

Getting IPv6 & Securing your Routing

IPv6 Kongress
May 2012, Frankfurt

Sandra Brás & Ferenc Csorba, RIPE NCC



Schedule

- IPv4 exhaustion
- IPv6 address space
- European IPv6 deployment statistics
- BGP multihoming
- Routing Registry and the RIPE Database
- Resource Certification

RIPE / RIPE NCC

RIPE

- Open community
- Develops addressing policies
- Working group mailing lists

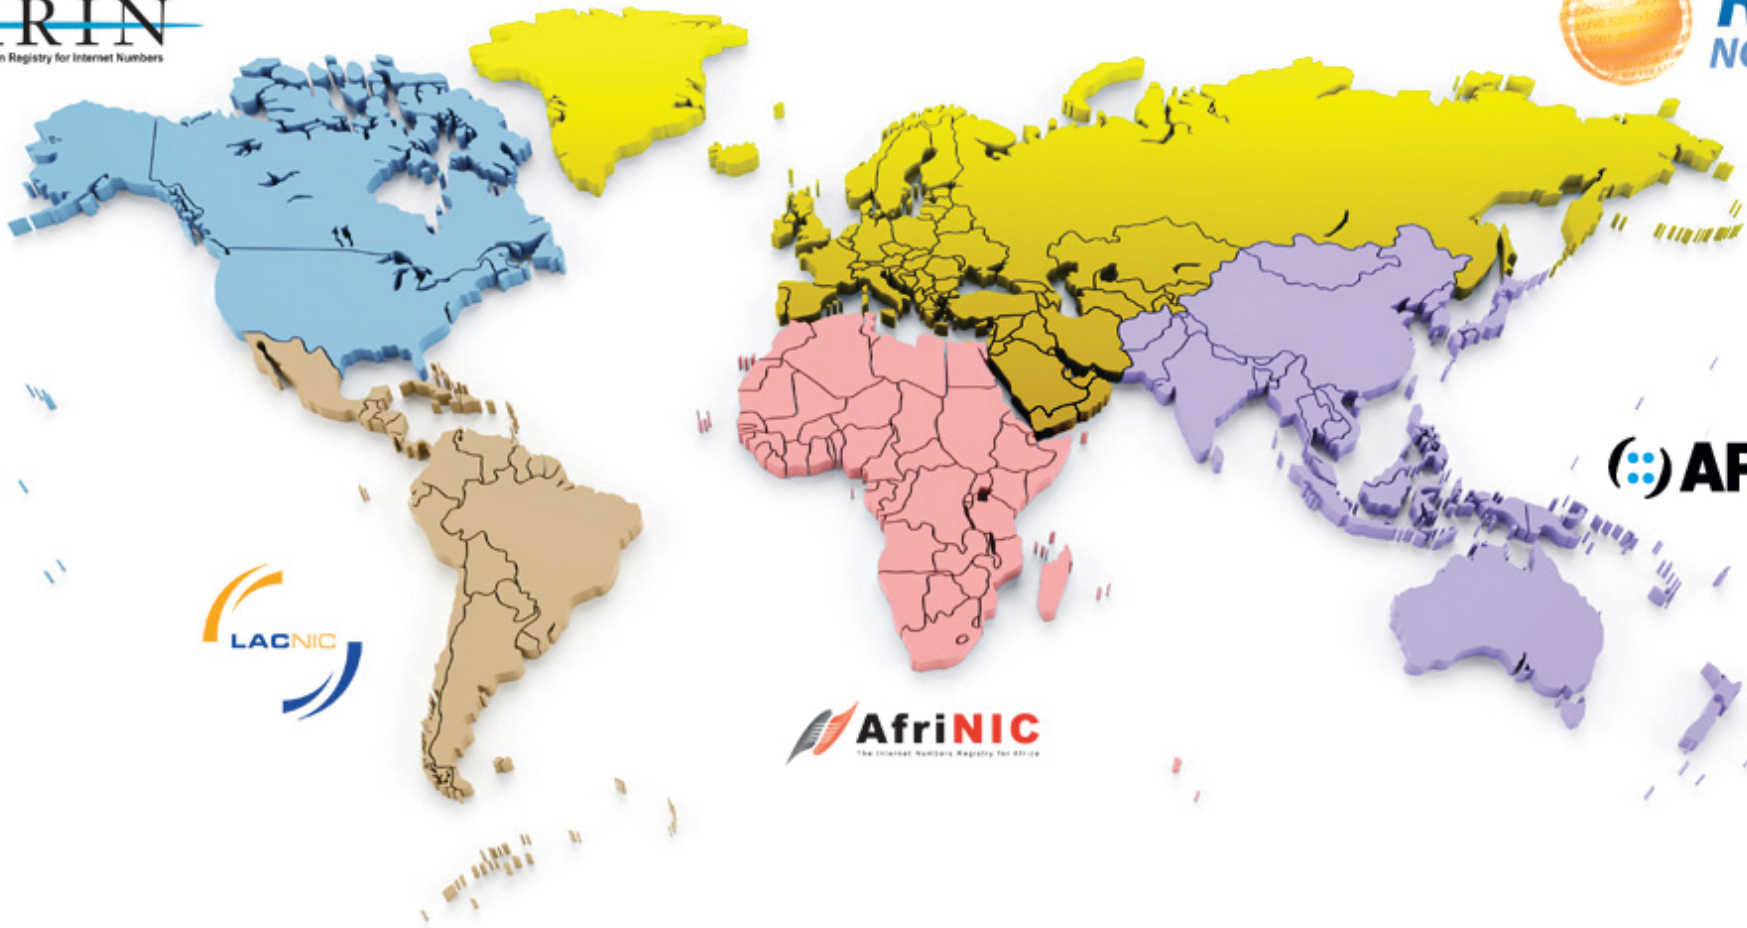
RIPE NCC

- Located in Amsterdam
- Not for profit membership organisation
- One of five RIRs

The five RIRs

ARIN
American Registry for Internet Numbers

 **RIPE**
NCC



 **LACNIC**

 **AfriNIC**
The Internet Numbers Registry for Africa

 **APNIC**

 **RIPE**
NCC

Internet Number Resources

- IP addresses
 - IPv4 eg. 193.0.0.203
 - IPv6 eg. 2001:610:240:11::c100:1319
- Autonomous System Numbers (ASN)
- Other public services
 - Training Services
 - RIPE Database
 - K-root name server
 - Measurement tools
 - E-learning
 - RIPE Labs
 - RIPE Stat
 - RIPE Atlas



Registration

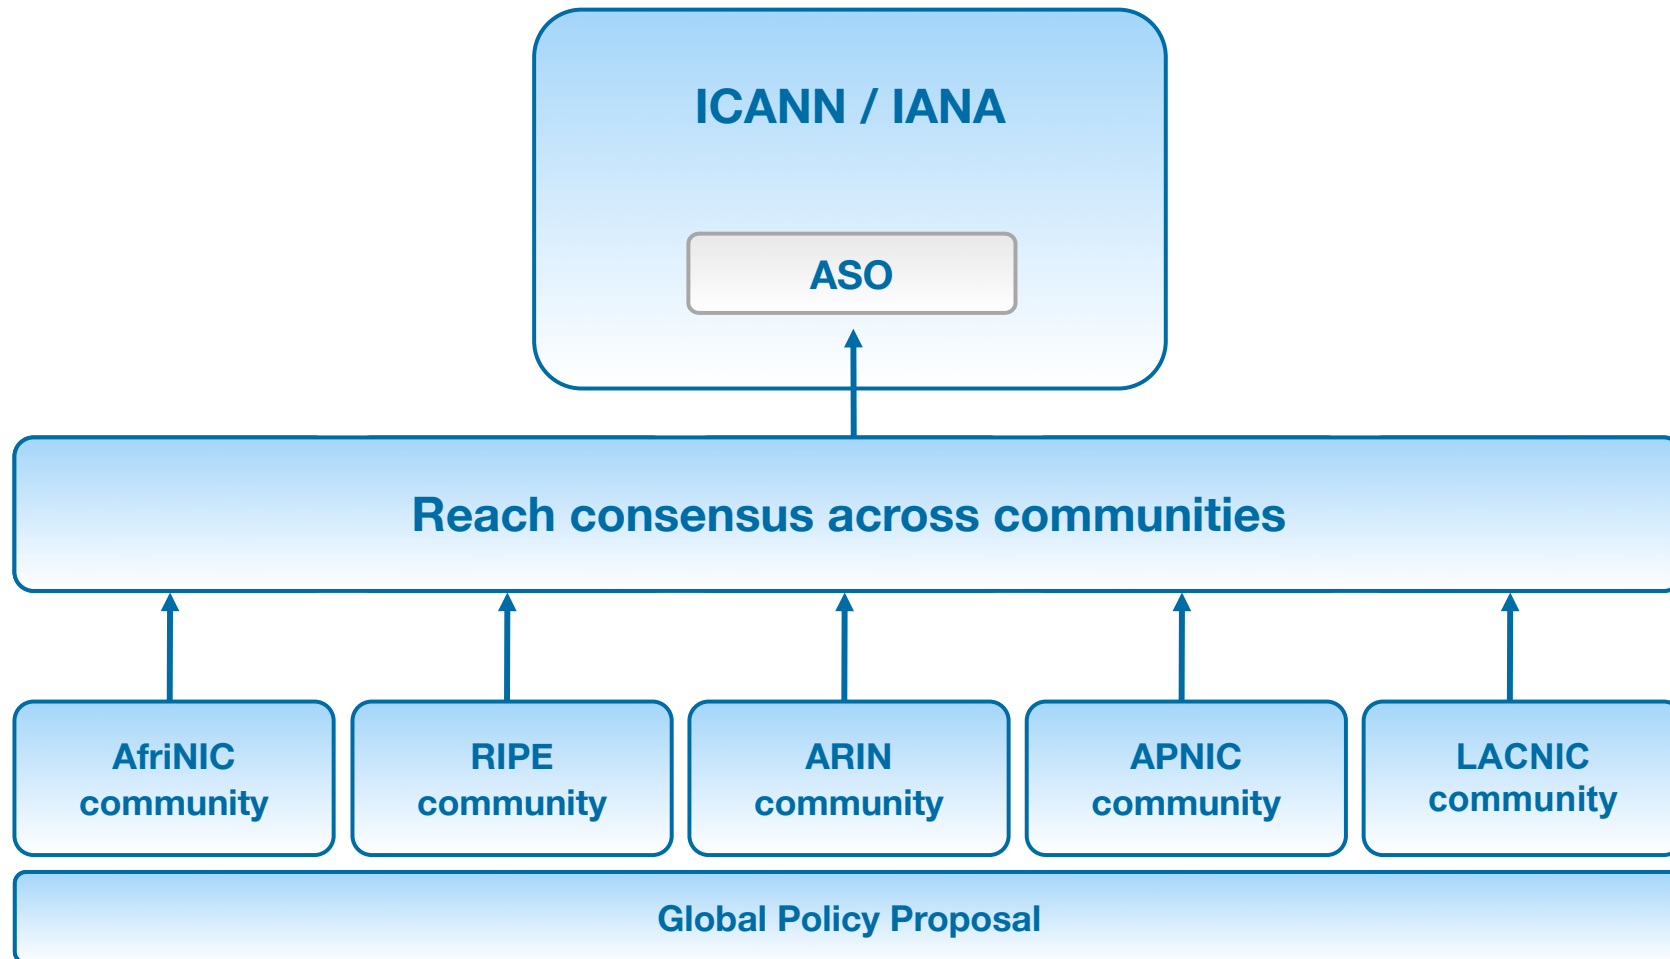


Conservation



Aggregation

Who makes policies?



Why would you want to participate?

- Policy determines how you run your business
- Over 8000 LIRs
- Only a fraction are active participants in the PDP


How can you participate?

- Working Group mailing lists
 - RIPE website → RIPE → Mailing Lists
- Come to the RIPE Meetings
 - Two free tickets for new LIRs
 - Remote participation is also possible


193.0.193.0
40:0:80:10
93.0.19.21.15
240:11::c100:13
0:1315 193.0.0.1
:240:0:53::193
93 193.0.0.1

IPv4 Address Pool Exhaustion

IANA IPv4 pool




IPv4 address distribution




 Allocation  PA Assignment  PI Assignment

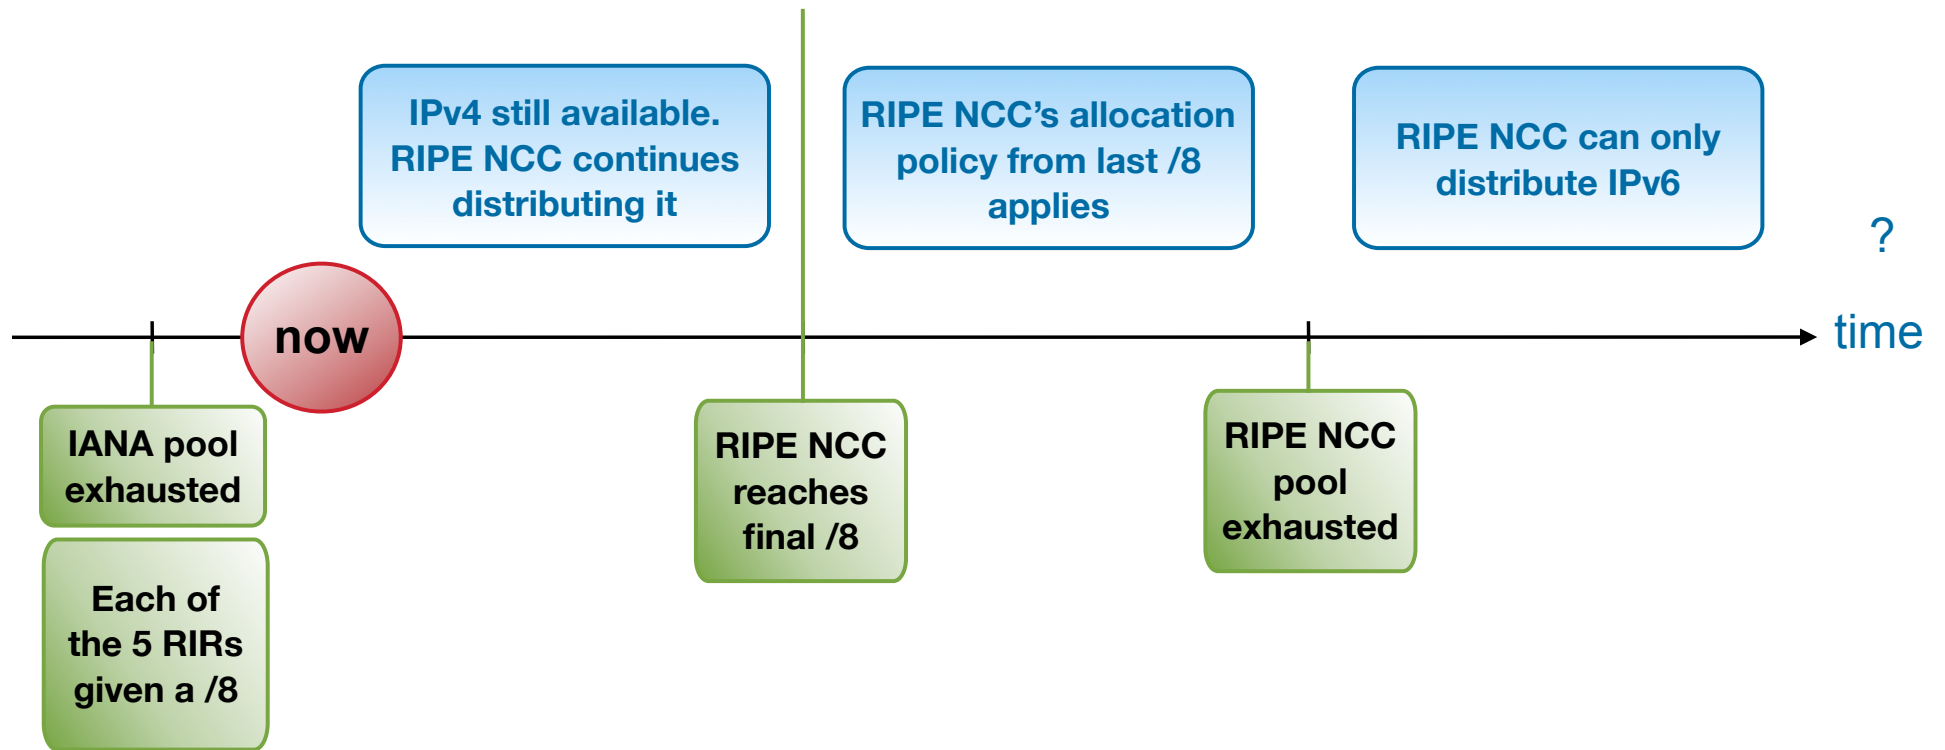
IANA and RIRs IPv4 pool



Our slice of the IPv4 pie



IPv4 exhaustion phases



Run Out Fairly (of IPv4)

- Gradually reduced allocation / assignment periods
- Needs for “Entire Period” of up to...
 - 12 months (January 2010)
 - 9 months (July 2010)
 - 6 months (January 2011)
 - 3 months (July 2011)
- 50% has to be used up by half-period

RIPE NCC's last /8


- We do things differently!
- Ensures IPv4 access for all members
 - 16000+ /22s in a /8
 - members can get **one /22** (=1024 addresses)
 - must already hold IPv6
 - must qualify for allocation
- /16 set aside for unforeseen situations
 - if unused, will be distributed
- No PI

Transfer of IPv4 Allocations

- Policy 2007-08: Allocation Transfer Policy
 - Don't buy your IPv4 on eBay!
 - Transfer unused allocations to another LIR
 - Minimum allocation size /21
 - Evaluated by RIPE NCC
 - Update in RIPE Database


<http://www.ripe.net/lir-services/resource-management/listing>

IPv4 Depletion in the RIPE NCC's Region




- <https://www.ripe.net/internet-coordination/ipv4-exhaustion>

IPv4 Allocation Rate



IPv4 Depletion Worldwide

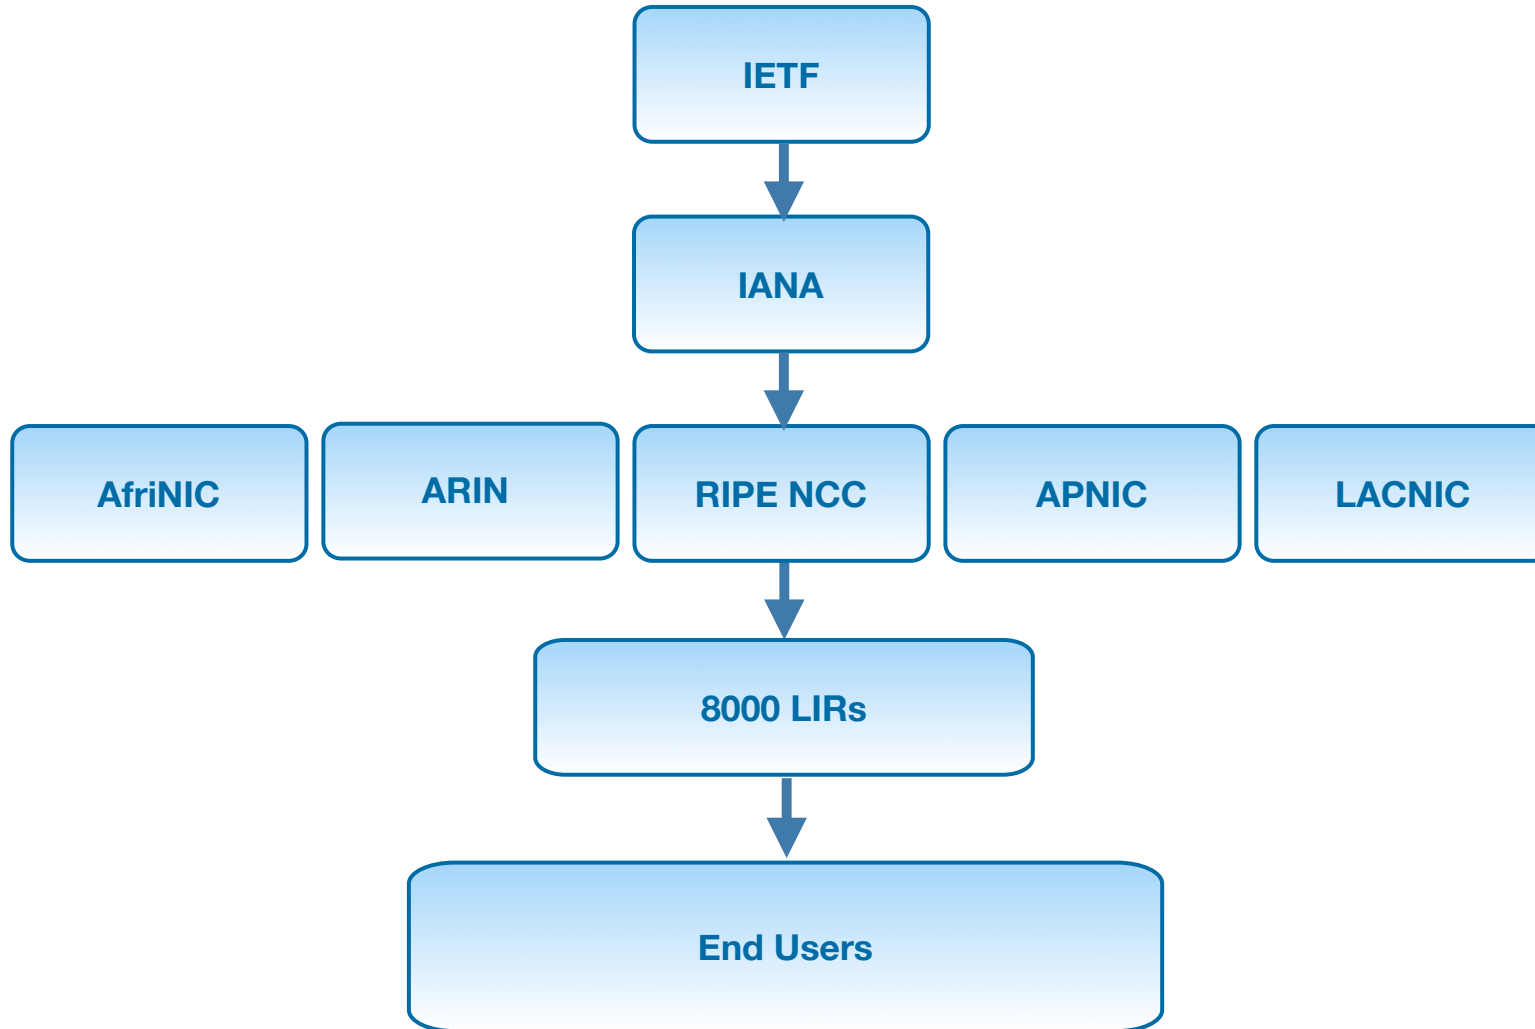


193.0.193.0
40:0:80:10
93.0.19.21.15
240:11::c100:13
0:1315 193.0.0.1
:240:0:53::193
93 193.0.0.1


IPv6 Address Space






Where do all the addresses come from?



IPv6 address distribution



 PA Allocation  Provider Aggregatable Assignment  PI Assignment

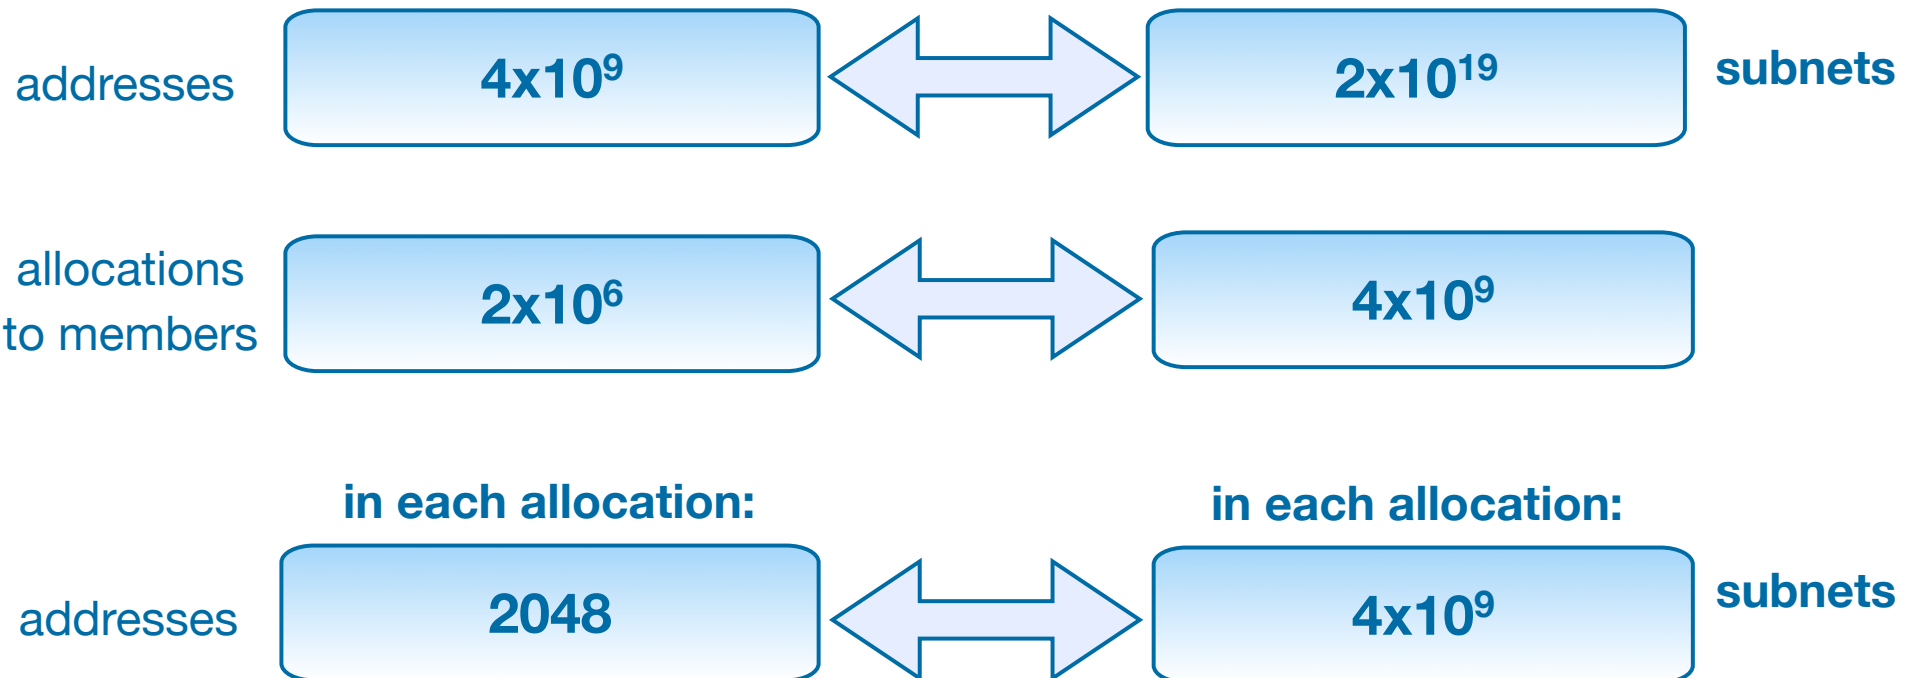
IPv6 basics

- IPv6 address: 128 bits
 - 32 bits in IPv4
- Every subnet should be a /64
- Customer assignments (sites) between:
 - /64 (1 subnet)
 - /48 (65536 subnets)
- Minimum allocation size /32
 - 65536 /48's
 - 16777216 /56's

IPv4 vs IPv6 (rounded off)

IPv4

IPv6



Classless Inter-Domain Routing (CIDR)

IPv6 Chart

Prefix	/48s	/56s	/64s	Bits
/24	16M	4G	1T	104
/25	8M	2G	512G	103
/26	4M	1G	256G	102
/27	2M	512M	128G	101
/28	1M	256M	64G	100
/29	512K	128M	32G	99
/30	256K	64M	16G	98
/31	128K	32M	8G	97
/32	64K	16M	4G	96
/33	32K	8M	2G	95
/34	16K	4M	1G	94
/35	8K	2M	512M	93
/36	4K	1M	256M	92
/37	2K	512K	128M	91
/38	1K	256K	64M	90
/39	512	128K	32M	89
/40	256	64K	16M	88
/41	128	32K	8M	87
/42	64	16K	4M	86
/43	32	8K	1M	85
/44	16	4K	1M	84
/45	8	2K	512K	83
/46	4	1K	256K	82
/47	2	512	128K	81
/48	1	256	64K	80
/49		128	32K	79
/50		64	16K	78
/51		32	8K	77
/52		16	4K	76
/53		8	2K	75
/54		4	1K	74
/55		2	512	73
/56		1	256	72
/57			128	71
/58			64	70
/59			32	69
/60			16	68
/61			8	67
/62			4	66
/63			2	65
/64			1	64

K = 1,024 • M = 1,048,576 • G = 1,073,741,824 • T = 1,099,511,627,776

RIPE NCC

IPv4 CIDR Chart

RIPE NCC

IP Addresses	Bits	Prefix	Subnet Mask
1	0	/32	255.255.255.255
2	1	/31	255.255.255.254
4	2	/30	255.255.255.252
8	3	/29	255.255.255.248
16	4	/28	255.255.255.240
32	5	/27	255.255.255.224
64	6	/26	255.255.255.192
128	7	/25	255.255.255.128
256	8	/24	255.255.255.0
512	9	/23	255.255.254.0
1 K	10	/22	255.255.252.0
2 K	11	/21	255.255.248.0
4 K	12	/20	255.255.240.0
8 K	13	/19	255.255.224.0
16 K	14	/18	255.255.192.0
32 K	15	/17	255.255.128.0
64 K	16	/16	255.255.0.0
128 K	17	/15	255.254.0.0
256 K	18	/14	255.252.0.0
512 K	19	/13	255.248.0.0
1 M	20	/12	255.240.0.0
2 M	21	/11	255.224.0.0
4 M	22	/10	255.192.0.0
8 M	23	/9	255.128.0.0
16 M	24	/8	255.0.0.0
32 M	25	/7	254.0.0.0
64 M	26	/6	252.0.0.0
128 M	27	/5	248.0.0.0
256 M	28	/4	240.0.0.0
512 M	29	/3	224.0.0.0
1024 M	30	/2	192.0.0.0
2048 M	31	/1	128.0.0.0
4096 M	32	/0	0.0.0.0

K = 1,024 • M = 1,048,576

Contact Registration Services:
hostmaster@ripe.net • lir-help@ripe.net


www.ripe.net

Address Notation

2001:0db8:003e:ef11:0000:0000:c100:004d

2001:0db8:003e:ef11:0000:0000:c100:004d

2001:db8:3e:ef11:0:0:c100:4d



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
- Announcement as a single prefix recommended

RIPE Policy Proposal 2011-04

- Extension of the Minimum Size for IPv6 Initial Allocation
 - Proposes initial allocation up to a /29
 - For example, for small LIRs to deploy IPv6 via 6RD (RFC 5969)

UNDER DISCUSSION

- Proposal currently in Review Phase
 - The RIPE NCC is working on impact analysis

What does the first IPv6 allocation cost?

FREE

- for all
- pending General Meeting decision

or:

FREE

- for approximately 97% of the LIRs
- more points, but not higher category!

Why Create an IPv6 Addressing Plan?



- Mental health during implementation(!)
- Easier implementation of security policies
- Efficient addressing plans are scalable
- More efficient route aggregation


IPv6 Subnetting

IPv6 Subnetting

2001:0db8:0000:0000:0000:0000:0000:0000

64 bits interface ID

/32 = 65536 /48
/48 = 65536 /64
/52 = 4096 /64
/56 = 256 /64
/60 = 16 /64
/64

 **RIPE**
NCC

Contact Training Services: training@ripe.net
Follow us on Twitter: www.twitter.com/TrainingRIPENCC
www.ripe.net/training

Make an addressing plan (I)

- Number of hosts is irrelevant
- Multiple /48s per pop can be used
 - separate blocks for infrastructure and customers
 - document address needs for allocation criteria
- /64 for all subnets
 - autoconfiguration works
 - renumbering easier
 - less typo errors because of simplicity

Make an addressing plan (II)

- Use one /64 block (per site) for loopbacks
 - One /128 per device
 - One /64 contains enough /128s for
18.446.744.073.709.551.616 devices

More On Addressing Plans for ISPs

- For private networks, look at ULA
- For servers you want manual configuration
- Use port numbers for addresses
 - pop server 2001:db8:1::110
 - dns server 2001:db8:1::53
 - etc...

Point-to-Point Connections

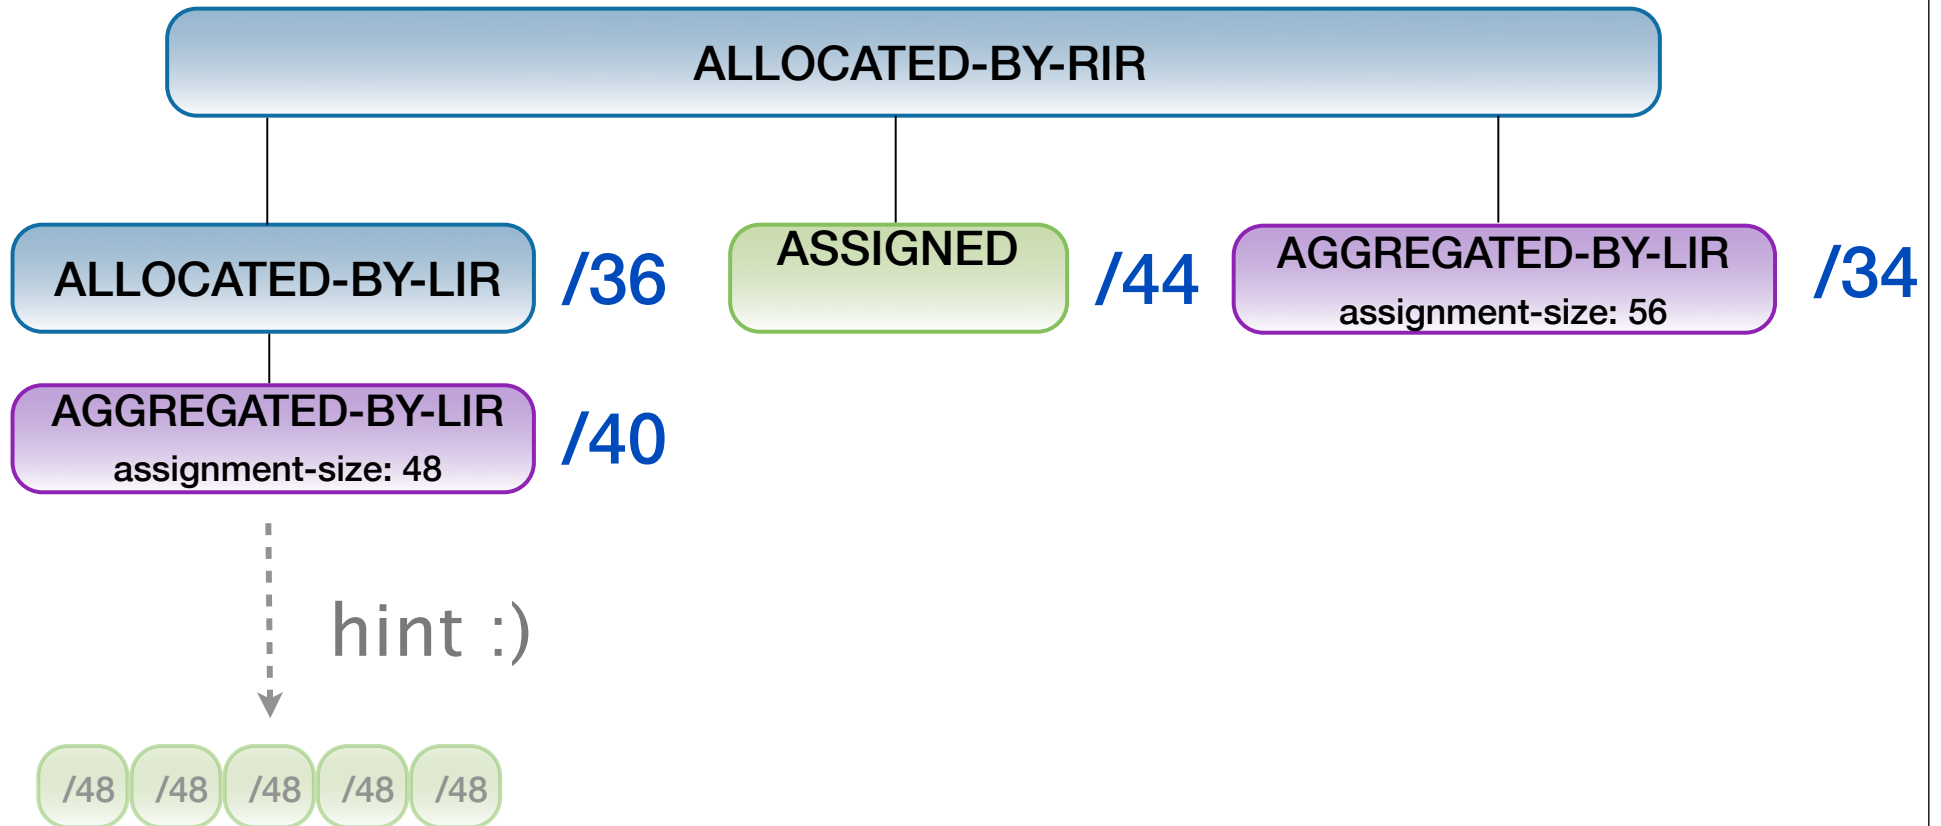
- How much space for point-to-point connections?
 - RFC4291: Interface IDs are required to be /64
 - RFC3627: Use of /127 between routers considered harmful
 - RFC6547: RFC3627 to Historic Status
 - RFC6164: Using /127 on Inter-Router links
- Be safe: reserve a /64, assign a /127 per point-to-point connection

Customer assignments

- Give your customers enough addresses
 - Up to a /48
- For more addresses, send in request form
 - Alternatively, make a sub-allocation
- Every assignment must now be registered in the RIPE database



Using AGGREGATED-BY-LIR




Group assignments in the RIPE DB

inet6num: 2001:db8:1000::/36
netname: Bluelight
descr: We want more Bluelight B.V.
descr: Colocation services
country: NL
admin-c: BN649-RIPE
tech-c: BN649-RIPE
status: AGGREGATED-BY-LIR
assignment-size: 48
mnt-by: BLUELIGHT-MNT
notify: noc@example.net
changed: noc@example.net 20110218
source: RIPE



Customers And Their /48

- Customers have no idea how to handle 65536 subnets!
- Provide them with information
 - https://www.ripe.net/lir-services/training/material/IPv6-for-LIRs-Training-Course/IPv6_addr_plan4.pdf



IPv6 Address Management

- Your Excel sheet might not scale
 - There are 65.536 /48s in a /32
 - There are 65.536 /64s in a /48
 - There are **16.777.216** /56s in a /32

- Find a suitable IPAM solution

Getting IPv6 PI address space

- To qualify, an organisation must:
 - Meet the contractual requirements for provider independent resources
- Minimum assignment size /48

Reverse DNS

2001:db8:3e:ef11::c100:4d

Reverse DNS

2001:0db8:003e:ef11:0000:0000:c100:004d

.ip6.arpa

d.4.0.0.0.1.c.0.0.0.0.0.0.0.1.1.f.e.e.


3.0.0.8.b.d.0.1.0.0.2.ip6.arpa PTR

yourname.domain.tld


d.4.0.0.0.1.c.0.0.0.0.0.0.0.1.1.f.e.e.3.0.0.8.b.d.0.1.0.0.2.ip6.arpa PTR yourname.domain.tld

IPv6 Deployment Statistics

IPv6 Allocation Rate




IPv6 PI Assignments



Members with IPv6 and IPv4

7853 members with IP resources




● IPv4 only ● IPv6 only ● IPv6 and IPv4

IPv6 Ripeness


- Rating system:
 - One star if the LIR has an IPv6 allocation
 - Additional stars if:
 - IPv6 Prefix is announced on router
 - A route6 object is in the RIPE Database
 - Reverse DNS is set up
 - A list of all 4 star LIRs: <http://ripeness.ripe.net/>

IPv6 RIPv6ness: 8097 LIRs (5 May 2012)


● 1 star ● 2 stars ● 3 stars ● 4 stars ● No IPv6




IPv6 RIPEness – countries (5 May 2012)



IPv6 RIPEness – relative (5 May 2012)



IPv6 enabled ASes in global routing



World IPv6 Launch



- <http://www.worldipv6launch.org/>

Useful information

Websites

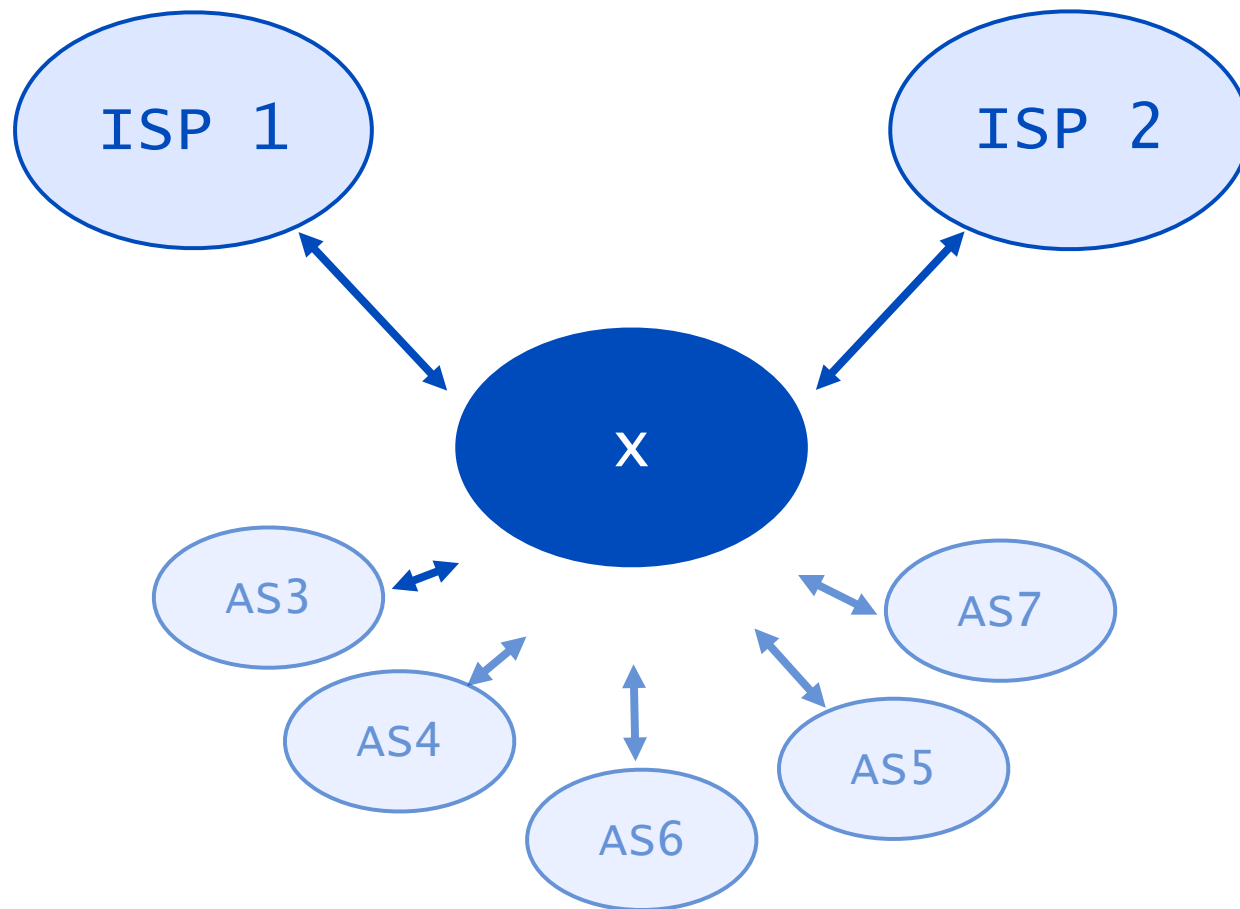
- <http://www.getipv6.info/>
- <http://www.ipv6actnow.org>
- <http://datatracker.ietf.org/wg/v6ops/>
- <http://www.ripe.net/ripe/docs/ripe-501.html>

Mailing lists

- <http://lists.cluenet.de/mailman/listinfo/ipv6-ops>
- <http://www.ripe.net/mailman/listinfo/ipv6-wg>


Multihomed BGP Routing Setup

Scenario 1: LIR = PA allocation + ASN




- Can make assignments to End Users

Scenario 2: End User = PI + ASN




- Can NOT sub-assign further!!!
 - (in IPv4 can still use PI for xDSL, broadband...)

Scenario 3: LIR = PI + ASN



- Can NOT sub-assign further!!!
 - (in IPv4 can still use PI for xDSL, broadband...)

Scenario 4: PI End User, not multihomed



- Part of LIR's AS number
 - does not want to / can not run BGP
 - still wants “portable” addresses

How to get an AS Number

- Assignment requirements
 - Address space
 - Multihoming
 - One AS Number per network
- For LIR itself
- For End User
 - Sponsoring LIR requests it for End User
 - Direct Assignment User requests it for themselves

32-bit AS Numbers and you

- New format: “AS4192351863”
- Act now!
- Prepare for 32-bit ASNs in your organisation:
 - Check if hardware is compatible;
if not, contact hardware vendor
 - Check if upstream uses compatible hardware;
if not, they should upgrade!

193.0.193.0
40:0:80:10
93.0.19.21.15
240:11::c100:13
0:1315 193.0.0.1
:240:0:53::193
93 193.0.0.1

RIPE Database



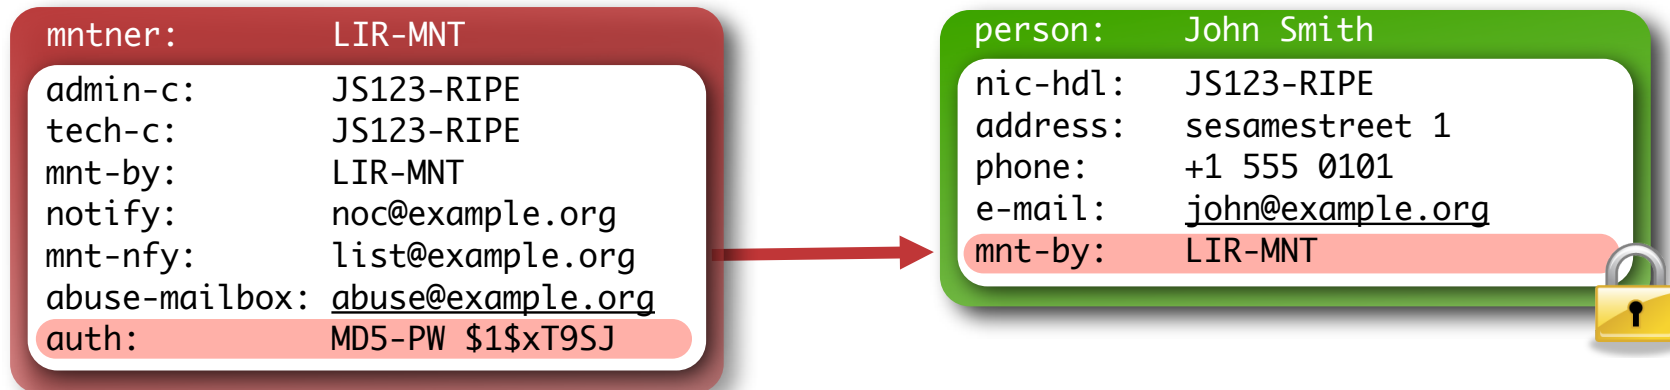
RIPE Database

- Public Internet resource and routing registry database

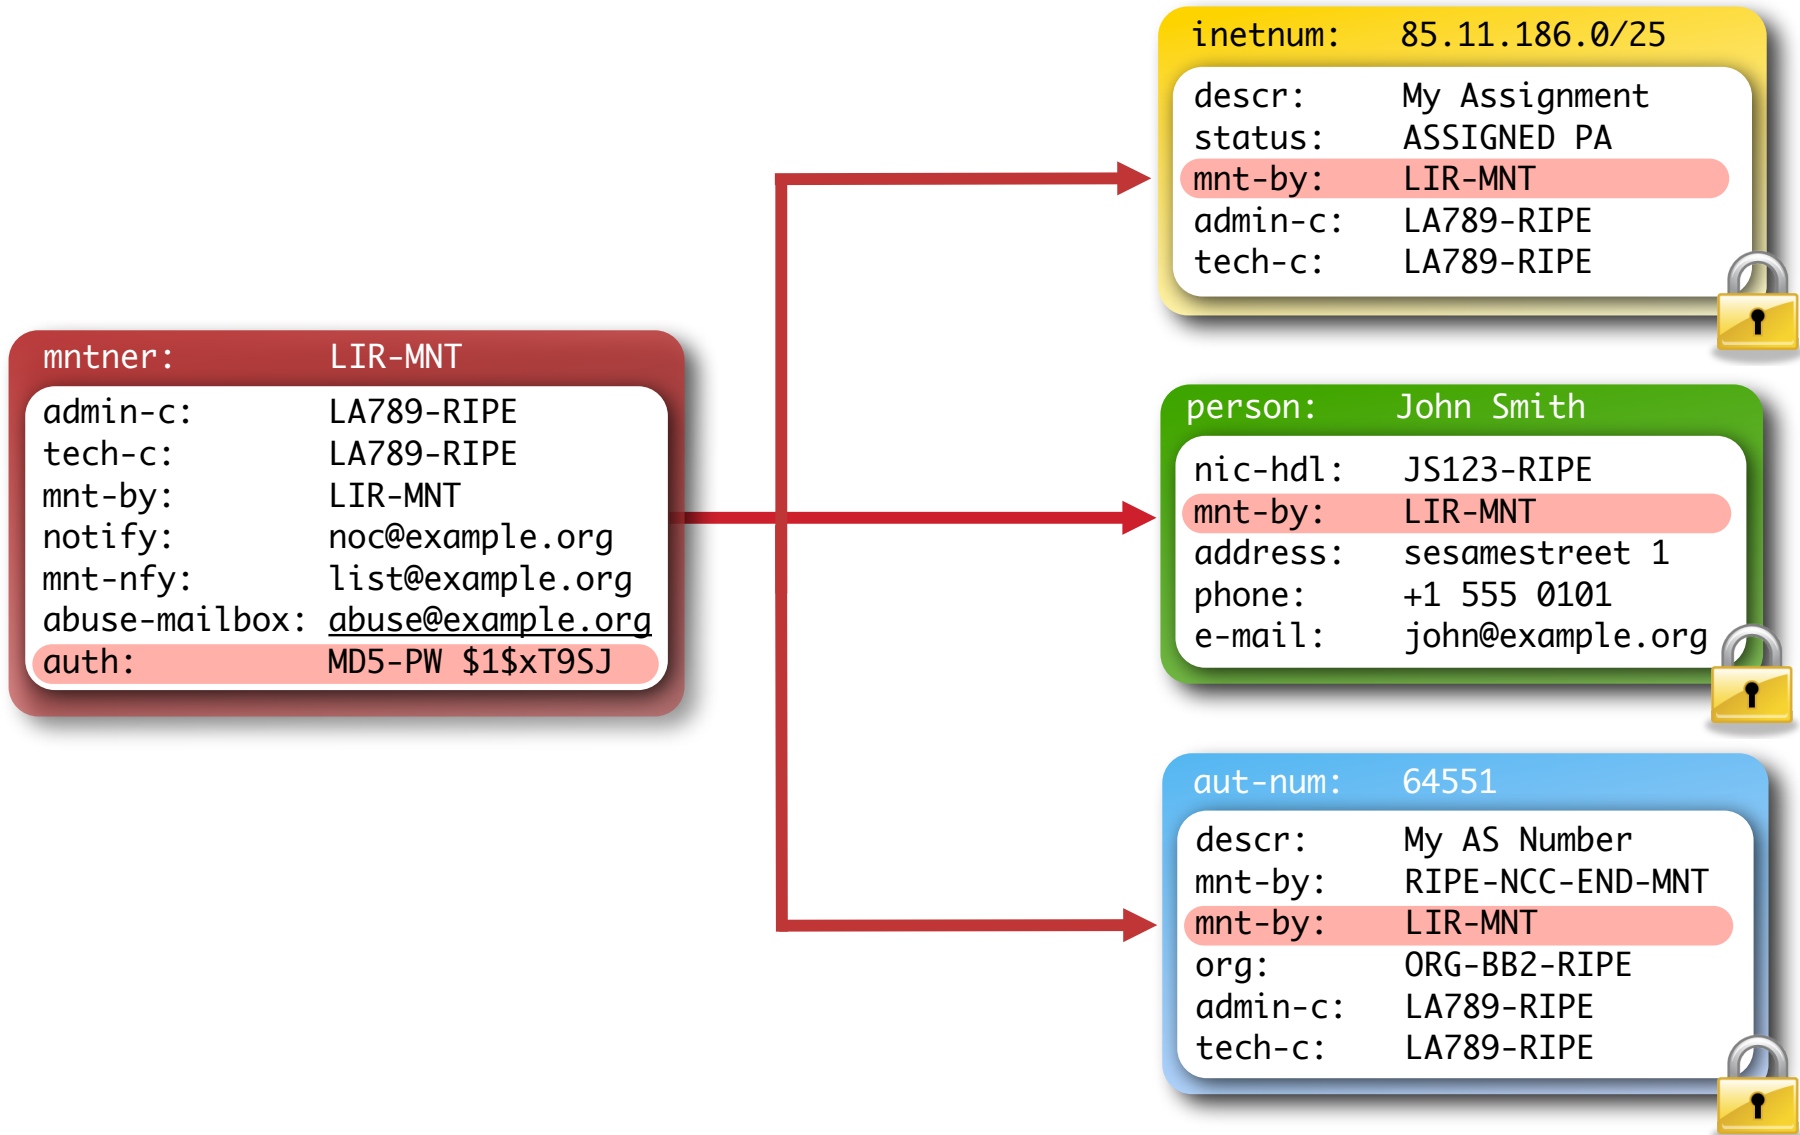
RIPE Database objects

- Resources
 - inetnum, inet6num, aut-num, domain
- Routing
 - route, route6, aut-num
- Security
 - mntner
- Contact
 - organisation, person, role

Protection



Protection



193.0.193.0
40:0:80:10
93.0.19.21.15
240:11::c100:13
0:1315 193.0.0.1
:240:0:53::193
93 193.0.0.1

Routing Registry



What is “Internet Routing Registry”

- Distributed databases with public routing policy information, mirroring each other: irr.net
 - APNIC, RADB, Level3, SAVVIS...
- RIPE NCC operates “RIPE Routing Registry”
- Big operators make use of it
 - AS286 (KPN), AS5400 (BT), AS1299 (Telia), AS8918 (Carrier1), AS2764 (Connect), AS3561 (Savvis), AS3356 (Level 3)...

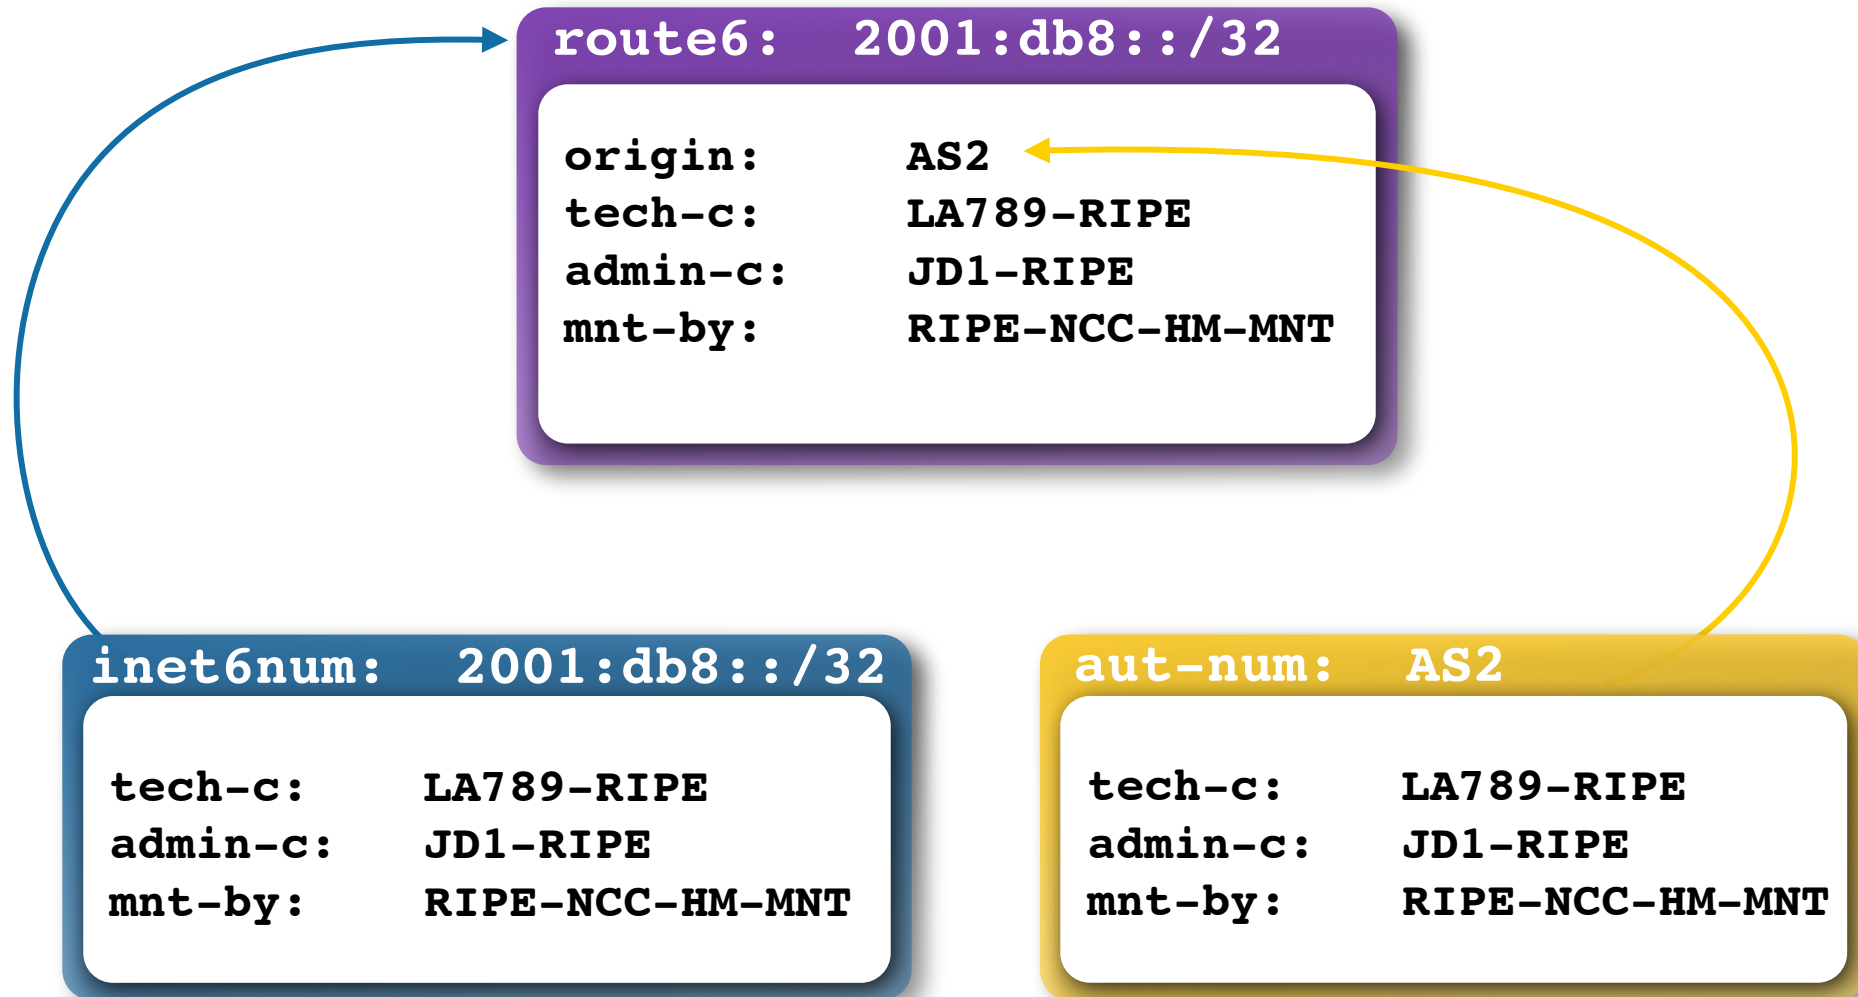
Publishing routing policy in IRR

- Required by some Transit Providers & IXPs
 - they use it for prefix-based filtering
- Allows for automated generation of prefix filters
 - and router configuration commands, based on RR
- Contributes to routing security
 - prefix filtering based on IRR registered routes
 - prevents accidental leaks and route hijacking
- Good housekeeping

RIPE RR is part of the RIPE Database

- route[6] object creation is responsibility of LIR
 - every time you receive a new allocation, do create a route or route6 object
- route and route6 objects represent routed prefix
 - address space being announced by an AS number

Route and route6 object



IPv6 in the Routing Registry

Route6 object:

```
route6:    2001:DB8::/32
origin:    AS65550
```

Aut-num object:

```
aut-num:   AS65550
mp-import: afi ipv6.unicast from AS64496 accept ANY
mp-export: afi ipv6.unicast to AS64496 announce AS65550
```

Automation of router configuration

- Describing routing policy in aut-num enables generation of route-maps for policy routing
- Tools can read your policy towards peers
 - translation from RPSL to router configuration commands
- Tools collect the data your peers have in RR
 - if their data changes, you only have to periodically run your scripts to collect updates

193.0.193.0
40:0:80:10
93.0.19.21.15
240:11::c100:13
0:1315 193.0.0.1
:240:0:53::193
93 193.0.0.1

Resource Certification



Limitations of the Routing Registry

- Many registries exist, operated by different parties:
 - Not all of them mirror each other
 - Do you trust the information they provide?
- The IRR system is far from complete
- Resulting filters are hard to maintain and can take a lot of router memory

The RIPE NCC involvement in RPKI

- The authority who is the holder of an Internet Number Resource in our region
 - IPv4 and IPv6 address ranges
 - Autonomous System Numbers
- Information is kept in the registry
- Accuracy and completeness are key

Digital resource certificates

- Issue digital certificates along with the registration of Internet Resources
- Two main purposes:
 - Make the registry more robust
 - Making Internet Routing more secure
- Validation is the added value



Using certificates

- Certification is a free, opt-in service
 - Your choice to request a certificate
 - Linked to your membership
 - Renewed every 12 months
- Certificate does not list any identity information
- Digital proof you are the holder of a resource



The PKI system

- The RIRs hold a self-signed root certificate for all the resources that they have in the registry
 - They are the trust anchor for the system
- That root certificate is used to sign a certificate that lists your resources
- You can issue child certificates for those resources to your customers
 - When making assignments or sub allocations

Certificate Authority (CA) Structure


Root CA (RIPE NCC)



Member CA (LIR)



Customer CA



Which resources are certified?

- Provider Aggregatable (PA) IP addresses
- Provider Independent (PI) addresses marked as “Infrastructure”
- Other resources will be added over time:
 - PI addresses for which we have a contract
 - ERX resources

Route Origination Authorisation (ROA)

- Next to the prefix and the ASN which is allowed to announce it, the ROA contains:
 - A minimum prefix length
 - A maximum prefix length
 - An expiry date
- Multiple ROAs can exist for the same prefix
- ROAs can overlap

Publication and validation

- ROAs are published in the same repositories as the certificates and their keys
- You can download them and use software to verify all the cryptographic signatures are valid
 - Was this really the owner of the prefix?
- You will end up with a list of prefixes and the ASN that is expected to originate them
 - And you can be sure the information comes from the holder of the resources

Validator

ROA Validation

- You can download all the certificates, public keys and ROAs which form the RPKI
- Software running on your own machine can retrieve and then verify the information
 - Cryptographic tools can check all the signatures
- The result is a list of all valid combinations of ASN and prefix, the “validated cache”

Reasons for a ROA to be invalid

- The start date is in the future
 - Actually this is flagged as an error
- The end date is in the past
 - It is expired and the ROA will be ignored
- The signing certificate or key pair has expired or has been revoked
- It does not validate back to a configured trust anchor

The Decision Process

- When you receive a BGP announcement from one of your neighbors you can compare this to the validated cache
- There are three possible outcomes:
 - **Unknown**: there is no covering ROA for this prefix
 - **Valid**: a ROA matching the prefix and ASN is found
 - **Invalid**: There is a ROA but it does not match the ASN or the prefix length

Modifying the Validated Cache

- The RIPE NCC Validator allows you to manually override the validation process
- Adding an ignore filter will ignore all ROAs for a given prefix
 - The end result is the validation state will be “unknown”
- Creating a whitelist entry for a prefix and ASN will locally create a valid ROA
 - The end result is the validation state becomes “valid”

The Decision is Yours

- The Validator is a tool which can help you making informed decisions about routing
- Using it properly can enhance the security and stability of the Internet
- It is your network and you make the final decision

Public Testbeds

- A few people allow access to routers that run RPKI and allow you to have a look at it
- RIPE NCC has a Cisco:
 - Telnet to rpkirtr.ripe.net
 - User: ripe, no password
- Eurotransit has a Juniper:
 - Telnet to 193.34.50.25 or 193.34.50.26
 - Username: rпки, password: testbed

<http://www.ripe.net/lir-services/resource-management/certification/tools-and-resources>

Roadmap

- Support for non-hosted is still under development by the RIPE NCC
 - Expected release will be third quarter 2012
- We can give you access to beta test
 - Mail certification@ripe.net if you are interested
- More information will be published on the certification website
 - <http://www.ripe.net/certification>

Follow us!

twitter

@TrainingRIPENCC

Questions?

training@ripe.net



The End!

Край

Y Diwedd

النهاية

Соңы

ჟებრე

Fí

Finis

Ende

Finvezh

Liðugt

Кінець

Konec

Kraj

Ěnn

Fund

پایان

Lõpp

Beigas

Vége

Son

Край

An Críoch

הסוף

Fine

Endir

Sfârșit

Fin

Τέλος

Einde

Конец

Slut

Slutt

დასასრული

Pabaiga

Fim

Amaia

Loppu

Tmíem

Koniec