



**RIPE  
NCC**

# **RPKI and IPv6**

**Tutorial | September 2015**

---

**Yerevan Regional Meeting**

<b>09:00 - 11:00</b>	<b>RPKI</b>
11:00 - 11:30	Break
<b>11:30 - 12:30</b>	<b>RPKI</b>
12:30 - 14:00	Lunch
<b>14:00 - 15:30</b>	<b>IPv6</b>
15:30 - 16:00	Break
<b>16:30 - 17:30</b>	<b>IPv6</b>
17:30	End



# Introduction to the Routing Registry

## Section 3



**RIPE**  
NCC

To be able to answer the question:

**Is that ASN authorised to originate  
that address range?**

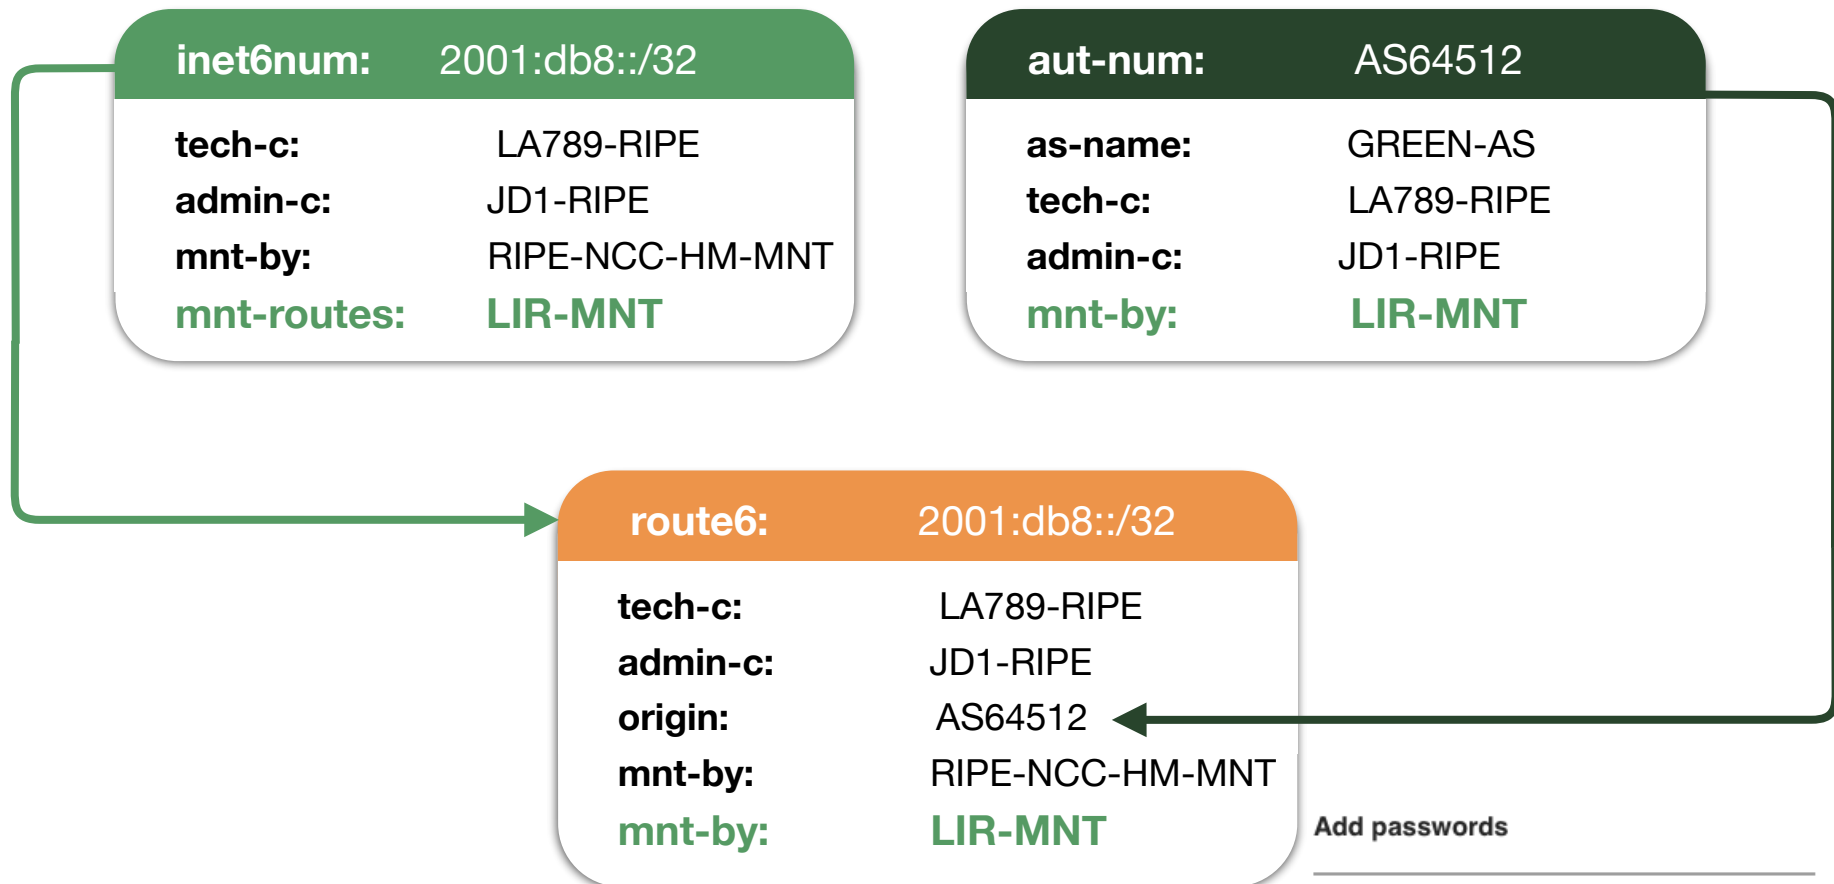


- **Number of public databases that contain routing policy information which mirror each other:**
  - RIPE, APNIC, RADB, JPIRR, Level3, ...
  - <http://www.irr.net>
- **RIPE NCC operates the RIPE Routing Registry**
  - Part of the RIPE Database
  - Part of the Internet Routing Registry

- **inetnum** = IPv4 address range
  - **inet6num** = IPv6 address range
  - **aut-num** = single AS number and routing policy
  - **route, route6** = glue between IP address range and an AS number announcing it
- 
- **person** = contact info for other objects
  - **role** = group of person objects
  - **maintainer** = protects all other objects

# Registering Routes

7



Add passwords

Session passwords

0 stored password(s)



12lir



- What prefixes do you announce?
- Who are your neighbours?
  - Peers, transits and customers
- Which prefixes do you accept from them?
- What are your preferences?

<b>aut-num:</b>	AS64512
<b>descr:</b>	RIPE NCC Training Services
<b>as-name:</b>	GREEN-AS
<b>tech-c:</b>	LA789-RIPE
<b>admin-c:</b>	JD1-RIPE
<b>import:</b>	from AS64444 accept ANY
<b>import:</b>	from AS64488 accept ANY
<b>export:</b>	to AS64444 announce AS64512
<b>export:</b>	to AS64488 announce AS64512
<b>mnt-by:</b>	LIR-MNT
<b>source:</b>	RIPE

- **Some transit providers and IXPs (Internet Exchange Points) require it**
  - They build their filters based on the routing registry
- **Contributes to routing security and stability**
  - Let people know about your intentions
- **Can help in troubleshooting**
  - Which parties are involved?

- **Close relation between registry information and routing policy**
  - The holder of the resources knows how they should be routed
- **The Routing Policy Specification Language (RPSL) originates from a RIPE Document**
  - Shares attributes with the RIPE Database

- **Accuracy and completeness**
- **Not every Routing Registry is linked directly to an Internet Registry**
  - Offline verification of the resource holder is needed
- **Different authorisation methods**
- **Mirrors are not always up to date**





# RPSL

## Section 4



**RIPE**  
NCC


- **A routing policy describes how a network works**
  - Who do you connect with
  - Which prefixes or routes do you announce
  - Which routes do you accept from others
  - What are your preferences
- **In your router, this is your BGP configuration**
  - neighbours
  - route-maps
  - localpref

- Language used by the IRRs
- Not vendor-specific
- Documented in RFC 2622
  - and RFC 2650 “Using RPSL in practice”
- Can be translated into router configuration

- **route or route6 object**
  - Connects a prefix to an origin AS
- **aut-num object**
  - Registration record of an AS Number
  - Contains the routing policy
- **Sets**
  - Objects can be grouped in sets, i.e. as-set, route-set
- **Keywords**
  - “ANY” matches every route

- **AS Numbers are written as ASxxx**
- **Prefixes are written in CIDR notation**
  - 193.0.4.0/24
- **Any value can be replaced by a list of values of the same type**
  - AS1 can be replaced by “AS1 AS2 AS3”
- **You can reference a set instead of a value**
  - “...announce AS1” or “...announce as-myname”

- **You can document your routing policy in your aut-num object in the RIPE Database:**
  - Import lines describe what routes you accept from a neighbour and what you do with them
  - Export lines describe which routes you announce to your neighbour



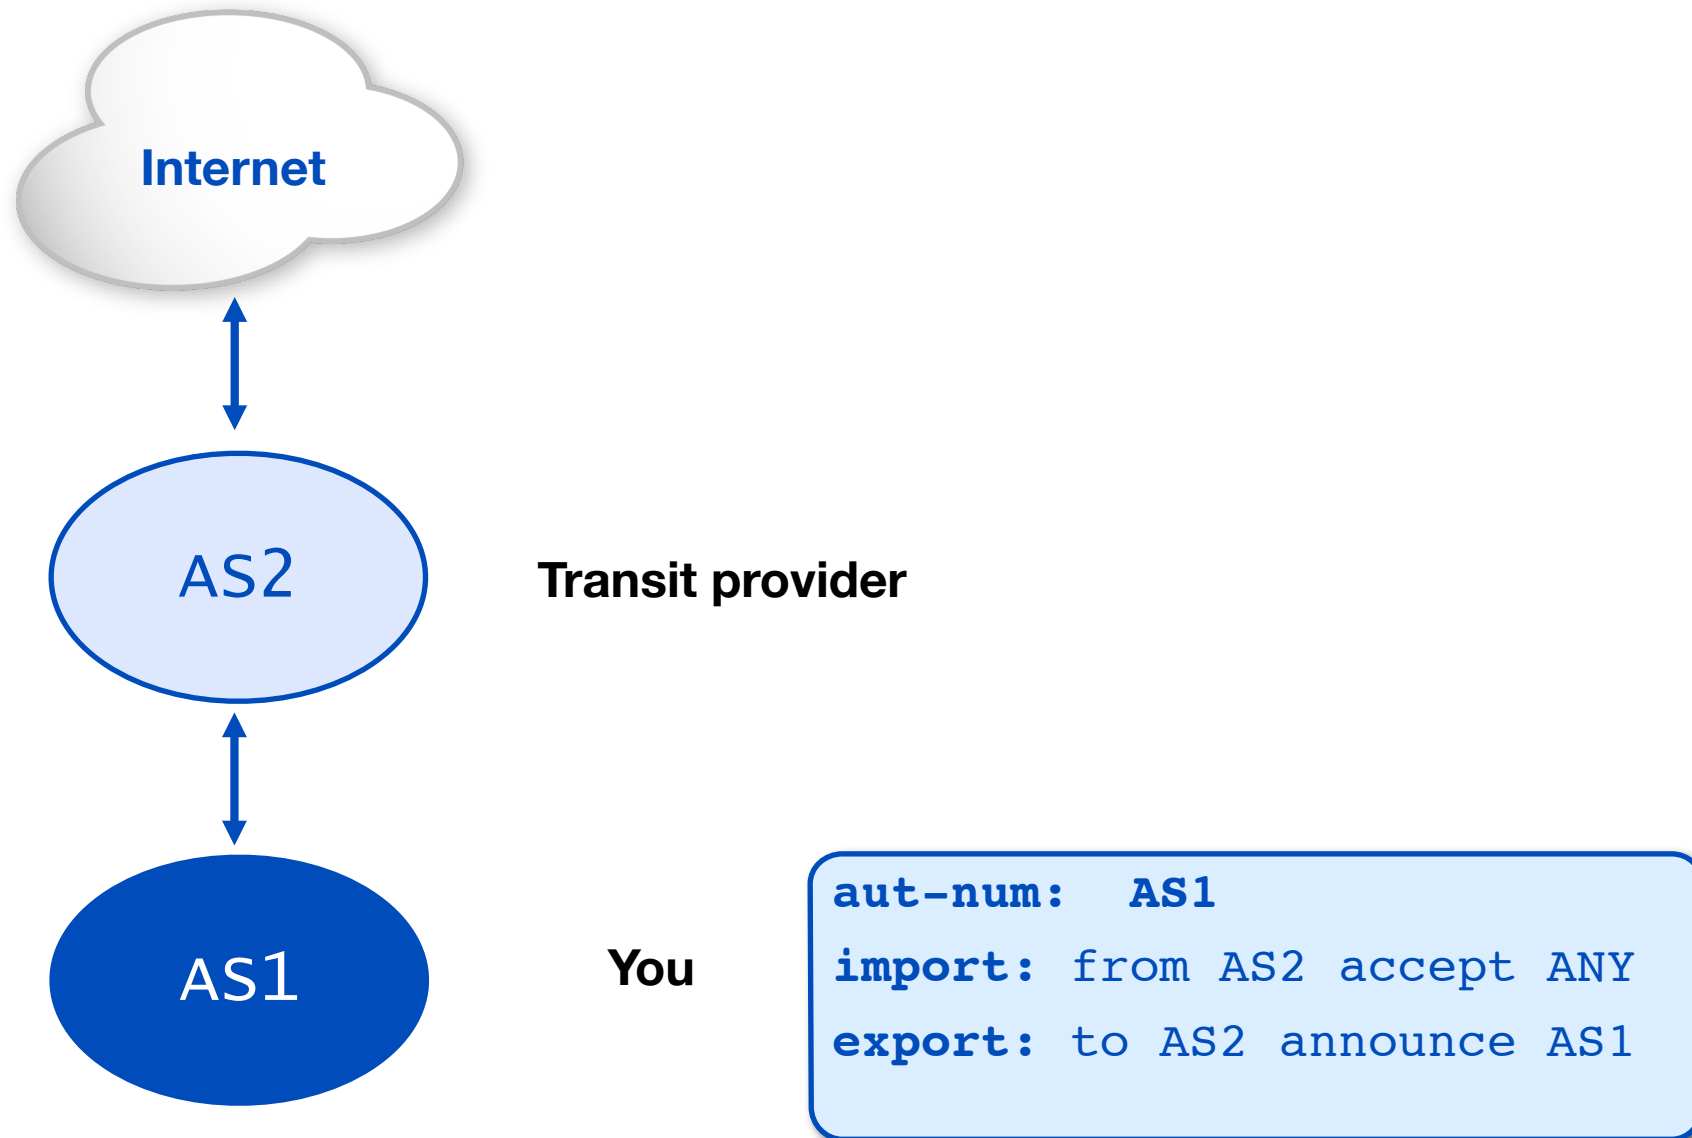
```
aut-num: AS1
```

```
import: from AS2 accept AS2
```

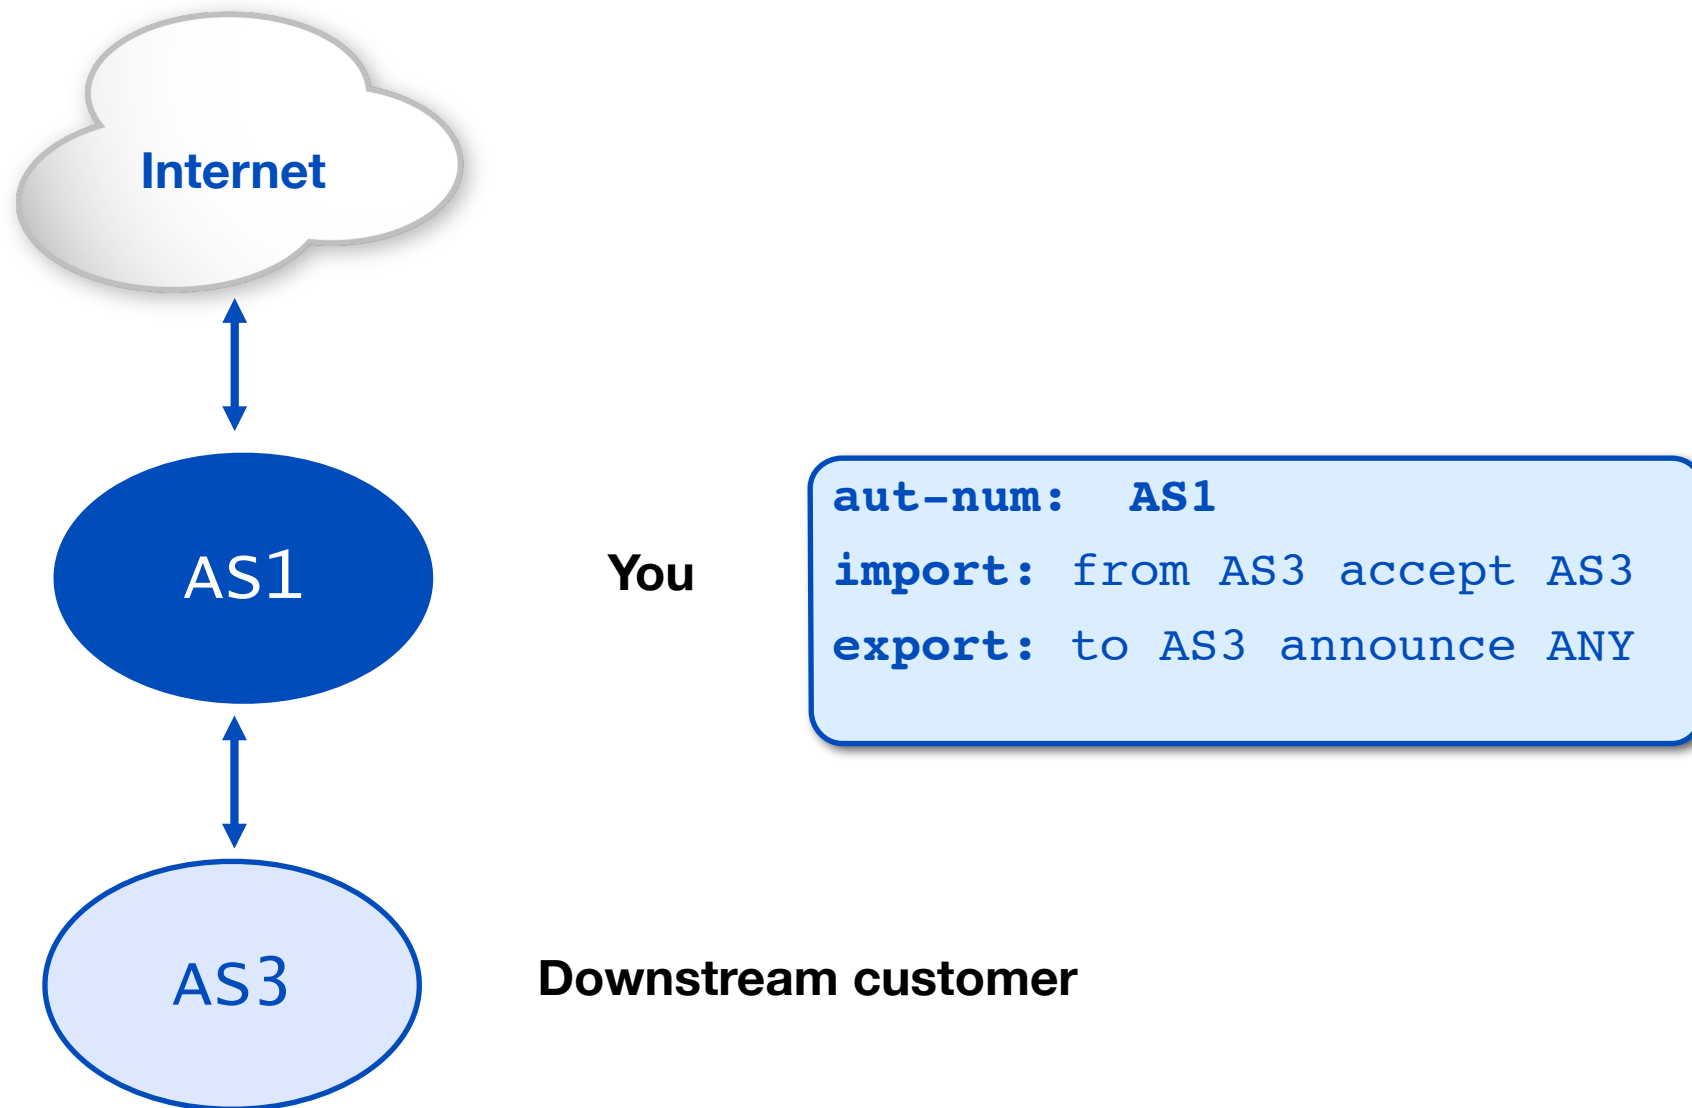
```
export: to AS2 announce AS1
```

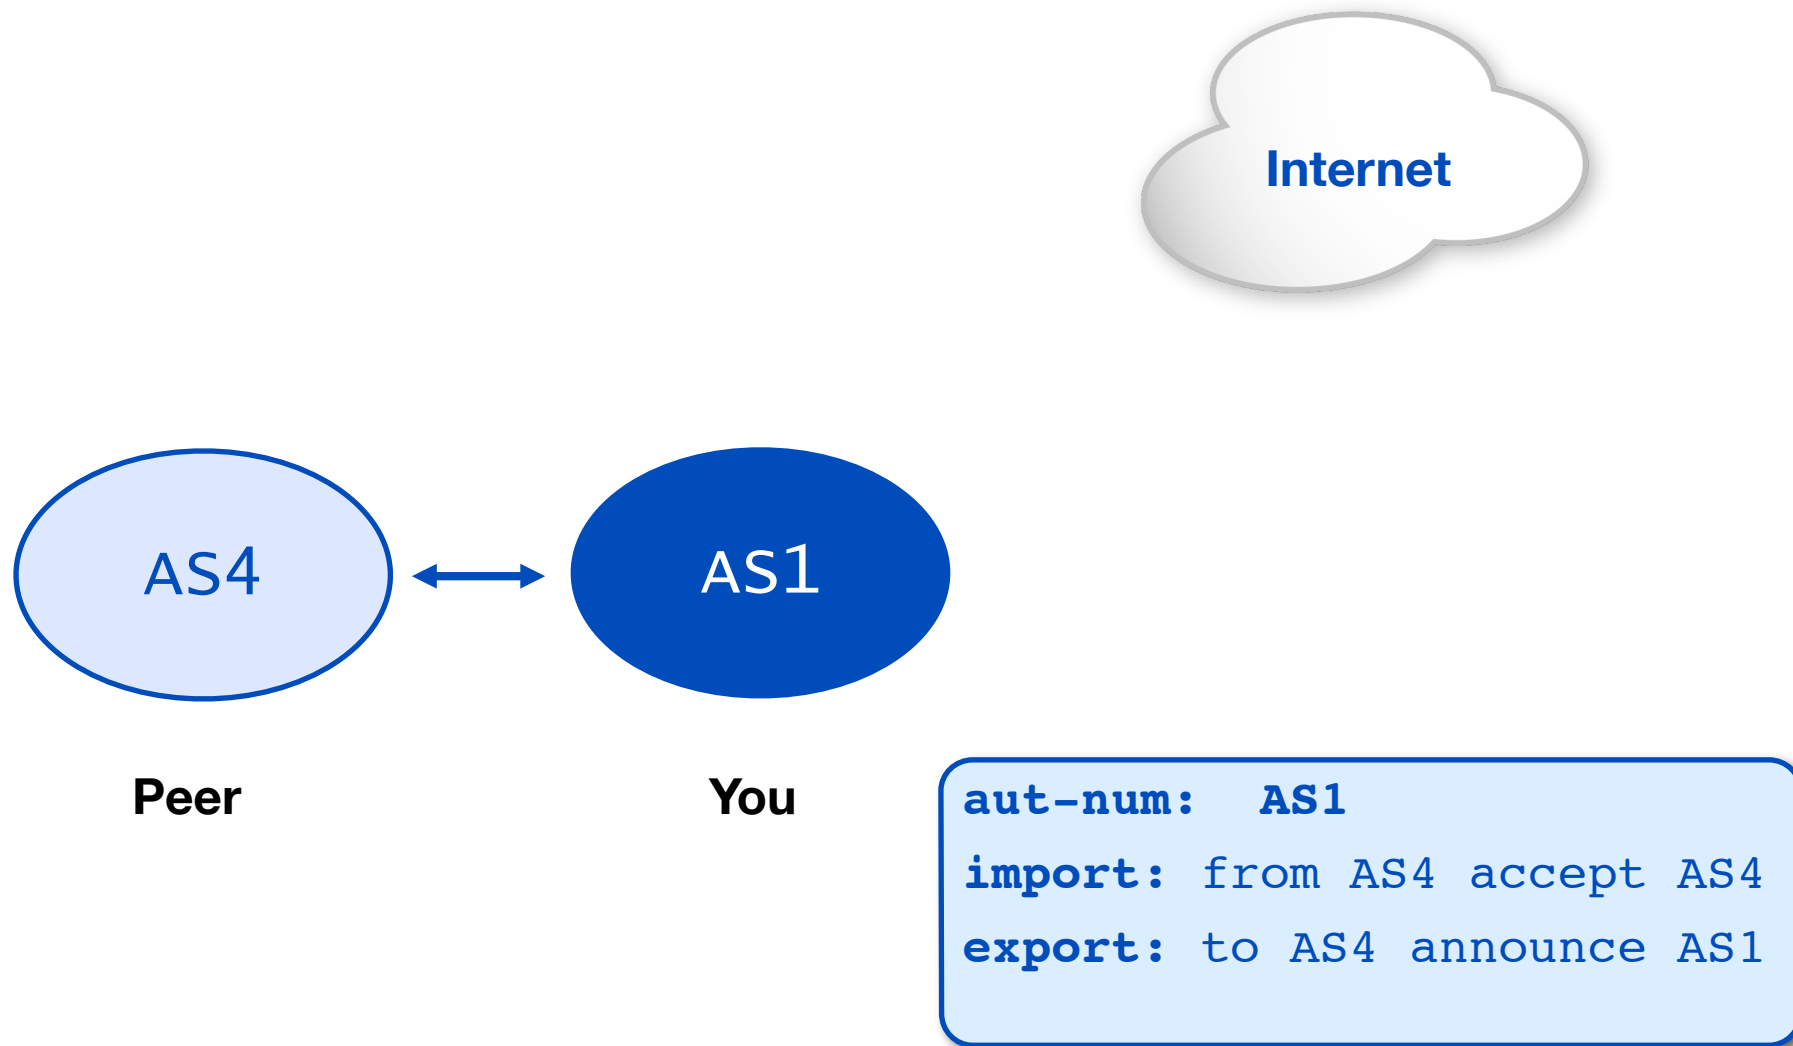
AS1 accepting those prefixes **from** AS2 that originate in AS2 so that the **outbound** traffic for AS2 can go **towards** the AS2

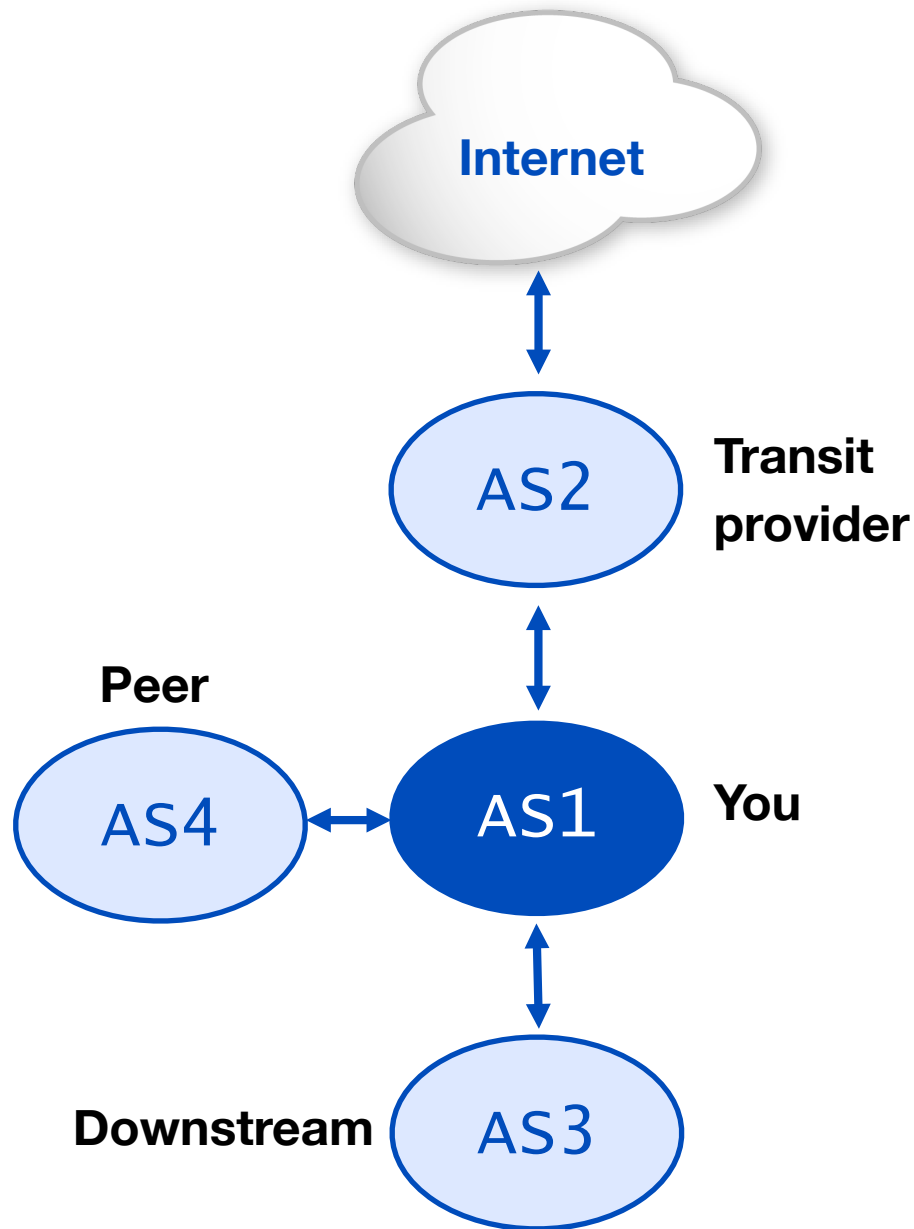
AS1 announcing prefixes (originating in AS1) **to** AS2, so that the **incoming** traffic for AS1 can flow **away** from the AS2











```
aut-num: AS1
import: from AS2 accept ANY
export: to AS2 announce AS1 AS3
import: from AS3 accept AS3
export: to AS3 announce ANY
import: from AS4 accept AS4
export: to AS4 announce AS1 AS3
```

- **Automation relies on the IRR being complete**
  - Not all resources are registered in an IRR
  - Not all information is correct
- **Small mistakes can have a big impact**
- **Check your output before using it**
  - Be prepared to make manual overrides
- **Help others by documenting your policy**

- You can compare the Routing Registry and the Internet routing table using <http://stat.ripe.net>

AS Routing Consistency (AS3333)

Prefixes Imports Exports

Show 10 entries Search:

Prefix	In RIS	RIPE IRR	Other IRRs
193.0.0.0/21	yes	yes	no
193.0.10.0/23	yes	yes	no
193.0.12.0/23	yes	yes	no
193.0.18.0/23	yes	yes	no
193.0.20.0/23	yes	yes	no
193.0.22.0/23	yes	yes	no
2001:67c:2e8::/48	yes	yes	no

Showing 1 to 7 of 7 entries

Showing results for AS3333 as of 2014-12-05 00:00:00 UTC

source data embed code permalink info





# Introduction the the RPKI

## Section 7



**RIPE**  
NCC

To be able to answer the question:

**Is that ASN authorised to originate  
that address range?**



- **Why yet another system?**
  - Lots of Routing Registries
  - Not all mirroring each other
  - Different levels of trustworthiness and authentication
- **RPKI replaces IRR or lives side by side?**
  - Side by side: different advantages
    - Security, almost real time, simple interface: RPKI
    - More info in: IRR

- **Useable toolset**
  - No installation required
  - Easy to configure manual overrides
- **Tight integration with routers**
  - Supported routers have awareness of RPKI validity states
- **Stepping stone for AS-Path Validation**
  - Prevent Attacks on BGP



# **RPKI**

## **Setting it up:**

## **The announcers side**

### **Section 8**

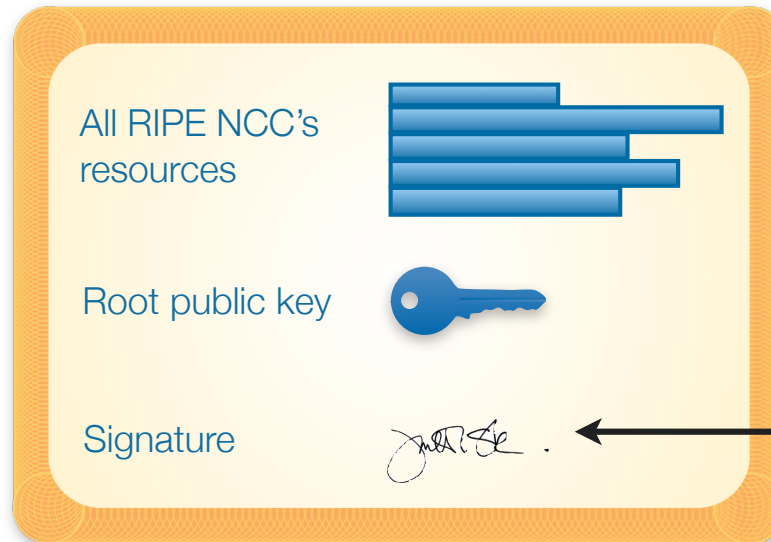


**RIPE**  
**NCC**

- **RIPE NCC issues digital certificates**
  - To LIRs
  - To all resource holders
  - Upon request
- **Certificate lists all resources held by the member**

- **RIPE NCC holds self-signed root certificate for all resources they have in the registry**
  - Signed by the root's private key
- **The root certificate is used to sign all certificates for members listing their resources**
  - Signed by the root's private key

## RIPE NCC's Root Certificate

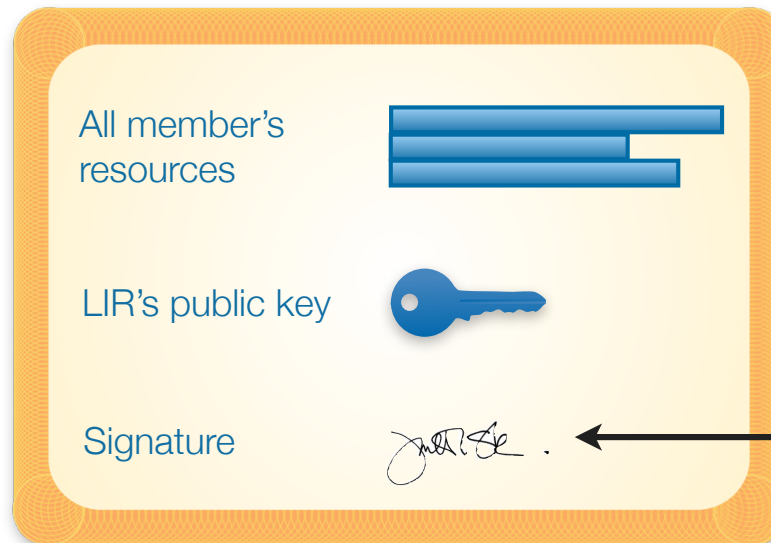


Root's (RIPE NCC)  
private key



sign

## LIR's Certificate



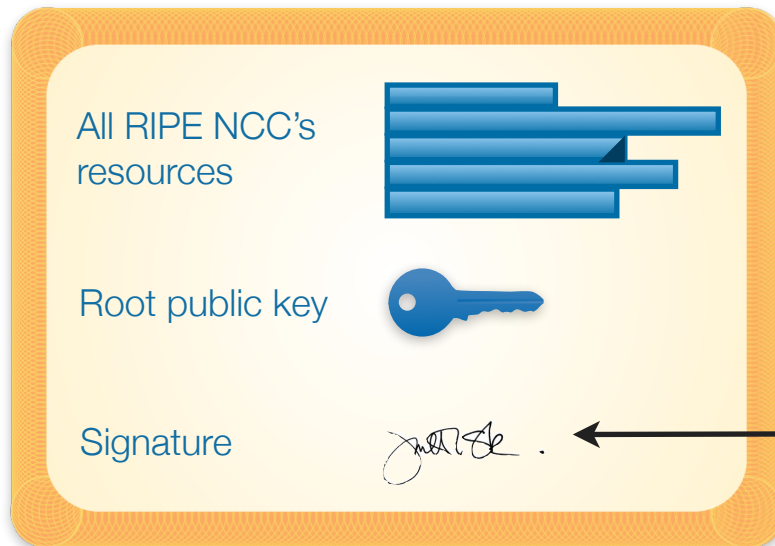
LIR's  
private key



sign

- **LIRs can use their certificate to create a ROA for each of their resources (IP address ranges)**
  - Signed by the root's private key
- **ROA states**
  - Address range
  - Which AS this is announced from (freely chosen)
  - Maximum length (freely chosen)
- **You can have multiple ROAs for an IP range**
- **ROAs can overlap**

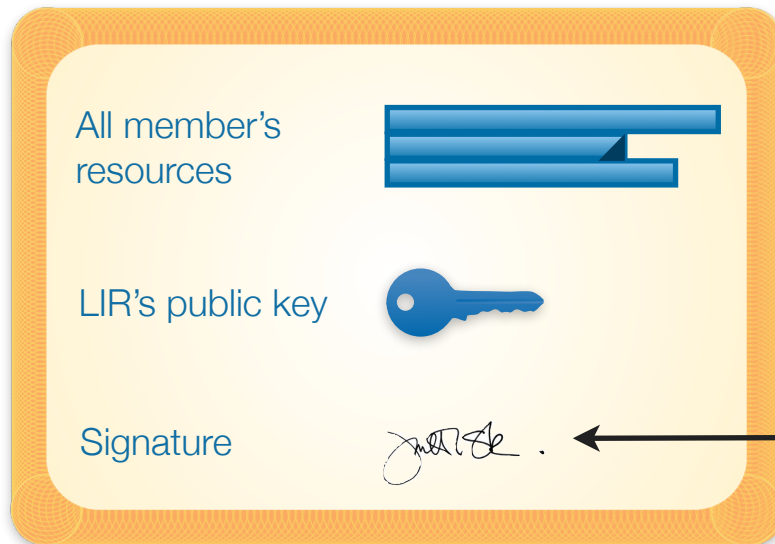
## RIPE NCC's Root Certificate



Root's (RIPE NCC)  
private key

sign



## LIR's Certificate



LIR's  
private key

sign

## ROA

IP Range	
AS Number	AS123
Max Length	/24
Signature	



## ROA

193.0.24.0/21

AS2121

Max Length: \_

193.0.24.0/21

193.0.24.0/22

193.0.30.0/23

## ROA

193.0.24.0/21

AS2121

Max Length: /23

193.0.24.0/21 ✓

193.0.24.0/22 ✓

193.0.28.0/22 ✓

193.0.24.0/23 ✓

193.0.26.0/23 ✓

193.0.28.0/23 ✓

193.0.30.0/23 ✓

# ROA (Route Origin Authorisation) Example

39

ROA

193.0.24.0/21

AS2121

Max Length: \_

193.0.24.0/21

193.0.24.0/22

193.0.28.0/22

ROA

193.0.24.0/23

AS2121

Max Length: /24

/23

/23

/23

/23

/24

/24

/24

/24

/24

/24

/24

/24

ROA

193.0.30.0/23

AS2121

Max Length: \_

- **RIPE NCC maintains a Certificate Repository containing**
  - All the certificates
  - All the public keys
  - All the ROAs



# **RPKI**

## **Setting it up: Certification**

### **Section 9**



**RIPE**  
**NCC**

**Edit Contact**

First name	Last name	Email	Status ?
Andrzej	Wolski	awolski@ripe.net	Active

**Comments**

**What this user is entitled to do:**

- ☒ Manage contacts and access all RIPE NCC services
- ☐ Access all RIPE NCC services
- ☐ Make payments and manage billing information

## 🌟 Create a Certificate Authority for zz.example

### RIPE NCC Certification Service Terms and Conditions

#### Introduction

This document will stipulate the Terms and Conditions for the RIPE NCC Certification Service. The RIPE NCC Certification Service is based on Internet Engineering Task Force (IETF) standards, in particular RFC3647, "Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework", RFC3779, "X.509 Extensions for IP Addresses and AS Identifiers", and the "Certificate Policy (CP) for the Resource PKI (RPKI)".

#### Article 1 – Definitions

In the Terms and Conditions, the following terms shall be understood to have the meanings assigned to them below:

**RIPE NCC** – Réseaux IP Européens Network Coordination Centre, a membership association under Dutch law, operating from its registered office in Amsterdam, the Netherlands.

**Certificate Holder** – A natural person or a legal entity that has entered into an agreement regarding the registration of their resources either with a sponsoring LIR or with the

By clicking on 'I accept' below you confirm that that you have read, understood and agree to the [RIPE NCC Certification Service Terms and Conditions](#).


☒ I accept. Create my Certificate Authority


<http://localcert.ripe.net>


RPKI Dashboard


9 CERTIFIED RESOURCES


NO ALERT EMAIL CONFIGURED


 **41** BGP Announcements


 **4** Valid

 **1** Invalid

 **36** Unknown

 **4** ROAs

 **3** OK


 **1** Causing problems


BGP Announcements


Route Origin Authorisations (ROAs)

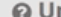
History







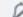

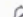

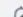

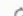

Search...

 Create ROAs for selected BGP Announcements

 Valid

 Invalid

 Unknown

<input type="checkbox"/>	Origin AS	Prefix	Current Status	
<input type="checkbox"/>	AS12654	2001:7fb:fe01::/48	UNKNOWN	 
<input type="checkbox"/>	AS12654	2001:7fb:fe0c::/48	UNKNOWN	 
<input type="checkbox"/>	AS12654	2001:7fb:fe0f::/48	UNKNOWN	 
<input type="checkbox"/>	AS12654	2001:7fb:ff00::/48	UNKNOWN	 
<input type="checkbox"/>	AS12654	2001:7fb:ff01::/48	UNKNOWN	 
<input type="checkbox"/>	AS12654	2001:7fb:ff02::/48	UNKNOWN	 
<input type="checkbox"/>	AS12654	2001:7fb:ff03::/48	UNKNOWN	 







# **RPKI**

## **Using it:**

## **Relying Party's side**

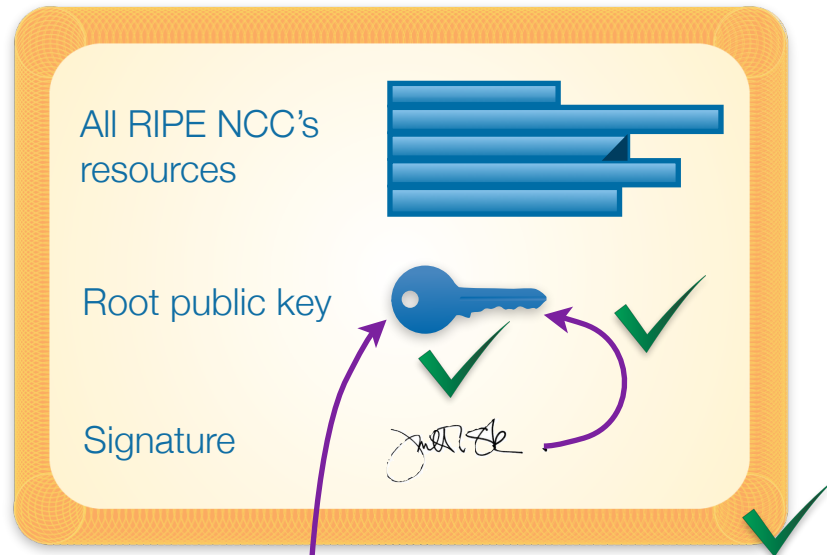
### **Section 10**



**RIPE**  
**NCC**

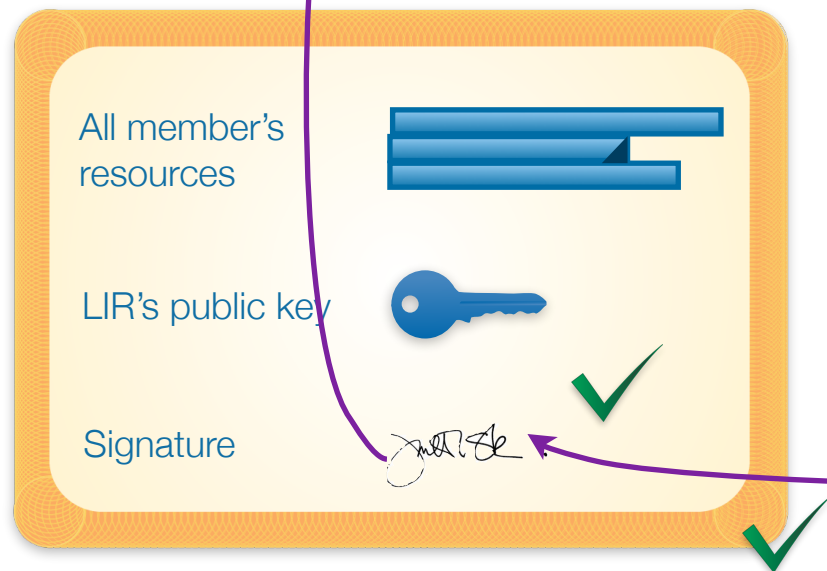
- The validator of the client can access RIPE NCC's Repository with all the certificates, public keys, ROAs
- It downloads everything and then performs validation, checking whether the certificates and ROAs are valid. Then it constructs a list of valid ROAs, which is its “validated cache”

## RIPE NCC's Root Certificate





Root's (RIPE NCC)  
private key

## LIR's Certificate




LIR's  
private key

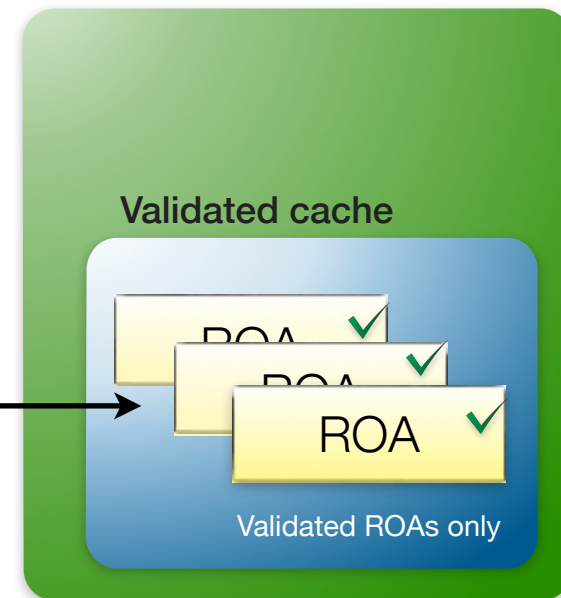
## ROA

IP Range	
AS Number	AS123
Max Length	/24
Signature	

## RIPE NCC's Repository



## Validator



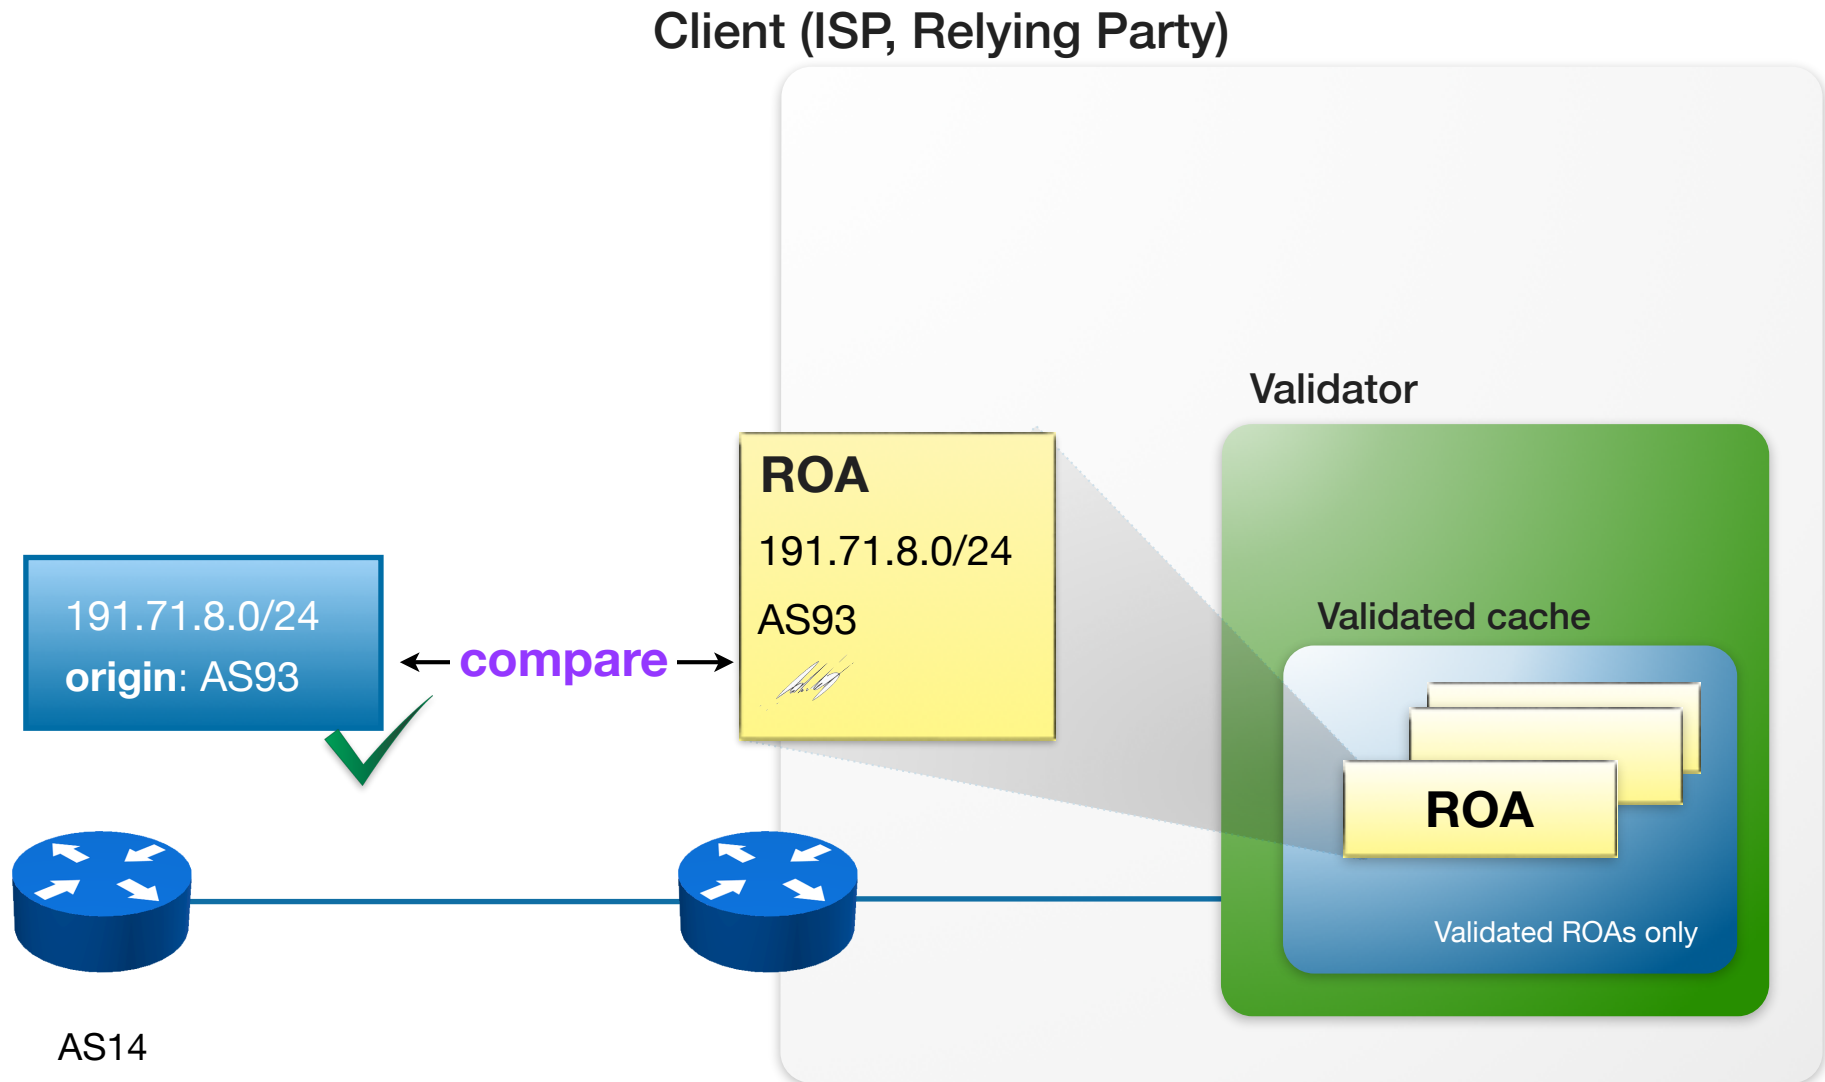
at the Relying Party's site

- Invalid ROAs are simply not included in the list of validate ROAs when the validator of the client computes them
- Reasons for a ROA to be invalid
  - The signing certificate or key pair has expired or has been revoked
  - It does not validate back to a configured trust anchor
    - The LIR's resource has been returned to the RIPE NCC

- **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 Relying Party's router can connect and download the cache from the validator**
- Router can then compare any BGP announcements to the list of valid ROAs in the validated cache





- **valid**
  - There is a ROA in the validated cache that matches the BGP announcement of the peer, size matches too
- **unknown**
  - There is no ROA for that prefix in the cache
- **invalid**
  - There is a ROA for the prefix, but for a different AS
  - The size doesn't match

- **Invalid ROA**
  - The ROA in the repository cannot be validated by the client (ISP) so it is not included in the validated cache
- **Invalid BGP announcement**
  - There is a ROA in validated cache for that prefix but for a different AS.
  - Or the max length doesn't match.
  - Remember if no ROA in the cache then announcement is “unknown”

- **As an announcer/LIR**
  - You choose if you want certification
  - You choose if you want to create ROAs
  - You choose AS, max length
- **As a Relying Party**
  - You can choose if you use the validator
  - You can override the lists of valid ROAs in the cache, adding or removing valid ROAs locally
  - You can choose to make any routing decisions based on the results of the BGP Verification (valid/invalid/unknown)



# **RPKI** **Using it:** **RIPE NCC Validator**

**Demo**



**RIPE**  
**NCC**

- <http://www.ripe.net/certification> -> Tools and Resources

## RIPE NCC RPKI Validator 2.17 (Updated 3 July 2014)

This application allows operators to download and validate the global RPKI data set for use in their [BGP decision making process](#) and [router configuration](#).

System requirements: a UNIX-like OS, Java 7, rsync and 2GB free memory.

To install, simply unpack the archive and run "rpki-validator.sh" from the base folder.

For more information, [view the release notes](#). You can also [contribute to the project on GitHub](#).

[Download Now](#)

- **Requires: UNIX, Java 7, rsync and 2GB free memory**
- **No Installation required**
  - Unzip the package
  - Run the program: `rpki-validator.sh start`
- **Interface available on localhost port 8080**

The screenshot shows a web browser window titled "RPKI Validator - Quick Overview of BGP Origin Validation". The address bar shows "http://127.0.0.1:8080/". The browser's search bar contains "Google". The website has a dark navigation bar with the following links: Home, Trust Anchors, ROAs, Ignore Filters, Whitelist, BGP Preview, Export, and Router Sessions. A settings icon is also present.

## Quick Overview of BGP Origin Validation

A diagram illustrates the validation process flow:

```
graph LR;
    TA[Trust Anchors] --> ROAs[ROAs];
    ROAs --> IF[Ignore Filters];
    IF --> W[Whitelist];
    W --> R[Router];
```

Below the diagram, a text box explains the role of Trust Anchors:

Trust Anchors are the entry points used for validation in any Public Key Infrastructure (PKI) system. This validator is intended for the validation of Resource PKI (RPKI) systems. It is pre-configured with Trust Anchors for all the RIRs who are running such systems now.

If you would like to add or change the Trust Anchors that are used by this validator, please see the README.txt file for details.

A "Feedback" button is located on the left side of the page.

At the bottom, the RIPE NCC logo is displayed next to the copyright notice: "Copyright ©2009-2012 the Réseaux IP Européens Network Coordination Centre RIPE NCC. All rights restricted. Version 2.0.4".

RPKI Validator – Configured Trust Anchors


http://127.0.0.1:8080/trust-anchors

RPKI Validator Home Trust Anchors ROAs Ignore Filters Whitelist BGP Preview Export Router Sessions

## Configured Trust Anchors

Trust anchor	Processed Items	Expires in	Last updated	Next update in	Update all
APNIC RPKI Root	1356 0 0	4 years and 2 months	7 minutes ago	3 hours	update
ARIN Test Lab	88 88 0	1 year and 2 months	8 minutes ago	3 hours	update
AfriNIC RPKI Root	80 0 1	4 years and 7 months	8 minutes ago	3 hours	update
LACNIC RPKI Root	216 0 0	10 months and 3 weeks	8 minutes ago	3 hours	update
RIPE NCC RPKI Root	3570 0 0	4 years and 9 months	7 minutes ago	3 hours	update

Feedback

 Copyright ©2009-2012 the Réseaux IP Européens Network Coordination Centre RIPE NCC. All rights restricted. Version 2.0.4



RPKI Validator – Validated ROAs

http://127.0.0.1:8080/roas

RPKI Validator Home Trust Anchors **ROAs** Ignore Filters Whitelist BGP Preview Export Router Sessions

## Validated ROAs

Validated ROAs from APNIC RPKI Root, ARIN Test Lab, AfriNIC RPKI Root, LACNIC RPKI Root, RIPE NCC RPKI Root.

Show 10 entries

Search:

ASN	Prefix	Maximum Length	Trust Anchor
1	10.0.1.0/24	24	ARIN Test Lab
1	192.168.1.0/24	24	ARIN Test Lab
1	61.45.250.0/23	23	APNIC RPKI Root
1	61.45.250.0/23	23	APNIC RPKI Root
21	10.4.0.0/16	16	ARIN Test Lab
22	10.255.1.0/24	24	ARIN Test Lab
42	2001:678:3::/48	48	RIPE NCC RPKI Root
42	194.0.17.0/24	24	RIPE NCC RPKI Root
174	89.207.56.0/21	21	RIPE NCC RPKI Root

**Prefix**

Insert the prefix and click “Add”

This locally creates a valid (but fake) ROA

Current filters		
Show <input type="text" value="10"/> entries	Search: <input type="text"/>	
Prefix	Filtered ROA prefixes	
193.0.24.0/21	1 prefix(es)	<input type="button" value="delete"/>
Showing 1 to 1 of 1 entries		

- The validator downloads a copy of the RIS
  - Allows you to get a hint of what would happen
  - RIS view might be different from your routing table

This page provides a **preview** of the likely rpk validity states your routers will associate with BGP announcements. This preview is based on:

- The [RIPE NCC Route Collector information](#) that was last updated 3 hours and 25 minutes ago.
- BGP announcements that are seen by 5 or more peers.
- Validation rules defined in the [IETF standard](#).
- The validated ROAs found by this validator after applying your filters and additional whitelist entries.

Please note that the actual validation of announcements is done in your routers and that the announcements that your routers see may differ from the announcements used here.

Feedback

Show 10 entries

Search:

ASN	Prefix	Validity
1	192.240.141.0/24	UNKNOWN
1	199.248.203.0/24	UNKNOWN
2	128.4.0.0/16	UNKNOWN
3	18.0.0.0/8	UNKNOWN
3	117.103.68.0/24	UNKNOWN
3	117.103.69.0/24	UNKNOWN
3	117.103.70.0/24	UNKNOWN

RPKI Validator – BGP Preview

http://127.0.0.1:8080/bgp-preview

RPKI Validator Home Trust Anchors ROAs Ignore Filters Whitelist **BGP Preview** Export Router Sessions

Show 10 entries Search: invalid

ASN	Prefix	Validity
14	2001:468:904::/48	
27	2001:468:c01::/48	
57	2001:468:1900::/40	
81	2001:468:1500::/40	
102	2001:468:c13::/48	INVALID
719	193.209.25.0/24	INVALID
1312	2001:468:c80::/48	INVALID
1312	2001:468:ce0::/44	INVALID
1351	2001:468:606::/48	INVALID
1406	2001:470:e::/48	INVALID

**Details**

ASN	Prefix	Length	Result
11537	2001:468::/32	48	INVALID

Feedback

First Previous 1 2 3 4 5 Next Last

Showing 1 to 10 of 1,043 entries (filtered from 428,362 total entries)



# **RPKI**

## **Router Integration**

### **Section 11**



**RIPE**  
NCC

- **Router sessions**
  - Validator listens on 8282 for RPKI-RTR Protocol
  - Routers can connect and download the cache
- **Export function**
  - Allows you to download a CSV with the cache
  - Can be integrated with your internal workflow
  - Use for statistics or spotting anomalies

**<http://www.ripe.net/certification> -> Tools and Resources**





## Afternoon: IPv6



**RIPE**  
NCC

- IPv4?
- IPv6 Addresses
- IPv6 in the RIPE Database
- IPv6 Addressing Plans
- Tips

## Section 1




**“On 14 September 2012, the RIPE NCC  
ran out of their regular pool of IPv4”**



**RIPE**  
NCC

- **Around 2.4 billion internet users now**
  - around 35% of all people
- **Mobile phones are becoming internet devices**
- **The Internet of things**
  - How will the Internet look like in 5 years?

- Extends the capacity of the IPv4 address space by sharing an IPv4 address between clients
- Fairly common technology, used everywhere
- Breaks the end to end connectivity model
- **It doesn't allow communication with IPv6!**
- You are probably going to need it in some form





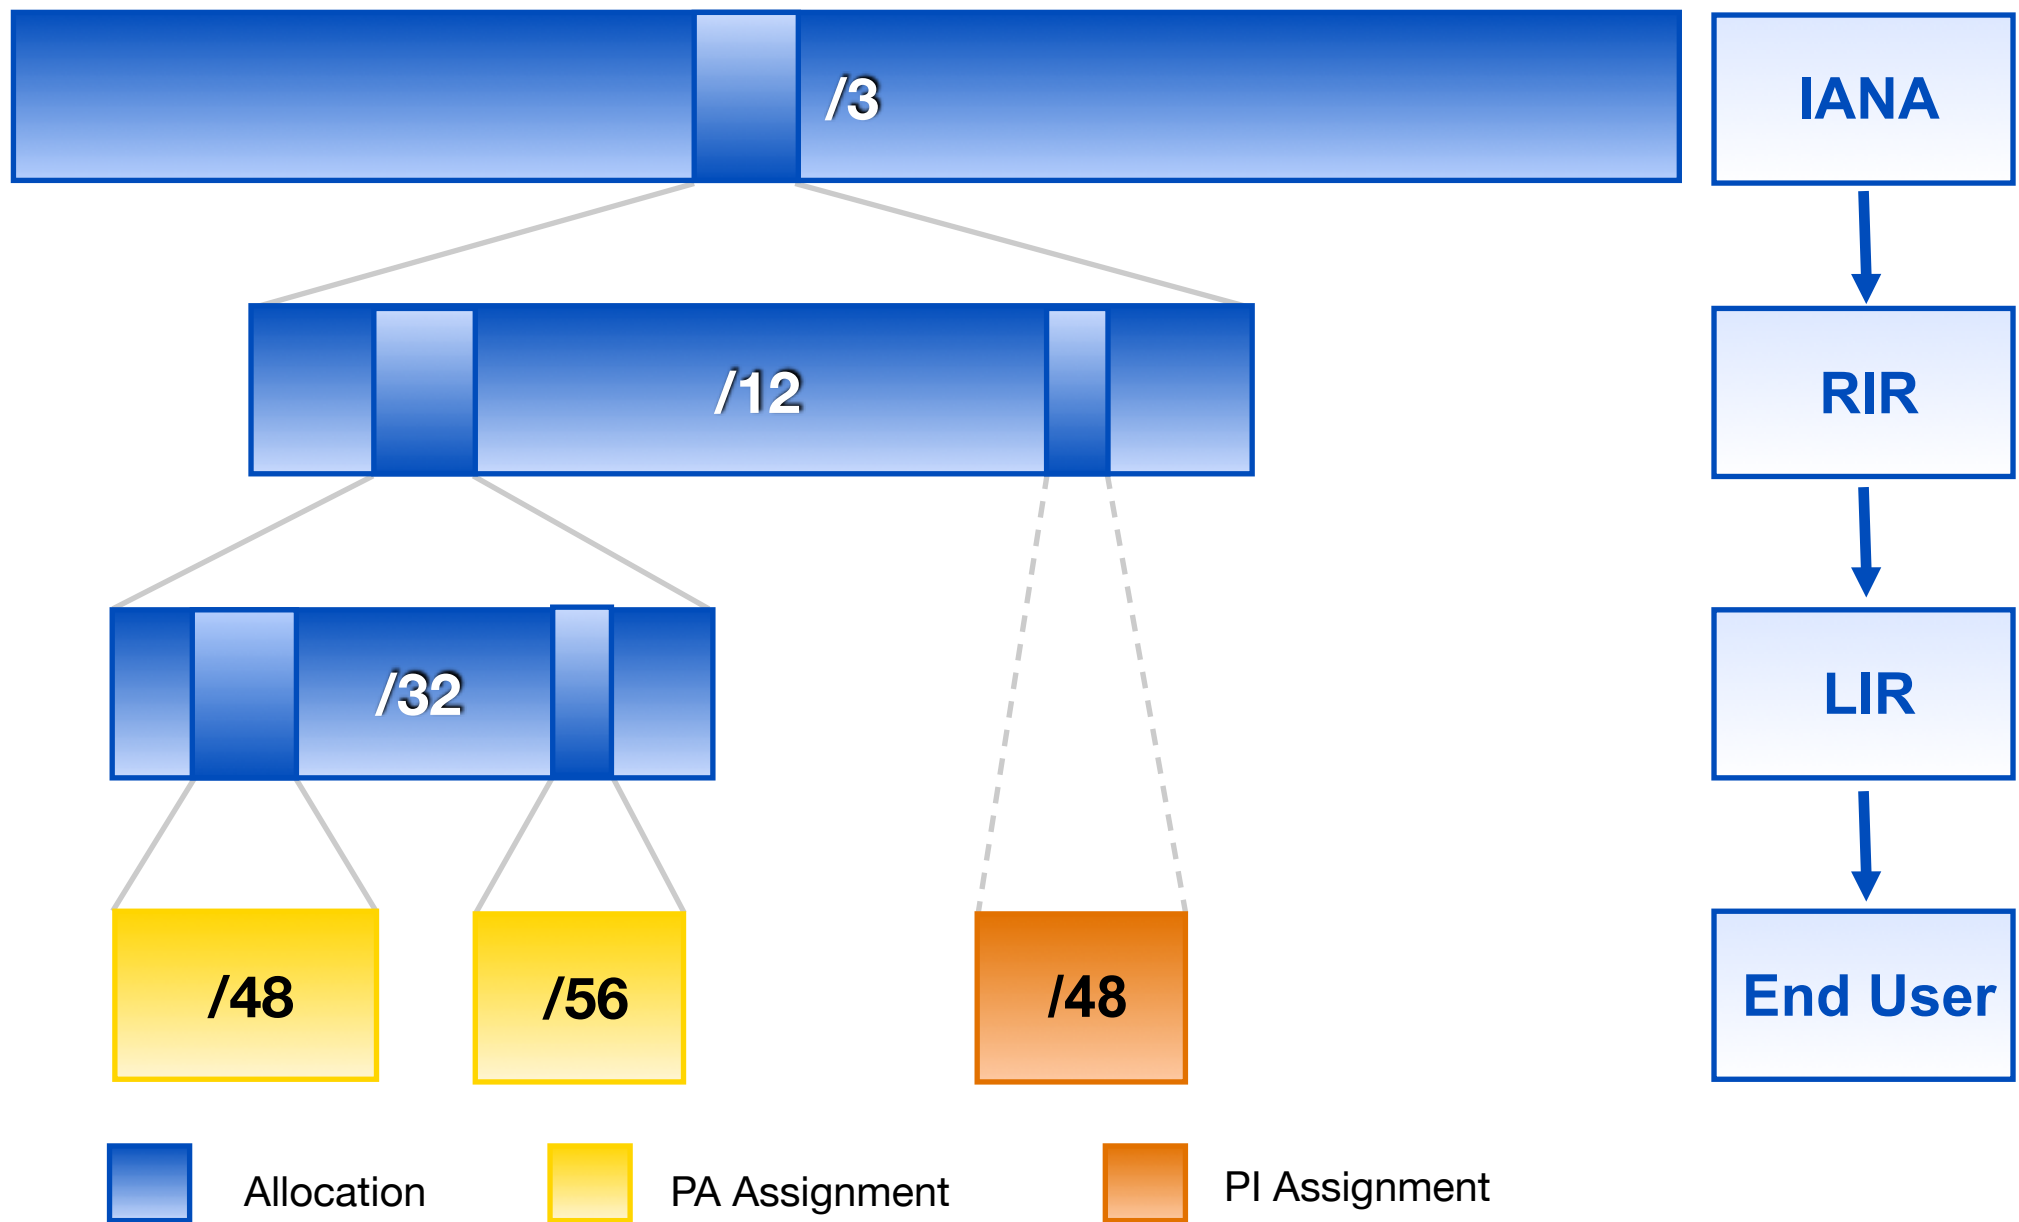
# IPv6 Address Basics

## Section 2



**RIPE**  
NCC





- IP addresses
  - IPv4
  - IPv6
- AS numbers

- **LIR Training Course**
- **RIPE Database Training course**
- **Routing Security Training Course**
- **DNSSEC Training Course**
- **IPv6 for LIRs Training Course**

# Most Difficult part of the IPv6 course?

80





~~How many IP addresses do I need?~~

How many subnets do I need?

Subnet always = /64


- **IPv6 address: 128 bits**
  - 32 bits in IPv4
- **Every subnet should be a /64**
- **Customer assignments (sites) between:**
  - /64 (1 subnet)
  - /48 (65,536 subnets)
- **Minimum allocation size /32**
  - 65,536 /48s
  - 16,777,216 /56s

## IPv6 Subnetting

2001:0DB8:0000:0000:0000:0000:0000:0000

64 bits interface ID

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



**RIPE**  
NCC


Contact Training Services: [ts@ripe.net](mailto:ts@ripe.net)  
Follow us on Twitter: [www.twitter.com/TrainingRIPENCC](https://www.twitter.com/TrainingRIPENCC)  
[www.ripe.net](http://www.ripe.net)



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

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

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



Addresses	Range	Scope
Loopback	::1	host
Link Local	fe80::/10	link
Unique Local	fc00::/7	global
Global Unicast	2000::/3	global
6to4	2002::/16	global
Teredo	2001::/32	global
Multicast	ff00::/8	variable



# IPv6 Address Notation

## Exercise



**RIPE**  
NCC



# IPv6 in the RIPE Database

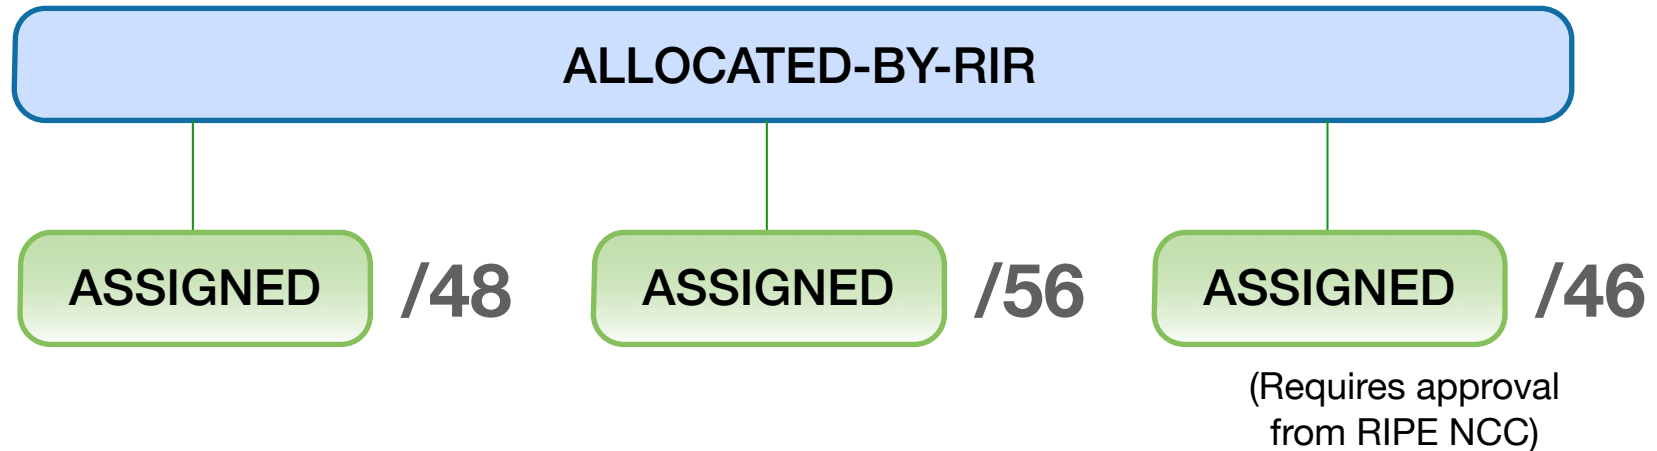
## Section 2



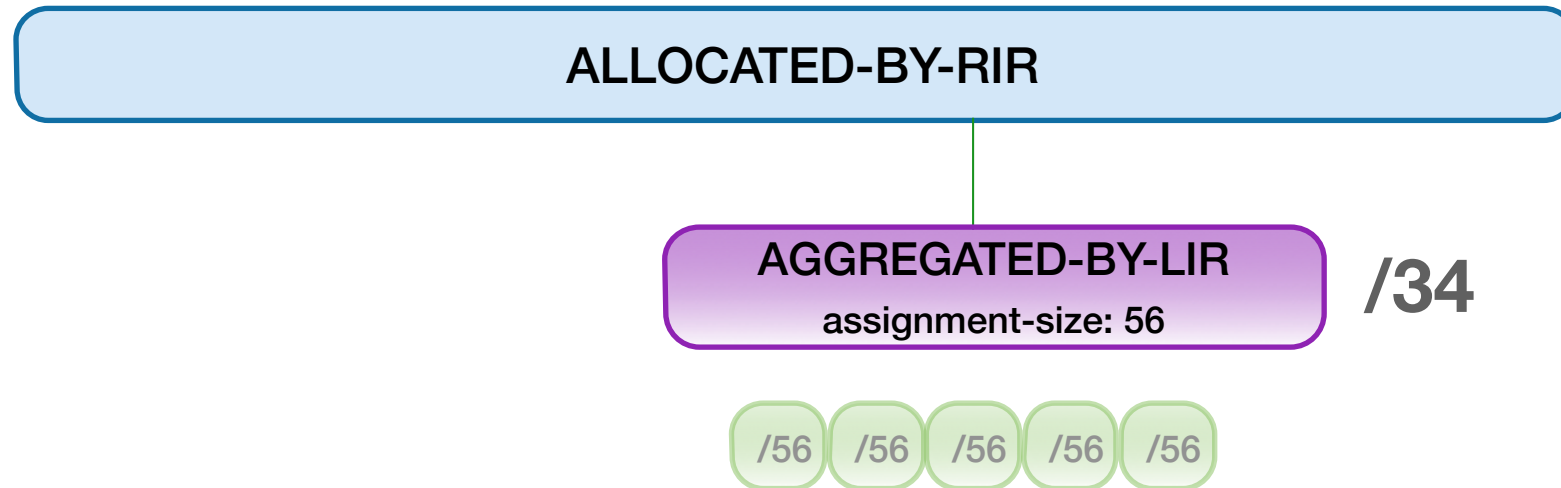
**RIPE**  
NCC

- **Give your customers enough addresses**
  - up to a /48
- **For more addresses send in request form**
  - alternatively, make a sub-allocation
- **Every assignment must be registered in the RIPE Database**

IPv4	IPv6
ALLOCATED PA	ALLOCATED-BY-RIR
ASSIGNED PA	ASSIGNED
ASSIGNED PA	AGGREGATED-BY-LIR
SUB-ALLOCATED PA	ALLOCATED-BY-LIR
ASSIGNED PI	ASSIGNED PI




- Status is **ASSIGNED**
- Minimum assignment size is a /64
- For more than a /48, send a request form



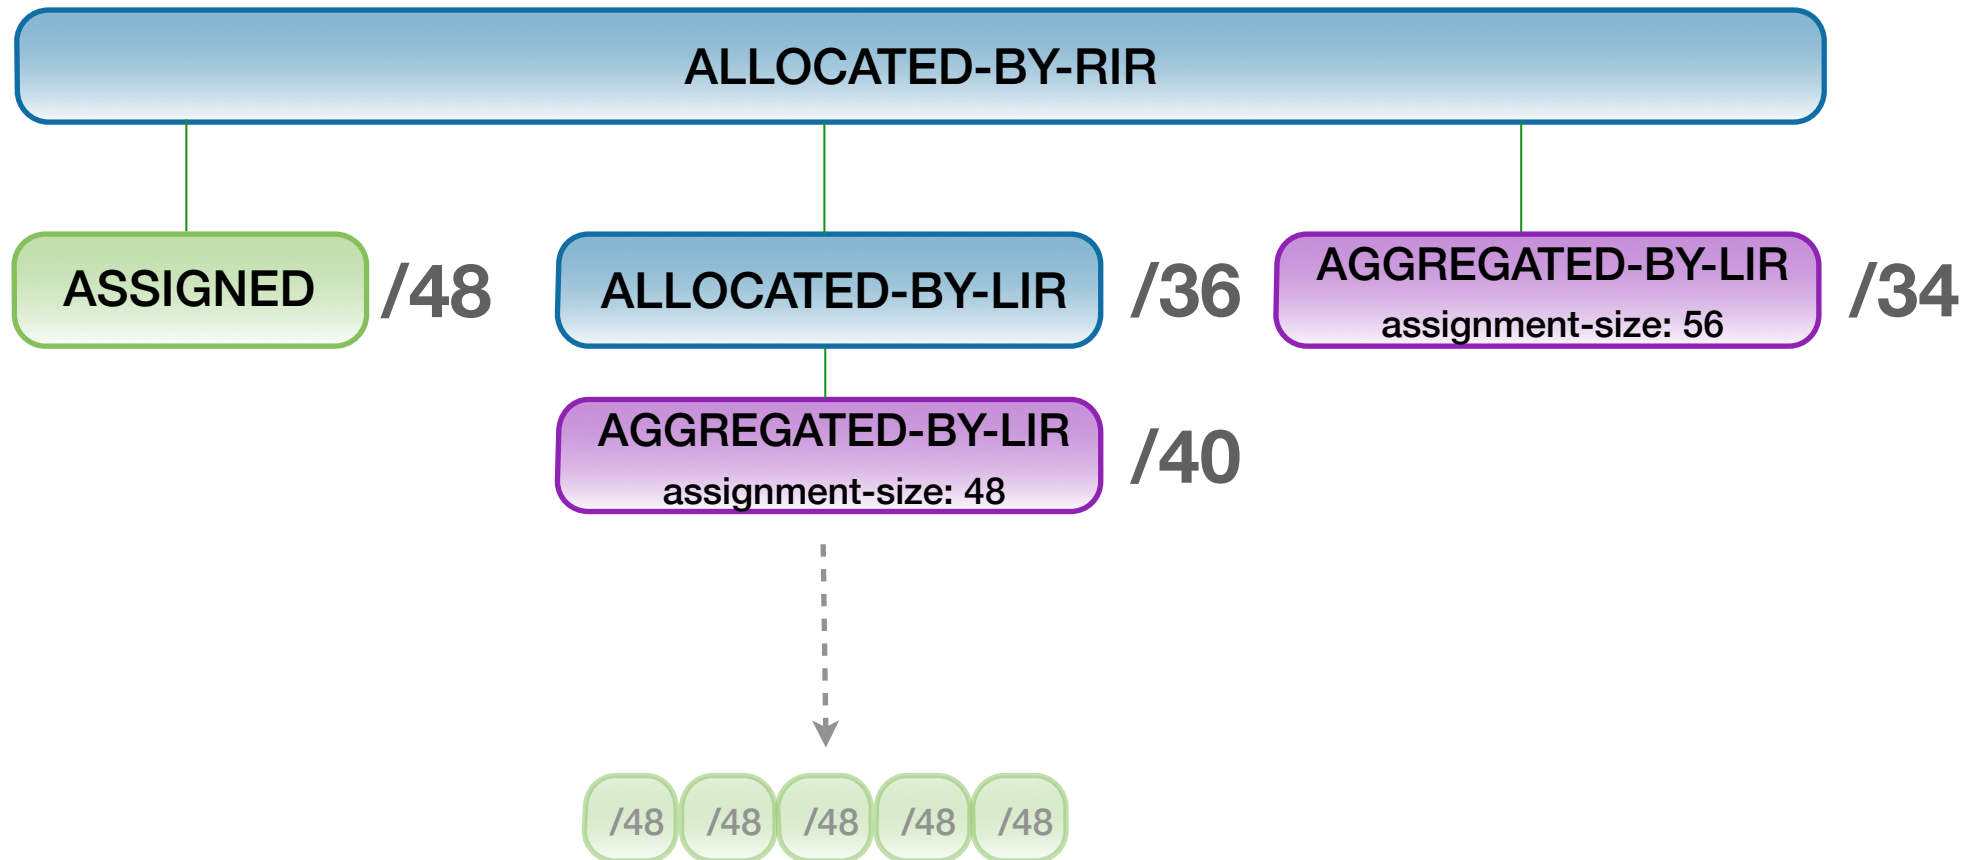
- Can be used to group customers
  - broadband, for example
- “assignment size” = assignment of each customer



inet6num:	2001:db8:1000::/36
netname:	Brightlife
descr:	Broadband services
country:	NL
admin-c:	BN649-RIPE
tech-c:	BN649-RIPE
<b>status:</b>	<b>AGGREGATED-BY-LIR</b>
<b>assignment-size:</b>	<b>48</b>
mnt-by:	BRIGHTLIFE-MNT
notify:	noc@example.net
changed:	noc@example.net 20130218
source:	RIPE



- Can be used for customers with potential for growth
  - or for your own infrastructure
  - or to delegate address space to a downstream ISP



- **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!

inet6num:	2001:db8:1000::/36
netname:	FREEZ
descr:	Freez Fridges
country:	NL
admin-c:	RM1204-RIPE
tech-c:	RM1204-RIPE
status:	<b>AGGREGATED-BY-LIR</b>
assignment-size:	<b>56</b>
mnt-by:	LIR-MNT
notify:	noc@lir-example.com
changed:	noc@lir-example.com 20110801
source:	RIPE



# IPv6 Addressing Plans


## Section 3



**RIPE**  
NCC

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

- IPv6 offers flexibility with addressing plans
- Network addressing can be done on 4-bit boundaries





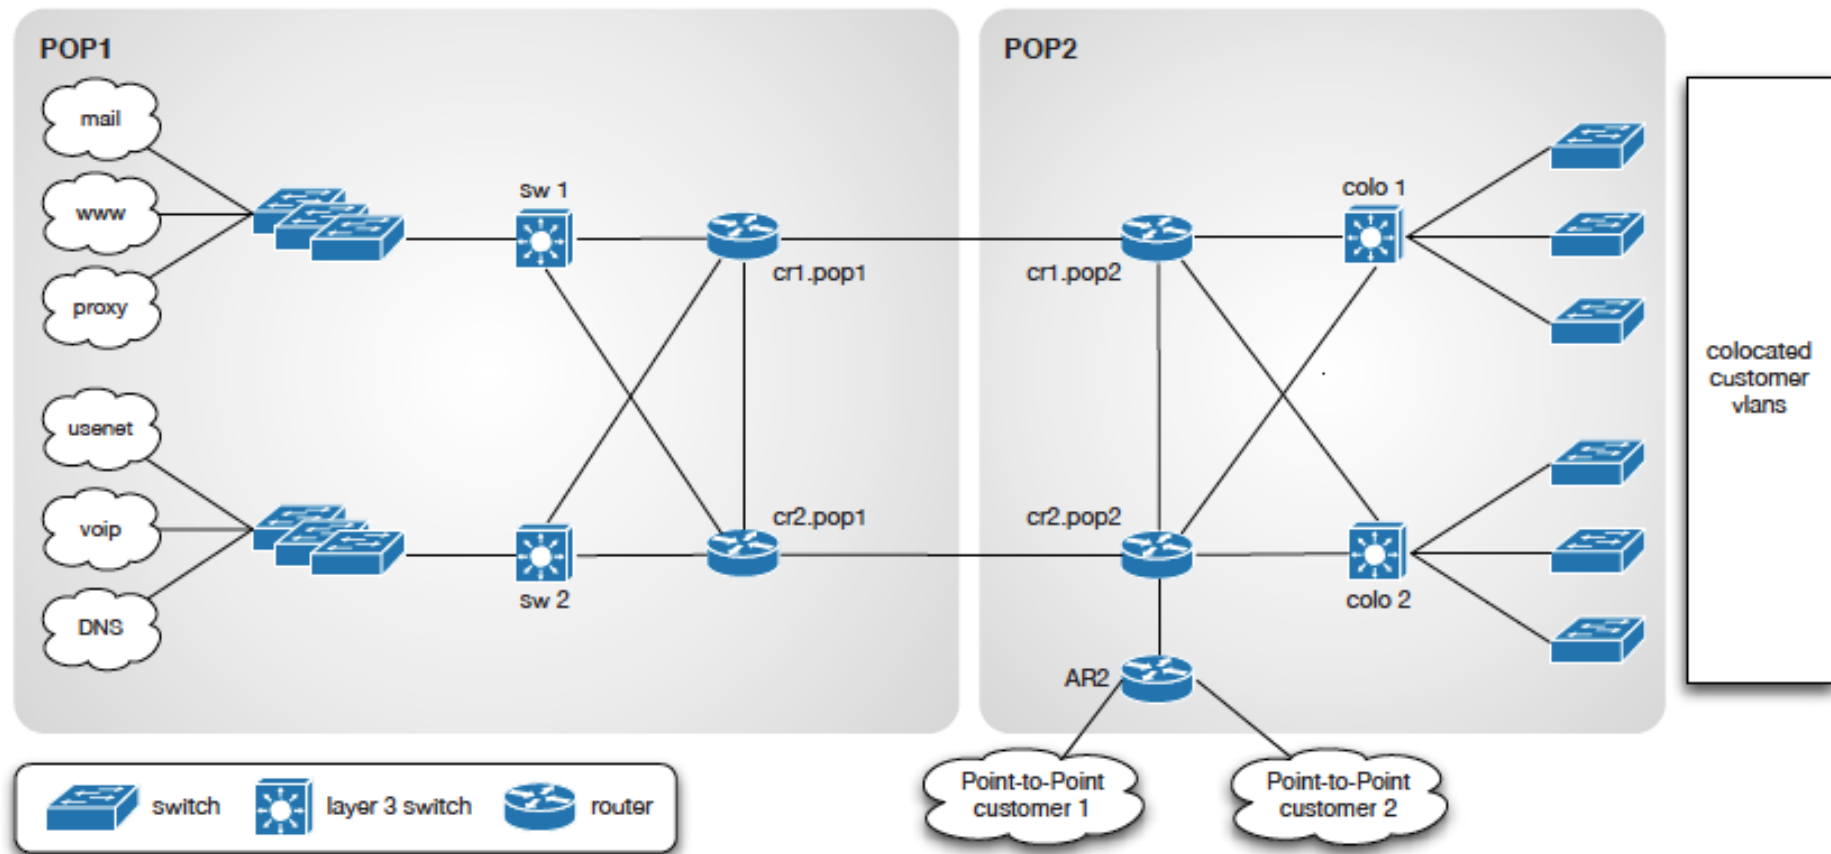
- **Your spreadsheet might not scale**
  - There are 512K /48s in a /29
  - There are 65.536 /48s in a /32
  - There are 65.536 /64s in a /48
  - There are **16.777.216** /56s in a /32
- **Find a suitable IPAM solution**

- “Every interface ID must be a /64” (RFC 4291)
- Because of SLAAC
- Other RFCs followed this
- The only exception is a /127 for point-to-point links

- What should an ISP addressing plan contain?
  - Address space for internal use
  - Loopback interfaces
  - Point-to-point connections
  - Servers, routers and other infrastructure at POPs
- Use a /48 per POP
- Address space for customers

- **One /128 per device**
  - One /64 contains enough addresses for all your manually configured loopback addresses
- **Take an easy to remember block for loopback addresses**

- **With old router operating software:**
  - One /64 per point-to-point connection
  - Reserve /64 per point-to-point link, but configure a /127
- **With new router operating software:**
  - RFC 6164
  - Configure a /127 per point-to-point connection

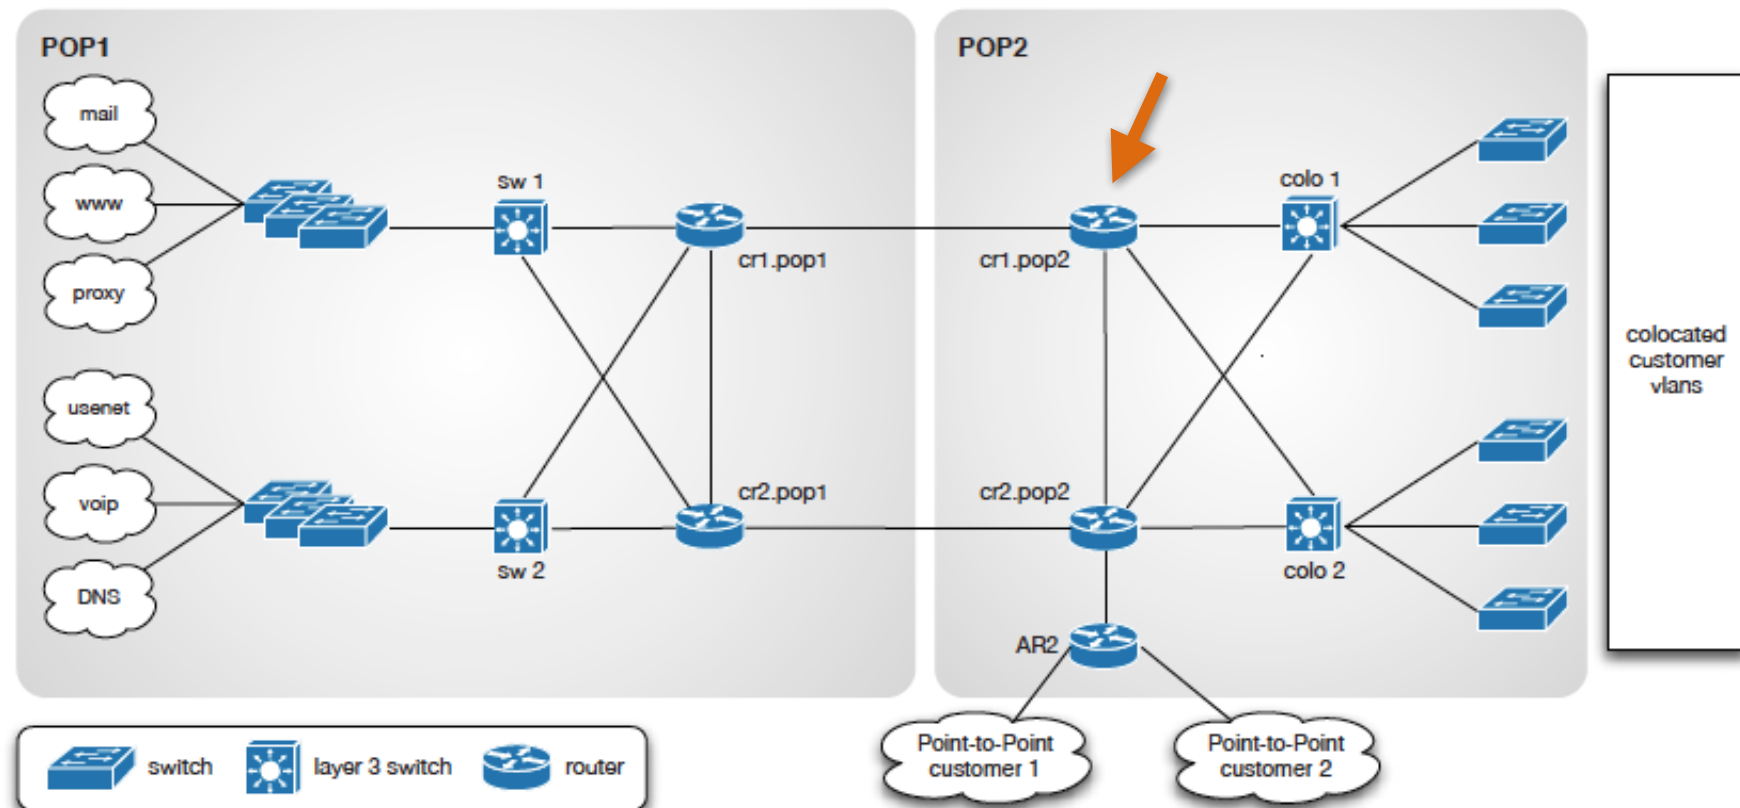


- We will assign a /48 per POP
- We will work on 4-bit boundary

Prefix	Number of /64 subnets
/48	65.536
/52	4096
/56	256
/60	16
/64	1

- Look at the number of point-to-point links
- Just to be sure, we reserve a /64 per link!

## How much would you assign to cr1.pop2?





- **In common cases:**
  - One /48 per POP
  - Calculate growth
  - Make it scalable

- Customers should get a large block of addresses
  - /48 for business customers
  - /48 or /56 for residential customers
- For more than a /48, send a request form
- Every assignment must be registered

- **A customer has 6 functions**
  - Servers
  - Office PCs
  - Network Engineer PCs
  - Guests
  - VPN (remote workers)
  - Infrastructure (point-to-point and loopbacks)

- **A customer has 3 locations**
  - Main building, floor 1
  - Main building, floor 2
  - Secondary office

- A customer receives 2001:0db8:1a2b::/48
- Work on 4-bit boundary
  - 6 functions (leaves room for 10 more functions)
  - 3 locations (leaves room for 13 more locations)
  - We still have 8 bits!
    - Room for 256 networks per function per location

- Putting this in the address:

2001:0db8:1a2b:**FLXX**::/64

- **F** = Function (0=infrastructure, 1=servers, 2=office, 3=engineers, 4=VPN, f=guests)
- **L** = Location (0=main building 1, 1=main building 2, 2=secondary office)
- **XX** = Number of network of type + location

- 2001:0db8:1a2b:1000::/64
  - Servers in Main Building, floor 1, network 0
- 2001:0db8:1a2b:1200::/64
  - Servers in the secondary office, network 0
- 2001:0db8:1a2b:f209::/64
  - Guest in secondary office, network 9

- 2001:0db8:1a2b:0000::1/128
  - Loopback address (infrastructure, location doesn't apply)
- 2001:0db8:1a2b:0102::/127
  - Point-to-point link (infrastructure, location doesn't apply)
- 2001:0db8:1a2b:41ab::/64
  - VPN in main office, floor 1, user 171



- The previous example is just an idea
  - ✓ Adapt as necessary
- 2001:0db8:1a2b:**FFLX**::/64
  - 256 functions
  - 16 locations
  - 16 networks per function per location

- **Tips:**
  - Work on 4-bit boundary
  - Group subnets by function
  - Group subnets by location
  - Make a scalable addressing plan

What is the IPv6 address for an engineer's PC,  
in the main building floor 2,  
for computer number 2?

Example:

2001:0db8:1a2b:FLXX::/64

- **F**= Function (0=infrastructure, 1=servers, 2=office, 3=engineers, 4=VPN, f=guests)
- **L**= Location (0=main building 1, 1=main building 2, 2=secondary office)
- **XX**= Number of network of type + location

What is the IPv6 address for an engineer's PC,  
in the main building floor 2,  
for computer number 2?

Example:

2001:0db8:1a2b:FLXX::/64

- **F**= Function (0=infrastructure, 1=servers, 2=office, 3=engineers, 4=VPN, f=guests)
- **L**= Location (0=main building 1, 1=main building 2, 2=secondary office)
- **XX**= Number of network of type + location

2001:0db8:3101::/64

or

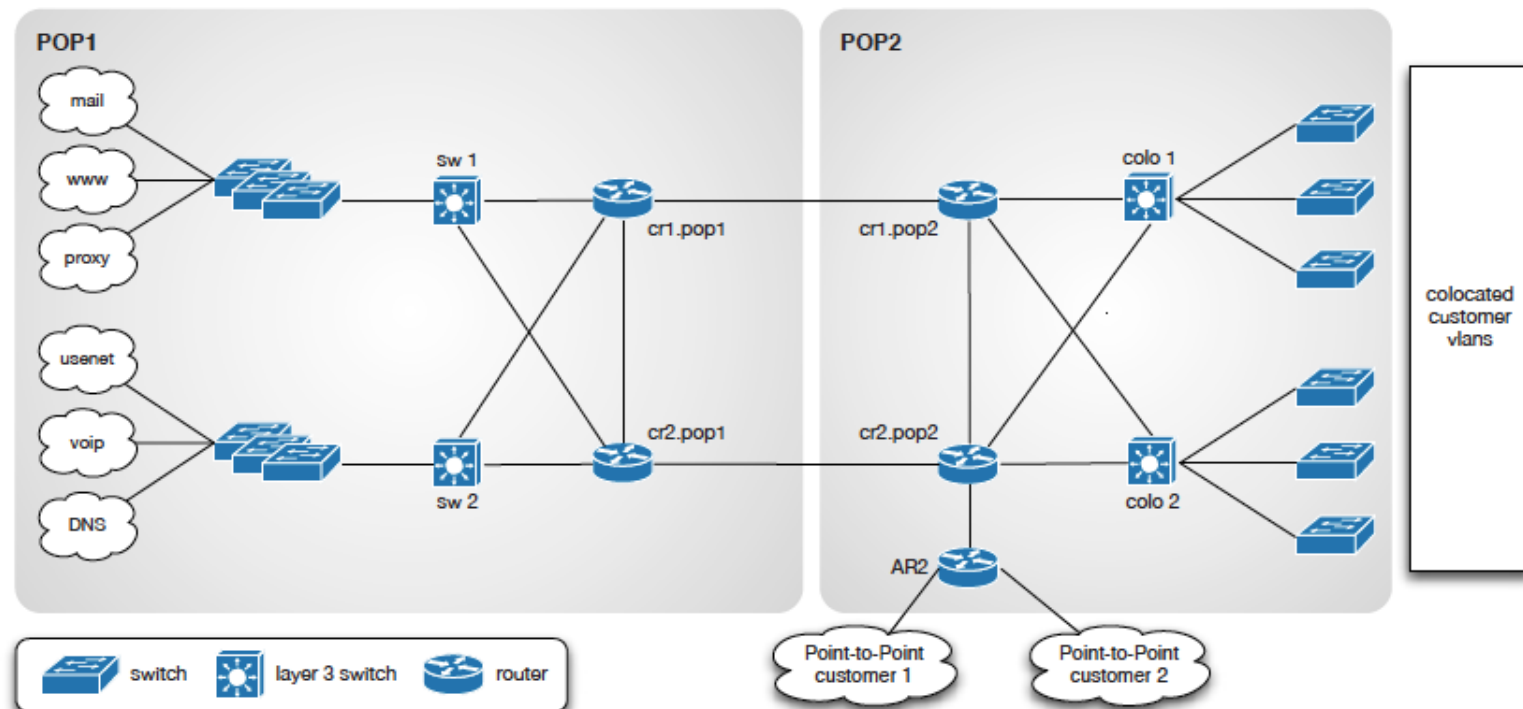
2001:0db8:3102::/64

# Exercise: Addressing plan

121

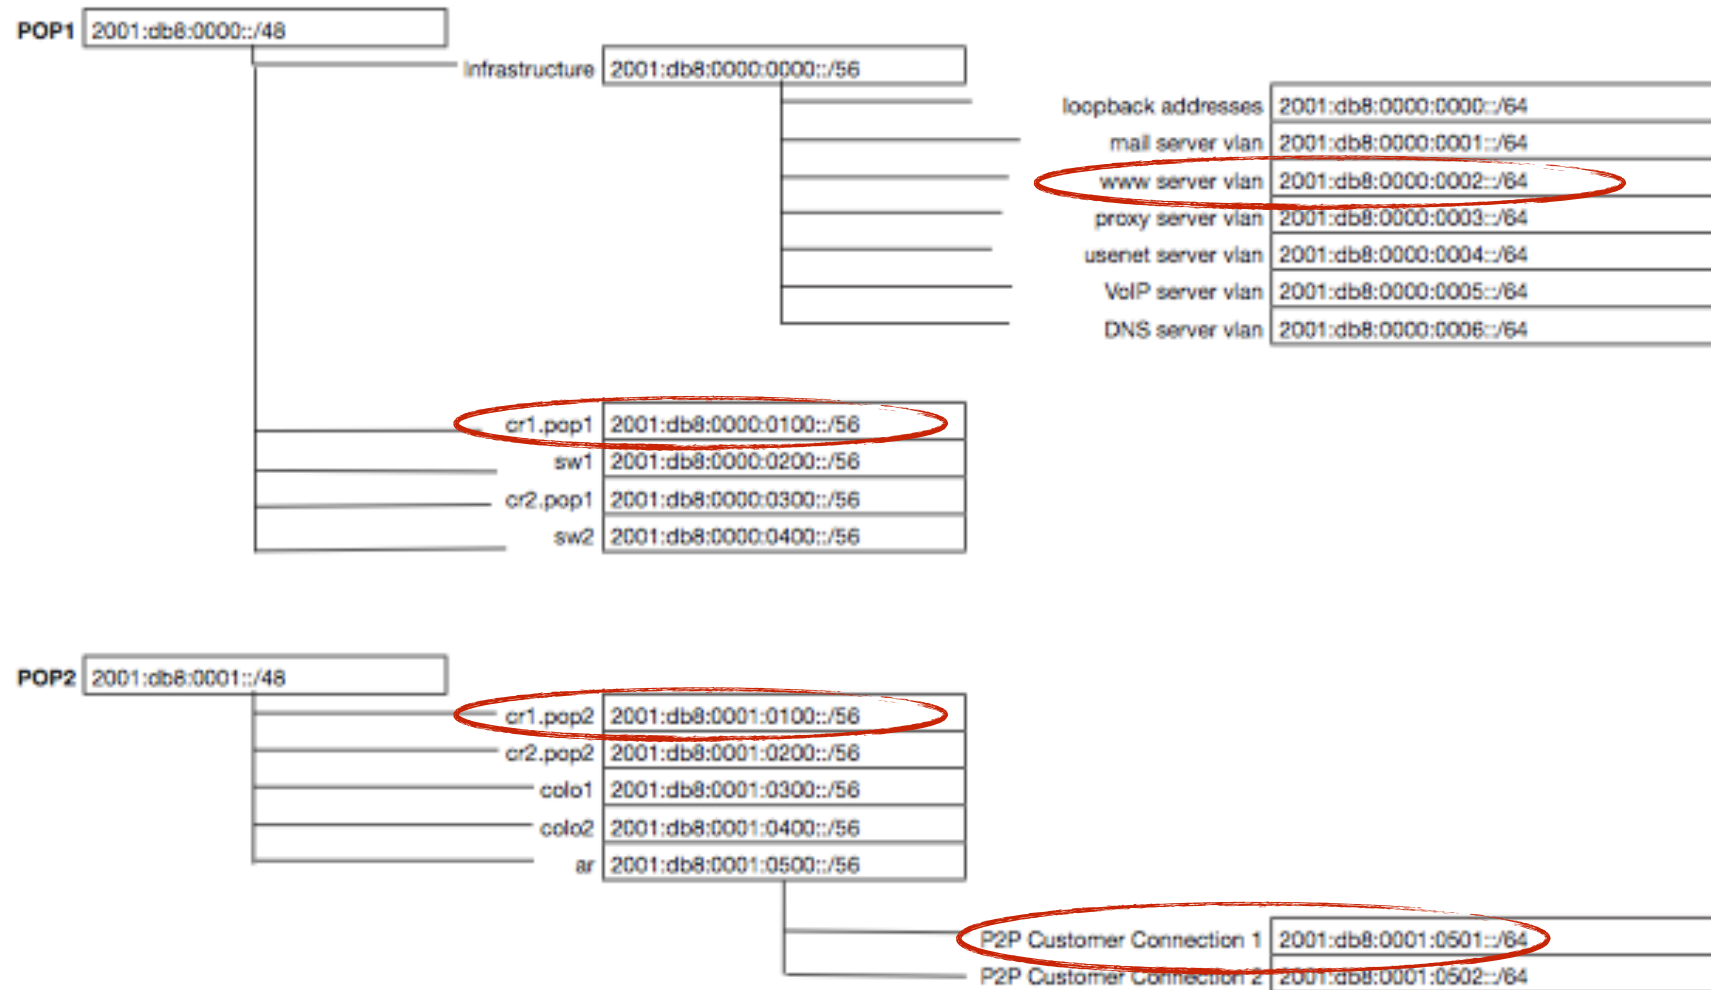
Assignments to:

- www VLAN
- colo 1 and colo 2: consider that there are 250 customers behind each colo.
- cr1.pop2 and cr1.pop1
- Point-to-Point customer 1



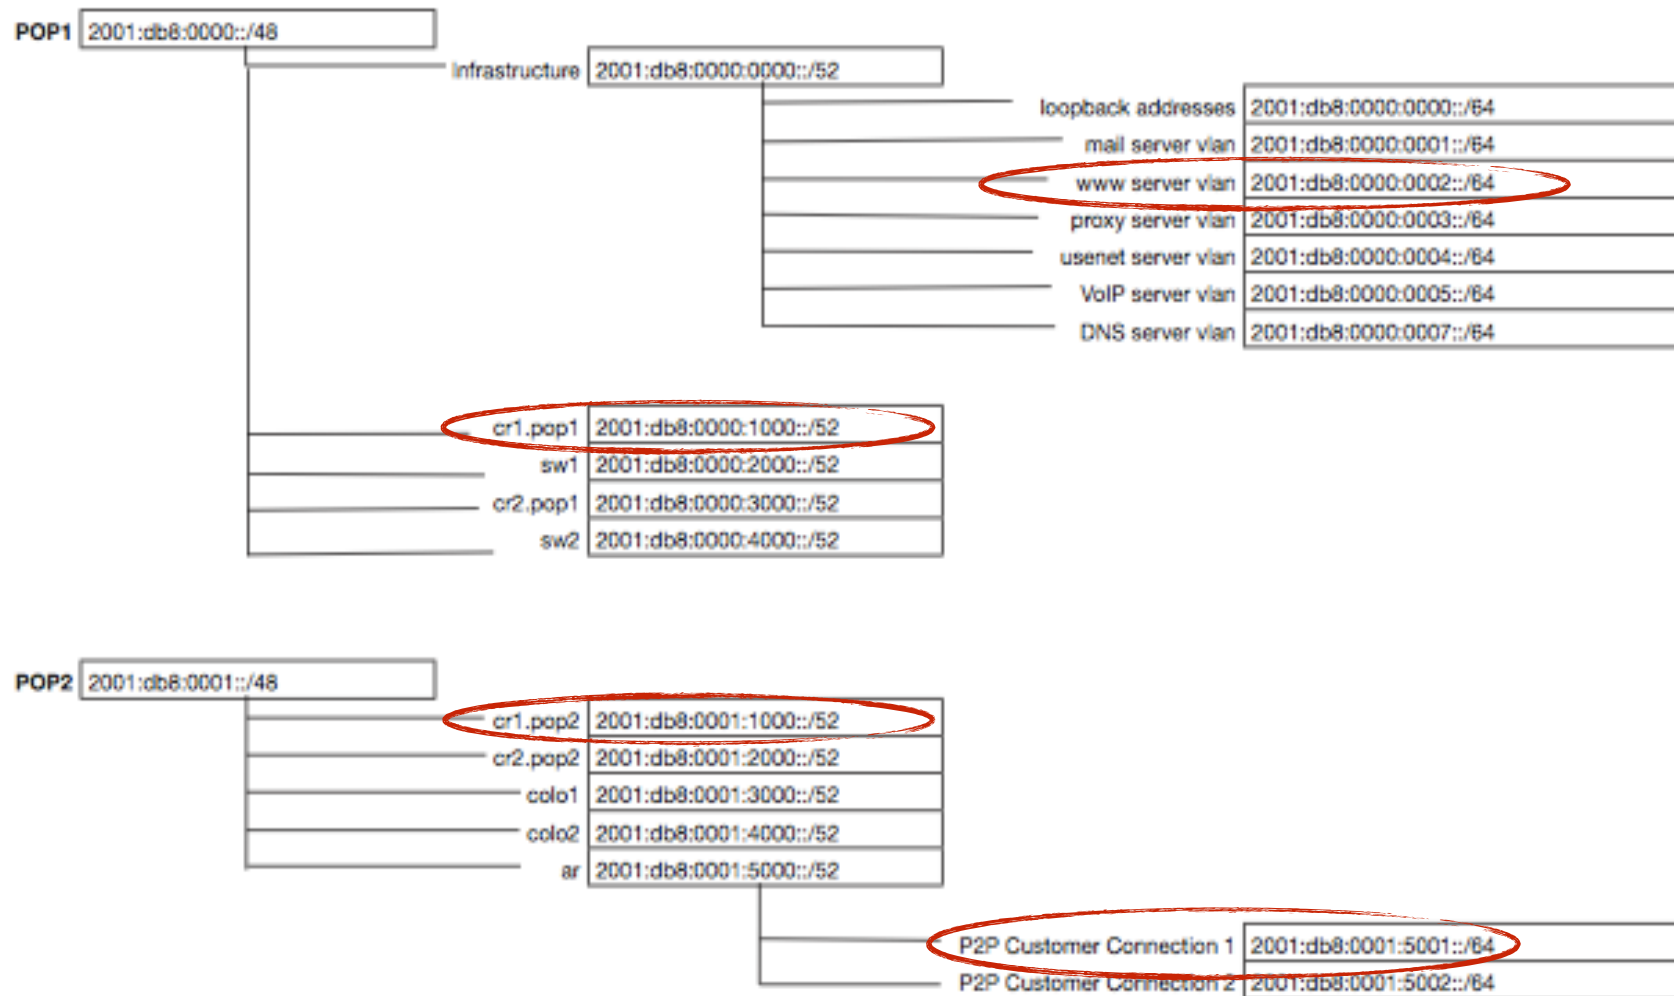
# Addressing plan: solution 1

122



# Addressing plan: solution 2

123



- **Number of hosts in a /64 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**



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

## Tips

### Section 9



**RIPE**  
NCC

- **“Requirements for IPv6 in ICT Equipment”**
  - Best Current Practice describing what to ask for when requesting IPv6 Support
  - Useful for tenders and RFPs
  - Originated by the Slovenian Government
  - Adopted by various others (Germany, Sweden)

- 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/Preparing-an-IPv6-Addressing-Plan.pdf>



- **Websites**

- <http://www.getipv6.info>
- <http://www.ipv6actnow.org>
- <http://datatracker.ietf.org/wg/v6ops/>
- <http://www.ripe.net/ripe/docs/ripe-554.html>

- **Mailing lists**

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



**RIPE**  
**NCC** ACADEMY

**Graduate to the next level!**

<http://academy.ripe.net>

twitter  
@TrainingRIPENCC





The End!

Край

Y Diwedd

Fí

Finis

Соңы

ჟღერჟ

Liðugt

النهاية

Ende

Finvezh

Кінець

Konec

Kraj

Ěnn

Fund

پایان

Кraj

Lõpp

Beigas

Vége

Son

An Críoch

הסוף

Fine

Endir

Sfârșit

Fin

Τέλος

Einde

Конец

Slut

Slutt

დასასრული

Pabaiga

Fim

Amaia

Loppu

Tmíem

Koniec