

Improving the peering business case with RPKI

Job Snijders

NTT Communications

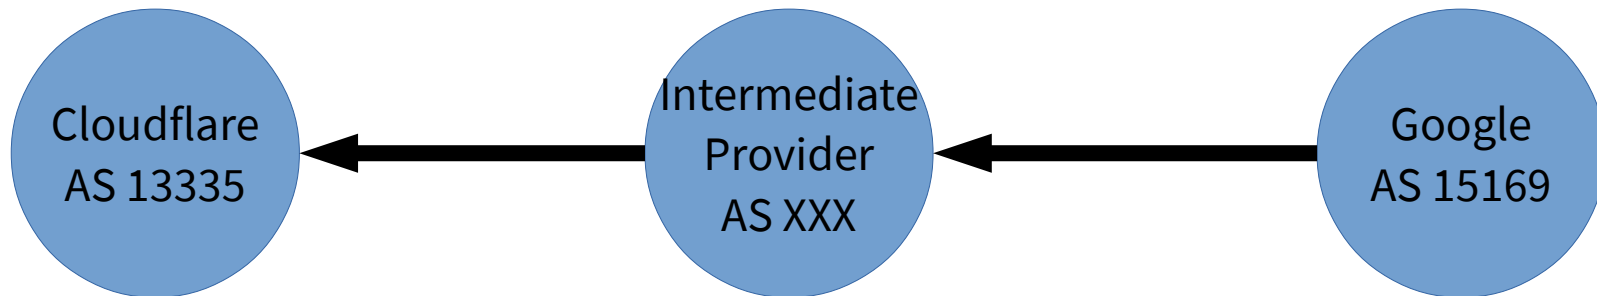
job@ntt.net

Agenda


- Traditional benefits of peering
- RPKI considerations for IP network operators
- RPKI considerations for IXP operators

Traditional benefits of peering

Scenario through transit, AS_PATH is 2 hops: XXX_15169



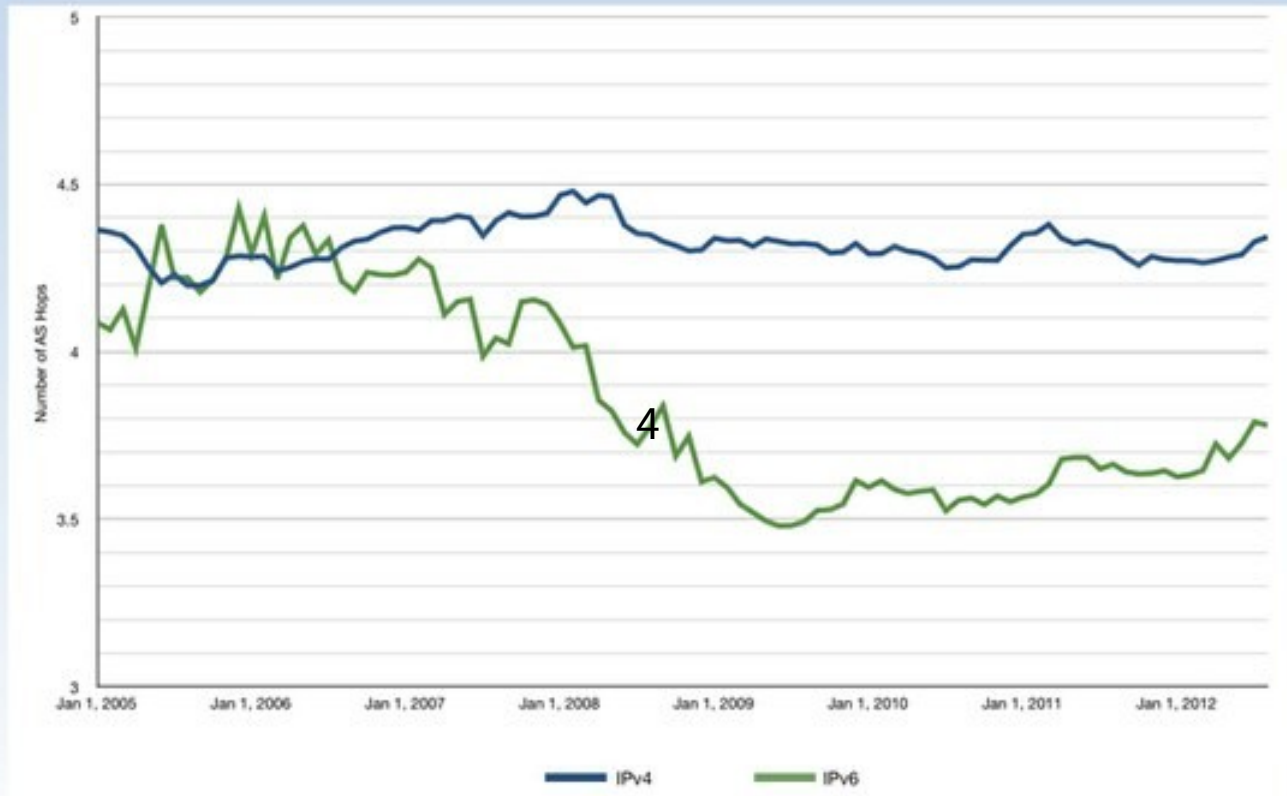
Scenario with direct peering: AS_PATH is 1 hop: _15169\$



- No dependency on the intermediate provider (simpler operations)
- Reduced cost (cross-connect often is cheaper than using an intermediate network if you are in the same building)
- Simplified capacity management
- Etc etc

The internet keeps connecting directly

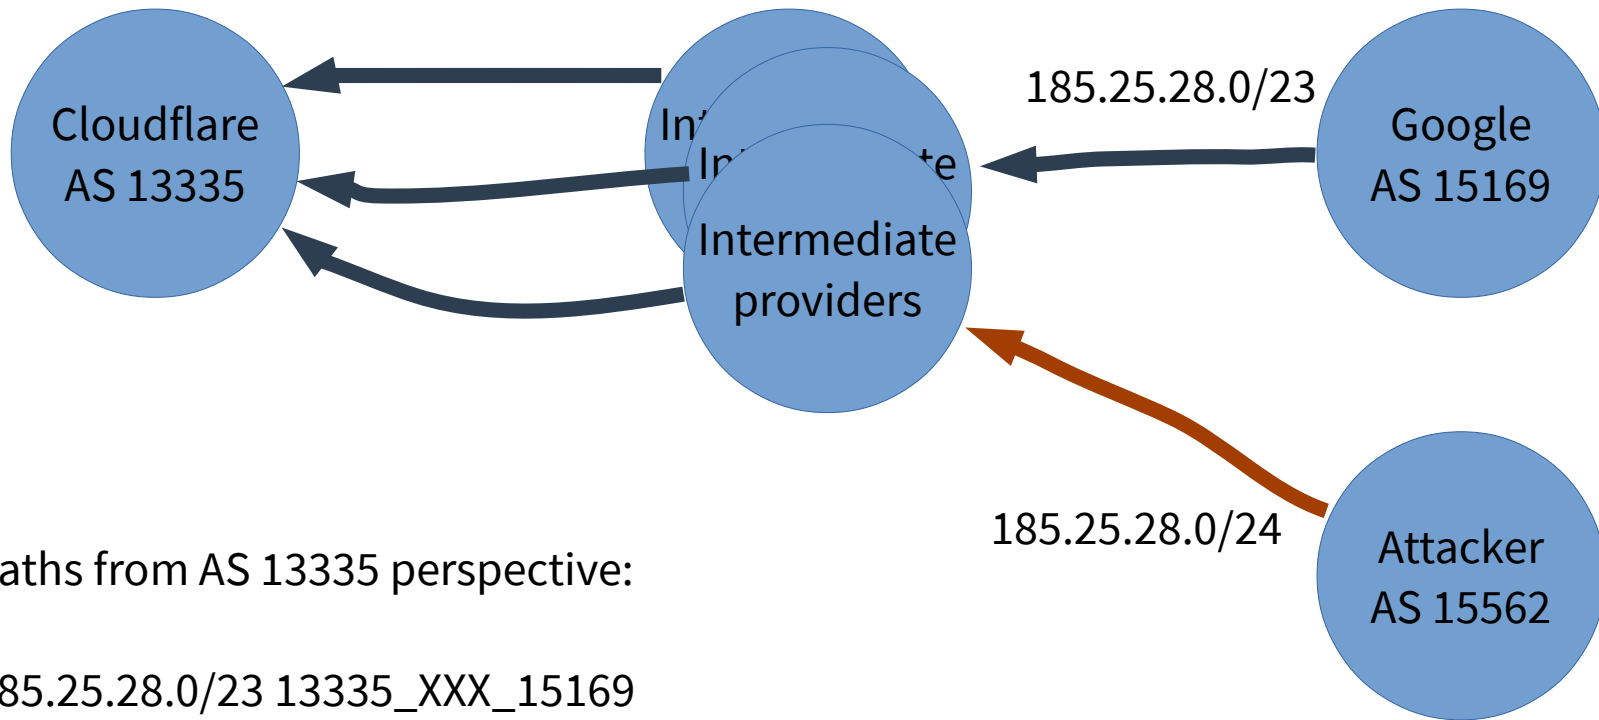
Average AS Path Length



2012 Source:

<https://labs.ripe.net/Members/mirjam/update-on-as-path-lengths-over-time>

Hijack / misconfiguration scenario



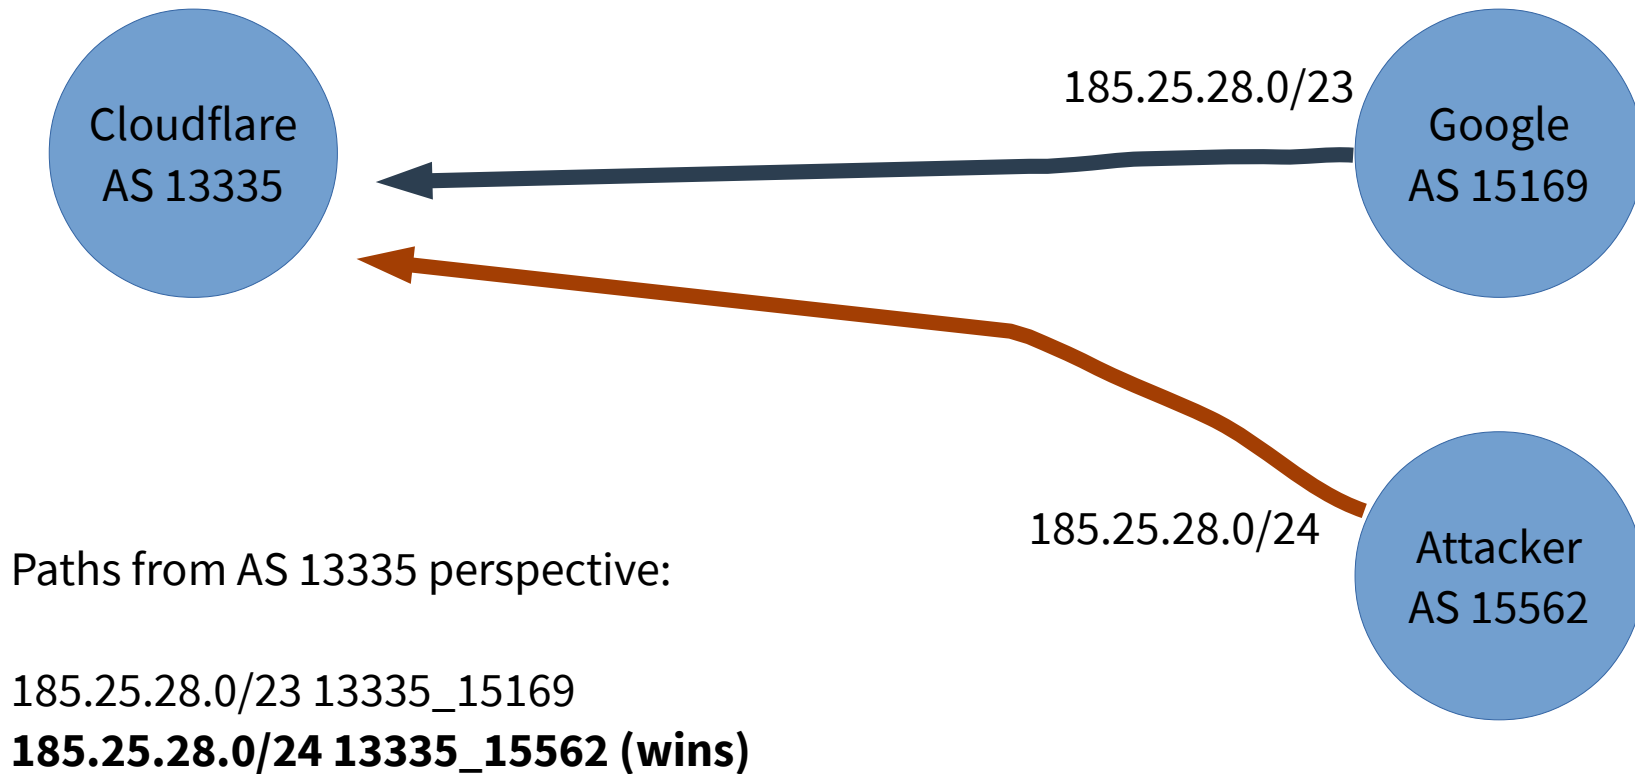
Paths from AS 13335 perspective:

185.25.28.0/23 13335_XXX_15169

185.25.28.0/23 13335_YYY_15169

185.25.28.0/24 13335_ZZZ_15562 (wins)

Hijack / misconfiguration scenario – direct peering




Enter RPKI ROAs

Prefix: 185.25.28.0/23
Prefix description: Google
Country code: CH
Origin AS: 15169
Origin AS Name: GOOGLE - Google LLC, US
RPKI status: ROA validation successful
MaxLength: 23
First seen: 2016-01-08
Last seen: 2019-02-26
Seen by #peers: 40

Hijack / misconfiguration scenario – RPKI ROA

Cloudflare applying “invalid == reject”



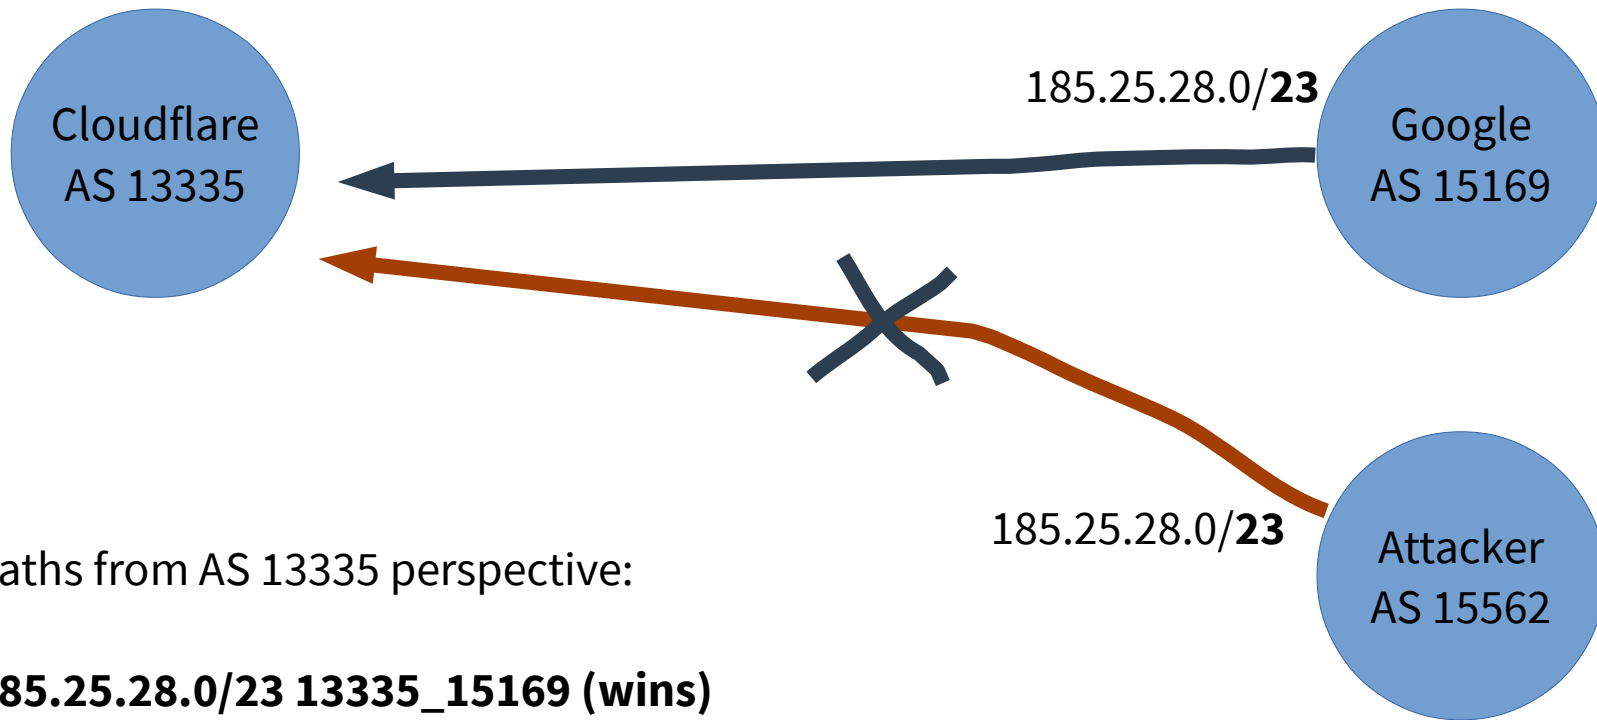
Paths from AS 13335 perspective:

185.25.28.0/23 13335_15169 (wins)

185.25.28.0/24 13335_15562 (rejected, wrong prefix length)

Change of tactics: announce same prefix

Cloudflare applying “invalid == reject”




Paths from AS 13335 perspective:

185.25.28.0/23 13335_15169 (wins)

185.25.28.0/23 13335_15562 (rejected, wrong Origin ASN)

Change of tactics: spoof origin – NOT EFFECTIVE!

Cloudflare applying “invalid == reject”



Paths from AS 13335 perspective:

185.25.28.0/23 13335_15169 (wins)

185.25.28.0/23 13335_15562_15169 (not shortest AS_PATH)

Summary for Network Operators

- RPKI based BGP Origin Validation protects against misconfigurations
- Origin Validation blocks out more-specifics (malicious or not)
- Shortest AS_PATH is now a security feature
- Direct peering, combined with RPKI, is extremely strong!

Considerations for IXP operators

Every IXP has a vulnerable spot: the Peering LAN Prefix

Imagine DE-CIX Frankfurt – 80.81.192.0/21

Imagine I announce 80.81.192.0/24 to the Default-Free Zone

If IX participants accept a more-specific of the Peering LAN Prefix, BGP packets may go the wrong place. Some BGP implementations follow the RIB/FIB for BGP packets

→ result – BGP sessions across IXP go down. BAD

IXP Operator problem space


Asking everyone to filter out the Peering LAN prefix is problematic:

- what if you grow the IXP peering lan? Email everyone?
- what about stale configurations, if people don't update?
- there are 1000+ IXPs world-wide, will all of them email everyone?

IXP Operator recommendation

- **Create RPKI ROAs for your Peering LAN Prefixes!**
- **Set MaxLength to exactly the same as the prefix length**
- **Use Origin “AS 0” or the IXP Services ASN**
- **Don’t announce the Peering LAN Prefix**

IXP Route Server considerations



Source: BKNIX <https://bknix.co.th>

IXP Route Server Considerations

Route Server intention is to mimick direct peering, and provide a convenience service to participants.

But if you have an ASN, you have a responsibility. Applies to IXPs too!

When networks create RPKI ROAs, we expect them to be honored...

Customers expect excellent leadership from IXP Operators!

IXP Route Server Considerations

We discussed the benefits of RPKI and direct peering.

Therefore, IXP Route Server must deploy “RPKI Invalid == Reject” policies!

No reason not to, all modern Route Server software supports RPKI:

- IXP Manager
- Arouteserver
- BIRD
- OpenBGPD

IXPs using RPKI based Origin Validation

- YYCIX
- INEX
- AMS-IX
- FranceIX
- DE-CIX (soon!)
- Netnod (soon!)
- Others?
- you?

More information: <http://peering.exposed/>

Summary for IXPs

IXP Operators must create RPKI ROAs for their Peering LAN Prefix

IXP Operators must enforce “invalid == reject” on Route Servers

Questions?

Ask at the microphone!

Or email me at job@ntt.net