Rocky Enterprise Linux 9.2 Manual Pages on command 'podman-kube-play.1'

$ man podman-kube-play.1

podman-kube-play(1) General Commands Manual podman-kube-play(1)

NAME

podman-kube-play - Create containers, pods and volumes based on Kuber?

netes YAML

SYNOPSIS

podman kube play [options] file.yml|-|https://website.io/file.yml

DESCRIPTION

podman kube play will read in a structured file of Kubernetes YAML. It

will then recreate the containers, pods or volumes described in the

YAML. Containers within a pod are then started and the ID of the new

Pod or the name of the new Volume is output. If the yaml file is speci?

fied as "-" then podman kube play will read the YAML file from stdin.

Using the --down command line option, it is also capable of tearing

down the pods created by a previous run of podman kube play. Using the

--replace command line option, it will tear down the pods(if any) cre?

ated by a previous run of podman kube play and recreate the pods with

the Kubernetes YAML file. Ideally the input file would be one created

by Podman (see podman-kube-generate(1)). This would guarantee a smooth

import and expected results. The input can also be a URL that points to a YAML file such as https://podman.io/demo.yml. podman kube play will read the YAML from the URL and create pods and containers from it. Currently, the supported Kubernetes kinds are: - Pod - Deployment - PersistentVolumeClaim - ConfigMap

Kubernetes Pods or Deployments

Only three volume types are supported by kube play, the hostPath, emp? tyDir, and persistentVolumeClaim volume types. For the hostPath volume type, only the default (empty), DirectoryOrCreate, Directory, FileOr? Create, File, Socket, CharDevice and BlockDevice subtypes are sup? ported. Podman interprets the value of hostPath path as a file path when it contains at least one forward slash, otherwise Podman treats the value as the name of a named volume. When using a persistentVolume? Claim, the value for claimName is the name for the Podman named volume. When using an emptyDir volume, podman creates an anonymous volume that is attached the containers running inside the pod and is deleted once the pod is removed.

Note: The default restart policy for containers is always. You can change the default by setting the restartPolicy field in the spec.

Note: When playing a kube YAML with init containers, the init container will be created with init type value once. To change the default type, use the io.podman.annotations.init.container.type annotation to set the type to always.

Note: hostPath volume types created by kube play will be given an SELinux shared label (z), bind mounts are not relabeled (use chcon -t container_file_t -R <directory>).

Note: If the :latest tag is used, Podman will attempt to pull the image from a registry. If the image was built locally with Podman or Buildah, it will have localhost as the domain, in that case, Podman will use the image from the local store even if it has the :latest tag.

Note: The command podman play kube is an alias of podman kube play, and will perform the same function.

Note: The command podman kube down can be used to stop and remove pods

or containers based on the same Kubernetes YAML used by podman kube play to create them.

Kubernetes PersistentVolumeClaims

A Kubernetes PersistentVolumeClaim represents a Podman named volume. Only the PersistentVolumeClaim name is required by Podman to create a volume. Kubernetes annotations can be used to make use of the available options for Podman volumes.

- ? volume.podman.io/driver
- ? volume.podman.io/device
- ? volume.podman.io/type
- ? volume.podman.io/uid
- ? volume.podman.io/gid
- ? volume.podman.io/mount-options
- ? volume.podman.io/import-source

Use volume.podman.io/import-source to import the contents of the tar? ball (.tar, .tar.gz, .tgz, .bzip, .tar.xz, .txz) specified in the anno? tation's value into the created Podman volume

Kube play is capable of building images on the fly given the correct directory layout and Containerfiles. This option is not available for remote clients, including Mac and Windows (excluding WSL2) machines, yet. Consider the following excerpt from a YAML file:

```
apiVersion: v1
kind: Pod
metadata:
spec:
  containers:
  - command:
    - top
    - name: container
      value: podman
    image: foobar
```

If there is a directory named foobar in the current working directory with a file named Containerfile or Dockerfile, Podman kube play will

build that image and name it foobar. An example directory structure for this example would look like:

```
|- mykubefiles
   |- myplayfile.yaml
   |- foobar
      |- Containerfile
```

The build will consider foobar to be the context directory for the build. If there is an image in local storage called foobar, the image will not be built unless the --build flag is used. Use --build=false to completely disable builds.

Kubernetes ConfigMap

Kubernetes ConfigMap can be referred as a source of environment variables or volumes in Pods or Deployments. ConfigMaps aren't a standalone object in Podman; instead, when a container uses a ConfigMap, Podman will create environment variables or volumes as needed.

For example, the following YAML document defines a ConfigMap and then uses it in a Pod:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: foo
data:
   FOO: bar
---
apiVersion: v1
kind: Pod
metadata:
  name: foobar
spec:
  containers:
  - command:
    - top
    name: container-1
```

```
        image: foobar
        envFrom:
        - configMapRef:
            name: foo
            optional: false
```

and as a result environment variable FOO will be set to bar for con?
tainer container-1.

OPTIONS

--annotation=key=value

Add an annotation to the container or pod. This option can be set mul?
tiple times.

--authfile=path

Path of the authentication file. Default is ${XDG_RUNTIME_DIR}/contain?
ers/auth.json, which is set using podman login. If the authorization
state is not found there, $HOME/.docker/config.json is checked, which
is set using docker login.

Note: There is also the option to override the default path of the au?
thentication file by setting the REGISTRY_AUTH_FILE environment vari?
able. This can be done with export REGISTRY_AUTH_FILE=path.

--build

Build images even if they are found in the local storage. Use
--build=false to completely disable builds. (This option is not avail?
able with the remote Podman client)

--cert-dir=path

Use certificates at path (*.crt, *.cert, *.key) to connect to the reg?
istry. (Default: /etc/containers/certs.d) Please refer to containers-
certs.d(5) for details. (This option is not available with the remote
Podman client, including Mac and Windows (excluding WSL2) machines)

--configmap=path

Use Kubernetes configmap YAML at path to provide a source for environ?
ment variable values within the containers of the pod. (This option is
not available with the remote Podman client)

Note: The --configmap option can be used multiple times or a comma-sep?

arated list of paths can be used to pass multiple Kubernetes configmap YAMLs.

--context-dir=path

Use path as the build context directory for each image. Requires --build option be true. (This option is not available with the remote Podman client)

--creds=[username[:password]]

The [username[:password]] to use to authenticate with the registry, if required. If one or both values are not supplied, a command line prompt will appear and the value can be entered. The password is entered without echo.

--force

Tear down the volumes linked to the PersistentVolumeClaims as part of --down

--help, -h

Print usage statement

--ip=IP address

Assign a static ip address to the pod. This option can be specified several times when kube play creates more than one pod. Note: When joining multiple networks use the --network name:ip=<ip> syntax.

--log-driver=driver

Set logging driver for all created containers.

--log-opt=name=value

Logging driver specific options.

Set custom logging configuration. The following *name*s are supported:

path: specify a path to the log file

(e.g. --log-opt path=/var/log/container/mycontainer.json);

max-size: specify a max size of the log file

(e.g. --log-opt max-size=10mb);

tag: specify a custom log tag for the container

(e.g. --log-opt tag="{{.ImageName}}". It supports the same keys as podman inspect --format. This option is currently supported only by the journald log driver.

--mac-address=MAC address

Assign a static mac address to the pod. This option can be specified several times when kube play creates more than one pod. Note: When joining multiple networks use the --network name:mac=<mac> syntax.

--network=mode, --net

Set the network mode for the pod.

Valid mode values are:

? bridge[:OPTIONS,...]: Create a network stack on the default bridge. This is the default for rootful containers. It is pos? sible to specify these additional options:

? alias=name: Add network-scoped alias for the container.

? ip=IPv4: Specify a static ipv4 address for this container.

? ip=IPv6: Specify a static ipv6 address for this container.

? mac=MAC: Specify a static mac address for this container.

? interface_name: Specify a name for the created network in? terface inside the container.

For example to set a static ipv4 address and a static mac ad? dress, use --network bridge:ip=10.88.0.10,mac=44:33:22:11:00:99.

? <network name or ID>[:OPTIONS,...]: Connect to a user-defined network; this is the network name or ID from a network created by podman network create. Using the network name implies the bridge network mode. It is possible to specify the same op? tions described under the bridge mode above. Use the --network option multiple times to specify additional networks.

? none: Create a network namespace for the container but do not configure network interfaces for it, thus the container has no network connectivity.

? container:id: Reuse another container's network stack.

? host: Do not create a network namespace, the container will use the host's network. Note: The host mode gives the con? tainer full access to local system services such as D-bus and is therefore considered insecure.

? ns:path: Path to a network namespace to join.

? private: Create a new namespace for the container. This will use the bridge mode for rootful containers and slirp4netns for rootless ones.

? slirp4netns[:OPTIONS,...]: use slirp4netns(1) to create a user network stack. This is the default for rootless containers. It is possible to specify these additional options, they can also be set with network_cmd_options in containers.conf:

? allow_host_loopback=true|false: Allow slirp4netns to reach the host loopback IP (default is 10.0.2.2 or the second IP from slirp4netns cidr subnet when changed, see the cidr op? tion below). The default is false.

? mtu=MTU: Specify the MTU to use for this network. (Default is 65520).

? cidr=CIDR: Specify ip range to use for this network. (De? fault is 10.0.2.0/24).

? enable_ipv6=true|false: Enable IPv6. Default is true. (Re? quired for outbound_addr6).

? outbound_addr=INTERFACE: Specify the outbound interface slirp should bind to (ipv4 traffic only).

? outbound_addr=IPv4: Specify the outbound ipv4 address slirp should bind to.

? outbound_addr6=INTERFACE: Specify the outbound interface slirp should bind to (ipv6 traffic only).

? outbound_addr6=IPv6: Specify the outbound ipv6 address slirp should bind to.

? port_handler=rootlesskit: Use rootlesskit for port forward? ing. Default. Note: Rootlesskit changes the source IP ad? dress of incoming packets to an IP address in the container network namespace, usually 10.0.2.100. If the application requires the real source IP address, e.g. web server logs, use the slirp4netns port handler. The rootlesskit port han? dler is also used for rootless containers when connected to user-defined networks.

? port_handler=slirp4netns: Use the slirp4netns port forward?
ing, it is slower than rootlesskit but preserves the correct
source IP address. This port handler cannot be used for
user-defined networks.

? pasta[:OPTIONS,...]: use pasta(1) to create a user-mode net?
working stack.
This is only supported in rootless mode.
By default, IPv4 and IPv6 addresses and routes, as well as the
pod interface name, are copied from the host. If port forward?
ing isn't configured, ports will be forwarded dynamically as
services are bound on either side (init namespace or container
namespace). Port forwarding preserves the original source IP
address. Options described in pasta(1) can be specified as
comma-separated arguments.
In terms of pasta(1) options, --config-net is given by de?
fault, in order to configure networking when the container is
started, and --no-map-gw is also assumed by default, to avoid
direct access from container to host using the gateway ad?
dress. The latter can be overridden by passing --map-gw in the
pasta-specific options (despite not being an actual pasta(1)
option).
Also, -t none and -u none are passed if, respectively, no TCP
or UDP port forwarding from host to container is configured,
to disable automatic port forwarding based on bound ports.
Similarly, -T none and -U none are given to disable the same
functionality from container to host.
Some examples:
? pasta:--map-gw: Allow the container to directly reach the
host using the gateway address.
? pasta:--mtu,1500: Specify a 1500 bytes MTU for the tap in?
terface in the container.
? pasta:--ipv4-only,-a,10.0.2.0,-n,24,-g,10.0.2.2,--dns-for?
ward,10.0.2.3,-m,1500,--no-ndp,--no-dhcpv6,--no-dhcp, equiv?

alent to default slirp4netns(1) options: disable IPv6, as?
sign 10.0.2.0/24 to the tap0 interface in the container,
with gateway 10.0.2.3, enable DNS forwarder reachable at
10.0.2.3, set MTU to 1500 bytes, disable NDP, DHCPv6 and
DHCP support.

? pasta:-I,tap0,--ipv4-only,-a,10.0.2.0,-n,24,-g,10.0.2.2,--dns-
forward,10.0.2.3,--no-ndp,--no-dhcpv6,--no-dhcp, equivalent
to default slirp4netns(1) options with Podman overrides:
same as above, but leave the MTU to 65520 bytes

? pasta:-t,auto,-u,auto,-T,auto,-U,auto: enable automatic port
forwarding based on observed bound ports from both host and
container sides

? pasta:-T,5201: enable forwarding of TCP port 5201 from con?
tainer to host, using the loopback interface instead of the
tap interface for improved performance

NOTE: For backward compatibility reasons, if there is an exist?
ing network named pasta, Podman will use it instead of the pasta
mode."?

When no network option is specified and host network mode is not con?
figured in the YAML file, a new network stack is created and pods are
attached to it making possible pod to pod communication.

--no-hosts

Do not create /etc/hosts for the pod. By default, Podman will manage
/etc/hosts, adding the container's own IP address and any hosts from
--add-host. --no-hosts disables this, and the image's /etc/hosts will
be preserved unmodified.

This option conflicts with host added in the Kubernetes YAML.

--publish=[[ip:][hostPort]:]containerPort[/protocol]

Define or override a port definition in the YAML file.

The lists of ports in the YAML file and the command line are merged.
Matching is done by using the containerPort field. If containerPort
exists in both the YAML file and the option, the latter takes prece?
dence.

--quiet, -q

Suppress output information when pulling images

--replace

Tears down the pods created by a previous run of kube play and recre?

ates the pods. This option is used to keep the existing pods up to date

based upon the Kubernetes YAML.

--seccomp-profile-root=path

Directory path for seccomp profiles (default: "/var/lib/kubelet/sec?

comp"). (This option is not available with the remote Podman client,

including Mac and Windows (excluding WSL2) machines)

--start

Start the pod after creating it, set to false to only create it.

--tls-verify

Require HTTPS and verify certificates when contacting registries (de?

fault: true). If explicitly set to true, TLS verification will be

used. If set to false, TLS verification will not be used. If not

specified, TLS verification will be used unless the target registry is

listed as an insecure registry in containers-registries.conf(5)

--userns=mode

Set the user namespace mode for the container. It defaults to the POD?

MAN_USERNS environment variable. An empty value ("") means user name?

spaces are disabled unless an explicit mapping is set with the --uidmap

and --gidmap options.

This option is incompatible with --gidmap, --uidmap, --subuidname and

--subgidname.

Rootless user --userns=Key mappings:

??

?Key ? Host User ? Container User ?

??

?"" ? $UID ? 0 (Default User ac? ?

? ? ? count mapped to ?

? ? ? root user in con? ?

? ? ? tainer.) ?

??

?keep-id ? $UID ? $UID (Map user ac? ?

? ? ? count to same UID ?

? ? ? within container.) ?

??

?keep-id:uid=200,gid=210 ? $UID ? 200:210 (Map user ?

? ? ? account to speci? ?

? ? ? fied uid, gid value ?

? ? ? within container.) ?

??

?auto ? $UID ? nil (Host User UID ?

? ? ? is not mapped into ?

? ? ? container.) ?

??

?nomap ? $UID ? nil (Host User UID ?

? ? ? is not mapped into ?

? ? ? container.) ?

??

Valid mode values are:

auto[:OPTIONS,...]: automatically create a unique user namespace.

The --userns=auto flag requires that the user name containers be speci?

fied in the /etc/subuid and /etc/subgid files, with an unused range of

subordinate user IDs that Podman containers are allowed to allocate.

See subuid(5).

Example: containers:2147483647:2147483648.

Podman allocates unique ranges of UIDs and GIDs from the containers

subordinate user ids. The size of the ranges is based on the number of

UIDs required in the image. The number of UIDs and GIDs can be overrid?

den with the size option.

The rootless option --userns=keep-id uses all the subuids and subgids

of the user. Using --userns=auto when starting new containers will not

work as long as any containers exist that were started with

--userns=keep-id.

Valid auto options:

? gidmapping=CONTAINER_GID:HOST_GID:SIZE: to force a GID mapping to be present in the user namespace.

? size=SIZE: to specify an explicit size for the automatic user namespace. e.g. --userns=auto:size=8192. If size is not speci? fied, auto will estimate a size for the user namespace.

? uidmapping=CONTAINER_UID:HOST_UID:SIZE: to force a UID mapping to be present in the user namespace.

container:id: join the user namespace of the specified container.

host: run in the user namespace of the caller. The processes running in the container will have the same privileges on the host as any other process launched by the calling user (default).

keep-id: creates a user namespace where the current rootless user's UID:GID are mapped to the same values in the container. This option is not allowed for containers created by the root user.

Valid keep-id options:

? uid=UID: override the UID inside the container that will be used to map the current rootless user to.

? gid=GID: override the GID inside the container that will be used to map the current rootless user to.

nomap: creates a user namespace where the current rootless user's UID:GID are not mapped into the container. This option is not allowed for containers created by the root user.

ns:namespace: run the pod in the given existing user namespace.

EXAMPLES

Recreate the pod and containers as described in a file called demo.yml

```
$ podman kube play demo.yml
52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6
```

Recreate the pod and containers as described in a file demo.yml sent to stdin

```
$ cat demo.yml | podman kube play -
52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6
```

Teardown the pod and containers as described in a file demo.yml

$ podman kube play --down demo.yml

Pods stopped:

52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6

Pods removed:

52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6

Provide configmap-foo.yml and configmap-bar.yml as sources for environ?

ment variables within the containers.

$ podman kube play demo.yml --configmap configmap-foo.yml,configmap-bar.yml

52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6

$ podman kube play demo.yml --configmap configmap-foo.yml --configmap configmap-bar.yml

52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6

Create a pod connected to two networks (called net1 and net2) with a

static ip

$ podman kube play demo.yml --network net1:ip=10.89.1.5 --network net2:ip=10.89.10.10

52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6

Please take into account that networks must be created first using pod?

man-network-create(1).

Create and teardown from a URL pointing to a YAML file

$ podman kube play https://podman.io/demo.yml

52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6

$ podman kube play --down https://podman.io/demo.yml

Pods stopped:

52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6

Pods removed:

52182811df2b1e73f36476003a66ec872101ea59034ac0d4d3a7b40903b955a6

podman kube play --down will not work with a URL if the YAML file the

URL points to has been changed or altered.

Podman Kube Play Support

This document outlines the kube yaml fields that are currently sup?

ported by the podman kube play command.

Note: N/A means that the option cannot be supported in a single-node

Podman environment.

Pod Fields

??

?Field ? Support ?

??

?containers ? ? ?

??

?initContainers ? ? ?

??

?imagePullSecrets ? ?

??

?enableServiceLinks ? ?

??

?os.name ? ?

??

?volumes ? ?

??

?nodeSelector ? N/A ?

??

?nodeName ? N/A ?

??

?affinity.nodeAffinity ? N/A ?

??

?affinity.podAffinity ? N/A ?

??

?affinity.podAntiAffinity ? N/A ?

??

?tolerations.key ? N/A ?

??

?tolerations.operator ? N/A ?

??

?tolerations.effect ? N/A ?

??

?tolerations.tolerationSeconds ? N/A ?

??

?schedulerName ? N/A ?

??????????????????????????????????????

?runtimeClassName ? ?

??????????????????????????????????????

?priorityClassName ? ?

??????????????????????????????????????

?priority ? ?

??????????????????????????????????????

?topologySpreadCon? ? ?

?straints.maxSkew ? ?

??????????????????????????????????????

? ? N/A ?

??????????????????????????????????????

?topologySpreadCon? ? ?

?straints.topologyKey ? ?

??????????????????????????????????????

? ? N/A ?

??????????????????????????????????????

?topologySpreadCon? ? ?

?straints.whenUnsatisfiable ? ?

??????????????????????????????????????

? ? N/A ?

??????????????????????????????????????

?topologySpreadConstraints.la? ? ?

?belSelector ? ?

??????????????????????????????????????

? ? N/A ?

??????????????????????????????????????

?topologySpreadCon? ? ?

?straints.minDomains ? ?

??????????????????????????????????????

? ? N/A ?

??????????????????????????????????????

?restartPolicy ? ? ?

??????????????????????????????????????

?terminationGracePeriod ? ?

??????????????????????????????????????

?activeDeadlineSeconds ? ?

??????????????????????????????????????

?readinessGates.conditionType ? ?

??????????????????????????????????????

?hostname ? ? ?

??????????????????????????????????????

?setHostnameAsFQDN ? ?

??????????????????????????????????????

?subdomain ? ?

??????????????????????????????????????

?hostAliases.hostnames ? ? ?

??????????????????????????????????????

?hostAliases.ip ? ? ?

??????????????????????????????????????

?dnsConfig.nameservers ? ? ?

??????????????????????????????????????

?dnsConfig.options.name ? ? ?

??????????????????????????????????????

?dnsConfig.options.value ? ? ?

??????????????????????????????????????

?dnsConfig.searches ? ? ?

??????????????????????????????????????

?dnsPolicy ? ?

??????????????????????????????????????

?hostNetwork ? ? ?

??????????????????????????????????????

?hostPID ? ? ?

??????????????????????????????????????

?hostIPC ? ?

??

?shareProcessNamespace ? ? ?

??

?serviceAccountName ? ?

??

?automountServiceAccountToken ? ?

??

?securityContext.runAsUser ? ?

??

?securityContext.runAsNonRoot ? ?

??

?securityContext.runAsGroup ? ?

??

?securityContext.supplemental? ? ?

?Groups ? ?

??

? ? ?

??

?securityContext.fsGroup ? ?

??

?securityContext.fs? ? ?

?GroupChangePolicy ? ?

??

? ? ?

??

?securityContext.seccompPro? ? ?

?file.type ? ?

??

? ? ?

??

?securityContext.seccompPro? ? ?

?file.localhostProfile ? ?

??

? ? ?

??????????????????????????????????????

?securityContext.seLinuxOp? ? ?

?tions.level ? ?

??????????????????????????????????????

? ? ?

??????????????????????????????????????

?securityContext.seLinuxOp? ? ?

?tions.role ? ?

??????????????????????????????????????

? ? ?

??????????????????????????????????????

?securityContext.seLinuxOp? ? ?

?tions.type ? ?

??????????????????????????????????????

? ? ?

??????????????????????????????????????

?securityContext.seLinuxOp? ? ?

?tions.user ? ?

??????????????????????????????????????

? ? ?

??????????????????????????????????????

?securityContext.sysctls.name ? ?

??????????????????????????????????????

?securityContext.sysctls.value ? ?

??????????????????????????????????????

?securityContext.windowsOp? ? ?

?tions.gmsaCredentialSpec ? ?

??????????????????????????????????????

? ? ?

??????????????????????????????????????

?securityContext.windowsOp? ? ?

?tions.hostProcess ? ?

??

? ? ?

??

?securityContext.windowsOp? ? ?

?tions.runAsUserName ? ?

??

? ? ?

??

Container Fields

??

?Field ? Support ?

??

?name ? ? ?

??

?image ? ? ?

??

?imagePullPolicy ? ? ?

??

?command ? ? ?

??

?args ? ? ?

??

?workingDir ? ? ?

??

?ports.containerPort ? ? ?

??

?ports.hostIP ? ? ?

??

?ports.hostPort ? ? ?

??

?ports.name ? ? ?

??

?ports.protocol ? ? ?

??

?env.name   ? ?   ?

??

?env.value   ? ?   ?

??

?env.valueFrom.configMapKeyRef.key ?   ?

??

?   ? ?   ?

??

?env.valueFrom.configMapKeyRef.name ? ?   ?

??

?env.valueFrom.configMapKeyRef.op? ?   ?

?tional   ?   ?

??

?   ? ?   ?

??

?env.valueFrom.fieldRef   ? ?   ?

??

?env.valueFrom.resourceFieldRef   ? ?   ?

??

?env.valueFrom.secretKeyRef.key   ? ?   ?

??

?env.valueFrom.secretKeyRef.name   ? ?   ?

??

?env.valueFrom.secretKeyRef.op?   ?   ?

?tional   ?   ?

??

?   ? ?   ?

??

?envFrom.configMapRef.name   ? ?   ?

??

?envFrom.configMapRef.optional   ? ?   ?

??

?envFrom.prefix ? ?

??

?envFrom.secretRef.name ? ? ?

??

?envFrom.secretRef.optional ? ? ?

??

?volumeMounts.mountPath ? ? ?

??

?volumeMounts.name ? ? ?

??

?volumeMounts.mountPropagation ? ?

??

?volumeMounts.readOnly ? ? ?

??

?volumeMounts.subPath ? ?

??

?volumeMounts.subPathExpr ? ?

??

?volumeDevices.devicePath ? ?

??

?volumeDevices.name ? ?

??

?resources.limits ? ? ?

??

?resources.requests ? ? ?

??

?lifecycle.postStart ? ?

??

?lifecycle.preStop ? ?

??

?terminationMessagePath ? ?

??

?terminationMessagePolicy ? ?

??

?livenessProbe ? ? ?

??

?readinessProbe ? ?

??

?startupProbe ? ?

??

?securityContext.runAsUser ? ? ?

??

?securityContext.runAsNonRoot ? ?

??

?securityContext.runAsGroup ? ? ?

??

?securityContext.readOnlyRoot? ? ?

?Filesystem ? ?

??

? ? ? ?

??

?securityContext.procMount ? ?

??

?securityContext.privileged ? ? ?

??

?securityContext.allowPrivi? ? ?

?legeEscalation ? ?

??

? ? ? ?

??

?securityContext.capabilities.add ? ?

??

? ? ? ?

??

?securityContext.capabilities.drop ? ?

??

? ? ? ?

??

?securityContext.seccompPro? ? ?

?file.type ? ?

??

? ? ?

??

?securityContext.seccompProfile.lo? ? ?

?calhostProfile ? ?

??

? ? ?

??

?securityContext.seLinuxOp? ? ?

?tions.level ? ?

??

? ? ? ?

??

?securityContext.seLinuxOp? ? ?

?tions.role ? ?

??

? ? ? ?

??

?securityContext.seLinuxOp? ? ?

?tions.type ? ?

??

? ? ? ?

??

?securityContext.seLinuxOp? ? ?

?tions.user ? ?

??

? ? ? ?

??

?securityContext.windowsOptions.gm? ? ?

?saCredentialSpec ? ?

??

? ? ?

??

?securityContext.windowsOp? ? ?

?tions.hostProcess ? ?

??

? ? ?

??

?securityContext.windowsOp? ? ?

?tions.runAsUserName ? ?

??

? ? ?

??

?stdin ? ?

??

?stdinOnce ? ?

??

?tty ? ?

??

PersistentVolumeClaim Fields

????????????????????????????????

?Field ? Support ?

????????????????????????????????

?volumeName ? ?

????????????????????????????????

?storageClassName ? ? ?

????????????????????????????????

?volumeMode ? ?

????????????????????????????????

?accessModes ? ? ?

????????????????????????????????

?selector ? ?

??????????????????????????????

?resources.limits ? ?

??????????????????????????????

?resources.requests ? ? ?

??????????????????????????????

ConfigMap Fields

??????????????????????

?Field ? Support ?

??????????????????????

?binaryData ? ? ?

??????????????????????

?data ? ? ?

??????????????????????

?immutable ? ?

??????????????????????

Deployment Fields

??????????????????????????????????????

?Field ? Support ?

??????????????????????????????????????

?replicas ? ? ?

??????????????????????????????????????

?selector ? ? ?

??????????????????????????????????????

?template ? ? ?

??????????????????????????????????????

?minReadySeconds ? ?

??????????????????????????????????????

?strategy.type ? ?

??????????????????????????????????????

?strategy.rollingUp? ? ?

?date.maxSurge ? ?

??????????????????????????????????????

? ? ?

??????????????????????????????????

?strategy.rollingUp? ? ?

?date.maxUnavailable ? ?

??????????????????????????????????

? ? ?

??????????????????????????????????

?revisionHistoryLimit ? ?

??????????????????????????????????

?progressDeadlineSeconds ? ?

??????????????????????????????????

?paused ? ?

??????????????????????????????????

SEE ALSO

podman(1), podman-kube(1), podman-kube-down(1), podman-network-cre?

ate(1), podman-kube-generate(1), containers-certs.d(5)

podman-kube-play(1)