RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Your First Steps to IPv6

Announcements

A hands-on tutorial

M. Stucchi & N. Trenaman | 9 October 2017| ENOG14

Agenda

- Getting your IPv6 Addresses
- Configuring Interfaces & Loopbacks
- Creating a Route6 object
- MP-BGP explained
- Setting up a Transit
- Setting up Peering
- Filtering Best Practices

Getting Your IPv6 Addresses

Getting an IPv6 allocation

- To qualify, an organisation must:
  - Be an LIR
  - Have a plan for making assignments within two years
- Minimum allocation size /32
  - Up to a /29 without additional justification
  - More if justified by customer numbers and network extension
  - Additional bits based on hierarchical and geographical structure, planned longevity and security levels

Getting IPv6 PI address space

- To qualify, an organisation must:
  - Meet the contractual requirements for provider independent resources
  - LIRs must demonstrate special routing requirements
- Minimum assignment size: /48
- PI space can not be used for sub-assignments
  - not even 1 IP address!

Where to get your IPv6 Addresses

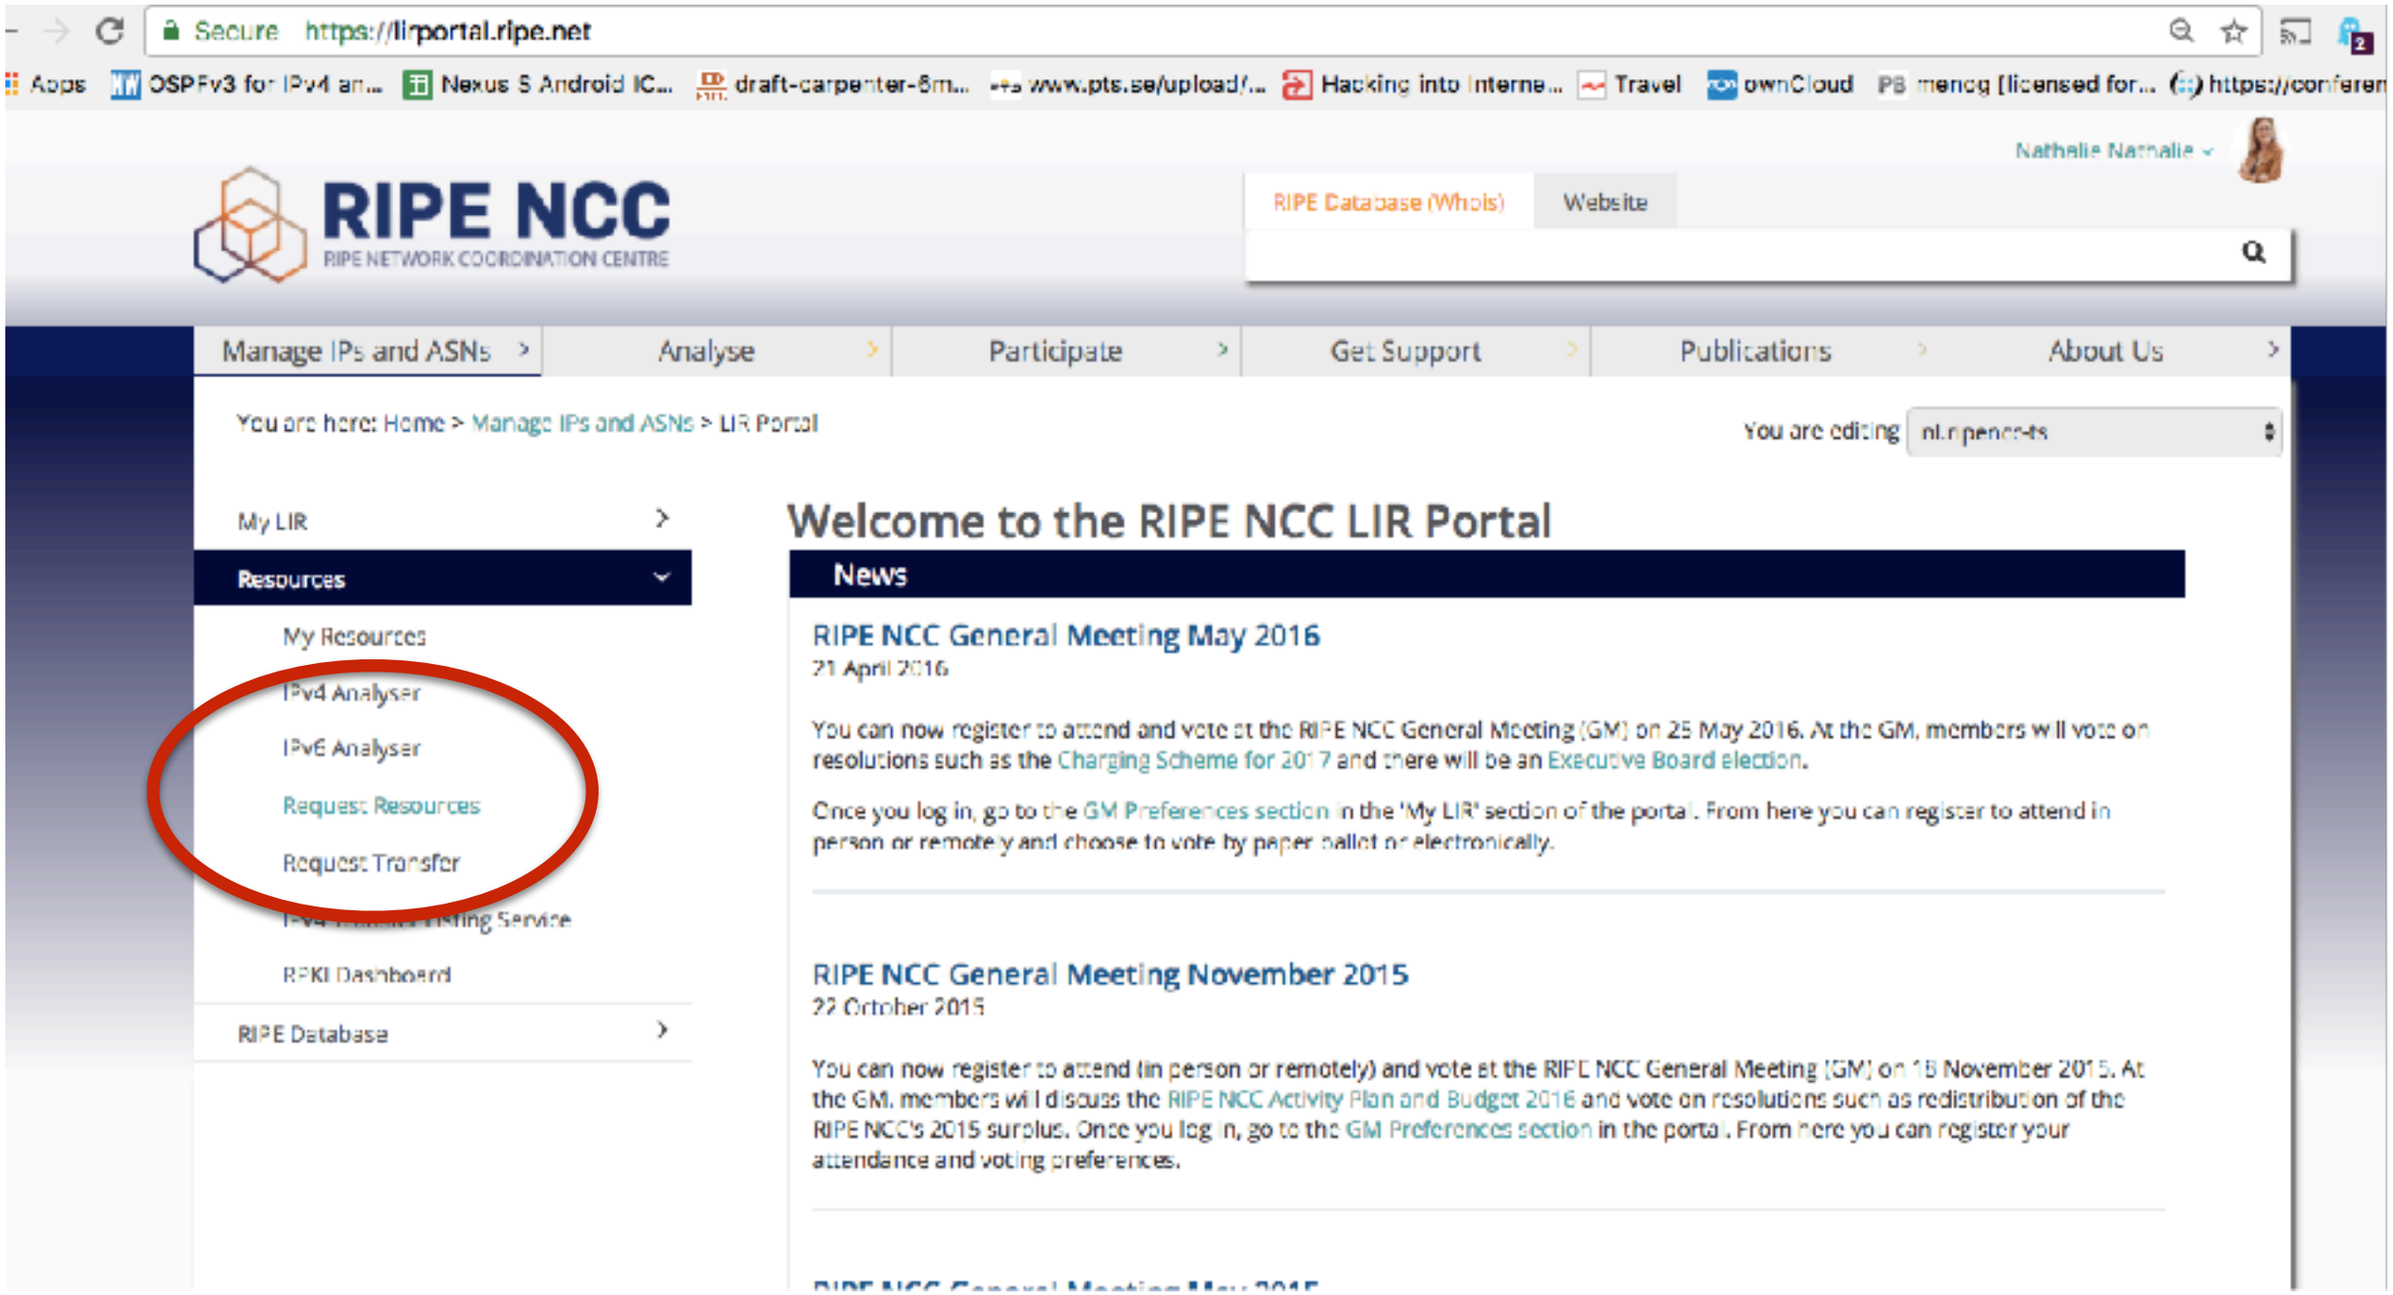

How to get your IPv6 Allocation

What would you like to request?

- (•) IPv6 Addresses
- ( ) IPv4 Addresses
- ( ) ASN (Autonomous System Number)
- ( ) IXP (Internet Exchange Point) Assignment

Continue

Not sure? Read the supporting notes.

What kind of IPv6 address space do you want to request?

Please request an IPv6 allocation before requesting an IPv6 End User Site Assignment.

- (•) IPv6 Allocation
- ( ) IPv6 Provider Independent (PI) Assignment
- ( ) Temporary IPv6 Provider Independent (PI) Assignment
- ( ) IPv6 End User Site Assignment
- ( ) IPv6 Anycast Assignment

< Continue

Not sure? Read the supporting notes.

How to get your IPv6 Allocation

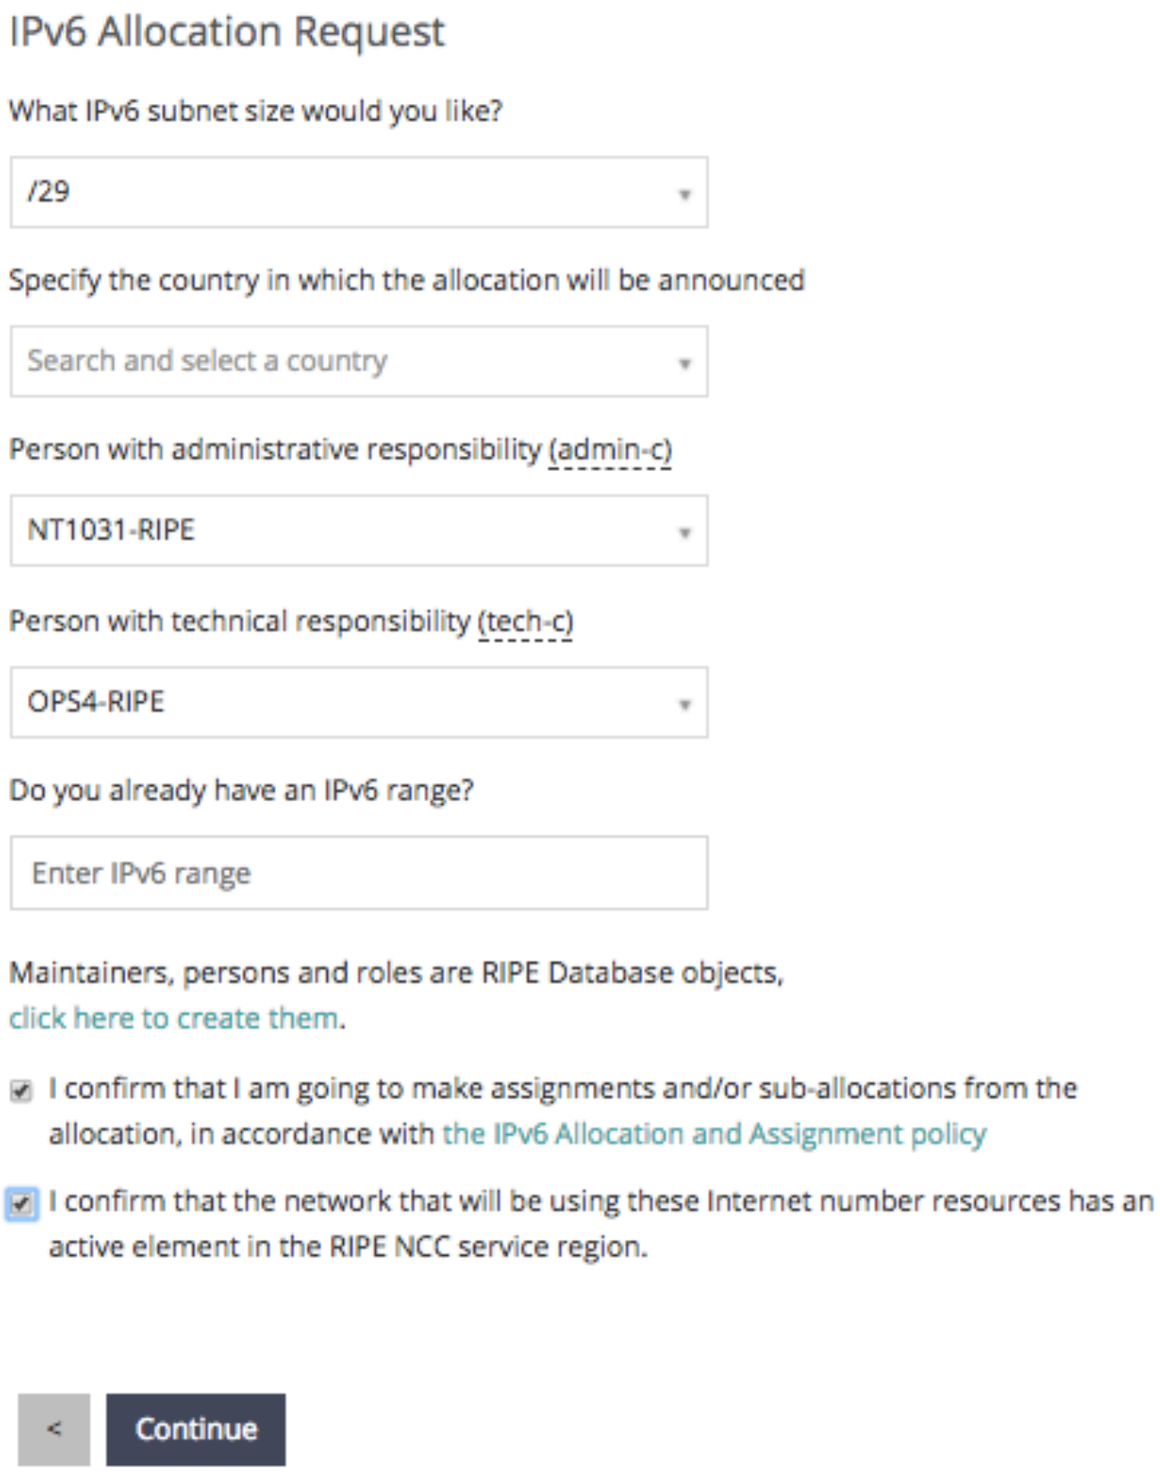

Your IPv6 Allocation - Example

```
inet6num:          2001:db8::/29
netname:           BY-ENOG-20170901
country:           BY
admin-c:           ADM321-RIPE
tech-c:            NOC123-RIPE
status:            ALLOCATED-BY-RIR
mnt-by:            RIPE-NCC-HM-MNT
mnt-lower:         BRANCH-OFFICE-MNT
notify:            noc@example.net
created:           2017-09-01T08:23:35Z
last-modified:     2017-09-01T08:23:35Z
source:            RIPE
```

Configuring interfaces and loopbacks

Addressing with IPv6

Where **X** is your number on the handout!

- Your allocation is 2001:ff**XX**::/32
- The point-to-point interface towards your transit has IPv6 address 2002:1**XX**::a/127
- The interface towards the IXP has address 2001:db8::**XX**/64

Network Diagram

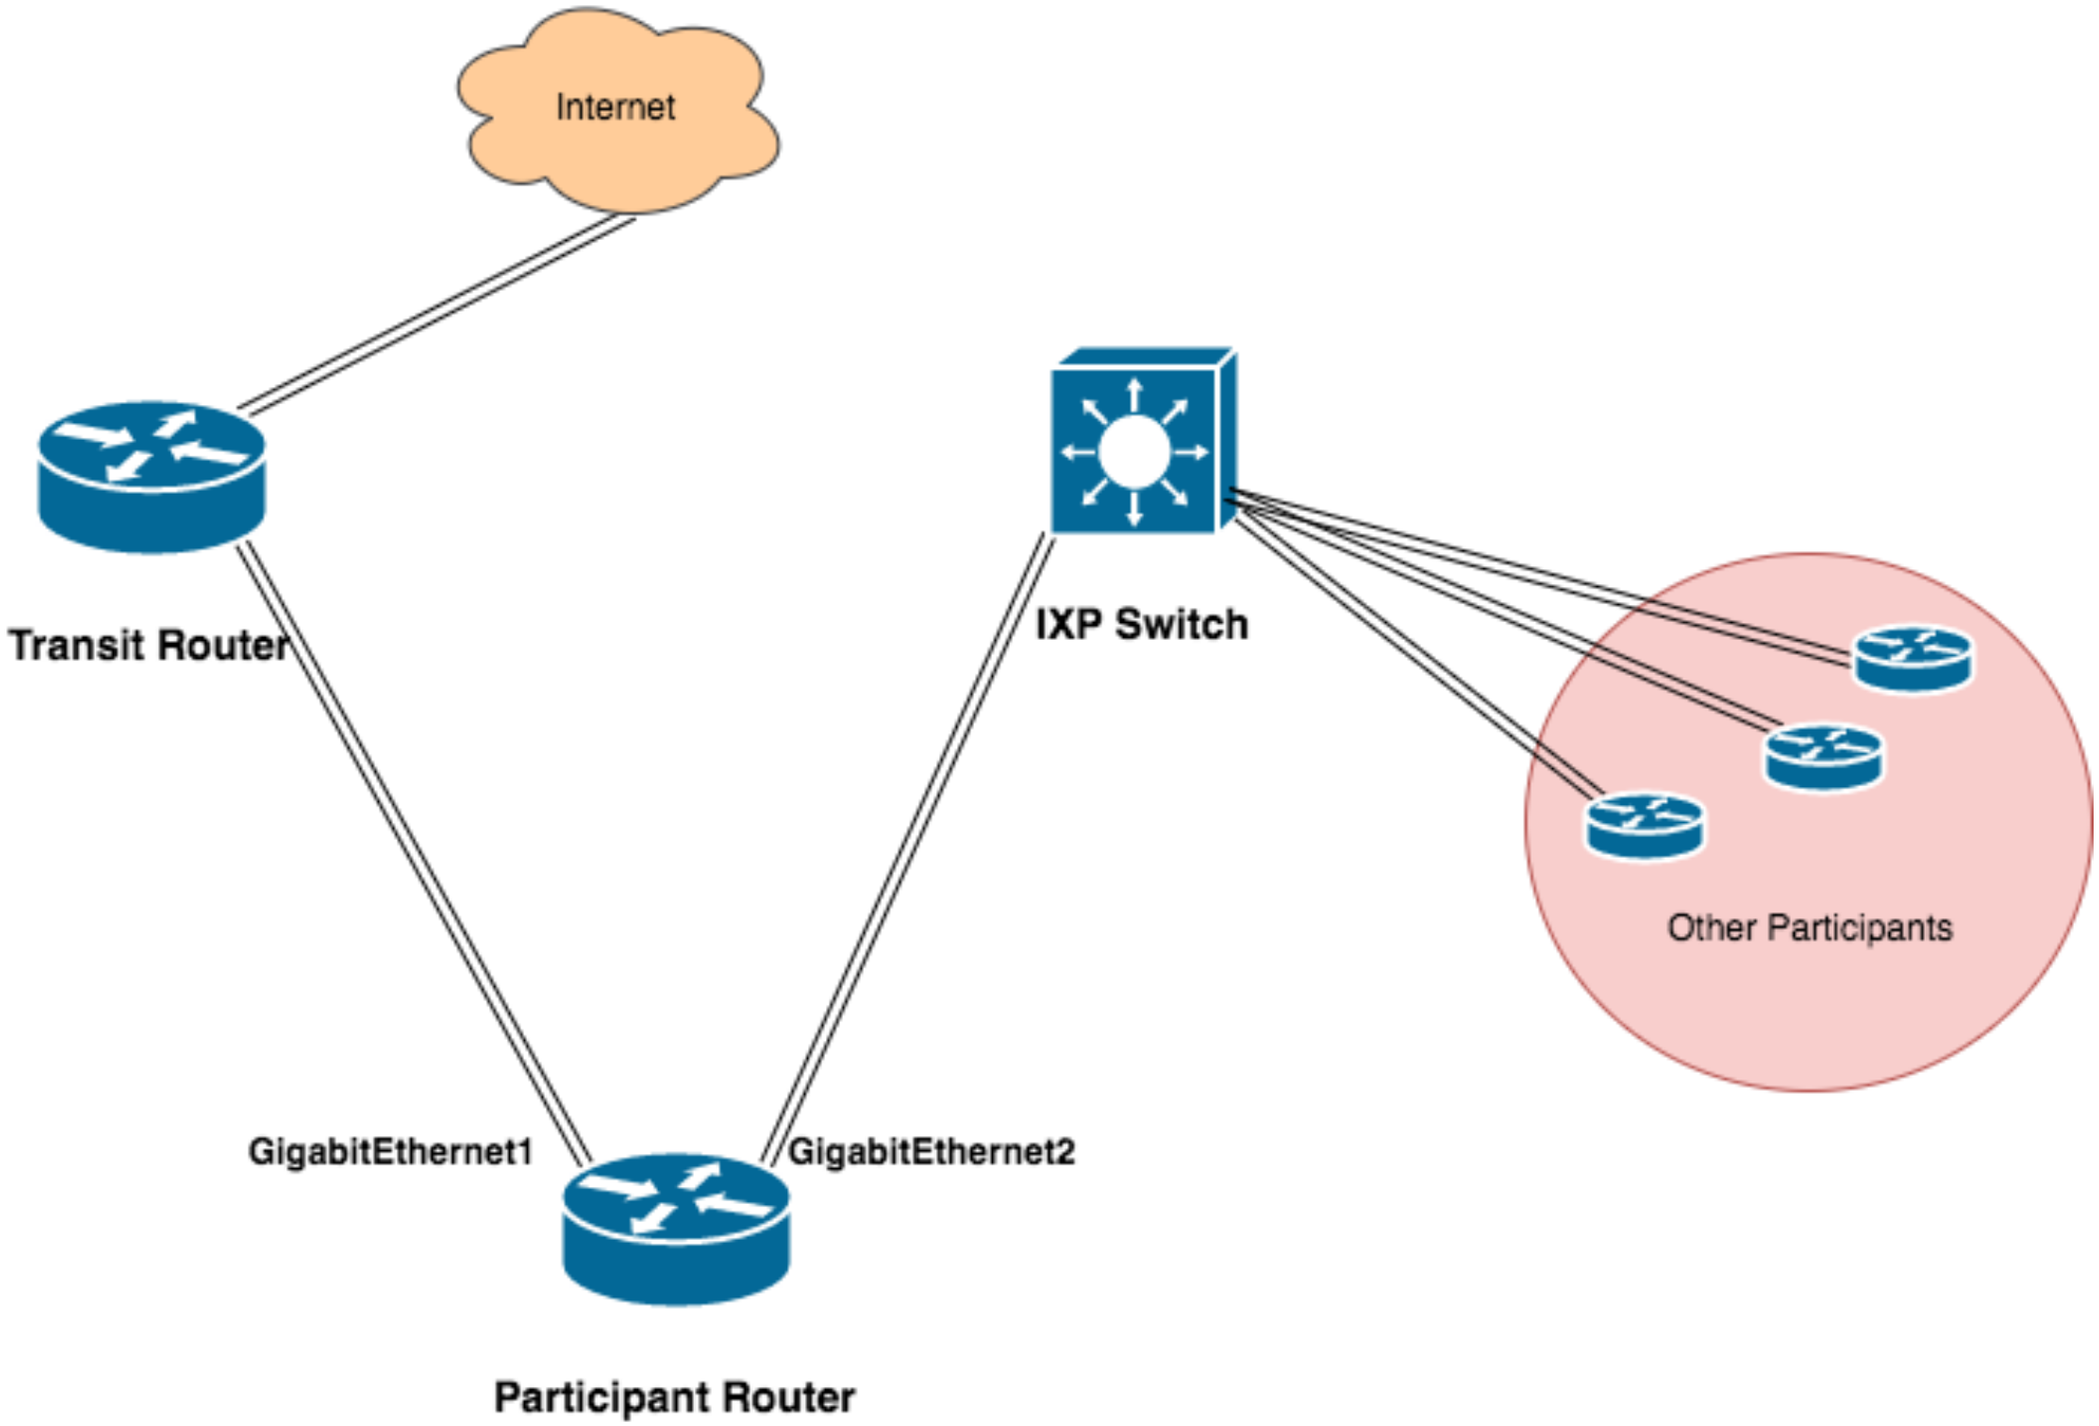

Discover the Network

- Make sure you have connectivity
- Go to: http://workbench.ripe.net
- Your login and password are provided by the trainers

- **Carefully read the instructions on the screen**

Interface IPv6 Settings Routers

- Disable Router Advertisements
  - On point-to-point interfaces
  - On LANs where unprepared devices are connected
- Otherwise they will suddenly be globally reachable over IPv6 without being configured, prepared and/or protected

Basic IPv6 Settings

- Before configuring IPv6 on your router interfaces, the basic IPv6 settings need to be set up on the router
- For R1:

```
(config)#ipv6 unicast-routing
```

Interface IPv6 Settings (on R1)

- Configure IPv6 on your IXP interface

```
(config)# interface Gi2
(config-if)# ipv6 address 2001:db8::XX/64
(config-if)# no ipv6 redirects
(config-if)# ipv6 nd ra suppress all
```

- Configure IPv6 on your Transit interface

```
(config)# interface Gi1
(config-if)# ipv6 address 2002:1XX::a/127
(config-if)# no ipv6 redirects
(config-if)# ipv6 nd ra suppress all
```

Verify

- Check everything is working by pinging the other end of the point to point link for the transit

```
# ping 2002:1XX::b
```

- And by pinging another host on the IXP Lan

```
# ping 2001:db8::3333
```

Creating a route6 Object

Assignment

- Create a **route6** object for your announcement

- Data needed
  - Your AS number
  - Your IPv6 allocation

What Are route(6) Objects?

- **route(6)** objects register which IPv4/IPv6 prefix will be announced by which AS number
- Used for creating BGP filters

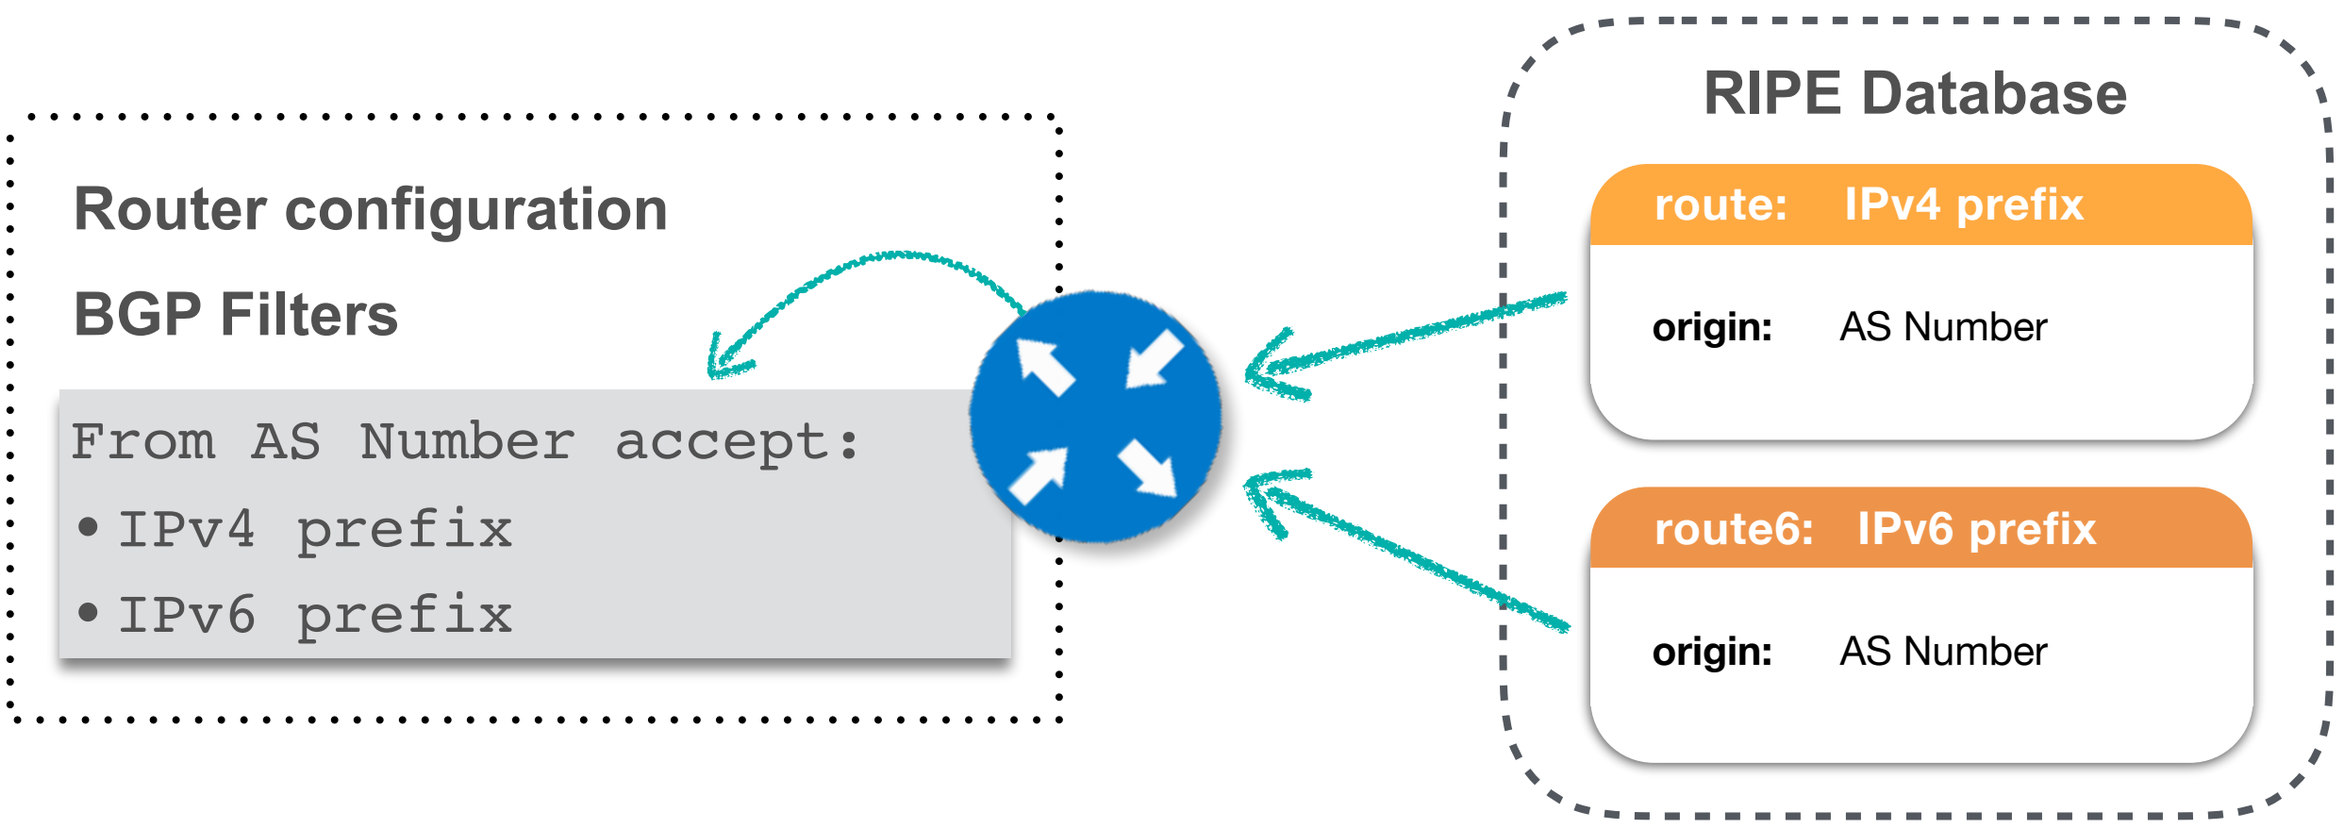

Registering IPv6 Routes

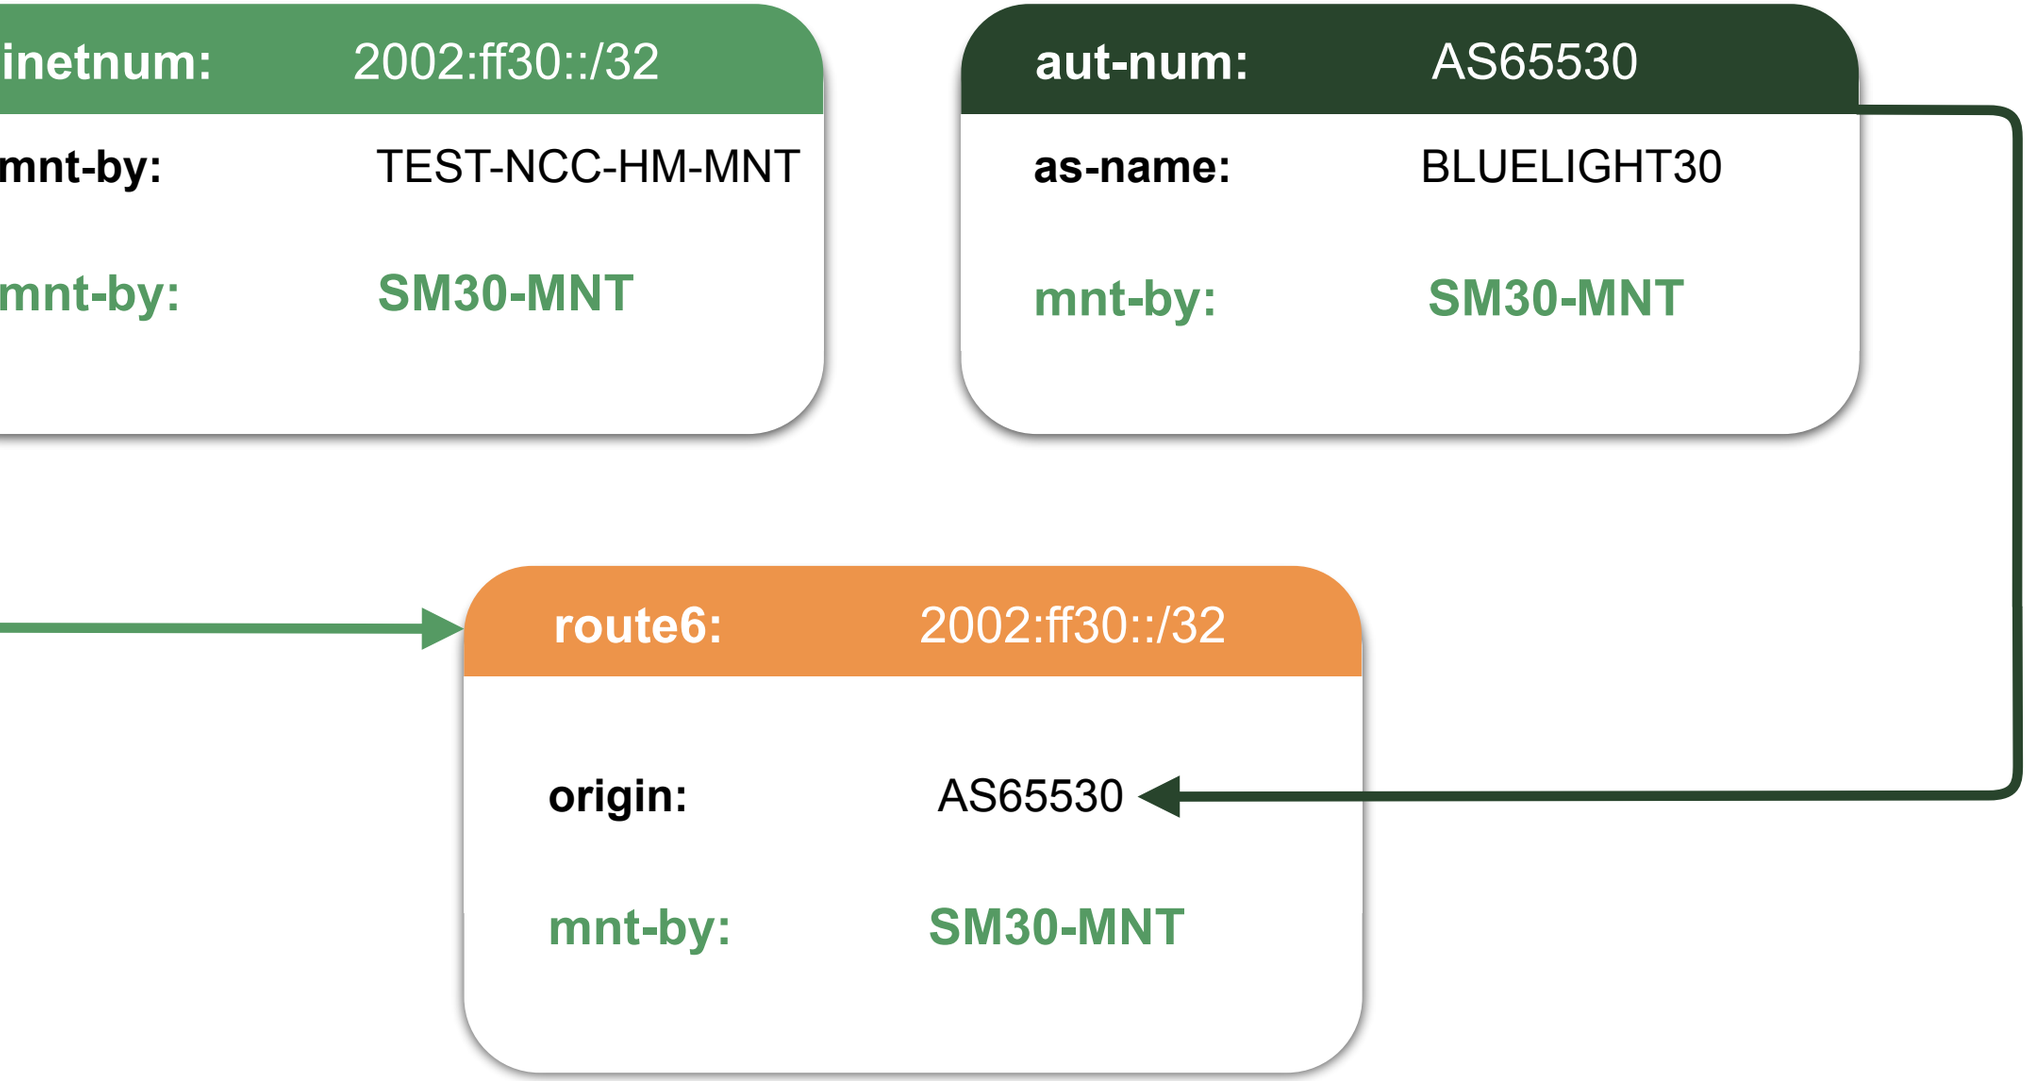

Preparation

- Create RIPE Access account (access.ripe.net)
- Using your number on the handout, identify your IPv6 allocations in RIPE TEST Database
- Find out your AS Number using the same method
- Find out the name and password of your maintainer object

Create route(6) Objects

- Open a browser and go to:
  - http://apps-test.db.ripe.net
- On the left side, click on “Create an object”
- Choose “**route6**” and click on [Create]

- Fill in the template:
  - Maintainer: CM**XX**-MNT
  - route6: 2001:ff**XX**::/32
  - origin: AS1**XX**
  - **The password will be provided by the trainers**

MP-BGP

BGP Overview

- Routing Protocol used to exchange routing information between networks
  - Exterior Gateway Protocol
- It is based on Path Vector Protocol
  - Similar to Distance Vector
- Each border router sends to its neigbors the full route to one destination, not just the distance

Autonomous Systems

- Collection of networks with the same routing policy
- Usually under single ownership and administrative control
  - Single routing policy
- Identified by 16 or 32 bit numbers
  - 16bit: 0 - 65,535
  - 32bit: 65,536 - 4,294,967,295

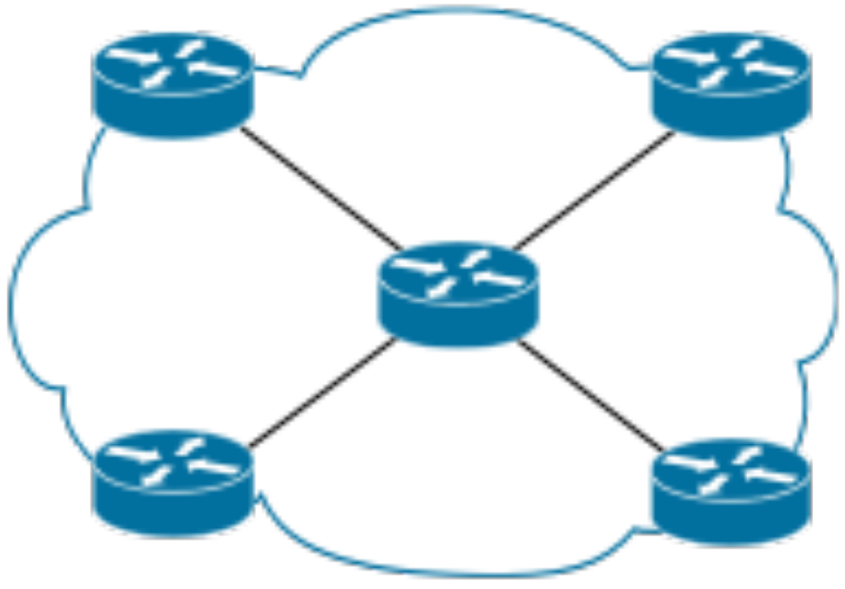

AS Path

- Sequence of ASes a route has traversed
  - Loop detection
  - Path selection (AS-PATH length)

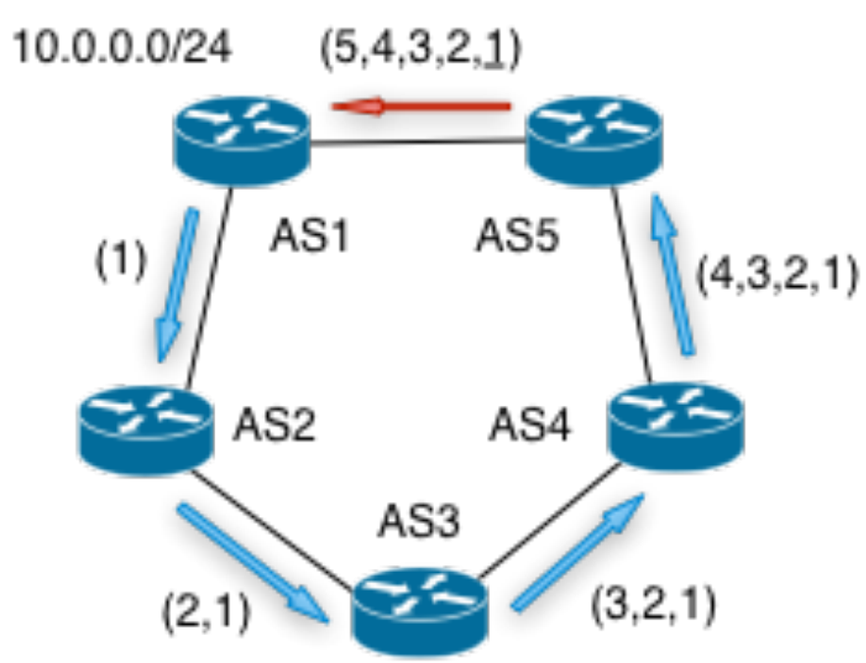

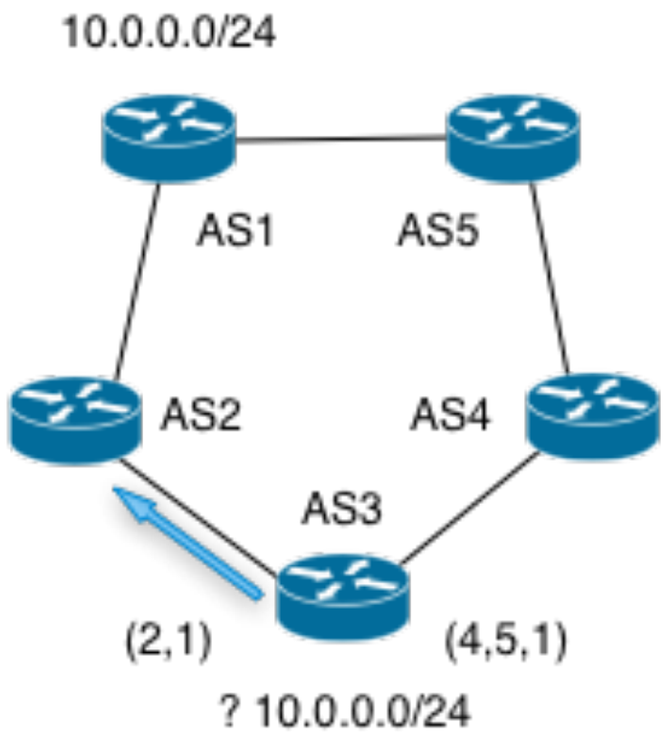

BGP Modes

- **eBGP**: Between BGP speakers in a different AS

- **iBGP**: Between BGP speakers within the same AS

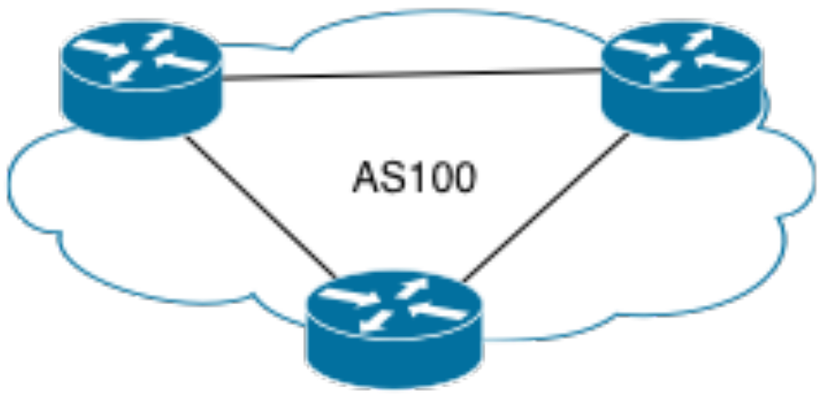

BGP Messages

- OPEN
  - opens the tcp session
- KEEPALIVE
  - keeps the session running
- NOTIFICATION
  - error handling
- UPDATE
  - actual route updates (NLRI, AS-path, AS-path attributes)

Multiprotocol BGP (MP-BGP)

- Extension to the BGP protocol
- Carries routing information about other protocols:
  - Multicast
  - MPLS VPN
  - IPv6
- Multi-Protocol NLRI exchange is negotiated at session set up (OPEN Message)

MP-BGP

- New features in OPEN Message:
  - **BGP Capabilities Advertisement:**
  - Address Family Identifier (**AFI**)
  - Subsequent Address Family Identifier (**SAFI**)
  - Multiprotocol Reachable Network Layer Reachability Information

AFI / SAFI

- Address Family Identifier (AFI)
  - Identifies Address Type
  - AFI = 1 (IPv4)
  - AFI = 2 (IPv6)
- Subsequent Address Family Identifier (SAFI)
  - Sub category for AFI Field
  - Address Family Identifier (AFI)
    - Sub-AFI = 1 (NLRI is used for unicast)
    - Sub-AFI = 2 (NLRI is used for multicast RPF check)
    - Sub-AFI = 3 (NLRI is used for both unicast and multicast RPF check)
    - Sub-AFI = 4 (label)
    - Sub-AFI = 128 (VPN)

MP-BGP Capabilities Negotiation

- BGP routers establish peering sessions through the OPEN message
- OPEN message contains optional parameters
- BGP session is terminated if OPEN parameters are not recognised
- A new optional parameter: CAPABILITIES containing its capabilities:
  - Multiprotocol extension (AFI/SAFI)
  - Route Refresh
  - Outbound Route Filtering

Managing Multiple Protocols

- Independent operation
  - One RIB per protocol
  - Distinct policies per protocol (IP address specific route maps and prefix lists must be adjusted)
  - Make separate route maps for IPv4 and IPv6
  - Prefix lists are always separate
  - It is common to use a **-v4** and a **-v6** suffix to names

Setting up your Transit

Assignment

- Enable Multiprotocol BGP
- Using IPv6
  - Connect your network to the Transit Provider
- Data needed
  - Your AS number
  - Your IPv6 address space
  - The AS number of your neighbors
  - The IPv6 address of your neighbors BGP routers

Preparation

- Insert static Null route
  - Before BGP advertised its network, it checks for an exact match of network number and mask on router's routing table

```
(config)# ipv6 route 2001:ffXX::/32 null0 250
```

Enable Multiprotocol BGP

- Enable MP-BGP

```
(config)# router bgp 1XX
(config-router)# no bgp default ipv4-unicast
(config-router)# bgp router-id 1XX.0.0.0
```

Create a filter

- BGP sends the best paths to all neighbours

```
(config)# ipv6 prefix-list TRANS-OUT-V6 seq 5 permit 2001:ffXX::/32
(config)# ipv6 prefix-list IXP-OUT-V6 seq 5 permit 2001:ffXX::/32
```

Configure Transit Session

- Configure BGP session with AS66

```
(config)# router bgp 1XX
(config-router)# neighbor 2002:1XX::b remote-as 66
(config-router)# address-family ipv6
(config-router-af)# neighbor 2002:1XX::b activate
(config-router-af)# neighbor 2002:1XX::b prefix-list
TRANS-OUT-V6 out
```

- Advertise route

```
(config-router-af)# network 2001:ffXX::/32
```

Verify

- Check session summary

```
# show bgp ipv6 unicast summary
```

- Check BGP and routing table

```
# show bgp ipv6
# show ipv6 route
```

- Verify reachability

```
# ping 2001:3c00::abcd
# ping <your colleague IPv6>
```

- Show logged events

```
# show logging
```

Questions

Setting up your IXP peering

Internet Exchanges (IX or IXP)

- A switch (or set of switches) that allows members to exchange traffic **directly**
  - Meeting point through BGP peering

- Many countries have at least one
  - AMS-IX, LINX, VIX, MIX, etc

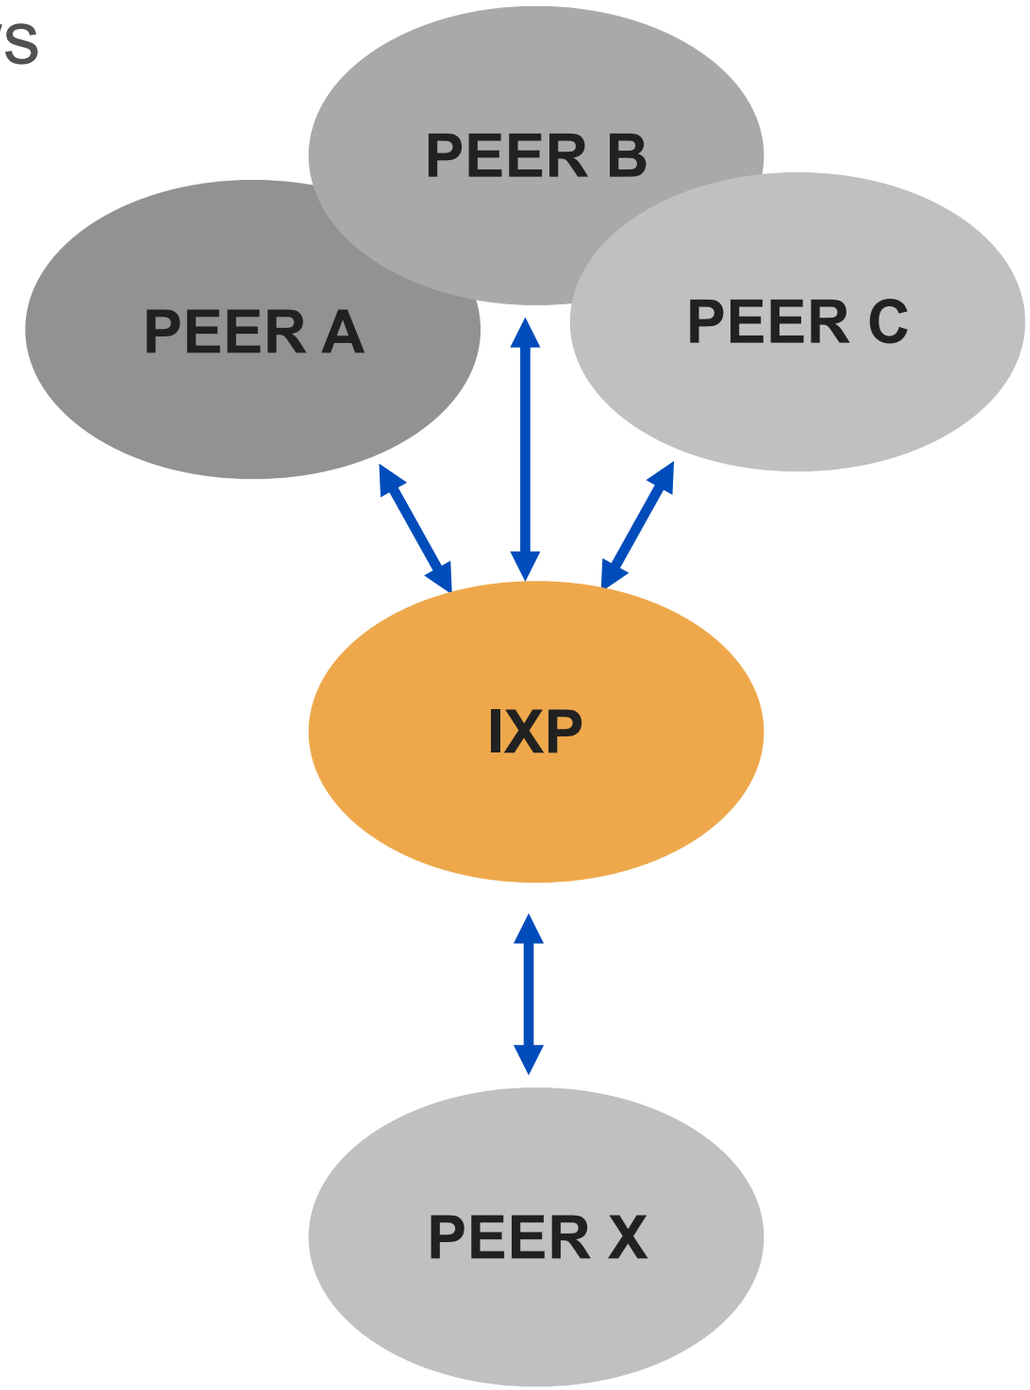

IXPs - Why

- IXPs enable traffic to remain local
  - Improves routing efficiency and fault-tolerance
  - Reduces the average per-bit delivery cost (no transit)

IXPs - Who

- Often non-profit, membership organisations
  - Cater to the local ISPs, Content Providers, Universities, Governments and others

IXPs - Architecture

- A switch, or a group of switches
  - Range is generally from 100Mb to 100Gb ports
- Switches are in colocation facilities
  - Easy to reach them
  - Can be spread in different facilities across a city or region
- Some IXes have two LANs for redundancy

IXPs - Route Servers

- A server running a BGP Daemon

- Helps networks who peer at many IXPs
  - Avoids setting up a meshed environment
  - Eases management

IXPs - Route Servers

- Sets **next-hop** as announcer, leaving itself out

- Traffic does not flow through the route server

Configure IXP Sessions

- Configure BGP sessions with AS127

```
(config)# ipv6 prefix-list AS127-IN-V6 permit 2001::ff27::/32
(config)# router bgp 1XX
(config-router)# neighbor 2001:db8::27 remote-as 127
(config-router)# address-family ipv6
(config-router-af)# neighbor 2001:db8::27 activate
(config-router-af)# neighbor 2001:db8::27 prefix-list IXP-OUT-
V6 out
(config-router-af)# neighbor 2001:db8::27 prefix-list AS127-
IN-v6 in
(config-router)# exit
```

Setting up your IXP Peers

BGP Peering at IXP

- Following the example BGP Session with AS127, set up BGP peerings with other participants in the room
- Remember:
  - Set up a prefix list first
  - Use the correct information (IP Address, AS Number)
  - Check for the BGP Session to come up

```
# sh bgp ipv6 unicast summary
```

Filtering best practises

Filtering Principles

- Filter as close to the edge as possible
- Filter as precisely as possible
- Filter both source and destination where possible

- Two filtering techniques:
  - Explicit Permit (permit then deny any)
  - Explicit Deny (deny then permit any)

Bogons

- Routes you shouldn't see in the routing table
  - Private addresses
  - Non-allocated space
  - Reserved space (Future use, Multicast, etc.)
- You should have filters applied so that these routes are not advertised to or propagated through the Internet
- Team Cymru provides list or BGP feed
  - http://www.team-cymru.org/bogon-reference-bgp.html

Prefix-lists

- Prefix lists are lists of routes you want to accept or announce
- Easy to use but not highly scalable
- You can create them manually or automatically
  - With data from RIPE DB or other Internet Routing Registry
- Or using a tool
  - Level3 Filtergen
  - bgpq3
  - IRRexplorer

Reverse Path Forwarding

- Called uRPF (Unicast Reverse Path Forwarding)
- Checks if an entry exists in the routing table before accepting the packet and forwarding it

- Two modes
  - Loose
  - Strict

Strict and Loose RPF

- Strict
  - Checks if the entry is in the routing table
  - and the route points to the receiving interface

- Loose
  - Simply checks that an entry exists for the route in the routing table

Best Current Practice 38

- Defines some steps to take in order to have a “cleaner” routing table
- Restricting forged traffic (TCP and UDP)
- Implies the use of:
  - Prefix filters
  - Bogon filters
  - uRPF
- http://tools.ietf.org/html/bcp38

Ingress filters

- Best Practices:
  - Don’t accept BOGON ASNs
  - Don’t accept BOGON prefixes
  - Don’t accept your own prefix
  - Don’t accept default (unless you requested it)
  - Don’t accept IPv4 prefixes longer than /24
  - Don’t accept IPv6 prefixes longer than /48
  - Consider Net Police Filtering

Preparation

- Examine your routing table

```
# show ipv6 route bgp
# show bgp ipv6
```

- Do you see anything strange?

Filter More Specifics (on R1)

- Filtering of the prefixes that are too specific

```
(config)# ipv6 prefix-list TRANS-IN-V6 seq 10 permit 2000::/3 le 48
(config)# ipv6 prefix-list IXP-IN-V6 seq 10 permit 2000::/3 le 48
```

Filter More Specifics

- Add incoming policy to the transit

```
(config)# router bgp 1XX
(config-router-af)# address-family ipv6
(config-router-af)# neighbor 2002:1XX::b prefix-list TRANS-
IN-V6 in
```

- For the BGP Peerings, we have already set up strict filtering

Clear the BGP Session

```
# clear bgp ipv6 unicast 2002:1XX::b soft in
```

Verify

- Check BGP and routing table

```
# show bgp ipv6 unicast
# show ipv6 route | include /64
# show ipv6 route | include /56
# show ipv6 route | include /49
# show ipv6 route | include /52
```

- None of these should appear

Questions

nathalie@ripe.net
mstucchi@ripe.net