

NAME

fanctl – fan bridge administration

SYNOPSIS

```
fanctl up [<options>...]
fanctl down [<options>...]
fanctl up -a
fanctl down -a
fanctl down -e
fanctl show
fanctl help
```

DESCRIPTION

fanctl is used to set up, tear down, and inspect Fan bridge mappings and devices in the linux kernel.

NOTE:

Future versions of the commands will make **fanatic** the primary interface for interacting with Fan networks. For maximum portability **fanctl** should only be used if absolutely necessary.

A network Fan is a mechanism for expanding the range of IP addresses available to a system. It is most useful for containers such as Docker and LXC/LXD, but it can be used in other contexts as well. Fan works by creating a bridge that uses a mathematical mapping between the host's (or underlay's) /16 address and the Fan's (or overlay's) /8 address. By mapping addresses in this way, a 253-fold increase in address space can be achieved. For example, if the host machine uses a subnet of 172.16.0.0/16 and assigns a 250.0.0.0/8 Fan to an IP address of 172.16.3.4, the host's Fan overlay addresses will be in the 250.3.4.0/24 subnet, where 250 is derived from the user defined overlay network prefix.

COMMAND SYNTAX**fanctl up -u <underlay> -o <overlay> [<options>]**

Sets up a new Fan bridge mapping **overlay** addresses to the corresponding **underlay** addresses on the local network. Using the example, the new bridge is named *fan-250* based on the overlay and underlay addresses specified (see ADDRESSING below). The **options** are described in the OPTIONS section below.

fanctl down -u <underlay> -o <overlay>

Tears down a previously-configured Fan bridge and associated mapping. This action will fail if the bridge is still in use.

fanctl up -a

Sets up all Fans defined in */etc/network/fan* if present.

fanctl down -a

Tears down all automatically defined Fan bridges on the system. These may be identified in the *fanctl show* output via the *auto* flag.

fanctl down -e

Tears down all defined Fan bridges on the system.

fanctl show

Lists all currently defined Fan bridges in the system; for instance:

```
# fanctl show
Bridge   Underlay   Overlay   Flags
fan-250  172.16.3.4/16  250.0.0.0/8  dhcp host-reserve 1
```

fanctl config set -u <underlay> -o <overlay> [<options>...]

Sets configuration options for the underlay/overlay combination specified (or clears any individual configuration if no options are specified).

To add options, the current set of options has to be obtained via **fanctl config show** and added to the list of new options.

fanctl config show -u <underlay> -o <overlay>

Displays any additional configuration options for the underlay/overlay combination specified.

fanctl config list

Displays all underlay/overlay combinations which have additional configuration.

fanctl help [<command>]

Displays basic usage information for **fanctl**. When used with a specific command limits output to that command.

ADDRESSING

The Fan mapping is defined by a combination of the **underlay** and **overlay** addresses. Each is defined as a CIDR network address. For example:

```
# fanctl up -u 172.16.3.4/16 -o 250.0.0.0/8
```

This example defines an overlay of *250.0.0.0/8* and an underlay of *172.16.3.4/16*. When mapping an address in the *250.0.0.0/8* subnet, we take the 16 bits of destination address starting at bit 8 and replace the bottom 16 bits of the underlay address with it. For example, attempting to talk to *250.3.4.15* will trigger the packet to be sent to *172.16.3.4* for delivery.

It is not always possible to know the local underlay address at the time the configuration is generated (such as when a common configuration is desired on all systems). In this case we can specify the underlay address using only the underlay prefix, or by reference to an interface.

For example, to bring up a Fan bridge slice for each address in the *172.16.0.0/16* subnet, the following example examines each interface as it is currently configured and configures a matching slice at that time:

```
# fanctl up -u 172.16.0.0/16 -o 250.0.0.0/8
```

To bring up Fan slices corresponding to the addresses on a specific interface we can substitute the interface name:

```
# fanctl up -u ens3/16 -o 250.0.0.0/8
```

To bring up Fan slices corresponding to the addresses on the primary network interface (the interface with the default route), the keyword *default* can be substituted:

```
# fanctl up -u default/16 -o 250.0.0.0/8
```

LIMITATIONS

Currently Fan can only apply overlay addresses with a /8 network mask, and underlay addresses with a /16 network mask. We expect to relax this limitation in a later update.

OPTIONS**-u <underlay>|--underlay=<underlay>**

Specify the underlay network address and mask.

-o <overlay>|--overlay=<overlay>

Specify the overlay network address and mask.

--type=<type>

Sets the encapsulation type for this Fan bridge. May be *ipip* or *vxlan* (default) only. You should not normally need to specify this.

--mode=<mode>

Sets the bridge mode. May be *compact* (default) or *sliced* only. In *compact* mode a single Fan bridge per overlay network is created. In *sliced* mode a Fan bridge is created for each local Fan slice. *sliced* mode is considered legacy. You should not normally need to specify this.

--dhcp Turns on automatic address allocation for the Fan bridge. A *dnsmasq* instance is started attached to the bridge allocating the unreserved addresses to entities attached to the Fan bridge.

--host-reserve=<count>

By default the *.1* address on the Fan bridge is allocated to the host, enabling it to communicate with entities on the Fan bridge. This option reserves additional addresses for host applications to

use. A *host-reserve 4* reserves .1 through .4 in the Fan bridge for host use.

--bridge=<name>

By default the bridge name is based on the overlay and underlay addresses specified (see ADDRESSING above). This option overrides the name to one you specify.

--enable

This option marks the Fan bridge to be automatically configured when the underlying interface comes up. This is primarily used in the local Fan persistent configuration to enable centrally configured Fan mappings on a specific host. (See PERSISTENT CONFIGURATION below).

PERSISTENT CONFIGURATION

NOTE: The location and content of the configuration files should still be considered unstable as they could change in future. Use of **fanatic** is highly recommended as this command will make any necessary changes to the right config files.

Fan mappings are configured via */etc/network/fan* by pairs of local (underlay) and overlay network addresses. The first element of each pair specifies the underlay range which should be mapped into the overlay network specified by the second element. For example:

```
# RFC1918 - we recommend you use these for easy interop with
# other FAN users on small private networks. They provide
# around 250 IP's per 192.168.0.0/16 address, or 16 per
# 172.16.0.0/12 address.

# local    overlay
192.168.0.0/16 250.0.0.0/8
172.16.0.0/12 251.0.0.0/8
```

Comments are introduced via a hash (#), and blank lines are ignored.

Note, that the local (underlay) range does not need to match the subnet mask of the interface as long as all addresses of the larger scope are routeable. If for example the interface is configured as 192.168.122.2/24 but can reach all 192.168.0.0/16 addresses that will be part of the Fan network, the local (underlay) address of 192.168.0.0/16 can be used regardless.

It is expected that the */etc/network/fan* configuration is globally managed ensuring that all hosts have consistent overlay to underlay mappings. Local deviation is managed via the **fanctl config** subcommand. This allow a local host to record additional flags against a specific overlay/underlay combination. For example:

```
# fanctl config set -u 172.16.0.0/16 -o 250.0.0.0/8 --enable
```

will set the **--enable** option to the local host configuration, triggering this Fan to be configured when the host interface is configured.

FAN BRIDGES

By default each Fan bridge represents a Fan overlay network which is expressed locally on the machine. The Fan bridge will have the various slice addresses mapped to it.

In legacy sliced mode each Fan bridge represents a slice of a Fan overlay network which is expressed locally on the machine. The Fan bridge will have the overlay addresses representing one local IP address mapped to it. A machine may have more than one local address on the underlay network, enabling it to have more than one such slice mapped. It may also have more than one overlay range defined for each local IP address.

Each Fan bridge is a separate broadcast domain, with routing between the bridges both locally and globally within the Fan.

Each Fan bridge appears as a bridge on the system, named for the overlay subnet hosted by that particular Fan bridge and the underlay address prefix for which it carries traffic. For our *250.0.0.0/8* on *172.16.3.4* example, the bridge would be named *fan-250* and would carry all traffic for *250.3.4.0/24*.

SEE ALSO

fanatic(8), `/usr/share/doc/ubuntu-fan/README`

AUTHOR(s)

Andy Whitcroft <apw@canonical.com>,
Stefan Bader <stefan.bader@canonical.com>