

**NAME**

docker-container-run - Create and run a new container from an image

**SYNOPSIS**

**docker container run [OPTIONS] IMAGE [COMMAND] [ARG...]**

**DESCRIPTION**

Alias for **docker run**.

**OPTIONS**

- add-host=** Add a custom host-to-IP mapping (host:ip)
- annotation=map[]** Add an annotation to the container (passed through to the OCI runtime)
- a, --attach=** Attach to STDIN, STDOUT or STDERR
- blkio-weight=0** Block IO (relative weight), between 10 and 1000, or 0 to disable (default 0)
- blkio-weight-device=[]** Block IO weight (relative device weight)
- cap-add=** Add Linux capabilities
- cap-drop=** Drop Linux capabilities
- cgroup-parent=""** Optional parent cgroup for the container
- cgroupns=""** Cgroup namespace to use (host|private) 'host': Run the container in the Docker host's cgroup namespace 'private': Run the container in its own private cgroup namespace '' : Use the cgroup namespace as configured by the default-cgroupns-mode option on the daemon (default)
- cidfile=""** Write the container ID to the file
- cpu-count=0** CPU count (Windows only)
- cpu-percent=0** CPU percent (Windows only)
- cpu-period=0** Limit CPU CFS (Completely Fair Scheduler) period
- cpu-quota=0** Limit CPU CFS (Completely Fair Scheduler) quota
- cpu-rt-period=0** Limit CPU real-time period in microseconds
- cpu-rt-runtime=0** Limit CPU real-time runtime in microseconds

|                               |                                                              |
|-------------------------------|--------------------------------------------------------------|
| <b>-c, --cpu-shares=0</b>     | CPU shares (relative weight)                                 |
| <b>--cpus=</b>                | Number of CPUs                                               |
| <b>--cpuset-cpus=""</b>       | CPUs in which to allow execution (0-3, 0,1)                  |
| <b>--cpuset-mems=""</b>       | MEMs in which to allow execution (0-3, 0,1)                  |
| <b>-d, --detach[=false]</b>   | Run container in background and print container ID           |
| <b>--detach-keys=""</b>       | Override the key sequence for detaching a container          |
| <b>--device=</b>              | Add a host device to the container                           |
| <b>--device-cgroup-rule=</b>  | Add a rule to the cgroup allowed devices list                |
| <b>--device-read-bps=[]</b>   | Limit read rate (bytes per second) from a device             |
| <b>--device-read-iops=[]</b>  | Limit read rate (IO per second) from a device                |
| <b>--device-write-bps=[]</b>  | Limit write rate (bytes per second) to a device              |
| <b>--device-write-iops=[]</b> | Limit write rate (IO per second) to a device                 |
| <b>--dns=</b>                 | Set custom DNS servers                                       |
| <b>--dns-option=</b>          | Set DNS options                                              |
| <b>--dns-search=</b>          | Set custom DNS search domains                                |
| <b>--domainname=""</b>        | Container NIS domain name                                    |
| <b>--entrypoint=""</b>        | Overwrite the default ENTRYPOINT of the image                |
| <b>-e, --env=</b>             | Set environment variables                                    |
| <b>--env-file=</b>            | Read in a file of environment variables                      |
| <b>--expose=</b>              | Expose a port or a range of ports                            |
| <b>--gpus=</b>                | GPU devices to add to the container ('all' to pass all GPUs) |
| <b>--group-add=</b>           | Add additional groups to join                                |
| <b>--health-cmd=""</b>        | Command to run to check health                               |

**--health-interval=0s** Time between running the check (ms|s|m|h) (default 0s)

**--health-retries=0** Consecutive failures needed to report unhealthy

**--health-start-interval=0s** Time between running the check during the start period (ms|s|m|h) (default 0s)

**--health-start-period=0s** Start period for the container to initialize before starting health-retries countdown (ms|s|m|h) (default 0s)

**--health-timeout=0s** Maximum time to allow one check to run (ms|s|m|h) (default 0s)

**--help[=false]** Print usage

**-h, --hostname=""** Container host name

**--init[=false]** Run an init inside the container that forwards signals and reaps processes

**-i, --interactive[=false]** Keep STDIN open even if not attached

**--io-maxbandwidth=0** Maximum IO bandwidth limit for the system drive (Windows only)

**--io-maxiops=0** Maximum IOps limit for the system drive (Windows only)

**--ip=** IPv4 address (e.g., 172.30.100.104)

**--ip6=** IPv6 address (e.g., 2001:db8::33)

**--ipc=""** IPC mode to use

**--isolation=""** Container isolation technology

**-l, --label=** Set meta data on a container

**--label-file=** Read in a line delimited file of labels

**--link=** Add link to another container

**--link-local-ip=** Container IPv4/IPv6 link-local addresses

**--log-driver=""** Logging driver for the container

**--log-opt=** Log driver options

**--mac-address=""** Container MAC address (e.g., 92:d0:c6:0a:29:33)

|                                   |                                                                                       |
|-----------------------------------|---------------------------------------------------------------------------------------|
| <b>-m, --memory=0</b>             | Memory limit                                                                          |
| <b>--memory-reservation=0</b>     | Memory soft limit                                                                     |
| <b>--memory-swap=0</b>            | Swap limit equal to memory plus swap: '-1' to enable unlimited swap                   |
| <b>--memory-swappiness=-1</b>     | Tune container memory swappiness (0 to 100)                                           |
| <b>--mount=</b>                   | Attach a filesystem mount to the container                                            |
| <b>--name=""</b>                  | Assign a name to the container                                                        |
| <b>--network=</b>                 | Connect a container to a network                                                      |
| <b>--network-alias=</b>           | Add network-scoped alias for the container                                            |
| <b>--no-healthcheck[=false]</b>   | Disable any container-specified HEALTHCHECK                                           |
| <b>--oom-kill-disable[=false]</b> | Disable OOM Killer                                                                    |
| <b>--oom-score-adj=0</b>          | Tune host's OOM preferences (-1000 to 1000)                                           |
| <b>--pid=""</b>                   | PID namespace to use                                                                  |
| <b>--pids-limit=0</b>             | Tune container pids limit (set -1 for unlimited)                                      |
| <b>--platform=""</b>              | Set platform if server is multi-platform capable                                      |
| <b>--privileged[=false]</b>       | Give extended privileges to this container                                            |
| <b>-p, --publish=</b>             | Publish a container's port(s) to the host                                             |
| <b>-P, --publish-all[=false]</b>  | Publish all exposed ports to random ports                                             |
| <b>--pull="missing"</b>           | Pull image before running ("always", "missing", "never")                              |
| <b>-q, --quiet[=false]</b>        | Suppress the pull output                                                              |
| <b>--read-only[=false]</b>        | Mount the container's root filesystem as read only                                    |
| <b>--restart="no"</b>             | Restart policy to apply when a container exits                                        |
| <b>--rm[=false]</b><br>exits      | Automatically remove the container and its associated anonymous volumes when it exits |
| <b>--runtime=""</b>               | Runtime to use for this container                                                     |

|                                 |                                                |
|---------------------------------|------------------------------------------------|
| <b>--security-opt=</b>          | Security Options                               |
| <b>--shm-size=0</b>             | Size of /dev/shm                               |
| <b>--sig-proxy[=true]</b>       | Proxy received signals to the process          |
| <b>--stop-signal=""</b>         | Signal to stop the container                   |
| <b>--stop-timeout=0</b>         | Timeout (in seconds) to stop a container       |
| <b>--storage-opt=</b>           | Storage driver options for the container       |
| <b>--sysctl=map[]</b>           | Sysctl options                                 |
| <b>--tmpfs=</b>                 | Mount a tmpfs directory                        |
| <b>-t, --tty[=false]</b>        | Allocate a pseudo-TTY                          |
| <b>--ulimit=[]</b>              | Ulimit options                                 |
| <b>--use-api-socket[=false]</b> | Bind mount Docker API socket and required auth |
| <b>-u, --user=""</b>            | Username or UID (format: [:])                  |
| <b>--users=""</b>               | User namespace to use                          |
| <b>--uts=""</b>                 | UTS namespace to use                           |
| <b>-v, --volume=</b>            | Bind mount a volume                            |
| <b>--volume-driver=""</b>       | Optional volume driver for the container       |
| <b>--volumes-from=</b>          | Mount volumes from the specified container(s)  |
| <b>-w, --workdir=""</b>         | Working directory inside the container         |

**SEE ALSO****docker-container(1)**