NAME

dockerd − Enable daemon mode

SYNOPSIS

dockerd [**--add-runtime**[*=[]*]] [**--authorization-plugin**[*=[]*]] [**-b**|**--bridge**[*=BRIDGE*]] [**--bip**[*=BIP*]] [**--bip6**[*=BIP*]] [**--bridge-accept-fwmark**[*=[]*]] [**--cgroup-parent**[*=[]*]] [**--config-file**[*=path*]] [**--containerd**[*=SOCKET-PATH*]] [**--data-root**[*=/var/lib/docker*]] [**-D**|**--debug**] [**--default-cgroupns-mode**[*=host*]] [**--default-gateway**[*=DEFAULT-GATEWAY*]] [**--default-gateway-v6**[*=DEFAULT-GATEWAY-V6*]] [**--default-address-pool**[*=DEFAULT-ADDRESS-POOL*]] [**--default-network-opt**[*=DRIVER=OPT=VALUE*]] [**--default-runtime**[*=runc*]] [**--default-ipc-mode**=*MODE*] [**--default-shm-size**[*=64MiB*]] [**--default-ulimit**[*=[]*]] [**--dns**[*=[]*]] [**--dns-opt**[*=[]*]] [**--dns-search**[*=[]*]] [**--exec-opt**[*=[]*]] [**--exec-root**[*=/var/run/docker*]] [**--experimental**[=**false**]] [**--feature**[*=NAME*[=**true**|**false**]] [**--fixed-cidr**[*=FIXED-CIDR*]] [**--fixed-cidr-v6**[*=FIXED-CIDR-V6*]] [**-G**|**--group**[*=docker*]] [**--help**] [**-H**|**--host**[*=[]*]] [**--host-gateway-ip**[*=HOST-GATEWAY-IP*]] [**--http-proxy**[*""*]] [**--https-proxy**[*""*]] [**--icc**[=**true**]] [**--init**[=**false**]] [**--init-path**[*=""*]] [**--insecure-registry**[*=[]*]] [**--ip**[*=0.0.0.0*]] [**--ip-forward**[=**true**]] [**--ip-forward-no-drop**[=**true**]] [**--ip-masq**[=**true**]] [**--iptables**[=**true**]] [**--ipv6**] [**--isolation**[*=default*]] [**-l**|**--log-level**[*=info*]] [**--label**[*=[]*]] [**--live-restore**[=**false**]] [**--log-driver**[*=json-file*]] [**--log-format**="*text*|*json*"] [**--log-opt**[*=map[]*]] [**--mtu**[*=0*]] [**--max-concurrent-downloads**[*=3*]] [**--max-concurrent-uploads**[*=5*]] [**--max-download-attempts**[*=5*]] [**--no-proxy**[*""*]] [**--node-generic-resources**[*=[]*]] [**-p**|**--pid-file**[*=/var/run/docker.pid*]] [**--raw-logs**] [**--registry-mirror**[*=[]*]] [**-s**|**--storage-driver**[*=STORAGE-DRIVER*]] [**--seccomp-profile**[*=SECCOMP-PROFILE-PATH*]] [**--selinux-enabled**] [**--shutdown-timeout**[*=15*]] [**--storage-opt**[*=[]*]] [**--swarm-default-advertise-addr**[*=IP|INTERFACE*]] [**--tls**] [**--tlscacert**[*=~/.docker/ca.pem*]] [**--tlscert**[*=~/.docker/cert.pem*]] [**--tlskey**[*=~/.docker/key.pem*]] [**--tlsverify**] [**--userland-proxy**[=**true**]] [**--userland-proxy-path**[*=""*]] [**--userns-remap**[*=default*]] [**--validate**]

DESCRIPTION

dockerd is used for starting the Docker daemon (i.e., to command the daemon to manage images, containers etc). So **dockerd** is a server, as a daemon.

To run the Docker daemon you can specify **dockerd**. You can check the daemon options using **dockerd --help**. Daemon options should be specified after the**dock erd** keyword in the following format.

dockerd [OPTIONS]

OPTIONS

--add-runtime=[]
Runtimes can be registered with the daemon either via the configuration file or using the **--add-runtime** command line argument.

The following is an example adding 2 runtimes via the configuration:

```
{
      "default-runtime": "runc",
      "runtimes": {
            "runc": {
                  "path": "runc"
            },
            "custom": {
                  "path": "/usr/local/bin/my-runc-replacement",
```

```
                "runtimeArgs": [
                      "--debug"
                ]
          }
      }
}
```

This is the same example via the command line:

```
$ sudo dockerd --add-runtime runc=runc --add-runtime custom=/usr/local/bin/my-runc-r
```

Note: defining runtime arguments via the command line is not supported.

--authorization-plugin=""
Set authorization plugins to load

-b, **--bridge**=""
Attach containers to a pre-existing network bridge, instead of docker0; use 'none' to disable the default bridge network

--bip=""
Use the provided CIDR notation IPv4 address for the default bridge network; Mutually exclusive of -b

--bip6=""
Use the provided CIDR notation IPv6 address for the default bridge network; Mutually exclusive of -b

--bridge-accept-fwmark="" Bridge networks will accept packets with this firewall mark/mask.

--cgroup-parent=""
Set parent cgroup for all containers. Default is "/docker" for fs cgroup driver and "system.slice" for systemd cgroup driver.

--config-file=*"/etc/docker/daemon.json"*
Specifies the JSON file path to load the configuration from. Default is */etc/docker/daemon.json*.

--containerd=""
Path to containerd socket.

--data-root=""
Path to the directory used to store persisted Docker data such as configuration for resources, swarm cluster state, and filesystem data for images, containers, and local volumes. Default is */var/lib/docker*.

-D, **--debug**=**true**|**false**
Enable debug mode. Default is **false**.

--default-cgroupns-mode="**host**|**private**"
Set the default cgroup namespace mode for newly created containers. The argument can either be **host** or **private**. If unset, this defaults to **host** on cgroup v1, or **private** on cgroup v2.

--default-gateway=""
IPv4 default gateway for the default bridge network; this address must be part of the bridge subnet (which is defined by -b or --bip)

--default-gateway-v6=""
IPv6 default gateway for the default bridge network

--default-address-pool=""
Default address pool from which IPAM driver selects a subnet for the networks. Example: base=172.30.0.0/16,size=24 will set the default address pools for the selected scope networks to {172.30.[0-255].0/24}

--default-network-opt=*DRIVER=OPT=VALUE*
Default network driver options

--default-runtime=*"runtime"*
Set default runtime if there’re more than one specified by **--add-runtime**.

--default-ipc-mode="**private**|**shareable**"
Set the default IPC mode for newly created containers. The argument can either be **private** or **shareable**.

--default-shm-size=*size*
Set the daemon-wide default shm *size* for containers. Default is **64MiB**.

--default-ulimit=[]
Default ulimits for containers.

--dns=""
Force Docker to use specific DNS servers.

--dns-opt=""
DNS options to use.

--dns-search=[]
DNS search domains to use.

--exec-opt=[]
Set runtime execution options. See RUNTIME EXECUTION OPTIONS.

--exec-root=""
Path to use as the root of the Docker execution state files. Default is **/var/run/docker**.

--experimental=""
Enable the daemon experimental features.

--feature=*NAME*[=**true**|**false**]
Enable or disable a feature in the daemon. This option corresponds with the "features" field in the daemon.json 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.
Usage example: **--feature containerd-snapshotter** or **--feature containerd-snapshotter=true**.

--fixed-cidr=""
IPv4 subnet for the default bridge network (e.g., 10.20.0.0/16); this subnet must be nested in the bridge subnet (which is defined by -b or --bip).

--fixed-cidr-v6=""
IPv6 subnet for the default bridge network (e.g., 2001:db8::/64).

-G, **--group**=""
Group to assign the unix socket specified by -H when running in daemon mode. use ’’ (the empty string) to disable setting of a group. Default is **docker**.

--help
Print usage statement

-H, **--host**=[*unix:///var/run/docker.sock*]: tcp://[host:port] to bind or unix:///[/path/to/socket] to use.
The socket(s) to bind to in daemon mode specified using one or more tcp://host:port, unix:///path/to/socket, fd://* or fd://socketfd.

--host-gateway-ip=[*2001:db8::1234*]
Supply host addresses to substitute for the special string host-gateway in --add-host options. Addresses from the docker0 bridge are used by default. Two of these options are allowed, one IPv4 and one IPv6 address.

--http-proxy*""*
Proxy URL for HTTP requests unless overridden by NoProxy.

--https-proxy*""*
Proxy URL for HTTPS requests unless overridden by NoProxy.

--icc=**true**|**false**
Allow unrestricted inter-container communication in the default bridge network. If disabled, containers can still be linked together using the **--link** option (see **docker-run**(1)). Default is **true**.

--init
Run an init process inside containers for signal forwarding and process reaping.

--init-path
Path to the docker-init binary.

--insecure-registry=[]
Enable insecure registry communication, i.e., enable un-encrypted and/or untrusted communication.

List of insecure registries can contain an element with CIDR notation to specify a whole subnet. Insecure registries accept HTTP and/or accept HTTPS with certificates from unknown CAs.

Enabling **--insecure-registry** is 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 using **--insecure-registry**.

--ip=""
Default host IP address to use when publishing container ports from the default bridge network. Default is **0.0.0.0**.

--ip-forward=**true**|**false**
Enables IP forwarding on the Docker host. The default is **true**. This flag interacts with the IP forwarding setting on your host system's kernel. If your system has IP forwarding disabled, this setting enables it. If your system has IP forwarding enabled, setting this flag to **false** has no effect.

This setting will also enable IPv6 forwarding if you have both **--ip-forward=true** and an IPv6 enabled bridge network. Note that this may reject Router Advertisements and interfere with the host's existing IPv6 configuration. For more information, consult the documentation about "Advanced Networking - IPv6".

--ip-forward-no-drop=**true**|**false**
When **false**, the default, if Docker enables IP forwarding itself (see **--ip-forward**), and **--iptables** or **--ip6tables** are enabled, it also sets the default policy for the FORWARD chain in the iptables or ip6tables filter table to DROP.

When **true**, and when IP forwarding is already enabled, Docker does not modify the default policy of the FORWARD chain.

--ip-masq=**true**|**false**
Enable IP masquerading for default bridge's IP range. Default is **true**.

--iptables=**true**|**false**
Enable Docker's addition of iptables rules. Default is **true**.

--ipv6=**true**|**false**

Enable IPv6 support on the default bridge network. Default is **false**.
By default, Docker will generate a ULA subnet, **--default-addresss-pool** or **--fixed-cidr-v6** can be used to override the subnet. IPv6 forwarding will be enabled if not used with **--ip-forward=false**. This may collide with your host's current IPv6 settings. For more information consult the documentation about "Advanced Networking - IPv6".

--isolation="*default*"
Isolation specifies the type of isolation technology used by containers. Note that the default on Windows server is **process**, and the default on Windows client is **hyperv**. Linux only supports **default**.

-l, **--log-level**="*debug*|*info*|*warn*|*error*|*fatal*"
Set the logging level. Default is **info**.

--label="[]"
Set key=value labels to the daemon (displayed in **docker info**)

--live-restore=false
Enable live restore of running containers when the daemon starts so that they are not restarted. This option is applicable only for docker daemon running on Linux host.

--log-driver="**json-file**|**syslog**|**journald**|**gelf**|**fluentd**|**awslogs**|**splunk**|**etwlogs**|**gcplogs**|**none**"
Default driver for container logs. Default is **json-file**.
Warning: **docker logs** command works only for **json-file** logging driver.

--log-format="*text*|*json*"
Set the format for logs produced by the daemon. Default is "text".

--log-opt=[]
Logging driver specific options.

--mtu=*0*
Set the network MTU for the default bridge network. Default is **0**.

--max-concurrent-downloads=*3*
Set the max concurrent downloads. Default is **3**.

--max-concurrent-uploads=*5*
Set the max concurrent uploads. Default is **5**.

--max-download-attempts=*5*
Set the max download attempts for each pull. Default is **5**.

--no-proxy=*""*
Comma-separated values specifying hosts that should be excluded from proxying.

--node-generic-resources=*[]*

Advertise user-defined resource. Default is **[]**.
Use this if your swarm cluster has some nodes with custom resources (e.g: NVIDIA GPU, SSD, ...) and you need your services to land on nodes advertising these resources.
Usage example: **--node-generic-resources "NVIDIA-GPU=UUID1" --node-generic-resources "NVIDIA-GPU=UUID2"**

-p, **--pidfile**="*path*"
Path to use for daemon PID file. Default is */var/run/docker.pid*.

--raw-logs
Output daemon logs in full timestamp format without ANSI coloring. If this flag is not set, the daemon outputs condensed, colorized logs if a terminal is detected, or full ("raw") output otherwise.

--registry-mirror=*://*
Prepend a registry mirror to be used for image pulls. May be specified multiple times.

-s, **--storage-driver**=""
Force the Docker runtime to use a specific storage driver.

--seccomp-profile=""
Path to seccomp profile.

--selinux-enabled=**true**|**false**
Enable selinux support. Default is **false**.

--shutdown-timeout=*seconds*
Set the shutdown timeout value in seconds. Default is **15**.

--storage-opt=[]
Set storage driver options. See STORAGE DRIVER OPTIONS.

--swarm-default-advertise-addr=*IP|INTERFACE*
Set default address or interface for swarm to advertise as its externally-reachable address to other cluster members. This can be a hostname, an IP address, or an interface such as **eth0**. A port cannot be specified with this option.

--tls=**true**|**false**
Use TLS; implied by **--tlsverify**. Default is **false**.

--tlscacert=*~/.docker/ca.pem*
Trust certs signed only by this CA.

--tlscert=*~/.docker/cert.pem*
Path to TLS certificate file.

--tlskey=*~/.docker/key.pem*
Path to TLS key file.

--tlsverify=**true**|**false**
Use TLS and verify the remote (daemon: verify client, client: verify daemon).
Default is **false**.

--allow-direct-routing=**true**|**false**
Allow remote access to published ports on container IP addresses.
Default is **false**.

--userland-proxy=**true**|**false**
Rely on a userland proxy implementation for inter-container and
outside-to-container loopback communications. Default is **true**.

--userland-proxy-path=""
Path to the userland proxy binary.

--userns-remap=*default*|*uid:gid*|*user:group*|*user*|*uid*
Enable user namespaces for containers on the daemon. Specifying "default"
will cause a new user and group to be created to handle UID and GID range
remapping for the user namespace mappings used for contained processes.
Specifying a user (or uid) and optionally a group (or gid) will cause the
daemon to lookup the user and group's subordinate ID ranges for use as the
user namespace mappings for contained processes.

--validate
Validate daemon configuration and exit.

STORAGE DRIVER OPTIONS

Docker uses storage backends (known as "storage drivers" in the Docker internals) to create writable containers from images. Many of these backends use operating system level technologies and can be configured.

Specify options to the storage backend with **--storage-opt** flags. The backends that currently take options are **zfs** and **btrfs**. Options for**zfs** start with **zfs.**, and options for **btrfs** start with **btrfs.**.

ZFS options

zfs.fsname

Set zfs filesystem under which docker will create its own datasets. By default docker will pick up the zfs filesystem where docker graph (**/var/lib/docker**) is located.

Example use: **dockerd -s zfs --storage-opt zfs.fsname=zroot/docker**

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** cannot be smaller than **btrfs.min_space**.

Example use: **docker daemon -s btrfs --storage-opt btrfs.min_space=10G**

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.

```
dockerd --authorization-plugin=plugin1 --authorization-plugin=plugin2,...
```

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, see access authorization plugin ⟨https://docs.docker.com/engine/extend/plugins_authorization/⟩ section in the Docker extend section of this documentation.

RUNTIME EXECUTION OPTIONS

You can configure the runtime using options specified with the **--exec-opt** flag. All the flag's options have the **native** prefix. A single **native.cgroupdriver** option is available.

The **native.cgroupdriver** option specifies the management of the container's cgroups. 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**:

```
$ sudo dockerd --exec-opt native.cgroupdriver=systemd
```

Setting this option applies to all containers the daemon launches.

HISTORY

Sept 2015, Originally compiled by Shishir Mahajan shishir.mahajan@redhat.com ⟨mailto:shishir.mahajan@redhat.com⟩ based on docker.com source material and internal work.