Route Leak Detection with Community

Alexander Azimov, Yandex
Sriram Kotikalapudi, US NIST
Doug Montgomery, US NIST
Brian Dickson, Independent
Andrei Robachevsky, Internet Society
Eugeniu Bogomazov, Qrator
Randy Bush, Internet Initiative Japan
Keyur Patel, Arrcus

The Signal(s)

Leak prevention:

If a route is received from provider, RS or peer it MUST not be sent to another provider or peer. The signal is set on ingress.

Leak detection:

If a route is sent to customer, peer or RS-client it also MUST follow 'only down' rule. The signal is set on egress.

Attribute vs Community

Attribute:

- More reliable signal;
- Memory efficient;
- Reserved for specific use.

Community:

- Easy to implement!

Choose between? No, we'll take both!

DO Community

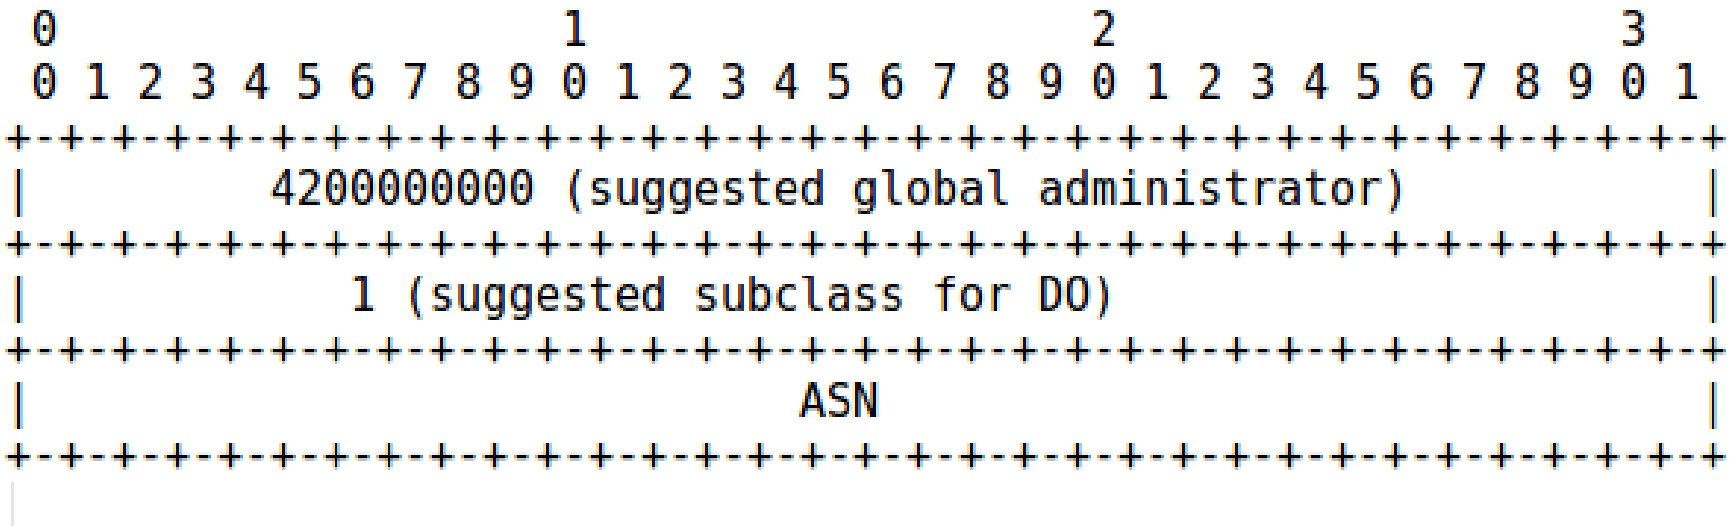

4200000000:* – a suggested reserved class for well-know transit communities.

DO: Setting

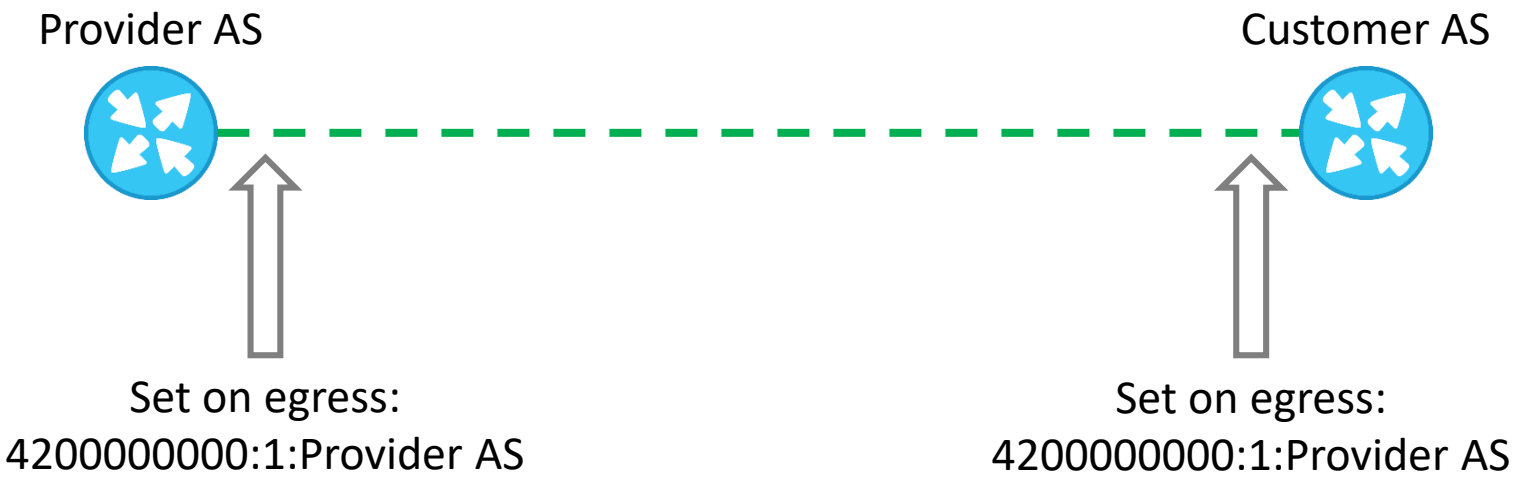

No matter who sets the signal – the value is the same.
The signal indicates that route can be sent only to customers!

Community: Prevention & Detection

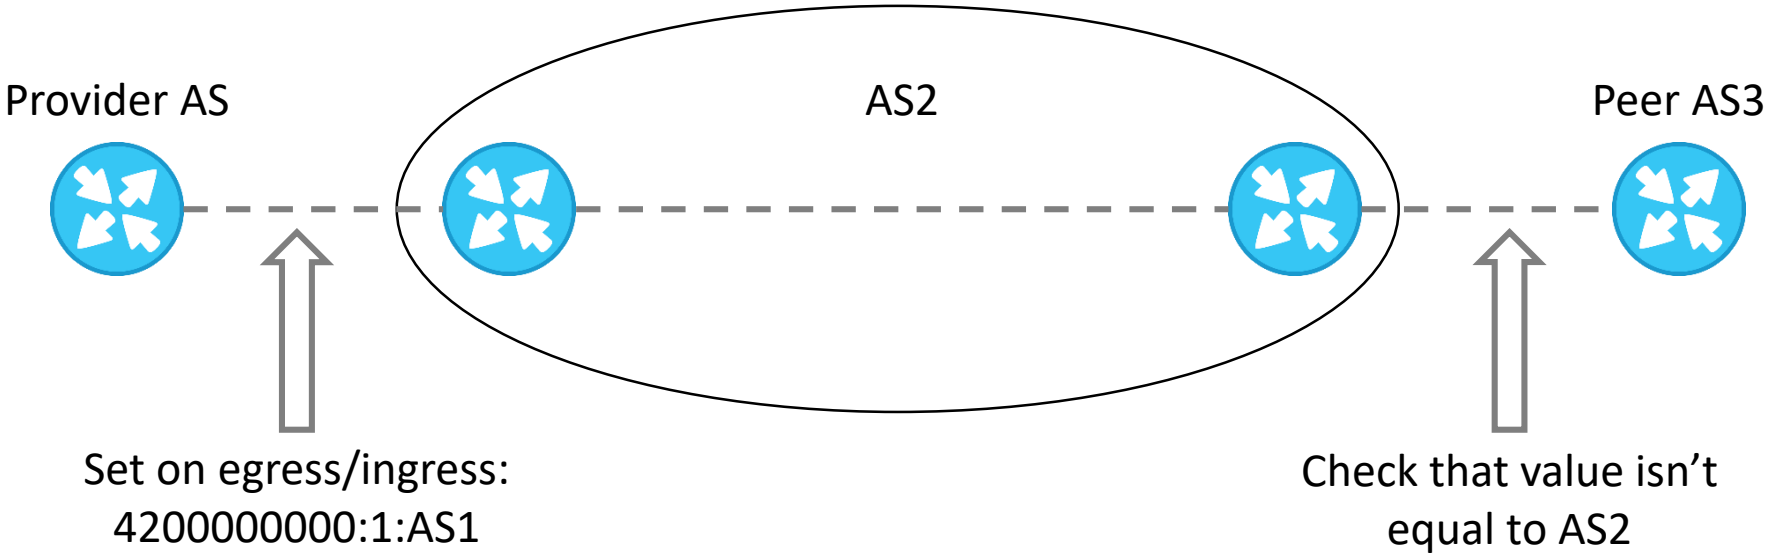

Ingress rule for peers: value != neighbor_as

Match 4200000000:1:neighbor_as – accept
Match 4200000000:1:* – route leak

What Should We Do with Route Leaks?

- Drop route leaks;
- Set LOCAL_PREF = 0;
- Pass it on.

LOCAL_PREF = 0: Original Consideration

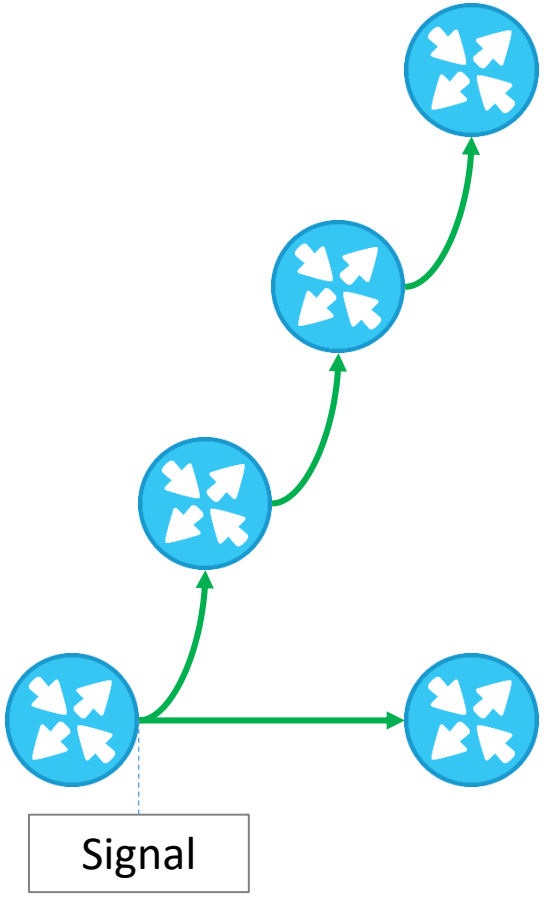

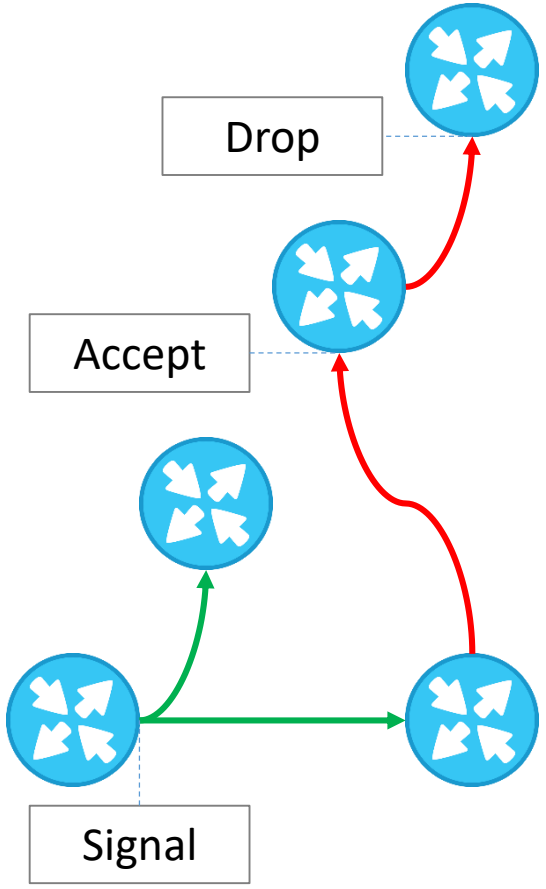

LOCAL_PREF = 0: Original Consideration

- The prefix must have a single path to Tier1;
- The prefix must not have a less specific prefix with multiple paths to Tier1;
- The leak must affect this single path, so it must be inside single path to Tier1;
- The prefix must be marked with leak detection signal;
- Immediate ISP that accepts leak must learn it from its customer;
- Immediate ISP that accepts leak must prefer leaked prefix;
- Immediate ISP that accepts leak should not use leak detection;
- Immediate ISP must keep detection signal;
- The upper provider is dropping leaked prefixes.

LOCAL_PREF=0: Problem with BGP Wedgies

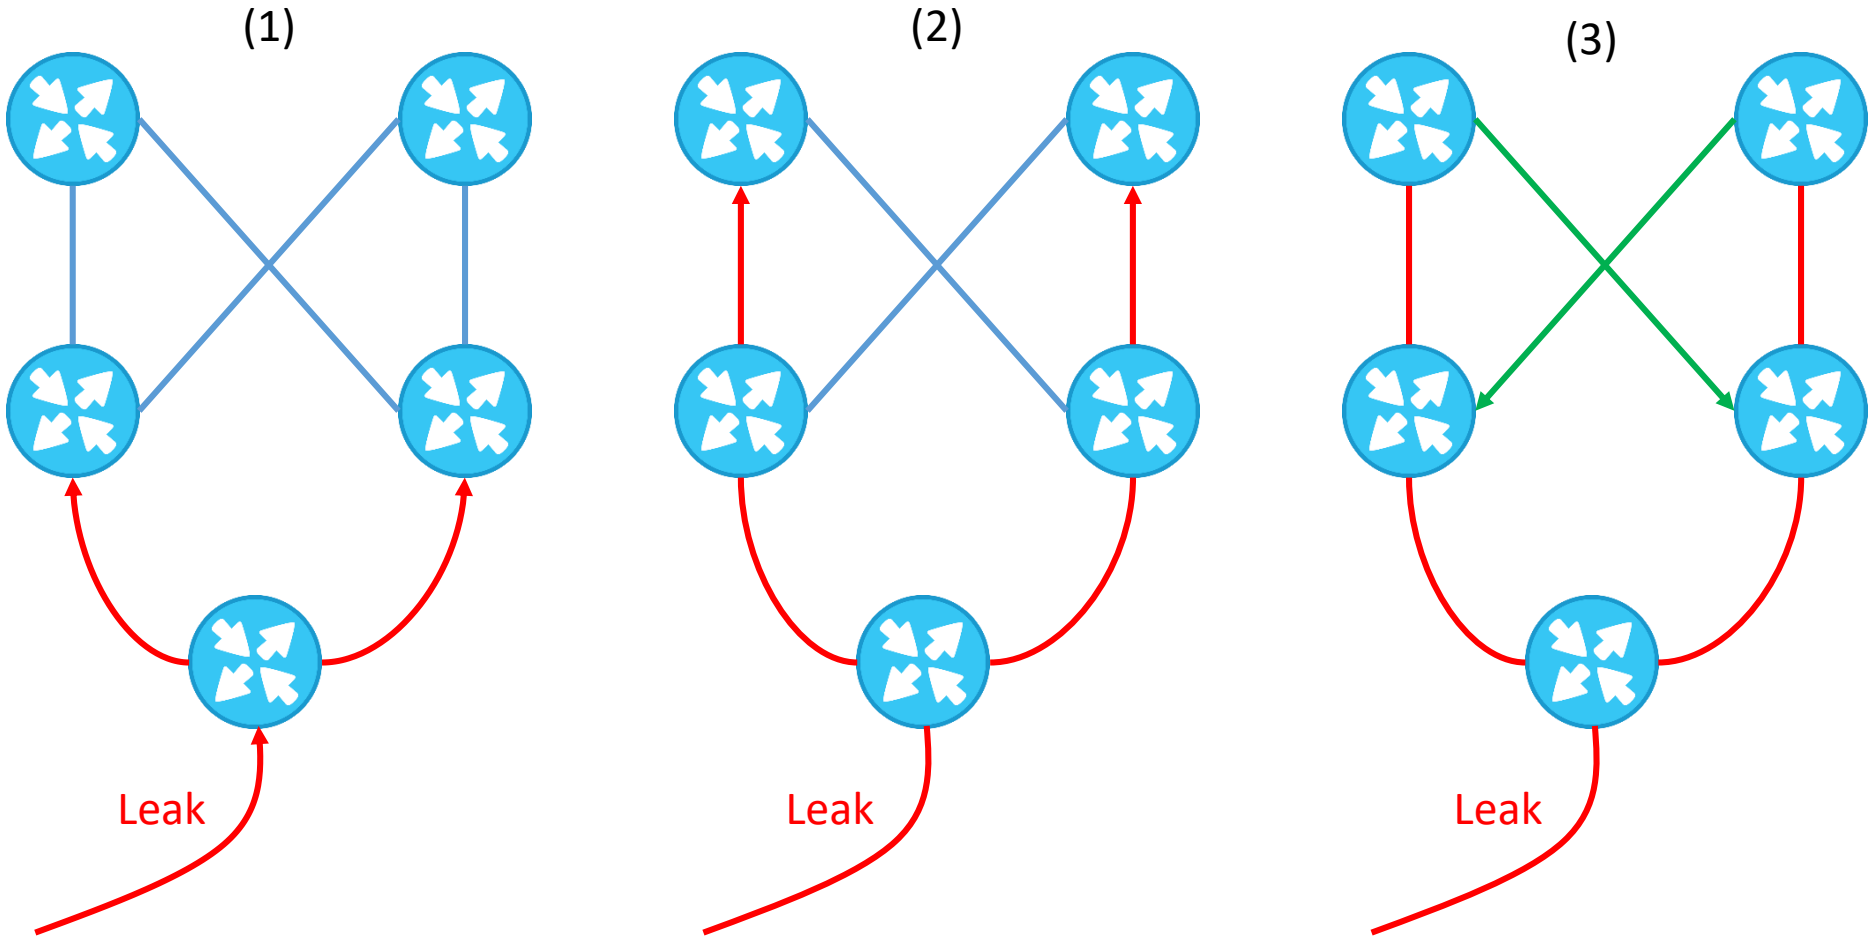

LOCAL_PREF=0: Problem with CDNs

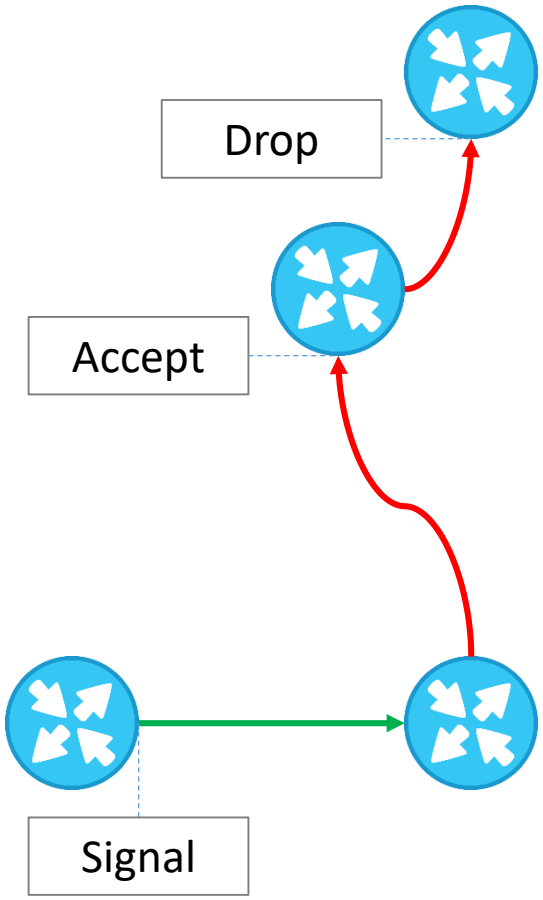

The leaked prefix might be:

- Poor customer with unreliable topology;
- CDN that send more specific to or peer or through IX.

Prefix Visibility Distribution

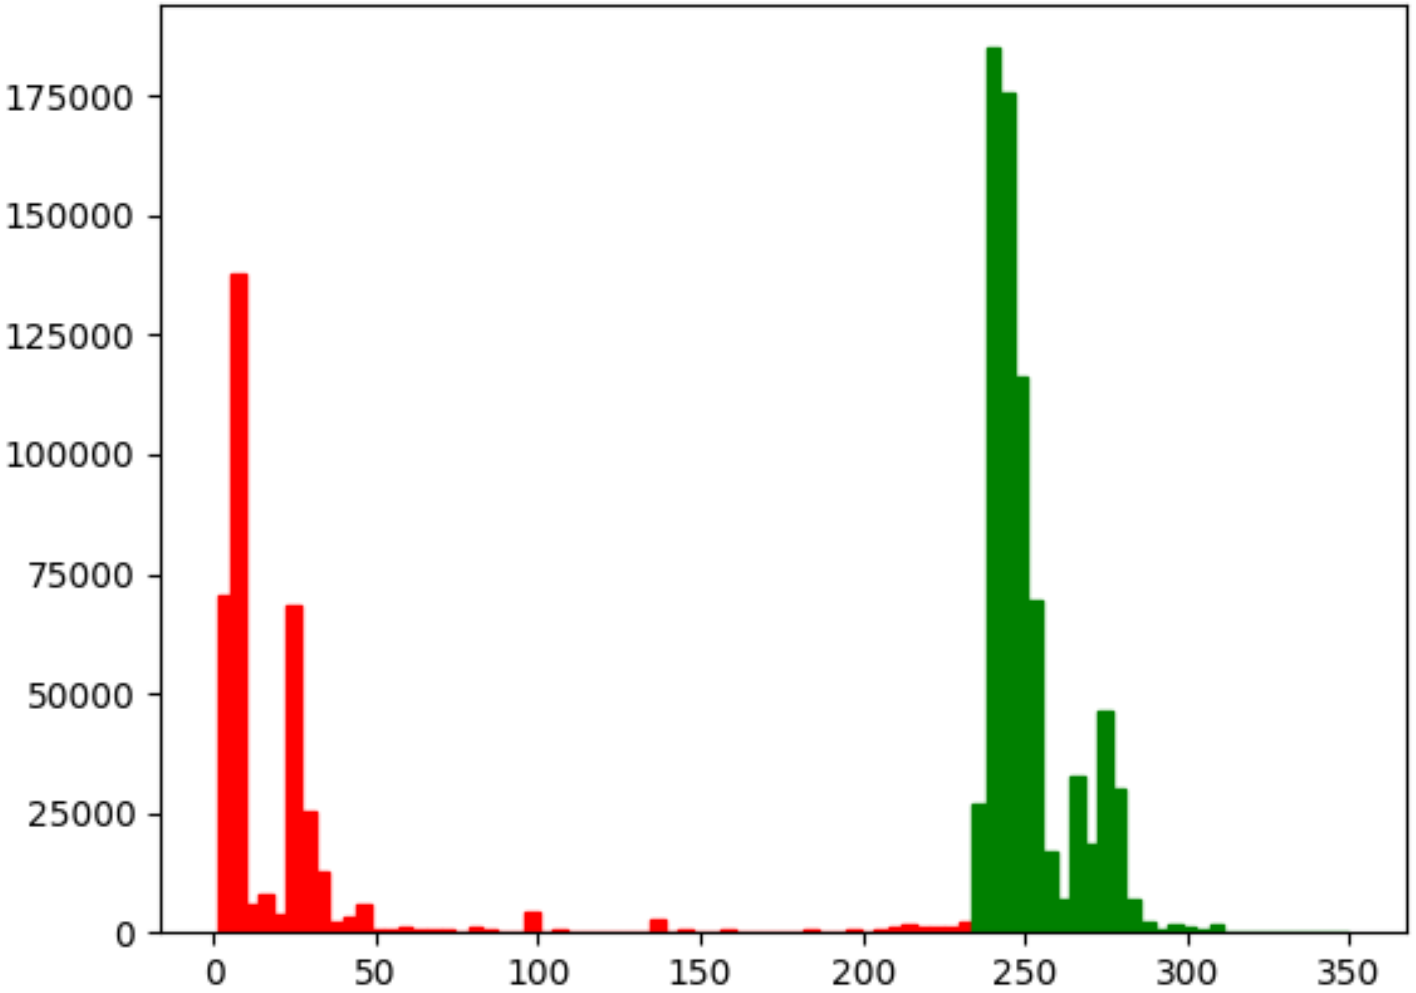

RIPE RIS data source.

What Should We Do with Route Leaks?

The only acceptable mitigation policy – route leaks MUST be rejected.

Pass On: With Caution

Ingress rule for peers: value != neighbor_as

Match 4200000000:1:neighbor_as – permit
Match 4200000000:1:* – deny

Early adopters MAY want to collect data before applying drop policy.

Early adopters MAY want to mark before applying drop policy.

MUST not mark on egress interface with peers without dropping.

What Should We Do with Route Leaks?

The only acceptable mitigation policy – route leaks MUST be rejected.

This mitigation policy SHOULD be used.

Next Steps

- Get your feedback!
- IANA allocation 4200000000:* for well-known transit communities;
- IANA allocation 4200000000:1:* for leak prevention/detection;
- WGLC.