



Reference

DOCUMENTATION SETUP CONCEPTS TASKS TUTORIALS **REFERENCE** CONTRIBUTE



kubectl Cheat Sheet



See also: [Kubectl Overview](#) and [JsonPath Guide](#).

This page is an overview of the `kubectl` command.

- [Kubectl Autocomplete](#)
- [Kubectl Context and Configuration](#)
- [Creating Objects](#)
- [Viewing, Finding Resources](#)
- [Updating Resources](#)
- [Patching Resources](#)
- [Editing Resources](#)
- [Scaling Resources](#)
- [Deleting Resources](#)
- [Interacting with running Pods](#)
- [Interacting with Nodes and Cluster](#)
- [What's next](#)

kubectl - Cheat Sheet

Kubectl Autocomplete

BASH

```
source <(kubectl completion bash) # setup autocomplete in bash into the current shell
echo "source <(kubectl completion bash) > ~/.bashrc # add autocomplete permanently"
```

ZSH

```
source <(kubectl completion zsh) # setup autocomplete in zsh into the current shell
echo "if [ \$commands[kubectl] ]; then source <(kubectl completion zsh); fi" >> ~/.zshrc
```

Kubectl Context and Configuration

Set which Kubernetes cluster `kubectl` communicates with and modifies configuration information. See [Authenticating Across Clusters with kubeconfig](#) documentation for detailed config file information.

```
kubectl config view # Show Merged kubeconfig settings.

# use multiple kubeconfig files at the same time and view merged config
KUBECONFIG=~/.kube/config:~/.kube/kubconfig2 kubectl config view

# Get the password for the e2e user
kubectl config view -o jsonpath='{.users[?(@.name == "e2e")].user.password}'

kubectl config current-context # Display the current-context
kubectl config use-context my-cluster-name # set the default context to my-cluster

# add a new cluster to your kubeconf that supports basic auth
kubectl config set-credentials kubeuser/foo.kubernetes.com --username=kubeuser --password=kubeuser

# set a context utilizing a specific username and namespace.
kubectl config set-context gce --user=cluster-admin --namespace=foo \
  && kubectl config use-context gce
```

Creating Objects

Kubernetes manifests can be defined in json or yaml. The file extension `.yaml`, `.yml`, and `.json` can be used.

```
kubectl create -f ./my-manifest.yaml  # create resource(s)
kubectl create -f ./my1.yaml -f ./my2.yaml # create from multiple files
kubectl create -f ./dir # create resource(s) in all manifests
kubectl create -f https://git.io/vPieo # create resource(s) from url
kubectl run nginx --image=nginx # start a single instance of nginx
kubectl explain pods,svc # get the documentation for pod and service
```

Create multiple YAML objects from stdin

```
cat <<EOF | kubectl create -f -
```

```
apiVersion: v1
```

```
kind: Pod
```

```
metadata:
```

```
  name: busybox-sleep
```

```
spec:
```

```
  containers:
```

- name: busybox
- image: busybox
- args:
- sleep
- "1000000"

```
---
```

```
apiVersion: v1
```

```
kind: Pod
```

```
metadata:
```

```
  name: busybox-sleep-less
```

```
spec:
```

```
  containers:
```

- name: busybox
- image: busybox
- args:
- sleep
- "1000"

```
EOF
```

Create a secret with several keys

```
cat <<EOF | kubectl create -f -
```

```
apiVersion: v1
```

```
kind: Secret
```

```
metadata:
```

```
  name: mysecret
```

```
type: Opaque
```

```
data:
```

```
  password: $(echo -n "s33msi4" | base64)
  username: $(echo -n "jane" | base64)
```

```
EOF
```

Viewing, Finding Resources

Get commands with basic output

```
kubectl get services
kubectl get pods --all-namespaces
kubectl get pods -o wide
kubectl get deployment my-dep
kubectl get pods --include-uninitialized
```



```
# List all services in the namespace
# List all pods in all namespaces
# List all pods in the namespace, with wide output
# List a particular deployment
# List all pods in the namespace, including uninitialized
```

Describe commands with verbose output

```
kubectl describe nodes my-node
kubectl describe pods my-pod
```

```
kubectl get services --sort-by=.metadata.name # List Services Sorted by Name
```

List pods Sorted by Restart Count

```
kubectl get pods --sort-by='.status.containerStatuses[0].restartCount'
```

Get the version label of all pods with label app=cassandra

```
kubectl get pods --selector=app=cassandra rc -o \
  jsonpath='{.items[*].metadata.labels.version}'
```

Get all running pods in the namespace

```
kubectl get pods --field-selector=status.phase=Running
```

Get ExternalIPs of all nodes

```
kubectl get nodes -o jsonpath='{.items[*].status.addresses[?(@.type=="ExternalIP")]'}
```

List Names of Pods that belong to Particular RC

```
# "jq" command useful for transformations that are too complex for jsonpath, it can be used to select specific fields from the output of kubectl
sel=${$(kubectl get rc my-rc --output=json | jq -j '.spec.selector | to_entries | .[] | select(.value=="my-rc") | .key'}
echo $(kubectl get pods --selector=$sel --output=jsonpath='{.items..metadata.name}')
```

Check which nodes are ready

```
JSONPATH='{range .items[*]}{@.metadata.name}:{range @.status.conditions[*]}{@.type}{"Ready=True"}'
&& kubectl get nodes -o jsonpath="$JSONPATH" | grep "Ready=True"
```

List all Secrets currently in use by a pod

```
kubectl get pods -o json | jq '.items[].spec.containers[].env[]?.valueFrom.secretKeyRef'
```

List Events sorted by timestamp

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

Updating Resources

```

kubect1 rolling-update frontend-v1 -f frontend-v2.json # Rolling update
kubect1 rolling-update frontend-v1 frontend-v2 --image=image:v2 # Change the name
kubect1 rolling-update frontend --image=image:v2 # Update the pods
kubect1 rolling-update frontend-v1 frontend-v2 --rollback # Abort existing
cat pod.json | kubect1 replace -f - # Replace a pod

# Force replace, delete and then re-create the resource. Will cause a service outage
kubect1 replace --force -f ./pod.json

# Create a service for a replicated nginx, which serves on port 80 and connects to
kubect1 expose rc nginx --port=80 --target-port=8000

# Update a single-container pod's image version (tag) to v4
kubect1 get pod mypod -o yaml | sed 's/\(image: myimage\):.*$/\1:v4/' | kubect1 re

kubect1 label pods my-pod new-label=awesome # Add a Label
kubect1 annotate pods my-pod icon-url=http://goo.gl/XXBTWq # Add an annotation
kubect1 autoscale deployment foo --min=2 --max=10 # Auto scale a deployment

```

Patching Resources

```

kubect1 patch node k8s-node-1 -p '{"spec":{"unschedulable":true}}' # Partially update
# Update a container's image; spec.containers[*].name is required because it's a name
kubect1 patch pod valid-pod -p '{"spec":{"containers":[{"name":"kubernetes-serve-h

# Update a container's image using a json patch with positional arrays
kubect1 patch pod valid-pod --type='json' -p='[{"op": "replace", "path": "/spec/co

# Disable a deployment livenessProbe using a json patch with positional arrays
kubect1 patch deployment valid-deployment --type json -p='[{"op": "remove", "pa

# Add a new element to a positional array
kubect1 patch sa default --type='json' -p='[{"op": "add", "path": "/secrets/1", "v

```

Editing Resources

The edit any API resource in an editor.

```

kubect1 edit svc/docker-registry # Edit the service named docker-registry
KUBE_EDITOR="nano" kubect1 edit svc/docker-registry # Use an alternative editor

```

Scaling Resources



```
kubectl scale --replicas=3 rs/foo # Scale a replic
kubectl scale --replicas=3 -f foo.yaml # Scale a resour
kubectl scale --current-replicas=2 --replicas=3 deployment/mysql # If the deploym
kubectl scale --replicas=5 rc/foo rc/bar rc/baz # Scale multiple
```

Deleting Resources

```
kubectl delete -f ./pod.json # Delete
kubectl delete pod,service baz foo # Delete
kubectl delete pods,services -l name=myLabel # Delete
kubectl delete pods,services -l name=myLabel --include-uninitialized # Delete
kubectl -n my-ns delete po,svc --all # Delete
```

Interacting with running Pods

```
kubectl logs my-pod # dump pod logs (stdout)
kubectl logs my-pod --previous # dump pod logs (stdout) for a
kubectl logs my-pod -c my-container # dump pod container logs (std
kubectl logs my-pod -c my-container --previous # dump pod container logs (std
kubectl logs -f my-pod # stream pod logs (stdout)
kubectl logs -f my-pod -c my-container # stream pod container logs (s
kubectl run -i --tty busybox --image=busybox -- sh # Run pod as interactive shell
kubectl attach my-pod -i # Attach to Running Container
kubectl port-forward my-pod 5000:6000 # Listen on port 5000 on the l
kubectl exec my-pod -- ls / # Run command in existing pod
kubectl exec my-pod -c my-container -- ls / # Run command in existing pod
kubectl top pod POD_NAME --containers # Show metrics for a given pod
```

Interacting with Nodes and Cluster

```

kubect1 cordon my-node
kubect1 drain my-node
kubect1 uncordon my-node
kubect1 top node my-node
kubect1 cluster-info
kubect1 cluster-info dump
kubect1 cluster-info dump --output-directory=/path/to/cluster-state

# If a taint with that key and effect already exists, its value is replaced as spe
kubect1 taint nodes foo dedicated=special-user:NoSchedule

```



Resource types

List all supported resource types along with their shortnames, [API group](#), whether they are [namespaced](#), and [Kind](#):

```
kubect1 api-resources
```

Other operations for exploring API resources:

```

kubect1 api-resources --namespaced=true      # All namespaced resources
kubect1 api-resources --namespaced=false    # All non-namespaced resources
kubect1 api-resources -o name                # All resources with simple output (j
kubect1 api-resources -o wide                # All resources with expanded (aka "w
kubect1 api-resources --verbs=list,get      # All resources that support the "lis
kubect1 api-resources --api-group=extensions # All resources in the "extensions" A

```

Formatting output

To output details to your terminal window in a specific format, you can add either the `-o` or `-output` flags to a supported `kubect1` command.

Output format	Description
<code>-o=custom-columns=<spec></code>	Print a table using a comma separated list of custom columns
<code>-o=custom-columns-file=<filename></code>	Print a table using the custom columns template in the <code><filename></code> file
<code>-o=json</code>	Output a JSON formatted API object

Output format	Description
<code>-o=jsonpath=<template></code>	Print the fields defined in a jsonpath expression
<code>-o=jsonpath-file=<filename></code>	Print the fields defined by the jsonpath expression in the <code><filename></code> file
<code>-o=name</code>	Print only the resource name and nothing else
<code>-o=wide</code>	Output in the plain-text format with any additional information, and for pods, the node name is included
<code>-o=yaml</code>	Output a YAML formatted API object

Kubectl output verbosity and debugging

Kubectl verbosity is controlled with the `-v` or `--v` flags followed by an integer representing the log level. General Kubernetes logging conventions and the associated log levels are described [here](#).

Verbosity	Description
<code>--v=0</code>	Generally useful for this to ALWAYS be visible to an operator.
<code>--v=1</code>	A reasonable default log level if you don't want verbosity.
<code>--v=2</code>	Useful steady state information about the service and important log messages that may correlate to significant changes in the system. This is the recommended default log level for most systems.
<code>--v=3</code>	Extended information about changes.
<code>--v=4</code>	Debug level verbosity.
<code>--v=6</code>	Display requested resources.
<code>--v=7</code>	Display HTTP request headers.
<code>--v=8</code>	Display HTTP request contents.
<code>--v=9</code>	Display HTTP request contents without truncation of contents.

What's next

- Learn more about [Overview of kubectl](#).
- See [kubectl](#) options.