

IPv6 for LIRs & the Routing Registry

ENOG/RIPE NCC Regional Meeting
June 2011, Moscow

Ferenc Csorba



Schedule

- IPv4 exhaustion
- IPv6 address space
- Russian and regional IPv6 deployment statistics
- BGP multihoming
- Routing & the RIPE Database

RIPE / RIPE NCC

RIPE

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

RIPE NCC

- Located in Amsterdam
- Not for profit membership organisation
- One of five RIRs - distributes IP & ASN


How can you influence addressing policies

- Take part in email discussions
 - RIPE website → RIPE → Mailing Lists
- Come to the RIPE Meetings
 - Amsterdam was in May, Vienna in October
 - Two free tickets for new LIRs
 - Remote participation 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

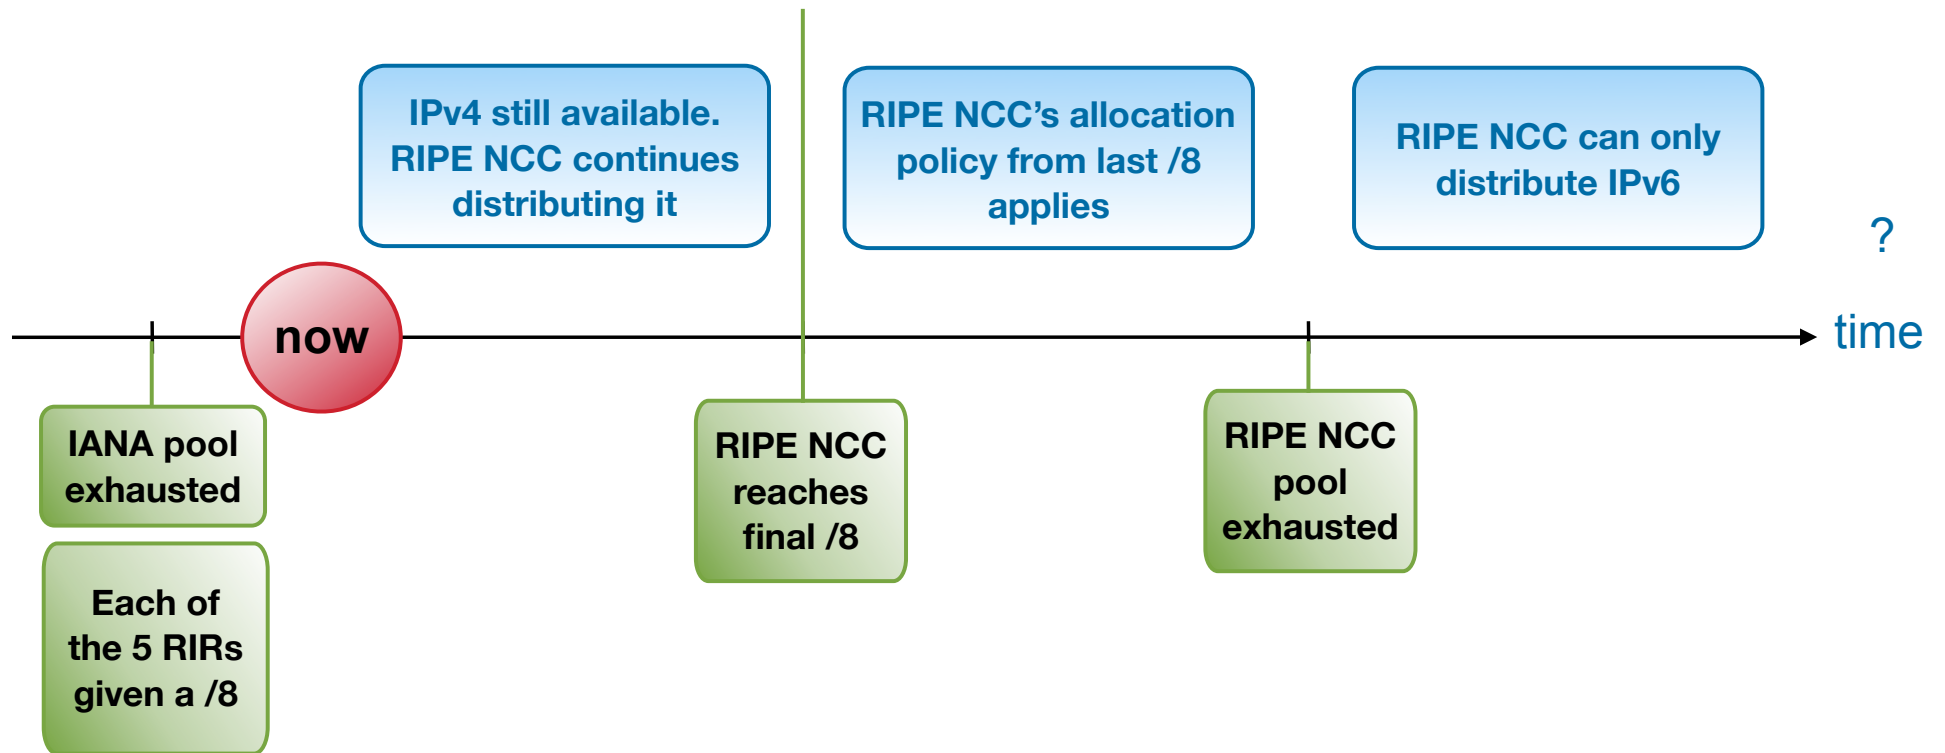


RIPE NCC's IPv4 Pool



<http://www.ripe.net/internet-coordination/ipv4-exhaustion/ipv4-available-pool-graph>

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

How will we evaluate your requests?

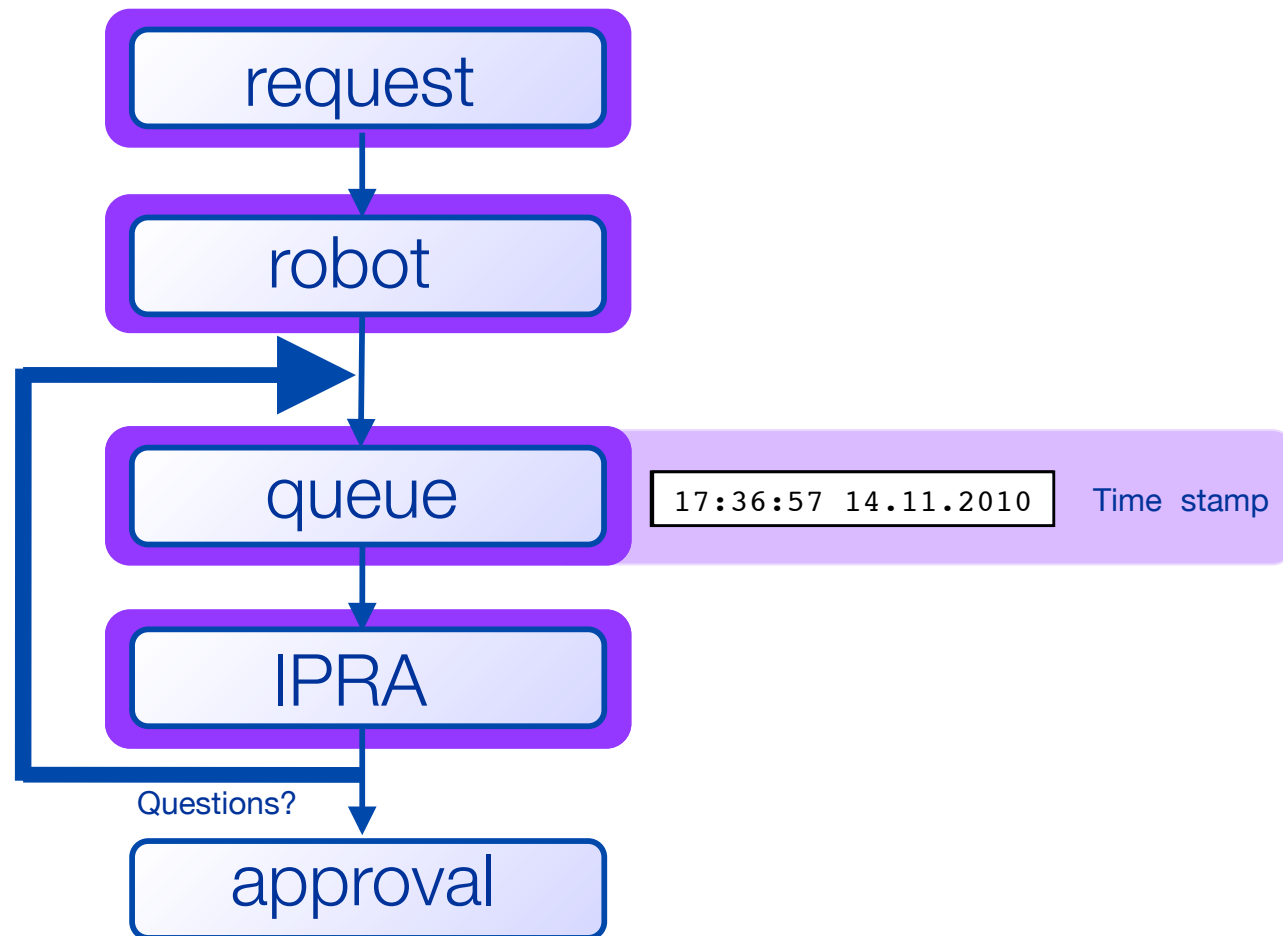
- Find all criteria at:

‘IPv4 Evaluation Procedures’ page

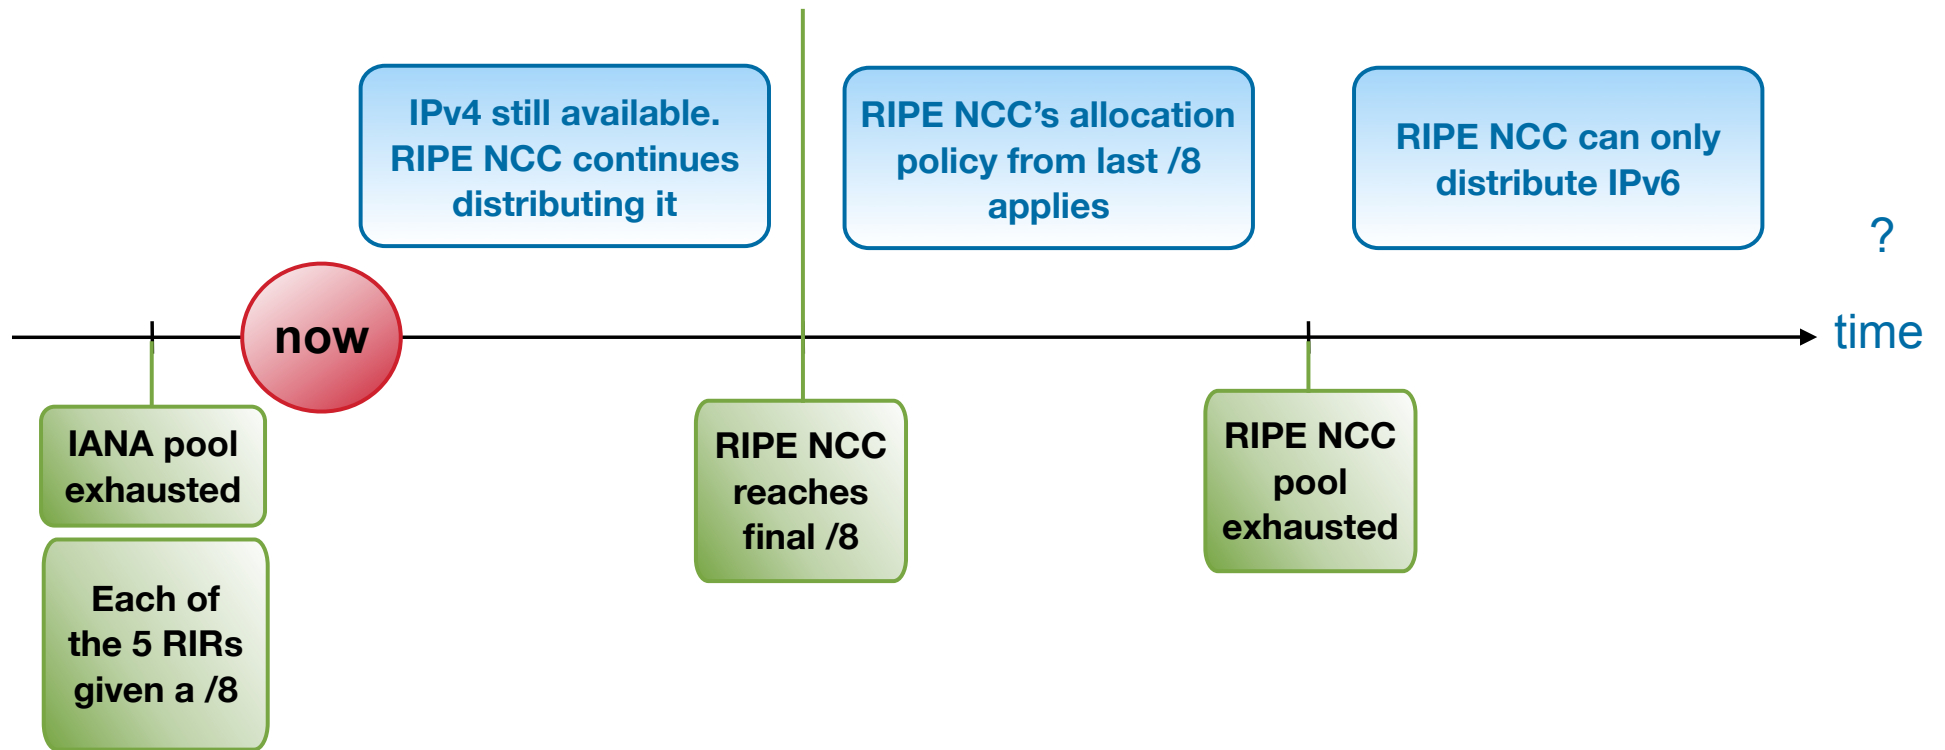
[http://www.ripe.net/lir-services/resource-management/
contact/ipv4-evaluation-procedures](http://www.ripe.net/lir-services/resource-management/contact/ipv4-evaluation-procedures)

New: All IPv4 Requests in one queue

New and ongoing requests. Every email: new time stamp



IPv4 exhaustion phases



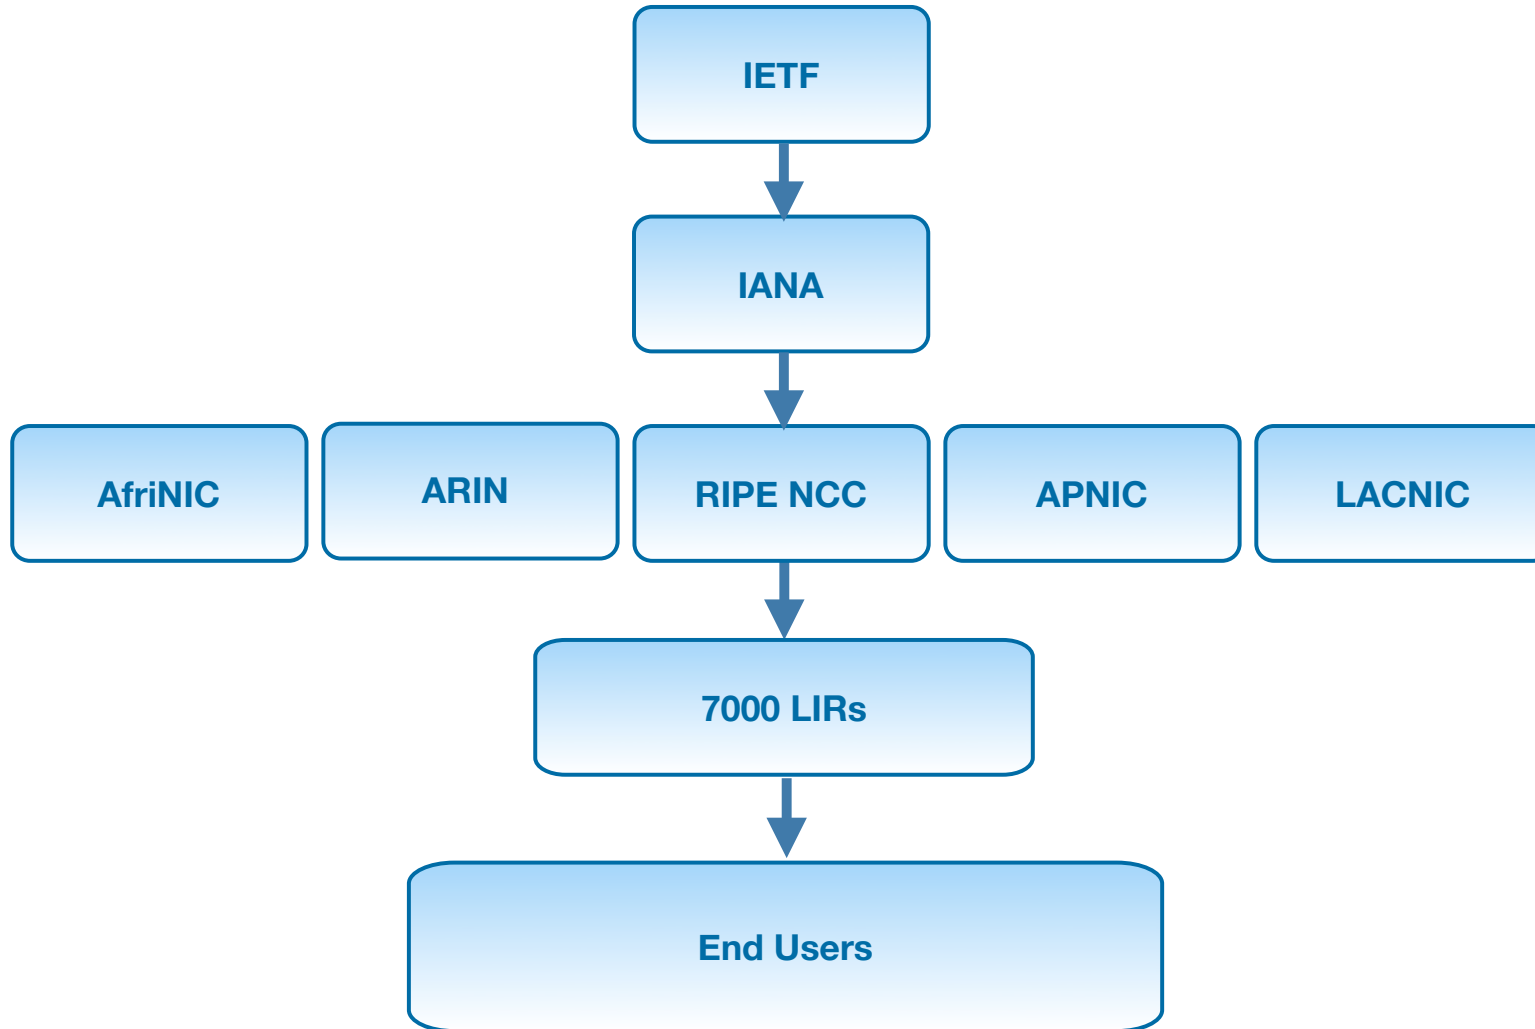
RIPE NCC's last /8

- 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**

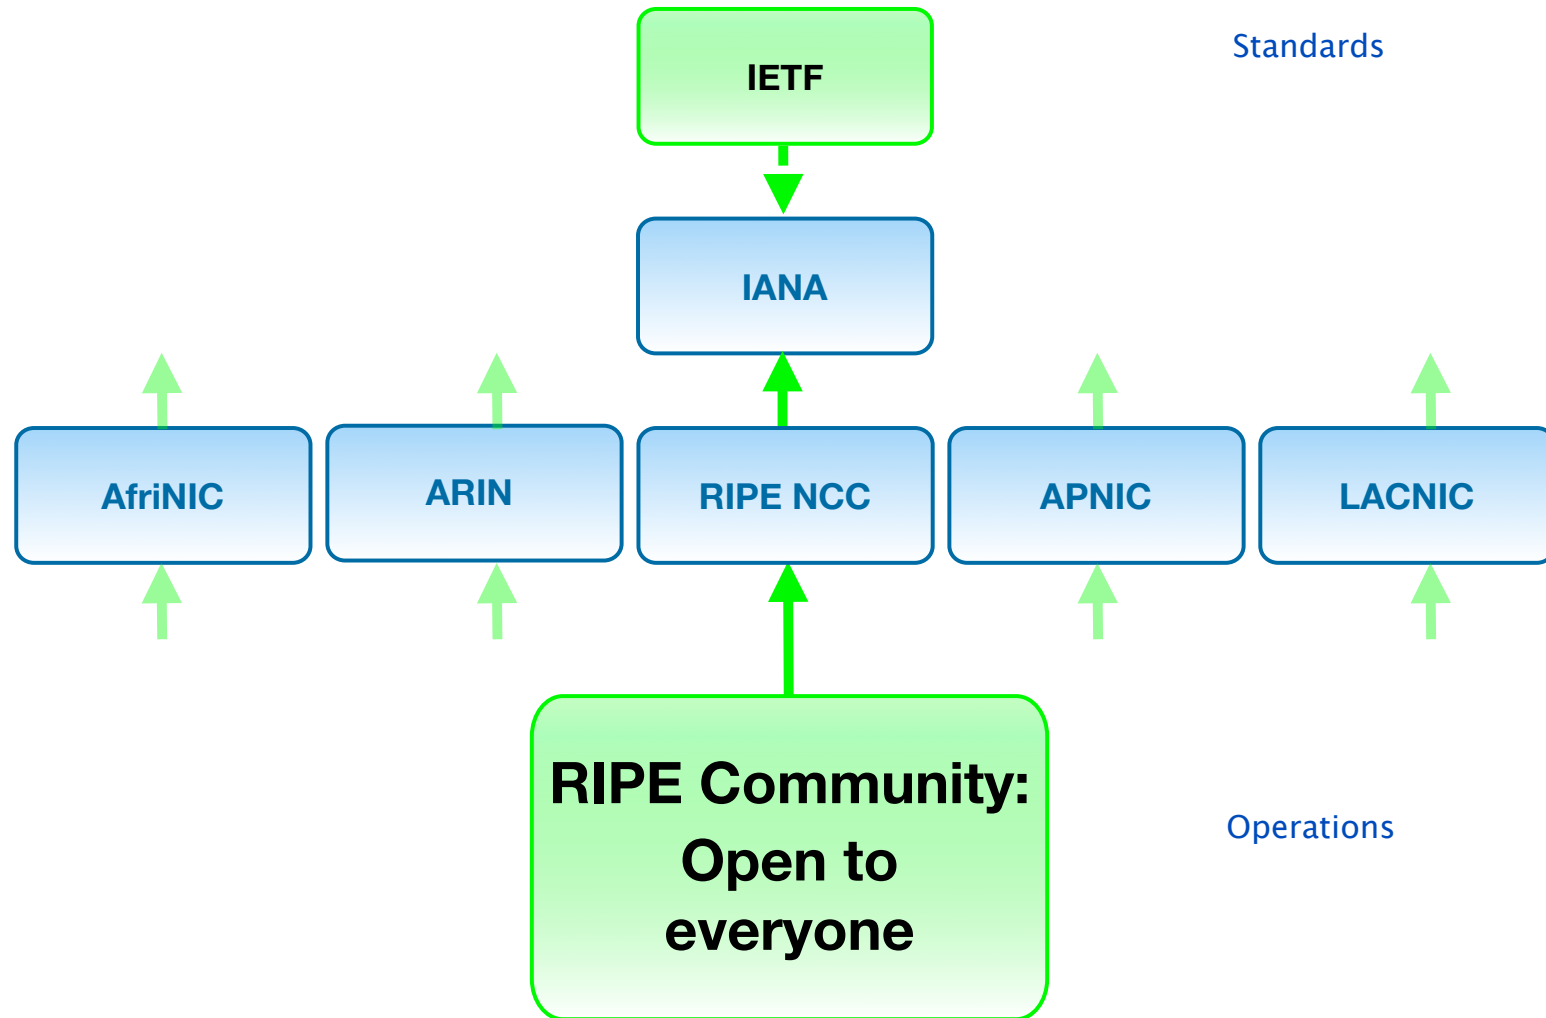
193.0.193.0
40:0:80:10
93.0.19.21.15
240:11::c100:13
0:1315 193.0.0.1
:240:0:53::193
93 193.0.0.1

IPv6 Address Space

Where do all the addresses come from?



Policy process: decision making





Registration



Conservation



Aggregation

Governing principles of addressing policy

- Registration (in RIR whois databases)
 - Ensure uniqueness of Internet number resources
 - Provide contact information for users of Internet number resources
- Aggregation
 - Introduction of Classless Inter Domain Routing (CIDR)
 - Provide scalable routing solution for Internet
- Conservation
 - Policies to ensure fair usage
 - Number resources are distributed based on need

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

IPv6 address distribution



 PA Allocation  Provider Aggregatable Assignment  PI Assignment

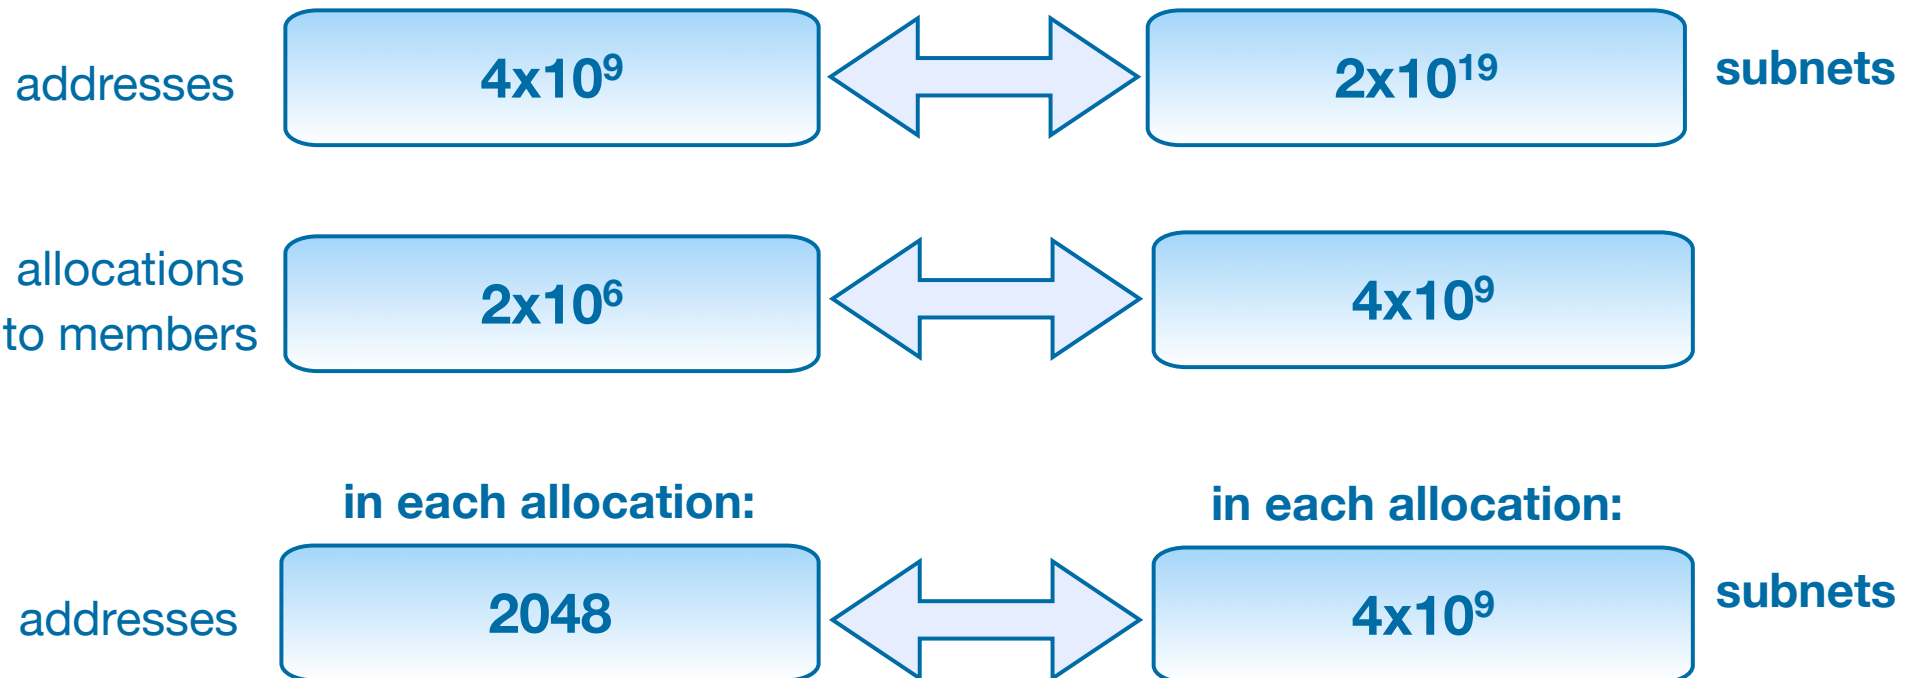
IPv6 basics

- 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

IPv4 vs IPv6 (rounded off, theoretically)

IPv4

IPv6



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

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!

Making addressing plans

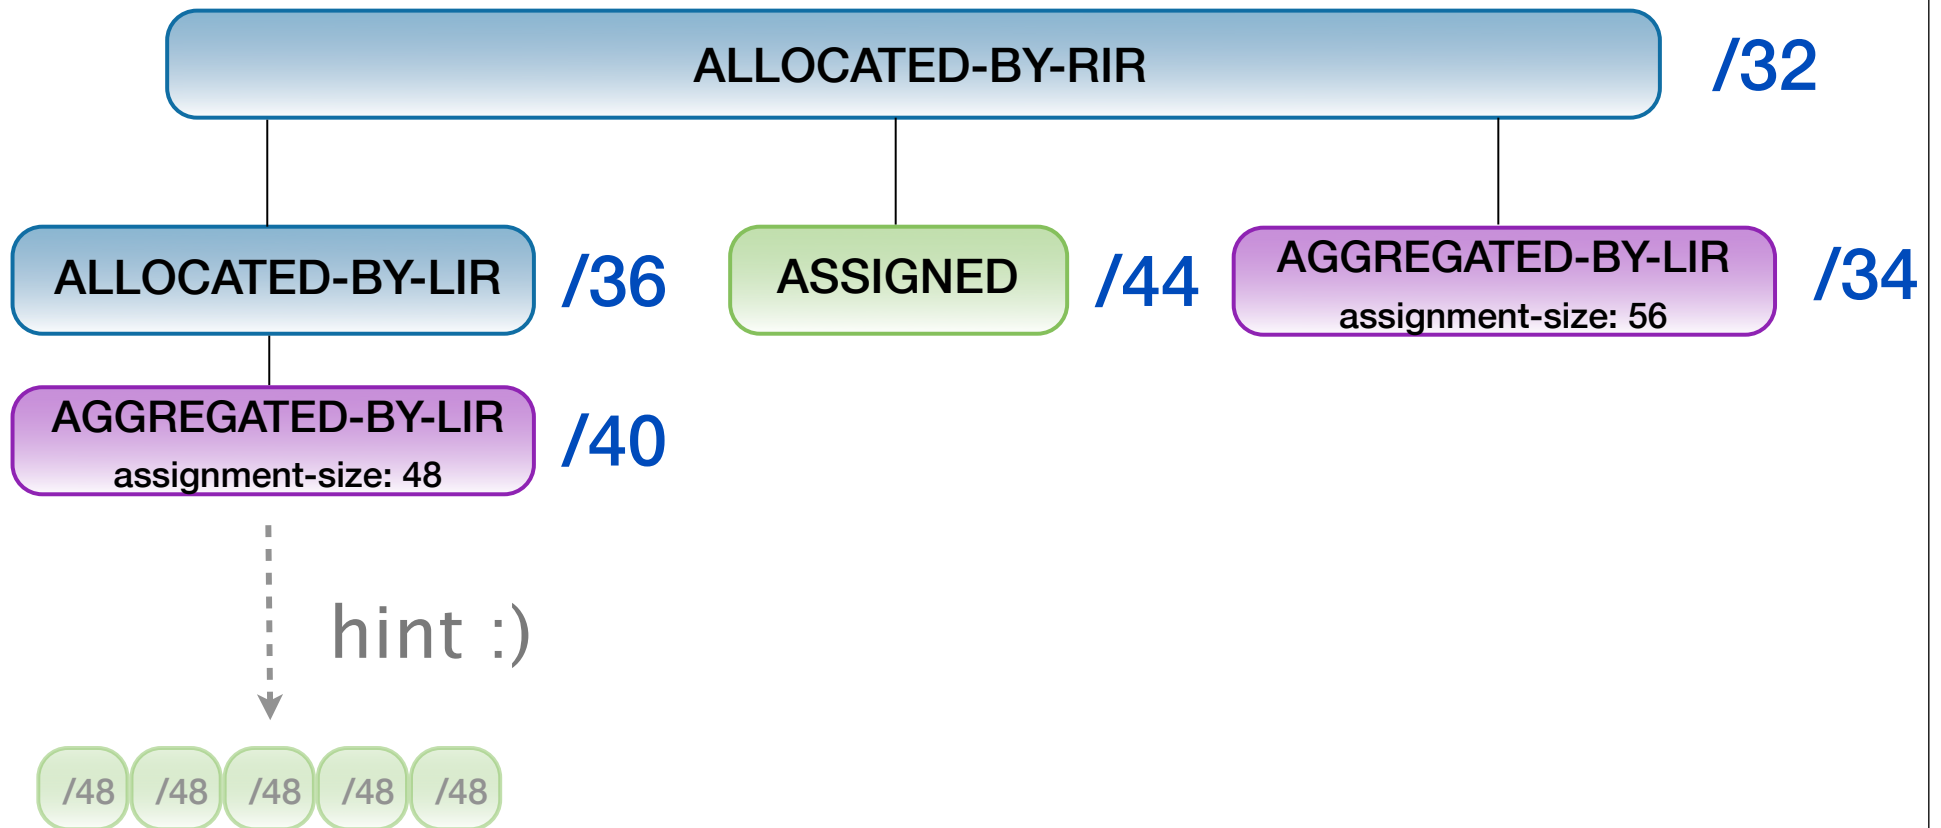
- Number of hosts is irrelevant
- Multiple /48s per pop can be used
 - separate blocks for infrastructure and customers
 - document address needs for allocation criteria
- Use one /64 block per site for loopbacks
- /64 for all subnets
 - autoconfiguration works
 - renumbering easier
 - less typo errors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

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



Getting IPv6 PI address space

- To qualify, an organisation must:
 - Demonstrate it will multihome
 - Meet the contractual requirements for provider independent resources

- Minimum assignment size /48

Getting IPv6 PI address space for an LIR


- To qualify, an organisation must:
 - Demonstrate it will multihome
 - 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

LIR's IPv6 PI **cannot** be used for

- DSL, cable, GPRS customers
- Webhosting, if IP addresses not shared

IPv6 and IPv4 compatibility?

- IPv6 is a different protocol from IPv4
- IPv6 hosts cannot talk to IPv4 hosts directly
- Transition mechanisms
 - NAT64 and DNS64
 - Tools like 6rd and other tunnelling options
 - ...



Dual Stack while you can


IPv6 Deployment Statistics

IPv6 Ripeness

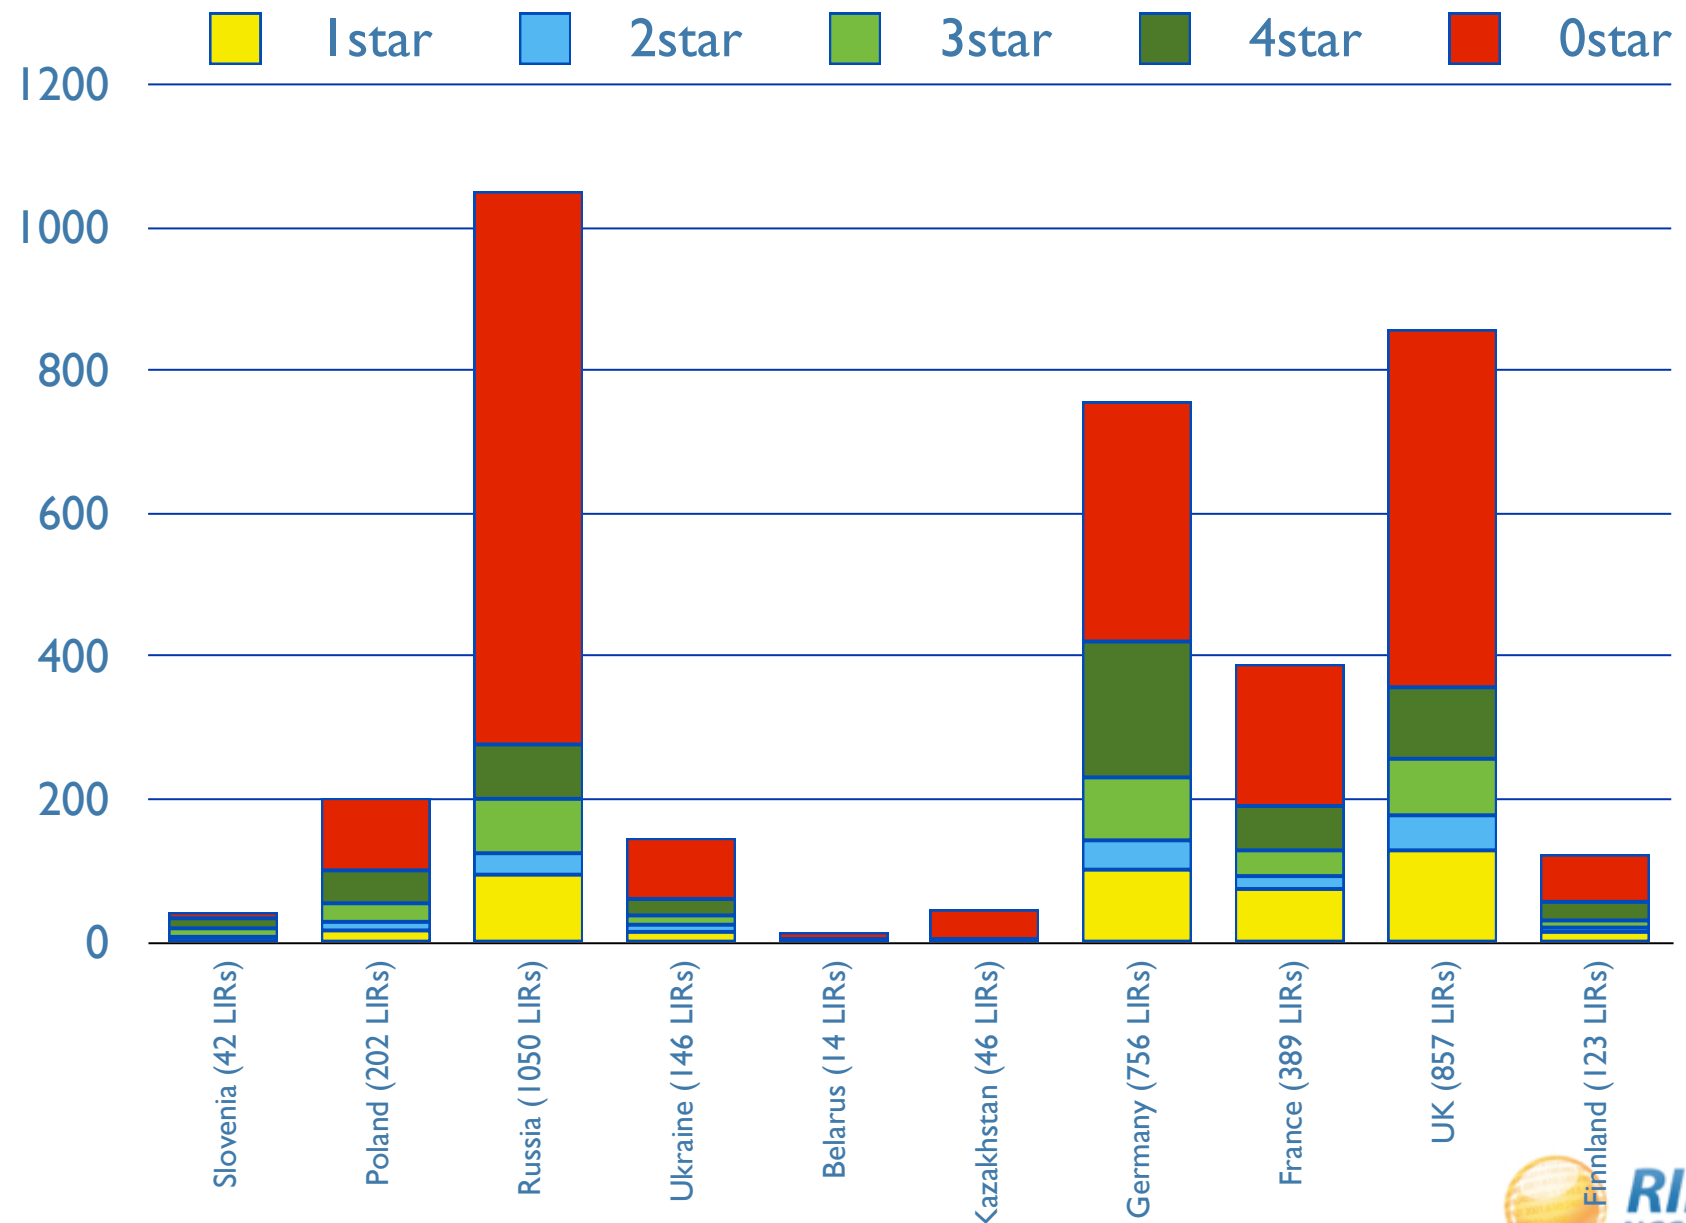
- Rating system:
 - One star if the LIR has an IPv6 allocation
 - Additional stars if:
 - IPv6 Prefix is announced on router
 - A route6 object is in the RIPE Database
 - Reverse DNS is set up
 - A list of all 4 star LIRs: <http://ripeness.ripe.net/>

IPv6 RIPEness: 7512 LIRs (31 May 2011)


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




IPv6 RIPEness – countries (31 May 2011)



IPv6 RIPEness – relative (31 May 2011)



IPv6 enabled ASes in global routing (31.05)



World IPv6 Day

- 8 June 2011
- Initiated by ISOC
- 0:00 GMT - 23:59 GMT
- Top 500 websites
 - Google
 - Facebook
 - Yahoo
 - and you?
- Great test opportunity



RIPE NCC and World IPv6 Day

- RIPE NCC Measurements

- Measuring connectivity to World IPv6 Day participants
- Testing connectivity and performance using TTM
- Monitoring performance of 6to4 versus native IPv6

- Coordinated events

- Amsterdam
- Moscow



- Live reports on <http://www.ripe.net/worldipv6day>

RIPE NCC @ World IPv6 Day

- All of our content and services over IPv6
- IPV6 Eyechart <http://ipv6eyechart.ripe.net/>
- IPv6 Day Measurements <http://v6day.ripe.net/>

Eye Chart for IPv6 Day

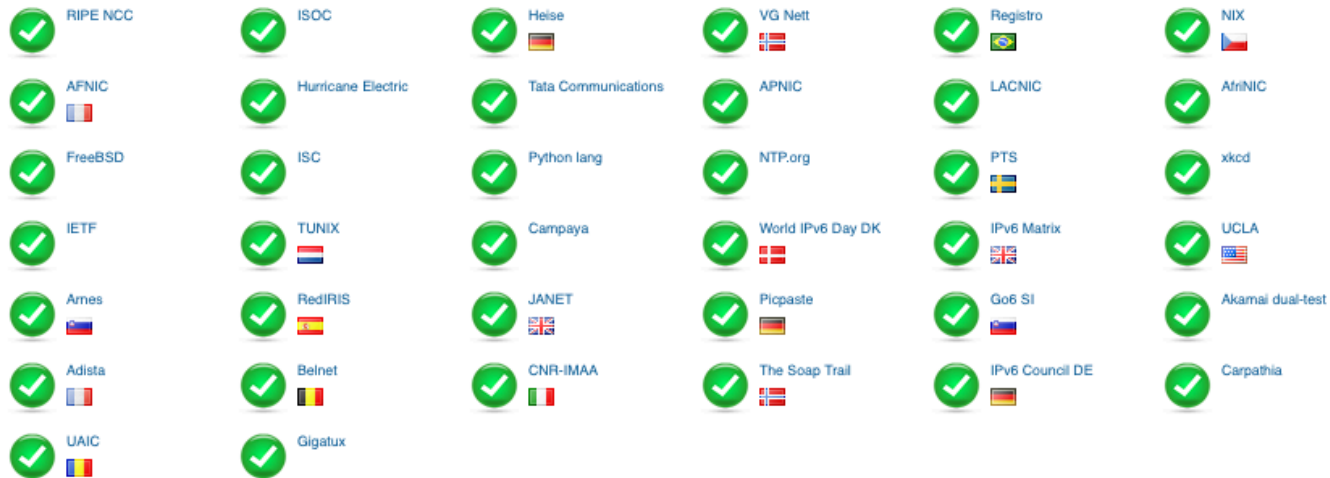


RIPE
NCC
RIPE NETWORK COORDINATION CENTRE


You are here: [Home](#) > [Data & Tools](#) > [Statistics & Analytics](#) > IPv6 Eye Chart

World IPv6 Day Connectivity Chart

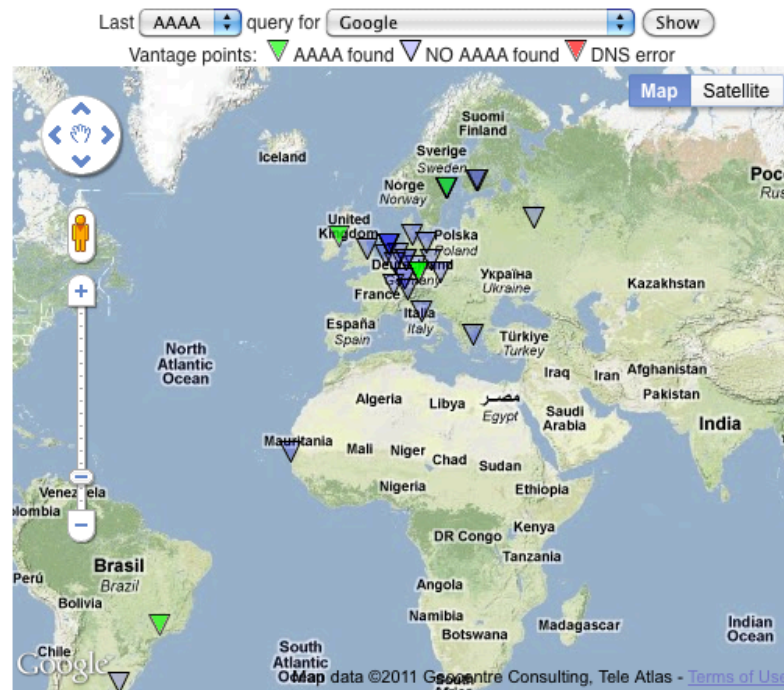

Dual stacked websites (dual stack = accessible over both IPv4 and IPv6)





World IPv6 day participants (currently IPv4 only, will be dual-stacked in 5 days)



Measurements for IPv6 Day



[See more DNS details...](#)

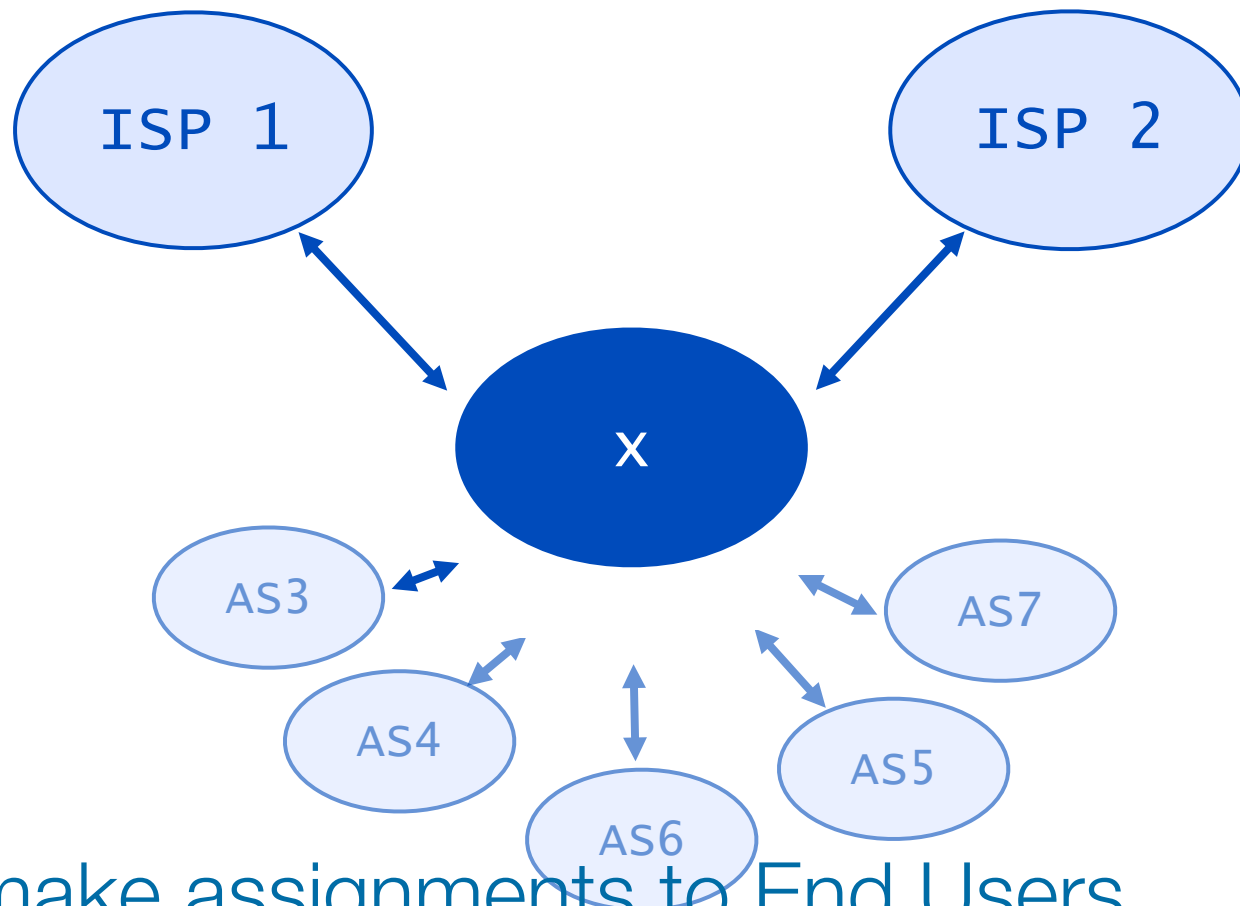


Multihomed BGP Routing Setup

To be or not to be an LIR


Type	Contract with:	Fee 2010 / 2011	Space	Member of RIPE NCC	Can influence RIPE policies
End User	LIR	PI = € 50 ASN = € 50	PI	No	Yes
LIR	RIPE NCC	Start-up fee + yearly fee XS = € 1300 + PI / ASN	PA allocations + PI	Yes	Yes
Direct Assignment User	RIPE NCC	Start-up fee + € 1300 + PI / ASN	PI	No	Yes

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 or DAU = PI + ASN



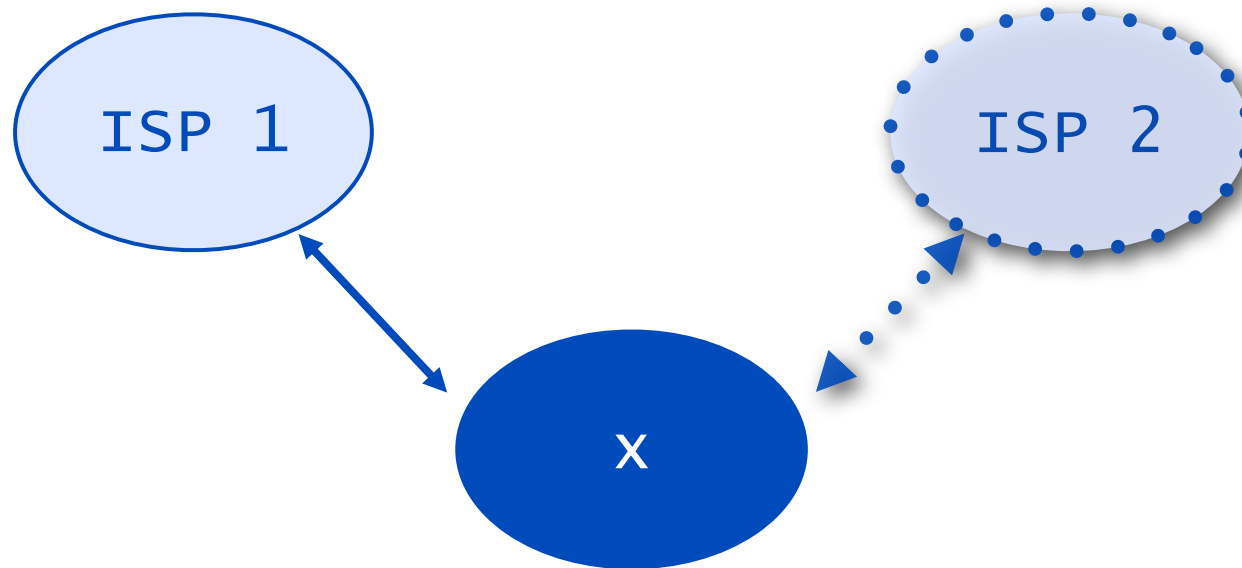
- Can NOT sub-assign further!!!
 - (in IPv4 can still use PI for xDSL, broadband...)

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

Scenario 5: PA assignment, multihomed



- Very rare and complicated
 - more specific PA prefix announced, to multiple ISPs
 - technically challenging, but “cheap”

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

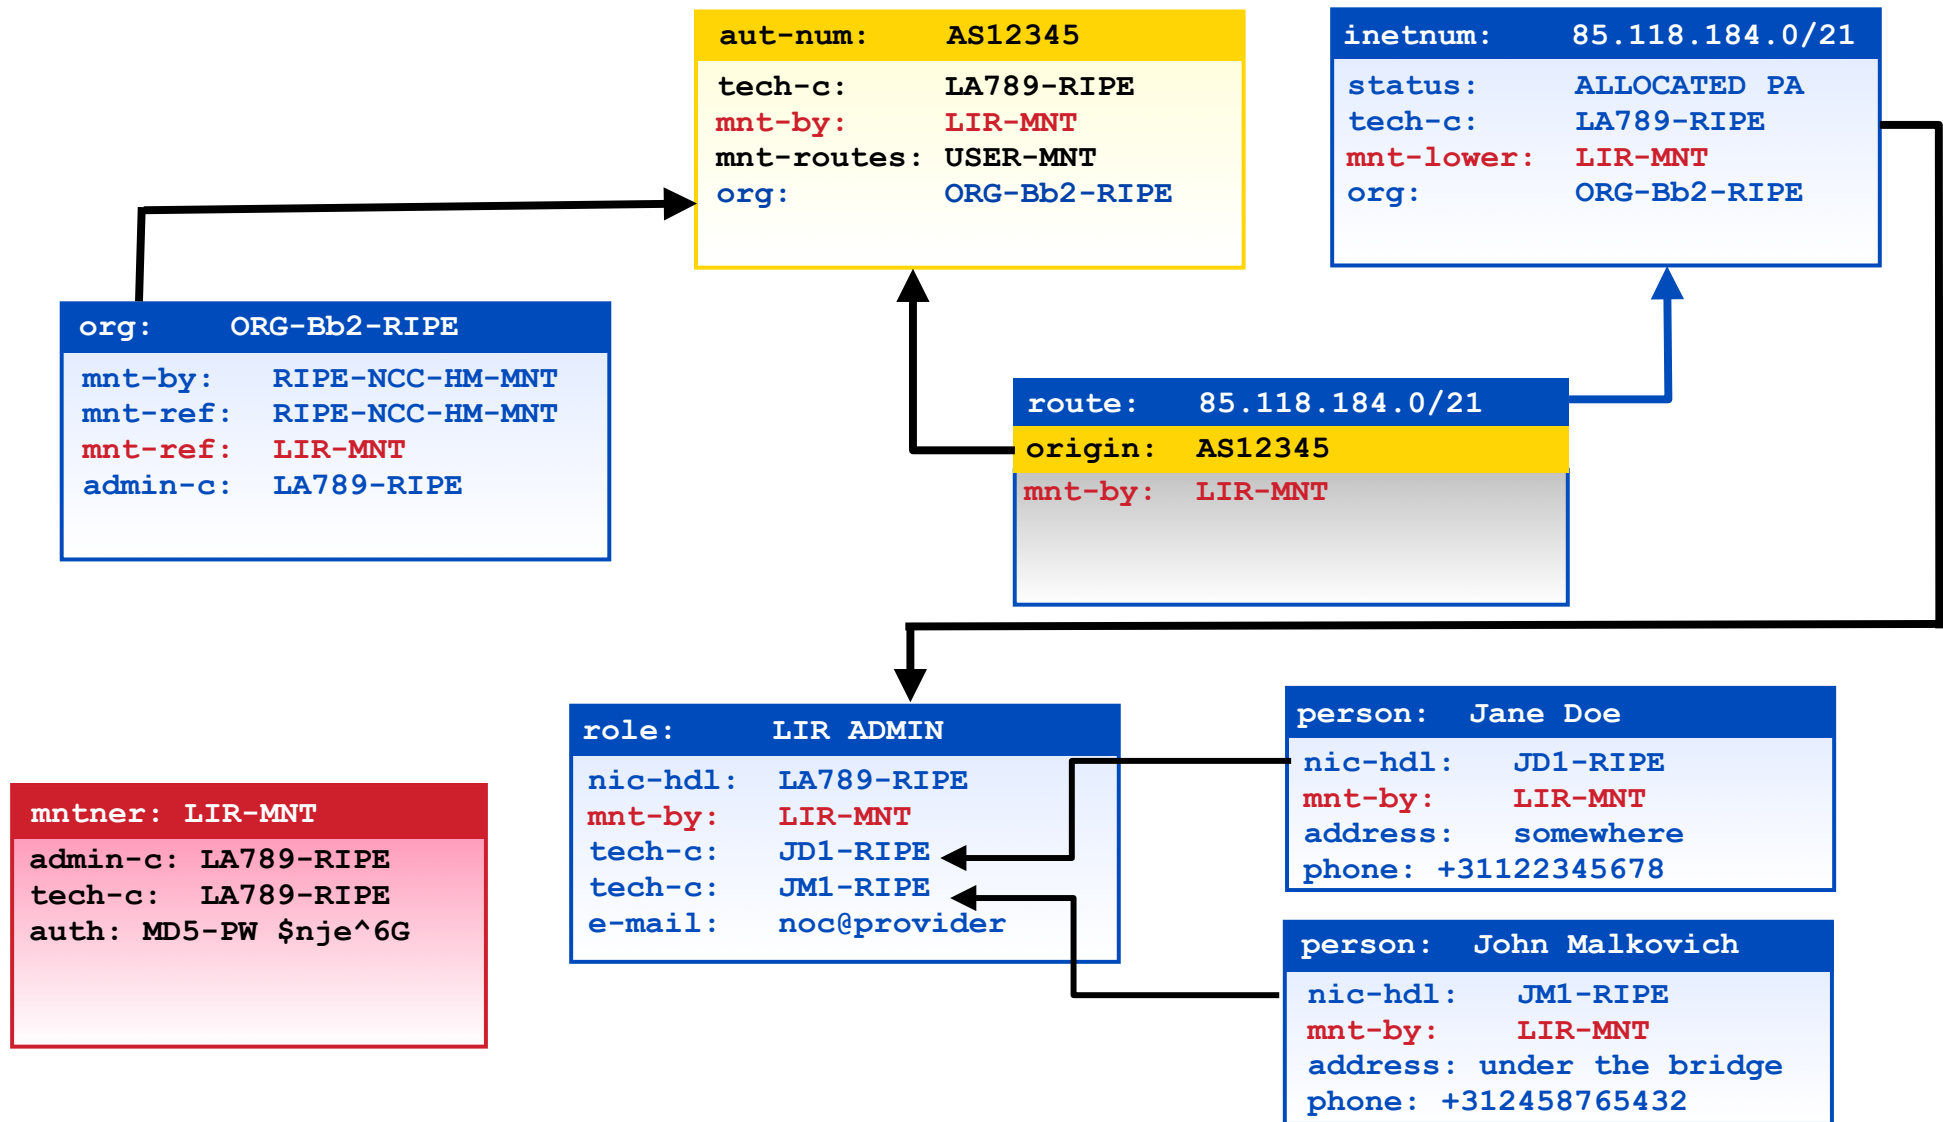
RIPE DB



Registration: RIPE Database

- Public Internet resources database
- All LIRs objects are there:
 - Address space: inetnum & inet6num
 - AS Number: aut-num
 - Contact details: person, role, organisation,
 - Strong protection: maintainer (key-cert, irt)

Connection between objects



Finding and changing objects

- Querying the RIPE Database
 - Command-line client
 - Web interface
 - Free text search (Glimpse)
 - & <http://lab.db.ripe.net/portal/free-text/search.htm>

- Updating = creating, modifying, deleting
 - Web, sync, email

Protection

```
mntner: LIR-MNT
```

```
auth: MD5-PW $1$o93Ux
```



```
person: John Smith
```

```
nic-hdl: JS1-RIPE
```



```
mnt-by: LIR-MNT
```

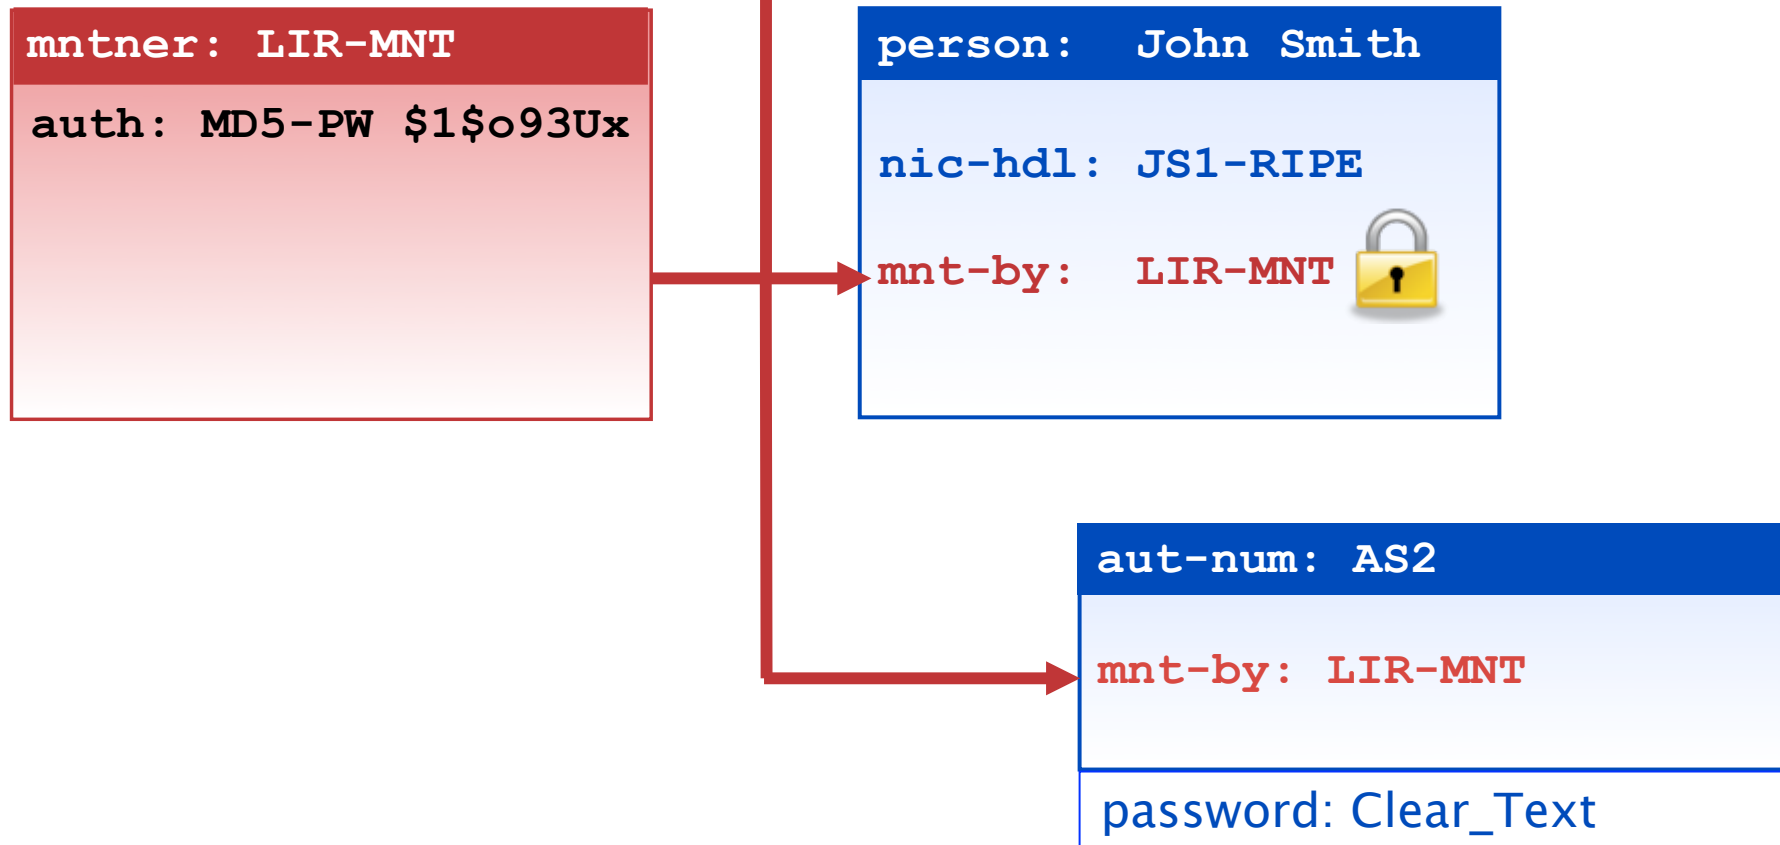
```
password: Clear_Text
```

Strong authentication

- Password (MD5-PW)
- Private key / public key
 - PGPKEY-<id> and key-cert object
 - x.509-<id> and key-cert object



Protection



Routing & 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 (Sawvis), 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
- Consistent information between neighbors
- Good housekeeping

85% match between BGP/RIS & RR

- According to the RIPE Labs article



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
 - those are two primary keys
 - only the holder of both address space and AS number can authorize creation of route[6] object

Authenticating a route6 object for an LIR

```
inet6num: 2001:db8::/32
```

```
status: ALLOCATED-BY-RIR  
mnt-by: RIPE-NCC-HM-MNT  
mnt-routes: LIR-MNT
```

```
aut-num: AS2
```

```
mnt-by: LIR-MNT
```

```
route6: 2001:db8::/32
```

```
origin: AS2  
mnt-by: LIR-MNT
```

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

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
```


RIPE NCC Resource Quality Assistance

- Address distribution - no claims about routability
 - but assistance in case of filtering issues:

[http://www.ripe.net/lir-services/resource-management/
ripe-ncc-resource-quality-assistance](http://www.ripe.net/lir-services/resource-management/ripe-ncc-resource-quality-assistance)

Questions?

training@ripe.net



The End!

Край

Y Diwedd

النهاية

Соңы

վերջ

Fí

Finis

Ende

Finvezh

Liðugt

Кінець

Konec

Kraj

Ěnn

Fund

پایان

Край

Lõpp

Beigas

Vége

Son

An Críoch

הסוף

Fine

Endir

Sfârșit

Fin

Τέλος

Einde

Конец

Slut

Slutt

დასასრული

Pabaiga

Fim

Amaia

Loppu

Tmíem

Koniec