

Lab — iptables and nftables, hands-on

Lesson 01 · ~10 min · Single host, no extra VMs

Warning

Every command in this lab requires `sudo` and modifies your live firewall. Run on a throwaway VM if you can. To undo everything, `iptables -F` and `nft flush ruleset`.

Step 1 — Read your current state

Before you change anything, see what's there:

```
sudo iptables -L -v -n
sudo iptables -t nat -L -v -n
```

Most distros ship with empty default chains and policy `ACCEPT`.

Step 2 — Drop ICMP, watch ping die

```
sudo iptables -A INPUT -p icmp -j DROP
ping -c 3 127.0.0.1
# Result: 100% packet loss
```

Step 3 — Remove the rule, ping works again

```
sudo iptables -D INPUT -p icmp -j DROP
ping -c 3 127.0.0.1
# Result: 0% packet loss
```

Step 4 — Same scenario in nftables

```
sudo nft add table inet f
sudo nft 'add chain inet f input { type filter hook input priority 0 ; }'
sudo nft add rule inet f input ip protocol icmp drop
ping -c 3 127.0.0.1
# Result: dropped

sudo nft flush ruleset
ping -c 3 127.0.0.1
# Result: works again
```

Step 5 — Translate iptables to nftables

```
iptables-translate -A INPUT -p tcp --dport 22 -j ACCEPT
# Output:
# nft add rule ip filter INPUT tcp dport 22 counter accept
```

Cleanup

`sudo iptables -F && sudo nft flush ruleset` resets you to a clean slate.

Self-check

- ✓ You can read the iptables rules table.
- ✓ You can add and remove a rule and observe its effect.
- ✓ You have written the same rule in both iptables and nftables.
- ✓ You have used `iptables-translate`.