Network Basics

System and Network Administration

Revision 2 (2020/21)

Pierre-Philipp Braun <pbraun@nethence.com>

Table of contents

MWE Routing

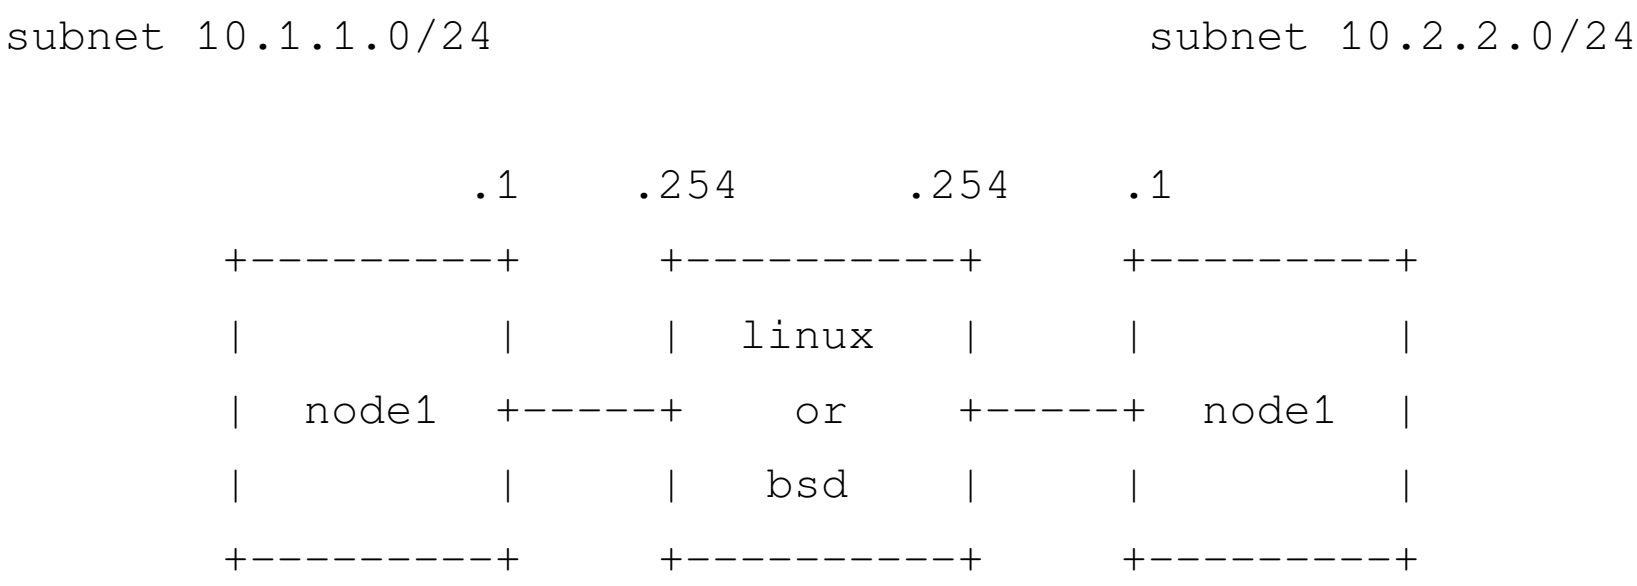

How to turn a UNIX system into a router?...

==> enable IP forwarding

GNU/Linux

```
#sysctl -w net.ipv4.ip_forward=1
echo 1 > /proc/sys/net/ipv4/ip_forward

#echo net.ipv4.ip_forward=1 >> /etc/sysctl.conf
#sysctl -p
```

BSD

```
sysctl -w net.inet.ip.forwarding=1
sysctl net.inet.ip.forwarding

echo net.inet.ip.forwarding=1 >> /etc/sysctl.conf
```

What to do next for the two subnets to talk to each other?…

==> enable static (or default) routes

configure the nodes to use the box/router as

- static route
- –or– default route

Note: both nodes need to be tweaked – otherwise the there would be no return path for e.g. an ICMP reply

Note: that works only from the next hop (not through the public network)

What's the most common scenario for a public network gw?…

==> NAT

Translating source or destination

- SNAT – outbound
  - traffic coming from internal subnet is translated to front-facing IP
  - not supposed to be reachable
- DNAT – inbound (port-forwarding)
  - traffic coming to front-facing IP gets translated to internal subnet
  - reachable by design

Do we absolutely need to enable firewalling for NAT to work?…

==> technically speaking, no

- Forwarding + SNAT is enough
- …and it is *almost* ok, as long as the gateway itself is clean
- …meaning it is not listening on any port on the front-facing interface

==> but sometimes, yes

- In case you need a firewall anyways to handle the internal network
- The gateway is ideal place to do it

And if you really need to enable Firewalling…

DO NOT *FULLY* DISABLE ICMP – IT IS USEFUL

```
==> /var/log/debug <==
Jan 16 06:30:17 slack9 dhcpd: ICMP Echo reply while lease
    10.1.1.145 valid.

==> /var/log/syslog <==
Jan 16 06:30:17 slack9 dhcpd: Abandoning IP address
    10.1.1.145: pinged before offer
```

Netfilter with IPTABLES

Second, *SNAT* on a static and front-facing IP

```
iptables -t nat -A POSTROUTING -o FACING-NIC -s INTERNAL-CIDR
    -j SNAT --to-source FACING-IP
```

–or– on a changing and front-facing IP

```
iptables -t nat -A POSTROUTING -o FACING-NIC -s INTERNAL/CIDR
    -j MASQUERADE
```

check

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

Netfilter with NFTABLES

with a STATIC IP

```
vi /etc/nftables.conf

...
    #SNAT
    chain postrouting {
            type nat hook postrouting priority 100;
            oifname eth0 masquerade
    }
```

with a DYNAMIC IP

```
...
    #SNAT
        chain my_masquerade {
           type nat hook postrouting priority srcnat; policy accept;
                oifname ”ppp0” masquerade
        }

systemctl reload nftables
```

NetBSD Packet Filter (NPF)

```
vi /etc/npf.conf

group default {
        pass in all
        pass out all
}

#SNAT
map xennet0 dynamic 10.1.1.0/24 -> 188.130.155.62

/etc/rc.d/npf reload
```

Besides, NPF is not vulnerable to NAT pivoting.

Now consider your home router, and let's say you want to do some peer-to-peer.

What do you need to enable here and what is it called?…

==> DNAT aka PORT-FORWARDING

DNAT with IPTABLES

```
iptables -t nat -A PREROUTING -i eth0 -p tcp --dport 80 -j DNAT --to-destination INTERNAL-IP
```

Note eventually against another port with `INTERNAL-IP:PORT`

DNAT with NFTABLES

```
vi /etc/nftables.conf

...
    #DNAT
    chain prerouting {
            type nat hook prerouting priority -100;
            iifname eth0 tcp dport 80 dnat x.x.x.x
    }

systemctl reload nftables
```

DNAT with NPF

```
vi /etc/npf.conf

map xennet0 dynamic proto tcp 10.1.1.x port xxxxx <-
188.130.155.62 port xxxxx

/etc/rc.d/npf reload
```

eBPF

LAB // dig into eBPF and PoC

// Questions on mwe routing?

STP & VLAN

What's the difference between perimeter and DMZ?…

==> front-facing vs NAT

network topology

Perimeter

- got public IP
- typically this is where your gateway lives
- using a default route against your ISP's gw…

DMZ

- DNAT & SNAT routed
- DNAT & isolated

Tag vs. physical VLAN

let's split our switch!

Terminology depends on hardware vendor

HPE

`tagged / untagged`

Cisco

`trunk / access mode`

Spanning tree?

(show on asciiflow…)

LAB // PoC & sniff STP on Linux bridge vs OpenvSwitch

LAB // Evaluate the 30 seconds delay caused by STP and try to remediate

Switching products & network emulation

- Packet Tracer
- GNS3
- EVE-NG
- VirtualBox
- DIY & Linux Bridge
- DIY & OpenvSwitch

Commercial router systems

- Cisco IOS, IOS XR
- Juniper …
- Mikrotik RouterOS (free-of-charge and Linux-based)
- …

Open Source router systems

Enterprise-class

- OpenBGPD
- OpenOSPFD
- Cumulus Linux (got ASIC & DPDK!)
- (FreeBSD got ASIC & DPDK?)

End-user / mid-size

- m0n0wall –> pfSense –> OPNsense (FreeBSD)
- OpenWRT/LEDE (Linux)

// Questions on stp and vlan?

Linux Bonding

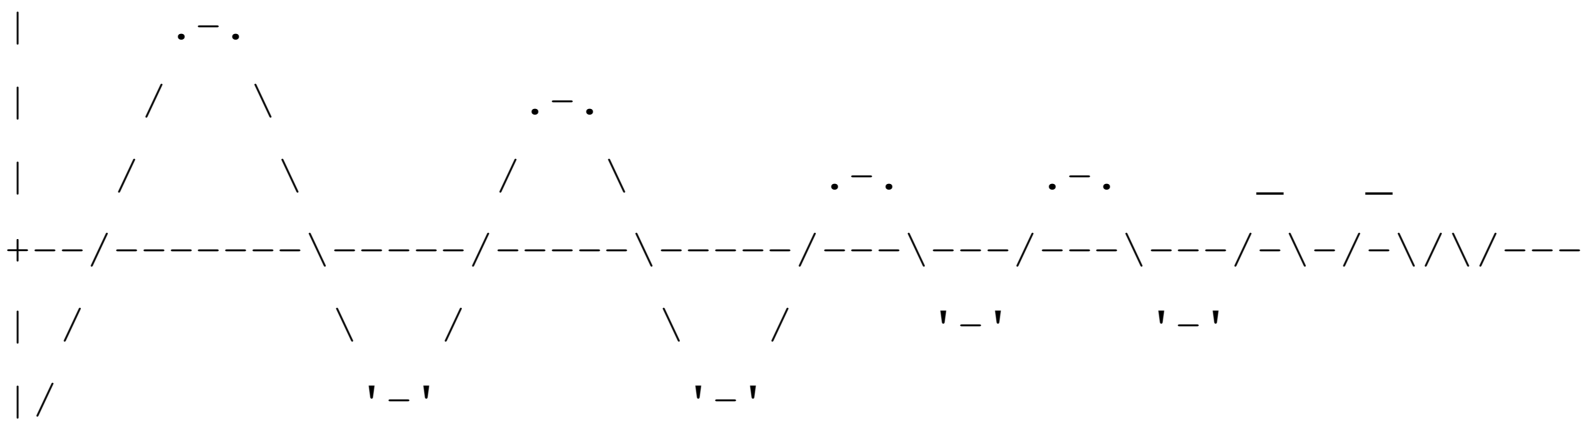

Managed vs. un-managed switch

Static port trunk

`balance-rr`

`balance-xor`

Dynamic port trunk

`802.3ad`

Un-managed switch is fine for those

`balance-tlb`

`balance-alb (also RX)`

Linux Bonding modes

```
0 balance-rr          lbs & ha
1 active-backup       active/passive
2 balance-xor         lbs/xmit & ha
3 broadcast       ha
4 802.3ad         lbs & ha
5 Balance-tlb         lbs & ?
6 balance-alb         lbs & ?
```

BONUS QUESTION // how come round-robin and XOR provide HA here?

Linux Bonding - the deprecated way

```
#vi /etc/modprobe.conf
vi /etc/modprobe.d/bonding.conf

alias bond0 bonding
options bond0 miimon=100 mode=X <other option=...>

ifenslave bond0 eth0
ifenslave bond0 eth1
```

check

```
ifenslave -a
```

Linux Bonding - the new way

```
modprobe bonding
echo 100 > /sys/class/net/bond0/bonding/miimon
echo 200 > /sys/class/net/bond0/bonding/downdelay
echo 200 > /sys/class/net/bond0/bonding/updelay
echo X > /sys/class/net/bond0/bonding/mode
echo ... > /sys/class/net/bond0/bonding/other_option
#echo layer3+4 > /sys/class/net/bond0/bonding/xmit_hash_policy
echo +eth0 > /sys/class/net/bond0/bonding/slaves
echo +eth1 > /sys/class/net/bond0/bonding/slaves
```

Status

```
cat /sys/class/net/bonding_masters
cat /proc/net/bonding/bond0
cat /sys/class/net/bond0/bonding/miimon
cat /sys/class/net/bond0/bonding/downdelay
cat /sys/class/net/bond0/bonding/updelay
cat /sys/class/net/bond0/bonding/mode
cat /sys/class/net/bond0/bonding/other_option
cat /sys/class/net/bond0/bonding/xmit_hash_policy
```

Acceptance testing

How to validate

- unplug / replug…
- iPerf3 (does upload/download)
- UDP vs TCP

What about max bandwidth

- multiple iPerf3 instances…

Linux Teaming

- alternative to Bonding
- user-space daemon

LAB // try-out and validate Linux Teaming

LAB // benchmark Linux Teaming vs. Bonding

FDX vs. HDX

- Full-duplex – dedicated cable for TX/RX
- Half-duplex – only one cable

LAB // search and dig into Half-duplex driver modes

FDX validation

Through a 100Mbit/s poor switch

```
94.1 Mbit/s if only one direction
91.5 Mbit/s with both direction at the same time
```

With direct 1Gbit/s link

```
940 Mbit/s if only one direction
930 Mbit/s with both directions at the same time
```

// Questions on linux bonding?

SSH Tips & Tricks

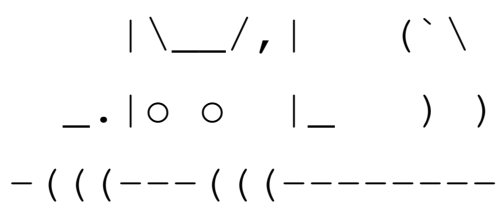

Daemon tuning

Public network

- define your port outside the top 1000 range so attack's quick discoveries won't find your daemon
- disable password authentication
- specify a single and enhanced host key
- many other options – further tune it like hell

Internal network

- listen only on the mgmt/backup VLAN
- same goes for a DIY gateway – listen only on the internal interface

SSH hardening is a good thing

```
Oct 11 13:17:59 pro5s2 sshd[28085]: Unable to negotiate with 51.91.7
rsa,ssh-dss [preauth]
Oct 11 13:19:32 pro5s2 sshd[28095]: Unable to negotiate with 51.91.7
rsa,ssh-dss [preauth]
Oct 11 13:21:03 pro5s2 sshd[28098]: Unable to negotiate with 51.91.7
rsa,ssh-dss [preauth]
Oct 11 13:22:36 pro5s2 sshd[28101]: Unable to negotiate with 51.91.7
rsa,ssh-dss [preauth]
Oct 11 13:24:08 pro5s2 sshd[28104]: Unable to negotiate with 51.91.7
rsa,ssh-dss [preauth]
Oct 11 13:25:40 pro5s2 sshd[28106]: Unable to negotiate with 51.91.7
rsa,ssh-dss [preauth]
Oct 11 13:27:10 pro5s2 sshd[28111]: Unable to negotiate with 51.91.7
rsa,ssh-dss [preauth]
Oct 11 13:27:25 pro5s2 sshd[28113]: Connection closed by 68.183.181.
Oct 11 13:27:44 pro5s2 sshd[28115]: Connection closed by 93.2.194.25
```

Public network

```
AllowUsers root user1 ...
AuthenticationMethods publickey
ChallengeResponseAuthentication no
HostKey /etc/ssh/ssh_host_ed25519_key
MaxAuthTries 3
PasswordAuthentication no
PermitEmptyPasswords no
PermitRootLogin without-password
Port SOME-EXOTIC-PORT-NOT-TOP-1000
PrintMotd no
Protocol 2
StrictModes yes
UseDNS no
UsePAM no
Subsystem sftp /usr/libexec/sftp-server
X11Forwarding no
```

Internal network or gateway – listen only there

```
AddressFamily inet
ListenAddress x.x.x.x
#AllowUsers root@CLIENT-IP gollum@CLIENT2 *@CIDR
```

Client tuning

```
vi /etc/ssh/ssh_config

Host *
        HashKnownHosts no
        GSSAPIAuthentication no
        VisualHostKey yes
```

Note on virtualization

Don’t forget to re-generate host keys

- When deploying guest templates
- When bootstrapping / terraforming / …

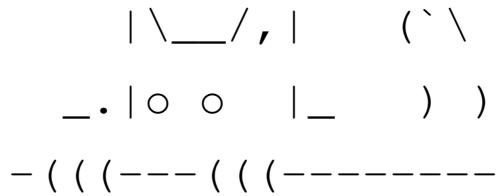

// Questions on those tips & tricks?

Standard LAB

in case you didn't spot any opportunity

VLAN setup **with hardware**

- reset & fw update
- tagged/untagged
- one person per switch –> PoC STP

don't forget to validate and show proof