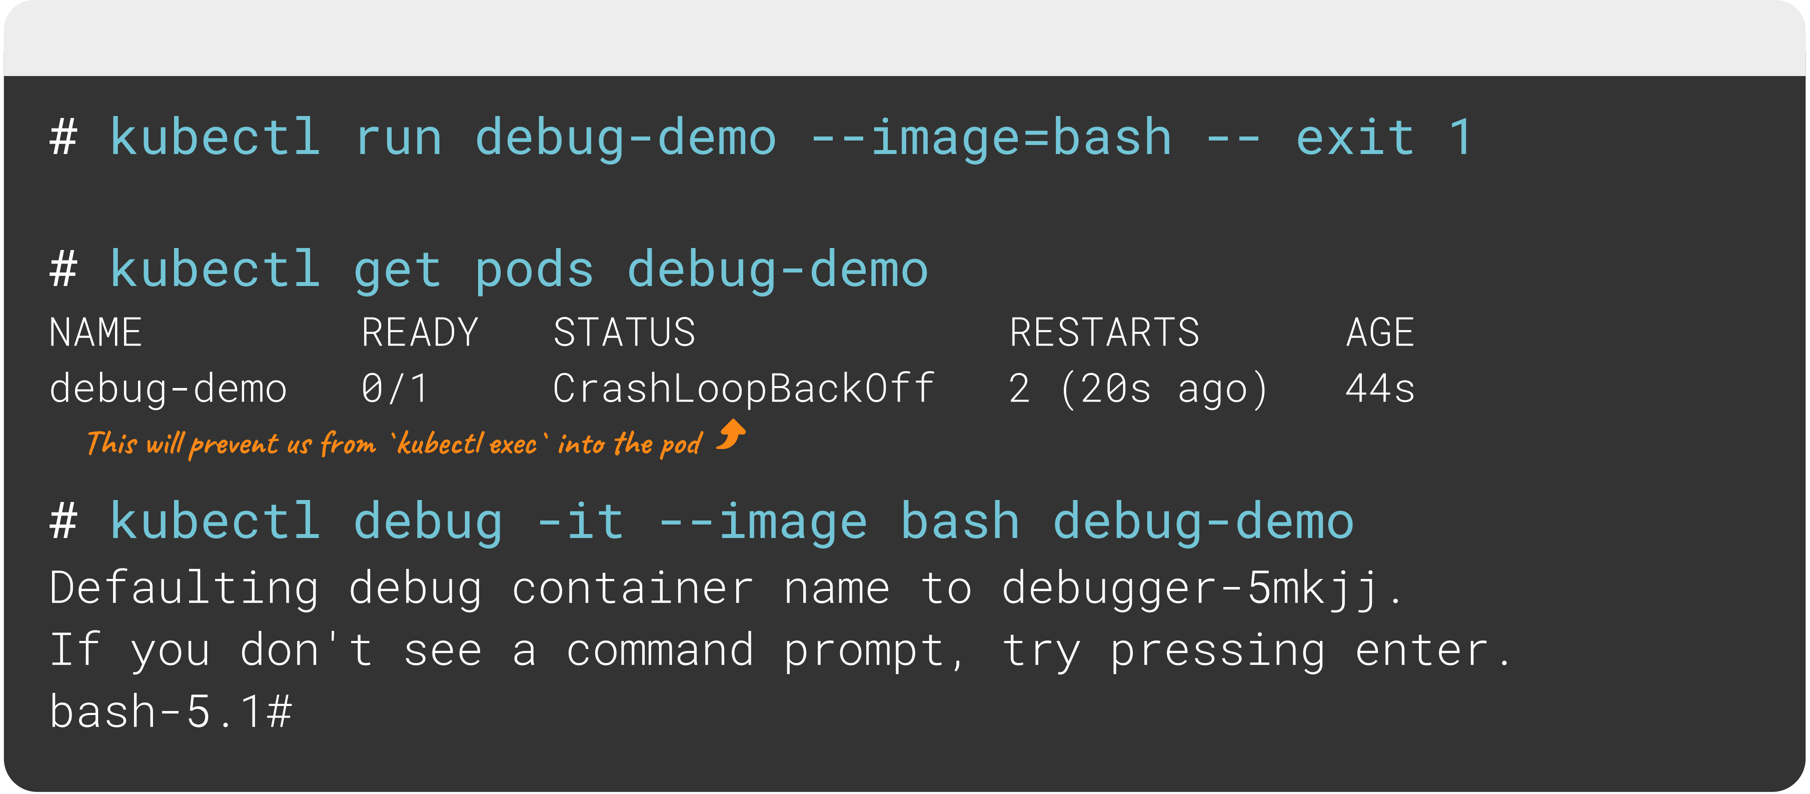

```
# kubectl run debug-demo --image=bash -- exit 1

# kubectl get pods debug-demo
NAME          READY   STATUS             RESTARTS        AGE
debug-demo    0/1     CrashLoopBackOff   2 (20s ago)     44s
```

This will prevent us from `kubectl exec` into the pod

```
# kubectl debug -it --image bash debug-demo
Defaulting debug container name to debugger-5mkjj.
If you don't see a command prompt, try pressing enter.
bash-5.1#
```

#6 - Pod Debugging

When to use what:

	`kshell`	`kubectl exec`	`kubectl debug`
Multiple workloads experiencing network issues	✅		
Workload not running as expected but not CrashLooping and isn't a stripped down image (e.g. not Scratch / Distroless)		✅	
Workload not running as expected but not CrashLooping and has an image based on Scratch / Distroless or similar			✅
Workload is CrashLooping			✅

#7 - Node Debugging

#7 - Node Debugging: kubectl debug (again)

★ Requires Kubernetes v1.23

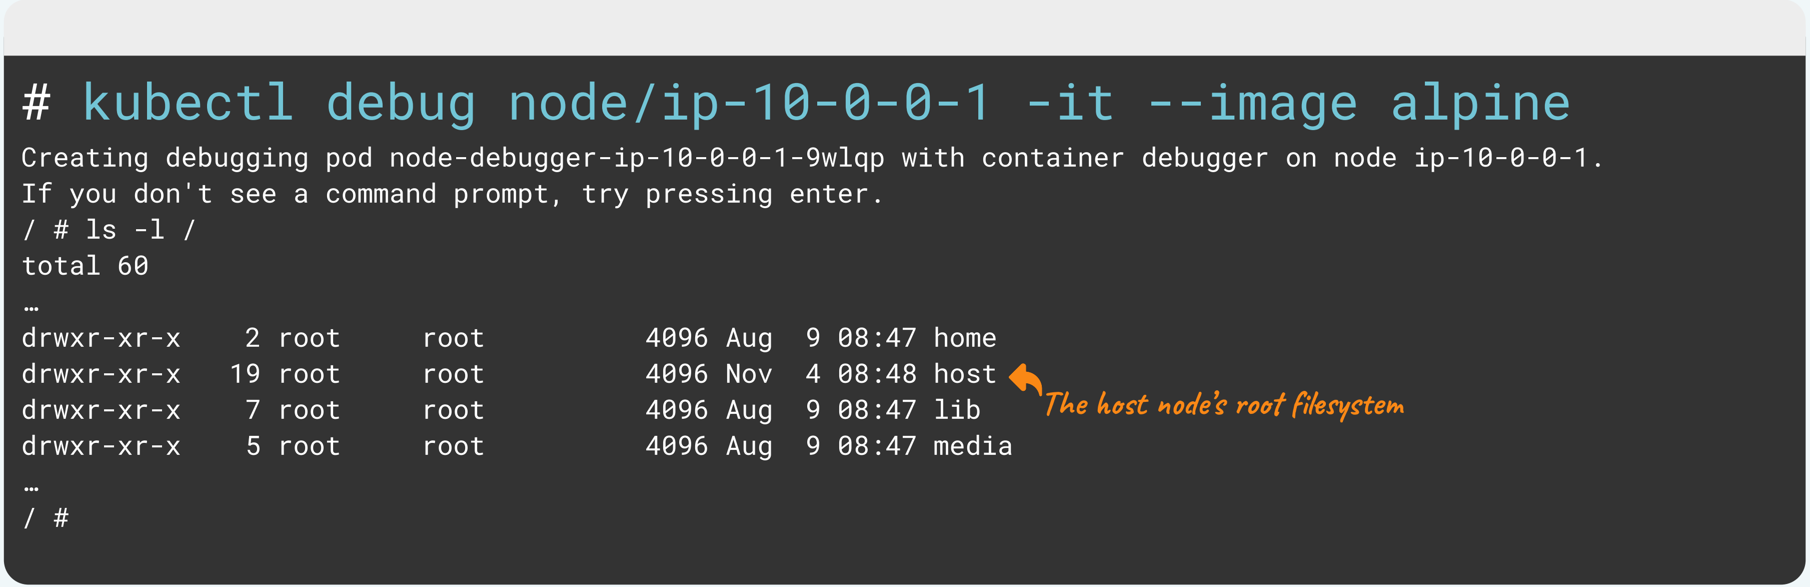

```
# kubectl debug node/ip-10-0-0-1 -it --image alpine
Creating debugging pod node-debugger-ip-10-0-0-1-9wlqp with container debugger on node ip-10-0-0-1.
If you don't see a command prompt, try pressing enter.
/ # ls -l /
total 60
…
drwxr-xr-x    2 root     root          4096 Aug  9 08:47 home
drwxr-xr-x   19 root     root          4096 Nov  4 08:48 host
drwxr-xr-x    7 root     root          4096 Aug  9 08:47 lib
drwxr-xr-x    5 root     root          4096 Aug  9 08:47 media
…
/ #
```

The host node's root filesystem

Why not SSH? - I prefer to use ephemeral instances with the minimal needed to run Kubernetes, no sshd, no port 22 open etc. but there are times when you just need to check what's actually going on with the underlying host machine.

#7 - Node Debugging: nsenter alternative

★ For older clusters before Kubernetes v1.23

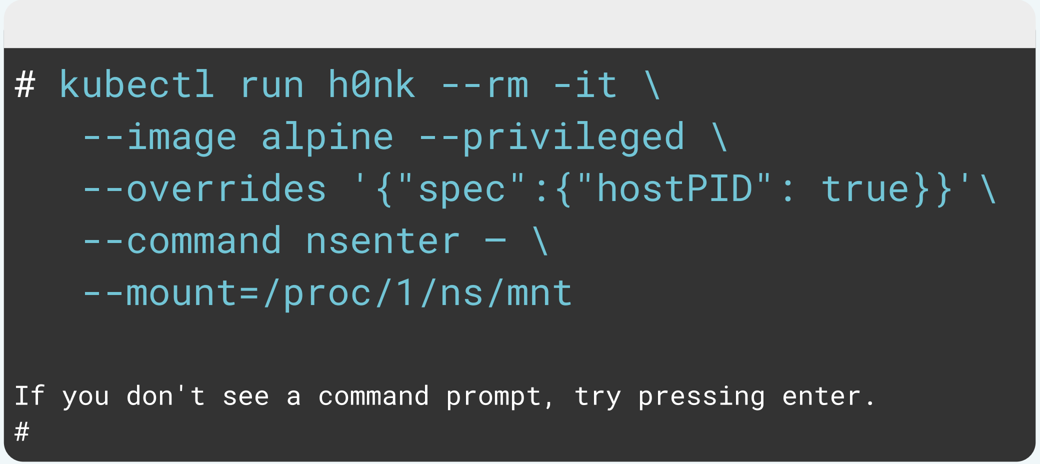

```
# kubectl run h0nk --rm -it \
    --image alpine --privileged \
    --overrides '{"spec":{"hostPID": true}}'\
    --command nsenter – \
    --mount=/proc/1/ns/mnt


If you don't see a command prompt, try pressing enter.
#
```

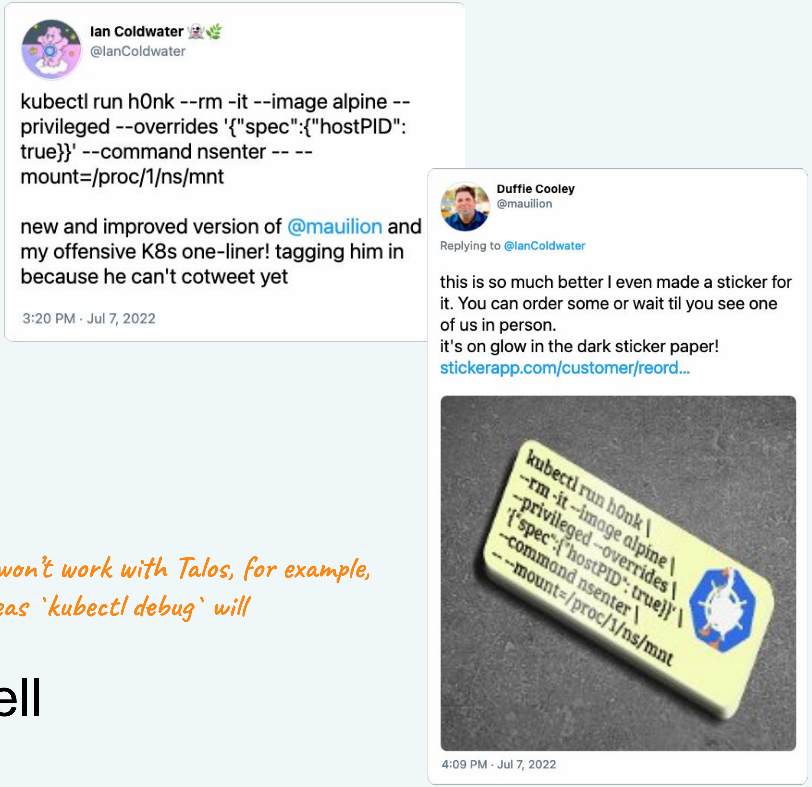

Ian Coldwater 👻🌿
@IanColdwater

kubectl run h0nk --rm -it --image alpine --privileged --overrides '{"spec":{"hostPID": true}}' --command nsenter -- --mount=/proc/1/ns/mnt

new and improved version of @mauilion and my offensive K8s one-liner! tagging him in because he can't cotweet yet

3:20 PM · Jul 7, 2022

Duffie Cooley
@mauilion

Replying to @IanColdwater

this is so much better I even made a sticker for it. You can order some or wait til you see one of us in person.
it's on glow in the dark sticker paper!
stickerapp.com/customer/reord...

kubectl run h0nk \
--rm -it --image alpine \
--privileged --overrides \
'{"spec":{"hostPID": true}}' \
--command nsenter \
--mount=/proc/1/ns/mnt

4:09 PM · Jul 7, 2022

★ Alternatives:
github.com/AverageMarcus/kube-ssh
github.com/giantswarm/kubectl-enter

This won't work with Talos, for example, whereas `kubectl debug` will

★ *Note*: Underlying host needs a valid shell

Summary

My 10 tips for working with Kubernetes

✅ ~~**#1 → #5**~~

~~Anyone can start using these today~~

✅ ~~**#6 → #7**~~

~~Good to know a little old-skool ops first~~

#8 → #10

Good have some programming knowledge

#8 - Webhooks

Webhooks

★ Implement more advanced access control than is possible with RBAC. [Restricting cluster-admin permissions]

★ Add defaulting logic to Kubernetes resources

★ Enforce company policies such as not using `latest` as an image tag or ensuring all workloads have resource requests/limits specified.

★ "Hotfix" for security issues (e.g. injecting env var to prevent Log4Shell exploit). [Log4Shell Mitigation]

⚠️ **Be careful using webhooks as it's easy to introduce cluster-breaking configurations!** 😱 [Webhooks Talk]

Tools:

★ Kyverno - Kubernetes native policy management.

★ OPA Gatekeeper - Policy management built on top of Open Policy Agent

#9 - Kubernetes API

Kubernetes API

Resources:

- kubernetes/client-go - the official Golang module for interacting with the Kubernetes API
- Kubernetes Provider for Terraform (actually uses the above Go module under the hood)
- kubernetes-client org on GitHub has many official clients in different languages

Where is this useful?

★ Building our own CLI / desktop tooling (e.g. k9s, Lens).

★ Cluster automation - resources managed by CI, CronJobs, etc.

★ Building our own operators to extend Kubernetes.

#10 - CRDs & Operators

CRDs & Operators

Extend Kubernetes' built-in API and functionality with your own Custom Resource Definitions (CRDs) and business logic (operators).

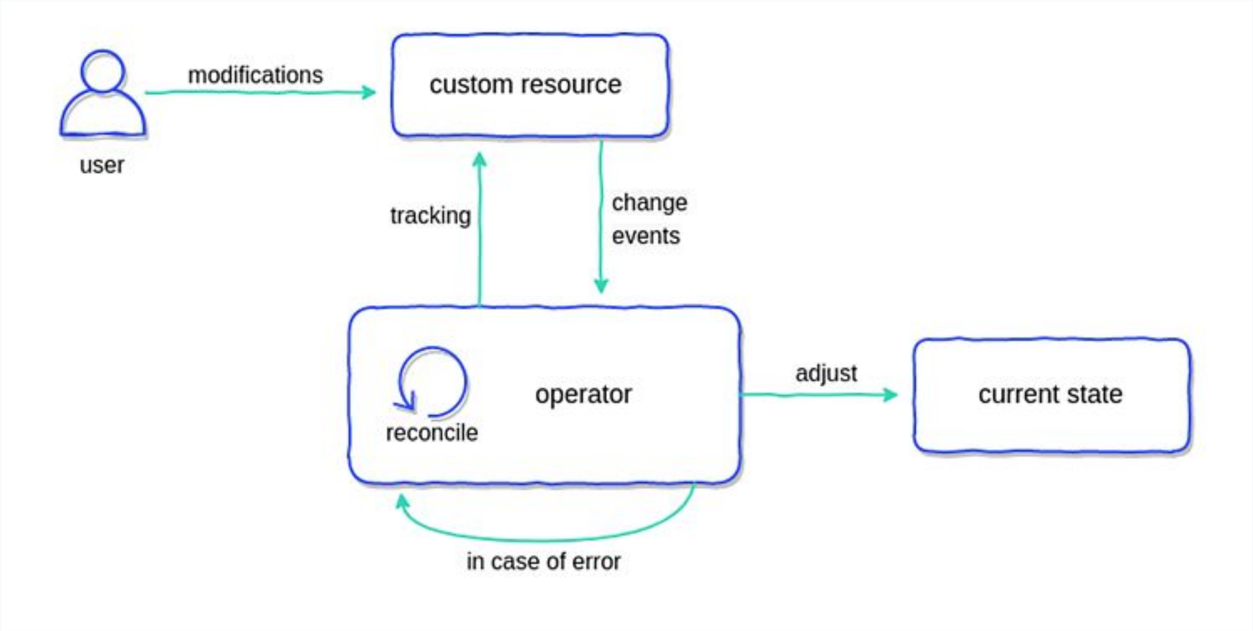

Image credits: Container Solutions
https://blog.container-solutions.com/kubernetes-operators-explained

#10 - CRDs & Operators

Frameworks

kubebuilder OPERATOR FRAMEWORK KUDO Metacontroller

References

- https://kubernetes.io/docs/concepts/extend-kubernetes/operator/
- https://blog.container-solutions.com/kubernetes-operators-explained
- https://operatorhub.io/ - Directory of existing operators

Videos

Summary

My 10 tips for working with Kubernetes

✅ **~~#1 → #5~~**

~~Anyone can start using these today~~

✅ **~~#6 → #7~~**

~~Good to know a little old-skool ops first~~

✅ **~~#8 → #10~~**

~~Good have some programming knowledge~~

Recap

#1 - Love your terminal

#2 - Learn to love kubectl

#3 - Multiple kubeconfigs

#4 - k9s / OpenLens

#5 - Kubectl plugins

#6 - Pod Debugging

#7 - Node Debugging

#8 - Webhooks

#9 - Kubernetes API

#10 - CRDs & Controllers

Wrap-up

Slides and resources available at:

https://go-get.link/dddem23

Thoughts, comments and feedback:

feedback@marcusnoble.co.uk

https://k8s.social/@Marcus

Thank you