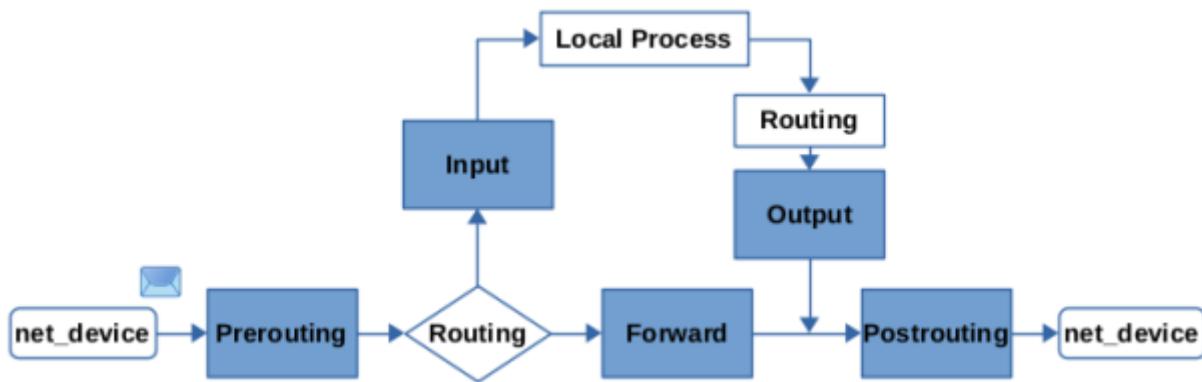


# NFTables

## Hooks

Every packet that enters a system, whether incoming or outgoing will trigger some hooks as it traverses through the Linux kernel's networking stack. Those five hooks have been present in the Linux kernel for a very long time. Linux kernel allows rules that are associated with these hooks to interact with the network traffic. Nftables has five hooks including prerouting, input, output, post routing, forward, and ingress.



When traffic flow goes into a local machine, first, it faces the prerouting hook and then input hook. Next, the traffic generated by the local machine's processes follows the output hook and then the postrouting hook as shown in the next figure. The packets destined to your network but are not addressed to the local node will face the forward hook after following prerouting and then postrouting path. Ingress hook, however, as a new hook in nftables, is a hook that is placed before all the hooks behind the prerouting hook and can filter traffic on layer 2 OSI model. With this hook, therefore, early filtering policies can be defined (2019).

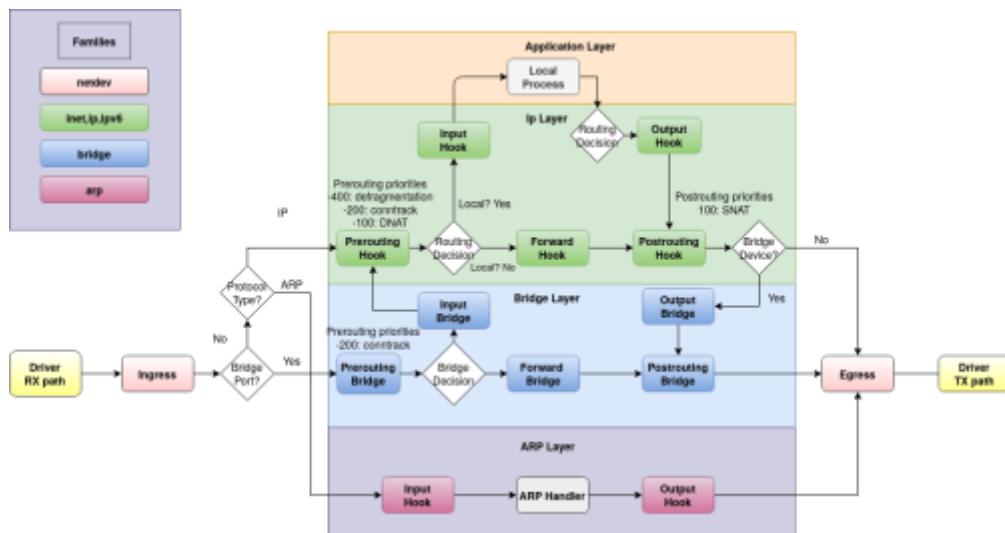
## Address Families

Address families determine the type of incoming and outgoing packets processed by nftables. For each address family, the Linux kernel contains specific hooks at different stages of the packet processing paths, which invoke nftables to decide either allow or drop a packet only if relevant rules for these hooks such as input or output are defined. These address families are as follows:

ip:	IPv4 address family.
ip6:	IPv6 address family.
inet:	Supports Both IPv4 and IPv6 address families.
arp:	ARP address family, handling IPv4 ARP packets at layer 2 OSI model.
bridge:	Bridge address family, handling packets traversing a bridge device at layer 2.
netdev:	Netdev address family, handling packets from ingress hook working before layer 3.

Table with netdev family, by means of ingress hook, allows early filtering traffic before they reach

other filters below layer 3 on the OSI model. netdev family with ingress hook is an ideal stage to drop packets that result from DDOS attacks since this hook works very early in the packet path of networking. There are different hooks for different “family” types. The following schematic shows packet flows through Linux networking:



## Priority

Nftables requires you to specify a priority value when creating a base chain. You can specify integer values, but the newer versions of Nftables also define placeholder names for several discrete priority values analog to the mentioned enums in Netfilter. The following table lists those placeholder names<sup>12</sup>).

Name	Priority Value
raw	-300
mangle	-150
conntrack <sup>13</sup>	-200
dstnat	-100
filter	0
security	50
srcnat	100

When creating a base chain, you can e.g. specify priority filter which translates into priority 0. The following example creates a table named myfilter in the ip address family (IPv4). It then creates two base chains named foo and bar, registering them with the Netfilter IPv4 hook input, but each with different priority.

```

nft create table ip myfilter
nft create chain ip myfilter foo {type filter hook input priority 0\;}
nft create chain ip myfilter bar {type filter hook input priority 50\;}

# alternatively you could create the same chains using named priority
# values:
nft create chain ip myfilter foo {type filter hook input priority filter\;}
nft create chain ip myfilter bar {type filter hook input priority
security\;}

```

## Links

- [wiki.nftables.org](https://wiki.nftables.org)

From:  
<https://wiki.oscardegroot.nl/> - **HomeWiki**

Permanent link:  
<https://wiki.oscardegroot.nl/doku.php?id=networking:nftables&rev=1693727239>

Last update: **2023/09/03 07:47**

