



Red Hat Enterprise Linux Release 9.2 Manual Pages on 'podman-pod-create.1' command

\$ man podman-pod-create.1

podman-pod-create(1) General Commands Manual podman-pod-create(1)

NAME

podman-pod-create - Create a new pod

SYNOPSIS

podman pod create [options] [name]

DESCRIPTION

Creates an empty pod, or unit of multiple containers, and prepares it to have containers added to it. The pod can be created with a specific name. If a name is not given a random name is generated. The pod id is printed to STDOUT. You can then use `podman create --pod <pod_id|pod_name> ...` to add containers to the pod, and `podman pod start <pod_id|pod_name>` to start the pod.

The operator can identify a pod in three ways: UUID long identifier
(?f78375b1c487e03c9438c729345e54db9d20cfa2ac1fc3494b6eb60872e74778?)

UUID short identifier (?f78375b1c487?) Name (?jonah?)

podman generates a UUID for each pod, and if a name is not assigned to the container with `--name` then a random string name will be generated for it. This name is useful to identify a pod.

Note: resource limit related flags work by setting the limits explicitly in the pod's cgroup parent for all containers joining the pod. A container can override the resource limits when joining a pod. For example, if a pod was created via `podman pod create --cpus=5`, specifying `podman container create --pod=<pod_id|pod_name> --cpus=4` causes the

container to use the smaller limit. Also, containers which specify their own cgroup, such as `--cgroupns=host`, do NOT get the assigned pod level cgroup resources.

OPTIONS

`--add-host=host:ip`

Add a custom host-to-IP mapping (host:ip)

Add a line to `/etc/hosts`. The format is `hostname:ip`. The `--add-host` option can be set multiple times. Conflicts with the `--no-hosts` option.

The `/etc/hosts` file is shared between all containers in the pod.

`--blkio-weight=weight`

Block IO relative weight. The weight is a value between 10 and 1000.

This option is not supported on cgroups V1 rootless systems.

`--blkio-weight-device=device:weight`

Block IO relative device weight.

`--cgroup-parent=path`

Path to cgroups under which the cgroup for the pod will be created. If the path is not absolute, the path is considered to be relative to the cgroups path of the init process. Cgroups will be created if they do not already exist.

`--cpu-shares, -c=shares`

CPU shares (relative weight).

By default, all containers get the same proportion of CPU cycles. This proportion can be modified by changing the container's CPU share weighting relative to the combined weight of all the running containers. Default weight is 1024.

The proportion will only apply when CPU-intensive processes are running. When tasks in one container are idle, other containers can use the left-over CPU time. The actual amount of CPU time will vary depending on the number of containers running on the system.

For example, consider three containers, one has a `cpu-share` of 1024 and two others have a `cpu-share` setting of 512. When processes in all three containers attempt to use 100% of CPU, the first container would receive 50% of the total CPU time. If a fourth container is added with a

cpu-share of 1024, the first container only gets 33% of the CPU. The remaining containers receive 16.5%, 16.5% and 33% of the CPU.

On a multi-core system, the shares of CPU time are distributed over all CPU cores. Even if a container is limited to less than 100% of CPU time, it can use 100% of each individual CPU core.

For example, consider a system with more than three cores. If the container C0 is started with `--cpu-shares=512` running one process, and another container C1 with `--cpu-shares=1024` running two processes, this can result in the following division of CPU shares:

```
????????????????????????????????????????????????????????????
?PID ? container ? CPU ? CPU share  ?
????????????????????????????????????????????????????????????
?100 ? C0      ? 0 ? 100% of CPU0 ?
????????????????????????????????????????????????????????????
?101 ? C1      ? 1 ? 100% of CPU1 ?
????????????????????????????????????????????????????????????
?102 ? C1      ? 2 ? 100% of CPU2 ?
????????????????????????????????????????????????????????????
```

On some systems, changing the resource limits may not be allowed for non-root users. For more details, see <https://github.com/containers/podman/blob/main/troubleshooting.md#26-running-containers-with-resource-limits-fails-with-a-permissions-error>

This option is not supported on cgroups V1 rootless systems.

`--cpus=amount`

Set the total number of CPUs delegated to the pod. Default is 0.000 which indicates that there is no limit on computation power.

`--cpuset-cpus=number`

CPUs in which to allow execution. Can be specified as a comma-separated list (e.g. 0,1), as a range (e.g. 0-3), or any combination thereof (e.g. 0-3,7,11-15).

On some systems, changing the resource limits may not be allowed for non-root users. For more details, see <https://github.com/containers/podman/blob/main/troubleshooting.md#26-running-containers-with-re?>

source-limits-fails-with-a-permissions-error

This option is not supported on cgroups V1 rootless systems.

`--cpuset-mems=nodes`

Memory nodes (MEMs) in which to allow execution (0-3, 0,1). Only effective on NUMA systems.

If there are four memory nodes on the system (0-3), use `--cpuset-mems=0,1` then processes in the container will only use memory from the first two memory nodes.

On some systems, changing the resource limits may not be allowed for non-root users. For more details, see <https://github.com/containers/podman/blob/main/troubleshooting.md#26-running-containers-with-resource-limits-fails-with-a-permissions-error>

This option is not supported on cgroups V1 rootless systems.

`--device=host-device[:container-device][:permissions]`

Add a host device to the pod. Optional permissions parameter can be used to specify device permissions by combining r for read, w for write, and m for mknod(2).

Example: `--device=/dev/sdc:/dev/xvdc:rwm`.

Note: if host-device is a symbolic link then it will be resolved first.

The pod will only store the major and minor numbers of the host device.

Podman may load kernel modules required for using the specified device.

The devices that Podman will load modules for when necessary are: `/dev/fuse`.

In rootless mode, the new device is bind mounted in the container from the host rather than Podman creating it within the container space. Because the bind mount retains its SELinux label on SELinux systems, the container can get permission denied when accessing the mounted device.

Modify SELinux settings to allow containers to use all device labels via the following command:

```
$ sudo setsebool -P container_use_devices=true
```

Note: the pod implements devices by storing the initial configuration passed by the user and recreating the device on each container added to the pod.

`--device-read-bps=path:rate`

Limit read rate (in bytes per second) from a device (e.g. `--device-read-bps=/dev/sda:1mb`).

On some systems, changing the resource limits may not be allowed for non-root users. For more details, see <https://github.com/containers/podman/blob/main/troubleshooting.md#26-running-containers-with-resource-limits-fails-with-a-permissions-error>

This option is not supported on cgroups V1 rootless systems.

`--device-write-bps=path:rate`

Limit write rate (in bytes per second) to a device (e.g. `--device-write-bps=/dev/sda:1mb`).

On some systems, changing the resource limits may not be allowed for non-root users. For more details, see <https://github.com/containers/podman/blob/main/troubleshooting.md#26-running-containers-with-resource-limits-fails-with-a-permissions-error>

This option is not supported on cgroups V1 rootless systems.

`--dns=ipaddr`

Set custom DNS servers in the `/etc/resolv.conf` file that will be shared between all containers in the pod. A special option, "none" is allowed which disables creation of `/etc/resolv.conf` for the pod.

`--dns-option=option`

Set custom DNS options in the `/etc/resolv.conf` file that will be shared between all containers in the pod.

`--dns-search=domain`

Set custom DNS search domains in the `/etc/resolv.conf` file that will be shared between all containers in the pod.

`--exit-policy=continue | stop`

Set the exit policy of the pod when the last container exits. Supported policies are:

??

?Exit Policy ? Description ?

??

?continue ? The pod continues running, ?

? by keeping its infra container alive, when the last container exits. Used by default.

??

?stop The pod (including its infrastructure container) is stopped when the last container exits. Used in kube play.

??

--gidmap=pod_gid:host_gid:amount

GID map for the user namespace. Using this flag will run all containers in the pod with user namespace enabled. It conflicts with the --users and --subgidname flags.

--help, -h

Print usage statement.

--hostname=name

Set a hostname to the pod.

--infra

Create an infra container and associate it with the pod. An infra container is a lightweight container used to coordinate the shared kernel namespace of a pod. Default: true.

--infra-command=command

The command that will be run to start the infra container. Default: "/pause".

--infra-common-pidfile=file

Write the pid of the infra container's common process to a file. As common runs in a separate process than Podman, this is necessary when using systemd to manage Podman containers and pods.

--infra-image=image

The custom image that will be used for the infra container. Unless specified, Podman builds a custom local image which does not require pulling down an image.

`--infra-name=name`

The name that will be used for the pod's infra container.

`--ip=ipv4`

Specify a static IPv4 address for the pod, for example 10.88.64.128.

This option can only be used if the pod is joined to only a single network

work - i.e., `--network=network-name` is used at most once - and if the

pod is not joining another container's network namespace via `--network?`

`work=container:id`. The address must be within the network's IP address

pool (default 10.88.0.0/16).

To specify multiple static IP addresses per pod, set multiple networks

using the `--network` option with a static IP address specified for each

using the `ip` mode for that option.

`--ip6=ipv6`

Specify a static IPv6 address for the pod, for example

fd46:db93:aa76:ac37::10. This option can only be used if the pod is

joined to only a single network - i.e., `--network=network-name` is used

at most once - and if the pod is not joining another container's network

namespace via `--network=container:id`. The address must be within

the network's IPv6 address pool.

To specify multiple static IPv6 addresses per pod, set multiple networks

using the `--network` option with a static IPv6 address specified

for each using the `ip6` mode for that option.

`--label, -l=key=value`

Add metadata to a pod.

`--label-file=file`

Read in a line-delimited file of labels.

`--mac-address=address`

Pod network interface MAC address (e.g. 92:d0:c6:0a:29:33) This option

can only be used if the pod is joined to only a single network - i.e.,

`--network=network-name` is used at most once - and if the pod is not

joining another container's network namespace via `--network=con?`

`tainer:id`.

Remember that the MAC address in an Ethernet network must be unique.

The IPv6 link-local address will be based on the device's MAC address according to RFC4862.

To specify multiple static MAC addresses per pod, set multiple networks using the `--network` option with a static MAC address specified for each using the `mac` mode for that option.

`--memory, -m=number[unit]`

Memory limit. A unit can be `b` (bytes), `k` (kibibytes), `m` (mebibytes), or `g` (gibibytes).

Allows the memory available to a container to be constrained. If the host supports swap memory, then the `-m` memory setting can be larger than physical RAM. If a limit of 0 is specified (not using `-m`), the container's memory is not limited. The actual limit may be rounded up to a multiple of the operating system's page size (the value would be very large, that's millions of trillions).

This option is not supported on cgroups V1 rootless systems.

`--memory-swap=number[unit]`

A limit value equal to memory plus swap. A unit can be `b` (bytes), `k` (kibibytes), `m` (mebibytes), or `g` (gibibytes).

Must be used with the `-m` (`--memory`) flag. The argument value should always be larger than that of

`-m` (`--memory`) By default, it is set to double the value of `--memory`.

Set number to `-1` to enable unlimited swap.

This option is not supported on cgroups V1 rootless systems.

`--name, -n=name`

Assign a name to the pod.

`--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 possible 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 interface inside the container.

For example to set a static ipv4 address and a static mac address, 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 options 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 container 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` option 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. (Default is 10.0.2.0/24).

? enable_ipv6=true|false: Enable IPv6. Default is true. (Required 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 forwarding. Default. Note: Rootlesskit changes the source IP address 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 handler is also used for rootless containers when connected to user-defined networks.

? port_handler=slirp4netns: Use the slirp4netns port forwarding, 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 networking 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 forwarding 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 default, 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 address. 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 interface in the container.

? `pasta:--ipv4-only,-a,10.0.2.0,-n,24,-g,10.0.2.2,--dns-forward,10.0.2.3,-m,1500,--no-ndp,--no-dhcpv6,--no-dhcp`, equivalent to default `slirp4netns(1)` options: disable IPv6, assign 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:-l,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 container to host, using the loopback interface instead of the tap interface for improved performance

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

Invalid if using --dns, --dns-option, or --dns-search with --network set to none or container:id.

--network-alias=alias

Add a network-scoped alias for the pod, setting the alias for all networks that the container joins. To set a name only for a specific network, use the alias option as described under the --network option. If the network has DNS enabled (podman network inspect -f {{.DNSEnabled}} <name>), these aliases can be used for name resolution on the given network. This option can be specified multiple times. NOTE: When using CNI a pod will only have access to aliases on the first network that it joins. This limitation does not exist with netavark/aardvark-dns.

--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 --add-host.

--pid=pid

Set the PID mode for the pod. The default is to create a private PID namespace for the pod. Requires the PID namespace to be shared via --share.

host: use the host's PID namespace for the pod

ns: join the specified PID namespace

private: create a new namespace for the pod (default)

--pod-id-file=path

Write the pod ID to the file.

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

Publish a container's port, or range of ports, within this pod to the host.

Both hostPort and containerPort can be specified as a range of ports.

When specifying ranges for both, the number of container ports in the range must match the number of host ports in the range.

If host IP is set to 0.0.0.0 or not set at all, the port will be bound on all IPs on the host.

By default, Podman will publish TCP ports. To publish a UDP port instead, give udp as protocol. To publish both TCP and UDP ports, set --publish twice, with tcp, and udp as protocols respectively. Rootful containers can also publish ports using the sctp protocol.

Host port does not have to be specified (e.g. podman run -p 127.0.0.1::80). If it is not, the container port will be randomly assigned a port on the host.

Use podman port to see the actual mapping: podman port \$CONTAINER \$CONTAINERPORT.

Note: You must not publish ports of containers in the pod individually, but only by the pod itself.

Note: This cannot be modified once the pod is created.

--replace

If another pod with the same name already exists, replace and remove it. The default is false.

--security-opt=option

Security Options

? apparmor=unconfined : Turn off apparmor confinement for the pod

? apparmor=alternate-profile : Set the apparmor confinement profile file for the pod

? label=user:USER: Set the label user for the pod processes

? label=role:ROLE: Set the label role for the pod processes

? label=type:TYPE: Set the label process type for the pod processes

? label=level:LEVEL: Set the label level for the pod processes

? label=filetype:TYPE: Set the label file type for the pod files

? label=disable: Turn off label separation for the pod

Note: Labeling can be disabled for all pods/containers by setting label=false in the containers.conf (/etc/containers/containers.conf or \$HOME/.config/containers/containers.conf) file.

? mask=/path/1:/path/2: The paths to mask separated by a colon.

A masked path cannot be accessed inside the containers within the pod.

? no-new-privileges: Disable container processes from gaining additional privileges.

? seccomp=unconfined: Turn off seccomp confinement for the pod.

? seccomp=profile.json: JSON file to be used as a seccomp filter. Note that the io.podman.annotations.seccomp annotation is set with the specified value as shown in podman inspect.

? proc-opts=OPTIONS : Comma-separated list of options to use for the /proc mount. More details for the possible mount options are specified in the proc(5) man page.

? unmask=ALL or /path/1:/path/2, or shell expanded paths (/proc/*): Paths to unmask separated by a colon. If set to ALL, it will unmask all the paths that are masked or made read-only by default. The default masked paths are /proc/acpi, /proc/kcore, /proc/keys, /proc/latency_stats, /proc/sched_debug, /proc/scsi, /proc/timer_list, /proc/timer_stats, /sys/firmware, and /sys/fs/selinux. The default paths that are read-only are /proc/asound, /proc/bus, /proc/fs, /proc/irq, /proc/sys, /proc/sysrq-trigger, /sys/fs/cgroup.

Note: Labeling can be disabled for all containers by setting label=false in the containers.conf(5) file.

--share=namespace

A comma-separated list of kernel namespaces to share. If none or "" is specified, no namespaces will be shared and the infra container will not be created unless explicitly specified via --infra=true. The name?

spaces to choose from are cgroup, ipc, net, pid, uts. If the option is prefixed with a "+" then the namespace is appended to the default list, otherwise it replaces the default list. Defaults matches Kubernetes default (ipc, net, uts)

--share-parent

This boolean determines whether or not all containers entering the pod will use the pod as their cgroup parent. The default value of this flag is true. Use the --share option to share the cgroup namespace rather than a cgroup parent in a pod.

Note: This options conflict with --share=cgroup since that would set the pod as the cgroup parent but enter the container into the same cgroupNS as the infra container.

--shm-size=number[unit]

Size of /dev/shm. A unit can be b (bytes), k (kibibytes), m (mebibytes), or g (gibibytes). If the unit is omitted, the system uses bytes. If the size is omitted, the default is 64m. When size is 0, there is no limit on the amount of memory used for IPC by the pod.

This option conflicts with --ipc=host.

--subgidname=name

Run the container in a new user namespace using the map with name in the /etc/subgid file. If running rootless, the user needs to have the right to use the mapping. See subgid(5). This flag conflicts with --userns and --gidmap.

--subuidname=name

Run the container in a new user namespace using the map with name in the /etc/subuid file. If running rootless, the user needs to have the right to use the mapping. See subuid(5). This flag conflicts with --userns and --uidmap.

--sysctl=name=value

Configure namespaced kernel parameters for all containers in the pod.

For the IPC namespace, the following sysctls are allowed:

? kernel.msgmax

? kernel.msgmnb

- ? kernel.msgmni
- ? kernel.sem
- ? kernel.shmall
- ? kernel.shmmax
- ? kernel.shmmni
- ? kernel.shm_rmid_forced
- ? Sysctls beginning with fs.mqueue.*

Note: if the ipc namespace is not shared within the pod, the above sysctls are not allowed.

For the network namespace, only sysctls beginning with net.* are allowed.

Note: if the network namespace is not shared within the pod, the above sysctls are not allowed.

--uidmap=container_uid:from_uid:amount

Run all containers in the pod in a new user namespace using the supplied mapping. This option conflicts with the --userns and --subuidname options. This option provides a way to map host UIDs to container UIDs. It can be passed several times to map different ranges.

--userns=mode

Set the user namespace mode for all the containers in a pod. It defaults to the PODMAN_USERNS environment variable. An empty value ("") means user namespaces are disabled.

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.) ?

??

?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 namespace. It is possible to specify these options to auto:
- ? 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. --users=auto:size=8192. If size is not specified, 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.
- ? 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.
- ? 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.

--uts=mode

Set the UTS namespace mode for the pod. The following values are sup?

ported:

? host: use the host's UTS namespace inside the pod.

? private: create a new namespace for the pod (default).

? ns:[path]: run the pod in the given existing UTS namespace.

--volume, -v=[[SOURCE-VOLUME|HOST-DIR:]CONTAINER-DIR[:OPTIONS]]

Create a bind mount. If -v /HOST-DIR:/CONTAINER-DIR is specified, Podman bind mounts /HOST-DIR from the host into /CONTAINER-DIR in the Podman container. Similarly, -v SOURCE-VOLUME:/CONTAINER-DIR will mount the named volume from the host into the container. If no such named volume exists, Podman will create one. If no source is given, the volume will be created as an anonymously named volume with a randomly generated name, and will be removed when the pod is removed via the --rm flag or the podman rm --volumes command.

(Note when using the remote client, including Mac and Windows (excluding WSL2) machines, the volumes will be mounted from the remote server, not necessarily the client machine.)

The OPTIONS is a comma-separated list and can be: [1] ?#Footnote1?

? rw|ro

? z|Z

? [O]

? [U]

? [no]copy

? [no]dev

? [no]exec

? [no]suid

? [r]bind

? [r]shared[[r]slave[[r]private[r]unbindable

? idmap[=options]

The CONTAINER-DIR must be an absolute path such as /src/docs. The volume will be mounted into the container at this directory.

If a volume source is specified, it must be a path on the host or the name of a named volume. Host paths are allowed to be absolute or relative; relative paths are resolved relative to the directory Podman is

run in. If the source does not exist, Podman will return an error.

Users must pre-create the source files or directories.

Any source that does not begin with a `.` or `/` will be treated as the name of a named volume. If a volume with that name does not exist, it will be created. Volumes created with names are not anonymous, and they are not removed by the `--rm` option and the `podman rm --volumes` command.

Specify multiple `-v` options to mount one or more volumes into a pod.

Write Protected Volume Mounts

Add `:ro` or `:rw` option to mount a volume in read-only or read-write mode, respectively. By default, the volumes are mounted read-write.

See examples.

Chowning Volume Mounts

By default, Podman does not change the owner and group of source volume directories mounted into containers. If a pod is created in a new user namespace, the UID and GID in the container may correspond to another UID and GID on the host.

The `:U` suffix tells Podman to use the correct host UID and GID based on the UID and GID within the pod, to change recursively the owner and group of the source volume. Chowning walks the file system under the volume and changes the UID/GID on each file, if the volume has thousands of inodes, this process will take a long time, delaying the start of the pod.

Warning use with caution since this will modify the host filesystem.

Labeling Volume Mounts

Labeling systems like SELinux require that proper labels are placed on volume content mounted into a pod. Without a label, the security system might prevent the processes running inside the pod from using the content. By default, Podman does not change the labels set by the OS.

To change a label in the pod context, add either of two suffixes `:z` or `:Z` to the volume mount. These suffixes tell Podman to relabel file objects on the shared volumes. The `z` option tells Podman that two or more pods share the volume content. As a result, Podman labels the content

with a shared content label. Shared volume labels allow all containers to read/write content. The Z option tells Podman to label the content with a private unshared label. Only the current pod can use a private volume. Relabeling walks the file system under the volume and changes the label on each file, if the volume has thousands of inodes, this process will take a long time, delaying the start of the pod. If the volume was previously relabeled with the z option, Podman is optimized to not relabel a second time. If files are moved into the volume, then the labels can be manually change with the `chcon -R container_file_t PATH` command.

Note: Do not relabel system files and directories. Relabeling system content might cause other confined services on the machine to fail. For these types of containers we recommend disabling SELinux separation. The option `--security-opt label=disable` disables SELinux separation for the pod. For example if a user wanted to volume mount their entire home directory into a pod, they need to disable SELinux separation.

```
$ podman pod create --security-opt label=disable -v $HOME:/home/user fedora touch /home/user/file
```

Overlay Volume Mounts

The `:O` flag tells Podman to mount the directory from the host as a temporary storage using the overlay file system. The pod processes can modify content within the mountpoint which is stored in the container storage in a separate directory. In overlay terms, the source directory will be the lower, and the container storage directory will be the upper. Modifications to the mount point are destroyed when the pod finishes executing, similar to a tmpfs mount point being unmounted.

For advanced users, the `overlay` option also supports custom non-volatile `upperdir` and `workdir` for the overlay mount. Custom `upperdir` and `workdir` can be fully managed by the users themselves, and Podman will not remove it on lifecycle completion. Example `:O,upperdir=/some/upper,workdir=/some/work`

Subsequent executions of the container will see the original source directory content, any changes from previous pod executions no longer exist.

ist.

One use case of the overlay mount is sharing the package cache from the host into the container to allow speeding up builds.

Note:

- The ``O`` flag conflicts with other options listed above.

Content mounted into the container is labeled with the private label.

On SELinux systems, labels in the source directory must be readable by the pod infra container label. Usually containers can read/execute `container_share_t` and can read/write `container_file_t`. If unable to change the labels on a source volume, SELinux container separation must be disabled for the pod or infra container to work.

- The source directory mounted into the pod with an overlay mount should not be modified, it can cause unexpected failures. It is recommended to not modify the directory until the container finishes running.

Mounts propagation

By default bind mounted volumes are private. That means any mounts done inside the pod will not be visible on host and vice versa. One can change this behavior by specifying a volume mount propagation property. Making a volume shared mounts done under that volume inside the pod will be visible on host and vice versa. Making a volume slave enables only one way mount propagation and that is mounts done on host under that volume will be visible inside container but not the other way around. [1] [Footnote1?](#)

To control mount propagation property of a volume one can use the `[r]shared`, `[r]slave`, `[r]private` or the `[r]unbindable` propagation flag.

Propagation property can be specified only for bind mounted volumes and not for internal volumes or named volumes. For mount propagation to work the source mount point (the mount point where source dir is mounted on) has to have the right propagation properties. For shared volumes, the source mount point has to be shared. And for slave volumes, the source mount point has to be either shared or slave. [1] [Footnote1?](#)

To recursively mount a volume and all of its submounts into a pod, use the `rbind` option. By default the `bind` option is used, and submounts of the source directory will not be mounted into the pod.

Mounting the volume with a `copy` option tells `podman` to copy content from the underlying destination directory onto newly created internal volumes. The copy only happens on the initial creation of the volume. Content is not copied up when the volume is subsequently used on different containers. The `copy` option is ignored on `bind` mounts and has no effect.

Mounting the volume with the `nosuid` options means that SUID applications on the volume will not be able to change their privilege. By default volumes are mounted with `nosuid`.

Mounting the volume with the `noexec` option means that no executables on the volume will be able to be executed within the pod.

Mounting the volume with the `nodev` option means that no devices on the volume will be able to be used by processes within the pod. By default volumes are mounted with `nodev`.

If the `HOST-DIR` is a mount point, then `dev`, `suid`, and `exec` options are ignored by the kernel.

Use `df HOST-DIR` to figure out the source mount, then use `findmnt -o TARGET,PROPAGATION source-mount-dir` to figure out propagation properties of source mount. If `findmnt(1)` utility is not available, then one can look at the mount entry for the source mount point in `/proc/self/mountinfo`. Look at the "optional fields" and see if any propagation properties are specified. In there, `shared:N` means the mount is shared, `master:N` means mount is slave, and if nothing is there, the mount is private. [1] ^{Footnote1?}

To change propagation properties of a mount point, use `mount(8)` command. For example, if one wants to bind mount source directory `/foo`, one can do `mount --bind /foo /foo` and `mount --make-private --make-shared /foo`. This will convert `/foo` into a shared mount point. Alternatively, one can directly change propagation properties of source mount. Say `/` is source mount for `/foo`, then use `mount --make-shared /` to con?

vert / into a shared mount.

Note: if the user only has access rights via a group, accessing the volume from inside a rootless pod will fail.

Idmapped mount

If `idmap` is specified, create an idmapped mount to the target user namespace in the container. The `idmap` option supports a custom mapping that can be different than the user namespace used by the container.

The mapping can be specified after the `idmap` option like:

`idmap=uids=0-1-10#10-11-10;gids=0-100-10`. For each triplet, the first value is the start of the backing file system IDs that are mapped to the second value on the host. The length of this mapping is given in the third value. Multiple ranges are separated with `#`.

`--volumes-from=CONTAINER[:OPTIONS]`

Mount volumes from the specified container(s). Used to share volumes between containers and pods. The options is a comma-separated list with the following available elements:

? rw|ro

? z

Mounts already mounted volumes from a source container onto another pod. `CONTAINER` may be a name or ID. To share a volume, use the `--volumes-from` option when running the target container. Volumes can be shared even if the source container is not running.

By default, Podman mounts the volumes in the same mode (read-write or read-only) as it is mounted in the source container. This can be changed by adding a `ro` or `rw` option.

Labeling systems like SELinux require that proper labels are placed on volume content mounted into a pod. Without a label, the security system might prevent the processes running inside the container from using the content. By default, Podman does not change the labels set by the OS.

To change a label in the pod context, add `z` to the volume mount. This suffix tells Podman to relabel file objects on the shared volumes. The `z` option tells Podman that two entities share the volume content. As a result, Podman labels the content with a shared content label. Shared

volume labels allow all containers to read/write content.

If the location of the volume from the source container overlaps with data residing on a target pod, then the volume hides that data on the target.

EXAMPLES

```
$ podman pod create --name test
```

```
$ podman pod create mypod
```

```
$ podman pod create --infra=false
```

```
$ podman pod create --infra-command /top toppod
```

```
$ podman pod create --publish 8443:443
```

```
$ podman pod create --network slirp4netns:outbound_addr=127.0.0.1,allow_host_loopback=true
```

```
$ podman pod create --network slirp4netns:cidr=192.168.0.0/24
```

```
$ podman pod create --network pasta
```

```
$ podman pod create --network net1:ip=10.89.1.5 --network net2:ip=10.89.10.10
```

SEE ALSO

podman(1), podman-pod(1), podman-kube-play(1), containers.conf(1),
cgroups(7)

HISTORY

July 2018, Originally compiled by Peter Hunt pehunt@redhat.com

[?mailto:pehunt@redhat.com?](mailto:pehunt@redhat.com)

FOOTNOTES

1: The Podman project is committed to inclusivity, a core value of open source. The master and slave mount propagation terminology used here is problematic and divisive, and should be changed. However, these terms are currently used within the Linux kernel and must be used as-is at this time. When the kernel maintainers rectify this usage, Podman will follow suit immediately.

podman-pod-create(1)