Prefixes and Autonomous Systems

BGP for networks who peer: Part 1

Wolfgang Tremmel
wolfgang.tremmel@de-cix.net

BGP Webinars Overview

- 01 - Prefixes and AS numbers
- 02 - BGP Introduction
- 03a - Setting up iBGP
- 03b - Setting up eBGP
- 04 - BGP best path selection
- ...

BGP Webinars Overview

- 01 - Prefixes and AS numbers
- 02 - BGP Introduction
- 03a - Setting up iBGP
- 03b - Setting up eBGP
- 04 - BGP best path selection
- ...

IP Prefixes

IPv4 Addresses

10.3.8.17

IPv4 Addresses

10.3.8.17

IPv4 Prefixes

10.3.8.0/22

IPv4 Prefixes

10.3.8.0/22

```
1  2  3  4    5  6  7  8    9  10 11 12   13 14 15 16   17 18 19 20   21 22 23 24   25 25 27 28   29 30 31 32
0000 1010 0000 0011 0000 1000 0000 0000
```

➔ IPv4 and IPv6 addresses have a network and a host part

IPv4 Prefixes

10.3.8.0/22

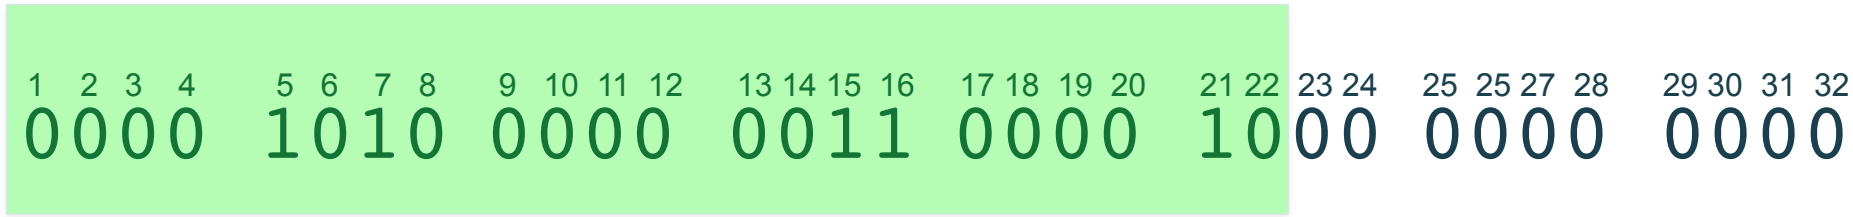

- IPv4 and IPv6 addresses have a network and a host part
- A **prefix** is just the network part + the **length** of the network part

IPv4 Prefixes

10.3.8.0/22

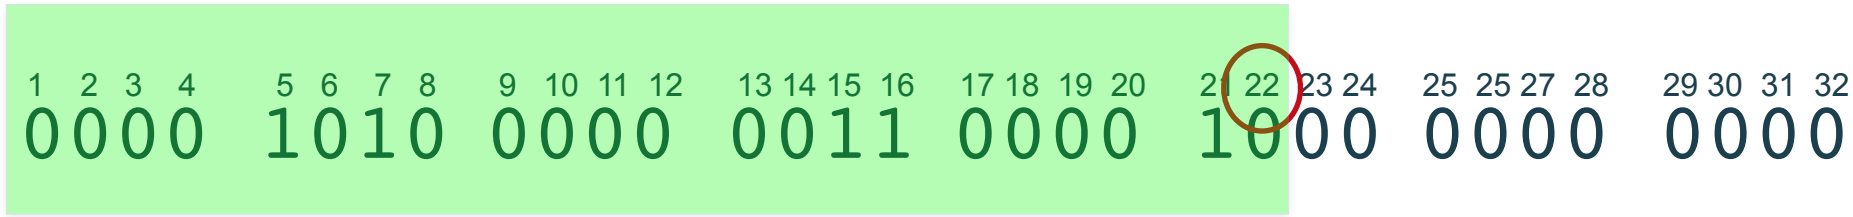

- IPv4 and IPv6 addresses have a network and a host part
- A **prefix** is just the network part + the **length** of the network part
- Important:
  - **The boundary between network and host can be anywhere!**

DE CIX

Characteristics of Prefixes: IPv4

10.3.8.0/22

Prefix-Length: 0-32

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 25 27 28 29 30 31 32

0000 1010 0000 0011 0000 1000 0000 0000

Notation:
- 4 Numbers 0-255
- Separated by "."
- a "/", followed by

Host-part all zero

32 Bits long

Characteristics of Prefixes: IPv4

10.3.8.0/22

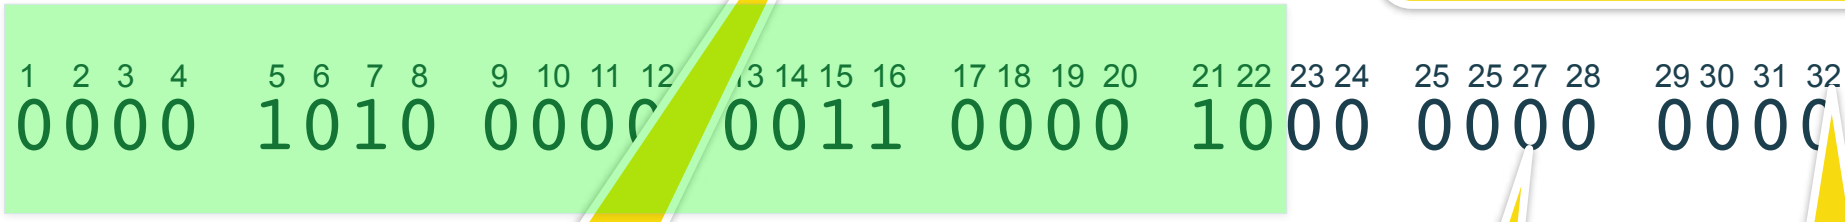

Prefix-Length: 0-32

Notation:
- 4 Numbers 0-255
- Separated by "."
- a "/", followed by

Host-part all zero

32 Bits long

IPv6 Addresses

2003:de:274f:400:226:b0ff:fed8:3d8a

IPv6 Addresses

2003:de:274f:400:226:b0ff:fed8:3d8a

IPv6 Prefixes

2003:de:274f:400::/64

Characteristics of Prefixes: IPv6

Prefix-Length: 0-128

2003:de:274f:400::/64

Notation:
- 4 digit hex numbers (0-9,a-f)
- Separated by ":"
- "::" = fill up with zeros

Host-part all zero

128 Bits long

IP Adresses and Prefixes

	IPv4	IPv6
Length	32 Bit	128 Bit
	0-32 Prefix Length	0-128 Prefix Length
Notation	4 Numbers, 0-255	8 Numbers, 0-fffff
Separator	.	:
Prefix: Host part	all zero	
Address: Host part	**not** all zero / **not** all one	
Example (Prefix)	198.51.100.0/24	2001:db8:4f30::/48

IP Adresses and Prefixes

Prefix or Not?

	IPv4	IPv6
Length	32 Bit	128 Bit
	0-32 Prefix Length	0-128 Prefix Length
Notation	4 Numbers, 0-255	8 Numbers, 0-fffff
Separator	.	:
Prefix: Host part	all zero	
Address: Host part	**not** all zero / **not** all one	
Example (Prefix)	198.51.100.0/24	2001:db8:4f30::/48

198.51.100.0/24

198.51.100.0/24

2001:db8:5669::/48

2001:db8:5669::/48

192.0.2.8/27

192.0.2.8/27

0000 1000

Host part not zero!

203.0.113.99/32

203.0.113.99/32

198.51.100.0/16

198.51.100.0/16

Host part not zero!

What is an Autonomous System?

And why do I need one?

Wolfgang Tremmel
academy@de-cix.net

A brief history of the Internet

According to the Internet Hall of Fame

➔ 1982 – Arpanet (successor of Internet)

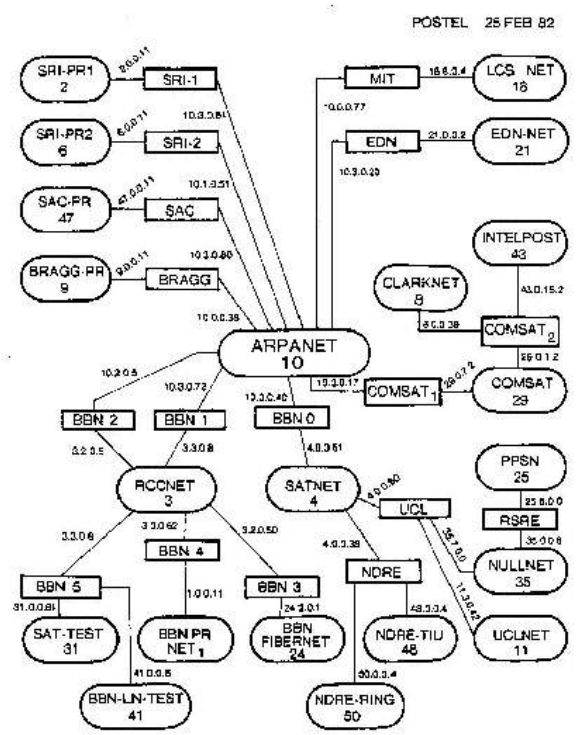

A brief history of the Internet

According to the Internet Hall of Fame

- 1982 – Arpanet (successor of Internet)
- 1982: RFC827 defines Exterior Gateway Protocol:

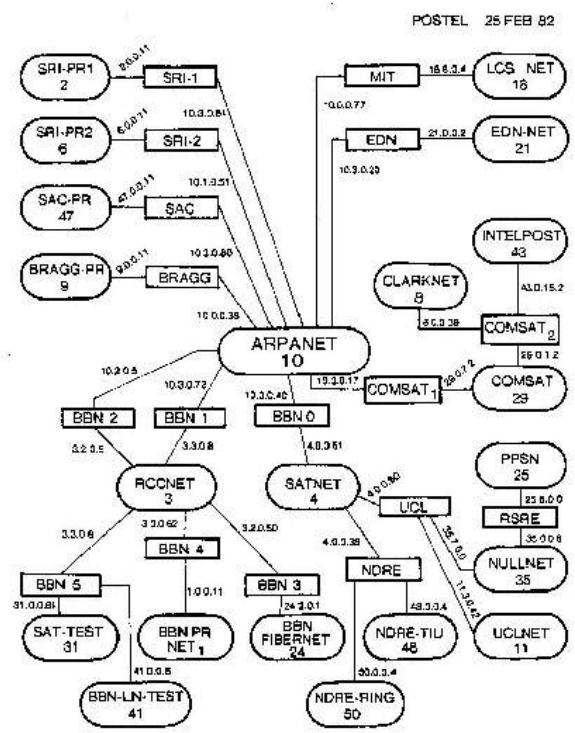

A brief history of the Internet

According to the Internet Hall of Fame

- 1982 – Arpanet (successor of Internet)
- 1982: RFC827 defines Exterior Gateway Protocol:

"Autonomous systems will be assigned 16-bit identification numbers (in much the same ways as network and protocol numbers are now assigned)"

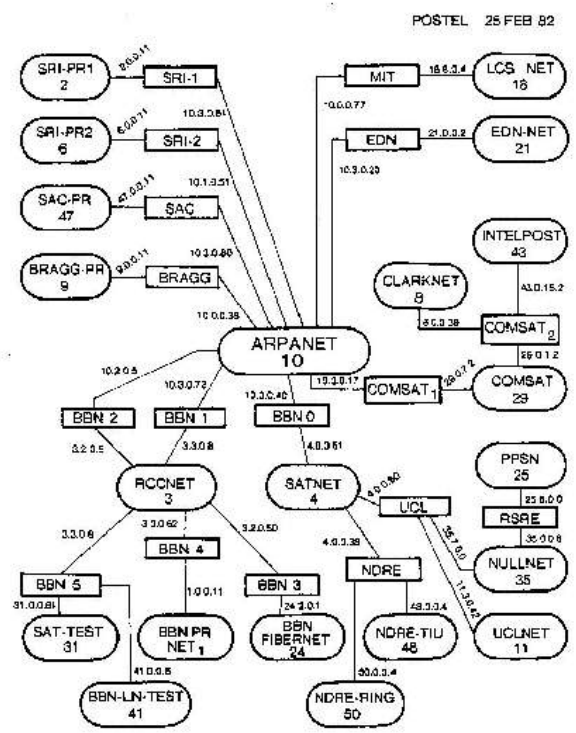

But what is an Autonomous System?

But what is an Autonomous System?

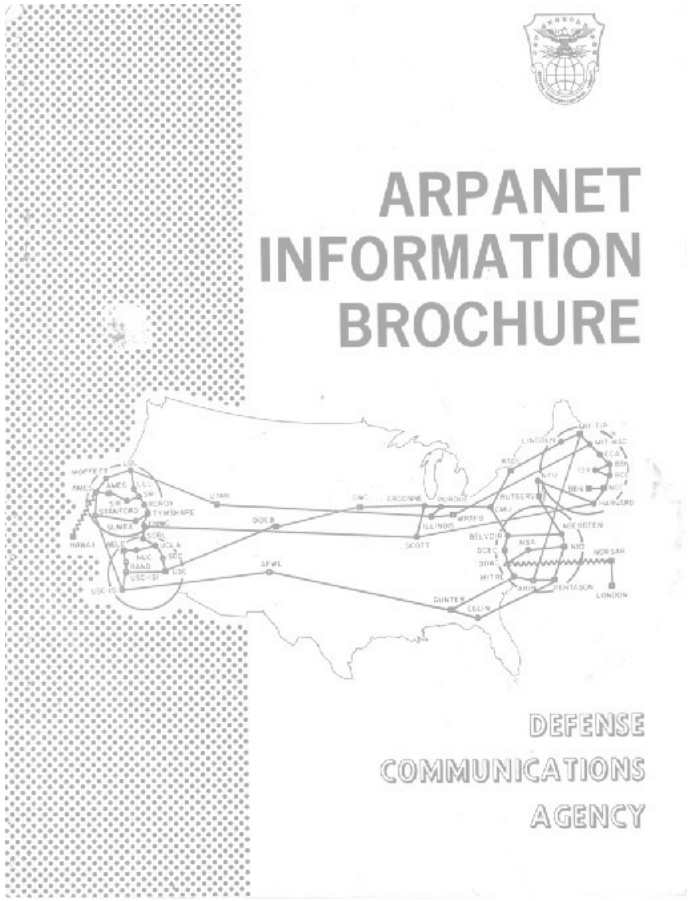

But what is an Autonomous System?

"An AS is a connected group of one or more IP prefixes run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy."

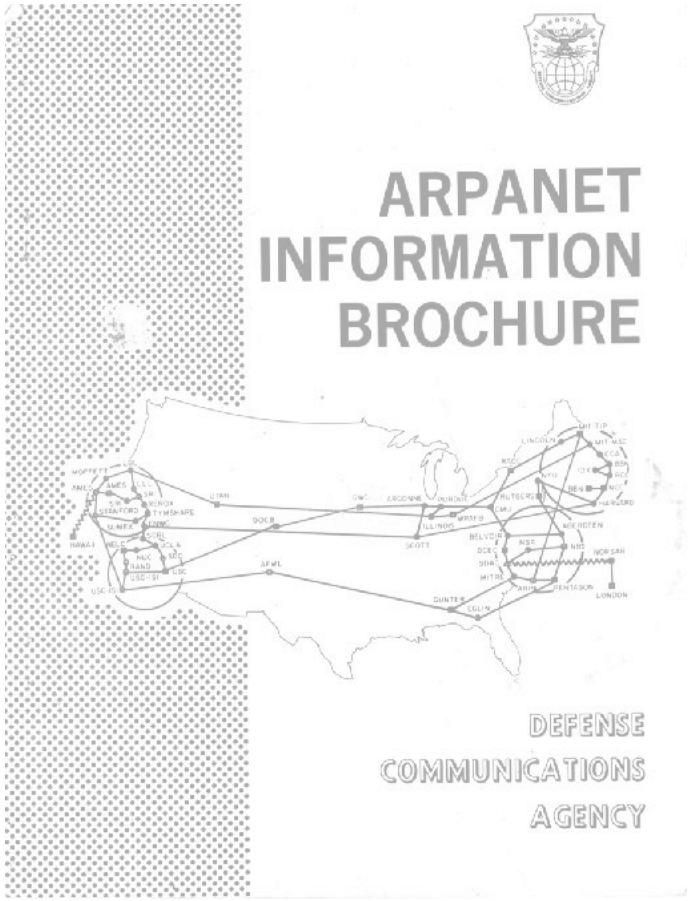

➔1996 – Defined in RFC1930 (earlier definitions exist)

➔Other routing protocols call these "areas" for example

➔The inside of an AS is "invisible" to the outside

But what is an Autonomous System?

"An AS is a **connected** group of one or more IP prefixes run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy."

But what is an Autonomous System?

"An AS is a **connected** group of one or more IP prefixes run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy."

→"connected": An autonomous system is continuous.
One AS cannot be in two different places.

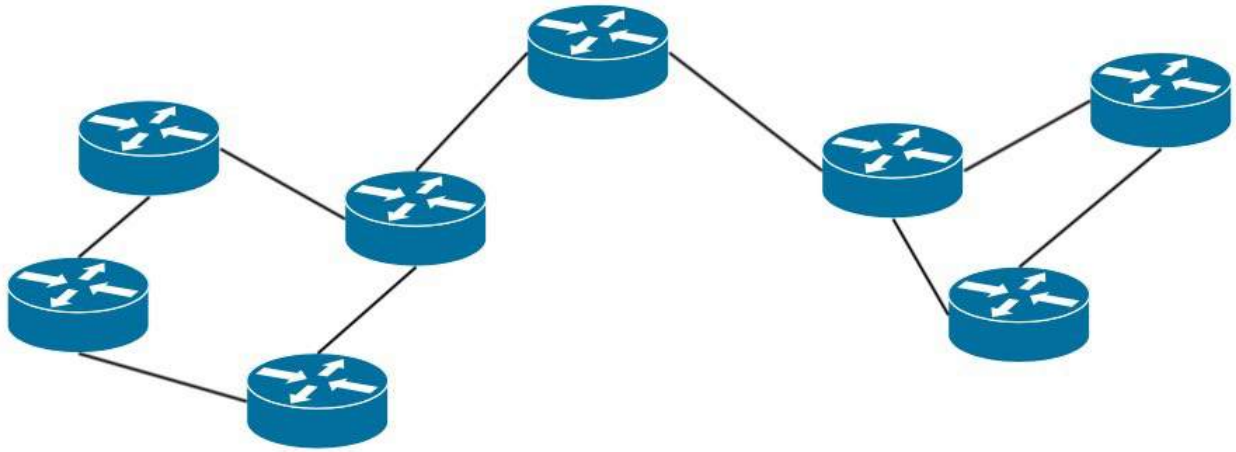

But what is an Autonomous System?

"An AS is a connected **group of one or more IP prefixes** run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy."

But what is an Autonomous System?

"An AS is a connected **group of one or more IP prefixes** run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy."

→"group of IP prefixes": This is about how IP prefixes are routed, not about devices. Routers are not even mentioned.

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	212.114.64.0/19	80.81.192.40	50	100		8859 i
*>	194.77.145.0/24	80.81.192.40	50	100		8859 i
*>	194.45.27.0/24	80.81.192.40	50	100		8859 i
*>	193.17.21.0/24	80.81.192.40	50	100		8859 i
*>	213.241.128.0/18	80.81.192.40	50	100		8859 i

But what is an Autonomous System?

"An AS is a connected group of one or more IP prefixes **run by one or more network operators** which has a SINGLE and CLEARLY DEFINED routing policy."

But what is an Autonomous System?

"An AS is a connected group of one or more IP prefixes **run by one or more network operators** which has a SINGLE and CLEARLY DEFINED routing policy."

→"run by one or more network operators": An AS does not have to be run by only one operator - no business model is enforced.

```
aut-num:        AS6695
as-name:        DECIX-AS
descr:          DE-CIX Management GmbH
descr:          DE-CIX, the German Internet Exchange
descr:          DE
org:            ORG-DtGI1-RIPE
status:         ASSIGNED
mnt-by:         RIPE-NCC-END-MNT
admin-c:        DXSU6695-RIPE
tech-c:         DXSU6695-RIPE
tech-c:         BH6695-RIPE
mnt-by:         DECIX-MNT
mnt-lower:      DECIX-MNT
```

But what is an Autonomous System?

"An AS is a connected group of one or more IP prefixes run by one or more network operators which **has a SINGLE and CLEARLY DEFINED routing policy**."

➔"has a SINGLE and CLEARLY DEFINED routing policy":

➔"routing policy": This is how routing decisions are made.

➔These policies are not defined for each single prefix, but for groups of prefixes.

➔Each AS defines this routing policy on its own.

So this is an Autonomous System!

"An AS is a connected group of one or more IP prefixes run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy."

- **So now you know:**
  - You do not need a router
  - However, you need prefixes to be routed
- **Most commonly:**
  - you do have a router
  - … or more than one
  - and it "belongs" to an AS if it runs BGP

What is an Autonomous System good for?

	If you have an AS	Without an AS
Redundancy	You can have multiple upstream ISPs and Peering	You only can have one upstream ISP
Control	You have full control over your outgoing traffic	Your upstream ISP controls your traffic
Cost	You can optimize your traffic for cost	You just pay your upstream ISP
Peering	You can setup your own peering policy and have full control	Your upstream ISP makes all decisions

Sounds good, I want an AS, where can I get one?

Sounds good, I want an AS, where can I get one?

→AS numbers are globally unique

Sounds good, I want an AS, where can I get one?

- →AS numbers are globally unique
- →So some sort of authority must exist for handing them out

Sounds good, I want an AS, where can I get one?

- →AS numbers are globally unique
- →So some sort of authority must exist for handing them out
- →This authority is IANA - the Internet Assigned Numbers Authority

DE CIX

Sounds good, I want an AS, where can I get one?

- →AS numbers are globally unique
- →So some sort of authority must exist for handing them out
- →This authority is IANA - the Internet Assigned Numbers Authority
- →But no, you cannot go to IANA and just ask for an AS – they delegated the task

Sounds good, I want an AS, where can I get one?

- →AS numbers are globally unique
- →So some sort of authority must exist for handing them out
- →This authority is IANA - the Internet Assigned Numbers Authority
- →But no, you cannot go to IANA and just ask for an AS – they delegated the task
  - →to five Regional Internet Registries (RIRs)

Sounds good, I want an AS, where can I get one?

- →AS numbers are globally unique
- →So some sort of authority must exist for handing them out
- →This authority is IANA - the Internet Assigned Numbers Authority
- →But no, you cannot go to IANA and just ask for an AS – they delegated the task
  - →to five Regional Internet Registries (RIRs)
  - →have a look at the map to see who is responsible for your region

Regional Internet Registries (RIRs)

- ➔Talking about everything what RIRs do would be beyond the scope of this training

Regional Internet Registries (RIRs)

- ➔Talking about everything what RIRs do would be beyond the scope of this training
- ➔So, let's focus on AS numbers

Regional Internet Registries (RIRs)

- ➔Talking about everything what RIRs do would be beyond the scope of this training
- ➔So, let's focus on AS numbers
- ➔And for now, let's focus on Europe

Regional Internet Registries (RIRs)

- ➔Talking about everything what RIRs do would be beyond the scope of this training
- ➔So, let's focus on AS numbers
- ➔And for now, let's focus on Europe
- ➔The RIR responsible for Europe, Russia and the Middle East is the RIPE NCC

Regional Internet Registries (RIRs)

- ➔Talking about everything what RIRs do would be beyond the scope of this training
- ➔So, let's focus on AS numbers
- ➔And for now, let's focus on Europe
- ➔The RIR responsible for Europe, Russia and the Middle East is the RIPE NCC
- ➔RIPE means Réseaux IP Européens – the founders wanted a French name

Regional Internet Registries (RIRs)

- → Talking about everything what RIRs do would be beyond the scope of this training
- → So, let's focus on AS numbers
- → And for now, let's focus on Europe
- → The RIR responsible for Europe, Russia and the Middle East is the RIPE NCC
- → RIPE means Réseaux IP Européens – the founders wanted a French name
- → NCC means Network Coordination Center

Regional Internet Registries (RIRs)

- Talking about everything what RIRs do would be beyond the scope of this training
- So, let's focus on AS numbers
- And for now, let's focus on Europe
- The RIR responsible for Europe, Russia and the Middle East is the RIPE NCC
- RIPE means Réseaux IP Européens – the founders wanted a French name
- NCC means Network Coordination Center
- RIPE is not the same as RIPE NCC, see the website for the difference.

Regional Internet Registries (RIRs)

- Talking about everything what RIRs do would be beyond the scope of this training
- So, let's focus on AS numbers
- And for now, let's focus on Europe
- The RIR responsible for Europe, Russia and the Middle East is the RIPE NCC
- RIPE means Réseaux IP Européens – the founders wanted a French name
- NCC means Network Coordination Center
- RIPE is not the same as RIPE NCC, see the website for the difference.
- Back to how to get an AS number …

Getting an AS number from RIPE NCC: The easy way

Getting an AS number from RIPE NCC: The easy way

- Just become a customer
  - You have to be a legal entity
  - Fill out the forms
  - Pay the sign-up fee (and annual fee)

Getting an AS number from RIPE NCC: The easy way

- Just become a customer
  - You have to be a legal entity
  - Fill out the forms
  - Pay the sign-up fee (and annual fee)
- Request your AS number
  - You have to be/want to be multi-homed (peering counts!)
  - RIPE Academy offers lots of online / offline trainings to help you get started.

Getting an AS number without becoming a RIPE NCC member

- ➔You can also get an AS from someone who already is a RIPE NCC member
- ➔This is called a "sponsoring LIR"
- ➔Basically they request the AS from RIPE NCC for you
- ➔… and may charge you for this

Now I have an AS – how can I route my IP prefix?

Now I have an AS – how can I route my IP prefix?

➔Hmm, this depends where you have your IP space from

Now I have an AS – how can I route my IP prefix?

- ➔Hmm, this depends where you have your IP space from
- ➔In general, IPv4 prefixes of /24 or larger are routable via BGP

Now I have an AS – how can I route my IP prefix?

- ➔Hmm, this depends where you have your IP space from
- ➔In general, IPv4 prefixes of /24 or larger are routable via BGP
- ➔In IPv6 you can route /48 or larger

Now I have an AS – how can I route my IP prefix?

- ➔Hmm, this depends where you have your IP space from
- ➔In general, IPv4 prefixes of /24 or larger are routable via BGP
- ➔In IPv6 you can route /48 or larger
- ➔If you have just become a new RIPE NCC member, you can also request IP space

Now I have an AS – how can I route my IP prefix?

- →Hmm, this depends where you have your IP space from
- →In general, IPv4 prefixes of /24 or larger are routable via BGP
- →In IPv6 you can route /48 or larger
- →If you have just become a new RIPE NCC member, you can also request IP space
  - →~~As there is not much IPv4 left, you get a /22 once (and not more)~~
  - →**IPv4 is out! No more IPv4 addresses (except by transfers)**

Now I have an AS – how can I route my IP prefix?

- →Hmm, this depends where you have your IP space from
- →In general, IPv4 prefixes of /24 or larger are routable via BGP
- →In IPv6 you can route /48 or larger
- →If you have just become a new RIPE NCC member, you can also request IP space
  - →~~As there is not much IPv4 left, you get a /22 once (and not more)~~
  - →**IPv4 is out! No more IPv4 addresses (except by transfers)**
  - →But plenty of IPv6 available...

DE CIX

Now I have an AS – how can I route my IP prefix?

- →Hmm, this depends where you have your IP space from
- →In general, IPv4 prefixes of /24 or larger are routable via BGP
- →In IPv6 you can route /48 or larger
- →If you have just become a new RIPE NCC member, you can also request IP space
  - →~~As there is not much IPv4 left, you get a /22 once (and not more)~~
  - →**IPv4 is out! No more IPv4 addresses (except by transfers)**
  - →But plenty of IPv6 available...
- →To check whether your current space is routable from your new AS, the best way is to check with whom you got that IP space from

DE-CIX Management GmbH | Lindleystr. 12 | 60314 Frankfurt | Germany
Phone + 49 69 1730 902 0 | sales@de-cix.net | www.de-cix.net

Thank you!

DE-CIX Management GmbH | Lindleystr. 12 | 60314 Frankfurt | Germany
Phone + 49 69 1730 902 0 | sales@de-cix.net | www.de-cix.net

Links and further reading

Interested in more webinars? Please subscribe to our mailing list at https://lists.de-cix.net/wws/subscribe/academy

Links visited during the webinar

- RFCs

  RFCs are Internet standards issued by the Internet Engineering Task Force (IETF)

  - RFC4632: Classless Inter-domain routing (CIDR)
  - RFC4291: IPv6 addressing architecture
  - RFC827: Exterior Gateway Architecture (EGP) (historical, obsolete)
  - RFC1930: Guidelines for creation, selection, and registration of an Autonomous System (AS)
  - RFC6793: BGP Support for Four-Octet Autonomous System (AS) Number Space

- AS Numbers
  - Giving AS numbers to the RIRs: iana.org
  - Requesting an AS number, links for:
    - ARIN
    - Lacnic
    - APNIC
    - RIPE NCC
    - Afrinic

DE CIX

Interested in more webinars? Please subscribe to our mailing list at https://lists.de-cix.net/wws/subscribe/academy