# Internet Resource Certification and Origin Validation

*An approach to more secure routing on the Internet*



Carlos Martinez – Cagnazzo
carlos @ lacnic.net

# Agenda

- Traffic routing on the Internet
- Route Hijacking
- Current counter-measures
- Resource certification
- Origin validation
- References

# Traffic Flow on the Internet

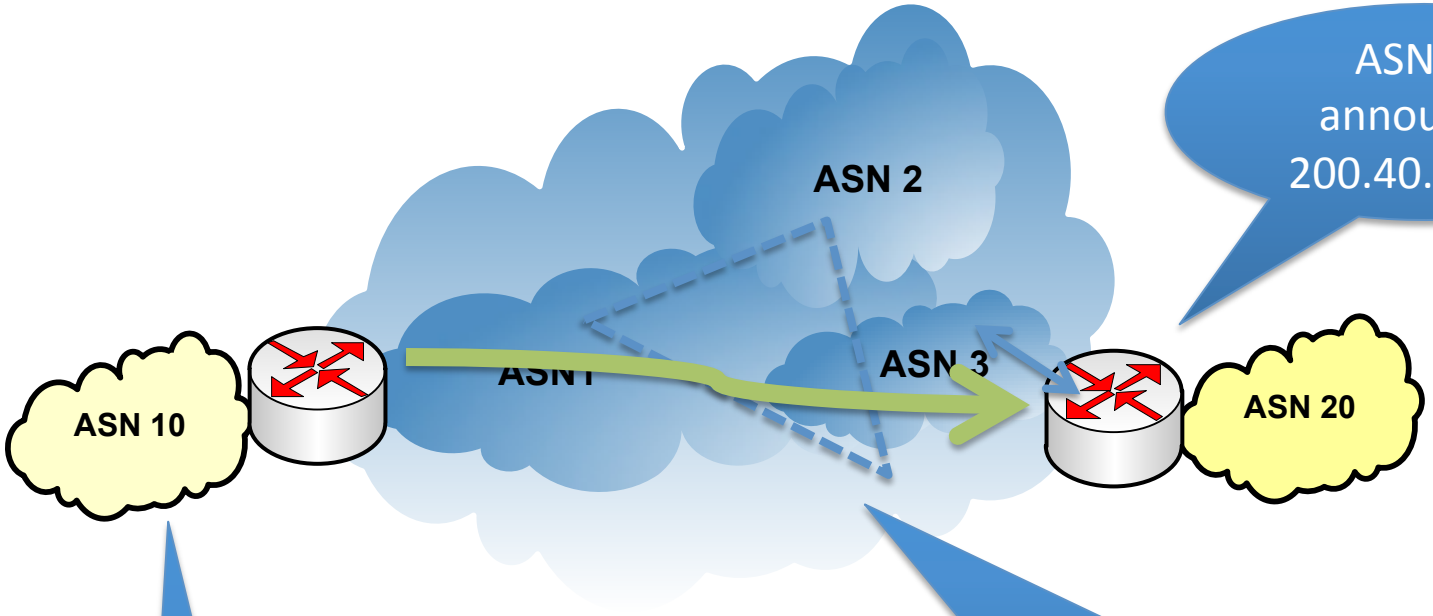


The 200.40.0.0/16 prefix propagates accross ASs (via BGP sessions)

ASN 10 receives announcement for 200.40.0.0/16

Each router applies a **decision algorithm**

Each announcement carries a set of attributes:
200.40.0.0/16 **AS_PATH** ASN1 ASN3 **ASN6057**

# Routing in the Internet (ii)

- BGP chooses routes using a **decision algorithm** and the values of the set of available **attributes**
- AS_PATH is a list of the autonomous systems a given UPDATE has traversed
  - The first entry is the AS originating the route (hence "origin-as")

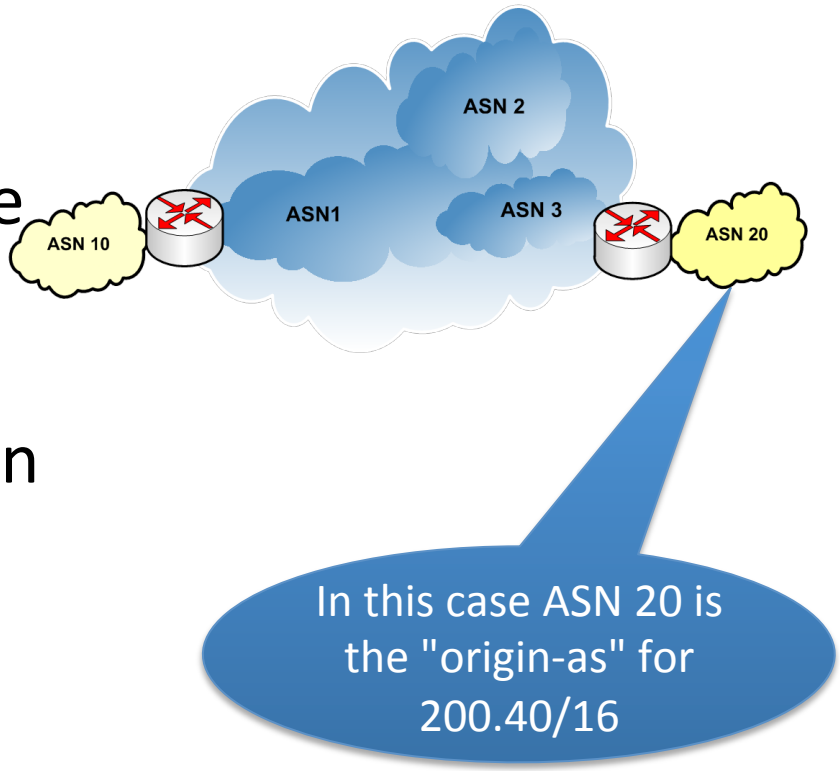

# Route Hijacking

- When an entity participating in Internet routing announces a prefix without authorization We face a *route hijack*
- Malicious or due to operational mistakes
  - Most of the time you just can't tell
- Some well-known cases:
  - Pakistan Telecom vs. You Tube (2008)
  - China Telecom capturing traffic to/from the U.S. (2010)
  - Google in Eastern Europe (various ASs, 2010)
  - Some occurrences in LACNIC's service region (January/ February 2011)
  - **One ongoing occurrence (CL – CO)**

# Route Hijacking (ii)

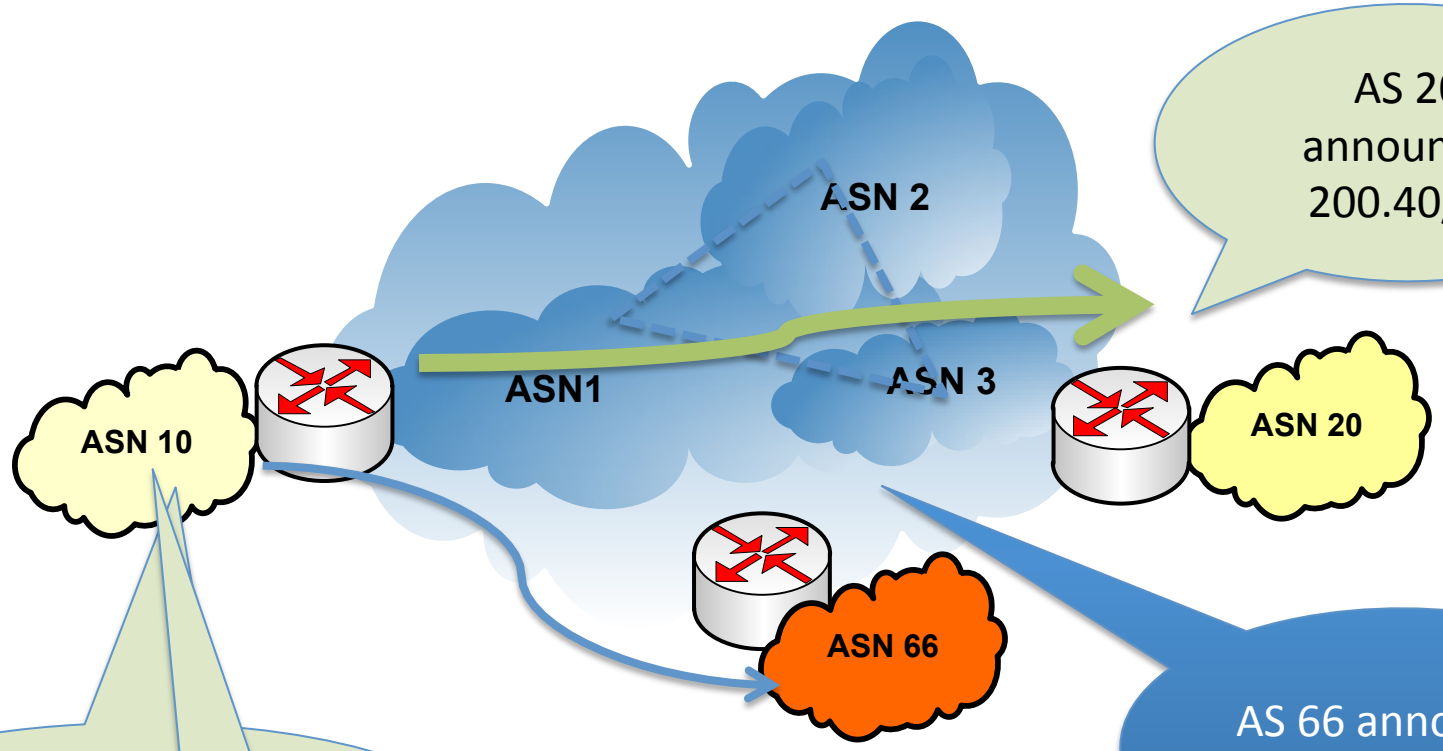


200.40.0.0/16 AS_PATH ASN1 ASN3 **ASN6057**
200.40.235.0/24 AS_PATH ASN1 ASN3 **ASN6057**

# Route Hijacking (iii)

- RIPE NCC Video of the YouTube incident
  - http://www.youtube.com/watch?v=IzLPKuAOe50

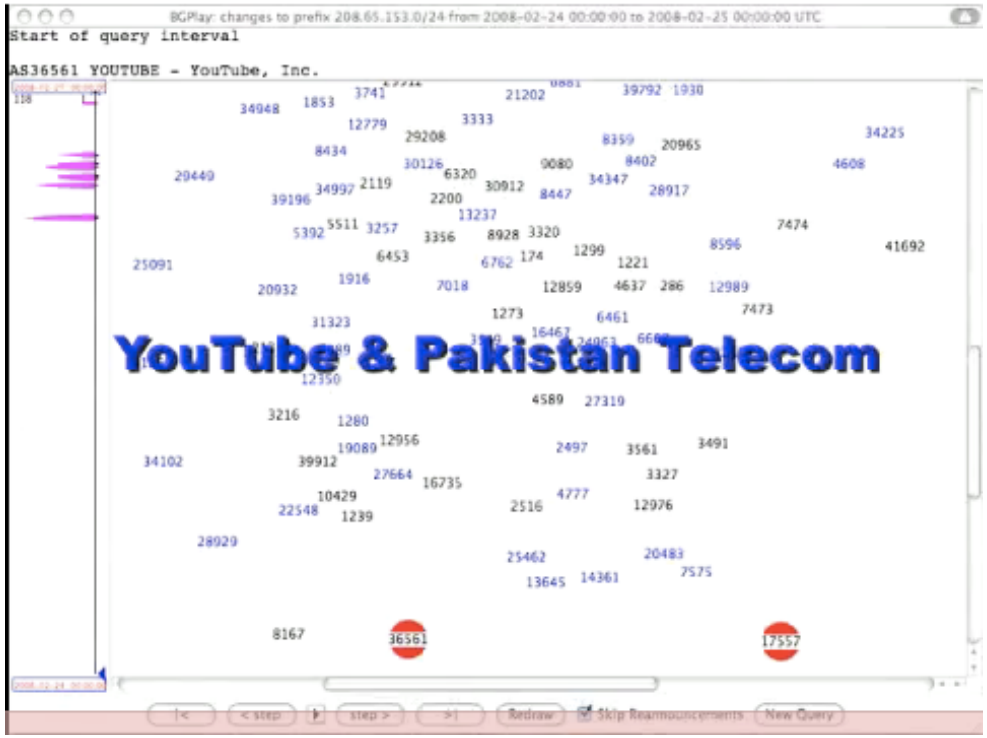

# Route Hijacking Mitigation Current Practices

# Peering Relationships

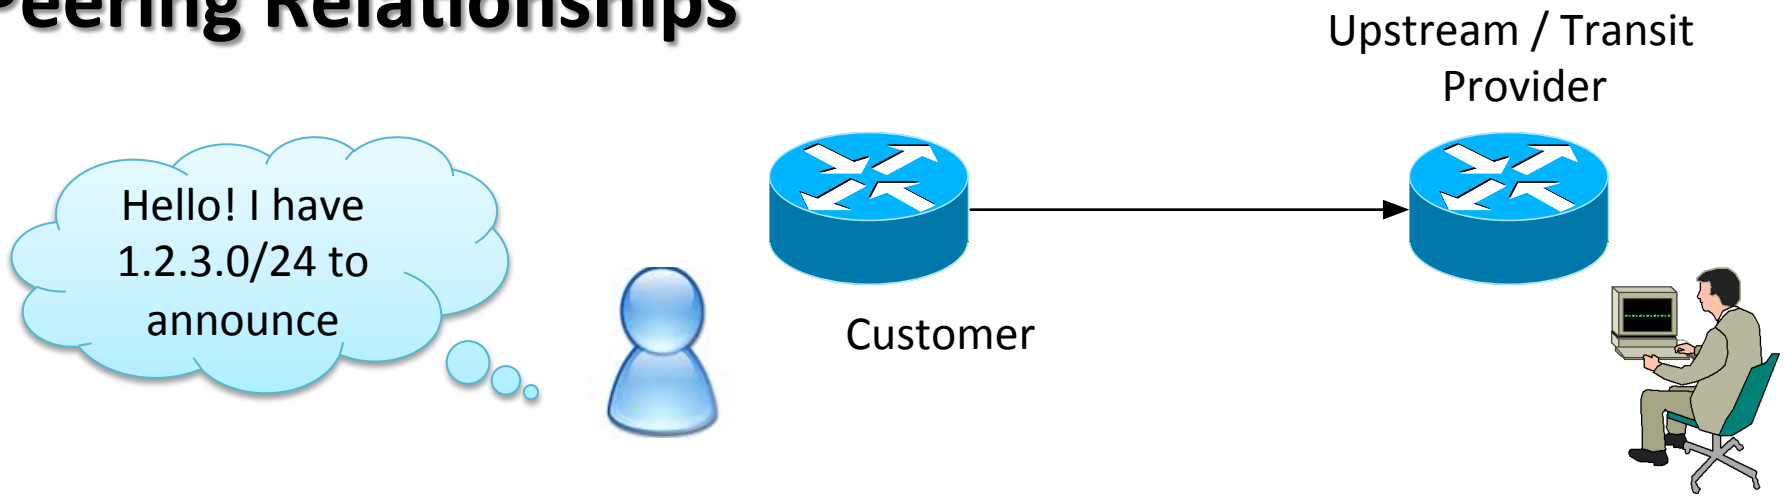


- Upstreams should check whether customers are authorized to announce resources
  - Some ask for an email to be sent to a specific address, others ask for a web form, others ask for entries in IRRs, others check WHOIS
  - Not consistent, varies from carrier to carrier
    - Sometimes *from customer to customer of the same carrier*

# Peering Relationships (ii)

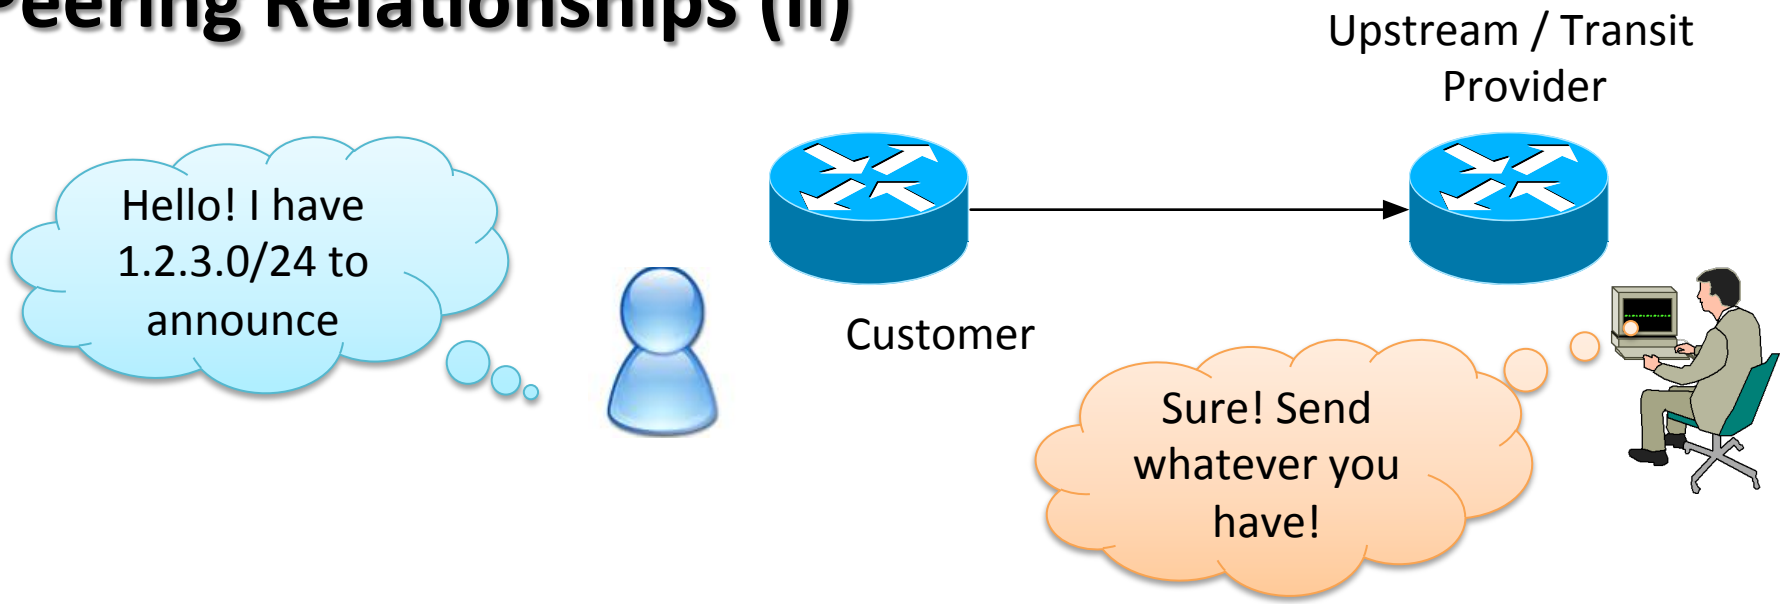


- In the end the integrity of the routing system depends on **ad-hoc trust relationships between peers**
- The problem lies in that
  - Checks are inconsistently applied
    - Sometimes no verification at all is performed
  - Current tools are ill-suited for automating this process

# Peering Relationships (iii)

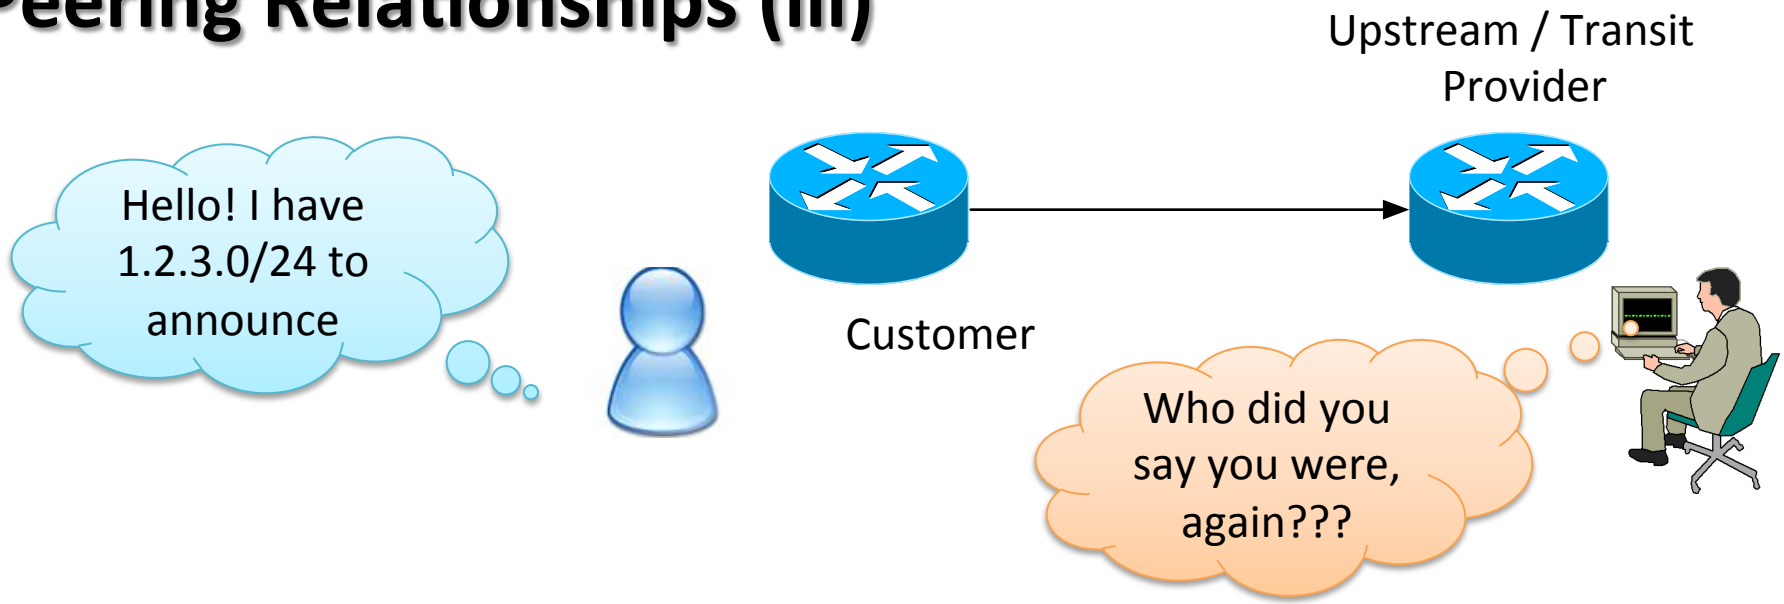


- Other recommended practices include
  - uRPF filtering where applicable
  - Routing protocol integrity
    - Peer authentication w/ MD5 passwords
- Filtering known-invalid routes
  - Filter RFC 1918 and other well-known bogons

# Resource Certification and Origin Validation

# Internet Number Resource Management

- What do we mean by resources?
  - IPv4, IPv6 Addresses, ASNs
- Five regional registries
  - AFRINIC, APINIC, ARIN, LACNIC
  - RIPE-NCC

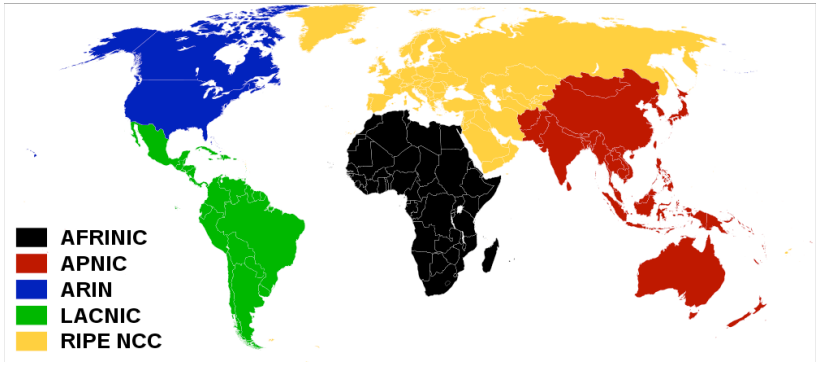


- One central pool: IANA
- Each RIR is the **authoritative source** on the relationship between **users/holders** and resources
  - Each RIR operates a registry database
  - Each RIR has a **contract** with the organizations receiving resources

# Internet Number Resource Management (ii)

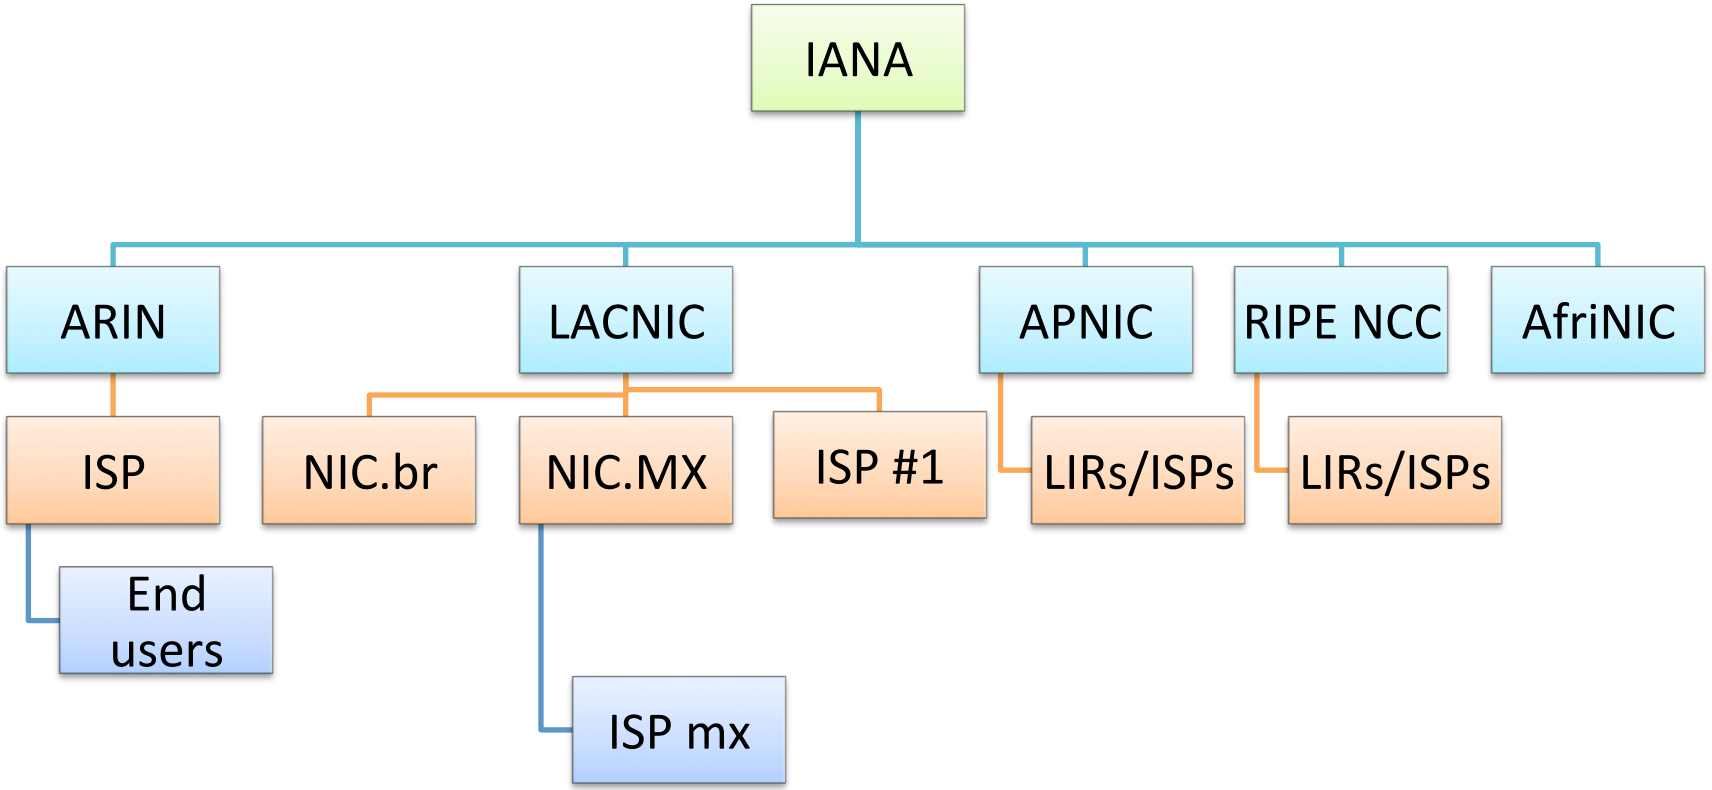

# Resource Public Key Infrastructure (RPKI)

- Goals:
  - Create cryptographic proofs (certificates) that serve as proof of resource holdership
  - Enable automatic verification of route announcements in routers
- High-level overview
  - Use of X.509 v3 certificates
  - Use RFC 3779 extensions on these certificates. These extensions allow Internet resources (IPv4/IPv6/ASNs) fields within certificates
  - ROAs: Signed objects that contain origin AS data.
  - Mechanisms to push validated data to routers and to automatically check the “origin-as” of a BGP UPDATE

# Resource PKI (ii)

- All RPKI signed objects are listed in public repositories
- After verification, these objects can be used to configure policy in routers
- Validation Process
  - Signed objects have references to the certificate used to sign them
  - The resources listed in a certificate MUST be valid subsets of the resources listed in its parent's certificate
  - The trust chain is traced to the trust anchor in two aspects:
    - Cryptographically
    - CIDR terms

# X.509 Certificates with RFC 3779 extensions

- "IP Delegation" Section
  - Special value: "INHERITED"
- "AS Delegation" Section
  - Special value: "INHERITED"
- Validation Process
  - Traditional crypto validation
    - Signature chain up to the trust anchor
  - Additionally involves validation of resources
    - CIDR (AKA subnetting) inclusion

# RPKI Trust Chain

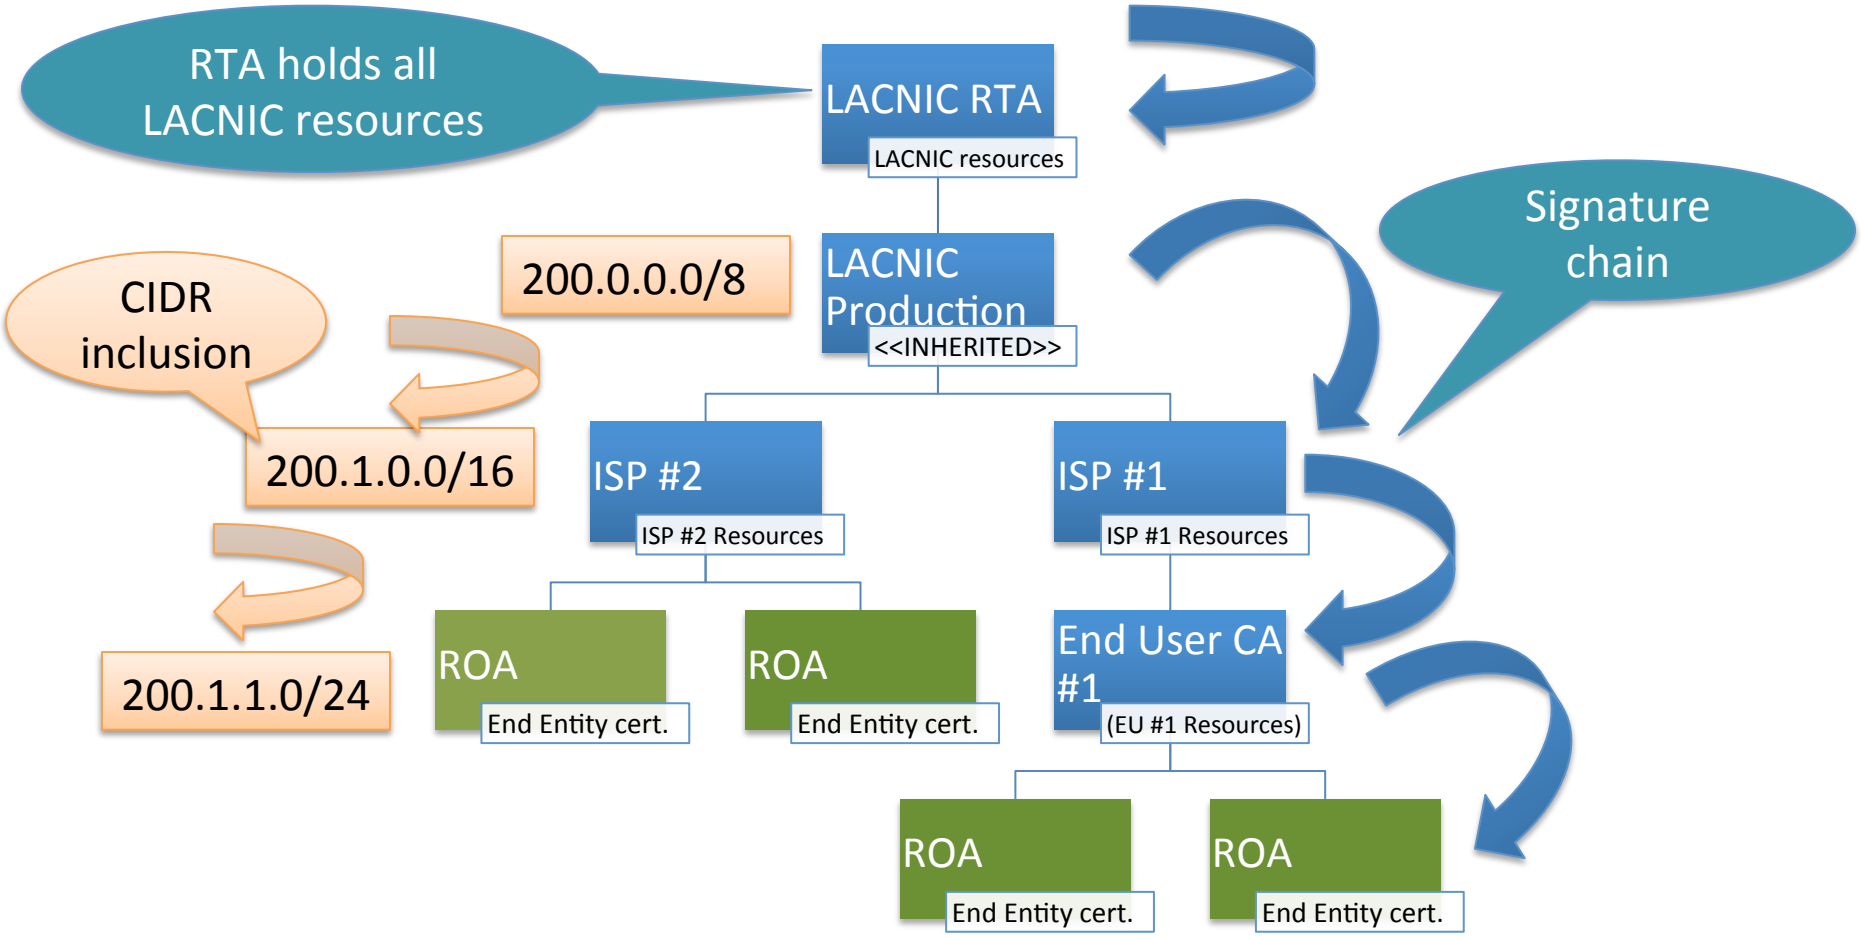

# Route Origin Authorizations

- A ROA provides a **signed** statement of **route origination**:

| Prefix | Max_Len | Origin_AS | Valid_Since | Valid_Until |
|---|---|---|---|---|
| 200.40.0.0/17 | 20 | 10 | 2011-01-02 | 2013-01-01 |
| 200.3.12.0/22 | 24 | 20 | 2011-01-02 | 2013-01-01 |

- The first ROAs states that:
  - *"The prefix 200.40.0.0/17 will be originated by ASN 10 and could be de-aggregated up to /20" "This statement is valid starting on Jan 2, 2011 until Jan 1, 2013"*
- ROAs also contain an EE certificate with the resources listed

# ROAs (ii) - Validation

- In order to validate a ROA three steps are performed
  - Crypto validation of the public keys and signatures included in the EE certificates inside each ROA
  - CIDR inclusion checking of resources listed in the EE certificate
  - CIDR inclusion checking of resources in the route origin attestations. These resources have to be included in the resources listed in the EE certificate

# RPKI in Action – The whole system

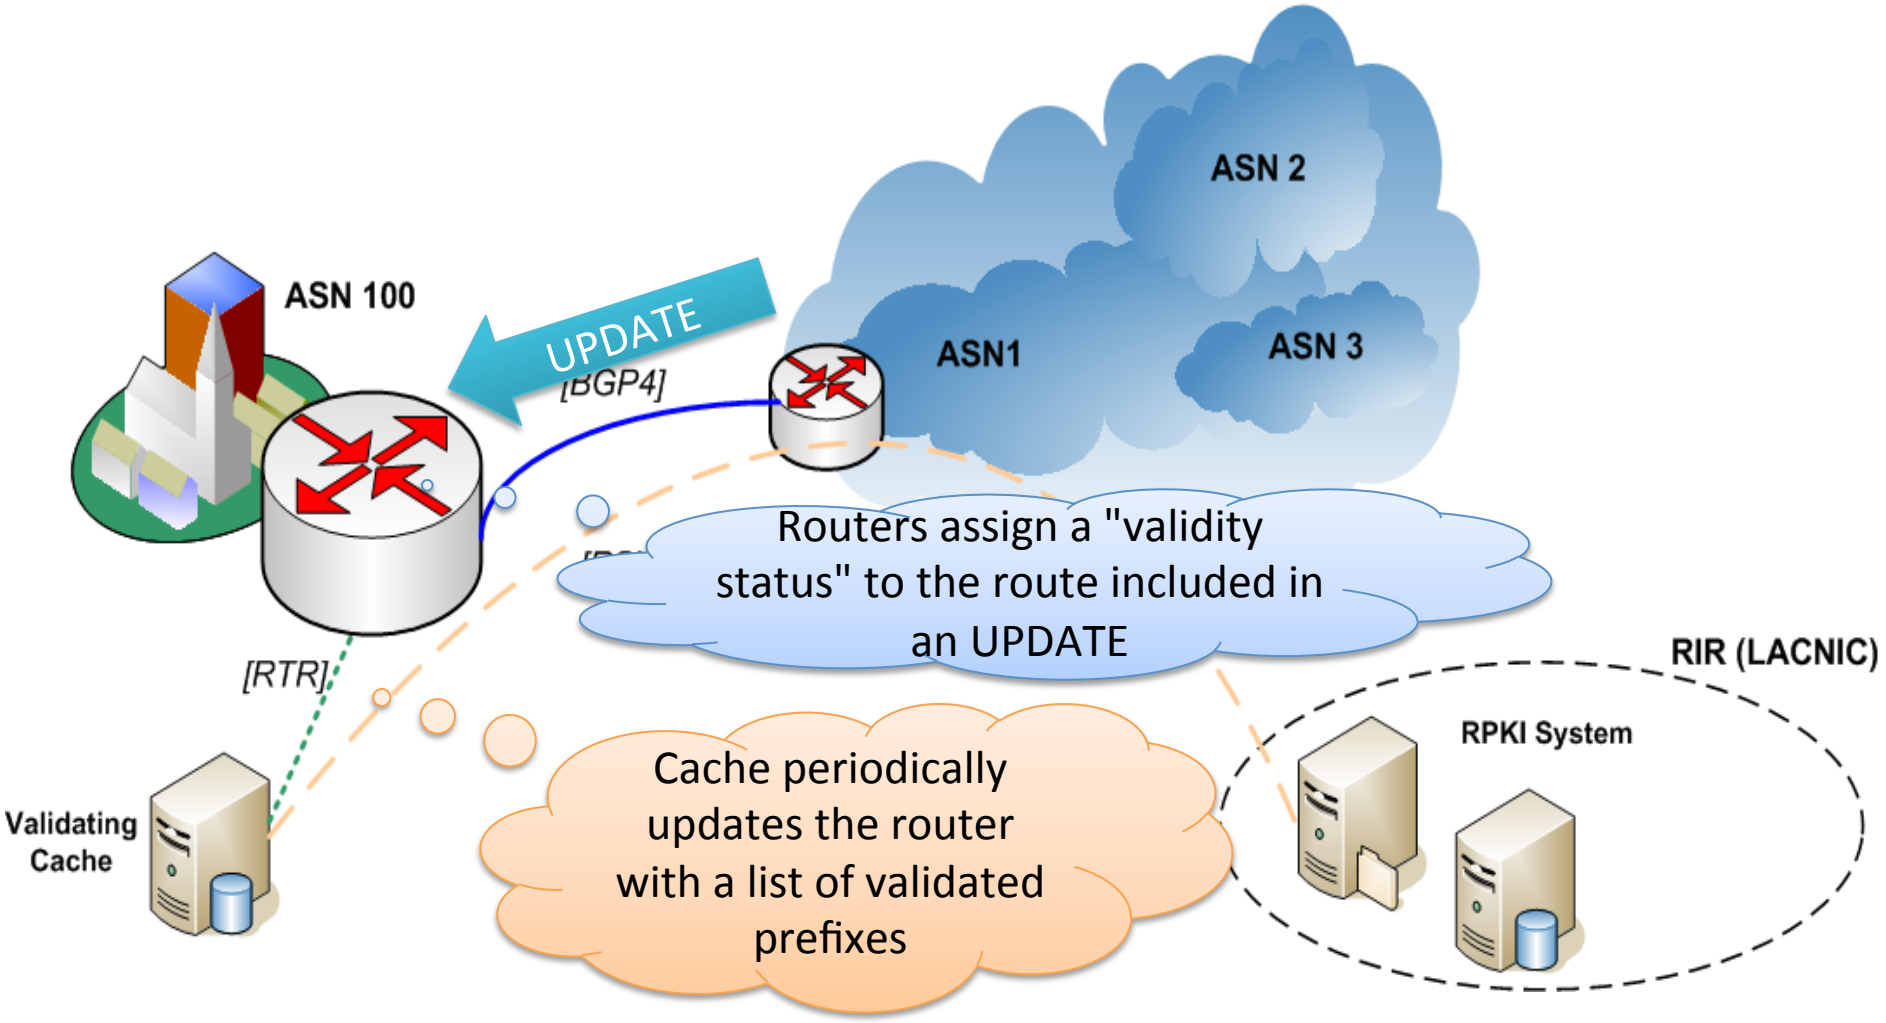

# BGP UPDATE Validation


| prefix/[min_len – max_len] | Origin AS |
|---|---|
| 172.16.0.0 / [16-20] | 10 |
| 200.0.0.0/[8-21] | 20 |

- If the "UPDATE pfx" is not covered by any entry in the DB -> "**not found**"
- If the "UPDATE pfx" is covered by at least one entry in the DB, and the origin-AS matches the ASNs in the DB -> "**valid**"
- If the origin-AS does NOT match -> "**invalid**"

# BGP UPDATE Validation (ii)

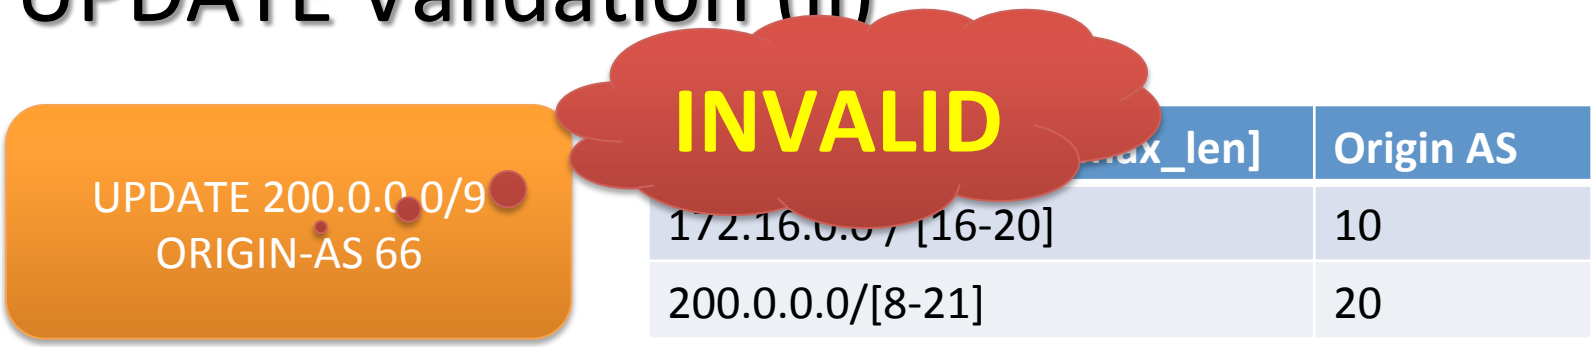


| [Max_len] | Origin AS |
|---|---|
| 172.16.0.0 / [16-20] | 10 |
| 200.0.0.0/[8-21] | 20 |

- If the "UPDATE pfx" is not covered by any entry in the DB -> "**not found**"
- If the "UPDATE pfx" is covered by at least one entry in the DB, and the origin-AS matches the ASNs in the DB -> "**valid**"
- If the origin-AS does NOT match -> "**invalid**"

# Links / References

- The LACNIC RPKI System
  - http://rpki.lacnic.net/
- LACNIC's RSYNC Repository
  - rsync://repository.lacnic.net/rpki/
- Listing the repository
  - rsync --list-only rsync://repository.lacnic.net/rpki/lacnic/
- Some RPKI Statistics
  - http://www.labs.lacnic.net/~rpki

# Thank You !

carlos @ lacnic.net