


RPKI Information Briefing For ICANN 66 DNSSEC & Security Workshop



Brief RPKI Description for DNS & DNSSEC Folk

Russ Mundy
Parsons

What is Routing Security?

- Border Gateway Protocol (BGP) & other routing specifications defined by IETF RFCs
 - Many routing protocols do not have any for recipients to know the info is correct
 - RFC4272 = BGP Security Vulnerabilities Analysis
 - IETF Secure Inter-Domain Routing (sidr) and SIDR Operations (sidrops) WGs publish RPKI specifications as RFCs
 - These RFCs provide IETF definition for routing security

Brief History of Routing Incidents

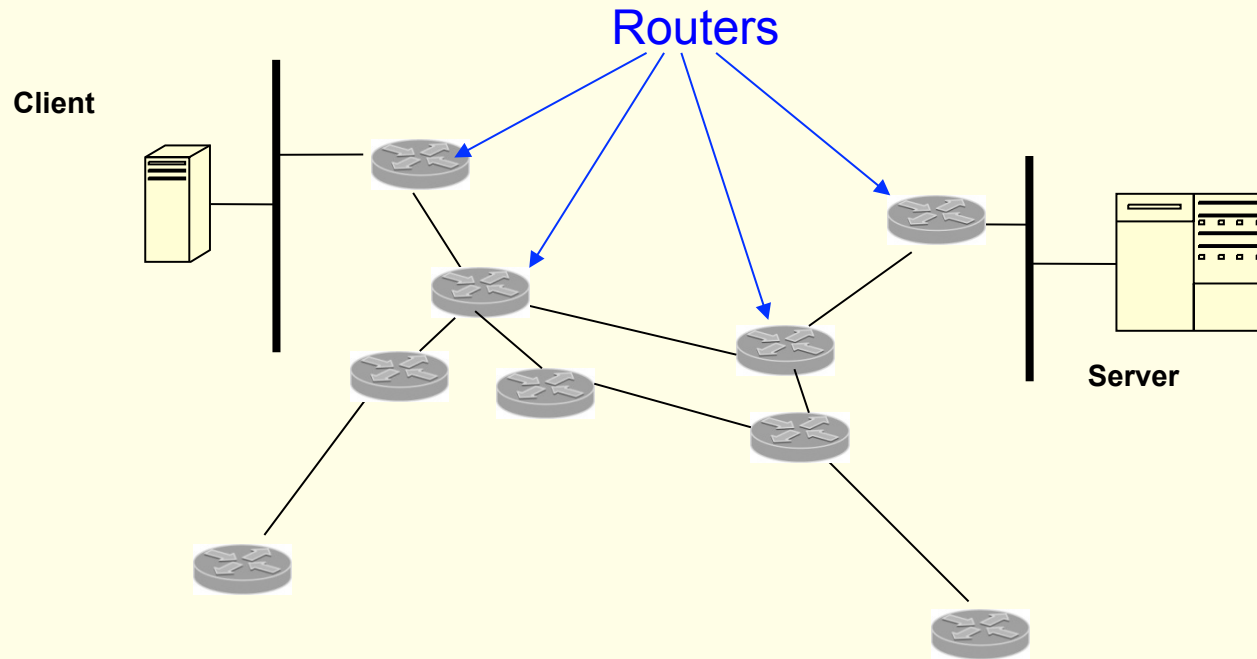
- **Apr 1997 – AS 7007 announced routes to all the Internet**
- Apr 1998 – AS 8584 mis-announced 100K routes
- Dec 1999 – AT&T's server network announced by another ISP – misdirecting their traffic (made the Wall Street Journal)
- May 2000 – Sprint addresses announced by another ISP
- Apr 2001 – Flag Telecom in London mis-announced 5K routes
- **Dec 24, 2004 – thousands of networks misdirected to Turkey**
- Feb 10, 2005: Estonian ISP announced a part of Merit address space
- **Sep 9, 2005 – AT&T, XO and Bell South (12/8, 64/8, 65/8) misdirected to Bolivia [the next day, Germany – prompting AT&T to deaggregate]**
- **Jan 22, 2006 – Many networks, including PANIX and Walrus Internet, misdirected to NY ISP (Con Edison)**
- Feb 26, 2006 - Sprint and Verio briefly passed along TTNET (Turkey again) announcements that it was the origin for 4/8, 8/8, and 12/8
- Jul 07, 2007 – Yahoo unreachable for an hour due to mis-origination to L3 from Hanaro Telecom
- **Feb 24, 2008 –Pakistan Telecom announces a part of YouTube's address blocks**
- Mar - Nov 2008 – various addresses within DoD address blocks announced by various ISPs (one in Russia, one in Argentina, others in Australia, Turkey, Indonesia, etc.) for periods up to 3 weeks
- Dec 2008 – Axtel in San Pedro, MX announces unallocated address block, and then sends a large amount of mail traffic (spam).
- Mar 2010 - For three weeks, the address of China's own internal version of the DNS root zone was advertised outside China. This made the altered China version of the root zone visible outside China (Asia, Chile, US, etc.)
- **April 2010 - China Telecom mis-originated about 15% of Internet address blocks**
- Jun 2010 – BGPmon reports bogon IPv6 announcements mis-originated by multiple ISPs to Cogent – no explanation
- Frequent full table leaks, e.g., Sep08 (Moscow), Nov08 (Brazil), Jan09(Russia), Jul 09 (Sweden)
- Jan 2012 – appeal for assistance in recovering address space from former customer

Steps To Improve Routing Security

Note: “Routing System Security” operators must provide for all BGP operators to be able to verify the following:

- Resources have been **allocated appropriately** (someone only uses what they’re supposed to or detect if not) - Step 1
- Resources are being **used where** they should be (only used at allowed Internet connection points or detect if not) - Step 2
- Verify **path taken** through the network – Step 3

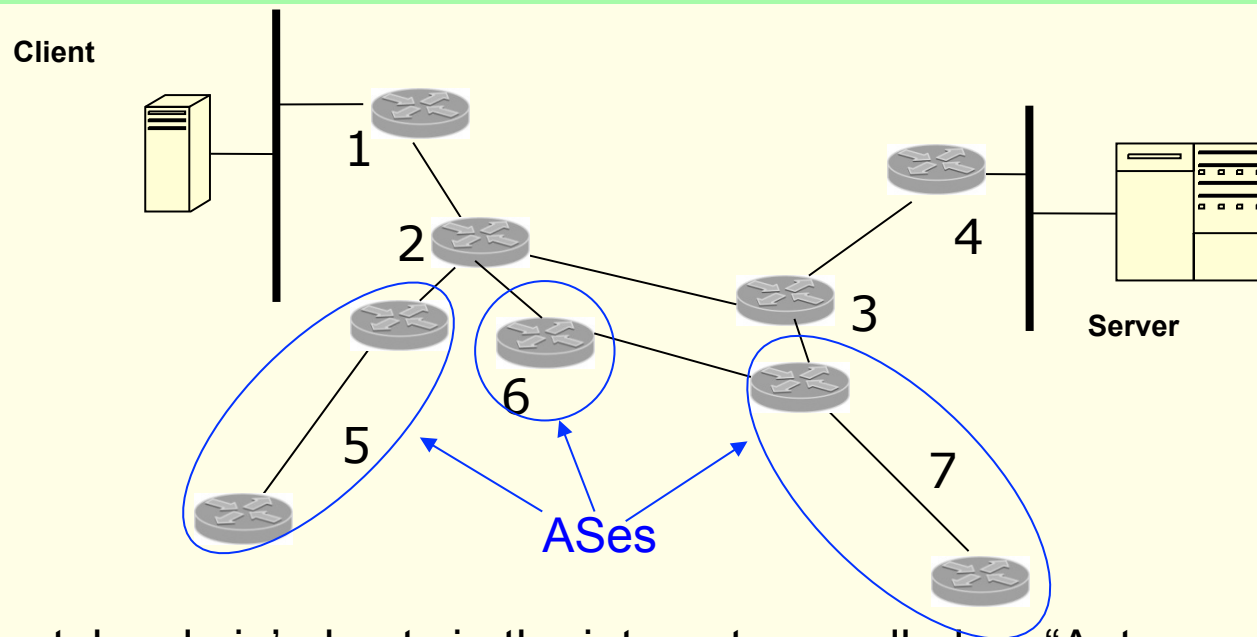
What Is Internet Routing?



The Internet is a long series of connected cables, with the following properties:

- “Routers” are responsible for sending information from one site to another
- There may be more than one way to send information
- Routers are responsible for picking the “best path”
- Routers advertise to their neighbors the routes they know and use
- Routers accept advertisements from neighbors and learn new paths

What Are Autonomous Systems?



Separately admin'ed nets in the internet are called an "Autonomous System" (AS)


- Each AS gets an AS number assigned to it
- Each AS may consist of one or more routers (some have hundreds)
- Each AS remembers how to get from one location to another
- Each router knows the best next AS to send data to heading for a destination
- Each router makes decisions independently from each other

Example: the best path from "Client" to "Server" is AS1 ► AS2 ► AS3 ► AS4

BGP Ties Everything Together

- BGP Route Advertisements:
 - Initiate from an “Origin” AS
 - Travel along each valid path
 - As a route advertisement passes through an AS:
 - Each adds its own AS number to the list
 - Each announcement is recorded as a path to the Origin
- But without added security, every AS believes every other AS
 - Including lies and configuration mistakes

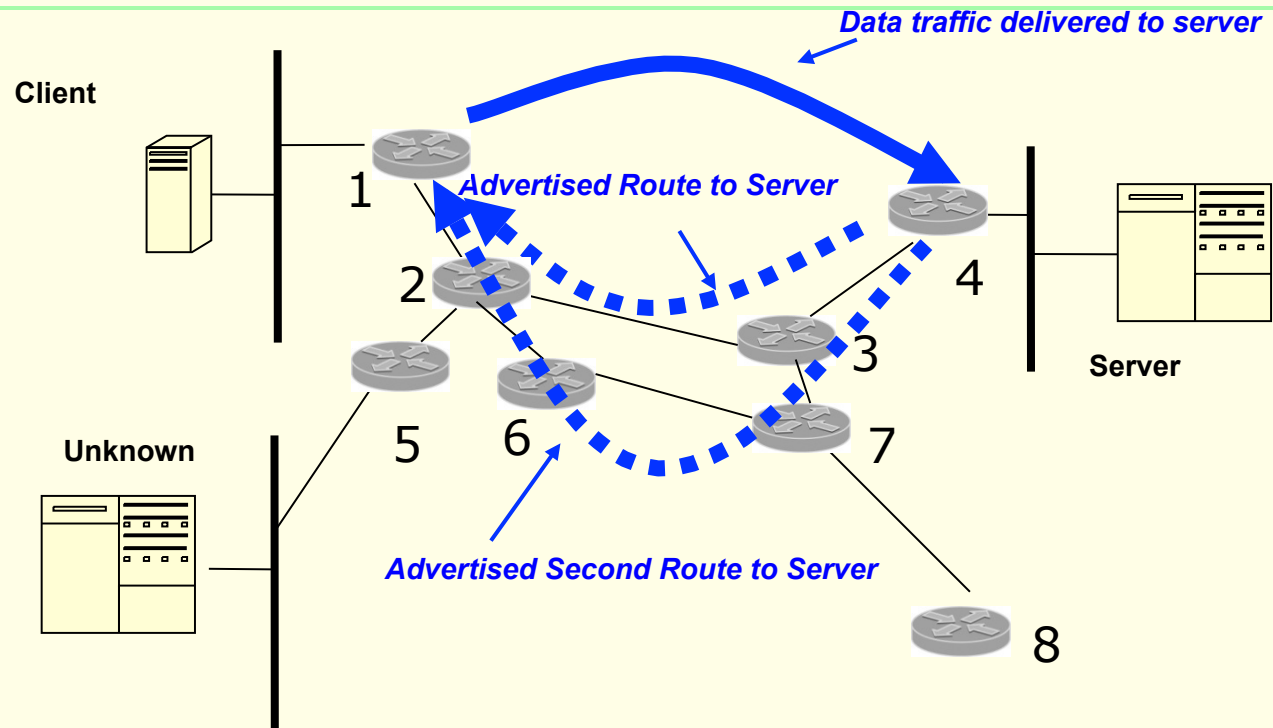
Client Needs A Route To Server



“Client” needs to contact “Server”

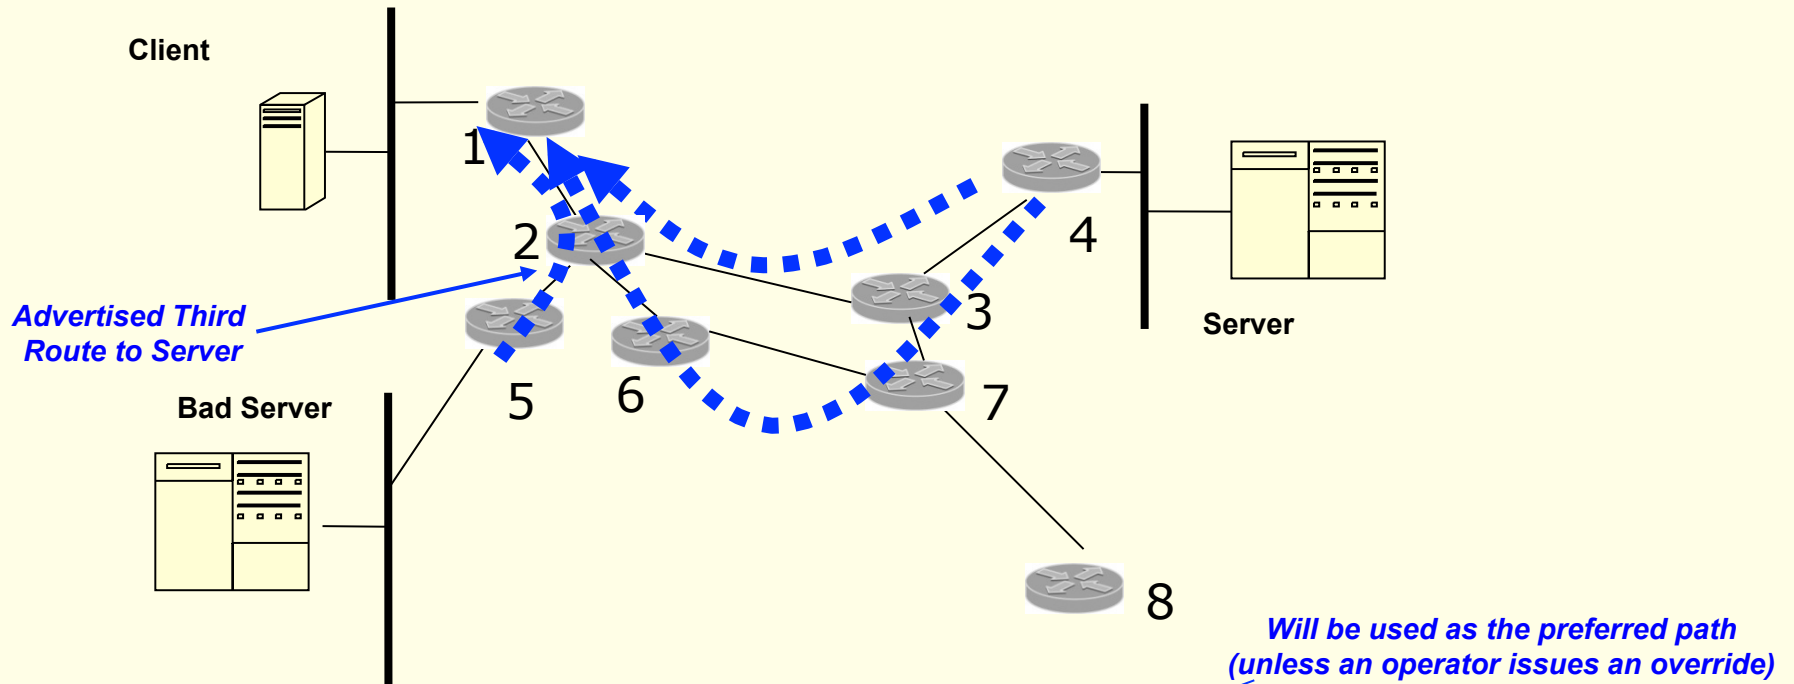
- Does its ISP know where to send traffic to?
- BGP is responsible for providing path information to every network component
- Each network component can then send packets to the next router in the path

Client Needs A Route To Server



- Client's ISP (AS1) believes there are multiple paths to Server:
 - Shortest Route: AS1 ▶ AS2 ▶ AS3 ▶ AS4
 - Longer (backup) Route: AS1 ▶ AS2 ▶ AS6 ▶ AS7 ▶ AS3 ▶ AS4
- AS1 will always deliver packets to AS2
- AS2 can use either route (via AS3 or AS6), but usually uses the shortest

What If AS2 Receives A 3rd Route?

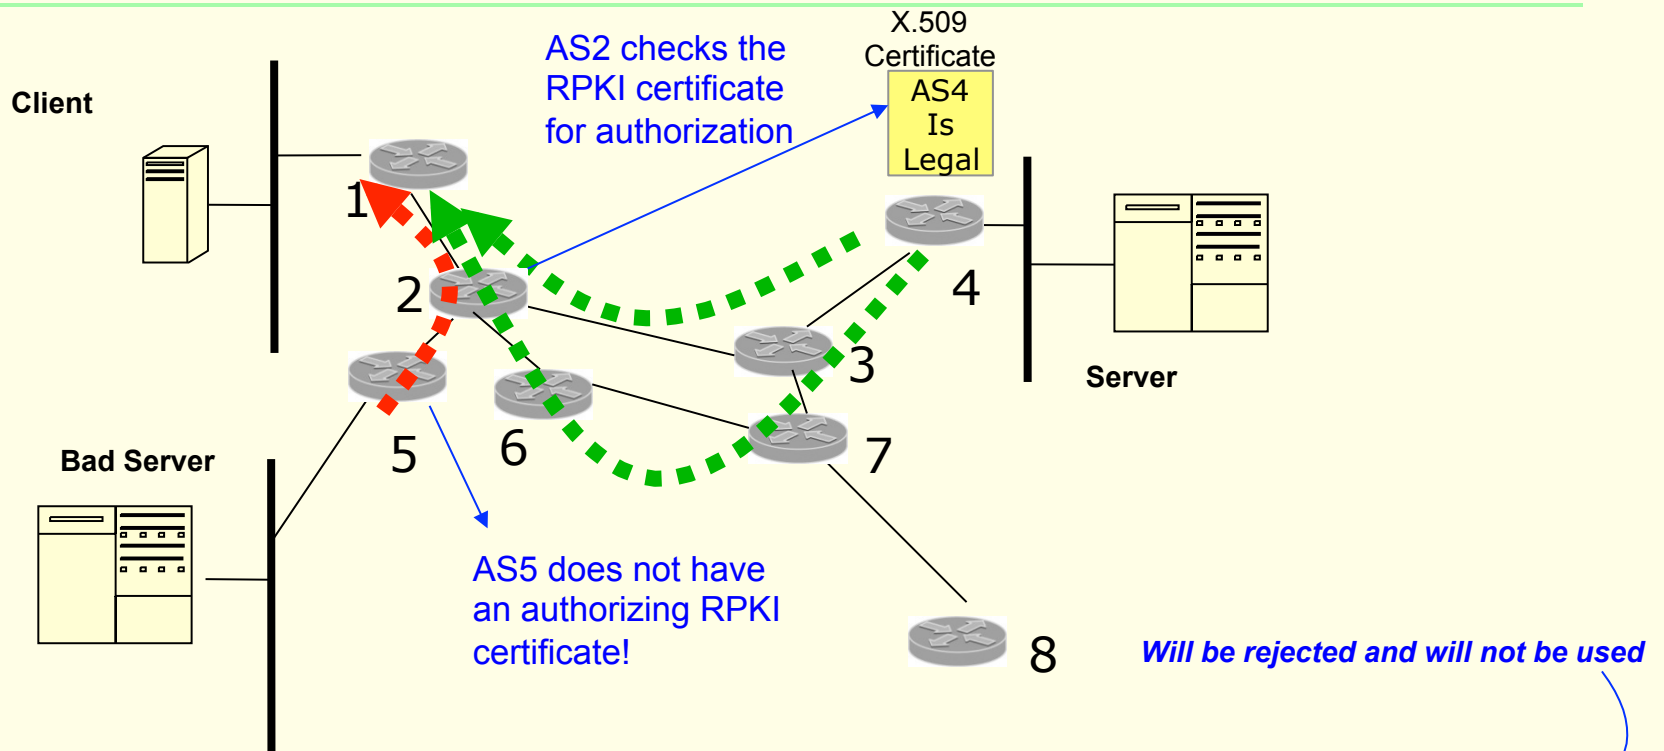


- AS2 now believes it has three paths to Server:
 - Shortest Route: AS2 ► AS5
 - Another Short Route: AS2 ► AS3 ► AS4
 - Longest Route: AS2 ► AS6 ► AS7 ► AS3 ► AS4
- AS2 has **NO** ability to realize that the shortest path is incorrect.
- AS2 will send traffic to AS5, even though it should pick AS3 or AS6

Where Does the RPKI 'Fit'?

- RPKI = Resource Public Key Infrastructure
- RPKI provides 2 principle functions:
 - Establishes resource holders right to use the address;
 - Provides the technical means for other operators (ISPs & others) to validate route origin
- The RPKI functions are similar to DNSSEC signing and DNSSEC validation
 - Many implementation details differ

RPKI Origin Authentication



RPKI Provides Origin Authentication:

- A cryptographically signed certificate authorizes to advertise Routes to Server
- **INVALID** (Doesn't Go To AS4): AS1 ► AS2 ► AS5
- **VALID** (Origin is AS4): AS1 ► AS2 ► AS3 ► AS4
- **VALID** (Origin is AS4): AS1 ► AS2 ► AS6 ► AS7 ► AS3 ► AS4

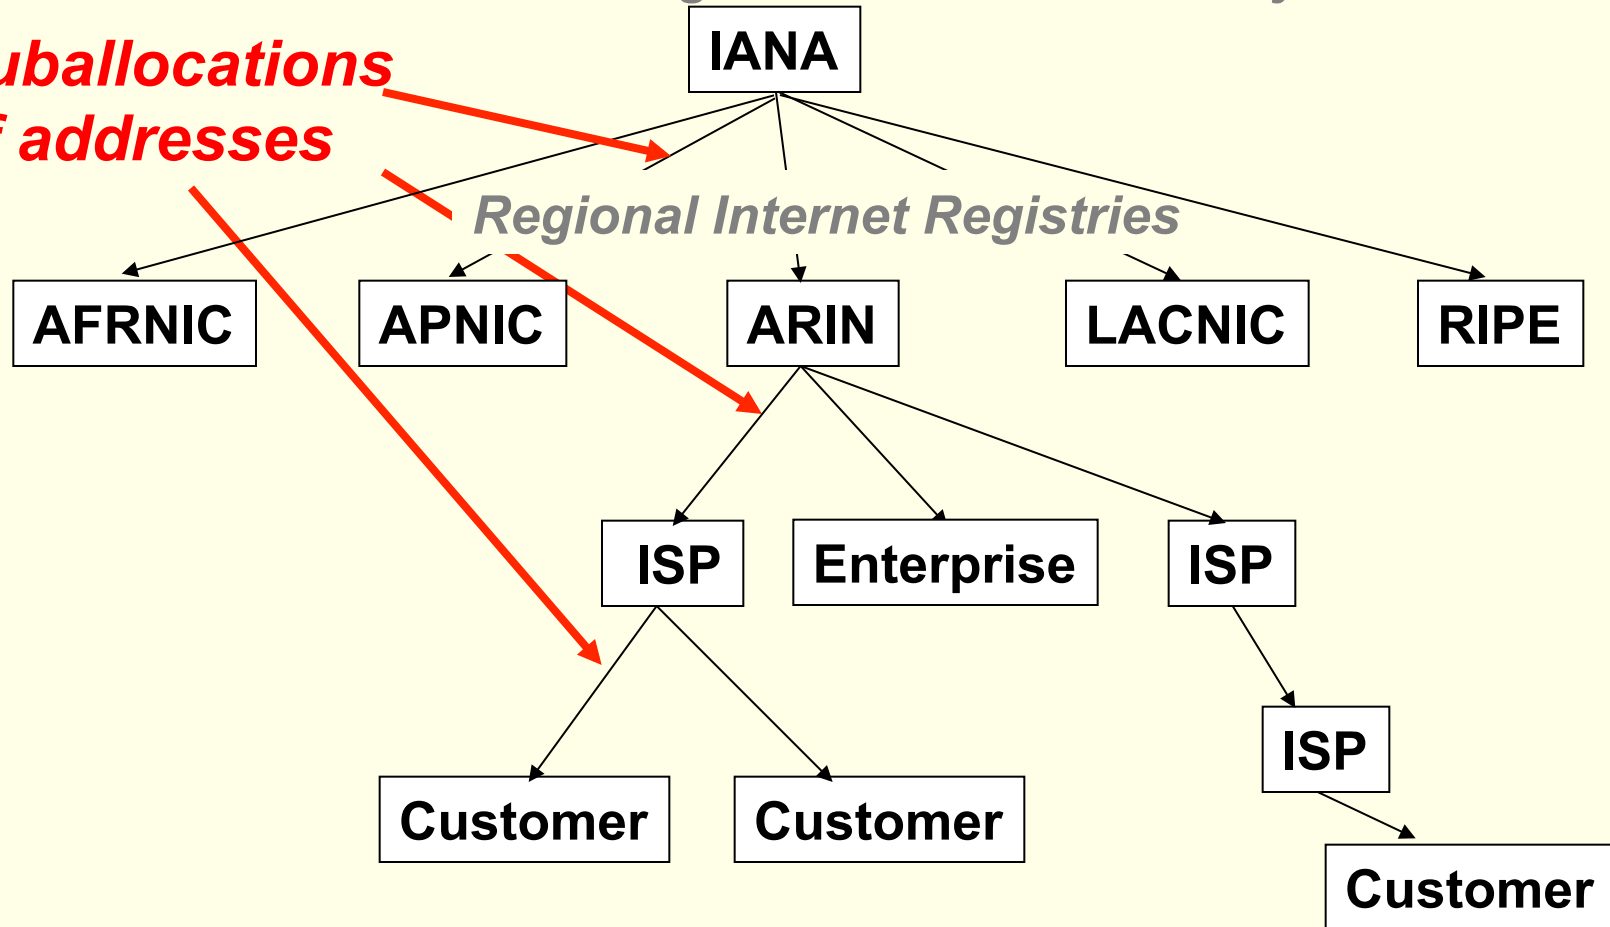
Steps To Improve Routing Security

- Step 1: Resource certification
 - Assurance of who holds rights to address space
- Step 2: Origin validation
 - Address space holder determines who is authorized to originate routes
 - ISPs/routers validate received routes
 - From list of authorized originating ASs
- Step 3: Path validation
 - Prevent valid origin from being attached to bogus path

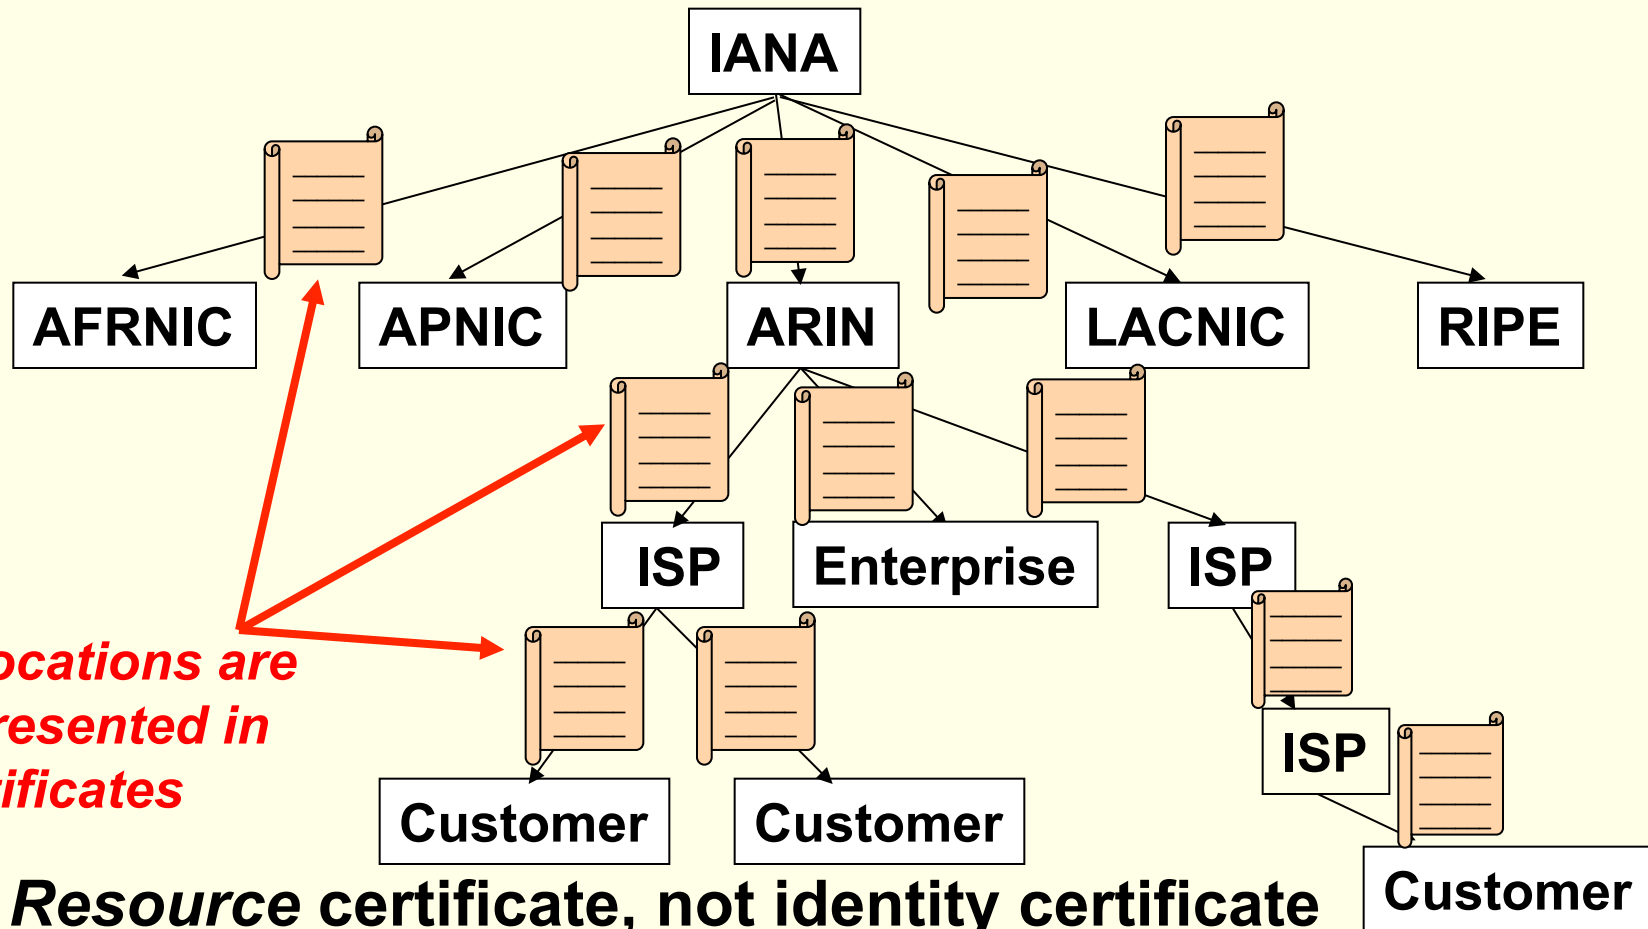
Just Who Does Hold an Address?

Internet Assigned Numbers Authority


**Suballocations
of addresses**



RPKI - Resource Certificates

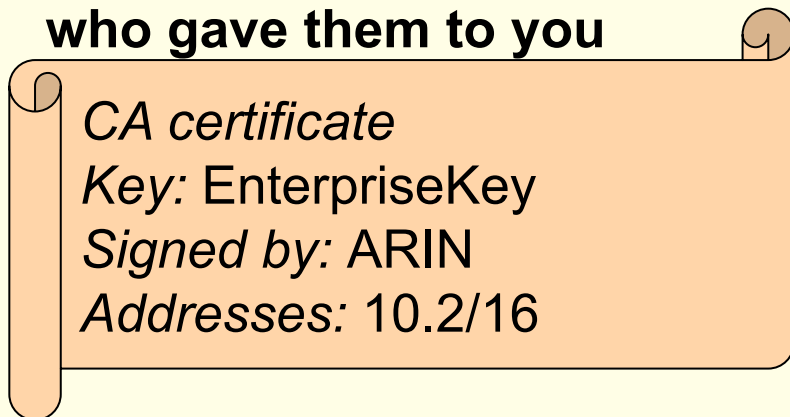


Certs (Step 1) & Origin Authorization (Step 2)



***Sign a Route Origin Authorization (ROA) for your address space
Your certificate validates the signature***

Certificate lists the addresses you hold and who gave them to you



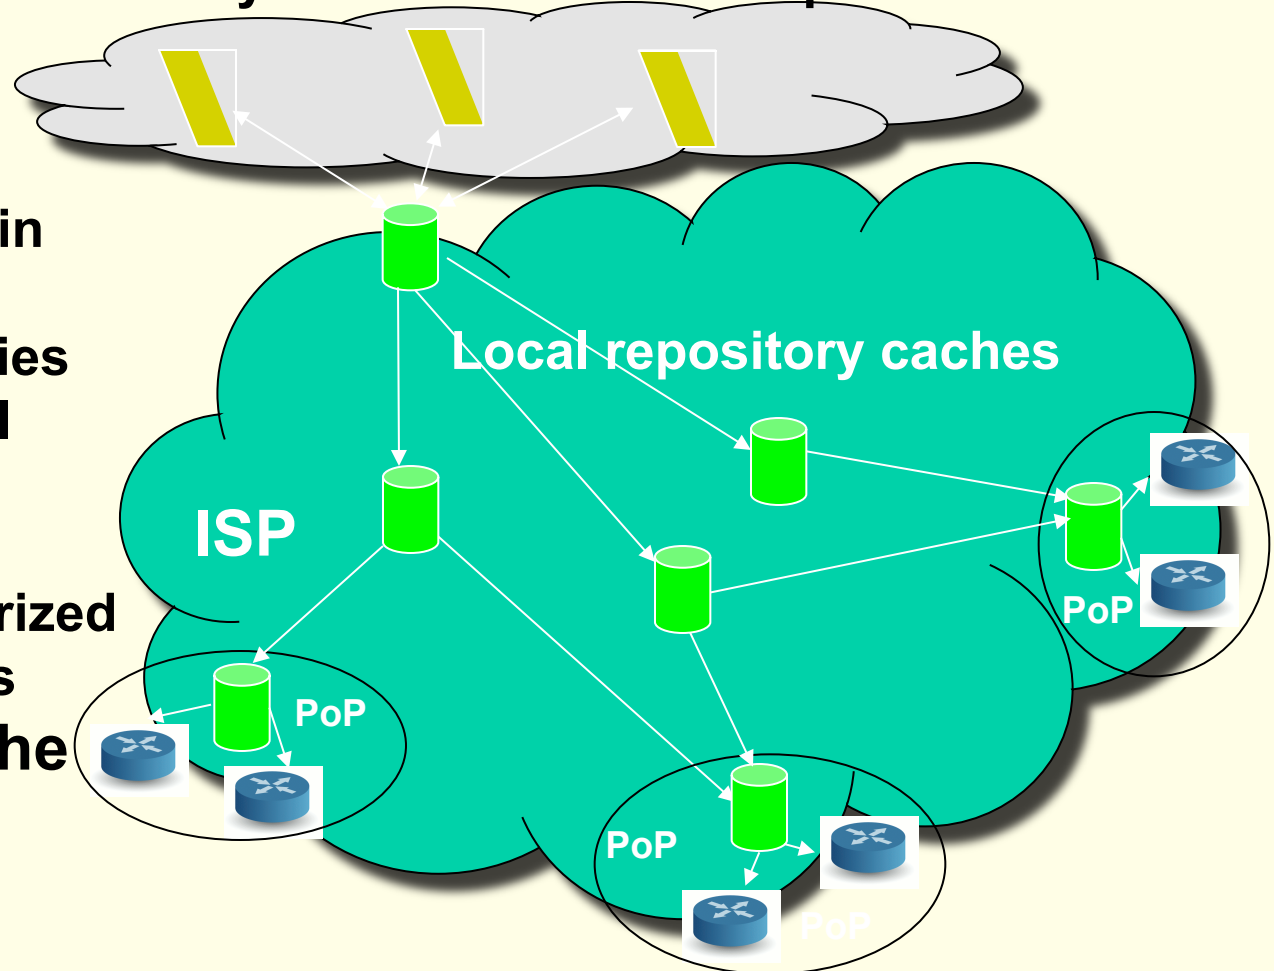
ROA Signed Object
Signed by: EnterpriseKey
Addresses: some part of 10.2/16
Valid Origin: some ASN

Each ROA lists the valid origins for those addresses

Origin Authorization (Step 2)

Global System of Individual Repositories

- Local cache is kept in sync with global distributed repositories
- Local cache does all needed crypto
- Routers need only receive list of (authorized origin, address) pairs
- *N*O* crypto in the routers



Regional Internet Registries (RIRs)

- Each RIR provides RPKI capabilities, status & reports
 - AFRINIC <https://afrinic.net>
 - APNIC <https://apnic.net>
 - ARIN <https://arin.net>
 - LACNIC <https://lacnic.net>
 - RIPE NCC <https://ripe.net>
- The Number Resource Organization (NRO) provides consolidated statistics
 - <https://www.nro.net/rir-statistics-rpki-adoption-reports-now-available/>