

dockerd

By Docker Inc

Published: 2026-03-23 · Archived: 2026-04-06 03:35:06 UTC

Options with [] may be specified multiple times.

`dockerd` is the persistent process that manages containers. Docker uses different binaries for the daemon and client. To run the daemon you type `dockerd` .

To run the daemon with debug output, use `dockerd --debug` or add `"debug": true` to [the daemon.json file](#).

Enabling experimental features

Enable experimental features by starting `dockerd` with the `--experimental` flag or adding `"experimental": true` to the `daemon.json` file.

Environment variables

The following list of environment variables are supported by the `dockerd` daemon. Some of these environment variables are supported both by the Docker Daemon and the `docker` CLI. Refer to [Environment variables](#) to learn about environment variables supported by the `docker` CLI.

Variable	Description
<code>DOCKER_CERT_PATH</code>	Location of your authentication keys. This variable is used both by the docker CLI and the <code>dockerd</code> daemon.
<code>DOCKER_DRIVER</code>	The storage driver to use.
<code>DOCKER_RAMDISK</code>	If set this disables <code>pivot_root</code> .
<code>DOCKER_TLS_VERIFY</code>	When set Docker uses TLS and verifies the remote. This variable is used both by the docker CLI and the <code>dockerd</code> daemon.
<code>DOCKER_TMPDIR</code>	Location for temporary files created by the daemon.
<code>HTTP_PROXY</code>	Proxy URL for HTTP requests unless overridden by NoProxy. See the Go specification for details.
<code>HTTPS_PROXY</code>	Proxy URL for HTTPS requests unless overridden by NoProxy. See the Go specification for details.
<code>MOBY_DISABLE_PIGZ</code>	Disables the use of <code>unpigz</code> to decompress layers in parallel when pulling images, even if it is installed.

Variable	Description
<code>NO_PROXY</code>	Comma-separated values specifying hosts that should be excluded from proxying. See the Go specification for details.

[Proxy configuration](#)

If you are behind an HTTP proxy server, for example in corporate settings, you may have to configure the Docker daemon to use the proxy server for operations such as pulling and pushing images. The daemon can be configured in three ways:

1. Using environment variables (`HTTP_PROXY` , `HTTPS_PROXY` , and `NO_PROXY`).
2. Using the `http-proxy` , `https-proxy` , and `no-proxy` fields in the [daemon configuration file](#) (Docker Engine version 23.0 or later).
3. Using the `--http-proxy` , `--https-proxy` , and `--no-proxy` command-line options. (Docker Engine version 23.0 or later).

The command-line and configuration file options take precedence over environment variables. Refer to [control and configure Docker with systemd](#) to set these environment variables on a host using `systemd` .

[Daemon socket option](#)

The Docker daemon can listen for [Docker Engine API](#) requests via three different types of Socket: `unix` , `tcp` , and `fd` .

By default, a `unix` domain socket (or IPC socket) is created at `/var/run/docker.sock` , requiring either `root` permission, or `docker` group membership.

If you need to access the Docker daemon remotely, you need to enable the `tcp` Socket. When using a TCP socket, the Docker daemon provides un-encrypted and un-authenticated direct access to the Docker daemon by default. You should secure the daemon either using the [built in HTTPS encrypted socket](#), or by putting a secure web proxy in front of it. You can listen on port `2375` on all network interfaces with `-H tcp://0.0.0.0:2375` , or on a particular network interface using its IP address: `-H tcp://192.168.59.103:2375` . It is conventional to use port `2375` for un-encrypted, and port `2376` for encrypted communication with the daemon.

If you're using an HTTPS encrypted socket, keep in mind that only TLS version 1.0 and higher is supported. Protocols SSLv3 and below are not supported for security reasons.

On `systemd` based systems, you can communicate with the daemon via [systemd socket activation](#), with `dockerd -H fd://` . Using `fd://` works for most setups, but you can also specify individual sockets: `dockerd -H fd:///3` . If the specified socket activated files aren't found, the daemon exits. You can find examples of using `systemd` socket activation with Docker and `systemd` in the [Docker source tree](#).

You can configure the Docker daemon to listen to multiple sockets at the same time using multiple `-H` options:

The example below runs the daemon listening on the default Unix socket, and on 2 specific IP addresses on this host:

The Docker client honors the `DOCKER_HOST` environment variable to set the `-H` flag for the client. Use **one** of the following commands:

Setting the `DOCKER_TLS_VERIFY` environment variable to any value other than the empty string is equivalent to setting the `--tlsverify` flag. The following are equivalent:

The Docker client honors the `HTTP_PROXY`, `HTTPS_PROXY`, and `NO_PROXY` environment variables (or the lowercase versions thereof). `HTTPS_PROXY` takes precedence over `HTTP_PROXY`.

The Docker client supports connecting to a remote daemon via SSH:

To use SSH connection, you need to set up `ssh` so that it can reach the remote host with public key authentication. Password authentication is not supported. If your key is protected with passphrase, you need to set up `ssh-agent`.

Bind Docker to another host/port or a Unix socket

Changing the default `docker` daemon binding to a TCP port or Unix `docker` user group introduces security risks, as it may allow non-root users to gain root access on the host. Make sure you control access to `docker`. If you are binding to a TCP port, anyone with access to that port has full Docker access; so it's not advisable on an open network.

With `-H` it's possible to make the Docker daemon to listen on a specific IP and port. By default, it listens on `unix:///var/run/docker.sock` to allow only local connections by the root user. You could set it to `0.0.0.0:2375` or a specific host IP to give access to everybody, but that isn't recommended because someone could gain root access to the host where the daemon is running.

Similarly, the Docker client can use `-H` to connect to a custom port. The Docker client defaults to connecting to `unix:///var/run/docker.sock` on Linux, and `tcp://127.0.0.1:2376` on Windows.

`-H` accepts host and port assignment in the following format:

For example:

- `tcp://` -> TCP connection to `127.0.0.1` on either port `2376` when TLS encryption is on, or port `2375` when communication is in plain text.
- `tcp://host:2375` -> TCP connection on host:2375
- `tcp://host:2375/path` -> TCP connection on host:2375 and prepend path to all requests
- `unix://path/to/socket` -> Unix socket located at `path/to/socket`

`-H`, when empty, defaults to the same value as when no `-H` was passed in.

`-H` also accepts short form for TCP bindings: `host:` or `host:port` or `:port`

Run Docker in daemon mode:

Download an `ubuntu` image:

You can use multiple `-H` , for example, if you want to listen on both TCP and a Unix socket

Daemon storage-driver

On Linux, the Docker daemon has support for several different image layer storage drivers: `overlay2` , `fuse-overlayfs` , `btrfs` , and `zfs` .

`overlay2` is the preferred storage driver for all currently supported Linux distributions, and is selected by default. Unless users have a strong reason to prefer another storage driver, `overlay2` should be used.

You can find out more about storage drivers and how to select one in [Select a storage driver](#).

On Windows, the Docker daemon only supports the `windowsfilter` storage driver.

Options per storage driver

Particular storage-driver can be configured with options specified with `--storage-opt` flags. Options for `zfs` start with `zfs` , and options for `btrfs` start with `btrfs` .

ZFS options

`zfs.fsname`

Specifies the ZFS filesystem that the daemon should use to create its datasets. By default, the ZFS filesystem in `/var/lib/docker` is used.

[Example](#)

Btrfs options

`btrfs.min_space`

Specifies the minimum size to use when creating the subvolume which is used for containers. If user uses disk quota for btrfs when creating or running a container with `--storage-opt size` option, Docker should ensure the **size** can't be smaller than **btrfs.min_space**.

[Example](#)

Overlay2 options

`overlay2.size`

Sets the default max size of the container. It is supported only when the backing filesystem is `trfs` and mounted with `quota` mount option. Under these conditions the user can pass any size less than the backing filesystem size.

[Example](#)

[Windowsfilter options](#)

`size`

Specifies the size to use when creating the sandbox which is used for containers. Defaults to 20G.

[Example](#)

[Runtime options](#)

The Docker daemon relies on a [OCI](#) compliant runtime (invoked via the `containerd` daemon) as its interface to the Linux kernel `namespaces`, `cgroups`, and `SELinux`.

[Configure container runtimes](#)

By default, the Docker daemon uses runc as a container runtime. You can configure the daemon to add additional runtimes.

`containerd` shims installed on `PATH` can be used directly, without the need to edit the daemon's configuration. For example, if you install the Kata Containers shim (`containerd-shim-kata-v2`) on `PATH`, then you can select that runtime with `docker run` without having to edit the daemon's configuration:

Container runtimes that don't implement `containerd` shims, or `containerd` shims installed outside of `PATH`, must be registered with the daemon, either via the configuration file or using the `--add-runtime` command line flag.

For examples on how to use other container runtimes, see [Alternative container runtimes](#)

[Configure runtimes using `daemon.json`](#)

To register and configure container runtimes using the daemon's configuration file, add the runtimes as entries under `runtimes`:

The key of the entry (`<runtime>` in the previous example) represents the name of the runtime. This is the name that you reference when you run a container, using `docker run --runtime <runtime>`.

The runtime entry contains an object specifying the configuration for your runtime. The properties of the object depends on what kind of runtime you're looking to register:

- If the runtime implements its own `containerd` shim, the object shall contain a `runtimeType` field and an optional `options` field.

See [Configure shims](#).

- If the runtime is designed to be a drop-in replacement for runc, the object contains a `path` field, and an optional `runtimeArgs` field.

See [Configure runc drop-in replacements](#).

After changing the runtimes configuration in the configuration file, you must reload or restart the daemon for changes to take effect:

[Configure containerd shims](#)

If the runtime that you want to register implements a containerd shim, or if you want to register a runtime which uses the runc shim, use the following format for the runtime entry:

`runtimeType` refers to either:

- A fully qualified name of a containerd shim.

The fully qualified name of a shim is the same as the `runtime_type` used to register the runtime in containerd's CRI configuration. For example, `io.containerd.runc.v1`.

- The path of a containerd shim binary.

This option is useful if you installed the containerd shim binary outside of `PATH`.

`options` is optional. It lets you specify the runtime configuration that you want to use for the shim. The configuration parameters that you can specify in `options` depends on the runtime you're registering. For most shims, the supported configuration options are `TypeUrl` and `ConfigPath`. For example:

You can configure multiple runtimes using the same `runtimeType`. For example:

The `options` field takes a special set of configuration parameters when used with `"runtimeType": "io.containerd.runc.v2"`. For more information about runc parameters, refer to the runc configuration section in [CRI Plugin Config Guide](#).

[Configure runc drop-in replacements](#)

If the runtime that you want to register can act as a drop-in replacement for runc, you can register the runtime either using the daemon configuration file, or using the `--add-runtime` flag for the `dockerd` cli.

When you use the configuration file, the entry uses the following format:

Where `path` is either the absolute path to the runtime executable, or the name of an executable installed on `PATH`:

And `runtimeArgs` lets you optionally pass additional arguments to the runtime. Entries with this format use the containerd runc shim to invoke a custom runtime binary.

When you use the `--add-runtime` CLI flag, use the following format:

Defining runtime arguments via the command line is not supported.

For an example configuration for a runc drop-in replacement, see [Alternative container runtimes > youki](#)

[Configure the default container runtime](#)

You can specify either the name of a fully qualified containerd runtime shim, or the name of a registered runtime. You can specify the default runtime either using the daemon configuration file, or using the `--default-runtime` flag for the `dockerd` cli.

When you use the configuration file, the entry uses the following format:

When you use the `--default-runtime` CLI flag, use the following format:

[Run containerd standalone](#)

By default, the Docker daemon automatically starts `containerd`. If you want to control `containerd` startup, manually start `containerd` and pass the path to the `containerd` socket using the `--containerd` flag. For example:

[Configure cgroup driver](#)

You can configure how the runtime should manage container cgroups, using the `--exec-opt native.cgroupdriver` CLI flag.

You can only specify `cgroupfs` or `systemd`. If you specify `systemd` and it is not available, the system errors out. If you omit the `native.cgroupdriver` option, `cgroupfs` is used on cgroup v1 hosts, `systemd` is used on cgroup v2 hosts with `systemd` available.

This example sets the `cgroupdriver` to `systemd`:

Setting this option applies to all containers the daemon launches.

[Configure container isolation technology_\(Windows\)](#)

For Windows containers, you can specify the default container isolation technology to use, using the `--exec-opt isolation` flag.

The following example makes `hyperv` the default isolation technology:

If no isolation value is specified on daemon start, on Windows client, the default is `hyperv`, and on Windows server, the default is `process`.

[Daemon DNS options](#)

To set the DNS server for all Docker containers, use:

To set the DNS search domain for all Docker containers, use:

[Insecure registries](#)

In this section, "registry" refers to a private registry, and `myregistry:5000` is a placeholder example of a private registry.

Docker considers a private registry either secure or insecure. A secure registry uses TLS and a copy of its CA certificate is placed on the Docker host at `/etc/docker/certs.d/myregistry:5000/ca.crt`. An insecure registry is either not using TLS (i.e., listening on plain text HTTP), or is using TLS with a CA certificate not known by the Docker daemon. The latter can happen when the certificate wasn't found under `/etc/docker/certs.d/myregistry:5000/`, or if the certificate verification failed (i.e., wrong CA).

By default, Docker assumes all registries to be secure, except for local registries. Communicating with an insecure registry isn't possible if Docker assumes that registry is secure. In order to communicate with an insecure registry, the Docker daemon requires `--insecure-registry` in one of the following two forms:

- `--insecure-registry myregistry:5000` tells the Docker daemon that `myregistry:5000` should be considered insecure.
- `--insecure-registry 10.1.0.0/16` tells the Docker daemon that all registries whose domain resolve to an IP address is part of the subnet described by the CIDR syntax, should be considered insecure.

The flag can be used multiple times to allow multiple registries to be marked as insecure.

If an insecure registry isn't marked as insecure, `docker pull`, `docker push`, and `docker search` result in error messages, prompting the user to either secure or pass the `--insecure-registry` flag to the Docker daemon as described above.

Local registries, whose IP address falls in the 127.0.0.0/8 range, are automatically marked as insecure as of Docker 1.3.2. It isn't recommended to rely on this, as it may change in the future.

Enabling `--insecure-registry`, i.e., allowing un-encrypted and/or untrusted communication, can be useful when running a local registry. However, because its use creates security vulnerabilities it should only be enabled for testing purposes. For increased security, users should add their CA to their system's list of trusted CAs instead of enabling `--insecure-registry`.

[Legacy Registries](#)

Operations against registries supporting only the legacy v1 protocol are no longer supported. Specifically, the daemon doesn't attempt to push, pull or sign in to v1 registries. The exception to this is `search` which can still be performed on v1 registries.

[Running a Docker daemon behind an HTTPS PROXY](#)

When running inside a LAN that uses an `HTTPS` proxy, the proxy's certificates replace Docker Hub's certificates. These certificates must be added to your Docker host's configuration:

1. Install the `ca-certificates` package for your distribution
2. Ask your network admin for the proxy's CA certificate and append them to `/etc/pki/tls/certs/ca-bundle.crt`
3. Then start your Docker daemon with `HTTPS_PROXY=http://username:password@proxy:port/ dockerd` . The `username:` and `password@` are optional - and are only needed if your proxy is set up to require authentication.

This only adds the proxy and authentication to the Docker daemon's requests. To use the proxy when building images and running containers, see [Configure Docker to use a proxy server](#)

Default `ulimit` settings

The `--default-ulimit` flag lets you set the default `ulimit` options to use for all containers. It takes the same options as `--ulimit` for `docker run` . If these defaults aren't set, `ulimit` settings are inherited from the Docker daemon. Any `--ulimit` options passed to `docker run` override the daemon defaults.

Be careful setting `nproc` with the `ulimit` flag, as `nproc` is designed by Linux to set the maximum number of processes available to a user, not to a container. For details, see [docker run reference](#).

Access authorization

Docker's access authorization can be extended by authorization plugins that your organization can purchase or build themselves. You can install one or more authorization plugins when you start the Docker `daemon` using the `--authorization-plugin=PLUGIN_ID` option.

The `PLUGIN_ID` value is either the plugin's name or a path to its specification file. The plugin's implementation determines whether you can specify a name or path. Consult with your Docker administrator to get information about the plugins available to you.

Once a plugin is installed, requests made to the `daemon` through the command line or Docker's Engine API are allowed or denied by the plugin. If you have multiple plugins installed, each plugin, in order, must allow the request for it to complete.

For information about how to create an authorization plugin, refer to the [authorization plugin](#) section.

Daemon user namespace options

The Linux kernel [user namespace support](#) provides additional security by enabling a process, and therefore a container, to have a unique range of user and group IDs which are outside the traditional user and group range utilized by the host system. One of the most important security improvements is that, by default, container processes running as the `root` user have expected administrative privileges it expects (with some restrictions) inside the container, but are effectively mapped to an unprivileged `uid` on the host.

For details about how to use this feature, as well as limitations, see [Isolate containers with a user namespace](#).

Configure host gateway IP

The Docker daemon supports a special `host-gateway` value for the `--add-host` flag for the `docker run` and `docker build` commands. This value resolves to addresses on the host, so that containers can connect to services running on the host.

By default, `host-gateway` resolves to the IPv4 address of the default bridge, and its IPv6 address if it has one.

You can configure this to resolve to a different IP using the `--host-gateway-ip` flag for the `dockerd` command line interface, or the `host-gateway-ip` key in the daemon configuration file.

To supply both IPv4 and IPv6 addresses on the command line, use two `--host-gateway-ip` options.

To supply addresses in the daemon configuration file, use `"host-gateway-ips"` with a JSON array, as shown below. For compatibility with older versions of the daemon, a single IP address can also be specified as a JSON string in option `"host-gateway-ip"`.

[Configure CDI devices](#)

Container Device Interface (CDI) is a [standardized](#) mechanism for container runtimes to create containers which are able to interact with third party devices.

CDI is currently only supported for Linux containers and is enabled by default since Docker Engine 28.3.0.

The Docker daemon supports running containers with CDI devices if the requested device specifications are available on the filesystem of the daemon.

The default specification directories are:

- `/etc/cdi/` for static CDI Specs
- `/var/run/cdi` for generated CDI Specs

[Set custom locations](#)

To set custom locations for CDI specifications, use the `cdi-spec-dirs` option in the `daemon.json` configuration file, or the `--cdi-spec-dir` flag for the `dockerd` CLI:

You can view the configured CDI specification directories using the `docker info` command.

[Disable CDI devices](#)

The feature is enabled by default. To disable it, use the `cdi` options in the `daemon.json` file:

To check the status of the CDI devices, run `docker info`.

[Daemon logging format](#)

The `--log-format` option or "log-format" option in the [daemon configuration file](#) lets you set the format for logs produced by the daemon. The logging format should only be configured either through the `--log-format` command line option or through the "log-format" field in the configuration file; using both the command-line

option and the "log-format" field in the configuration file produces an error. If this option is not set, the default is "text".

The following example configures the daemon through the `--log-format` command line option to use `json` formatted logs;

The following example shows a `daemon.json` configuration file with the "log-format" set;

[Miscellaneous options](#)

IP masquerading uses address translation to allow containers without a public IP to talk to other machines on the internet. This may interfere with some network topologies, and can be disabled with `--ip-masq=false`.

Docker supports soft links for the Docker data directory (`/var/lib/docker`) and for `/var/lib/docker/tmp` . The `DOCKER_TMPDIR` and the data directory can be set like this:

[Default cgroup parent](#)

The `--cgroup-parent` option lets you set the default cgroup parent for containers. If this option isn't set, it defaults to `/docker` for the cgroupfs driver, and `system.slice` for the systemd cgroup driver.

If the cgroup has a leading forward slash (`/`), the cgroup is created under the root cgroup, otherwise the cgroup is created under the daemon cgroup.

Assuming the daemon is running in cgroup `daemoncgroup` , `--cgroup-parent=/foobar` creates a cgroup in `/sys/fs/cgroup/memory/foobar` , whereas using `--cgroup-parent=foobar` creates the cgroup in `/sys/fs/cgroup/memory/daemoncgroup/foobar`

The systemd cgroup driver has different rules for `--cgroup-parent` . systemd represents hierarchy by slice and the name of the slice encodes the location in the tree. So `--cgroup-parent` for systemd cgroups should be a slice name. A name can consist of a dash-separated series of names, which describes the path to the slice from the root slice. For example, `--cgroup-parent=user-a-b.slice` means the memory cgroup for the container is created in `/sys/fs/cgroup/memory/user.slice/user-a.slice/user-a-b.slice/docker-<id>.scope` .

This setting can also be set per container, using the `--cgroup-parent` option on `docker create` and `docker run` , and takes precedence over the `--cgroup-parent` option on the daemon.

[Daemon metrics](#)

The `--metrics-addr` option takes a TCP address to serve the metrics API. This feature is still experimental, therefore, the daemon must be running in experimental mode for this feature to work.

To serve the metrics API on `localhost:9323` you would specify `--metrics-addr 127.0.0.1:9323` , allowing you to make requests on the API at `127.0.0.1:9323/metrics` to receive metrics in the [prometheus](#) format.

Port `9323` is the [default port associated with Docker metrics](#) to avoid collisions with other Prometheus exporters and services.

If you are running a Prometheus server you can add this address to your scrape configs to have Prometheus collect metrics on Docker. For more information, see [Collect Docker metrics with Prometheus](#).

[Node generic resources](#)

The `--node-generic-resources` option takes a list of key-value pair (`key=value`) that allows you to advertise user defined resources in a Swarm cluster.

The current expected use case is to advertise NVIDIA GPUs so that services requesting `NVIDIA-GPU=[0-16]` can land on a node that has enough GPUs for the task to run.

Example of usage:

[Enable feature in the daemon \(--feature\)](#)

The `--feature` option lets you enable or disable a feature in the daemon. This option corresponds with the "features" field in the [daemon.json configuration file](#). Features should only be configured either through the `--feature` command line option or through the "features" field in the configuration file; using both the command-line option and the "features" field in the configuration file produces an error. The feature option can be specified multiple times to configure multiple features. The `--feature` option accepts a name and optional boolean value. When omitting the value, the default is `true`.

The following example runs the daemon with the `cdi` and `containerd-snapshotter` features enabled. The `cdi` option is provided with a value;

The following example is the equivalent using the `daemon.json` configuration file;

[Daemon configuration file](#)

The `--config-file` option allows you to set any configuration option for the daemon in a JSON format. This file uses the same flag names as keys, except for flags that allow several entries, where it uses the plural of the flag name, e.g., `labels` for the `label` flag.

The options set in the configuration file must not conflict with options set using flags. The Docker daemon fails to start if an option is duplicated between the file and the flags, regardless of their value. This is intentional, and avoids silently ignore changes introduced in configuration reloads. For example, the daemon fails to start if you set daemon labels in the configuration file and also set daemon labels via the `--label` flag. Options that are not present in the file are ignored when the daemon starts.

The `--validate` option allows to validate a configuration file without starting the Docker daemon. A non-zero exit code is returned for invalid configuration files.

[On Linux](#)

The default location of the configuration file on Linux is `/etc/docker/daemon.json`. Use the `--config-file` flag to specify a non-default location.

The following is a full example of the allowed configuration options on Linux:

You can't set options in `daemon.json` that have already been set on daemon startup as a flag. On systems that use `systemd` to start the Docker daemon, `-H` is already set, so you can't use the `hosts` key in `daemon.json` to add listening addresses. See [custom Docker daemon options](#) for an example on how to configure the daemon using `systemd` drop-in files.

[On Windows](#)

The default location of the configuration file on Windows is `%programdata%\docker\config\daemon.json`. Use the `--config-file` flag to specify a non-default location.

The following is a full example of the allowed configuration options on Windows:

The `default-runtime` option is by default unset, in which case `dockerd` automatically detects the runtime. This detection is based on if the `containerd` flag is set.

Accepted values:

- `com.docker.hcsshim.v1` - This is the built-in runtime that Docker has used since Windows supported was first added and uses the v1 HCS API's in Windows.
- `io.containerd.runhcs.v1` - This is uses the `containerd runhcs` shim to run the container and uses the v2 HCS API's in Windows.

[Feature options](#)

The optional field `features` in `daemon.json` lets you enable or disable specific daemon features.

The list of feature options include:

- `containerd-snapshotter` : when set to `true`, the daemon uses `containerd` snapshotters instead of the classic storage drivers for storing image and container data. For more information, see [containerd storage](#).
- `windows-dns-proxy` : when set to `true`, the daemon's internal DNS resolver will forward requests to external servers. Without this, most applications running in the container will still be able to use secondary DNS servers configured in the container itself, but `nslookup` won't be able to resolve external names. The current default is `false`, it will change to `true` in a future release. This option is only allowed on Windows.

The `windows-dns-proxy` feature flag will be removed in a future release.

[Configuration reload behavior](#)

Some options can be reconfigured when the daemon is running without requiring to restart the process. The daemon uses the `SIGHUP` signal in Linux to reload, and a global event in Windows with the key `Global\docker-daemon-config-$PID`. You can modify the options in the configuration file, but the daemon still checks for

conflicting settings with the specified CLI flags. The daemon fails to reconfigure itself if there are conflicts, but it won't stop execution.

The list of currently supported options that can be reconfigured is this:

Option	Description
<code>debug</code>	Toggles debug mode of the daemon.
<code>labels</code>	Replaces the daemon labels with a new set of labels.
<code>live-restore</code>	Toggles live restore .
<code>max-concurrent-downloads</code>	Configures the max concurrent downloads for each pull.
<code>max-concurrent-uploads</code>	Configures the max concurrent uploads for each push.
<code>max-download-attempts</code>	Configures the max download attempts for each pull.
<code>default-runtime</code>	Configures the runtime to be used if not is specified at container creation.
<code>runtimes</code>	Configures the list of available OCI runtimes that can be used to run containers.
<code>authorization-plugin</code>	Specifies the authorization plugins to use.
<code>insecure-registries</code>	Specifies a list of registries that the daemon should consider insecure.
<code>registry-mirrors</code>	Specifies a list of registry mirrors.
<code>shutdown-timeout</code>	Configures the daemon's existing configuration timeout with a new timeout for shutting down all containers.
<code>features</code>	Enables or disables specific features.

[Run multiple daemons](#)

Running multiple daemons on a single host is considered experimental. You may encounter unsolved problems, and things may not work as expected in some cases.

This section describes how to run multiple Docker daemons on a single host. To run multiple daemons, you must configure each daemon so that it doesn't conflict with other daemons on the same host. You can set these options either by providing them as flags, or by using a [daemon configuration file](#).

The following daemon options must be configured for each daemon:

When your daemons use different values for these flags, you can run them on the same host without any problems. It is important that you understand the meaning of these options and to use them correctly.

- The `-b, --bridge=` flag is set to `docker0` as default bridge network. It is created automatically when you install Docker. If you aren't using the default, you must create and configure the bridge manually, or set it to 'none': `--bridge=none`
- `--exec-root` is the path where the container state is stored. The default value is `/var/run/docker` . Specify the path for your running daemon here.
- `--data-root` is the path where persisted data such as images, volumes, and cluster state are stored. The default value is `/var/lib/docker` . To avoid any conflict with other daemons, set this parameter separately for each daemon.
- `-p, --pidfile=/var/run/docker.pid` is the path where the process ID of the daemon is stored. Specify the path for your PID file here.
- `--host=[]` specifies where the Docker daemon listens for client connections. If unspecified, it defaults to `/var/run/docker.sock` .
- `--iptables=false` prevents the Docker daemon from adding iptables rules. If multiple daemons manage iptables rules, they may overwrite rules set by another daemon. Be aware that disabling this option requires you to manually add iptables rules to expose container ports. If you prevent Docker from adding iptables rules, Docker also doesn't add IP masquerading rules, even if you set `--ip-masq` to `true` . Without IP masquerading rules, Docker containers can't connect to external hosts or the internet when using network other than default bridge.
- `--config-file=/etc/docker/daemon.json` is the path where configuration file is stored. You can use it instead of daemon flags. Specify the path for each daemon.
- `--tls*` Docker daemon supports `--tlsverify` mode that enforces encrypted and authenticated remote connections. The `--tls*` options enable use of specific certificates for individual daemons.

Example script for a separate “bootstrap” instance of the Docker daemon without network:

[Default network options](#)

The `default-network-opts` key in the `daemon.json` configuration file, and the equivalent `--default-network-opt` CLI flag, let you specify default values for driver network driver options for new networks.

The following example shows how to configure options for the `bridge` driver using the `daemon.json` file.

This example uses the `bridge` network driver. Refer to the [bridge network driver page](#) for an overview of available driver options.

After changing the configuration and restarting the daemon, new networks that you create use these option configurations as defaults.

Note that changing this daemon configuration doesn't affect pre-existing networks.

Using the `--default-network-opt` CLI flag is useful for testing and debugging purposes, but you should prefer using the `daemon.json` file for persistent daemon configuration. The CLI flag expects a value with the following

format: `driver=opt=value` , for example:

Source: <https://docs.docker.com/engine/reference/commandline/dockerd/>