PFALZKOM
Pfalzwerke Gruppe

How to built a multi-tenant BGP/EVPN network

Our journey to a modern BGP core

NICOLA VON THADDEN
nicola.vonthadden@pfalzkom.de
@nicoduck

Who is Pfalzkom

- Formerly known as PfalzKom|MAnet
- Operate multiple DCs
- 12.000m² ruled area
- >1.400 km of fibre
- WDM, IP, Ethernet, …

Our old setup

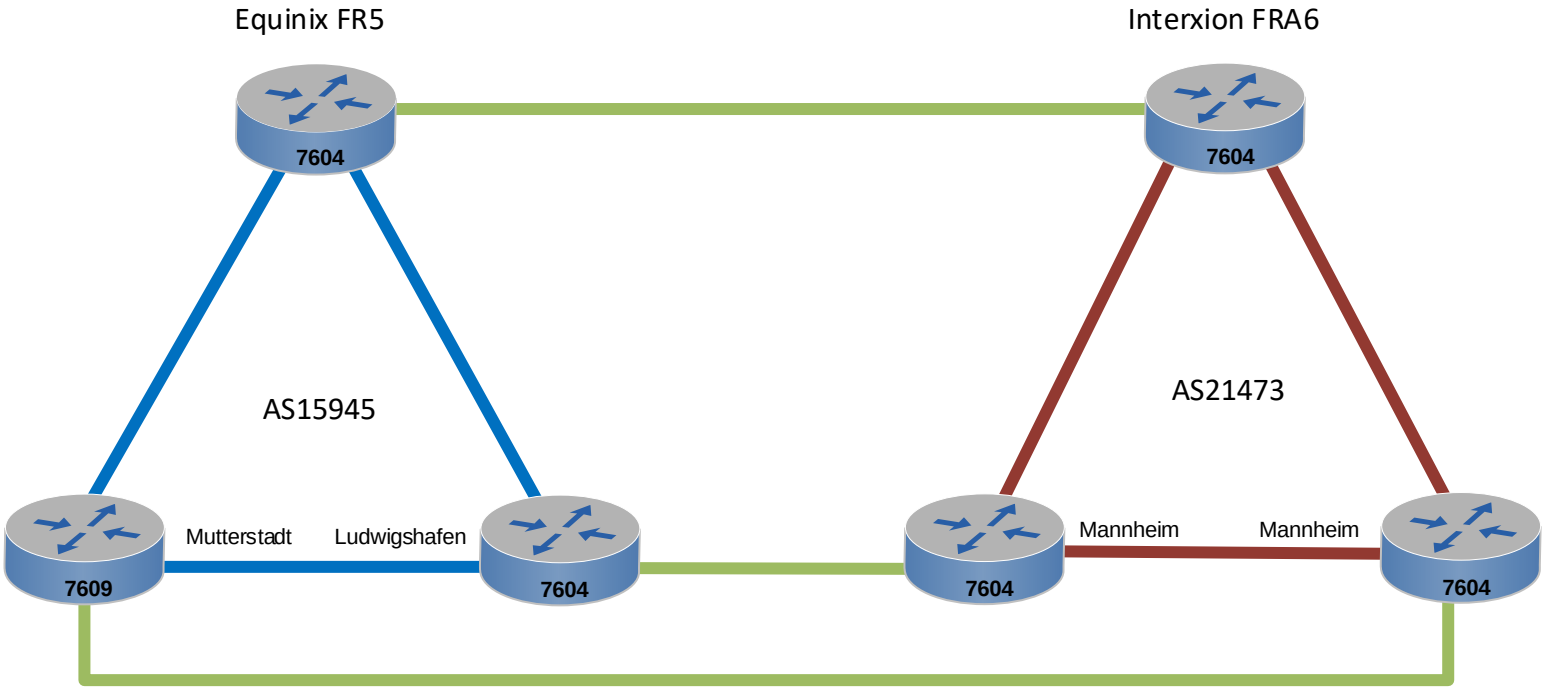

Our old setup

- Cisco 7600
- 10G max
- Aged CPU
- Session flaps at DE-CIX
  - Sorry ☹
- Limited routing table
- Only one AS per box

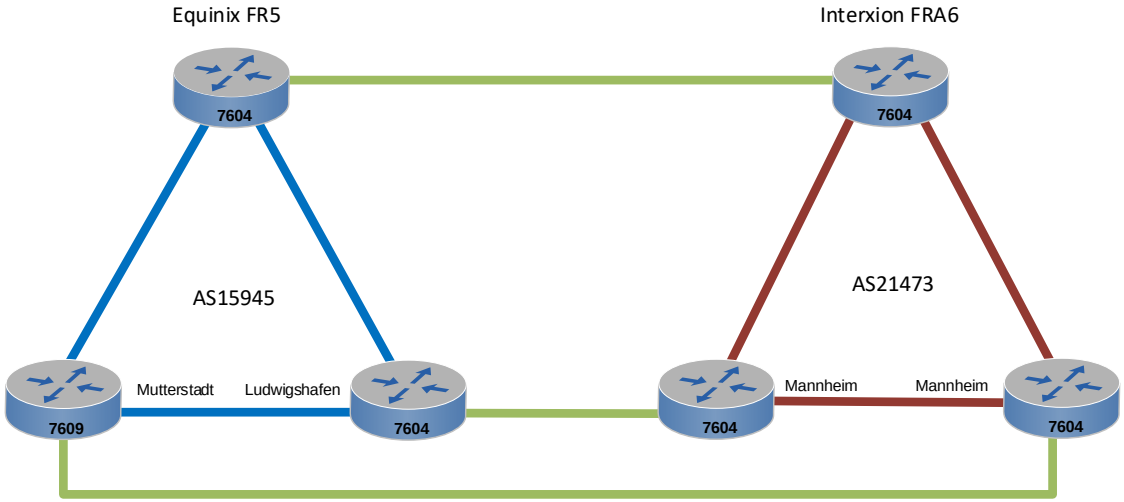

Requirements & considerations for future design

- 100G based
- End-of support not before 2029
- >2M IPv4 routes in FIB
- Option for DDoS mitigation (on or off-chip)
- EVPN support

Requirements & considerations for future design

- 100G based
- End-of support not before 2029
- >2M IPv4 routes in FIB
- Option for DDoS mitigation (on or off-chip)
- EVPN support

- Number of CVEs
- Location of vendors: HQ and production

Future BGP design

- Most customer lines migrate into one AS
- Second AS on separate hardware for special contracts
- Virtualize public routing in one L3-VPN (VPRN)
- Option for BGPaaS

New setup

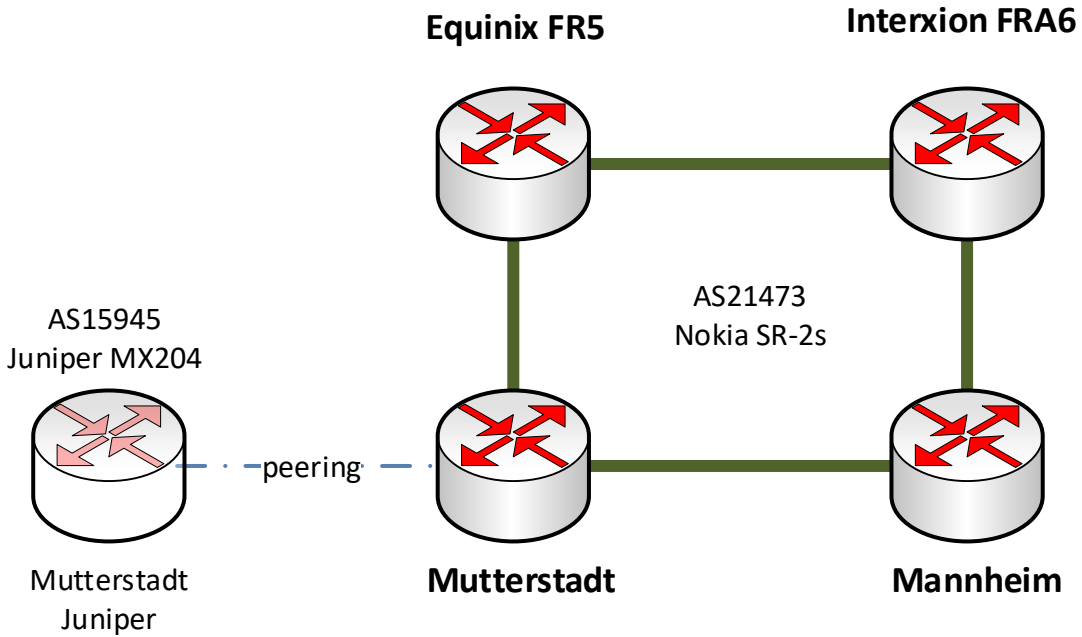
Equinix FR5
Interxion FRA6
AS15945
Juniper MX204
AS21473
Nokia SR-2s
peering
Mutterstadt
Juniper
Mutterstadt
Mannheim

What is EVPN

- Ethernet VPN
- New address family for BGP
- Distributes L2 or L3 addresses between routers (or switches)
- Works with MPLS or VXLAN as packet transport
- One EVPN session can carry information for multiple virtual routers

Topology overview

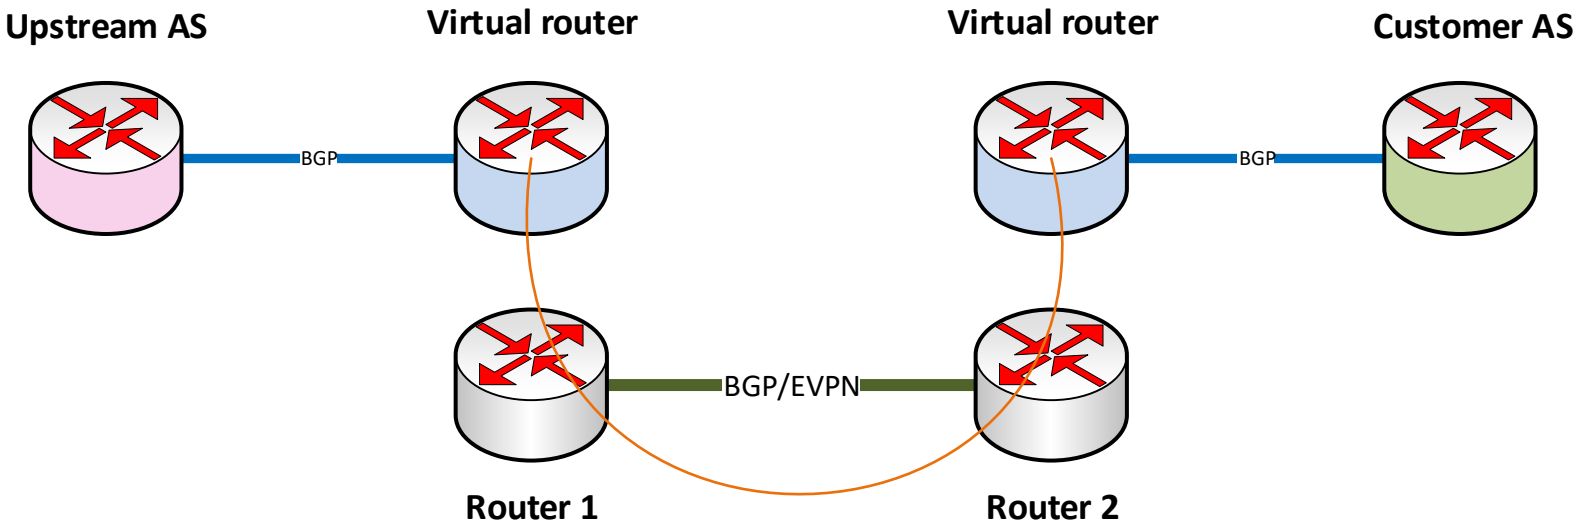

EVPN activation (underlay)

Nokia classic CLI

```
group "Core-MPiBGP"
      family vpn-ipv4 vpn-ipv6
      peer-as 65005
      neighbor xxxx
      exit
  exit
```

Industry standard CLI

```
neighbor xxxx
  remote-as 64601
  send-community extended
  !
  address-family ipv4 unicast
   no activate
  !
  address-family l2vpn evpn
   activate
!
```

Route im-/export (overlay)

Nokia classic CLI

```
configure service vprn 10
            vrf-import "imp_mp-ibpg2vprn"
            vrf-export "exp_vprn2mp-ibgp"
            router-id 192.0.2.152
            autonomous-system 21473
            route-distinguisher 21473:10
```

…

Differences to standard BGP

- Route selection happens on more than one layer
- Second address family to query when looking for routes
- Export policies are a bit different

Topology overview

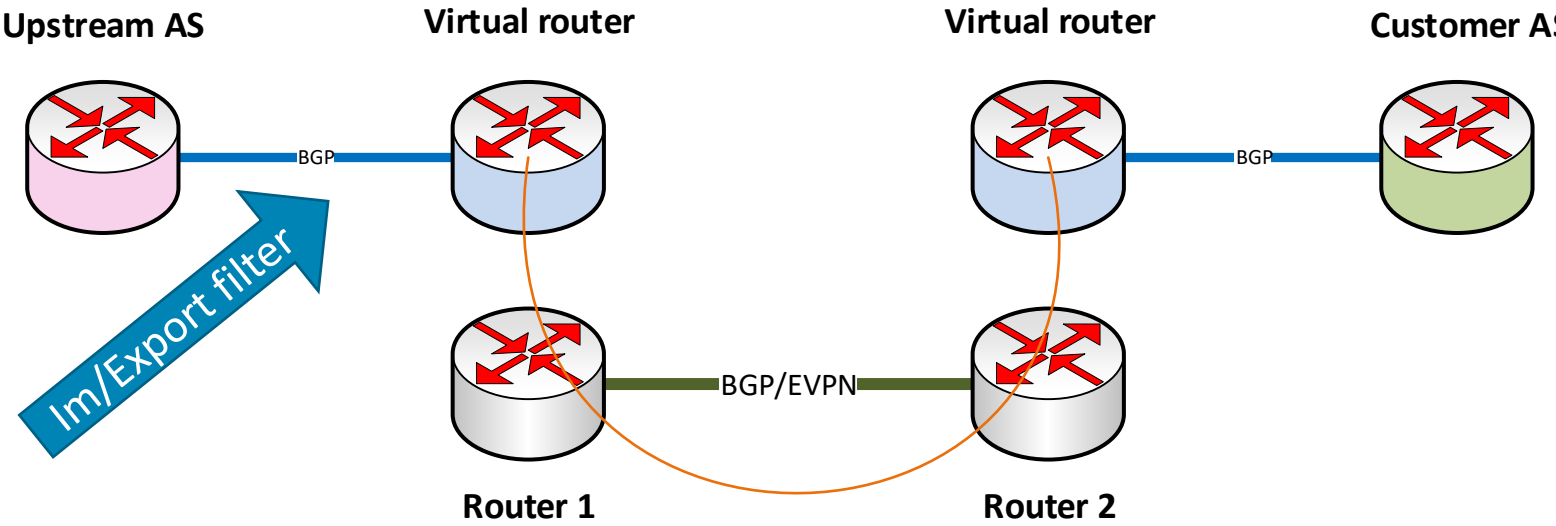

Topology overview

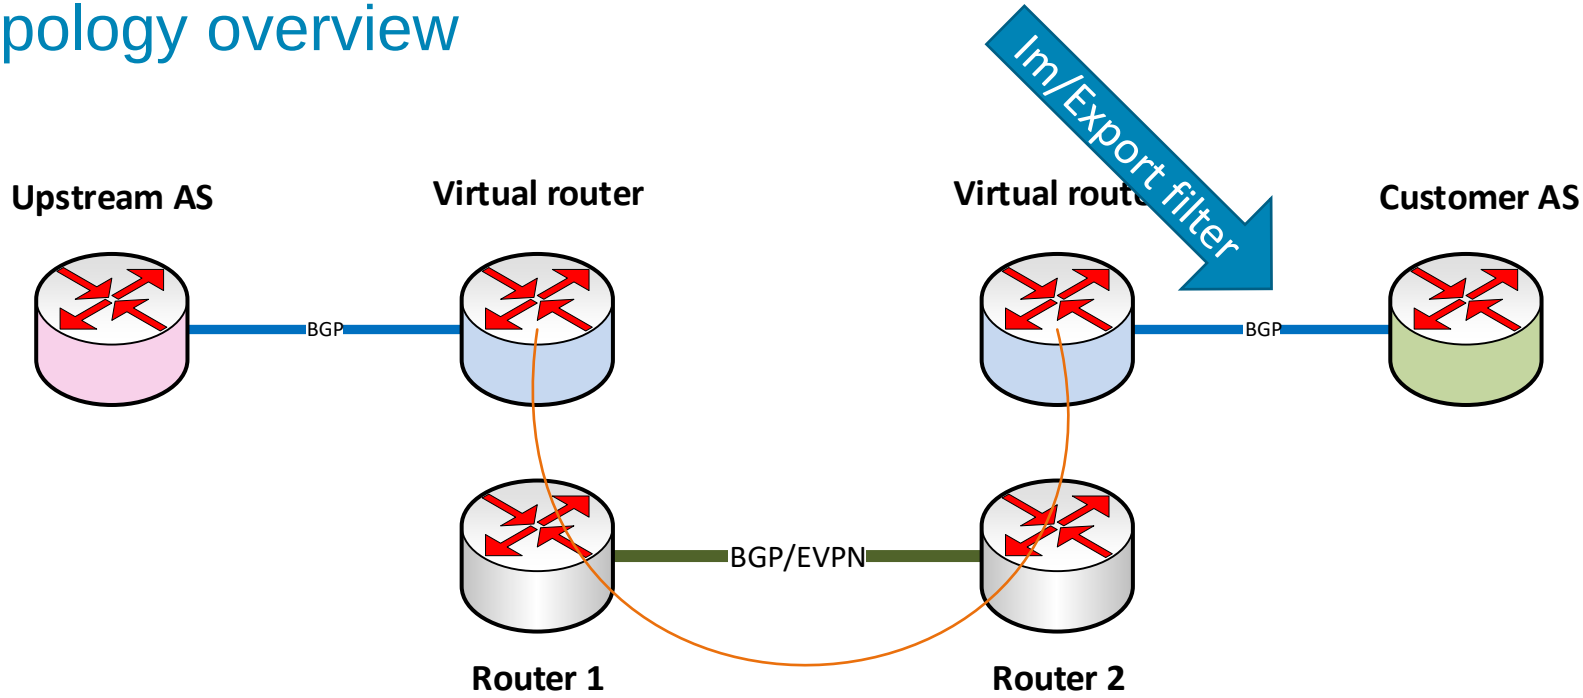

Topology overview

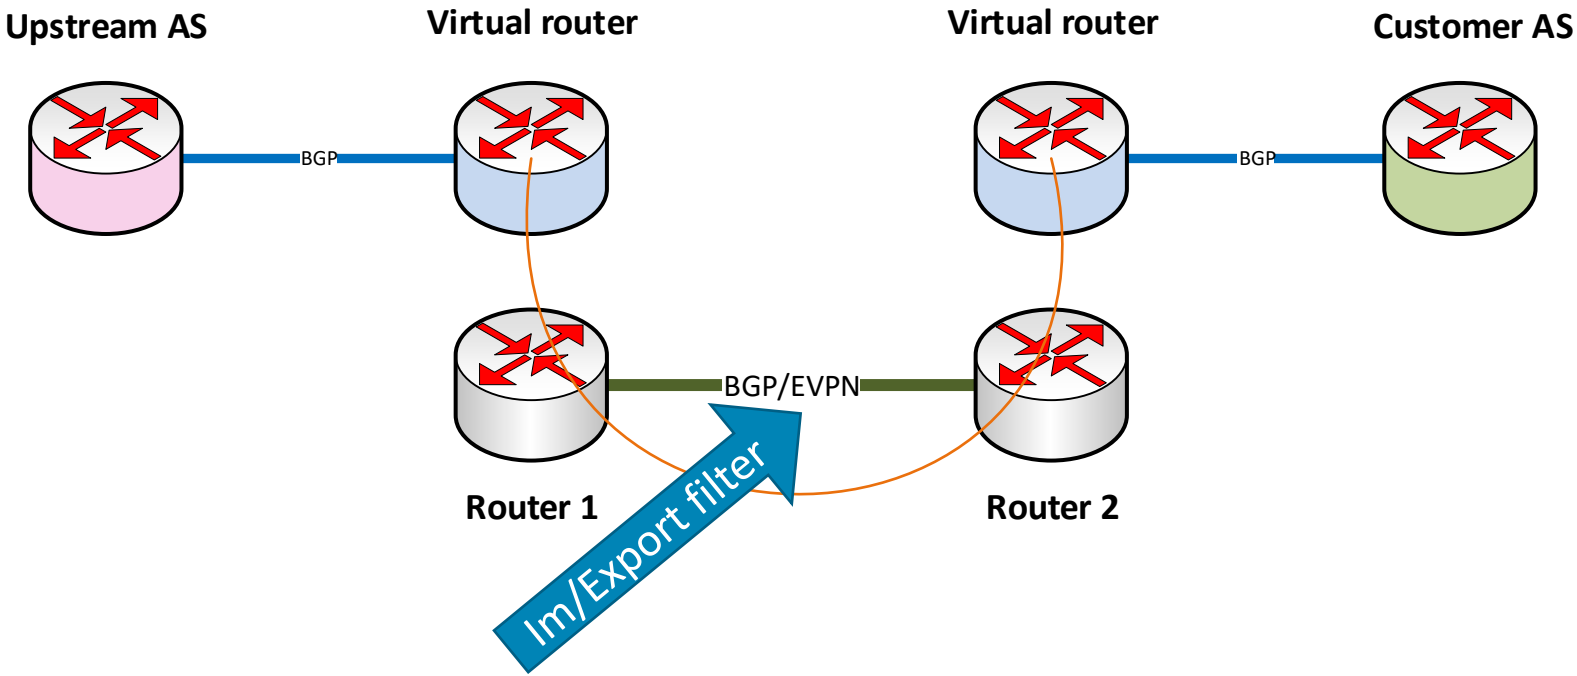

Example route

```
A:# show router 10 route-table 1.1.1.0/24

===============================================================================
Route Table (Service: 10)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
1.1.1.0/24                                      Remote  BGP VPN   16d12h55m  170
       192.0.2.154 (tunneled)                                          10
-------------------------------------------------------------------------------
```

No BGP route found

```
A:# show router 10 bgp routes 1.1.1.0/24
Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete
===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
No Matching Entries Found.
===============================================================================
```

Looking deeper

```
A:# show router bgp routes 1.1.1.0/24 vpn-ipv4
===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  21473:10:1.1.1.0/24                              200         20040
      192.0.2.154                                      None        10
      3356 13335                                                   524286
*i    21473:10:1.1.1.0/24                              200         None
      192.0.2.153                                      None        20
      1299 13335                                                   524286
i     21473:10:1.1.1.0/24                              200         None
      192.0.2.153                                      None        20
      1299 13335                                                   524286
```

Always-compare-med

```
A:# show router bgp routes 1.1.1.0/24 vpn-ipv4
===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     IGP Cost
      As-Path                                                      Label
-------------------------------------------------------------------------------
u*>i  21473:10:1.1.1.0/24                              200         20040
      192.0.2.154                                      None        10
      3356 13335                                                   524286
*i    21473:10:1.1.1.0/24                              200         None
      192.0.2.153                                      None        20
      1299 13335                                                   524286
i     21473:10:1.1.1.0/24                              200         None
      192.0.2.153                                      None        20
      1299 13335                                                   524286
```

MED comparison

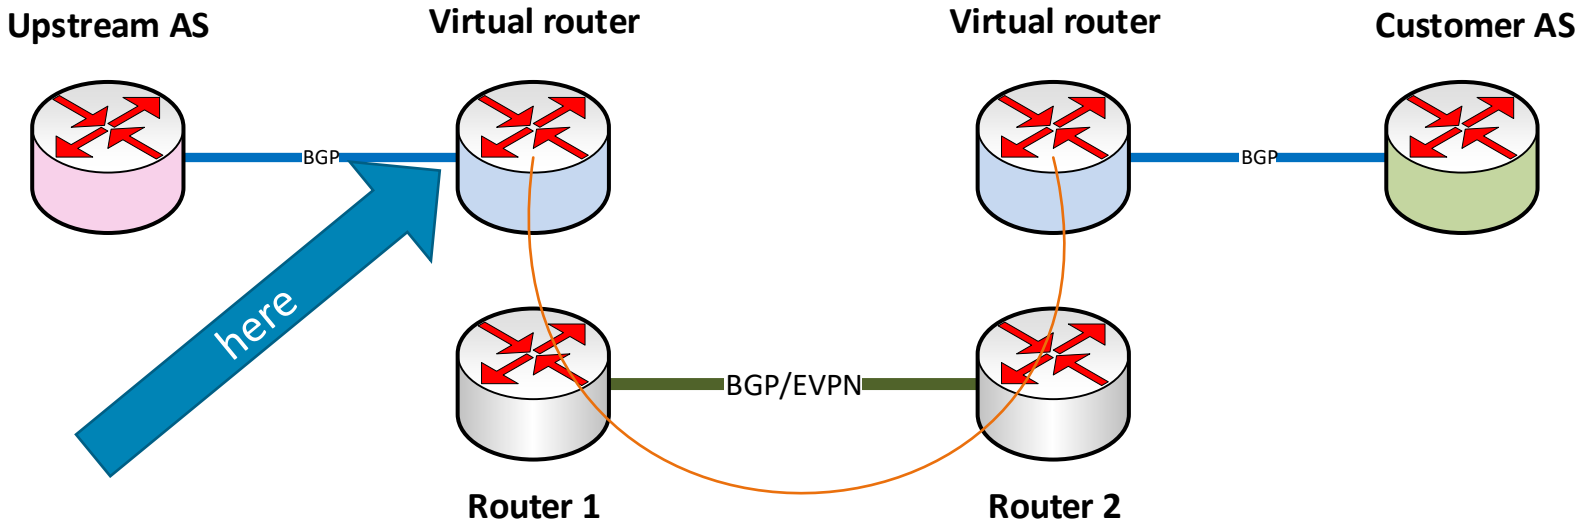

MED comparison

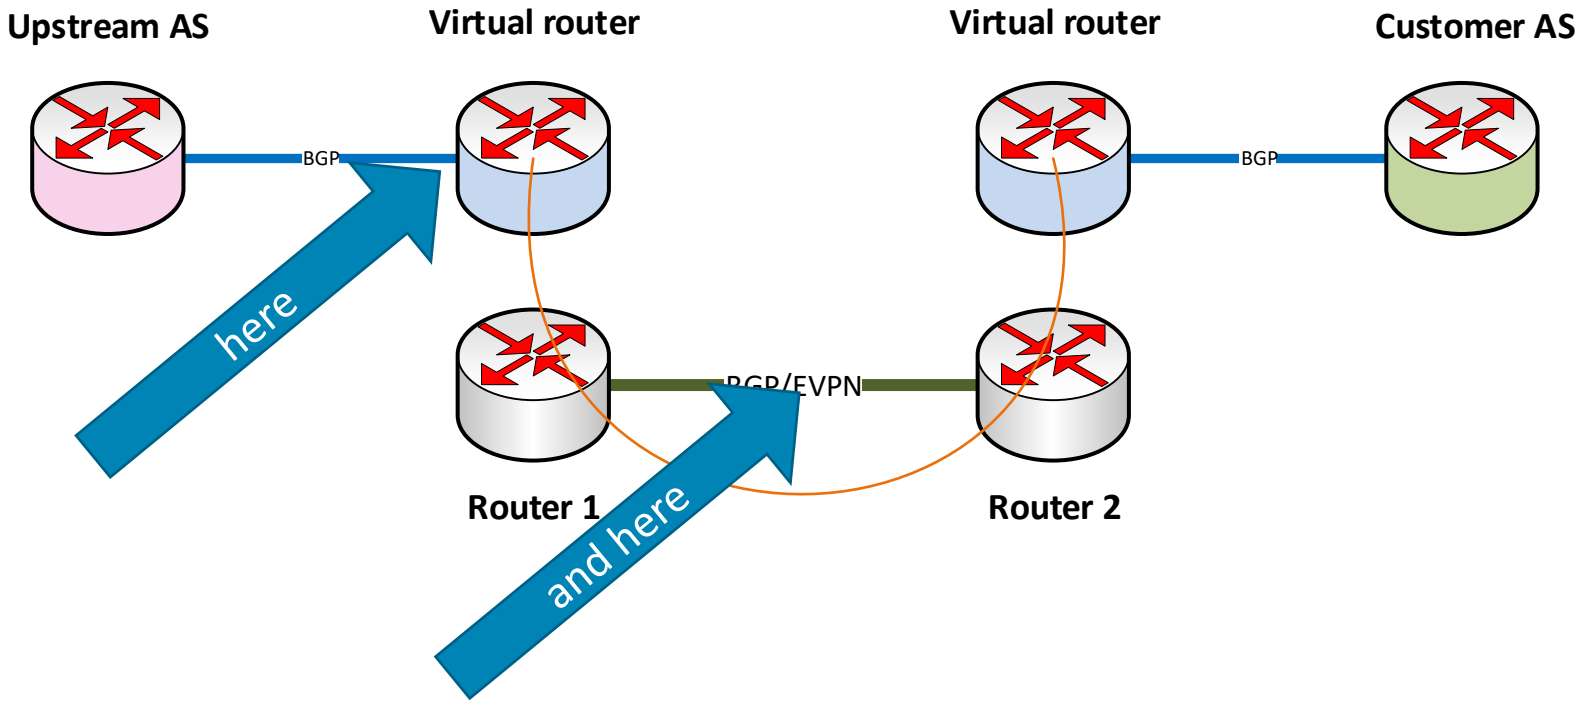

Export filtering

```
entry 10
        from
            protocol bgp-vpn
            prefix-list "AS21473_export-v6"
            community "as21473-250-customer"
        exit
        action accept
        exit
    exit

entry 20
        from
            protocol bgp
            prefix-list "AS21473_export-v6"
            community "as21473-250-customer"
        exit
        action accept
        exit
    exit
```

Customer connectivity

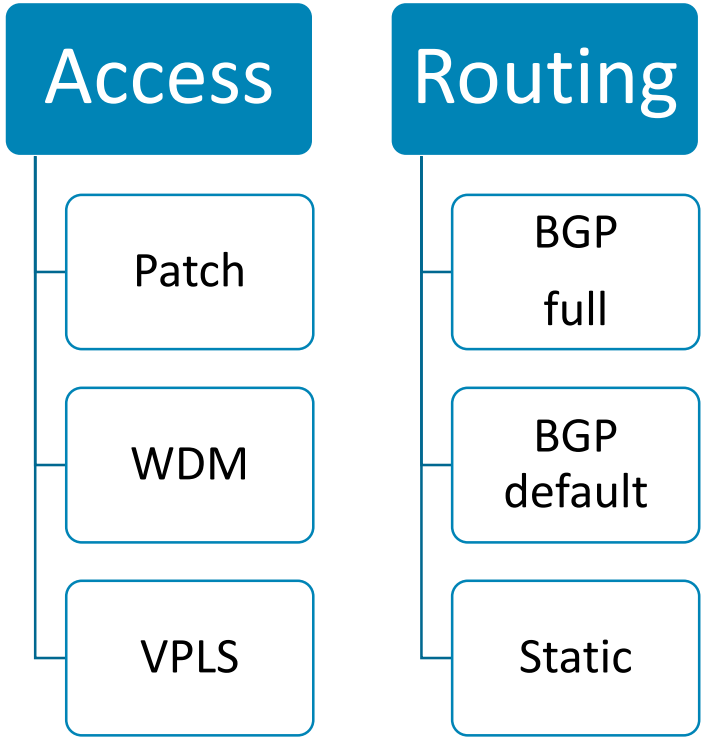

The big picture

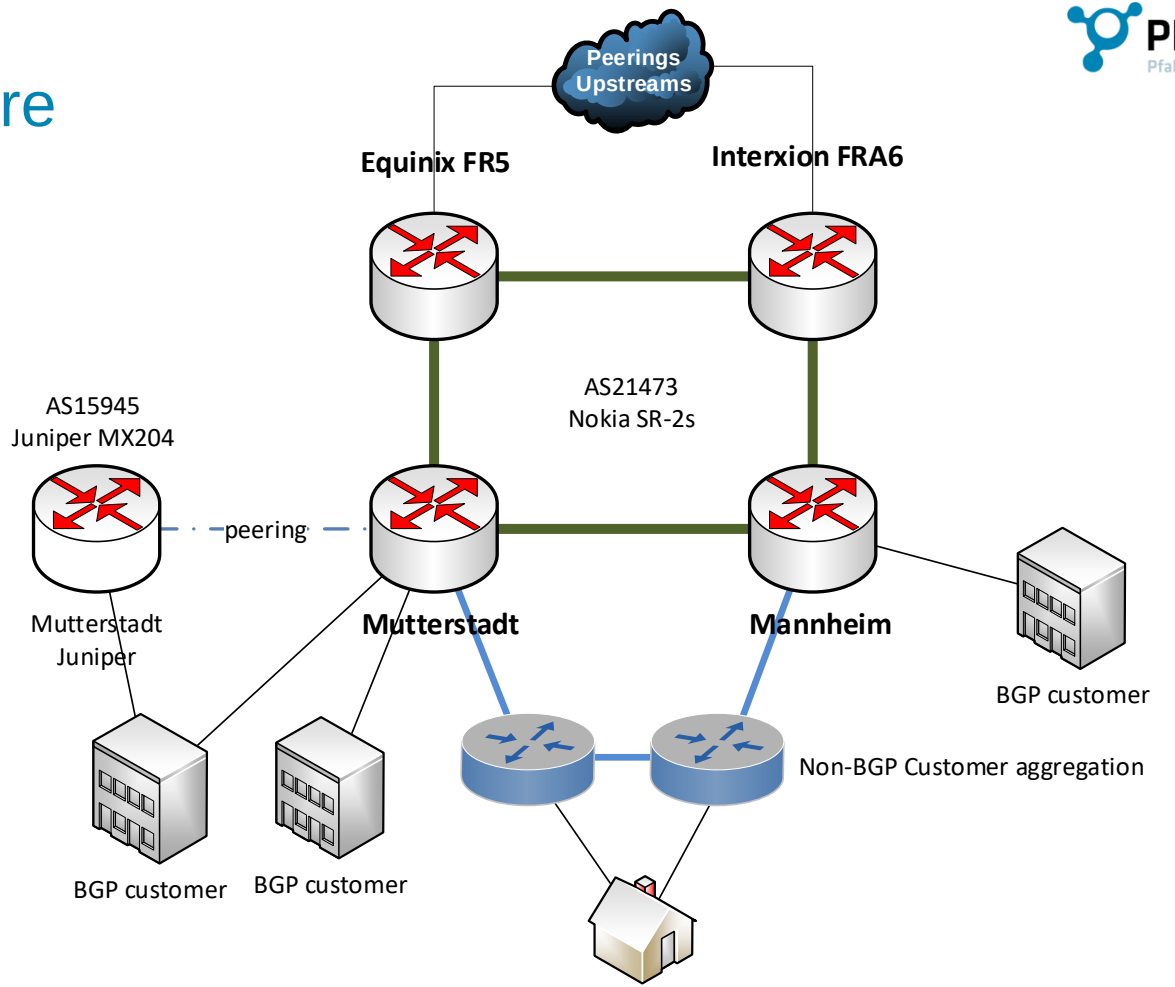

Current status

Network in production

1. AS21473 in VPRN
2. NGN-IC for SIP interconnect to other provider in another VPRN
3. BGP sessions for customers (AWS,…) in other VPRNs

Thank you