Międzyoperatorska wymiana ruchu IP z wykorzystaniem protokołu BGP

Aspekty organizacyjne i praktyczne przykłady.

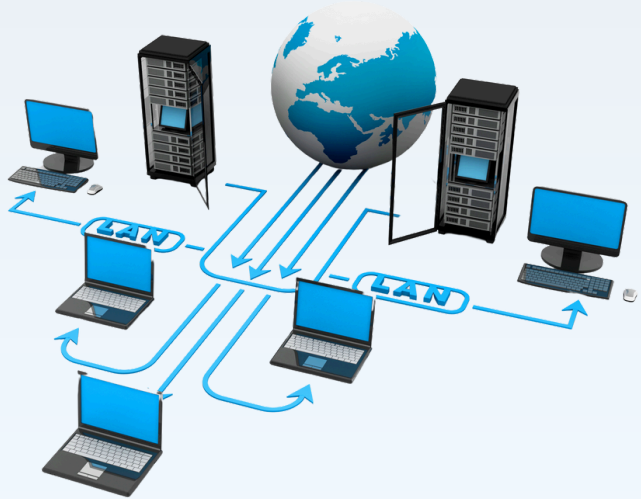

WACHOWIAK & SYN

Oferujemy rozwiązania firm:

appear

ASTRO GOING FUTURE TODAY.

ATTB Antennentechnik Bad Blankenburg

axing Competence in Communication Technologies

Barnfind® Technologies

PCTEL

RAFI

MEDIA LINKS® Media Defined Networking®

VideoFlow

Allied Telesis

HUAWEI

WACHOWIAK & SYN

Usługi wsparcia technicznego:

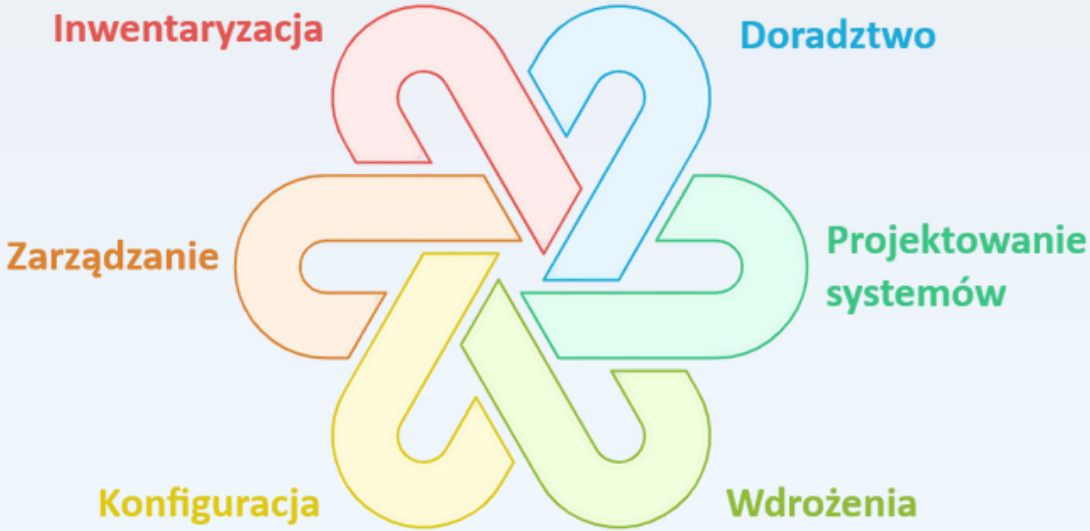

Winogradzka Telewizja Kablowa

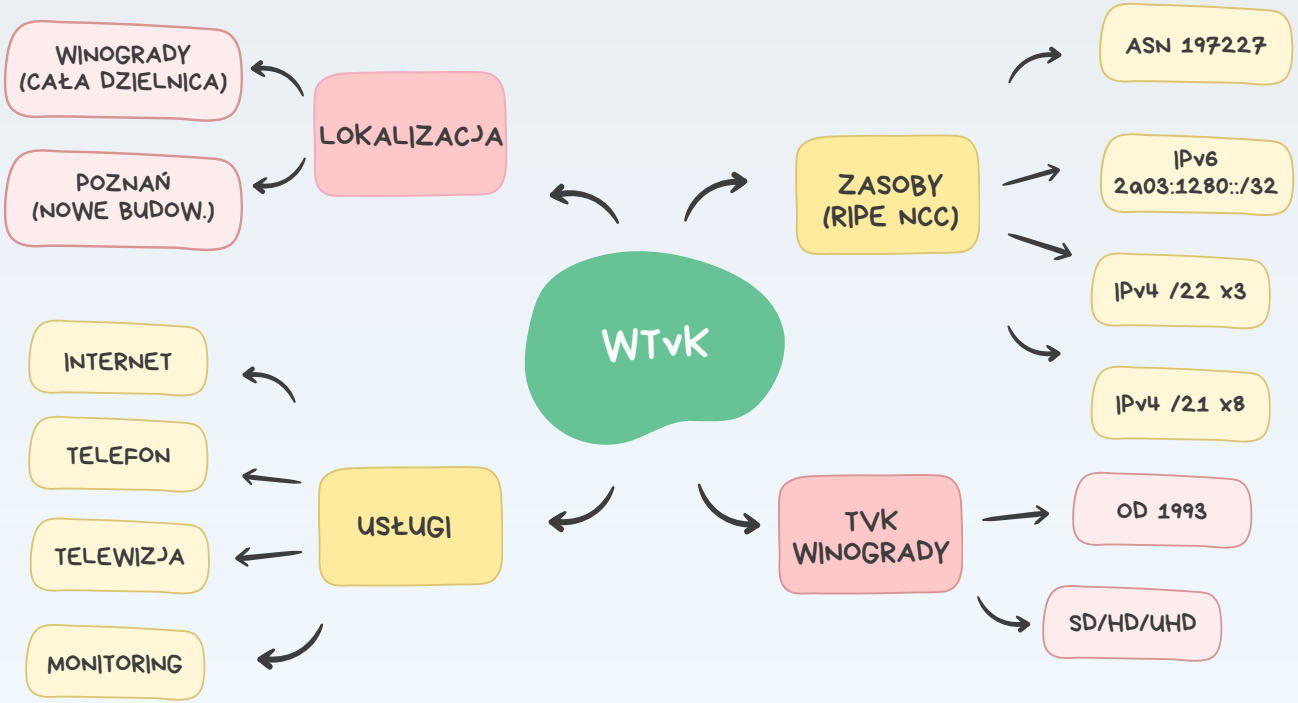

IANA (Internet Assigned Numbers Authority)

↳ RIR (Regional Internet Registry)

RIPE NCC (**RIPE** Network Coordination Centre)

≠ **RIPE** (Réseaux IP Européens)

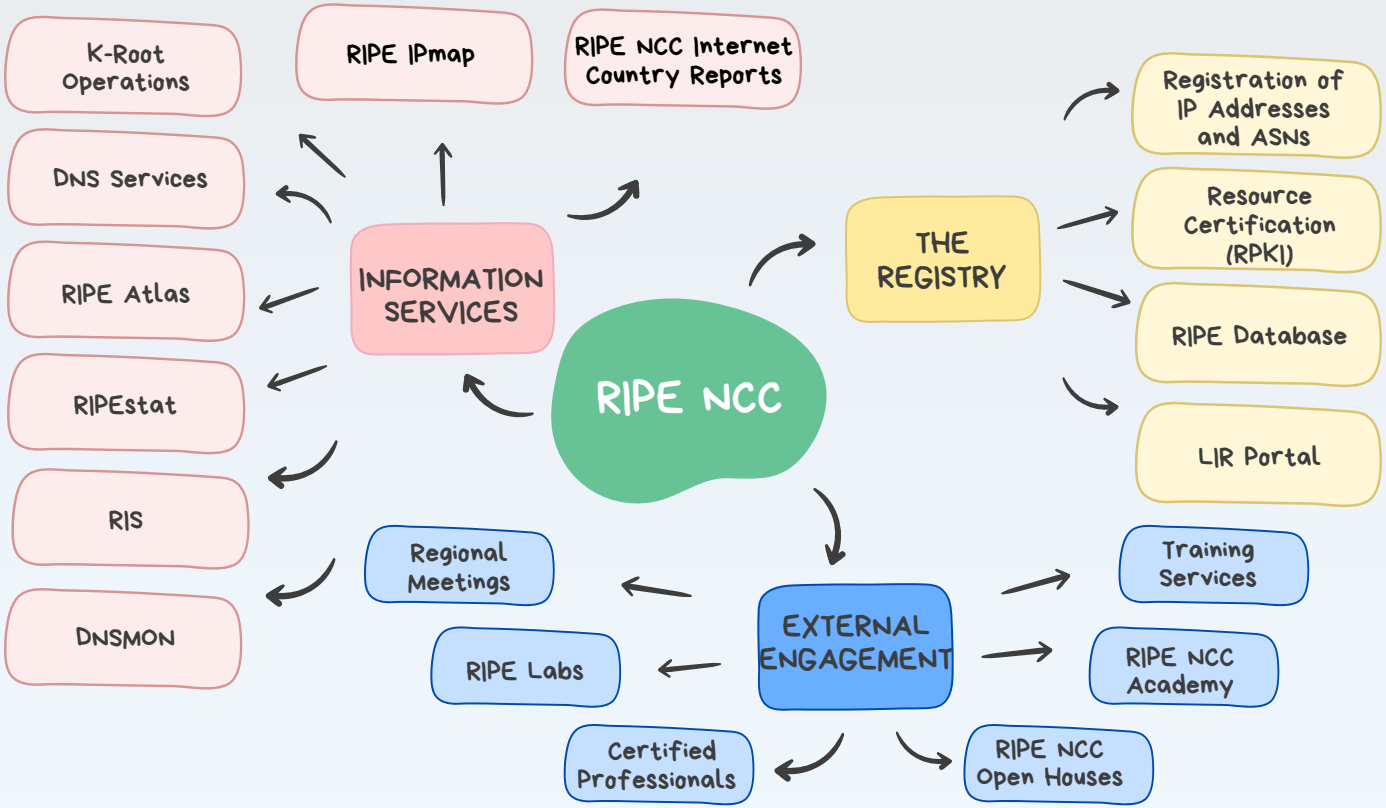
RIPE NCC
INFORMATION SERVICES
K-Root Operations
DNS Services
RIPE Atlas
RIPEstat
RIS
DNSMON
RIPE IPmap
RIPE NCC Internet Country Reports
THE REGISTRY
Registration of IP Addresses and ASNs
Resource Certification (RPKI)
RIPE Database
LIR Portal
EXTERNAL ENGAGEMENT
Regional Meetings
RIPE Labs
Certified Professionals
Training Services
RIPE NCC Academy
RIPE NCC Open Houses


LETS
GO!

Początek:
- jeden router
- jeden ISP
- zasoby od ISP

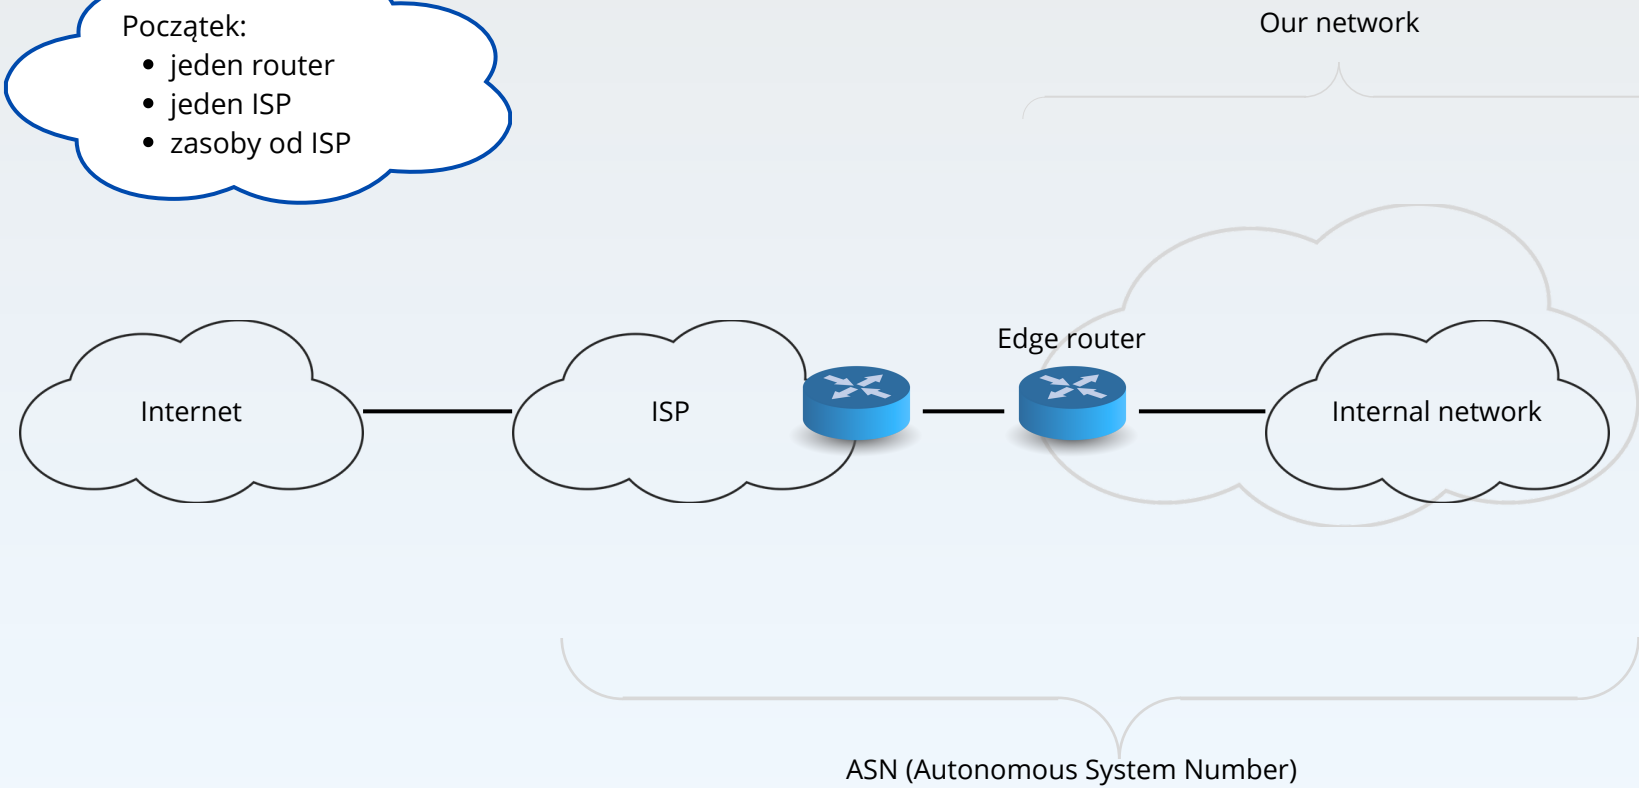

Redundantne łącze?
- własne adresy IP
- ASN

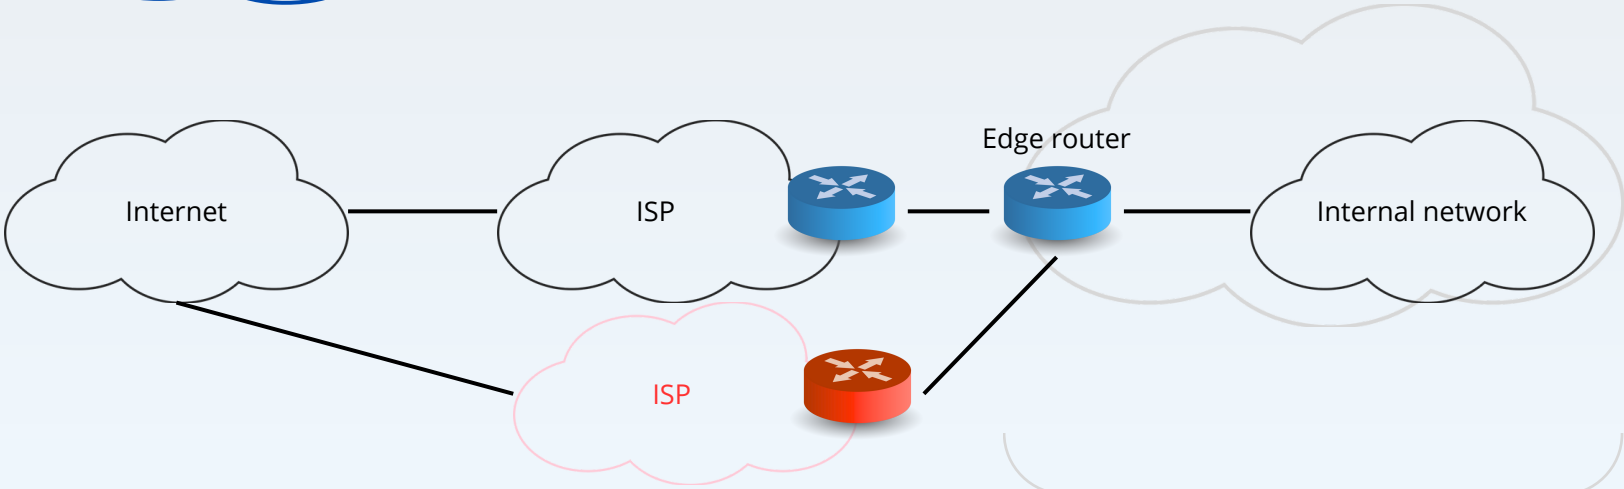

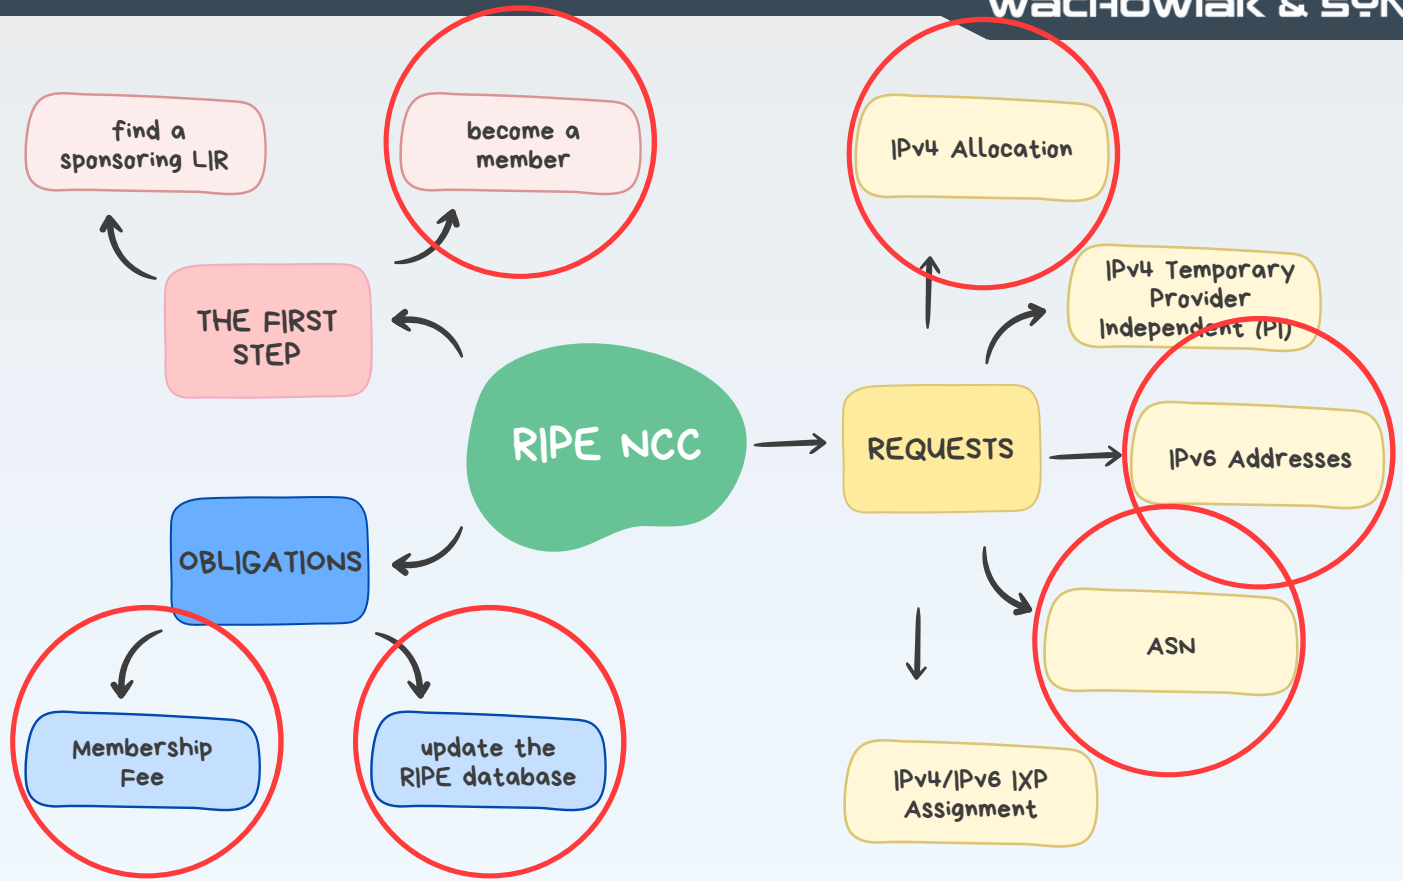
find a sponsoring LIR
become a member
THE FIRST STEP
RIPE NCC
OBLIGATIONS
Membership Fee
update the RIPE database
REQUESTS
IPv4 Allocation
IPv4 Temporary Provider Independent (PI)
IPv6 Addresses
ASN
IPv4/IPv6 IXP Assignment

become a member

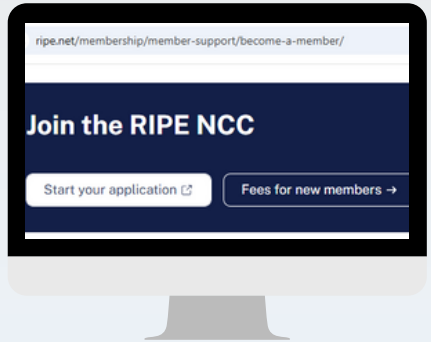

requests

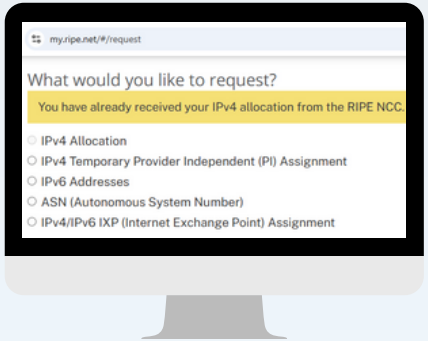

update the RIPE database

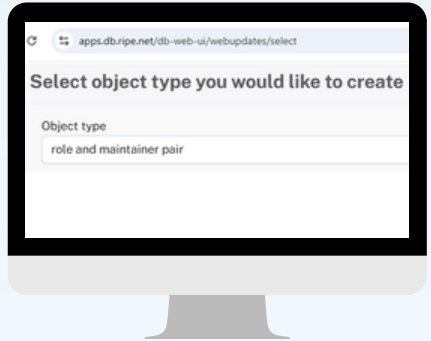

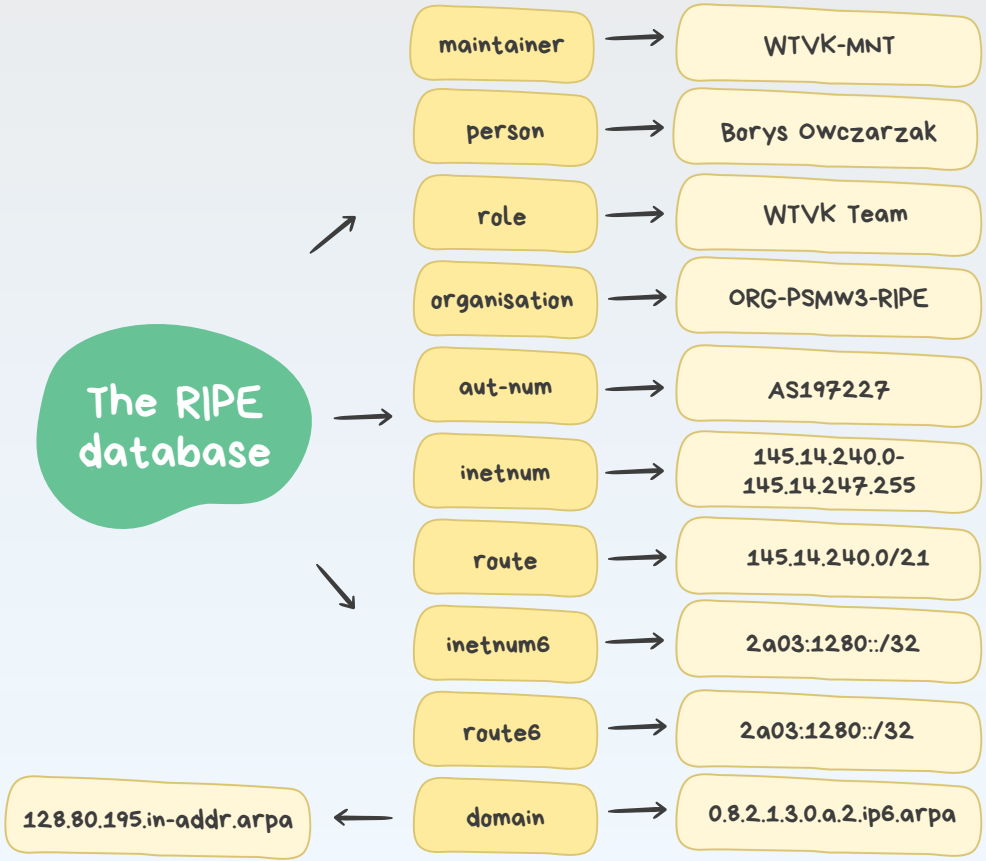
The RIPE database
maintainer
WTVK-MNT
person
Borys Owczarzak
role
WTVK Team
organisation
ORG-PSMW3-RIPE
aut-num
AS197227
inetnum
145.14.240.0-145.14.247.255
route
145.14.240.0/21
inetnum6
2a03:1280::/32
route6
2a03:1280::/32
domain
0.8.2.1.3.0.a.2.ip6.arpa
128.80.195.in-addr.arpa

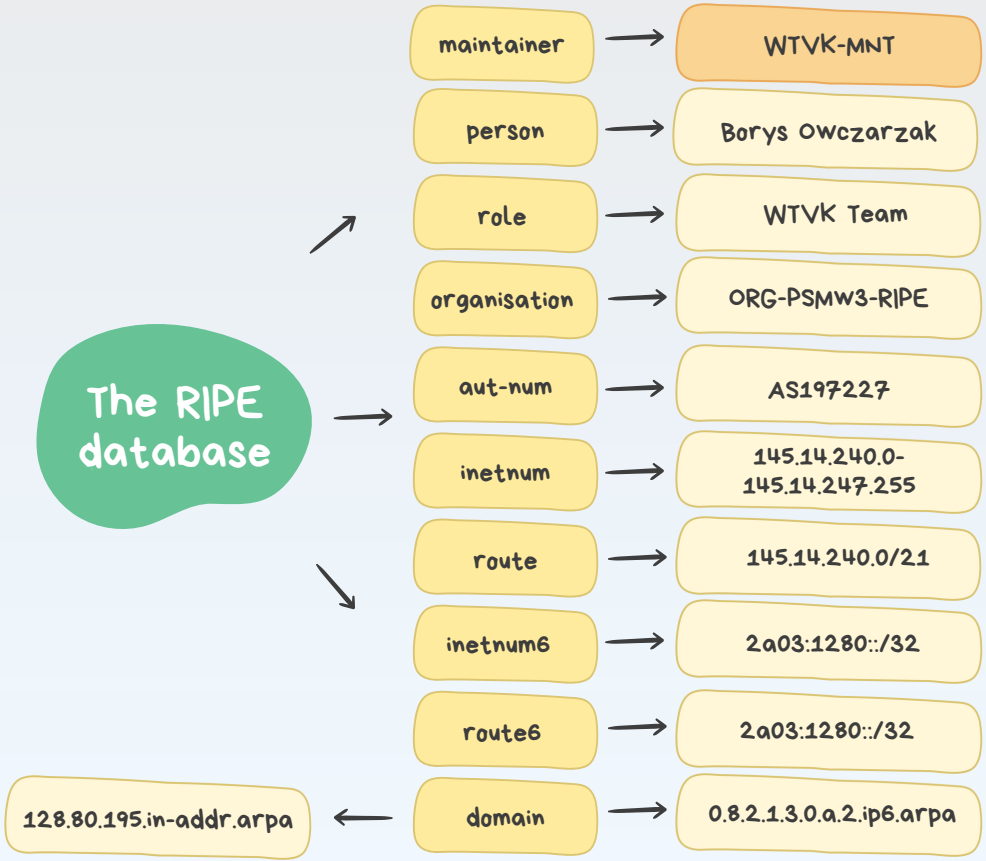

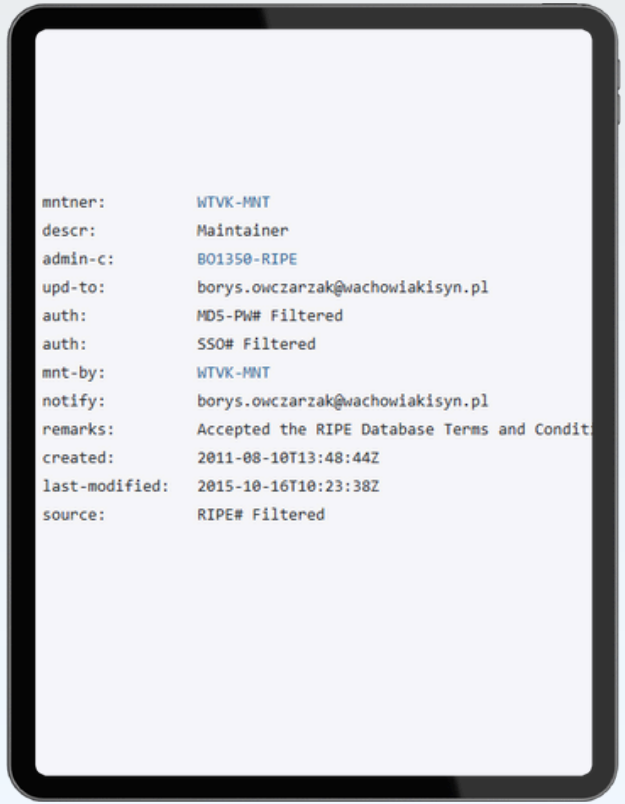

```
mntner:          WTVK-MNT
descr:           Maintainer
admin-c:         BO1350-RIPE
upd-to:          borys.owczarzak@wachowiakisyn.pl
auth:            MD5-PW# Filtered
auth:            SSO# Filtered
mnt-by:          WTVK-MNT
notify:          borys.owczarzak@wachowiakisyn.pl
remarks:         Accepted the RIPE Database Terms and Condit
created:         2011-08-10T13:48:44Z
last-modified:   2015-10-16T10:23:38Z
source:          RIPE# Filtered
```

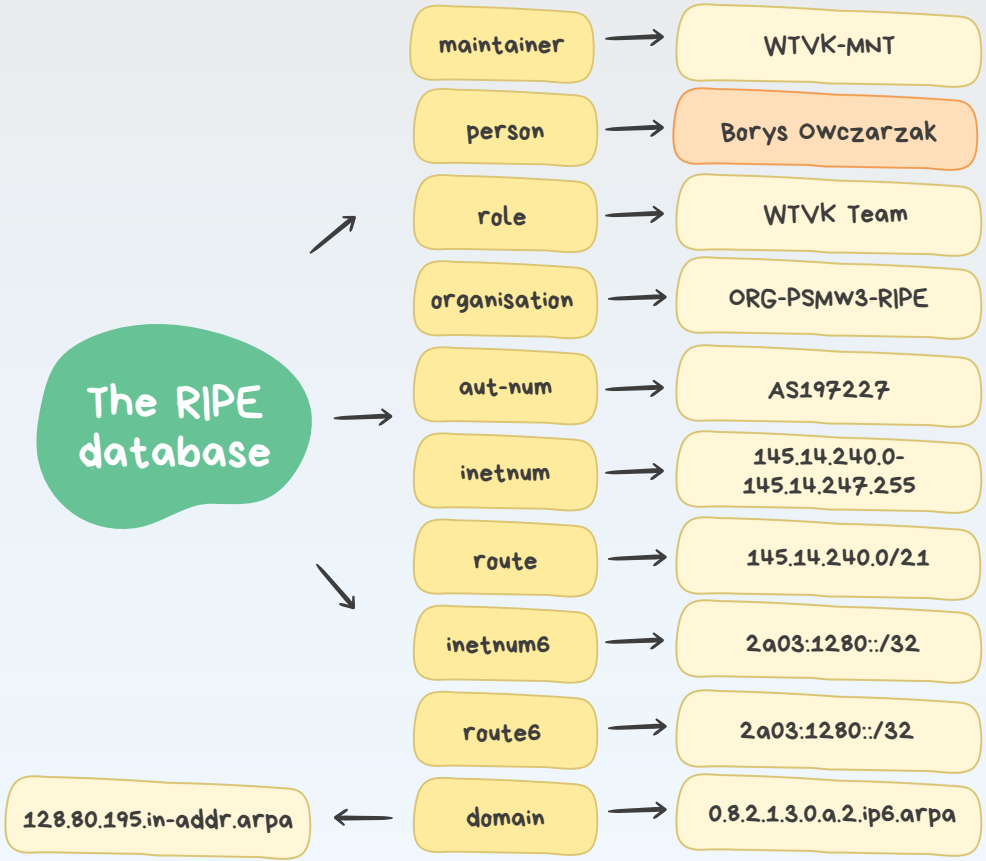

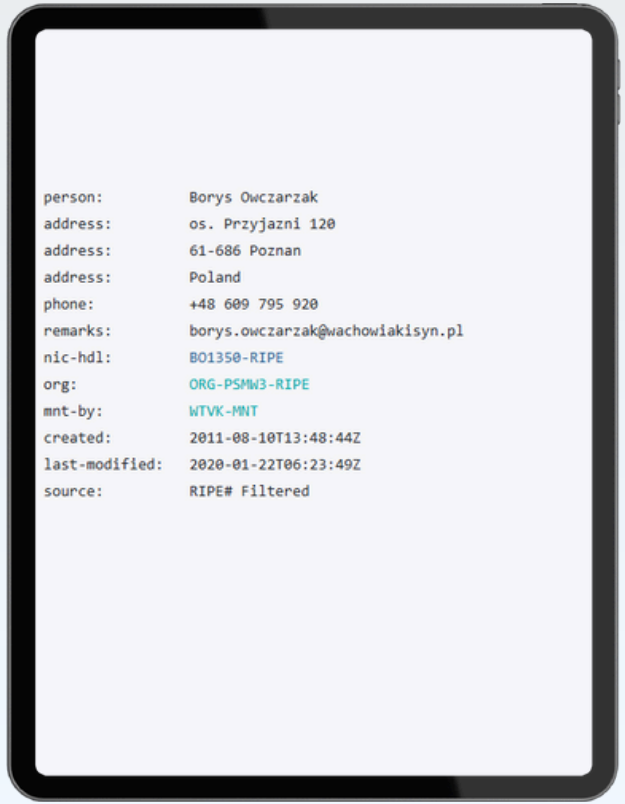

```
person:          Borys Owczarzak
address:         os. Przyjazni 120
address:         61-686 Poznan
address:         Poland
phone:           +48 609 795 920
remarks:         borys.owczarzak@wachowiakisyn.pl
nic-hdl:         BO1350-RIPE
org:             ORG-PSMW3-RIPE
mnt-by:          WTVK-MNT
created:         2011-08-10T13:48:44Z
last-modified:   2020-01-22T06:23:49Z
source:          RIPE# Filtered
```

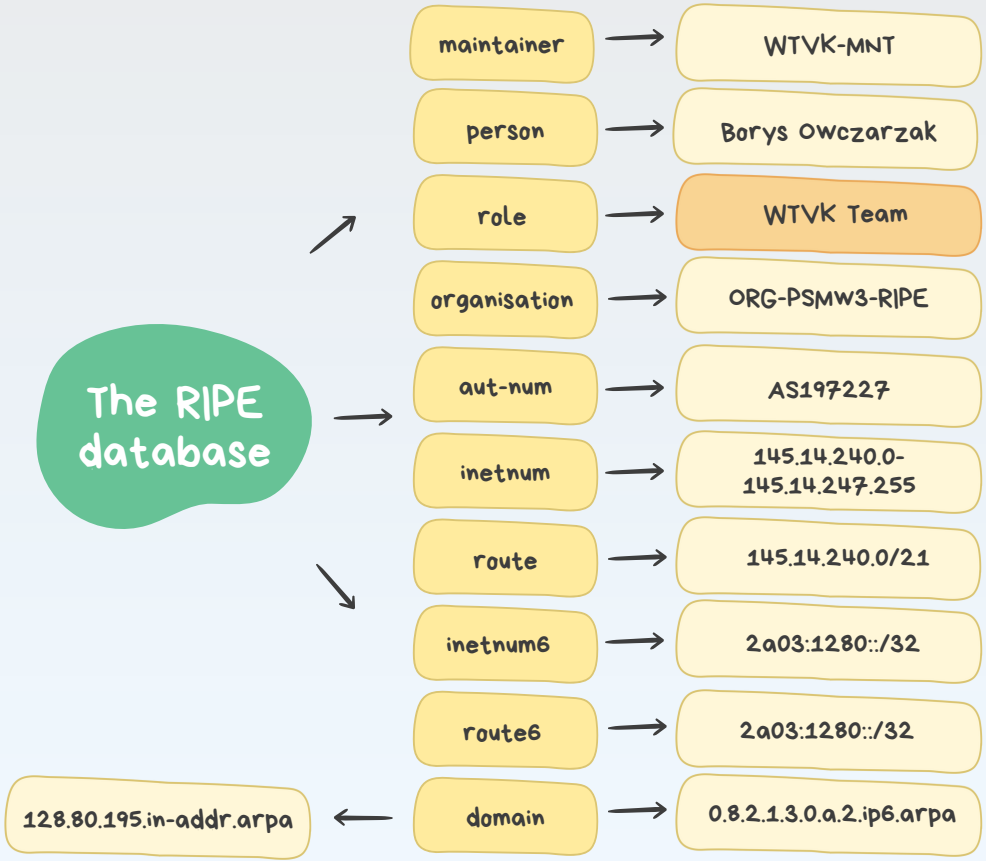

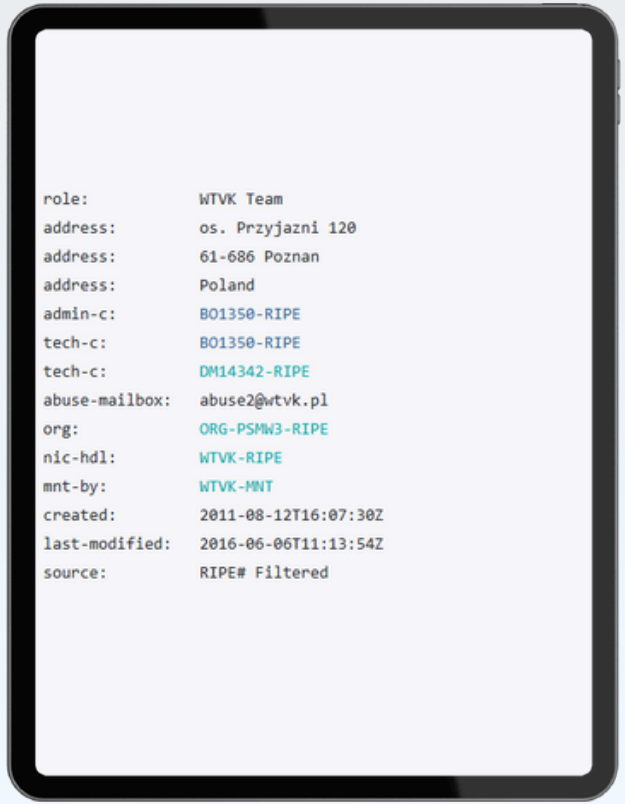

```
role:           WTVK Team
address:        os. Przyjazni 120
address:        61-686 Poznan
address:        Poland
admin-c:        BO1350-RIPE
tech-c:         BO1350-RIPE
tech-c:         DM14342-RIPE
abuse-mailbox:  abuse2@wtvk.pl
org:            ORG-PSMW3-RIPE
nic-hdl:        WTVK-RIPE
mnt-by:         WTVK-MNT
created:        2011-08-12T16:07:30Z
last-modified:  2016-06-06T11:13:54Z
source:         RIPE# Filtered
```

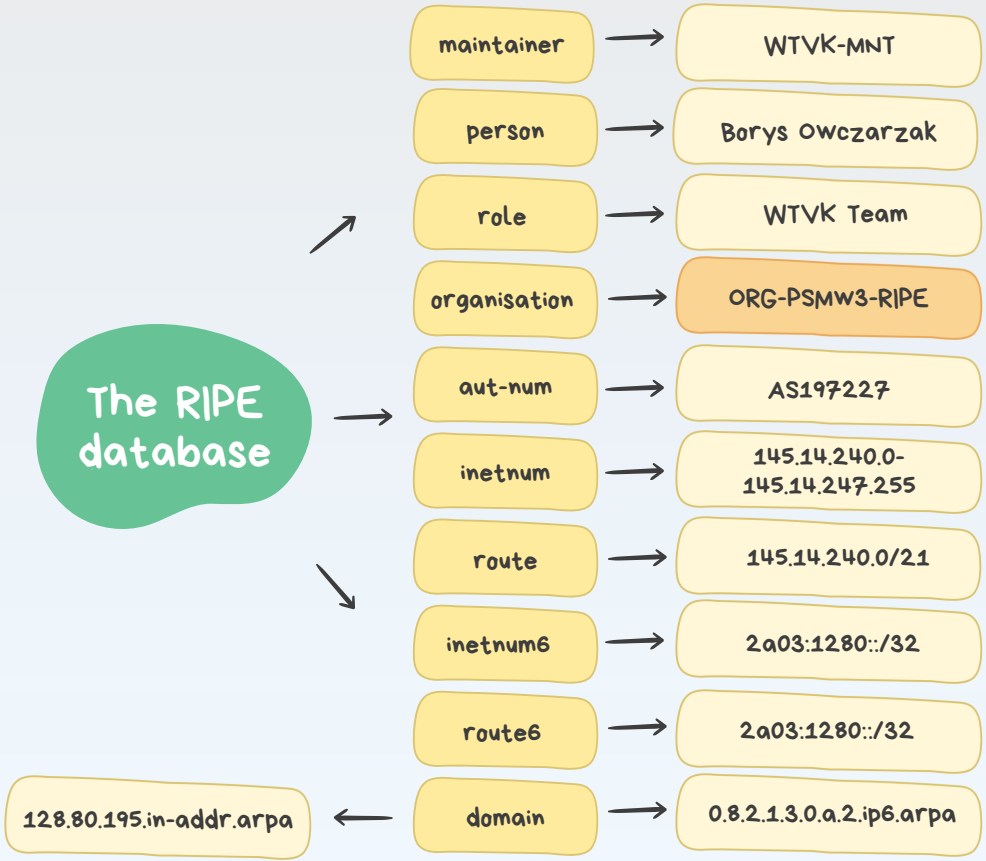


```
organisation:    ORG-PSMw3-RIPE
org-name:        Poznanska Spoldzielnia Mieszkaniowa "Winogrady" w
country:         PL
org-type:        LIR
address:         os. Przyjazni 125B
address:         61-686
address:         Poznan
address:         POLAND
phone:           +48616303273
fax-no:          +48616303270
e-mail:          pawel.suchorski@wtvk.pl
admin-c:         BO1350-RIPE
tech-c:          WTVK-RIPE
abuse-c:         WTVK-RIPE
mnt-ref:         RIPE-NCC-HM-MNT
mnt-ref:         WTVK-MNT
mnt-by:          RIPE-NCC-HM-MNT
mnt-by:          WTVK-MNT
created:         2011-08-05T13:23:32Z
last-modified:   2020-12-16T12:44:05Z
source:          RIPE
```

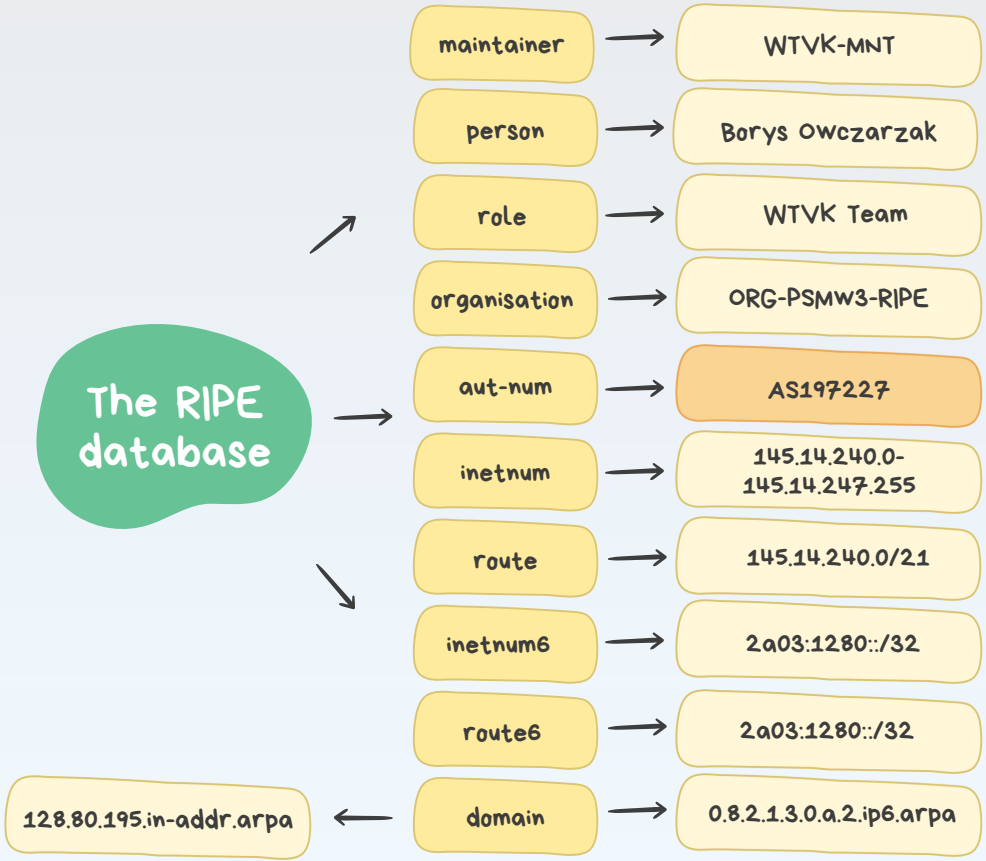


```
aut-num:        AS197227
as-name:        PSM_Winogrady
org:            ORG-PSMW3-RIPE
remarks:        ******************
remarks:        IPv4 tranzit
remarks:        ******************
import:         from AS12741 action pref=100; accept ANY
import:         from AS20552 action pref=100; accept ANY
import:         from AS6939 action pref=100; accept ANY
import:         from AS20804 action pref=100; accept ANY
export:         to AS12741 announce AS-WTVK
export:         to AS20552 announce AS-WTVK
export:         to AS6939 announce AS-WTVK
export:         to AS20804 announce AS197227
remarks:        ******************
remarks:        IPv4 IX
remarks:        ******************
import:         from AS24115 action pref=100; accept ANY
import:         from AS44896 action pref=100; accept AS-PIX
export:         to AS24115 announce AS197227
export:         to AS44896 announce AS197227
remarks:        ******************
remarks:        IPv6 tranzit
remarks:        ******************
mp-import:      afi ipv6.unicast from AS6939 action pref=100; accept ANY
mp-import:      afi ipv6.unicast from AS20552 action pref=100; accept ANY
mp-import:      afi ipv6.unicast from AS12741 action pref=100; accept ANY
mp-import:      afi ipv6.unicast from AS20804 action pref=100; accept ANY
mp-export:      afi ipv6.unicast to AS6939 announce AS197227
mp-export:      afi ipv6.unicast to AS20552 announce AS-WTVK
mp-export:      afi ipv6.unicast to AS12741 announce AS-WTVK
mp-export:      afi ipv6.unicast to AS20804 announce AS-WTVK
```

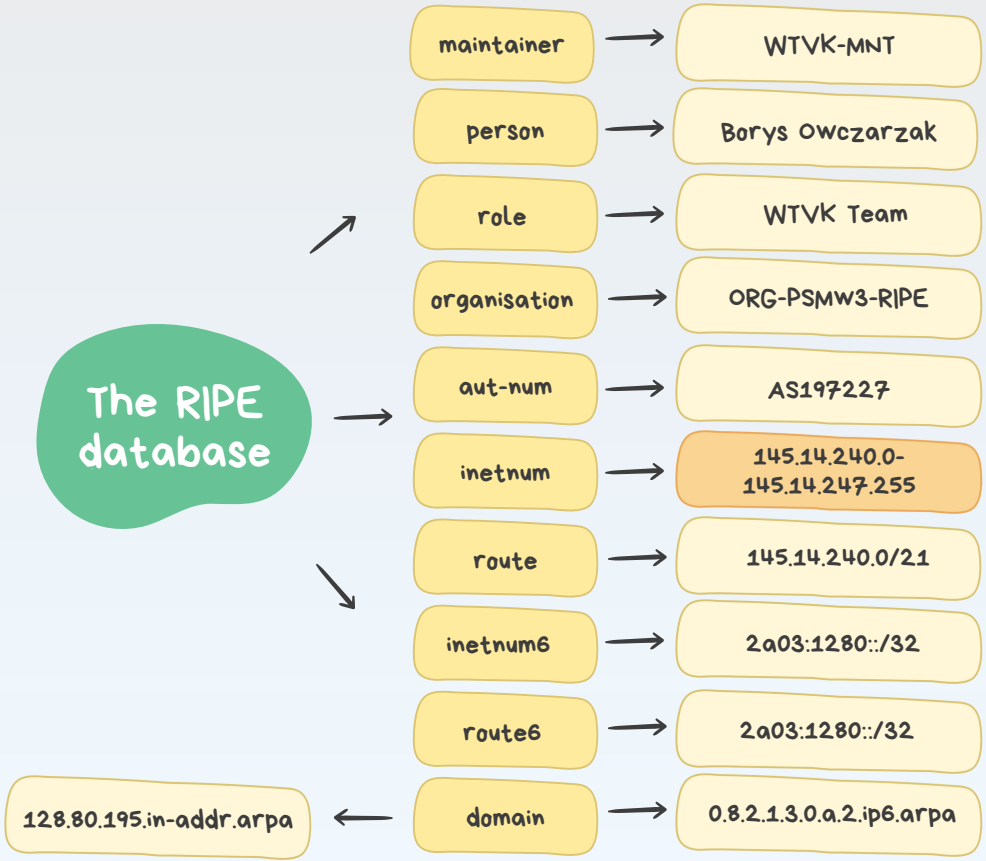


```
inetnum:        145.14.240.0 - 145.14.247.255
netname:        WINOGRADY-POZNAN
descr:          PSM "Winogrady" w Poznaniu
descr:          Os. Przyjazni 125B
descr:          Poznan
geoloc:         52.43148 16.92503
country:        PL
org:            ORG-PSMw3-RIPE
admin-c:        BO1350-RIPE
tech-c:         WTVK-RIPE
status:         LEGACY
mnt-by:         RIPE-NCC-LEGACY-MNT
mnt-routes:     WTVK-MNT
mnt-domains:    WTVK-MNT
mnt-by:         WTVK-MNT
created:        2017-03-15T11:09:58Z
last-modified:  2017-05-12T12:57:31Z
source:         RIPE
```

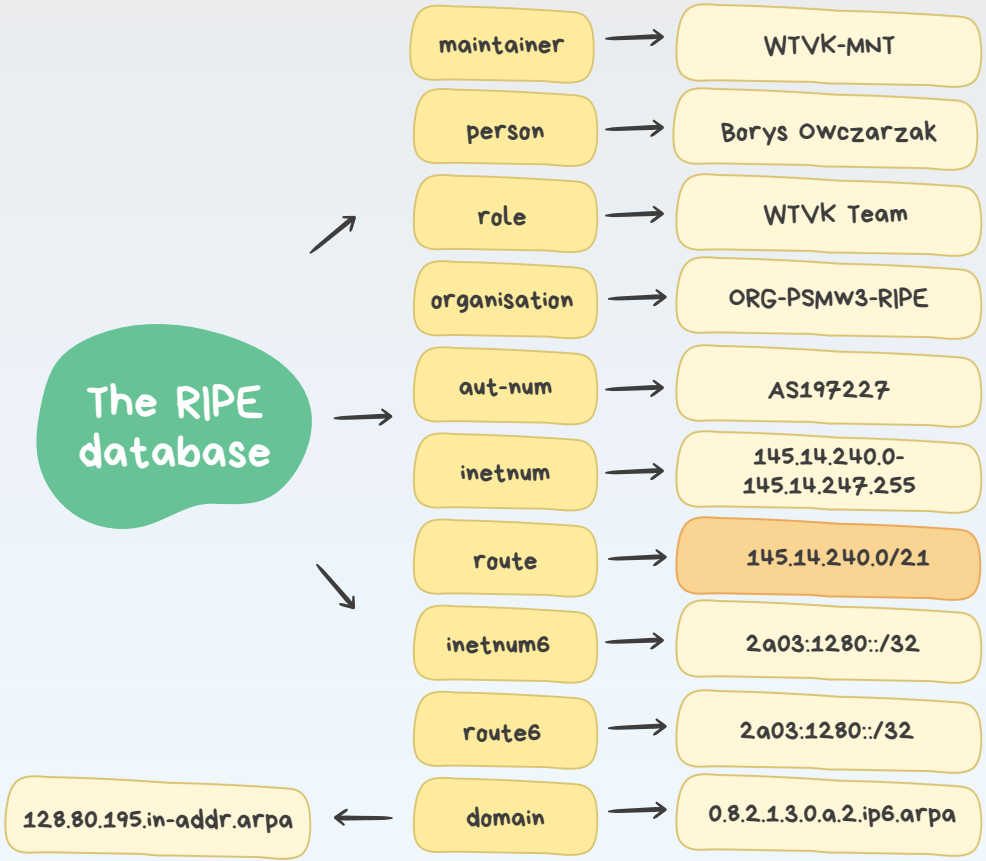
The RIPE database
maintainer
WTVK-MNT
person
Borys Owczarzak
role
WTVK Team
organisation
ORG-PSMW3-RIPE
aut-num
AS197227
inetnum
145.14.240.0-145.14.247.255
route
145.14.240.0/21
inetnum6
2a03:1280::/32
route6
2a03:1280::/32
domain
0.8.2.1.3.0.a.2.ip6.arpa
128.80.195.in-addr.arpa

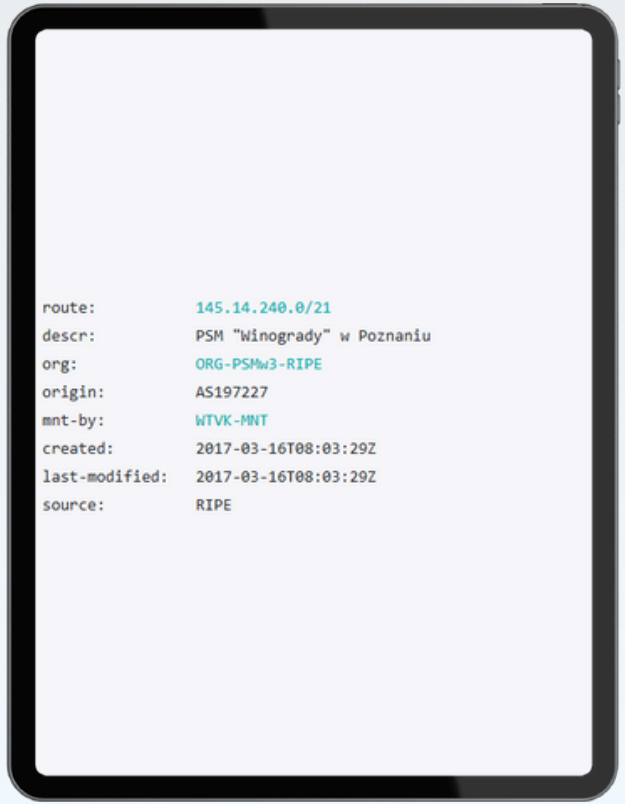
route: 145.14.240.0/21
descr: PSM "Winogrady" w Poznaniu
org: ORG-PSMW3-RIPE
origin: AS197227
mnt-by: WTVK-MNT
created: 2017-03-16T08:03:29Z
last-modified: 2017-03-16T08:03:29Z
source: RIPE

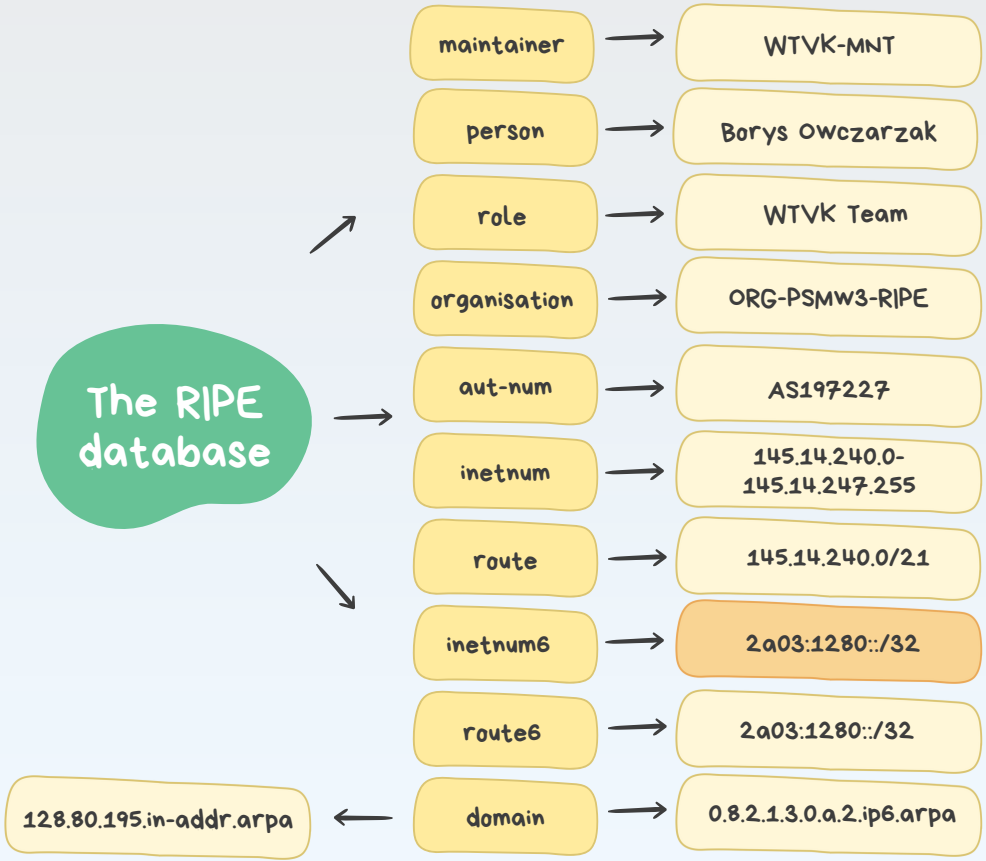


```
inet6num:       2a03:1280::/32
netname:        PL-WTVK-20110811
descr:          PSM "Winogrady" w Poznaniu
geoloc:         52.43148 16.92503
country:        PL
org:            ORG-PSMW3-RIPE
admin-c:        BO1350-RIPE
tech-c:         WTVK-RIPE
status:         ALLOCATED-BY-RIR
mnt-by:         RIPE-NCC-HM-MNT
mnt-by:         WTVK-MNT
mnt-lower:      WTVK-MNT
mnt-routes:     WTVK-MNT
mnt-domains:    WTVK-MNT
created:        2011-08-11T12:32:56Z
last-modified:  2017-04-04T09:36:57Z
source:         RIPE# Filtered
```

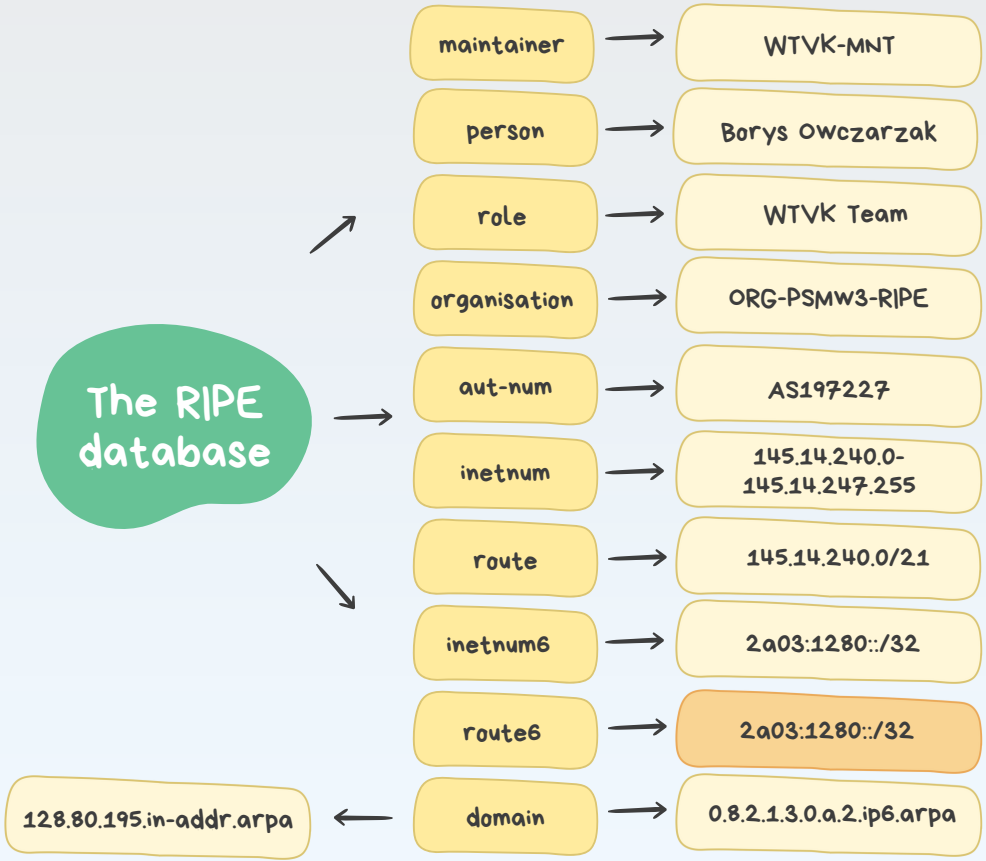
The RIPE database
maintainer
WTVK-MNT
person
Borys Owczarzak
role
WTVK Team
organisation
ORG-PSMW3-RIPE
aut-num
AS197227
inetnum
145.14.240.0-145.14.247.255
route
145.14.240.0/21
inetnum6
2a03:1280::/32
route6
2a03:1280::/32
domain
0.8.2.1.3.0.a.2.ip6.arpa
128.80.195.in-addr.arpa

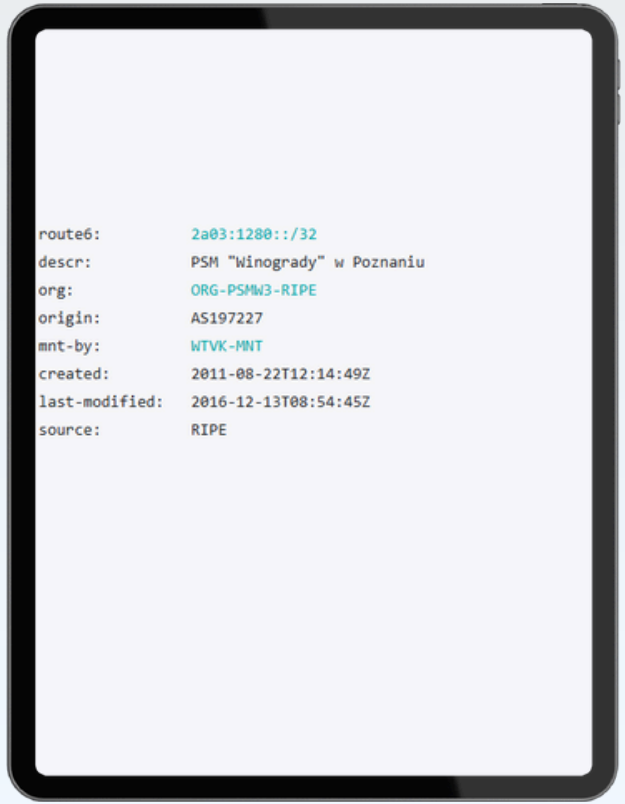
route6: 2a03:1280::/32
descr: PSM "Winogrady" w Poznaniu
org: ORG-PSMW3-RIPE
origin: AS197227
mnt-by: WTVK-MNT
created: 2011-08-22T12:14:49Z
last-modified: 2016-12-13T08:54:45Z
source: RIPE

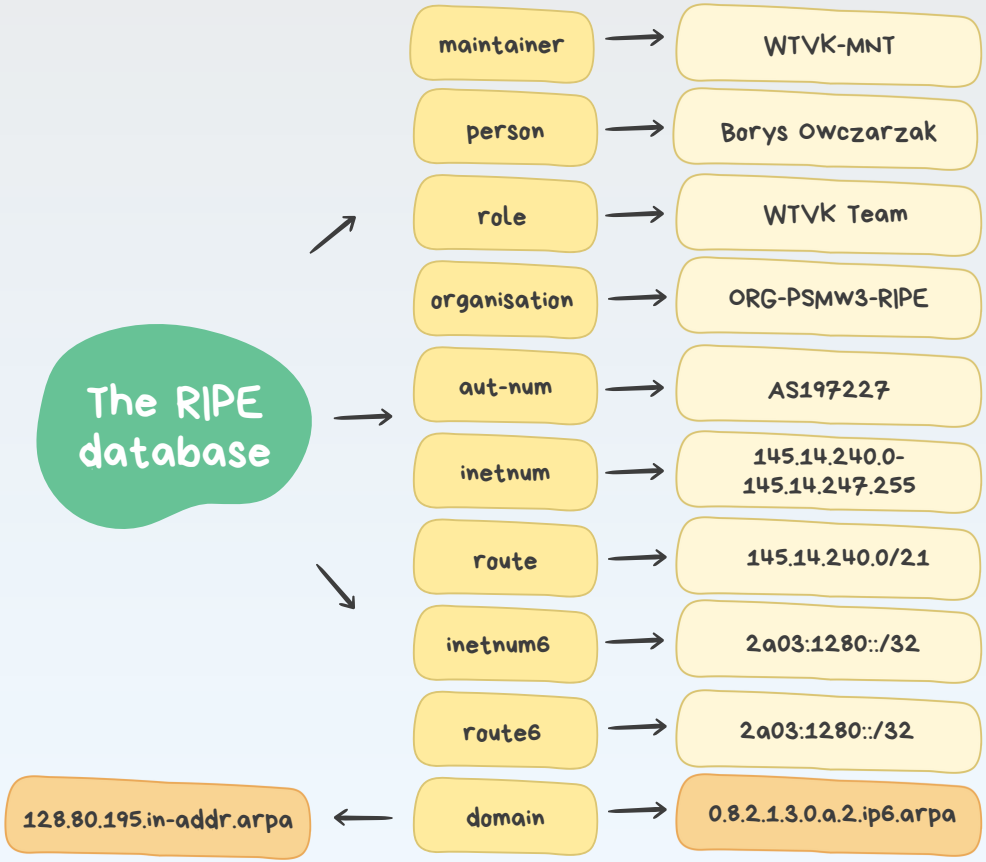


```
domain:          0.8.2.1.3.0.a.2.ip6.arpa
descr:           WINOGRADY-POZNAN
admin-c:         BO1350-RIPE
tech-c:          WTVK-RIPE
zone-c:          BO1350-RIPE
nserver:         ns1.wtvk.pl
nserver:         ns2.wtvk.pl
org:             ORG-PSMW3-RIPE
mnt-by:          WTVK-MNT
created:         2011-09-01T10:06:03Z
last-modified:   2016-03-06T12:23:52Z
source:          RIPE

domain:          128.80.195.in-addr.arpa
descr:           WINOGRADY-POZNAN
org:             ORG-PSMW3-RIPE
tech-c:          WTVK-RIPE
admin-c:         BO1350-RIPE
zone-c:          BO1350-RIPE
nserver:         ns1.wtvk.pl
nserver:         ns2.wtvk.pl
mnt-by:          WTVK-MNT
created:         2011-08-18T07:10:04Z
last-modified:   2015-10-18T07:46:22Z
source:          RIPE
```

Sieć operatorska gotowa:
- ASN
- prefiksy IPx4
- routery BGP
- obiekty z bazie RIPE

Ustalenia z ISP
- POP / medium / pasmo
- połączeniówka IPv4/v6
- routery BGP
- cena ;-)

Przekazać operatorowi
- ASN
- prefiksy IPv6
- prefiksy IPv4

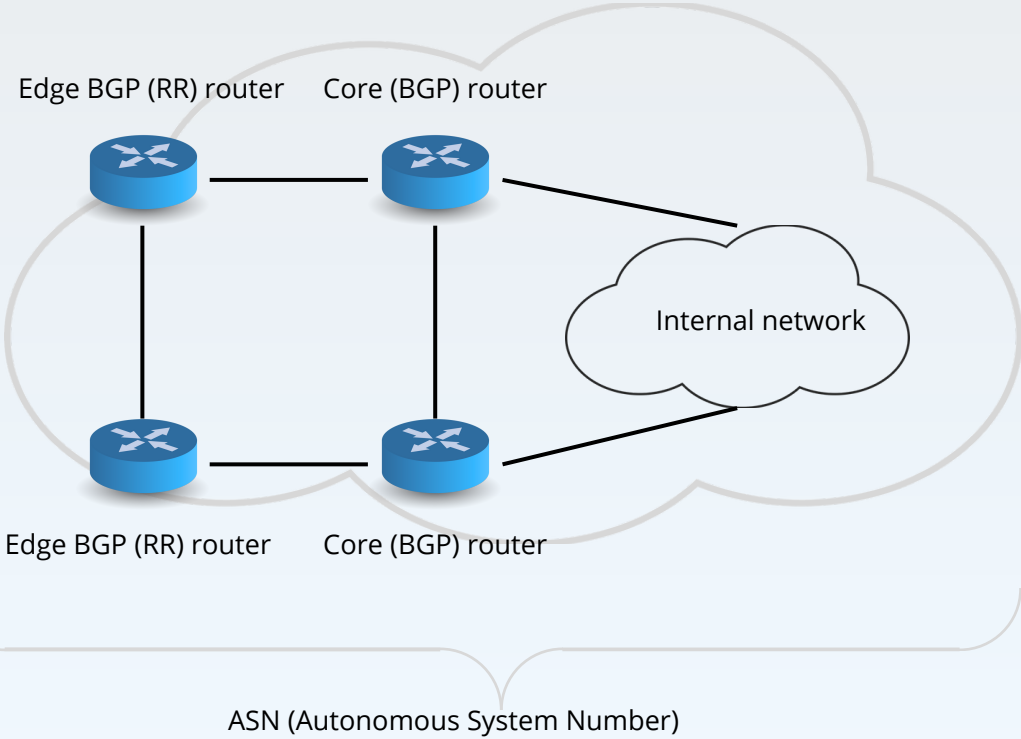

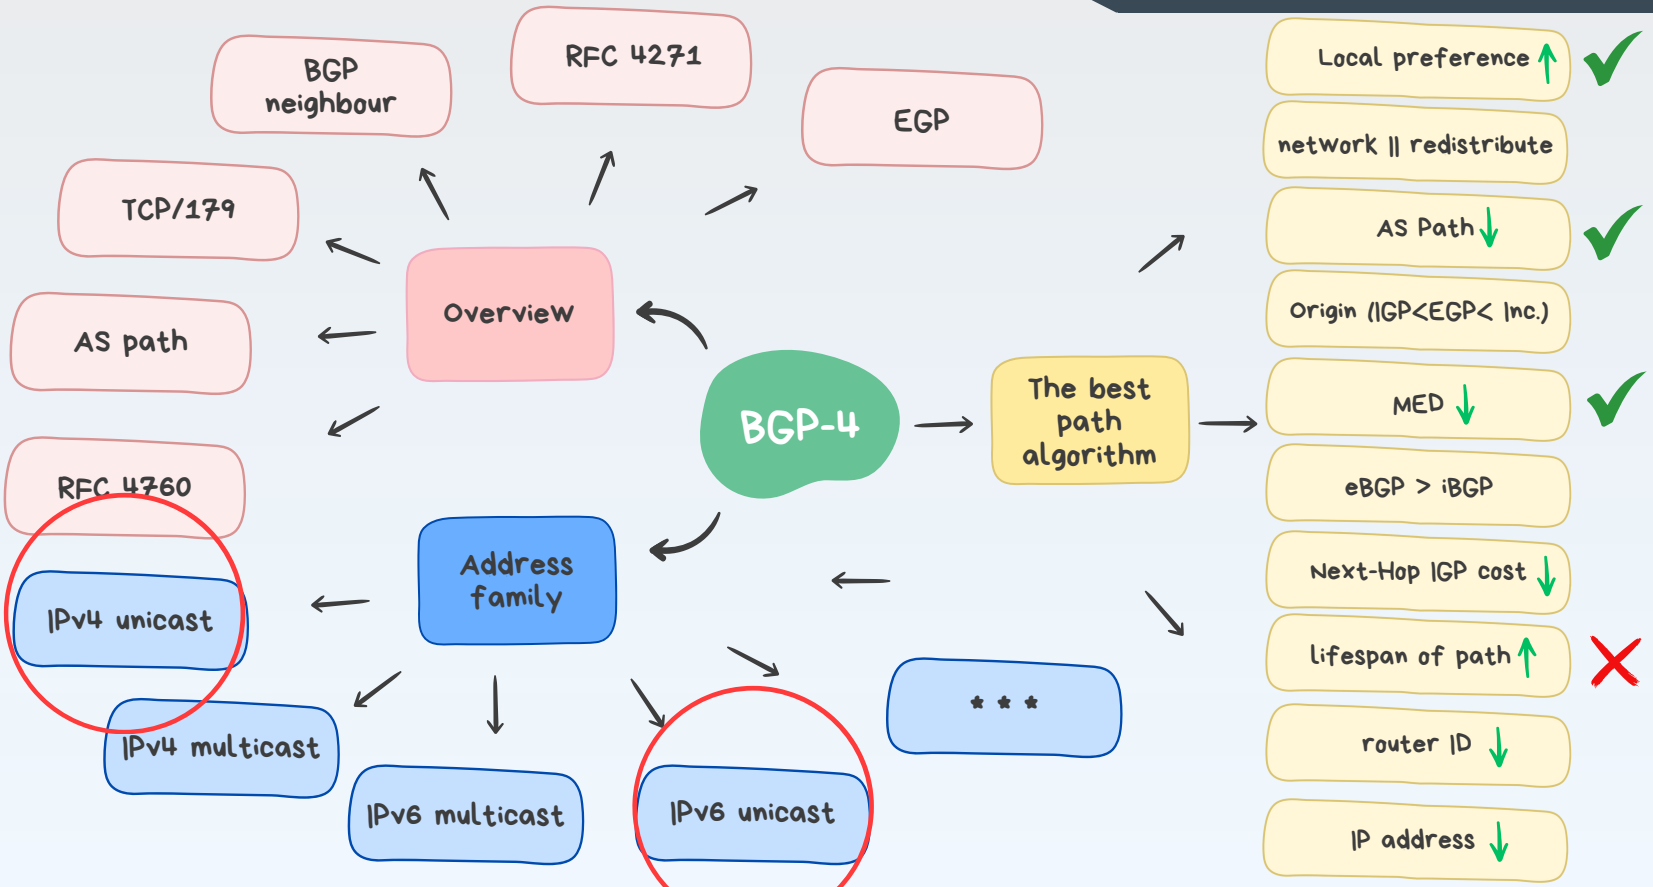
BGP neighbour
RFC 4271
EGP
TCP/179
Overview
AS path
RFC 4760
BGP-4
The best path algorithm
Address family
IPv4 unicast
IPv4 multicast
IPv6 multicast
IPv6 unicast
* * *
Local preference ↑
network || redistribute
AS Path ↓
Origin (IGP<EGP< Inc.)
MED ↓
eBGP > iBGP
Next-Hop IGP cost ↓
lifespan of path ↑
router ID ↓
IP address ↓

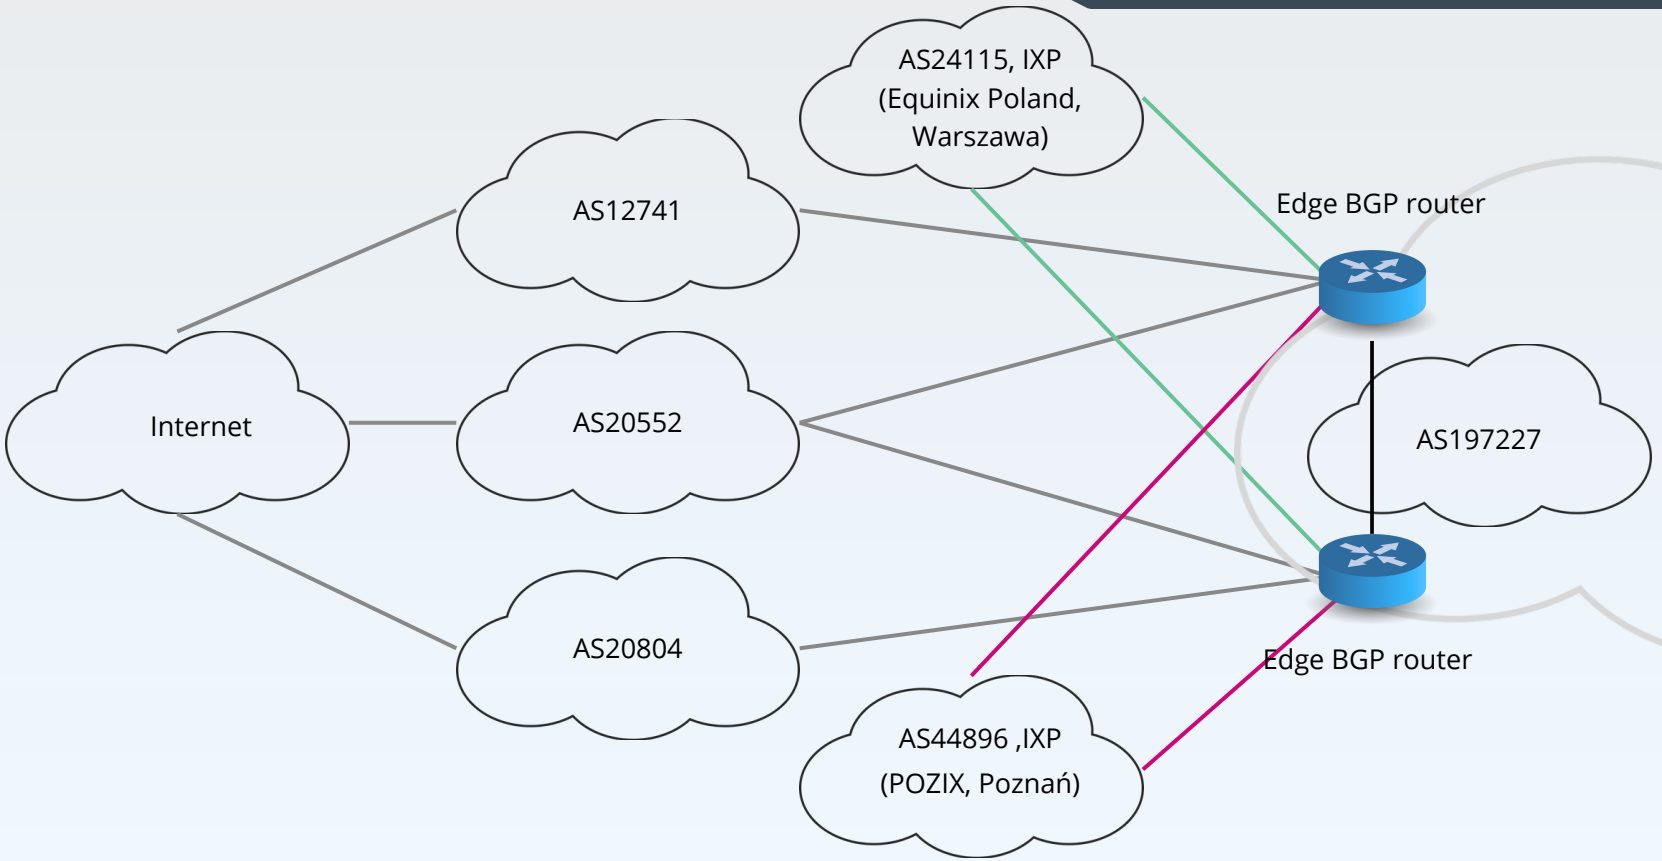
AS24115, IXP
(Equinix Poland,
Warszawa)
AS12741
Edge BGP router
Internet
AS20552
AS197227
AS20804
Edge BGP router
AS44896 ,IXP
(POZIX, Poznań)

```
router bgp 197227
 no bgp enforce-first-as
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 ...
 neighbor 10.xx.yy.1 remote-as 197227
 neighbor 10.xx.yy.1 update-source Loopback0
 neighbor 10.xx.yy.1 description ASRv2
 ...
 neighbor 81.210.ww.zz remote-as 12741
 neighbor 81.210.ww.zz description NETIA
 ...
 neighbor FDFD:BEEF:CAFE:1::1 remote-as 197227
 neighbor FDFD:BEEF:CAFE:1::1 description ASRv2
 neighbor FDFD:BEEF:CAFE:1::1 update-source Loopback0
 ...
 neighbor 2001:41B0:x:y::1 remote-as 12741
 neighbor 2001:41B0:x:y::1 description NETIA
 ...
 address-family ipv4
 ...
 exit-address-family
 address-family ipv6
 ...
 exit-address-family
exit
```

197227 - nasz numer AS
przyznany przez RIPE NCC

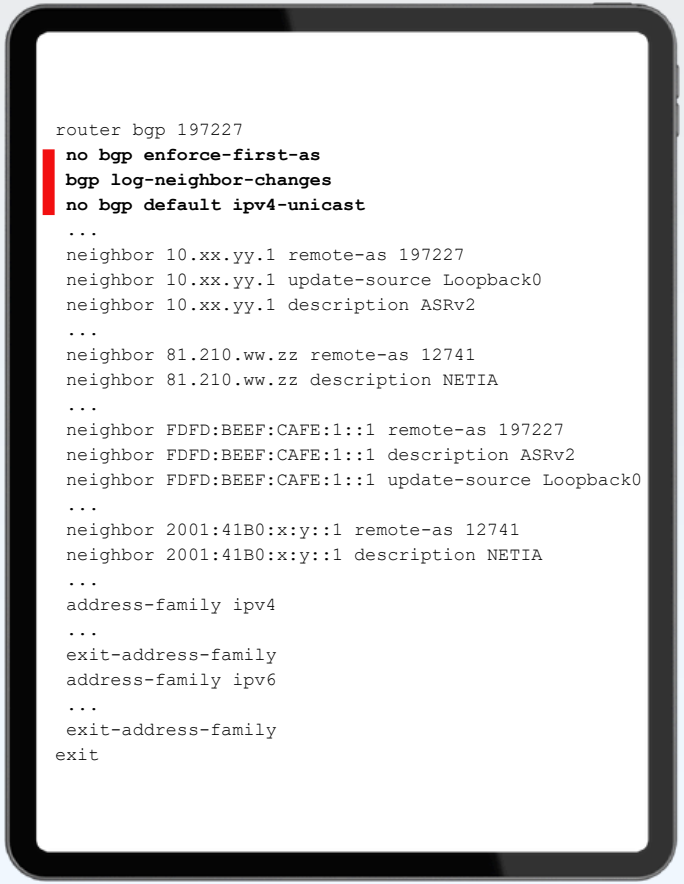

```
router bgp 197227
 no bgp enforce-first-as
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 ...
 neighbor 10.xx.yy.1 remote-as 197227
 neighbor 10.xx.yy.1 update-source Loopback0
 neighbor 10.xx.yy.1 description ASRv2
 ...
 neighbor 81.210.ww.zz remote-as 12741
 neighbor 81.210.ww.zz description NETIA
 ...
 neighbor FDFD:BEEF:CAFE:1::1 remote-as 197227
 neighbor FDFD:BEEF:CAFE:1::1 description ASRv2
 neighbor FDFD:BEEF:CAFE:1::1 update-source Loopback0
 ...
 neighbor 2001:41B0:x:y::1 remote-as 12741
 neighbor 2001:41B0:x:y::1 description NETIA
 ...
 address-family ipv4
 ...
 exit-address-family
 address-family ipv6
 ...
 exit-address-family
exit
```

no enforce-first-as - AS path vs AS neighbour

- wymagane przy IX-ach

log-neighbor-changes - logi, logi, logi

no bgp default ipv4-unicast - rodziny adresów

```
router bgp 197227
 no bgp enforce-first-as
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 ...
 neighbor 10.xx.yy.1 remote-as 197227
 neighbor 10.xx.yy.1 update-source Loopback0
 neighbor 10.xx.yy.1 description ASRv2
 ...
 neighbor 81.210.ww.zz remote-as 12741
 neighbor 81.210.ww.zz description NETIA
 ...
 neighbor FDFD:BEEF:CAFE:1::1 remote-as 197227
 neighbor FDFD:BEEF:CAFE:1::1 description ASRv2
 neighbor FDFD:BEEF:CAFE:1::1 update-source Loopback0
 ...
 neighbor 2001:41B0:x:y::1 remote-as 12741
 neighbor 2001:41B0:x:y::1 description NETIA
 ...
 address-family ipv4
 ...
 exit-address-family
 address-family ipv6
 ...
 exit-address-family
exit
```

Sesja iBGP IPv4 (z routerem wewnątrz sieci)

- **197227** - nasz AS
- **Loopback0** - interfejs sesji BGP (sąsiad BGP nie musi być sąsiadem IP)
- **description** - opisy, opisy, opisy

```
router bgp 197227
 no bgp enforce-first-as
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 ...
 neighbor 10.xx.yy.1 remote-as 197227
 neighbor 10.xx.yy.1 update-source Loopback0
 neighbor 10.xx.yy.1 description ASRv2
 ...
 neighbor 81.210.ww.zz remote-as 12741
 neighbor 81.210.ww.zz description NETIA
 ...
 neighbor FDFD:BEEF:CAFE:1::1 remote-as 197227
 neighbor FDFD:BEEF:CAFE:1::1 description ASRv2
 neighbor FDFD:BEEF:CAFE:1::1 update-source Loopback0
 ...
 neighbor 2001:41B0:x:y::1 remote-as 12741
 neighbor 2001:41B0:x:y::1 description NETIA
 ...
 address-family ipv4
 ...
 exit-address-family
 address-family ipv6
 ...
 exit-address-family
exit
```

Sesja eBGP IPv4 (poł. z routerem z innego AS-a)

```
router bgp 197227
 no bgp enforce-first-as
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 ...
 neighbor 10.xx.yy.1 remote-as 197227
 neighbor 10.xx.yy.1 update-source Loopback0
 neighbor 10.xx.yy.1 description ASRv2
 ...
 neighbor 81.210.ww.zz remote-as 12741
 neighbor 81.210.ww.zz description NETIA
 ...
 neighbor FDFD:BEEF:CAFE:1::1 remote-as 197227
 neighbor FDFD:BEEF:CAFE:1::1 description ASRv2
 neighbor FDFD:BEEF:CAFE:1::1 update-source Loopback0
 ...
 neighbor 2001:41B0:x:y::1 remote-as 12741
 neighbor 2001:41B0:x:y::1 description NETIA
 ...
 address-family ipv4
 ...
 exit-address-family
 address-family ipv6
 ...
 exit-address-family
exit
```

Sesja iBGP IPv6 (poł. z routerem z naszej sieci)

```
router bgp 197227
 no bgp enforce-first-as
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 ...
 neighbor 10.xx.yy.1 remote-as 197227
 neighbor 10.xx.yy.1 update-source Loopback0
 neighbor 10.xx.yy.1 description ASRv2
 ...
 neighbor 81.210.ww.zz remote-as 12741
 neighbor 81.210.ww.zz description NETIA
 ...
 neighbor FDFD:BEEF:CAFE:1::1 remote-as 197227
 neighbor FDFD:BEEF:CAFE:1::1 description ASRv2
 neighbor FDFD:BEEF:CAFE:1::1 update-source Loopback0
 ...
 neighbor 2001:41B0:x:y::1 remote-as 12741
 neighbor 2001:41B0:x:y::1 description NETIA
 ...
 address-family ipv4
 ...
 exit-address-family
 address-family ipv6
 ...
 exit-address-family
exit
```

Sesja eBGP IPv6 (poł. z routerem z innego AS-a)

```
router bgp 197227
 no bgp enforce-first-as
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 ...
 neighbor 10.xx.yy.1 remote-as 197227
 neighbor 10.xx.yy.1 update-source Loopback0
 neighbor 10.xx.yy.1 description ASRv2
 ...
 neighbor 81.210.ww.zz remote-as 12741
 neighbor 81.210.ww.zz description NETIA
 ...
 neighbor FDFD:BEEF:CAFE:1::1 remote-as 197227
 neighbor FDFD:BEEF:CAFE:1::1 description ASRv2
 neighbor FDFD:BEEF:CAFE:1::1 update-source Loopback0
 ...
 neighbor 2001:41B0:x:y::1 remote-as 12741
 neighbor 2001:41B0:x:y::1 description NETIA
 ...
 address-family ipv4
 ...
 exit-address-family
 address-family ipv6
 ...
 exit-address-family
exit
```

Konfiguracja BGP dla sesji IPv4 unicast

```
router bgp 197227
 no bgp enforce-first-as
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 ...
 neighbor 10.xx.yy.1 remote-as 197227
 neighbor 10.xx.yy.1 update-source Loopback0
 neighbor 10.xx.yy.1 description ASRv2
 ...
 neighbor 81.210.ww.zz remote-as 12741
 neighbor 81.210.ww.zz description NETIA
 ...
 neighbor FDFD:BEEF:CAFE:1::1 remote-as 197227
 neighbor FDFD:BEEF:CAFE:1::1 description ASRv2
 neighbor FDFD:BEEF:CAFE:1::1 update-source Loopback0
 ...
 neighbor 2001:41B0:x:y::1 remote-as 12741
 neighbor 2001:41B0:x:y::1 description NETIA
 ...
 address-family ipv4
 ...
 exit-address-family
 address-family ipv6
 ...
 exit-address-family
exit
```

Konfiguracja BGP dla sesji IPv6 unicast

```
address-family ipv4
 network 5.133.248.0 mask 255.255.252.0
 ...
 neighbor 10.xx.yy.1 activate
 neighbor 10.xx.yy.1 send-community
 neighbor 10.xx.yy.1 next-hop-self
 neighbor 10.xx.yy.1 prefix-list BGPfilter in
 neighbor 10.xx.yy.1 route-map ASRv1-in in
 neighbor 10.xx.yy.1 route-map ASRv1-out out
 ...
 neighbor 81.210.ww.zz activate
 neighbor 81.210.ww.zz send-community
 neighbor 81.210.ww.zz prefix-list BGPfilter in
 neighbor 81.210.ww.zz route-map NETIA-in in
 neighbor 81.210.ww.zz route-map NETIA-out out
 ...
exit-address-family
address-family ipv6
 network 2A03:1280::/32
 neighbor 2001:41B0:x:y::1 activate
 neighbor 2001:41B0:x:y::1 send-community
 neighbor 2001:41B0:x:y::1 prefix-list BGPfil6 in
 neighbor 2001:41B0:x:y::1 route-map NETIA-inv6 in
 neighbor 2001:41B0:x:y::1 route-map NETIA-outv6 out
 ...
 neighbor FDFD:BEEF:CAFE:1::1 activate
 neighbor FDFD:BEEF:CAFE:1::1 send-community
 neighbor FDFD:BEEF:CAFE:1::1 prefix-list BGPfil6 in
 neighbor FDFD:BEEF:CAFE:1::1 next-hop-self
 neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-inv6 in
 neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-outv6 out
exit-address-family
```

Definiujemy podsieć do rozgłoszenia.

- Pamiętamy, że musi istnieć w tablicy routingu
- Jak brak, to “static route to Null0”

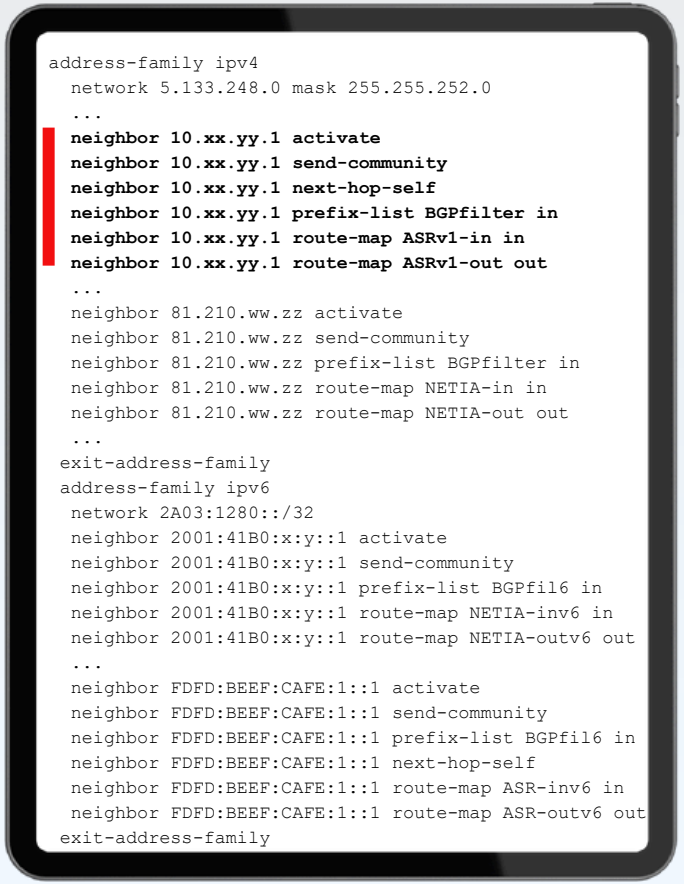

```
address-family ipv4
  network 5.133.248.0 mask 255.255.252.0
  ...
  neighbor 10.xx.yy.1 activate
  neighbor 10.xx.yy.1 send-community
  neighbor 10.xx.yy.1 next-hop-self
  neighbor 10.xx.yy.1 prefix-list BGPfilter in
  neighbor 10.xx.yy.1 route-map ASRv1-in in
  neighbor 10.xx.yy.1 route-map ASRv1-out out
  ...
  neighbor 81.210.ww.zz activate
  neighbor 81.210.ww.zz send-community
  neighbor 81.210.ww.zz prefix-list BGPfilter in
  neighbor 81.210.ww.zz route-map NETIA-in in
  neighbor 81.210.ww.zz route-map NETIA-out out
  ...
 exit-address-family
 address-family ipv6
  network 2A03:1280::/32
  neighbor 2001:41B0:x:y::1 activate
  neighbor 2001:41B0:x:y::1 send-community
  neighbor 2001:41B0:x:y::1 prefix-list BGPfil6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-inv6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-outv6 out
  ...
  neighbor FDFD:BEEF:CAFE:1::1 activate
  neighbor FDFD:BEEF:CAFE:1::1 send-community
  neighbor FDFD:BEEF:CAFE:1::1 prefix-list BGPfil6 in
  neighbor FDFD:BEEF:CAFE:1::1 next-hop-self
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-inv6 in
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-outv6 out
 exit-address-family
```

Konfiguracja iBGP dla IPv4 unicast

- **activate** - aktywacja konfiguracji
- **send-community** - możliwość wysłania commu..
- **next-hop-self** - zmiana bramy (eBGP > iBGP)
- **prefix-list (...) in** - filtracja prefiksów (u nas śmieci!)
- **route-map (...) in|out** - sterowanie ruchem

```
address-family ipv4
  network 5.133.248.0 mask 255.255.252.0
  ...
  neighbor 10.xx.yy.1 activate
  neighbor 10.xx.yy.1 send-community
  neighbor 10.xx.yy.1 next-hop-self
  neighbor 10.xx.yy.1 prefix-list BGPfilter in
  neighbor 10.xx.yy.1 route-map ASRv1-in in
  neighbor 10.xx.yy.1 route-map ASRv1-out out
  ...
  neighbor 81.210.ww.zz activate
  neighbor 81.210.ww.zz send-community
  neighbor 81.210.ww.zz prefix-list BGPfilter in
  neighbor 81.210.ww.zz route-map NETIA-in in
  neighbor 81.210.ww.zz route-map NETIA-out out
  ...
 exit-address-family
 address-family ipv6
  network 2A03:1280::/32
  neighbor 2001:41B0:x:y::1 activate
  neighbor 2001:41B0:x:y::1 send-community
  neighbor 2001:41B0:x:y::1 prefix-list BGPfil6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-inv6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-outv6 out
  ...
  neighbor FDFD:BEEF:CAFE:1::1 activate
  neighbor FDFD:BEEF:CAFE:1::1 send-community
  neighbor FDFD:BEEF:CAFE:1::1 prefix-list BGPfil6 in
  neighbor FDFD:BEEF:CAFE:1::1 next-hop-self
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-inv6 in
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-outv6 out
 exit-address-family
```

Konfiguracja eBGP dla IPv4 unicast

- **next-hop-self** - niepotrzebne

```
address-family ipv4
  network 5.133.248.0 mask 255.255.252.0
  ...
  neighbor 10.xx.yy.1 activate
  neighbor 10.xx.yy.1 send-community
  neighbor 10.xx.yy.1 next-hop-self
  neighbor 10.xx.yy.1 prefix-list BGPfilter in
  neighbor 10.xx.yy.1 route-map ASRv1-in in
  neighbor 10.xx.yy.1 route-map ASRv1-out out
  ...
  neighbor 81.210.ww.zz activate
  neighbor 81.210.ww.zz send-community
  neighbor 81.210.ww.zz prefix-list BGPfilter in
  neighbor 81.210.ww.zz route-map NETIA-in in
  neighbor 81.210.ww.zz route-map NETIA-out out
  ...
 exit-address-family
 address-family ipv6
  network 2A03:1280::/32
  neighbor 2001:41B0:x:y::1 activate
  neighbor 2001:41B0:x:y::1 send-community
  neighbor 2001:41B0:x:y::1 prefix-list BGPfil6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-inv6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-outv6 out
  ...
  neighbor FDFD:BEEF:CAFE:1::1 activate
  neighbor FDFD:BEEF:CAFE:1::1 send-community
  neighbor FDFD:BEEF:CAFE:1::1 prefix-list BGPfil6 in
  neighbor FDFD:BEEF:CAFE:1::1 next-hop-self
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-inv6 in
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-outv6 out
 exit-address-family
```

Definiujemy podsieć do rozgłoszenia.

- Pamiętamy, że musi istnieć w tablicy routingu
- Jak brak, to “static route to Null0”

```
address-family ipv4
  network 5.133.248.0 mask 255.255.252.0
  ...
  neighbor 10.xx.yy.1 activate
  neighbor 10.xx.yy.1 send-community
  neighbor 10.xx.yy.1 next-hop-self
  neighbor 10.xx.yy.1 prefix-list BGPfilter in
  neighbor 10.xx.yy.1 route-map ASRv1-in in
  neighbor 10.xx.yy.1 route-map ASRv1-out out
  ...
  neighbor 81.210.ww.zz activate
  neighbor 81.210.ww.zz send-community
  neighbor 81.210.ww.zz prefix-list BGPfilter in
  neighbor 81.210.ww.zz route-map NETIA-in in
  neighbor 81.210.ww.zz route-map NETIA-out out
  ...
 exit-address-family
 address-family ipv6
  network 2A03:1280::/32
  neighbor 2001:41B0:x:y::1 activate
  neighbor 2001:41B0:x:y::1 send-community
  neighbor 2001:41B0:x:y::1 prefix-list BGPfil6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-inv6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-outv6 out
  ...
  neighbor FDFD:BEEF:CAFE:1::1 activate
  neighbor FDFD:BEEF:CAFE:1::1 send-community
  neighbor FDFD:BEEF:CAFE:1::1 prefix-list BGPfil6 in
  neighbor FDFD:BEEF:CAFE:1::1 next-hop-self
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-inv6 in
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-outv6 out
 exit-address-family
```

Konfiguracja iBGP dla IPv6 unicast

- analogicznie jak dla IPv4

```
address-family ipv4
  network 5.133.248.0 mask 255.255.252.0
  ...
  neighbor 10.xx.yy.1 activate
  neighbor 10.xx.yy.1 send-community
  neighbor 10.xx.yy.1 next-hop-self
  neighbor 10.xx.yy.1 prefix-list BGPfilter in
  neighbor 10.xx.yy.1 route-map ASRv1-in in
  neighbor 10.xx.yy.1 route-map ASRv1-out out
  ...
  neighbor 81.210.ww.zz activate
  neighbor 81.210.ww.zz send-community
  neighbor 81.210.ww.zz prefix-list BGPfilter in
  neighbor 81.210.ww.zz route-map NETIA-in in
  neighbor 81.210.ww.zz route-map NETIA-out out
  ...
 exit-address-family
 address-family ipv6
  network 2A03:1280::/32
  neighbor 2001:41B0:x:y::1 activate
  neighbor 2001:41B0:x:y::1 send-community
  neighbor 2001:41B0:x:y::1 prefix-list BGPfil6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-inv6 in
  neighbor 2001:41B0:x:y::1 route-map NETIA-outv6 out
  ...
  neighbor FDFD:BEEF:CAFE:1::1 activate
  neighbor FDFD:BEEF:CAFE:1::1 send-community
  neighbor FDFD:BEEF:CAFE:1::1 prefix-list BGPfil6 in
  neighbor FDFD:BEEF:CAFE:1::1 next-hop-self
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-inv6 in
  neighbor FDFD:BEEF:CAFE:1::1 route-map ASR-outv6 out
 exit-address-family
```

Konfiguracja iBGP dla IPv6 unicast

- analogicznie jak dla IPv4

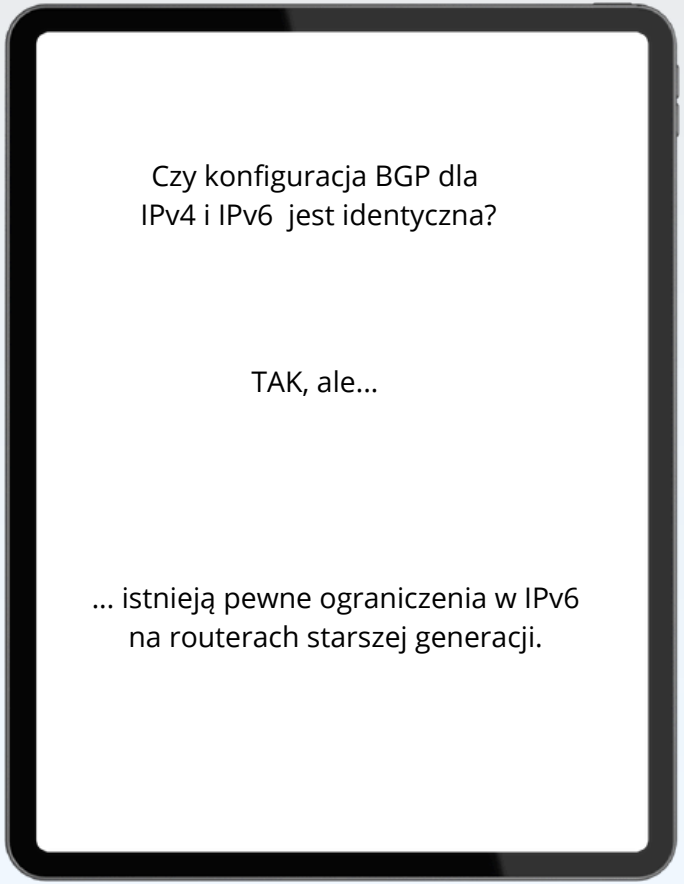

Brak wsparcia dla
“send-community”

Ograniczenia dla
“ACL IPv6”

Brak wsparcia dla
“prefix-list”

Ograniczenie w długości
maski: /126 /127
Akceptowalna: /64

Proporcje ruchu: IPv4 vs IPv6

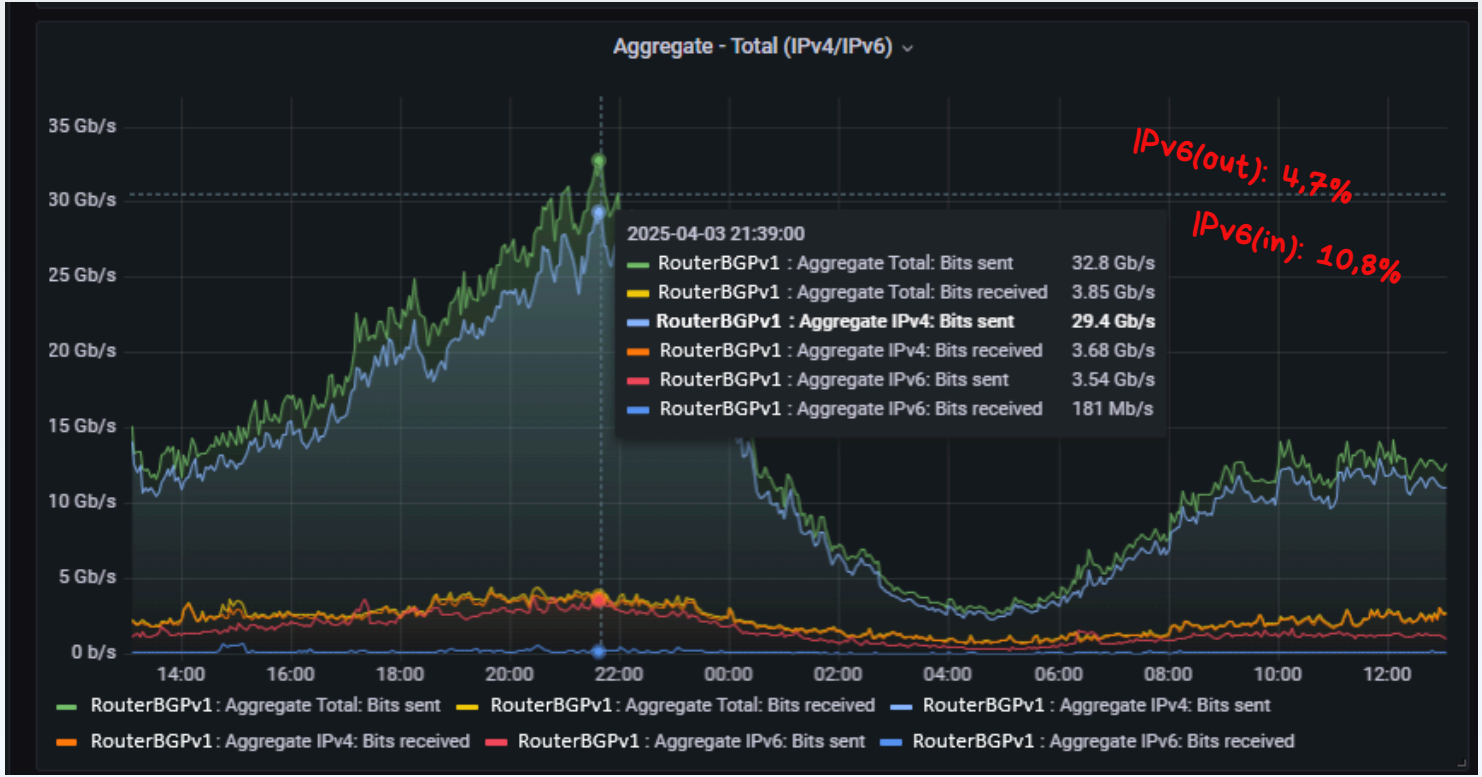

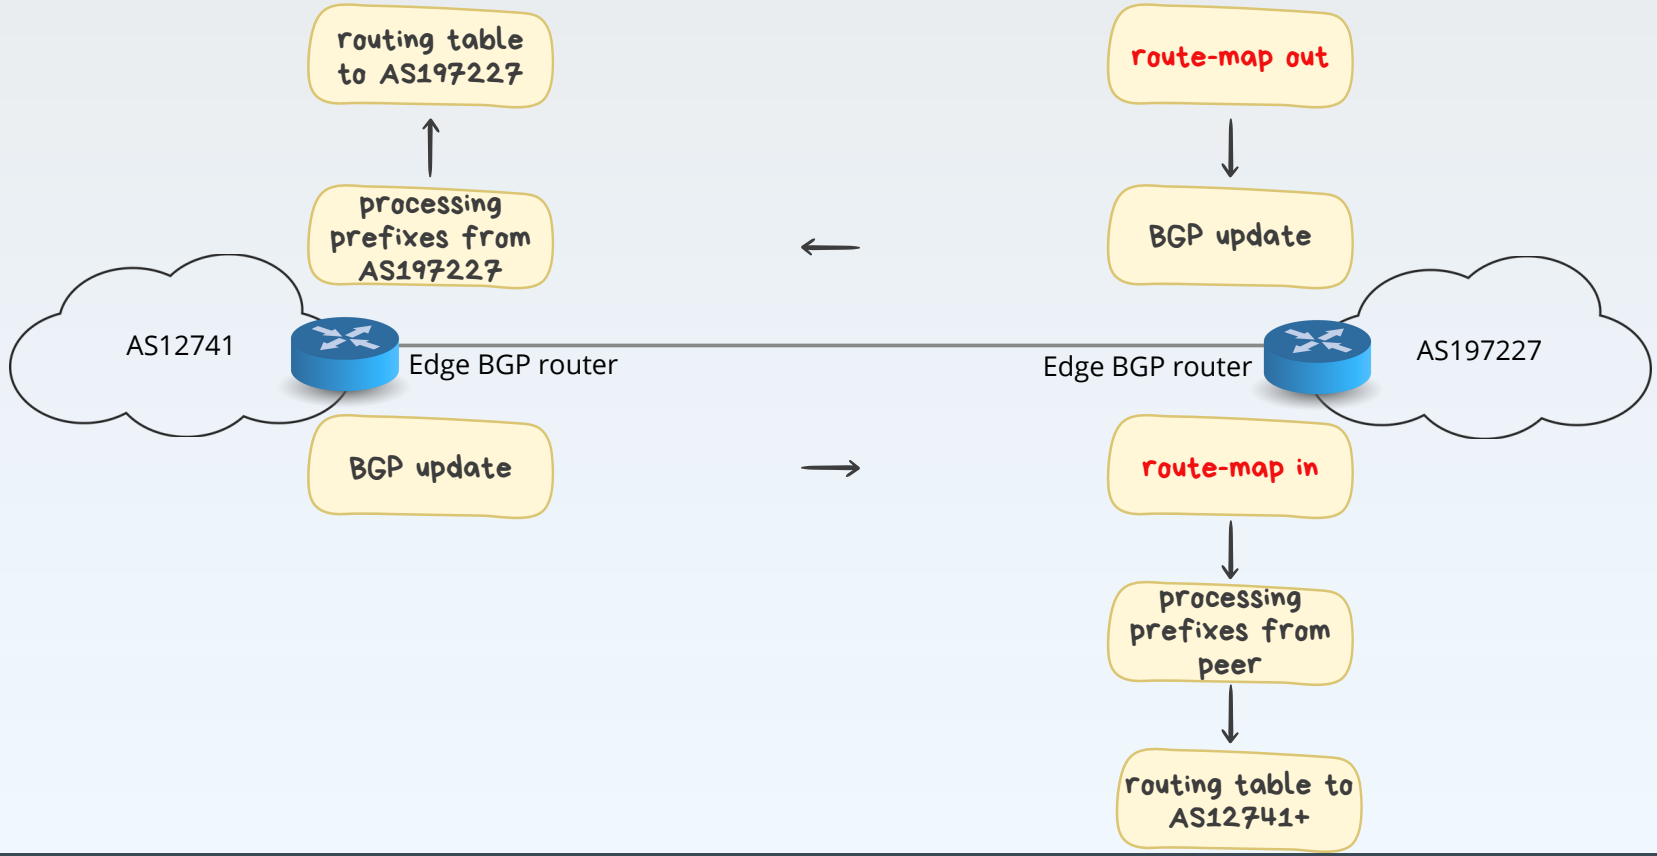
routing table to AS197227
route-map out
processing prefixes from AS197227
BGP update
AS12741
Edge BGP router
Edge BGP router
AS197227
BGP update
route-map in
processing prefixes from peer
routing table to AS12741+

```
!
route-map NETIA-out permit 1800
 match ip address prefix-list WTvK
 set as-path prepend 197227 197227
 set community 12741:31000
 !
route-map NETIA-out permit 2000
 description OurPrefix_default
 match ip address prefix-list WTvK
!
!
route-map NETIA-in permit 1800
 description JustAnExample
 match ip address prefix-list JustAnExample
 set metric 50
 set local-preference 200
 set community 1:101 additive
!
route-map NETIA-in permit 2000
 description ALL_default
 set local-preference 100
!
```

Domyślne działanie:

- **out** - rozgłoś **nasze** prefiksy
- **in** - przyjmij **wszystkie** prefiksy

```
!
route-map NETIA-out permit 1800
 match ip address prefix-list WTvK
 set as-path prepend 197227 197227
 set community 12741:31000
 !
route-map NETIA-out permit 2000
 description OurPrefix_default
 match ip address prefix-list WTvK
!
!
route-map NETIA-in permit 1800
 description JustAnExample
 match ip address prefix-list JustAnExample
 set metric 50
 set local-preference 200
 set community 1:101 additive
!
route-map NETIA-in permit 2000
 description ALL_default
 set local-preference 100
!
```

Rozgłoś **nasze** prefiksy i

- wydłuż trasę (AS path) o dwa ASy
- ustaw community na 12741:31000 (o tym później)

```
!
route-map NETIA-out permit 1800
 match ip address prefix-list WTvK
 set as-path prepend 197227 197227
 set community 12741:31000
 !
route-map NETIA-out permit 2000
 description OurPrefix_default
 match ip address prefix-list WTvK
!
!
route-map NETIA-in permit 1800
 description JustAnExample
 match ip address prefix-list JustAnExample
 set metric 50
 set local-preference 200
 set community 1:101 additive
!
route-map NETIA-in permit 2000
 description ALL_default
 set local-preference 100
!
```

Pobierz prefiksy, które pasują do prefix-list'y i

- ustaw MED na 50
- ustaw local-preference na 200
- **dodaj** community 1:101

match:

- local-preference
- community
- as-path
- .....

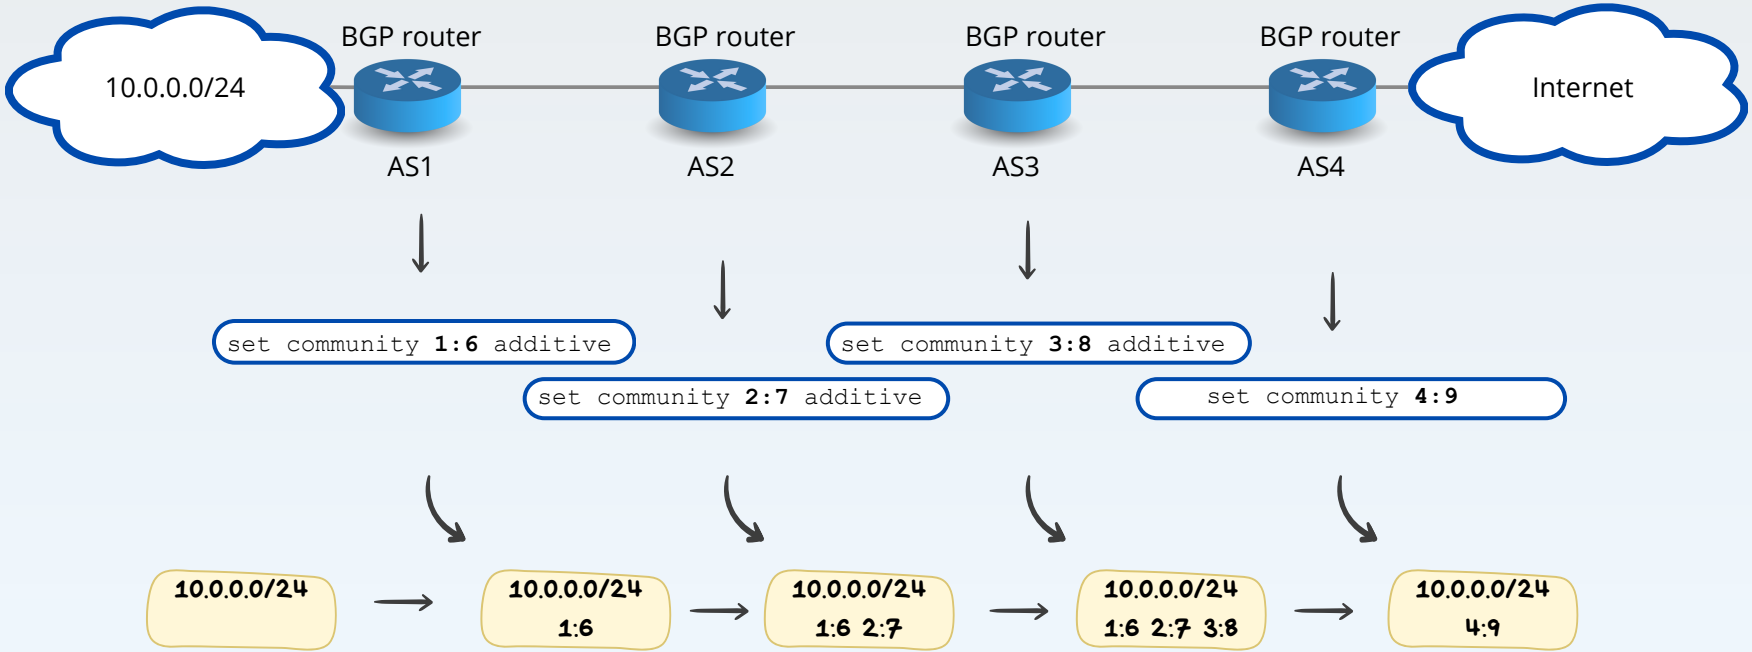
10.0.0.0/24
BGP router
BGP router
BGP router
BGP router
Internet
AS1
AS2
AS3
AS4
set community 1:6 additive
set community 2:7 additive
set community 3:8 additive
set community 4:9
10.0.0.0/24
10.0.0.0/24
1:6
10.0.0.0/24
1:6 2:7
10.0.0.0/24
1:6 2:7 3:8
10.0.0.0/24
4:9

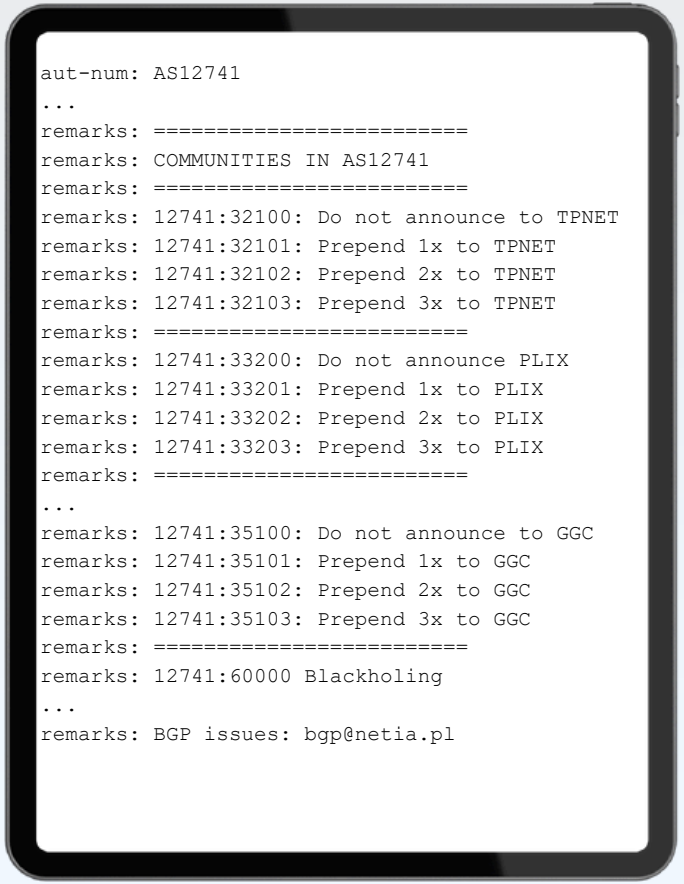

```
aut-num: AS12741
...
remarks: ========================
remarks: COMMUNITIES IN AS12741
remarks: ========================
remarks: 12741:32100: Do not announce to TPNET
remarks: 12741:32101: Prepend 1x to TPNET
remarks: 12741:32102: Prepend 2x to TPNET
remarks: 12741:32103: Prepend 3x to TPNET
remarks: ========================
remarks: 12741:33200: Do not announce PLIX
remarks: 12741:33201: Prepend 1x to PLIX
remarks: 12741:33202: Prepend 2x to PLIX
remarks: 12741:33203: Prepend 3x to PLIX
remarks: ========================
...
remarks: 12741:35100: Do not announce to GGC
remarks: 12741:35101: Prepend 1x to GGC
remarks: 12741:35102: Prepend 2x to GGC
remarks: 12741:35103: Prepend 3x to GGC
remarks: ========================
remarks: 12741:60000 Blackholing
...
remarks: BGP issues: bgp@netia.pl
```

Spis **community** w bazie RIPE NCC w obiekcie **asn-num**

Więcej community? Kontakt z **operatorem**

Sterowanie ruchem przychodzącym:

- route-map (...) out
- rozgłaszamy nasze prefiksy

Prepand - wydłużenie ścieżki: 1x = o 1AS

Blackholing:

- zazwyczaj tylko prefiksy **/32**
- przekierowanie na **Null0**

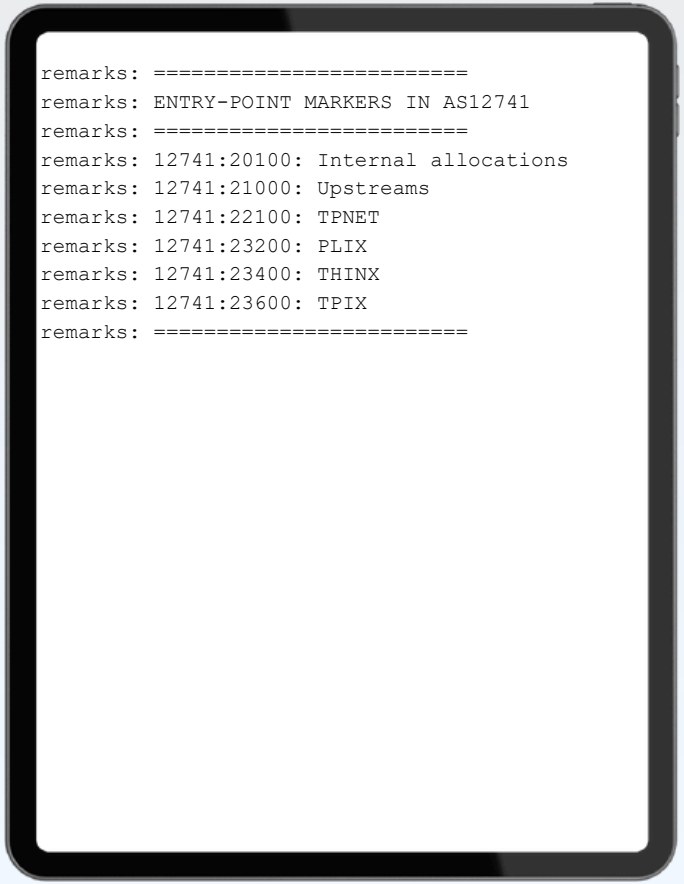

```
remarks: ========================
remarks: ENTRY-POINT MARKERS IN AS12741
remarks: ========================
remarks: 12741:20100: Internal allocations
remarks: 12741:21000: Upstreams
remarks: 12741:22100: TPNET
remarks: 12741:23200: PLIX
remarks: 12741:23400: THINX
remarks: 12741:23600: TPIX
remarks: ========================
```

Źródła prefiksów w sieci AS12741.

Sterowanie ruchem wychodzącym:

- route-map (...) in
- decyzja o ruchu wychodzącym

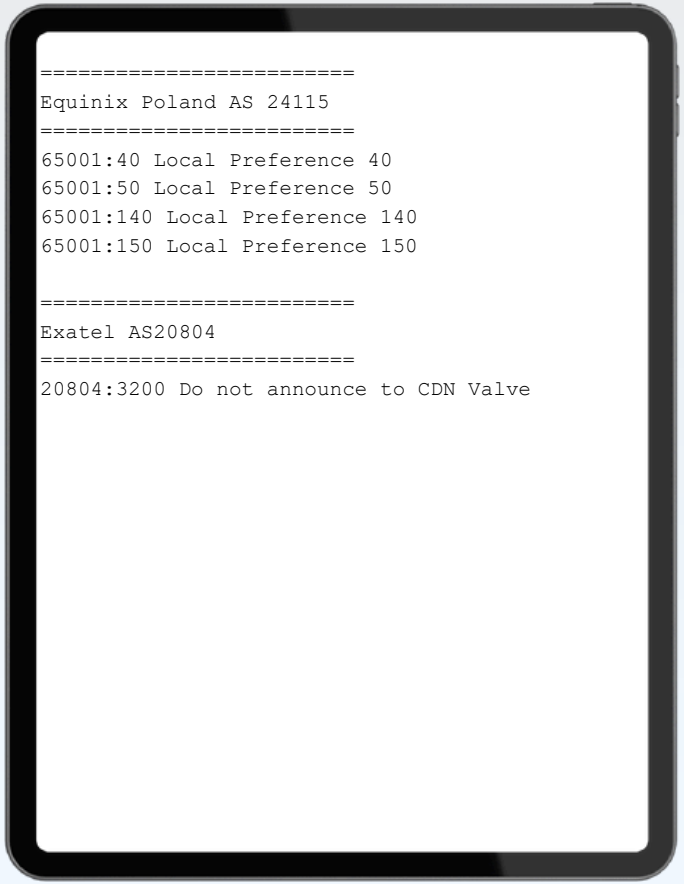

```
=========================
Equinix Poland AS 24115
=========================
65001:40 Local Preference 40
65001:50 Local Preference 50
65001:140 Local Preference 140
65001:150 Local Preference 150

=========================
Exatel AS20804
=========================
20804:3200 Do not announce to CDN Valve
```

Punkt wymiany ruchu Equinix Poland

Sterowanie ruchem wychodzącym:

- route-map (...) in
- decyzja o trasie

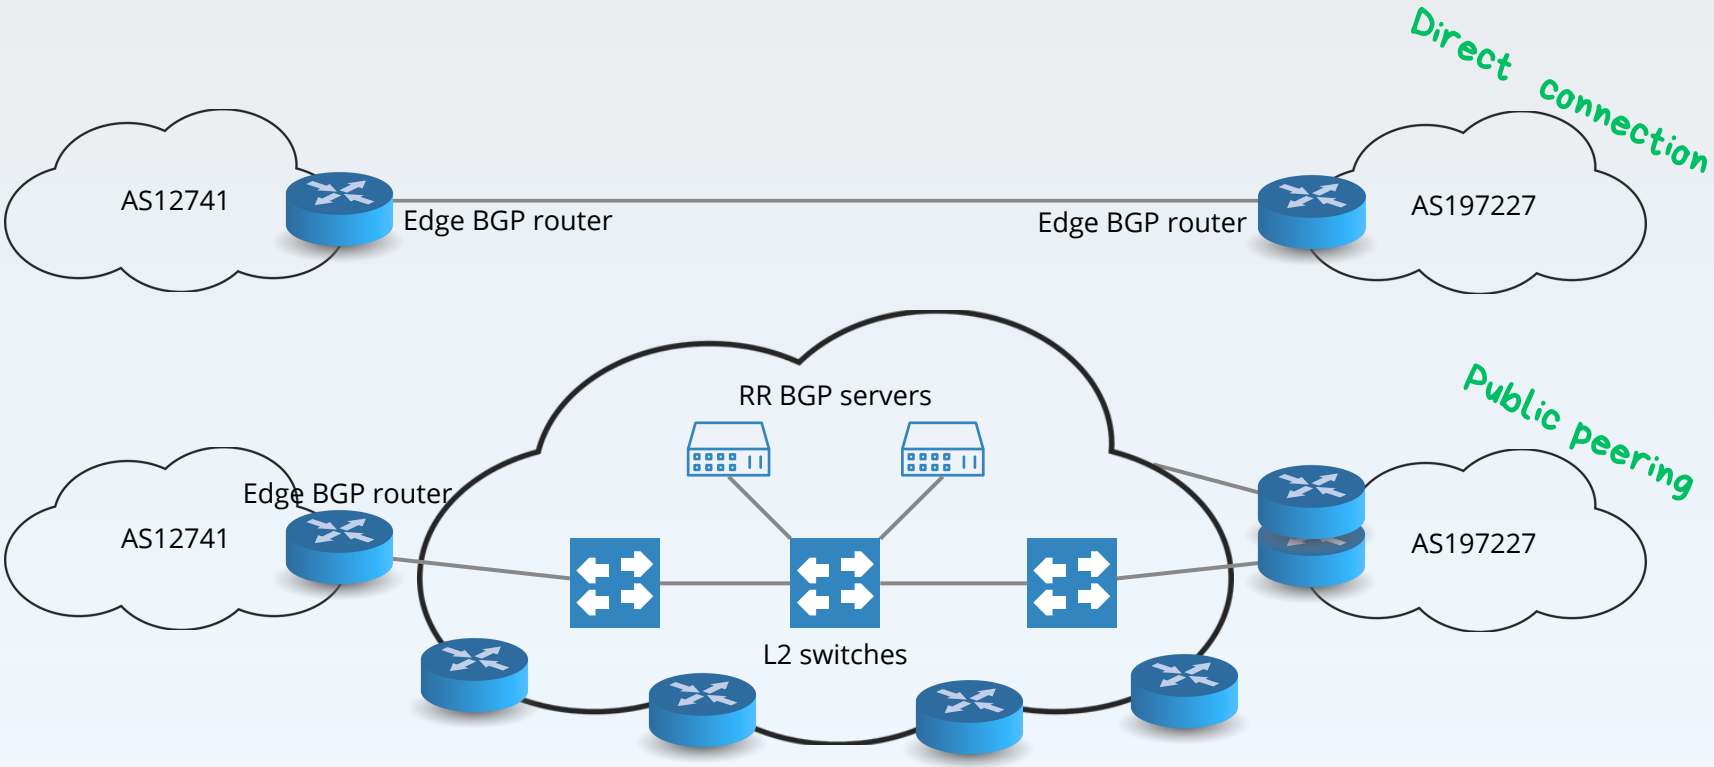
Direct connection
AS12741
Edge BGP router
Edge BGP router
AS197227
Public peering
RR BGP servers
Edge BGP router
AS12741
AS197227
L2 switches

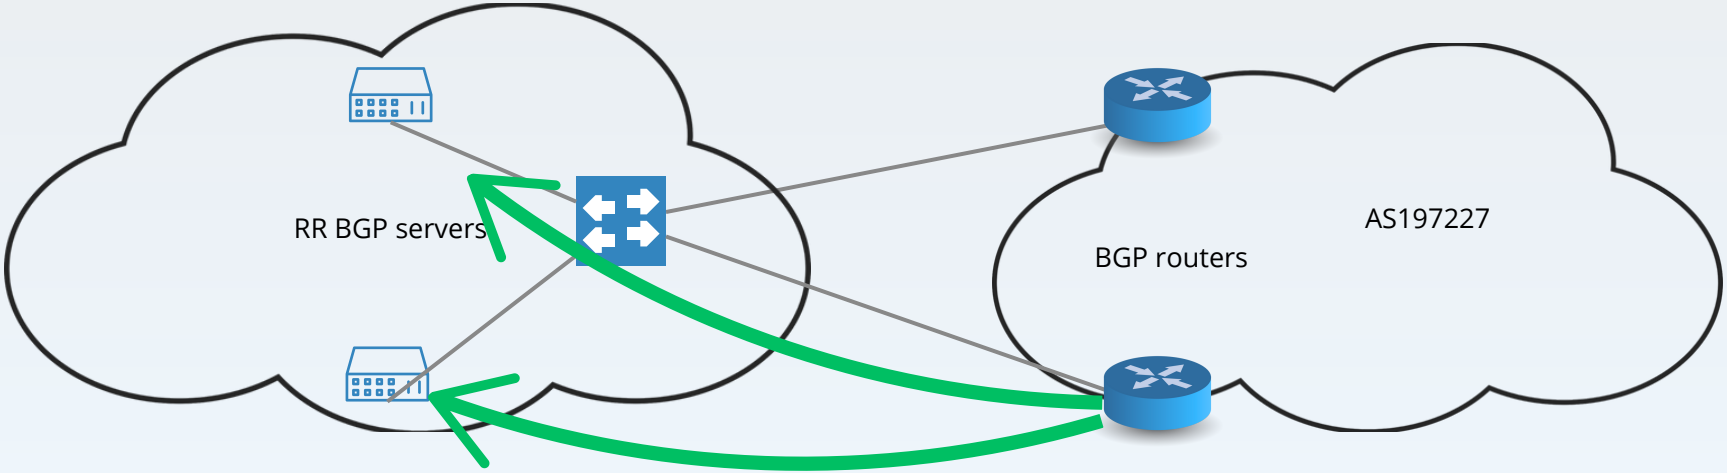

Sesje BGP: każdy BGP router - z każdym RR - IPv4/IPv6


Private peering
AS5617
własne CDNy
niższa cena
Odporny na DDoS
Public selective peering
AS6939
AS32934
only / on request
dodatkowe prefiksy
kontrola peerów

peeringdb.com

Organization	Netia S.A.
Also Known As	Netia
Long Name	
Company Website	https://www.netia.pl
ASN	12741
IRR as-set/route-set	RIPE::AS-NETIA
Route Server URL	
Looking Glass URL	http://lg.netia.pl/
Network Types	NSP
IPv4 Prefixes	3000
IPv6 Prefixes	100
Traffic Levels	1-5Tbps
Traffic Ratios	Mostly Inbound
Geographic Scope	Europe
Protocols Supported	✓ Unicast IPv4 ○ Multicast ✓ IPv6 ○ Never via route servers
Last Updated	2025-02-24T06:31:27Z
Public Peering Info Updated	2024-06-04T15:32:24Z
Peering Facility Info Updated	2025-02-24T06:30:56Z
Contact Info Updated	2018-04-03T07:41:20Z
Notes	
RIR Status	ok
RIR Status Updated	2024-06-26T04:47:55Z

Public Peering Exchange Points

Filter

Exchange A↕ IPv4	ASN IPv6	Speed Port Location	RS Peer	BFD Support
1-IX EU 185.1.254.11	12741 2001:7f8:115:1::11	100G	✓	○
BCIX 193.178.185.46	12741 2001:7f8:19:1::31c5:1	100G	✓	○
DE-CIX Frankfurt 80.81.192.193	12741 2001:7f8::31c5:0:1	100G	✓	○
EPIX Katowice 178.216.40.86	12741	20G	✓	○
EPIX Warszawa 89.46.144.212	12741 2001:678:3ac::212	100G	✓	○
Equinix Warsaw 195.182.218.13	12741 2001:7f8:42::a501:2741:1	100G	✓	○
MSK-IX Moscow 195.208.210.9	12741 2001:7f8:20:101::210:9	20G	✓	○
NIX.CZ 91.210.17.31	12741 2001:7f8:14::5c:1	50G	✓	○
NL-ix 193.239.118.72	12741 2001:7f8:13::a501:2741:1	100G	✓	○
POZIX 185.1.4.44	12741 2001:7f8:4b::2c	100G	✓	○
THINX Warsaw 212.91.0.252	12741 2001:7f8:60::1026:8	100G	✓	○
TPIX PL	12741	100G	✓	○

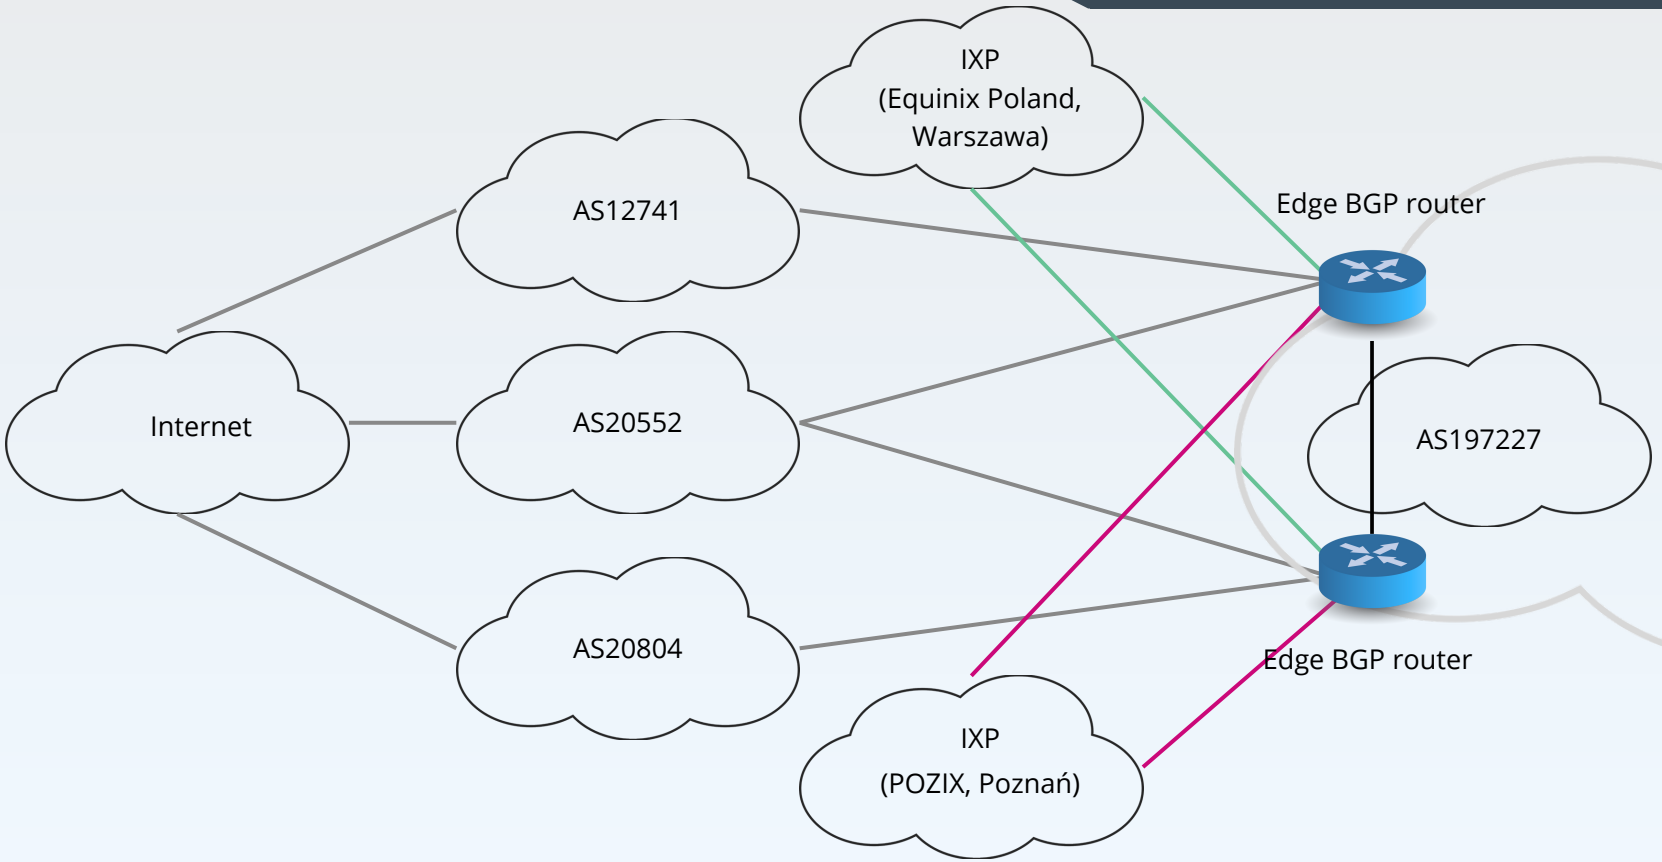
IXP
(Equinix Poland,
Warszawa)
AS12741
Edge BGP router
Internet
AS20552
AS197227
AS20804
Edge BGP router
IXP
(POZIX, Poznań)

Proporcje ruchu: przychodzący vs wychodzący

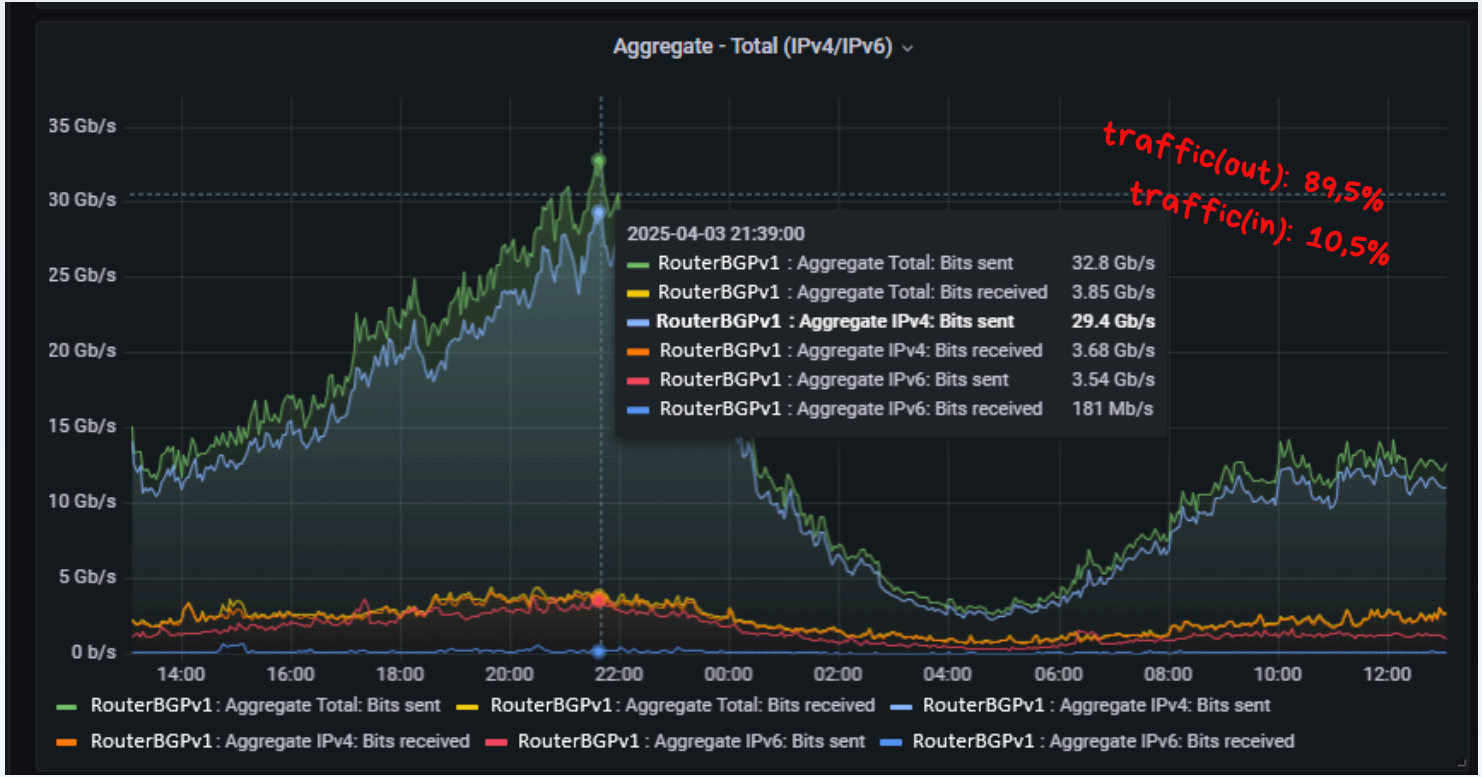

Case #1: General outband policy

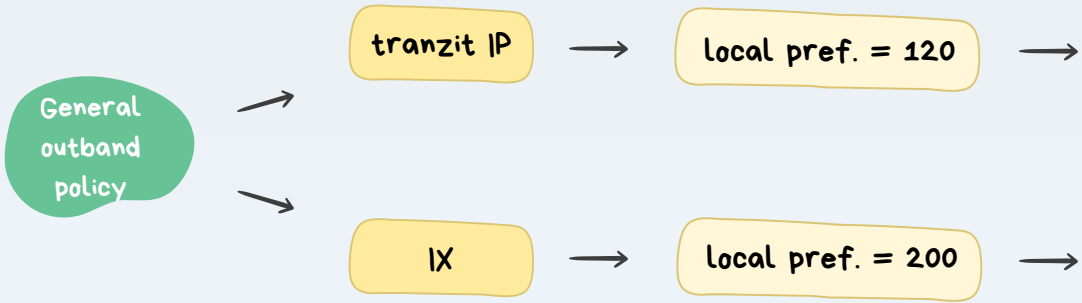


```
route-map NETIA-in permit 1900
 description ALL_PREF120
 set local-preference 120
route-map NETIA-in permit 2000
 description ALL_default
 set local-preference 100
```

```
route-map POZIX-rs1-in 1900
 description ALL_PREF200
 set local-preference 200
route-map POZIX-rs1-in permit 2000
 description ALL_default
 set local-preference 100
```

Case #2: VoIP to AS12741 - problem of quality

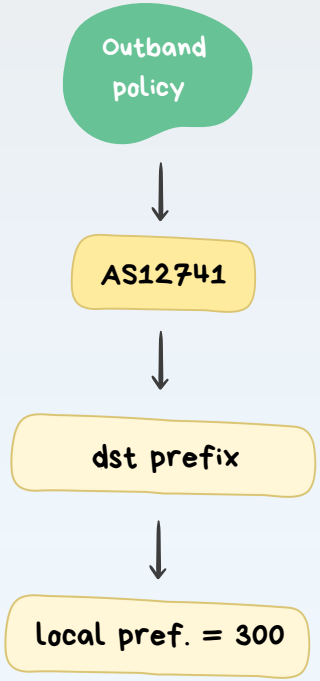


```
ip prefix-list Nvoip description DDmmYY Borys
ip prefix-list Nvoip seq 5 permit 84.40.236.0/24

route-map NETIA-in permit 1800
 description VoIP-DDmmYY
 match ip address prefix-list NETIAvoip
 set local-preference 300
route-map NETIA-in permit 1900
 description ALL_PREF120
 set local-preference 120
route-map NETIA-in permit 2000
 description ALL_default
 set local-preference 100
```

Case #3: Load Balancing of inbound traffic

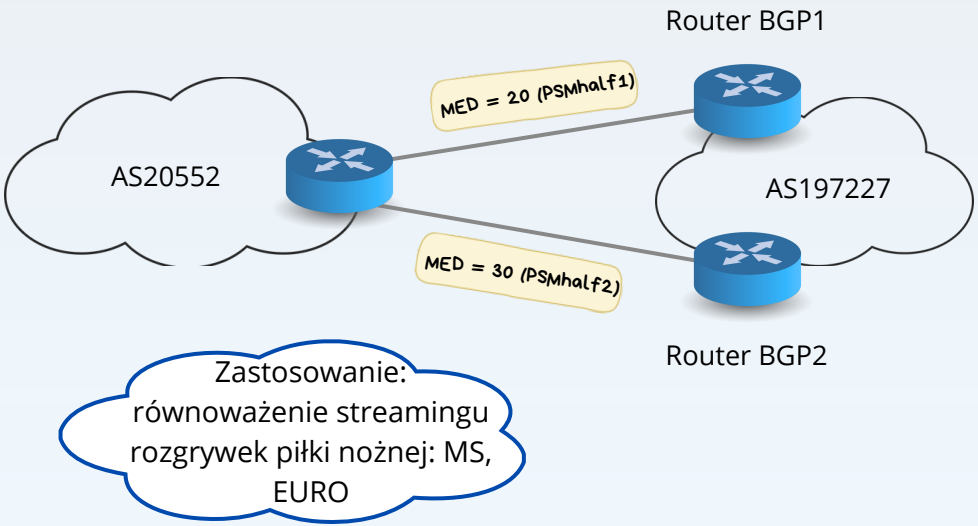


```
ip prefix-list PSMhalf1 seq (...)
ip prefix-list PSMhalf2 seq (...)
...
!
! Router BGP1
!
route-map HTI-out permit 1800
 description OurPref-Pref120Med20
 match ip address prefix-list PSMhalf1
 set metric 20
!
route-map HTI-out permit 1900
 description OurPref-Pref120Med50
 match ip address prefix-list PSMhalf2
 set metric 50
!
! Router BGP2
!
route-map HTI-out permit 1800
 description OurPref-Pref120Med40
 match ip address prefix-list PSMhalf1
 set metric 40
!
route-map HTI-out permit 1900
 description OurPref-Pref120Med30
 match ip address prefix-list PSMhalf2
 set metric 30
```

Case #4: Load Balancing of inbound traffic

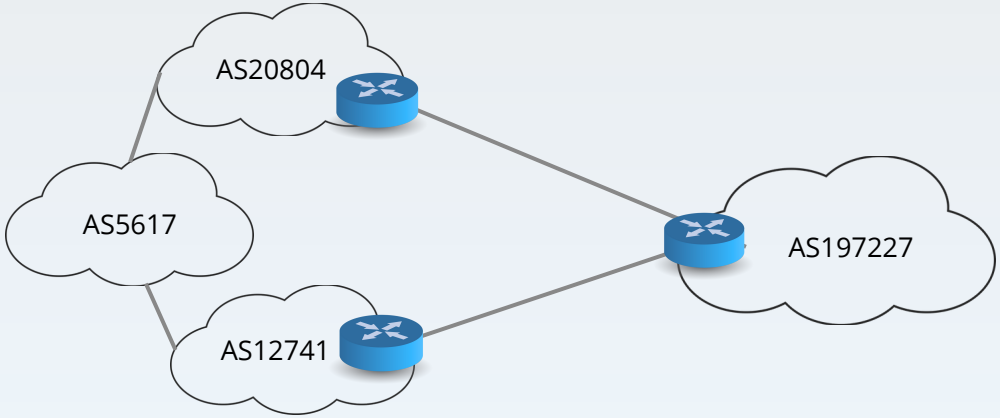


```
remarks: ============================
remarks: 12741:32100: Do not announce to TPNET
remarks: 12741:32101: Prepend 1x to TPNET
remarks: 12741:32102: Prepend 2x to TPNET
remarks: 12741:32103: Prepend 3x to TPNET
remarks:  ============================
```

```
route-map Exatel-out permit 2000
 description OurPref-default
 match ip address prefix-list PSM

route-map NETIA-out permit 1800
 description OurPref-costam-borys-250408
 match ip address prefix-list PSM
 set as-path prepend 197227 197227
 set community 12741:32102
!
route-map NETIA-out permit 2000
 description OurPref-default
 match ip address prefix-list PSM
```

Two paths:

- 20804 197227
- 12741 12741 12741 197227 197227 197227

Case #5: Forcing the traffic from Valve (Steam Platform)

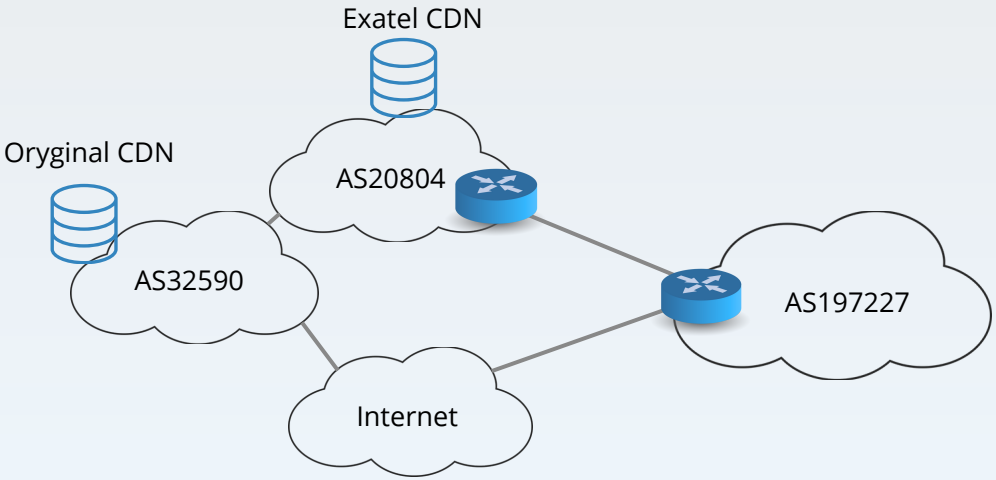


```
route-map Exatel-out permit 1500
 match ip address prefix-list PSM
 set community 20804:3200
```

Kontent dostarczony z CDN Exatel → Community 20804:3200 - nie wysyłaj kontentu z lokalnego CDNa → Kontent dostarczony z CDN Valve

Case #6: DDoS Blackholing

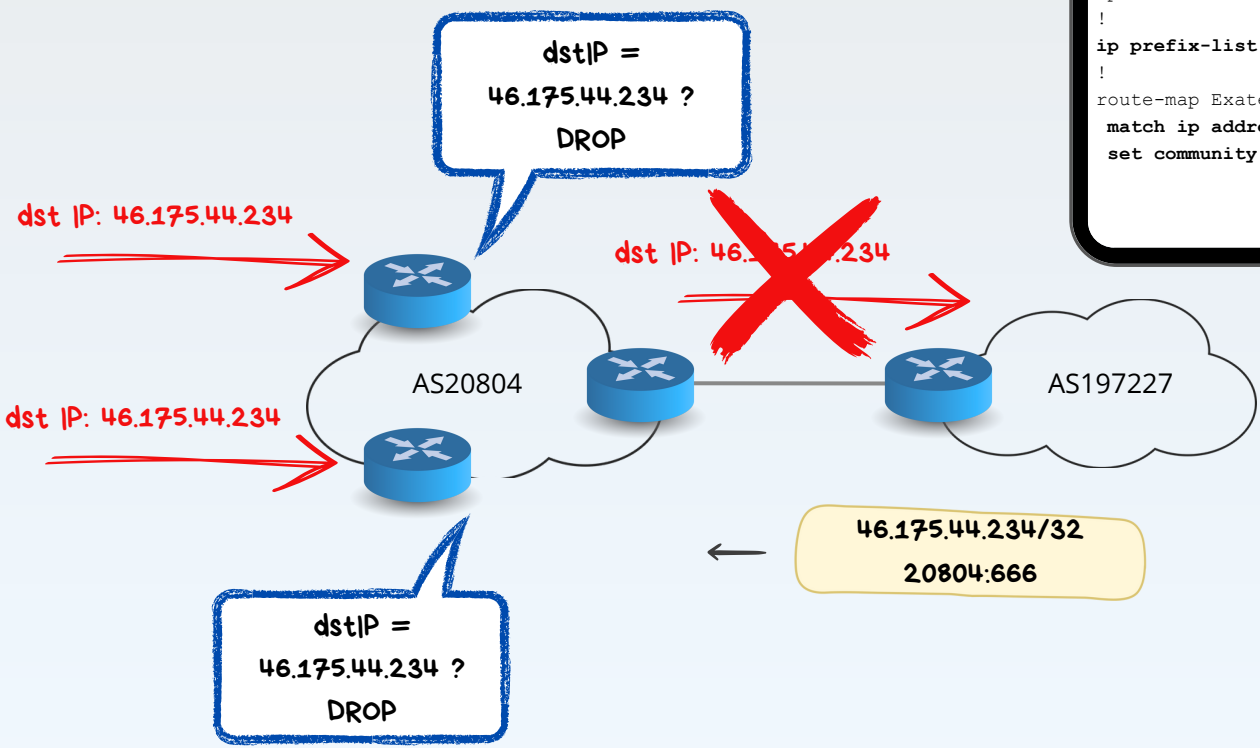


```
ip route 46.175.44.234 255.255.255.255 Null0
!
ip prefix-list DDoS seq 20 permit 46.175.44.234/32
!
route-map Exatel-out permit 1000
 match ip address prefix-list DDoS
 set community 20804:666
```

Case #7: Announcing smaller subnets

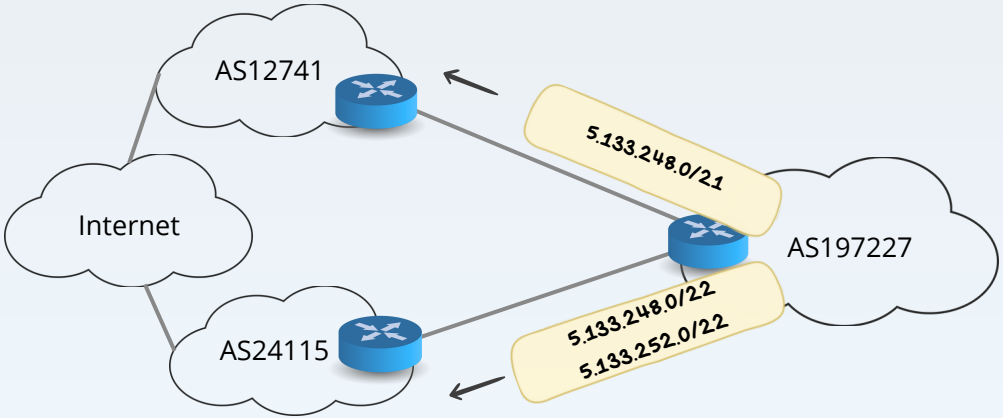


```
!
ip prefix-list PSM seq 20 permit 5.133.248.0/21
!
ip prefix-list PSMextra seq 20 permit 5.133.248.0/22
ip prefix-list PSMextra seq 30 permit 5.133.252.0/22
!
route-map NETIA-out permit 1200
 match ip address prefix-list PSM
!
route-map PLIX-rs1-out permit 1200
 match ip address prefix-list PSMextra
```

Case #7: Announcing smaller subnets

core2.abq1.he.net> show ip bgp routes detail 5.133.248.0

Matching Routes	2
Status Codes	**A** - Aggregate **B** - Best **b** - Not Install Best **C** - Confederation eBGP **D** - Damped **E** - eBGP **H** - History **I** - iBGP **L** - Local **M** - Multipath **S** - Suppressed **F** - Filtered **s** - Stale **x** - Best-External

Status	Network	Next Hop	Learned	Metric	LocPrf	MED	Weight	Path	Origin	ROA
BMI	5.133.248.0/22	195.182.218.94	216.218.252.24 (6939)	1355	100	0	0	197227	IGP	✓
MI	5.133.248.0/22	195.182.218.94	216.218.253.19 (6939)	1355	100	0	0	197227	IGP	✓

Case #7: Announcing smaller subnets

Prefix	Description	Visibility
5.133.248.0/21	Poznanska Spoldzielnia Mieszkaniowa "Winogrady" w Poznaniu	**100%** 638/638
5.133.248.0/22	Poznanska Spoldzielnia Mieszkaniowa "Winogrady" w Poznaniu	**23%** 145/638
5.133.252.0/22		**23%** 145/638
37.247.56.0/21	Poznanska Spoldzielnia Mieszkaniowa "Winogrady" w Poznaniu	**100%** 638/638
37.247.56.0/22		**23%** 145/638
37.247.60.0/22		**23%** 145/638
46.175.40.0/21	Poznanska Spoldzielnia Mieszkaniowa "Winogrady" w Poznaniu	**100%** 638/638
46.175.44.234/32		**1%** 6/638

Important links

peeringdb.com

stat.ripe.net

lg.he.net

bgp.he.net

Dziękuję za uwagę!