RPKI because of a tweet

Tomas Lynch
Senior Network Architect

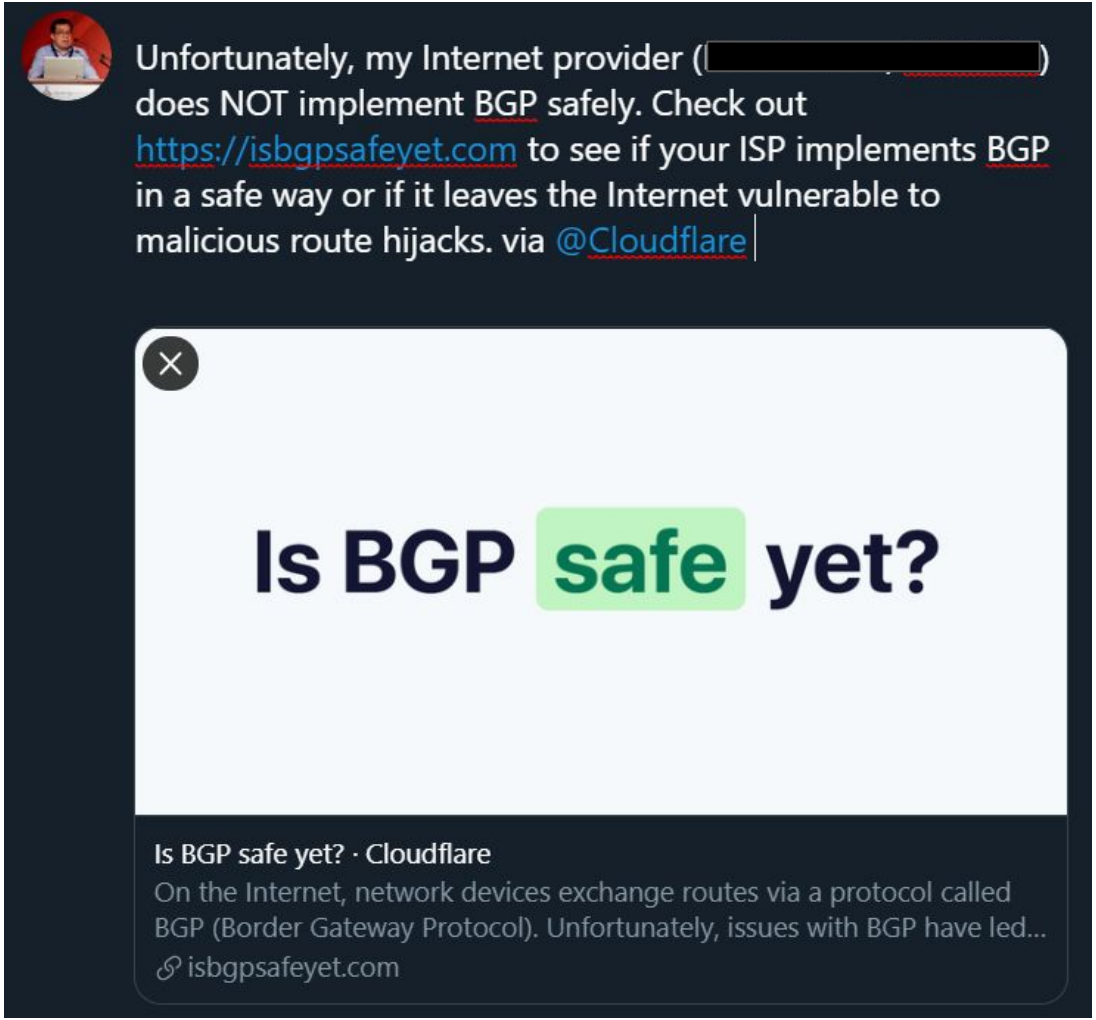
Unfortunately, my Internet provider ([redacted]) does NOT implement BGP safely. Check out https://isbgpsafeyet.com to see if your ISP implements BGP in a safe way or if it leaves the Internet vulnerable to malicious route hijacks. via @Cloudflare
Is BGP safe yet?
Is BGP safe yet? · Cloudflare
On the Internet, network devices exchange routes via a protocol called BGP (Border Gateway Protocol). Unfortunately, issues with BGP have led...
isbgpsafeyet.com

Is BGP safe yet? *No.*

Border Gateway Protocol (BGP) is the postal service of the Internet. It's responsible for looking at all of the available paths that data could travel and picking the best route.

Unfortunately, it isn't secure, and there have been some major Internet disruptions as a result. But fortunately there is a way to make it secure.

ISPs and other major Internet players (Comcast, Sprint, Verizon, and others) would need to implement a certification system, called RPKI.

Test your ISP | Read FAQ

FAILURE

Your ISP ([redacted]) does not implement BGP safely. It should be using RPKI to protect the Internet from BGP hijacks. Tweet this →

▶ Details

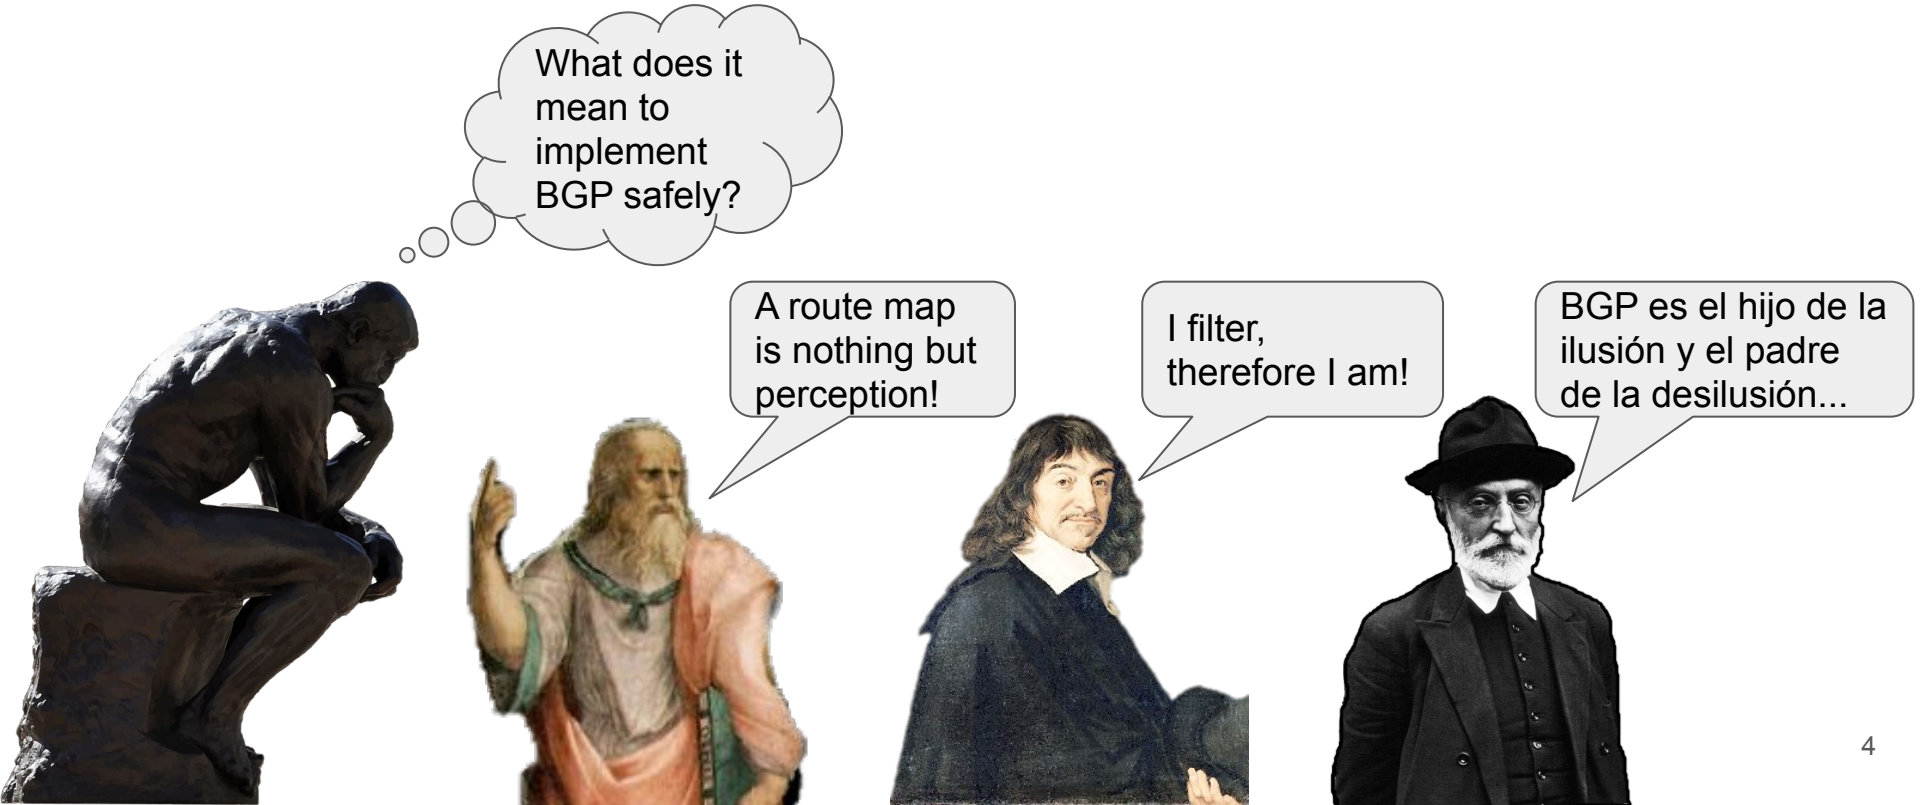
What does it mean to implement BGP safely?
A route map is nothing but perception!
I filter, therefore I am!
BGP es el hijo de la ilusión y el padre de la desilusión...

Let’s implement RPKI! (and join MANRS!)

Basic Topology and former RPKI validation

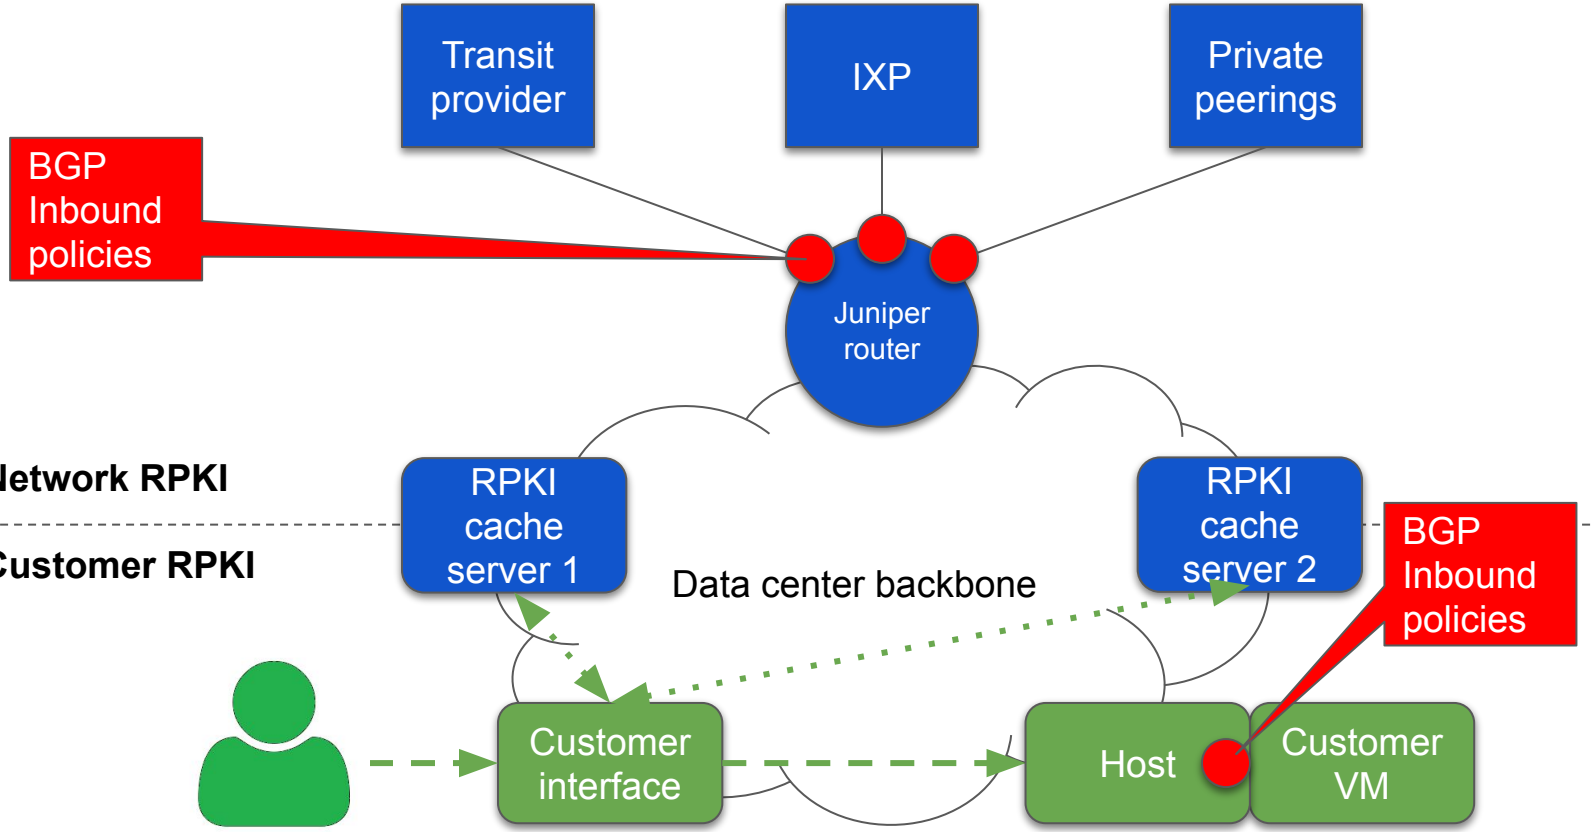

Customer RPKI (simplified)

Already implemented since day 0 using routinator

- Customer adds a prefix to their database
- If valid RPKI then customer prefix list at the host is updated
- If invalid RPKI then the prefix is rejected
- If unknown RPKI, confirmation email is sent to RIR POC
  - If reply is positive then customer prefix list at the host is updated
  - If reply is negative the prefix is rejected

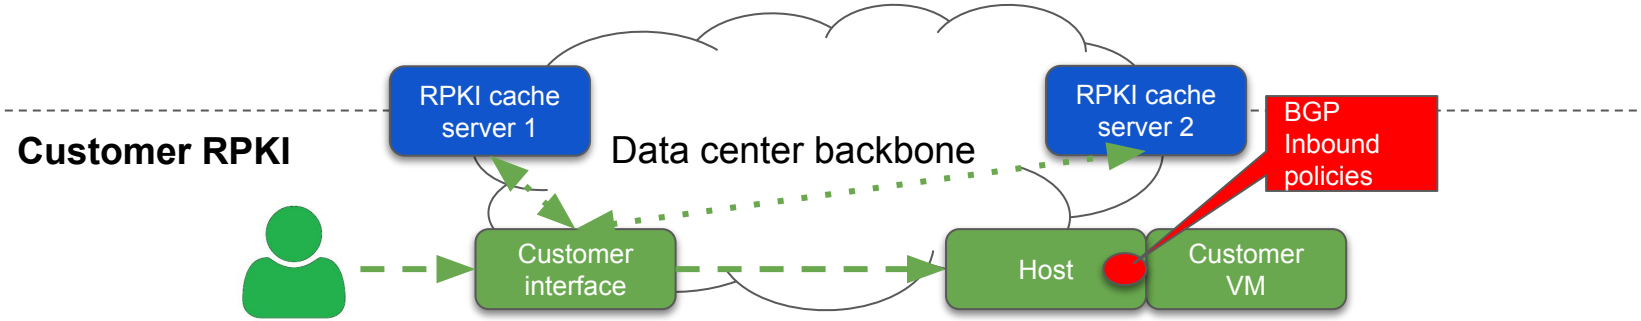

Basic Topology and current RPKI validation

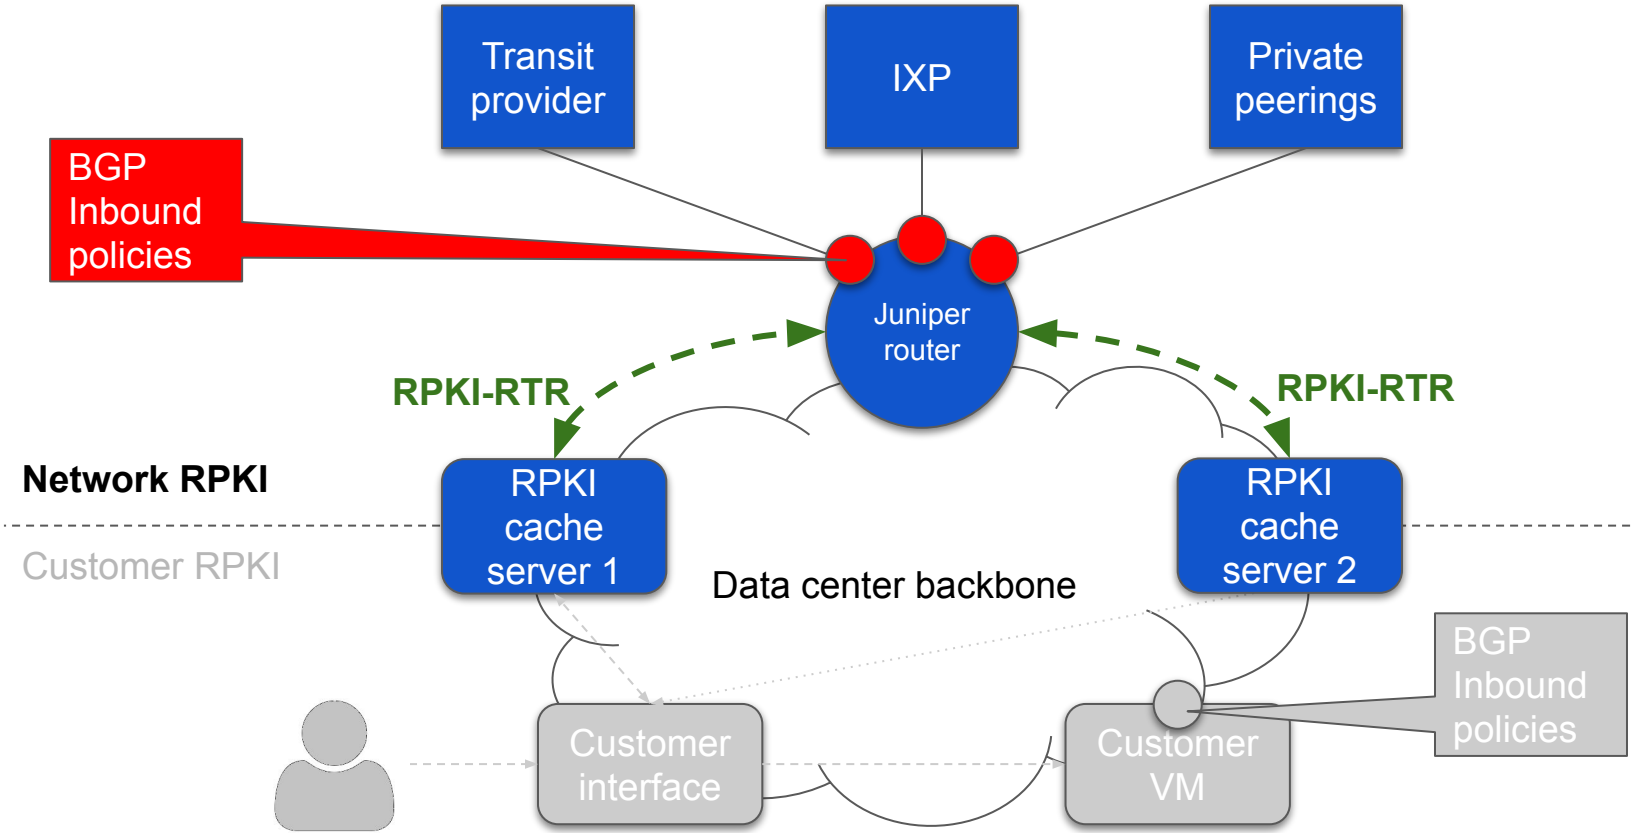

Configuration

Two main configurations

1) Sessions against RPKI validators

2) Policies against peer BGP sessions

Origin Validation for BGP

```
routing-options {
    validation {
        group VALIDATOR {
            session 2001:db8::1;
            session 2001:db8::2;
        }
    }
}
```

- Default port 2222
- Sessions are “dual stack”: information for v4 and v6 is received independent of the server IP address version

Route validation records

On the router, the database entries are formatted as route validation (RV) records

An RV record is a (prefix, maximum length, origin AS) triple. One per validator.

```
router> show validation database
RV database for instance master
Prefix                  Origin-AS Session       State   Mismatch
...
2001:db8::/32-48            64496 2001:db8::1   valid
2001:db8::/32-48            64496 2001:db8::2   valid
...
```

Validating routes

Validation against peers is done using BGP policies

Validation has three states:

> **Valid**: the prefix has an RV and matches the autonomous system
>
> **Invalid**: the prefix has an RV but does not match the autonomous system or the prefix length is longer than the maximum accepted
>
> **Unknown**: the prefix does not have an RV entry

A fourth state called **Unverified** is used when the prefix didn’t run on a validation policy

Basic policy configuration [1/3]

```
[edit policy-options policy-statement PEER_INBOUND_POLICY]
term RPKI_INVALID {
    from {
        protocol bgp;
        validation-database invalid;
    }
    then {
        validation-state invalid;
        community set origin-validation-state-invalid;
        reject;
    }
}
...
```

Basic policy configuration [2/3]

```
...
term RPKI_UNKNOWN {
    from {
        protocol bgp;
        validation-database unknown;
    }
    then {
        validation-state unknown;
        community set origin-validation-state-unknown;
        then accept;
    }
}
...
```

Basic policy configuration [3/3]

```
...
term RPKI_VALID {
    from {
        protocol bgp;
        validation-database valid;
    }
    then {
        validation-state valid;
        community set origin-validation-state-valid;
        then accept;
    }
}
```

Route validation state examples

```
router> show route validation-state valid

2001:db8::/32      *[BGP/170] 3d 03:43:17, MED 0, localpref 70
                      AS path: 64496 I, validation-state: valid
                    > to 2001:db8::10 via et-0/0/0

router> show route validation-state invalid hidden

2001:db8::/32       [BGP/170] 3d 03:43:17, MED 0, localpref 70
                      AS path: 64496 I, validation-state: invalid
                      Discard
```

Some data

Route validation per IP version

	IPv4		IPv6	
Invalid	539	0.07%	71	0.08%
Valid	175,891	21.98%	23,516	27.17%
Unknown	623,835	77.95%	62,962	72.75%
Total	800,265	100.00%	86,549	100.00%

Total project implementation time (aprox.)

17 datacenters

3 transit providers, 1 IXP, 3 private peers per datacenter (average)

Testing in one site: 5 days

Two routinators per datacenter: installation 1 day, router configuration 1 day

Rewrite policies 2 days, apply policies 5 days

Total project implementation time: 14 business days

Is BGP safe yet?

Is BGP safe yet? *No.*

Border Gateway Protocol (BGP) is the postal service of the Internet. It's responsible for looking at all of the available paths that data could travel and picking the best route.

Unfortunately, it isn't secure, and there have been some major Internet disruptions as a result. But fortunately there is a way to make it secure.

ISPs and other major Internet players (Comcast, Sprint, Verizon, and others) would need to implement a certification system, called RPKI.

Test your ISP | Read FAQ

SUCCESS

Your ISP ([redacted]) implements BGP safely. It correctly drops invalid prefixes. Tweet this →

▶ Details

Are you happy now?
Not yet, let’s join MANRS

MANRS for CDN and Cloud Providers

Action 1. Prevent propagation of incorrect routing information (mandatory)

Action 2. Prevent traffic with illegitimate source IP addresses (mandatory)

Action 3. Facilitate global operational communication and coordination (mandatory)

Action 4. Facilitate validation of routing information on a global scale (mandatory)

Action 5. Encourage MANRS adoption (mandatory)

Action 6. Provide monitoring and debugging tools to the peering partners (optional).

https://www.manrs.org/cdn-cloud-providers/

Did it help us to join MANRS?

Yes, it was not only a few press releases and tweets

For us, to discover an issue related to our customer prefixes and fix it

For ISOC, to find a small issue with one RIR and fix it

Is that enough?
Is my network secure now?
Can we get rid of those prefix lists?

No, of course not
but you are getting closer

Thank you!

Tomás Lynch
Senior network architect
Vultr, LLC.