



# Welcome to the

# RIPE NCC

# Training Seminar

Vesna Manojlovic  
BECHA@ripe.net



# Part 1: LIR Basics

- Being an LIR
- RIPE Database Security
- Making Assignments
- PI Address Space
- IPv6 Address Space




## Part 2: Routing Registry & News

- Routing Registry and Related Services
  - Documenting Routing Policy
  - AS Numbers
  - Practical Usage of RPSL & Existing Tools
  - The Routing Information Service (RIS)
    - The Routing Registry Consistency Check
  
- News
  - 32-Bit AS Numbers
  - The Routing Table Growth & Other Statistics
  - The Resource Certificates



# Being an LIR

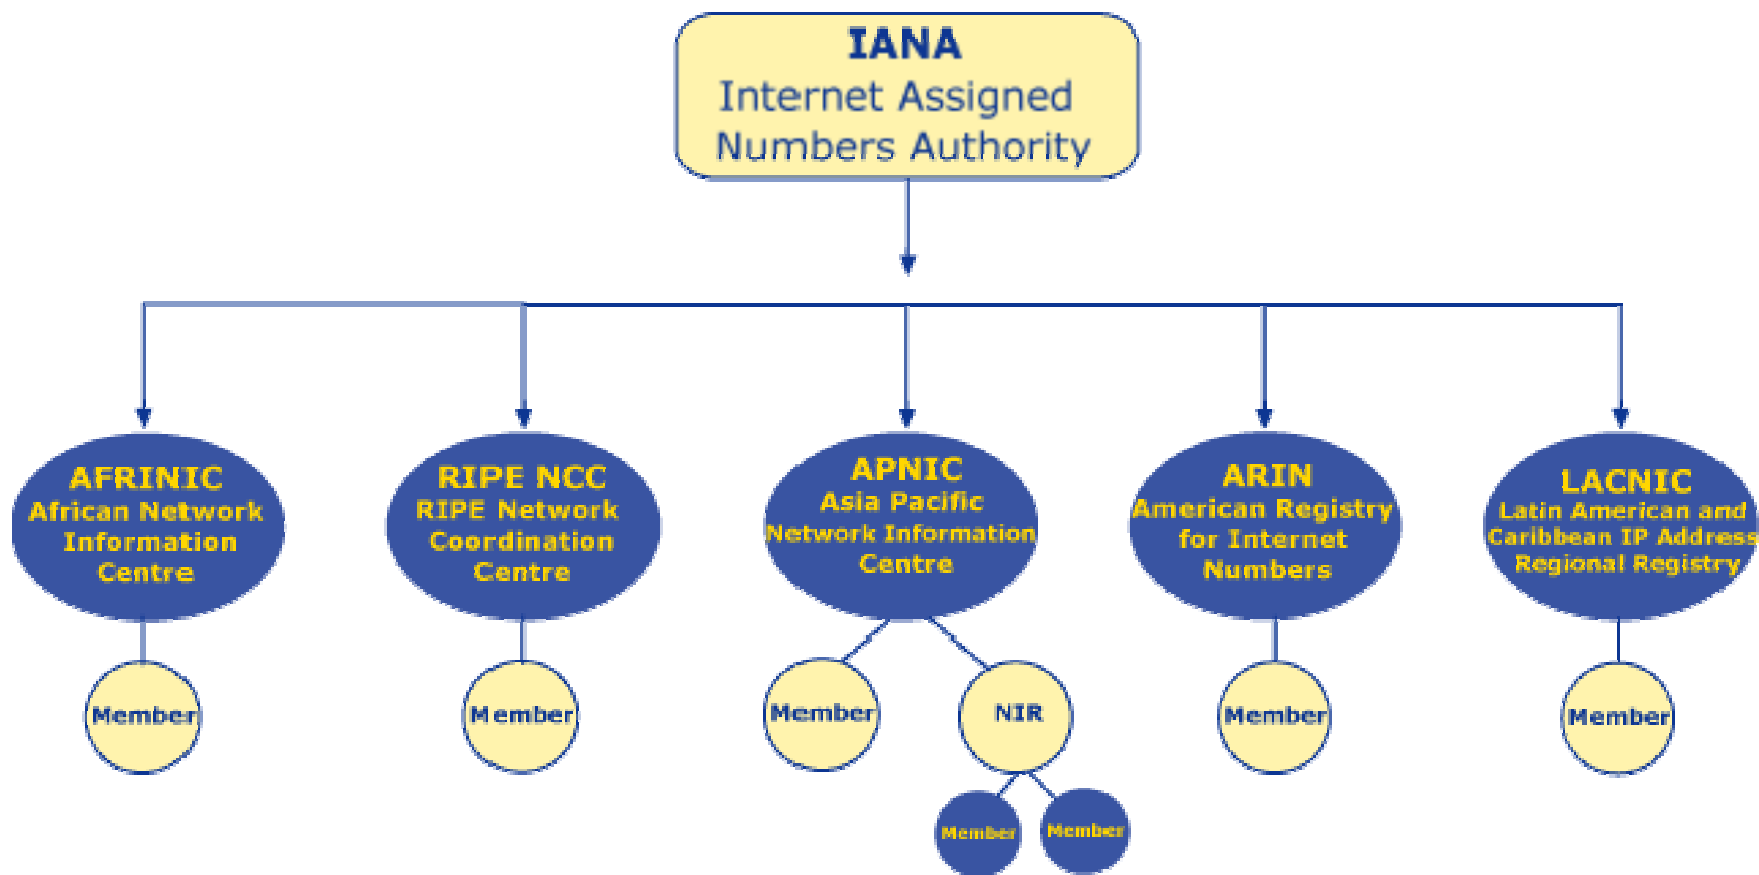


# RIPE and RIPE NCC

- Réseaux IP Européens (1989)
  - collaborative operators' community for co-ordinating IP infrastructure development
  - open to all
  - developing policies; input to the RIPE NCC
- RIPE Network Coordination Centre (1992)
  - independent not-for-profit membership organisation
  - one of five Regional Internet Registries (RIRs)

# Hierarchical IP Resources Distribution

## Internet Resource Allocation



# Internet Registry System Goals

Problem	Solution	Principle/Goal
Uniqueness and contact details	RIPE Database	<b>Registration</b>
Routing table growth	Scalable routing	<b>Aggregation</b>
Limited resource	Efficient use	<b>Conservation</b>

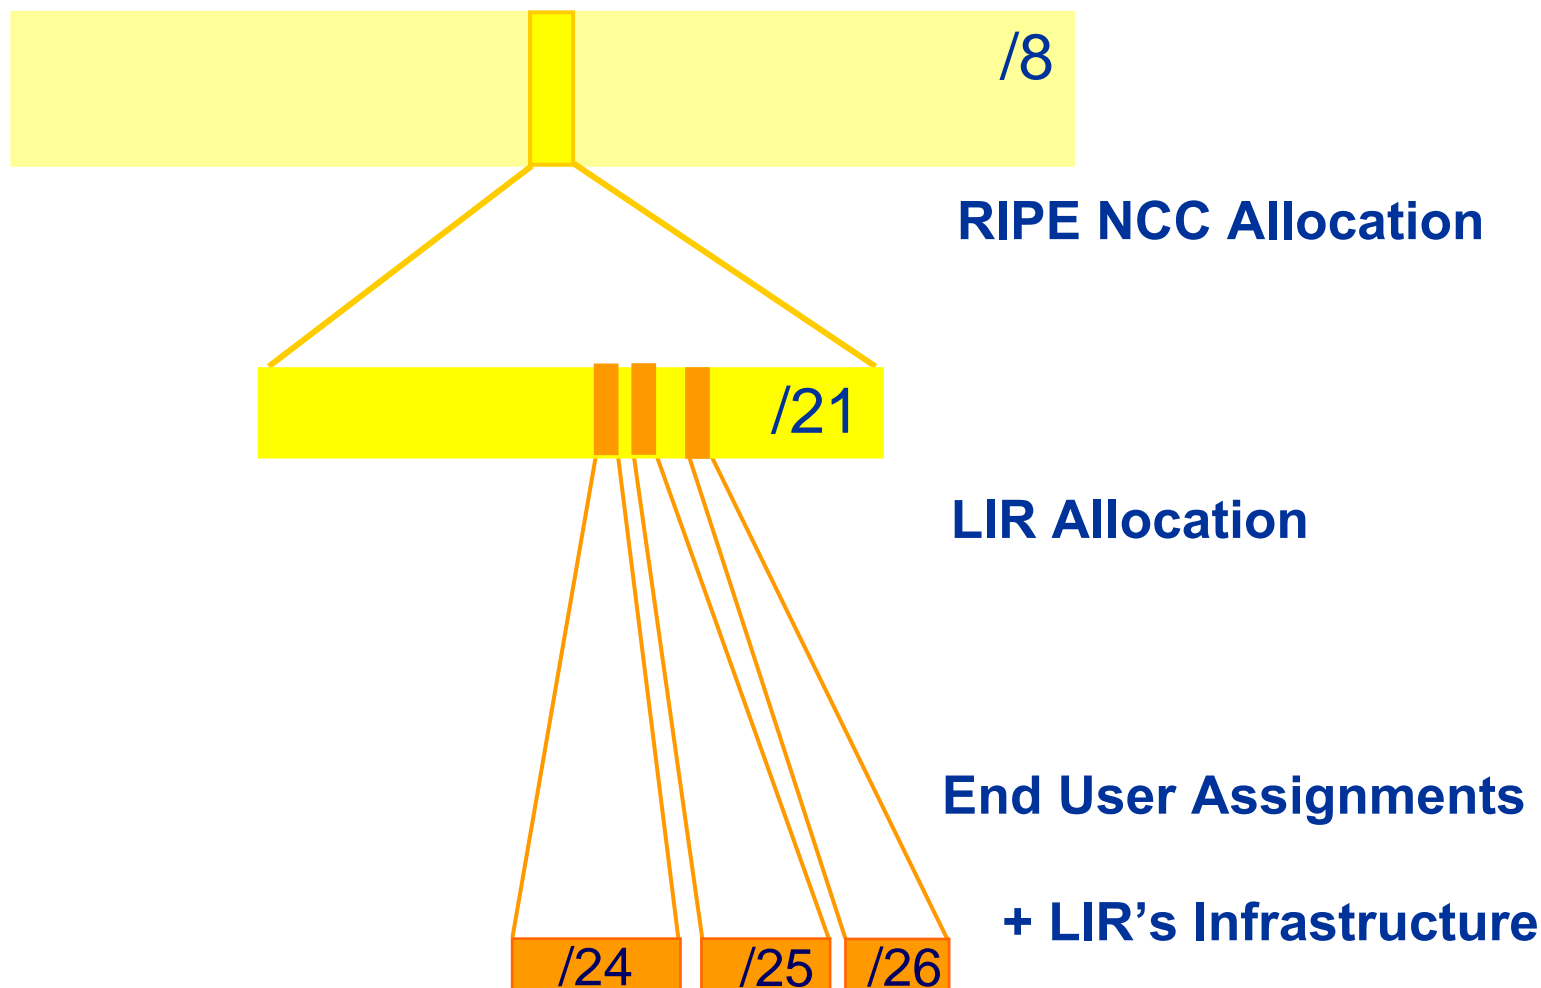


# Classless Addressing

- Classful addressing ('80-'93) now obsolete
  - waste of addresses; routing table growth
- '93: **Classless Inter Domain Routing (CIDR)**
  - flexible allocation / assignment sizes
  - w.x.y.z/nn notation
- CIDR implemented in all modern routing protocols
- CIDR used for address space distribution



# Definitions: Allocation and Assignment



# PI versus PA Assignments

**Provider Independent**  
(Portable Assignments)

No Aggregation

BGP Announcements (4)




Customer Assignments

**Provider Aggregatable**  
(Non-portable Assignments)

Aggregation

BGP Announcement (1)



Customer Assignments



# What is an LIR?

- **Local Internet Registry**
  - responsible for obtaining, distributing and registering IP resources, according to the RIPE policies
- **Member of the RIPE NCC**
  - receiving resources directly from the RIPE NCC
- **Benefits**
  - flexibility
  - independence (BGP multihoming)



# LIR Set-up Process

- Steps
  - read policy documents
  - apply for membership
    - RegID, contacts
  - pay the fees
  - sign the contract
- Next steps
  - LIR: register **person & role** in RIPE Database
  - RIPE NCC: “Reg” file, “organisation” object
  - LIR: activate LIR Portal account

# Sources of Contact Information

## LIR Portal

- RIPE NCC confidential
  - access only by “users”
- “admin” creates “users”
- “users” create “contacts”
- Use: Reg-ID, user, pwd

## Reg File

- RIPE NCC “contacts” can:
  - request resources
  - update contact info
- Use: Reg-ID, name

## RIPE Database

- Public info
  - access by anyone
  - updates by anyone
- Operational contacts
  - troubleshooting
- Responsibility over registered resources
- Use: nic-handle
- Additional authorisation: using “maintainer”



# LIR: Summary

- You are part of the global Registry System
- Think CIDR!
- LIRPortal.ripe.net : main interface

## Questions?



# RIPE Database Security

# Protection of DB Objects

- “mnt-by”: attribute refers to mntner object
  - Checked at every update
- Password:
  - CRYPT-PW **about to be deprecated!**
  - MD5-PW
    - <https://www.ripe.net/cgi-bin/crypt.cgi>
- Private key/Public key
  - PGPKEY-<id> & key-cert object
  - X.509-<id> & key-cert object
- Multiple auth / mnt-by / mntner-s are OR-ed



# Using mntner Object

```

role: Blue Light Internet Contacts
nic-hdl: BLIC1-RIPE
mnt-by: BLUELIGHT-MNT

```

*role object protected  
by **BLUELIGHT-MNT***

```

password: cleartxt_pass

```

*For updates:  
include clear  
text password  
of mntner*

```

mntner: BLUELIGHT-MNT
descr: maintainer for our objects
admin-c: JJ231-RIPE
auth: MD5-PW q5nd!~Sfhk0#a
upd-to: jan@bluelight.nl
mnt-nfy: auto-mnt@bluelight.nl
referral-by: BLUELIGHT-MNT
mnt-by: BLUELIGHT-MNT

```

*Encrypted password*

*If update fails*

*If update is  
successful*



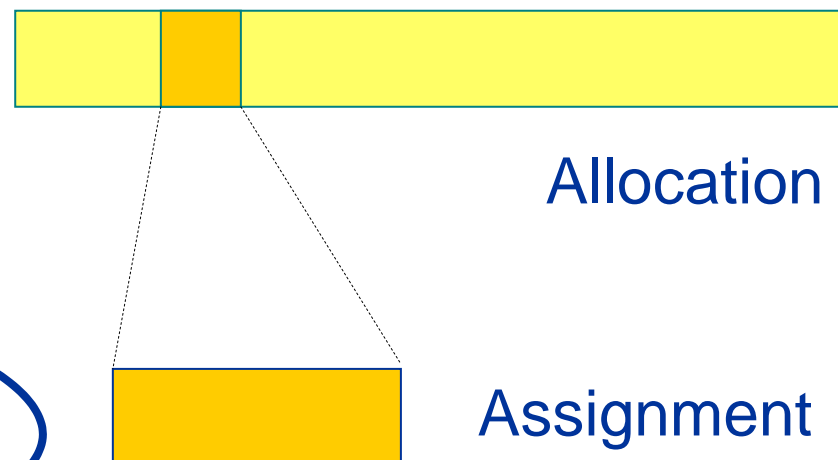
# auth: CRYPT-PW Deprecated Soon

- Phase One: Notification: 24 November 2006
  - Please replace with MD5-PW, PGPKEY or X509 !
- Phase Two: No new CRYPT-PW: 10 January 2007
  - YOU CAN STILL UPDATE YOUR MAINTAINERS!
- Phase Three: Blocked: 21 February 2007
  - Cannot update objects protected with CRYPT-PW any more
    - RIPE NCC will add MD5-PW with new random password
      - If other protection exists, we'll move CRYPT-PW to "remarks:"
  - To recover the locked **mntner**: use your old CRYPT password
    - Go to *RIPE DB > DB Support > Security*

# Hierarchical Authorisation

```
inetnum: 85.118.184.0/21
status: ALLOCATED PA
mnt-by: RIPE-NCC-HM-MNT
mnt-routes: LIR-MNT
mnt-domains: LIR-MNT
mnt-lower: LIR-MNT
```

```
inetnum: 85.118.186.0/24
status: ASSIGNED PA
mnt-by: LIR-MNT
```





# RIPE Database: Summary

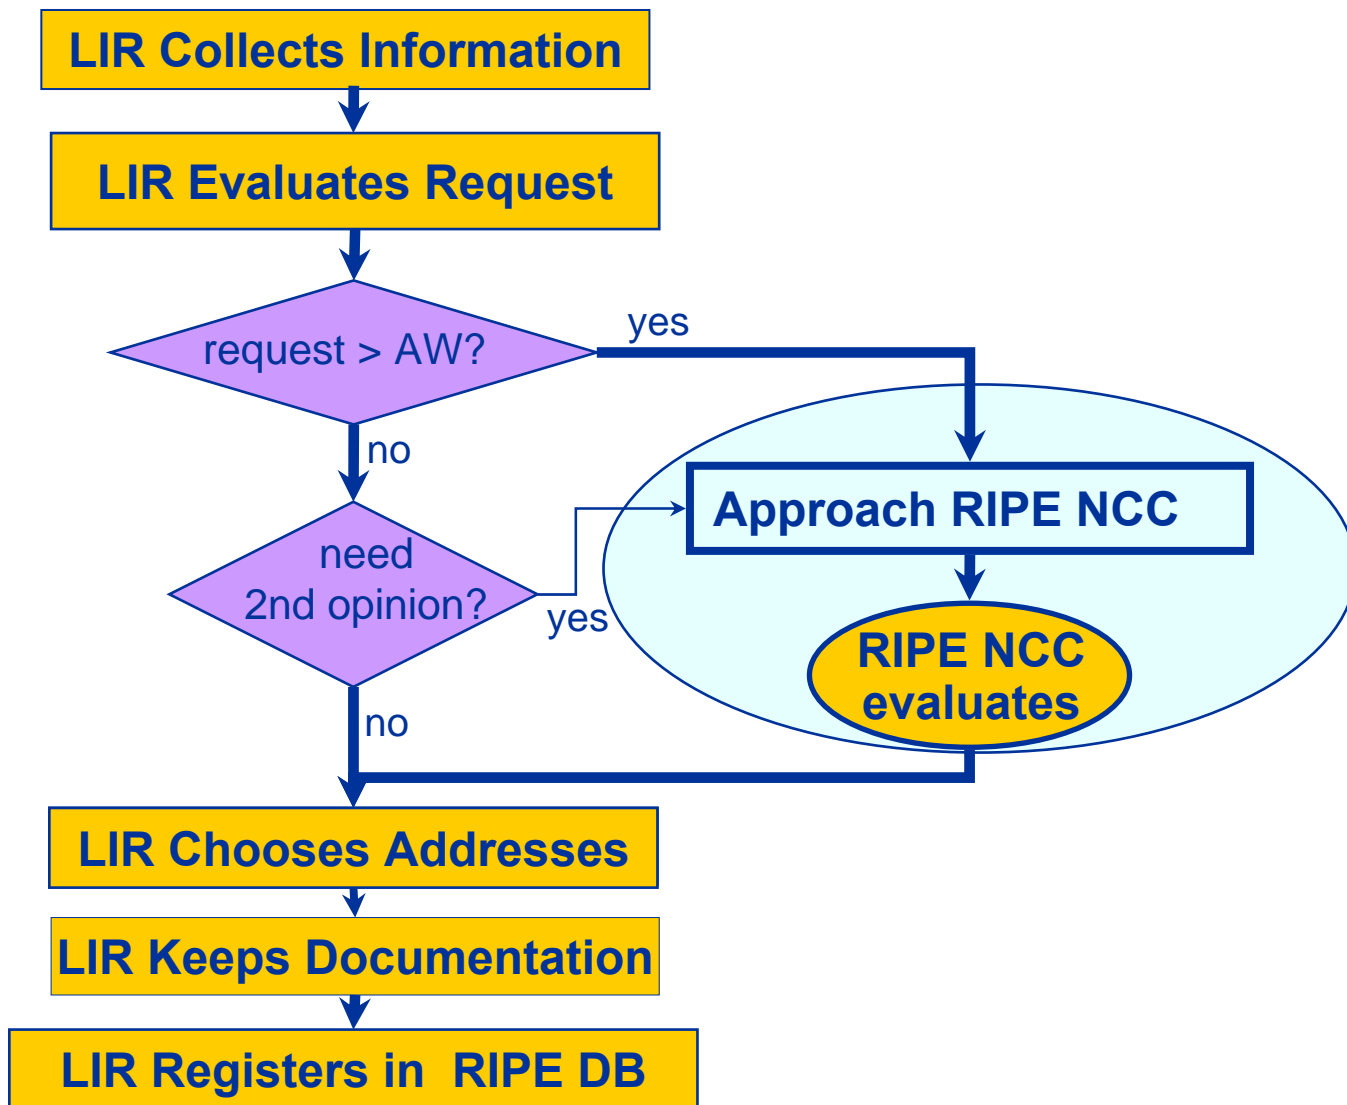
- Maintainers
- Hierarchical authorisation

## Questions?

# Making Assignments



# Assignment Process





# Assignment Window Concept

- Maximum number of IP addresses the LIR can assign without approval from the RIPE NCC
- For each End User, within any 12 months
- New LIR, AW = zero
- RIPE NCC increases AW gradually



# Infrastructure versus End User

- LIR / ISP infrastructure
  - **blocks** for co-location: server housing, web hosting
  - **blocks** for connection to End Users (dial-up, P2P)
  
  - multiple assignments to own infrastructure, within AW
  - **remarks: INFRA-AW**
  
- End User network
  - their equipment, their location
  - separate subnet(s)





# LIR Evaluation

- Collect information from End User
  - Confidential, local language
  
- Planning of growth two years ahead
  - Utilisation: 25% now, 50% in one year
  
- Assignment address range
  - Your choice
  - Any range from your allocation
  - You should make all subnets classless



# Ask for Approval if...

- Request is above AW:
  - This request and all previous assignments you made without the RIPE NCC's approval to the same End User in the last 12 months
- New LIR's  $AW=0$ 
  - Need approval for **every** assignment!
- Request form
  - Through the [LIRPortal.ripe.net](http://LIRPortal.ripe.net)
  - Or in the "Document Store"



# RIPE NCC Evaluation

- Based on “IPv4 Address Policies” document
  - Dynamic assigning encouraged
    - not static
  - More than /20: usage statistics verification
    - always-on technologies: xDSL, cable, GPRS...
  - Name-based virtual web hosting encouraged
    - not IP-based
    - exceptions: SSL, ftp & mail servers...



# Approval

- RIPE NCC sends approval message to LIR
  - Size
    - “netname:”
  - Date
    - ticket closed
- LIR keeps approval message
  - keep all original documents too
- Next steps
  - LIR chooses addresses
  - LIR creates `inetnum` object



# Registering End Users Separately

- Obligatory
- Must match:
  - Size, date & “netname:” of approval
  - Internal documentation: range, “admin-c:”, “tech-c:”
- Benefits:
  - Abuse complaints can go directly to End User
  - Network operators can block End User prefix



# Making Assignments: Summary

- Evaluate End User needs
- Assignment size  $>$  AW: send request for approval (New LIR: AW=0)
- Always register End Users separately

## Questions?



# PI Address Space



# PA versus PI Assignments

- **Provider Aggregatable assignments**
  - LIR assigns to End User
  - Must renumber when changing providers
  - Only way to effectively scale the Internet!
- **Provider Independent assignments**
  - Portable
  - Can be difficult to route
  - Next assignment not aggregatable
  - Affects yearly fee
  - Increases size of the routing tables





# Evaluation of PI requests

- Additional questions
  - Why does End User want PI (and not PA)?
  - Requesting extra address space for routing?
  - Aware of consequences?
- Same criteria as for PA assignments
  - Conservative estimates
  - Classless
- Assignment is only valid as long as original criteria remain valid



# PI Responsibilities

- RIPE NCC
  - Assigns to End User
  - Creates `inetnum`
- LIR
  - Sends request for PI space for End User
  - Makes contracts with End Users
  - Helps End Users with reverse DNS, `route` objects
  - Helps End Users if changing provider
- End User
  - Maintains objects
  - Must not assign further



# PI Addresses: Summary

- PA recommended
- LIR requests PI space for End User
- Shared responsibilities

## Questions?

# IPv6 Address Space



# First IPv6 Allocation

- If you
  - a) are an LIR
  - b) not an End Site
  - c) plan to provide IPv6 connectivity to aggregated 'customers', who are assigned /48s
  - d) plan to assign 200 /48s within two years
    - Criteria being discussed currently
  
- Send us "IPv6 first allocation request form"
  
- Minimum initial allocation size /32

# IPv6 Assignments

- Usual assignment size - /48 for each “site”
  - End User network
  - LIR infrastructure (per PoP)
  - No approval needed
- Smaller sizes:
  - /64 just one subnet
  - /128 just one device
- Multiple /48 for very large End Users
  - Approval needed
- Assignment policy being discussed



# New IPv6 Allocation

- HD ratio = 0.8 usage of previous allocation
  - 7132 /48s assignments in a /32
- Correct registrations (all /48s registered)
- New allocation's size: the same as the first
  - Resulting in IPv6 prefix one bit shorter
  - Or bigger if justified (sufficient for two years)




# IPv6 Allocations

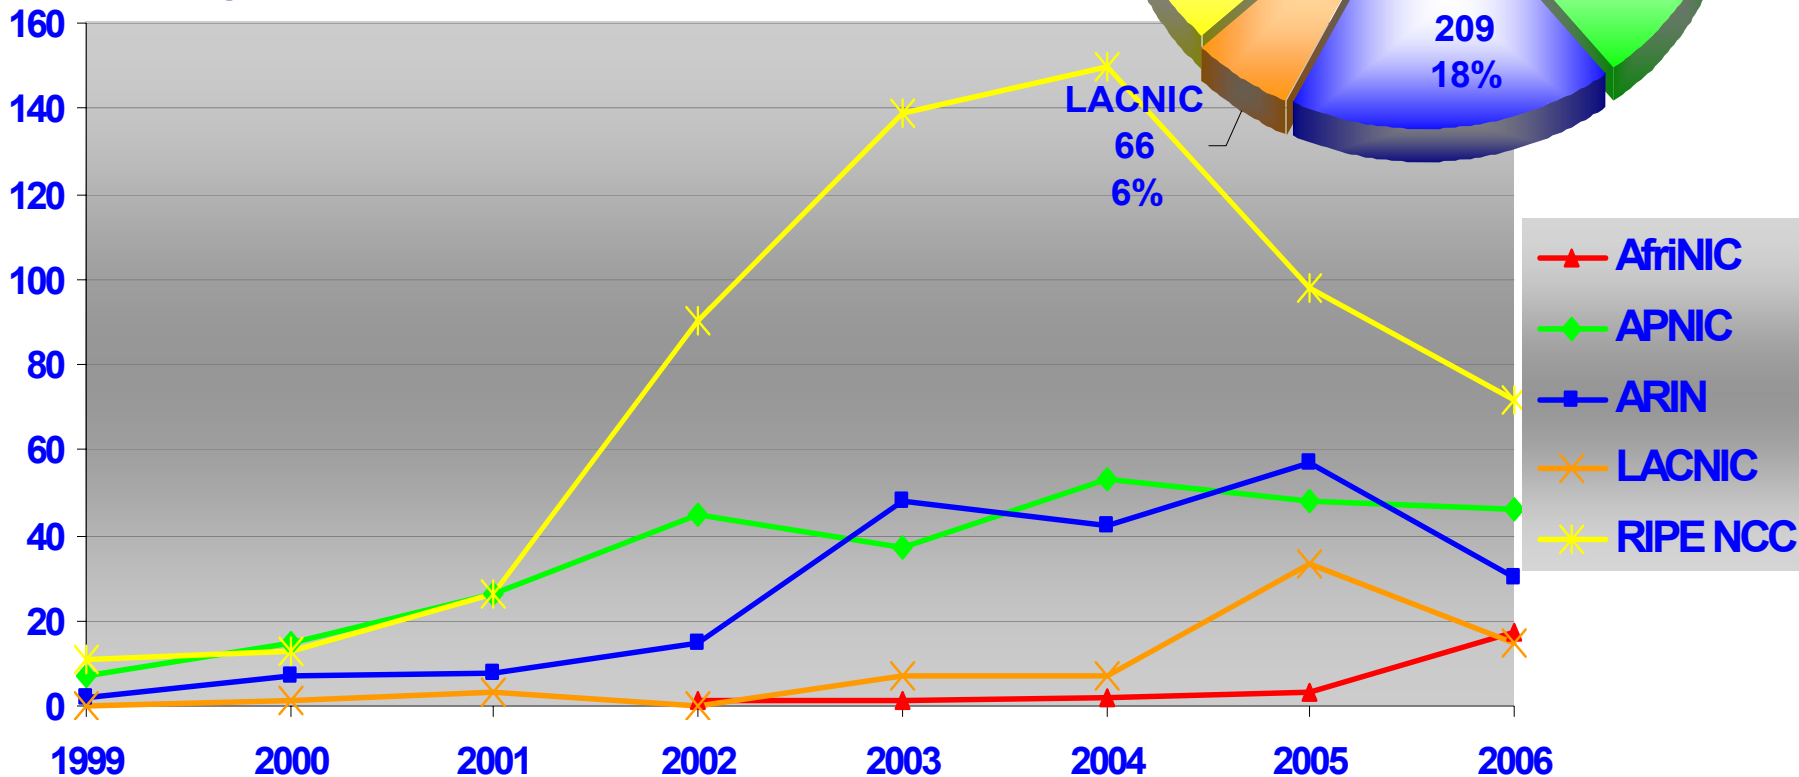
## RIRs to LIRs/ISPs

### Yearly Comparison

### Cumulative Total (1999–2006)




### Yearly Comparison







# Global IPv4 Usage in /8s



Source: <http://www.potaroo.net/>

October 2006



# IPv6: Summary

- IPv6 the future protocol?
- Be prepared for it!

## Questions?

# Routing Registry and Related Services

- Documenting Routing Policy
- AS Numbers
- Practical Usage of RPSL & Existing Tools
- The Routing Information Service (RIS)
  - The Routing Registry Consistency Check

# Documenting Routing Policy



# Elements of the Routing Policy

- Who are my BGP neighbours?
  - Customers / peers / upstreams
- What routes are
  - Originated by each neighbour?
  - Imported from / exported to each neighbour?
  - Preferred when multiple routes exist?
  - **How are routes treated (modifying routing parameters)**
- What to do when no route exists?



# Why Document Routing Policy?

- Recreate your policy in case of loss of hardware / administrators
  - Less downtime
- Scaling
- Troubleshooting



# Routing Policy Specification Language

- Abstract
  - Not vendor specific
  - Object-oriented: route, aut-num, as-set
- Global view, not router specific
  - Describes relations between BGP peers
- Established standard
  - RFC-2622, RFC-2725, Using RPSL in Practice (RFC-2650)
  - RFC-4012 (RPSLNg: for IPv6 and multicast)
- Tools available for translation to router configuration



# Why Publish Your Policy in IRR?

- Required by some Transit Providers
- Required by some Exchange Points
  
- Allows peers to automatically update filters
  - For your announcements
  - Consistent information between neighbours
  - Gives (limited) routing security
  
- Good housekeeping





# Why RIPE DB as Your RR?

- RIPE RR is part of the **Internet Routing Registry**
  - <http://www.irr.net/>
- Convenience – all objects already there
  - inetnum, aut-num, maintainer, person
- Strong security
- It's free!!!
- Database most likely used by your peers



# Documenting Policy: Summary

- Use RPSL to document your policy
- Use RIPE Routing Registry to publish your policy

## Questions?


# Autonomous System Numbers

# Requesting an ASN

- Autonomous System (RFC-1930)
  - “An AS is a connected group of ... IP prefixes ... which has a **single** and **clearly defined** routing policy.”
- LIR can request an ASN
  - For own network, or for another organisation
- Assignment criteria: multihomed
  - Unique routing policy
  - E-mail addresses of peers



# Building an Aut-num object



Internet

AS3

AS2

AS1

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

```
aut-num: AS1
```

```
export: to AS2
import: from AS3 action pref=100;
        accept ANY
import: from AS2 action pref=20;
        accept AS2
export: to AS3 announce AS1
import: from AS2 action pref=200;
        accept ANY
```

```
action aspath.prepend (AS1, AS1); announce AS1
```

```
aut-num: AS2
```

```
import: from AS1 accept AS1
export: to AS1 announce ANY
```



# aut-num Object

- RIPE NCC creates aut-num object
  - mnt-by: LIR-MNT
  - mnt-routes: End-User-MNT (or LIR)
  - org: ORG-End-User-RIPE (or LIR)
- When the peering is established, LIR should update routing policy
- AS Number assignment is only valid as long as the original criteria remain valid



# route Object

- **route** objects: part of Routing Registry
- LIR creates **route**/**route6** objects for any (new) allocations they announce
  - Both “**route:**” and “**origin:**” are primary key
  - Complex hierarchical authorisation for creation
  - Used for prefix filtering by some ISPs

# RIPE NCC Creating Route Objects

```
inetnum: 80.1.0.0 - 80.1.255.255
```

```
mnt-by: RIPE-NCC-HM-MNT
```

```
mnt-routes: LIR2-MNT
```

*allocation*

```
mntner: LIR2-MNT  
auth: MD5-PW $bla
```

```
aut-num: AS2
```

```
mnt-by: LIR2-MNT
```

```
route: 80.1.0.0/16
```

```
origin: AS2
```

```
mnt-by: LIR2-MNT
```







# ASN: Summary

- If you want to be multihomed
- Describe your routing policy
- **Route** object recommended

# Questions

# Practical Usage of RPSL & Existing Tools

# RR Policy Expression

- Aut-num
  - Lists neighbors (in import / export lines)
  - Defines filter rules for each neighbour
  - Defines route parameters modifications per prefix
  
- Route object
  - Represents address range originating by ASN
  
- Set objects
  - Grouping objects with similar policy / usage



# Preliminary Work

1. Create **person** and **mntner** objects
2. Create **route** objects in the database
3. Create various **as-set** objects, to group different categories of neighbours
4. Describe policy in your **aut-num** object
5. Create RtConfig (or other tool) template file(s)
6. Run RtConfig / tools / scripts periodically



# IRRToolSet: RtConfig & AOE

- Started as RAToolSet
- Changed to IRRToolset
  - first maintained by RIPE NCC
- Now maintained by ISC
  - <http://www.isc.org/index.pl?/sw/IRRToolSet/>
- Download: <ftp://ftp.isc.org/isc/IRRToolSet/>
- Installation needs: lex, yacc and C++ compiler



# RtConfig: Router Configuration

- RtConfig reads policy from the IRR
- Generates **parts** of the router configuration file
  - Creates access list, route-map and AS path filters
  - Vendor specific
- You need to use other scripts (built around it)!
- One of the tools in the IRRToolSet

# RtConfig: RR Integration

RPSL DB Objects  
(routing policy)

Commands in the  
Template/Input File



```
graph TD; A[RPSL DB Objects (routing policy)] --> C[RtConfig]; B[Commands in the Template/Input File] --> C; D[Flags, Env_Var] --> C; C --> E["(Partial) Router Configuration"]
```

**RtConfig**

Flags, Env\_Var

(Partial)  
Router Configuration



# Example RtConfig Template File

! setting up the Max\_Preference to 100

```
@RtConfig set cisco_max_preference = 100
```

```
send community
```

!

! Peering with OTHERCOMPANY (AS2)

```
@RtConfig set cisco_map_name = "AS%d-IMPORT-%d"
```

```
@RtConfig import AS1 10.0.0.1 AS2 10.0.0.2
```

!

```
@RtConfig set cisco_map_name = "AS%d-EXPORT-%d"
```

```
@RtConfig export AS1 10.0.0.1 AS2 10.0.0.2
```





# Additional Work: New Neighbour

- Your neighbour needs to:
  - Obtain and register an ASN
  - Create route objects for the new AS
  
- You need to:
  - Add the new AS to one of your as-set objects
  - Create RtConfig template for the peering
  - Run your scripts



# AOE: Aut-num Object Editor

- Makes **aut-num** editing easier
- Takes input from:
  - Your **aut-num** object
  - Your peer's **aut-num** object
  - BGP dump
  - Pre-defined and user-defined “templates”
- Sends mail with updated Aut-num object
  - Does not sign

# More RR-related Tools

- Nemecis

- *Georgos Siganos and Michalis Faloutsos*
- analyzes & checks for correctness and freshness of registered policy
- <http://ira.cs.ucr.edu:8080/Nemecis>

- RPSL Analysis Service

- *Giuseppe Di Battista, Tiziana Refice, Massimo Rimondini*
- [http://tocai.dia.uniroma3.it/~irr\\_analysis/](http://tocai.dia.uniroma3.it/~irr_analysis/)



# Even More RR-related Tools

- Routers configuration / BGP filters generator
  - *Marco d'Itri*
  - uses RPSL data, a local configuration file and a configuration template based on the popular Template::Toolkit perl package
  - <http://www.linux.it/~md/software/rpsltool-1.0.tgz>
- IRR Power Tool
  - *Richard Steenbergen*
  - a PHP and CVS-based framework for using IRR data to manage prefix-lists
  - <http://sourceforge.net/projects/irrpt/>

# Tools: Summary

- The quality of filters generated by tools strongly depends on the data you have in the RR!
  - It is crucial to maintain RR objects up-to-date
  
- Using the tools will help you to benefit from registering your data in RR, to achieve:
  - Automating access-list generation
  - Avoiding mistakes
  - Improving configuration / operation process

# Routing Information Service (RIS)

# Gigantic Looking-glass with History

- Database of information about routes and their development **over time**
- Populated by collecting & storing time-stamped BGP announcements, from the default-free core
- “Remote Route Collectors” at several major IXes
  - ~637 peers (~80 also IPv6)
- Aimed at: NOC and ISP engineers, research community
- Similar to routeviews
  - <http://www.routeviews.org/>



# RIS Applications

- Debugging
  - e.g. Checking why customer route **was** not available
  - Verify local policies vs router setup => correct errors
  - Prefix distribution: aggregation, correct filters
  
- Analysis
  - Routing table convergence times; route flaps
  - Comparing RR policies with actual announcements
  - More: <http://www.ripe.net/ris/analysis.html>





# Existing RIS-related Tools

- **BGPlay** visualises routing updates
- ASInUse / PrefixInUse
  - Last appearance of ASN / prefix in global routing table
- Search by AS / prefix
  - ASN / prefix activity during a particular time interval
- Looking Glass (also for IPv6)
- Weekly reports
  - BGP Traffic Hot Spots
    - prefixes with high activity
  - RIS Martians
- Raw data available

- Quick & summarised view of prefix visibility
- Lists which route collectors see the prefix
- Maps prefix to originating ASN
- `whois -h riswhois.ripe.net [ip/range]`

```
route:          212.3.64.0/19
origin:         AS8900
descr:          Global One Hungary Internet
upd-first:      2003-11-29 22:15Z  164.128.32.11@rrc09
upd-last:       2003-12-04 17:09Z  193.203.0.52@rrc05
seen-at:        rrc05,rrc09
source:         RISWHOIS
```



# MyASn

- Notification system for BGP
  - Generates alerts based on user settings (e.g. AS path)
- Alarms you when...
  - another AS announces your prefix; or
  - someone gives you transit while you didn't expect that.
  - Alarms through email or syslog
- GUI - configuration:
  - Hold-down Time, Hold-down Event, Time To Live
  - Alarms can be grouped, with group notification settings
- Sign up: <http://www.ris.ripe.net/myasn.html>



# Routing Registry Consistency Check

- The goal: making RR more accurate
  - Comparing “real” routing data (via RIS) with the RR
  - Spotting inconsistencies & suggest corrections
- Data output
  - Web interface for interactive lookups
  - Reports per mntner (requests to <auto-rrcc@ripe.net>)
  - Published on the web, reported to the routing-wg
  - Scripts available
- <http://www.ripe.net/rrcc/> , <rrcc@ripe.net>



# RIPE: Summary

- Use RIS to get global view of your prefixes
- <http://www.ripe.net/ris/>, [ris@ripe.net](mailto:ris@ripe.net)
- Use RRCC to find possible errors in RR

## Questions?

# News

- 32-Bit AS Numbers
- The Routing Table Growth & Other Statistics
- The Resource Certificates



# Running out of AS Numbers

- Consumption rate studies:
  - Rene Wilhelm: ASN-MIA, RIPE50
  - Geoff Huston: AS Numbers, RIPE51
  
- Assignment rate is 10-12/day
  
- ASN pool exhausted
  - Sometime between 2010 and 2013
  - September 2006:  $\approx$ 22,000 ASNs left

# More AS numbers: 32 bits ASNs

- Details: draft-ietf-idr-as4bytes-11.txt
  - Proposed standard, in IESG queue
  - Implementations exist (sort of)
- Transition mechanism exist
  - Existing BGP speakers continue to work
  - New BGP speakers will have to use 32 bit ASN
  - No flag date, mixed world
- No changes to policies or request procedure
  - Requirement: multihoming





# ASN32 policy: PDP 2005-12

- **1 January 2007 – 31 December 2008**
  - LIR can ask for a 16 or 32 bit AS Number
  - RIR will assign a 16 bit AS Number by default
    - 32 bit on request
- **1 January 2009 – 31 December 2009**
  - LIR can ask for a 16 or 32 bit AS Number
  - RIR will assign a 32 bit AS Number by default
    - 16 bit on request
- **After 1 January 2010**
  - RIR will always assign a 32 bit AS Number



# What should you do

- Start thinking about 32 bits AS Numbers
  - In your organisation
  
- Ask your vendor for support
  - or be prepared for a nasty surprise in 2009
  
- Don't wait until 2009!
  - Can you handle your new ASN **1.5432**?

# Let's request a 32 bit ASN!

## AS Number Request Form


```
#[GENERAL INFORMATION]#  
#[AS NUMBER USER]#  
#[ADDRESS SPACE TO BE ANNOUNCED]#  
#[PEERING CONTACTS]#  
#[DATABASE TEMPLATE(S)]#  
    aut-num: ASNEW  
  
#[INSERT SUPPLEMENTAL COMMENTS]#  
    I'd like a 32 bit ASN, please!  
  
#[END of REQUEST]#
```




# ASN Assignments

## Yearly Comparison

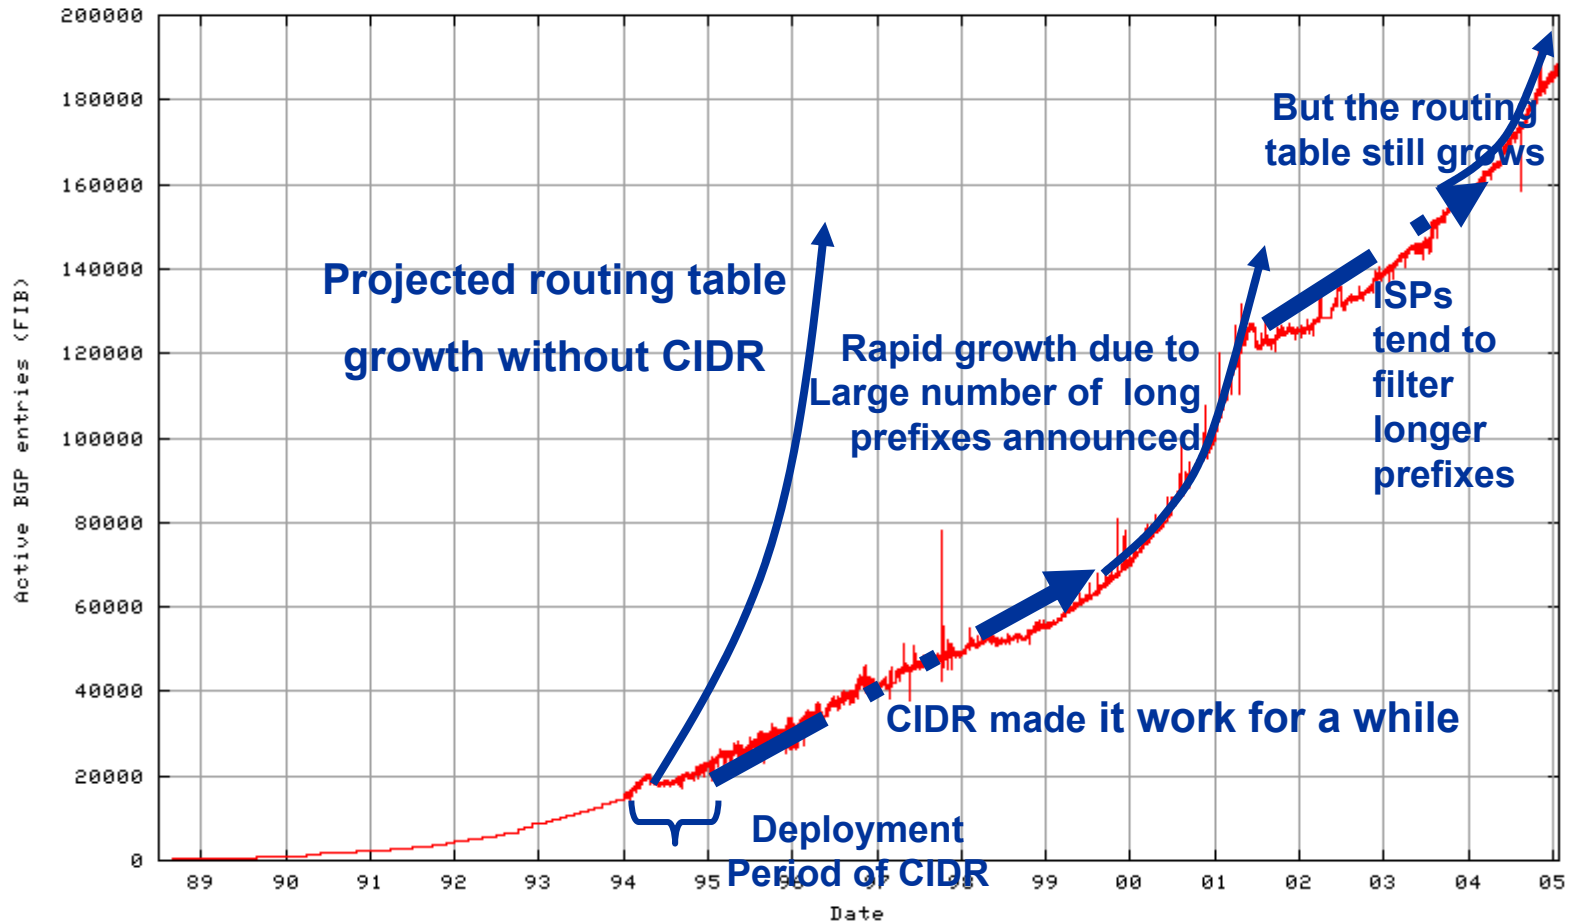
### Cumulative Total (1999–2006)



### Yearly Comparison



# Growth of Global Routing Table



last updated 27 January 2005



# A Friday the 13<sup>th</sup> NANOG E-mail

- October 2006: <http://www.cidr-report.org/>
- BGP routing table entries examined:
  - **200339**
- Prefixes after maximum aggregation:
  - 108814
- RIPE53: “RIPE Routing-WG Recommendations on Route Aggregation “



# But what about certificates?

- Resource certificates are coming
  - Secure BGP?
- Task force started at RIPE 53
- Too soon to predict the outcome
- Until then:
  - Routing Registry is here now!



# News: Summary

- Get ready for 32-bit AS Numbers
- Aggregate!
- Take part in the resource certification efforts

## Questions?





# Homework

- Subscribe to mailing lists:
  - db-wg, routing-wg, policy-announce [[@ripe.net](mailto:@ripe.net)], [irrtools@isc.org](mailto:irrtools@isc.org)
- Use Routing Registry to:
  - Create route objects for your allocations
  - Update your aut-num with the latest policy
- Use the LIR Portal to:
  - Add mnt-routes to your allocation
  - Get the resources you need
  - Keep your LIR info up-to-date
  - Register for RIPE NCC courses



# The End!

Y Diwedd

النهاية

Соңы

Край

Finis

Liõugt

Ende

Кінець

Konec

Kraj

Fund

Son

پایان

Край

Fine

Lõpp

Vége

An Críoch

අප  
මගේ  
අවසාන

Endir

Sfârșit

Fin

Τέλος

Einde

Конец

Amaia

Pabaiga

Slut

Slutt

Fim

Loppu

Tmíem

Koniec

# Bonus Track 1

## Bogon Filtering

# Security

- Problems:
  - Bogon address space used as source for spamming, DDoS, probes...
  - Leaking “martians” & bogons due to mis-configuration
- Definitions:
  - Martians – reserved ranges (rfc-1918)
    - <http://www.isi.edu/~bmanning/dsua.html>
  - Bogons – un-allocated (&reserved) address ranges
- Secure BGP Template
  - [www.cymru.com/Documents/secure-bgp-template.html](http://www.cymru.com/Documents/secure-bgp-template.html)



# Applying “bogon” Filters

- RPSL: add “AND NOT fltr-bogons” to all your import and export attribute *filter* rules

```
aut-num:      AS1
```

```
import:      from AS1:AS-CUSTOMERS accept  
              PeerAS AND NOT fltr-bogons
```

```
import:      from AS1:AS-UPSTREAMS accept  
              ANY AND NOT fltr-bogons
```

```
export:      to AS1:AS-CUSTOMERS announce  
              ANY AND NOT fltr-bogons
```

```
export:      to AS1:AS-UPSTREAMS announce  
              AS1 AS1:AS-CUSTOMERS AND NOT fltr-bogons
```



# Example “filter-set”: fltr-bogons

```
filter-set:    fltr-bogons
descr:        All bogon IPv4 prefixes.
filter:        fltr-unallocated OR fltr-martian
tech-c:        RTH32-ARIN
admin-c:       RTH32-ARIN
mnt-by:        MAINT-BOGON-FILTERS
changed:       radb@cymru.com 20040420
source:        RIPE
```

```
filter-set:    fltr-unallocated
filter:        {1.0.0.0/8^+, 2.0.0.0/8^+, [...] }
```



# Outdated Bogon Filters

- Bogon filters in place, but not kept up-to-date
  - Consequence: new networks unreachable
- Solutions:
  - Use fltr-bogons
  - Check for RIR announcements of new /8 blocks
  - Use a bogon route server



# Bogon Filtering: Summary

- Keep your bogon filters up-to-date!
- Use filter-set objects
- Add “AND NOT fltr-bogons”

## Questions?





# Bonus Track 2

## RPSL: Communities



```

# Multihomed customers, backup route
# match community 3007:30, pref=30, localpref=970
import: from AS3007:AS-BGP-CUSTOMERS
        action pref=30 ;
        accept
        community.contains (3007:30) } Filtering on the community
        AND PeerAS; } value set by customer

# Announce only to customers (not to peers)
import: from AS3007:AS-BGP-CUSTOMERS
        action community = {no_export}; Setting community
        accept
        community.contains (3007:440) } Filtering on the community
        AND PeerAS; } value set by customer

```

There is an implicit logical OR when combining filter rules in aut-num!

- Therefore, an explicit “AND” has to be used!

```
# Prepend own ASN 2 times when announcing to peers
# (preferred route is through another ISP)
export: to AS3007:AS-PEERS
        action aspath.prepend (AS3007,as3007);
announce
community.contains (3007:112)
AND AS3007:AS-BGP-CUSTOMERS
```

} Filtering on the community value set by customer



# Customer Setup

- Controlling in- & out-bound traffic using “pref” value and communities

```
import: from AS3007 accept
        AS3007
        AS3007:as-bgp-customers
        AS3007:AS-upstreams
export: to AS3007
        action community = {3007:20, 3007:332};
        announce AS2031
        # prepend AS3007 twice to EU peers
```

# Communities in RPSL: Summary

- Filtering on import / export routes:
  - `accept community.contains (3007:30) AND PeerAS;`
- Setting on export routes:
  - `action community = {no_export};`
  - `announce community.contains (3007:112) AND PeerAS`
  - `action community = {3007:20, 3007:331};`
- Cisco IP Journal: “Application of BGP Communities”
  - [http://www.cisco.com/web/about/ac123/ac147/archived\\_issues/ipj\\_6-2/bgp\\_communities.html](http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_6-2/bgp_communities.html)