NAME

docker-network-create - Create a network

SYNOPSIS

docker network create [OPTIONS] NETWORK

DESCRIPTION

Creates a new network. The **DRIVER** accepts **bridge** or **overlay** which are the built-in network drivers. If you have installed a third party or your own custom network driver you can specify that **DRIVER** here also. If you don't specify the **--driver** option, the command automatically creates a **bridge** network for you. When you install Docker Engine it creates a **bridge** network automatically. This network corresponds to the **docker0** bridge that Engine has traditionally relied on. When you launch a new container with **dock er run** it automatically connects to this bridge network. You cannot remove this default bridge network, but you can create new ones using the **network create** command.

```
$ docker network create -d bridge my-bridge-network
```

Bridge networks are isolated networks on a single Engine installation. If you want to create a network that spans multiple Docker hosts each running an Engine, you must enable Swarm mode, and create an **overlay** network. To read more about overlay networks with Swarm mode, see " ⟨https://docs.docker.com/network/overlay/⟩.

Once you have enabled swarm mode, you can create a swarm-scoped overlay network:

```
$ docker network create --scope=swarm --attachable -d overlay my-multihost-network
```

By default, swarm-scoped networks do not allow manually started containers to be attached. This restriction is added to prevent someone that has access to a non-manager node in the swarm cluster from running a container that is able to access the network stack of a swarm service.

The **--attachable** option used in the example above disables this restriction, and allows for both swarm services and manually started containers to attach to the overlay network.

Network names must be unique. The Docker daemon attempts to identify naming conflicts but this is not guaranteed. It is the user's responsibility to avoid name conflicts.

Overlay network limitations

You should create overlay networks with **/24** blocks (the default), which limits you to 256 IP addresses, when you create networks using the default VIP-based endpoint-mode. This recommendation addresses limitations with swarm mode ⟨https://github.com/moby/moby/issues/30820⟩. If you need more than 256 IP addresses, do not increase the IP block size. You can either use **dnsrr** endpoint mode with an external load balancer, or use multiple smaller overlay networks. See Configure service discovery ⟨https://docs.docker.com/engine/swarm/networking/#configure−service−discovery⟩ for more information about different endpoint modes.

Connect containers

When you start a container, use the **--network** flag to connect it to a network. This example adds the **busybox** container to the **mynet** network:

```
$ docker run -itd --network=mynet busybox
```

If you want to add a container to a network after the container is already running, use the **docker network connect** subcommand.

You can connect multiple containers to the same network. Once connected, the containers can communicate using only another container's IP address or name. For **overlay** networks or custom plugins that support multi-host connectivity, containers connected to the same multi-host network but launched from different Engines can also communicate in this way.

You can disconnect a container from a network using the **docker network disconnect** command.

Specify advanced options

When you create a network, Engine creates a non-overlapping subnetwork for the network by default. This subnetwork is not a subdivision of an existing network. It is purely for ip-addressing purposes. You can override this default and specify subnetwork values directly using the **--subnet** option. On a **bridge** network you can only create a single subnet:

```
$ docker network create --driver=bridge --subnet=192.168.0.0/16 br0
```

Additionally, you also specify the **--gateway --ip-range** and **--aux-address** options.

```
$ docker network create \
  --driver=bridge \
  --subnet=172.28.0.0/16 \
  --ip-range=172.28.5.0/24 \
  --gateway=172.28.5.254 \
  br0
```

If you omit the **--gateway** flag the Engine selects one for you from inside a preferred pool. For **overlay** networks and for network driver plugins that support it you can create multiple subnetworks. This example uses two **/25** subnet mask to adhere to the current guidance of not having more than 256 IPs in a single overlay network. Each of the subnetworks has 126 usable addresses.

```
$ docker network create -d overlay \
  --subnet=192.168.10.0/25 \
  --subnet=192.168.20.0/25 \
  --gateway=192.168.10.100 \
  --gateway=192.168.20.100 \
  --aux-address="my-router=192.168.10.5" --aux-address="my-switch=192.168.10.6" \
  --aux-address="my-printer=192.168.20.5" --aux-address="my-nas=192.168.20.6" \
  my-multihost-network
```

Be sure that your subnetworks do not overlap. If they do, the network create fails and Engine returns an error.

Bridge driver options

When creating a custom network, the default network driver (i.e. **bridge**) has additional options that can be passed. The following are those options and the equivalent docker daemon flags used for docker0 bridge:

Option	Equivalent	Description
com.docker.network.bridge.name	-	Bridge name to be used when creating the Linux bridge
com.docker.network.bridge.enable_ip_masquerade	**--ip-masq**	Enable IP masquerading
com.docker.network.bridge.enable_icc	**--icc**	Enable or Disable Inter Container Connectivity
com.docker.network.bridge.host_binding_ipv4	**--ip**	Default IP when binding container ports
com.docker.network.driver.mtu	**--mtu**	Set the containers network MTU
com.docker.network.container_iface_prefix	-	Set a custom prefix for container interfaces

The following arguments can be passed to **docker network create** for any network driver, again with their approximate equivalents to **docker daemon**.

Argument	Equivalent	Description
--gateway	-	IPv4 or IPv6 Gateway for the master subnet
--ip-range	**--fixed-cidr**	Allocate IPs from a range
--internal	-	Restrict external access to the network
--ipv6	**--ipv6**	Enable IPv6 networking
--subnet	**--bip**	Subnet for network

For example, let's use **-o** or **--opt** options to specify an IP address binding when publishing ports:

```
$ docker network create \
    -o "com.docker.network.bridge.host_binding_ipv4"="172.19.0.1" \
    simple-network
```

Network internal mode

By default, when you connect a container to an **overlay** network, Docker also connects a bridge network to it to provide external connectivity. If you want to create an externally isolated **overlay** network, you can specify the **--internal** option.

Network ingress mode

You can create the network which will be used to provide the routing-mesh in the swarm cluster. You do so by specifying **--ingress** when creating the network. Only one ingress network can be created at the time. The network can be removed only if no services depend on it. Any option available when creating an overlay network is also available when creating the ingress network, besides the **--attachable** option.

```
$ docker network create -d overlay \
  --subnet=10.11.0.0/16 \
  --ingress \
  --opt com.docker.network.driver.mtu=9216 \
  --opt encrypted=true \
  my-ingress-network
```

Run services on predefined networks

You can create services on the predefined docker networks **bridge** and **host**.

```
$ docker service create --name my-service \
  --network host \
  --replicas 2 \
  busybox top
```

Swarm networks with local scope drivers

You can create a swarm network with local scope network drivers. You do so by promoting the network scope to **swarm** during the creation of the network. You will then be able to use this network when creating services.

```
$ docker network create -d bridge \
  --scope swarm \
  --attachable \
  swarm-network
```

For network drivers which provide connectivity across hosts (ex. macvlan), if node specific configurations are needed in order to plumb the network on each host, you will supply that configuration via a configuration only network. When you create the swarm scoped network, you will then specify the name of the network which contains the configuration.

```
node1$ docker network create --config-only --subnet 192.168.100.0/24 --gateway 192.1
node2$ docker network create --config-only --subnet 192.168.200.0/24 --gateway 192.1
node1$ docker network create -d macvlan --scope swarm --config-from mv-config --atta
```

OPTIONS

--attachable[=false] Enable manual container attachment

--aux-address=map[] Auxiliary IPv4 or IPv6 addresses used by Network driver

--config-from="" The network from which to copy the configuration

--config-only[=false] Create a configuration only network

-d, **--driver**="bridge" Driver to manage the Network

--gateway=[] IPv4 or IPv6 Gateway for the master subnet

--ingress[=false] Create swarm routing-mesh network

--internal[=false] Restrict external access to the network

--ip-range=[] Allocate container ip from a sub-range

--ipam-driver="default" IP Address Management Driver

--ipam-opt=map[] Set IPAM driver specific options

--ipv4[=true] Enable or disable IPv4 address assignment

--ipv6[=false] Enable or disable IPv6 address assignment

--label= Set metadata on a network

-o, **--opt**=map[] Set driver specific options

--scope="" Control the network’s scope

--subnet=[] Subnet in CIDR format that represents a network segment

SEE ALSO

docker-network(1)