Detecting Inconsistent Advertisements from Neighboring ASes

Nick Feamster, M.I.T.
Z. Morley Mao, University of Michigan
Jennifer Rexford, AT&T Labs -- Research

Settlement-Free Peering

Settlement-Free Interconnection Policy

"3.1 Each Internet Network will announce consistent customer routes at all interconnection points, unless both Internet Networks mutually agree otherwise based on special circumstances."

A peer should:

1. Advertise routes at all peering points.
2. The routes must have equal:
   - AS path length
   - MED values

Advertising Consistent Routes is Important

Enables an AS to perform "hot potato" routing.

"Hot−potato" Routing

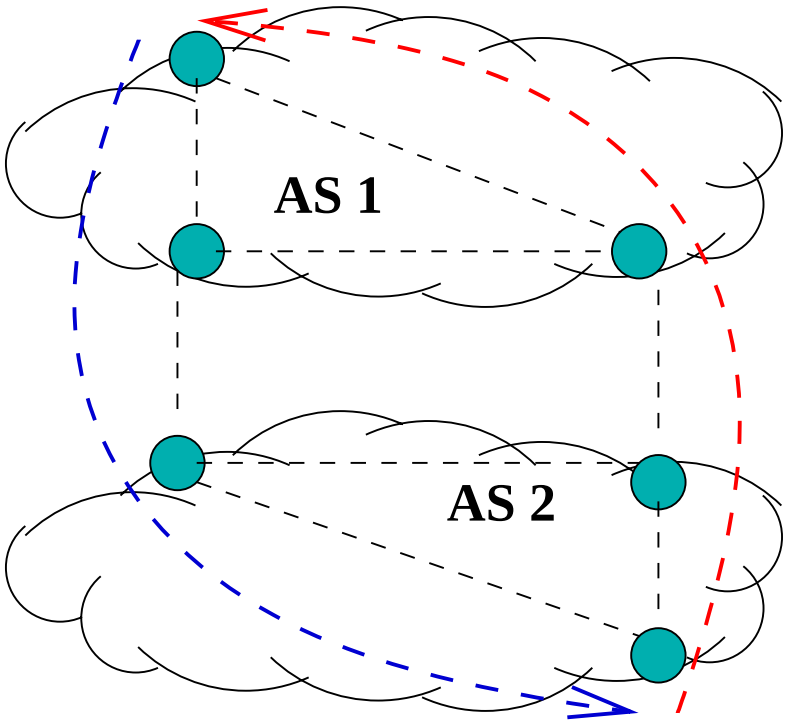

If AS 1 does not advertise a consistent route on this session, AS 2 cannot do "hot potato" routing!

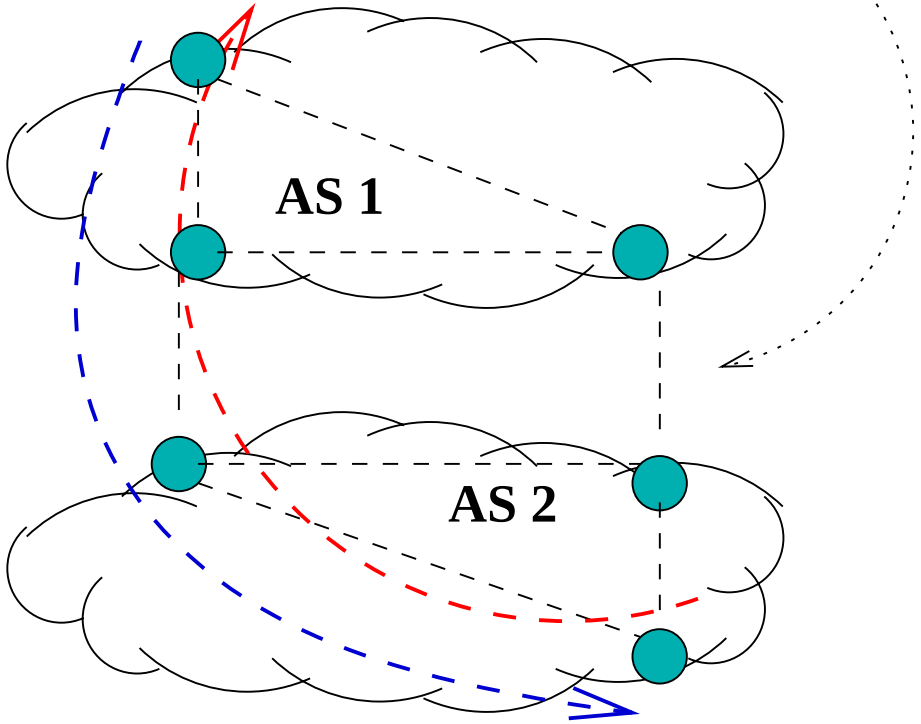

This Work Addresses Three Questions

- Why do inconsistent advertisements happen?
- How can an AS detect inconsistent advertisements?
- How often do they appear in practice?

Why do inconsistent advertisements happen?

- Intentional Cheating
- Misconfiguration
  - Inconsistent export policies to a peer
  - iBGP signaling partitions
  - Inconsistent routes from customer
- Fundamental Limitations
  - Mix of customer and peer routes
    (not all routers in the AS will select routes from the same AS)

Inconsistent Export Policies

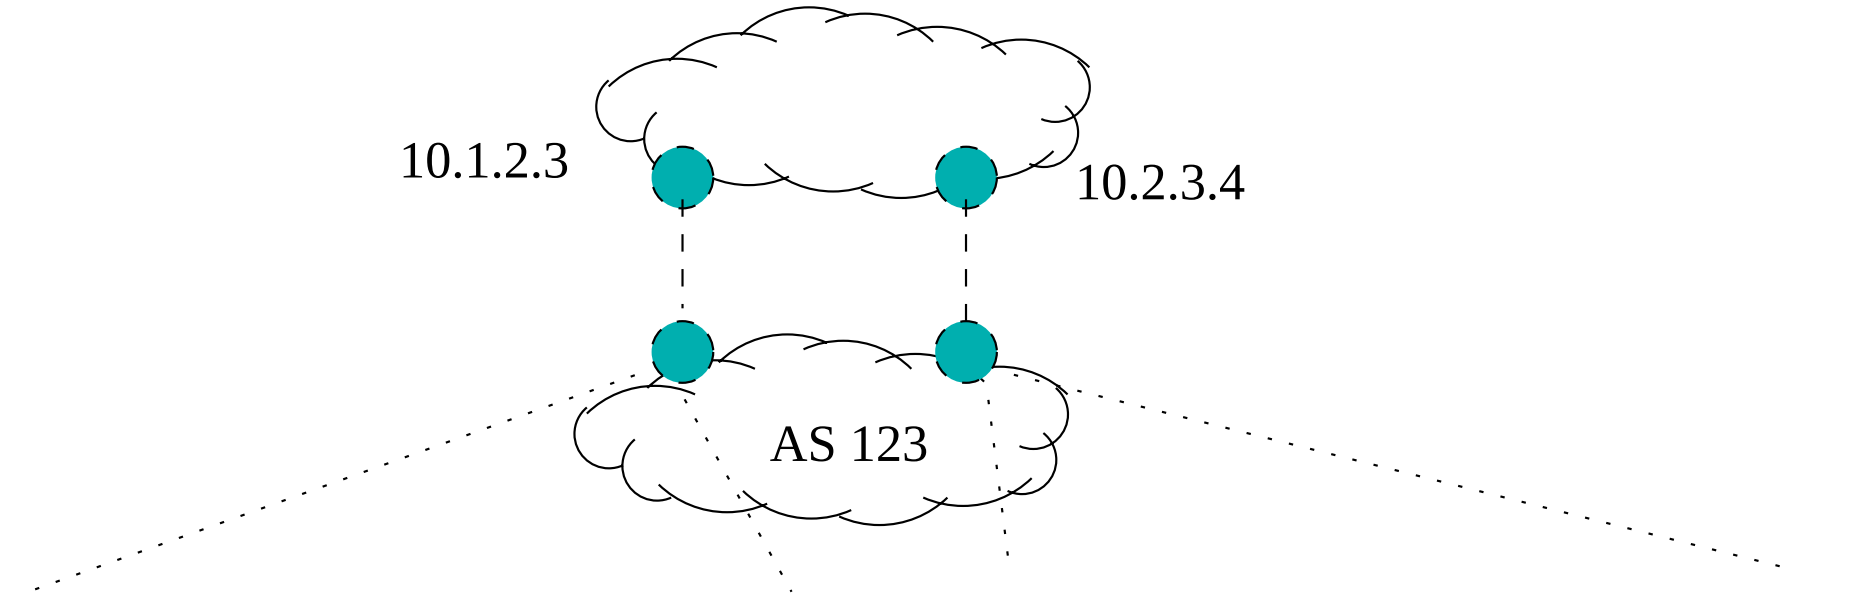

"rcc" detects inconsistent export policies.

http://nms.lcs.mit.edu/bgp/rcc/

iBGP Signaling Partitions

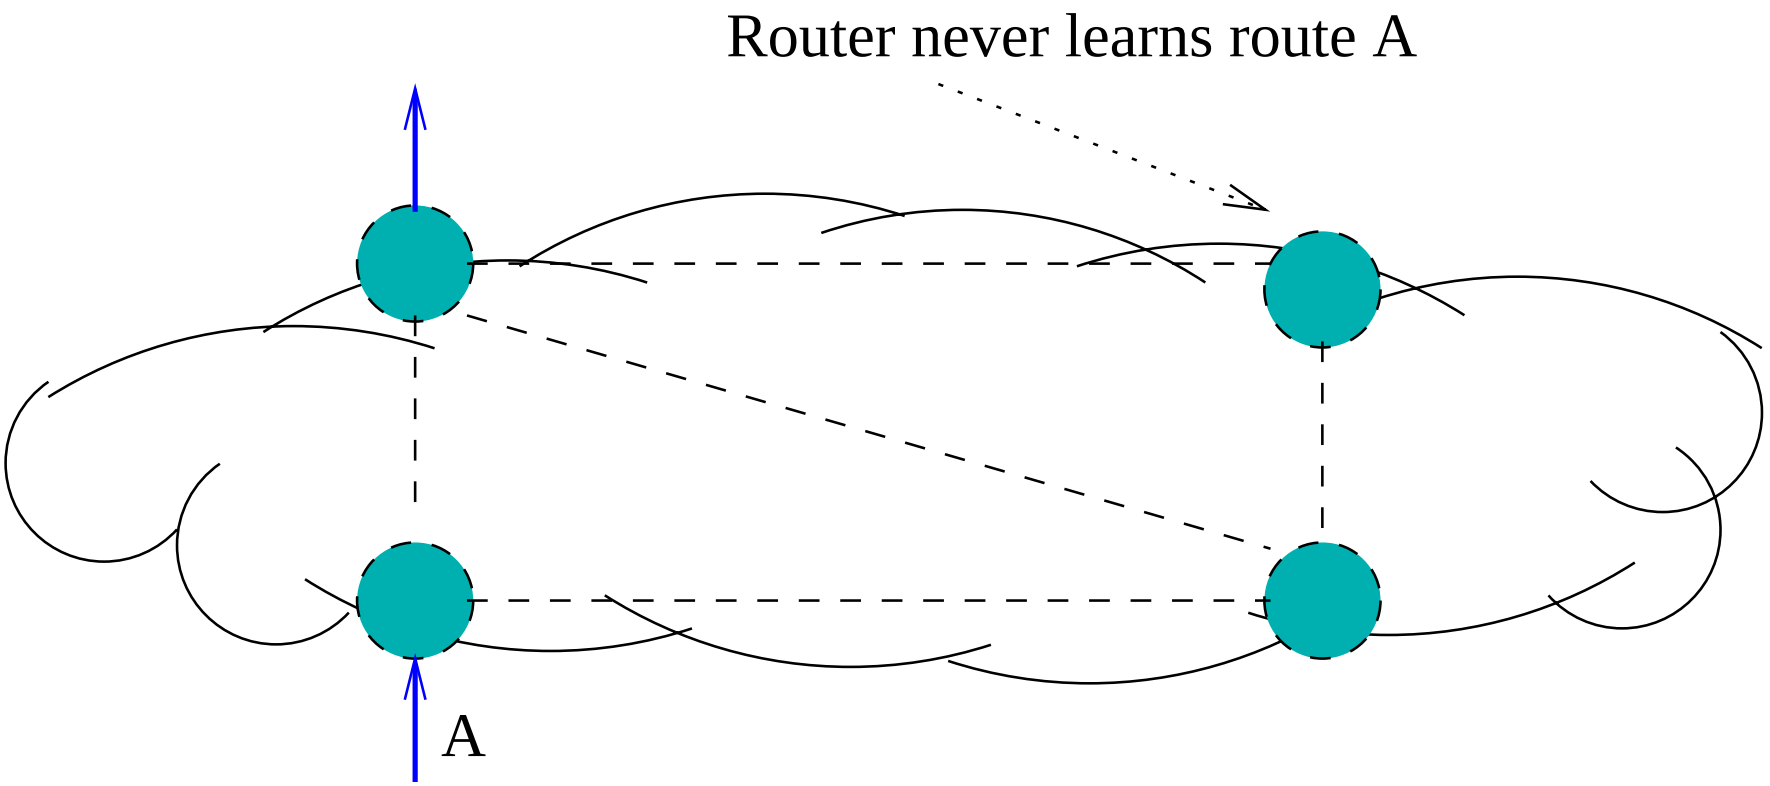

When iBGP is not configured correctly, all routers may not learn "equally good" routes.

iBGP Signaling Partitions

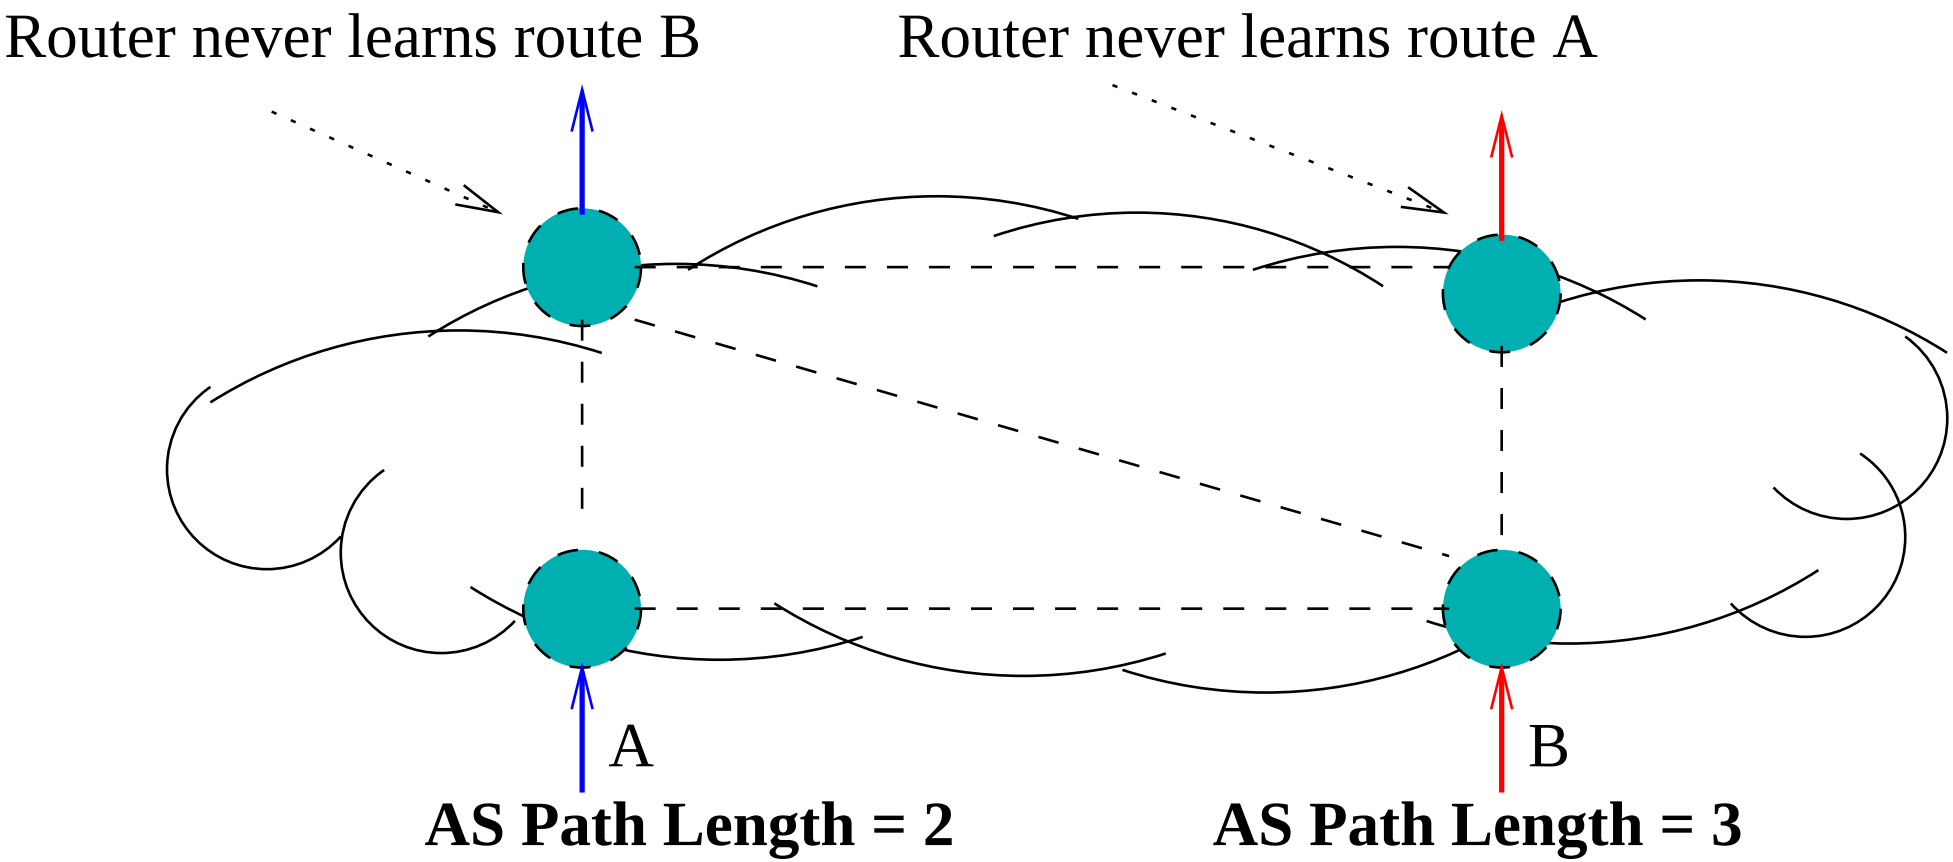

When iBGP is not configured correctly, all routers may not learn "equally good" routes.

iBGP Signaling Partitions

Route reflectors can cause more subtle partitions.

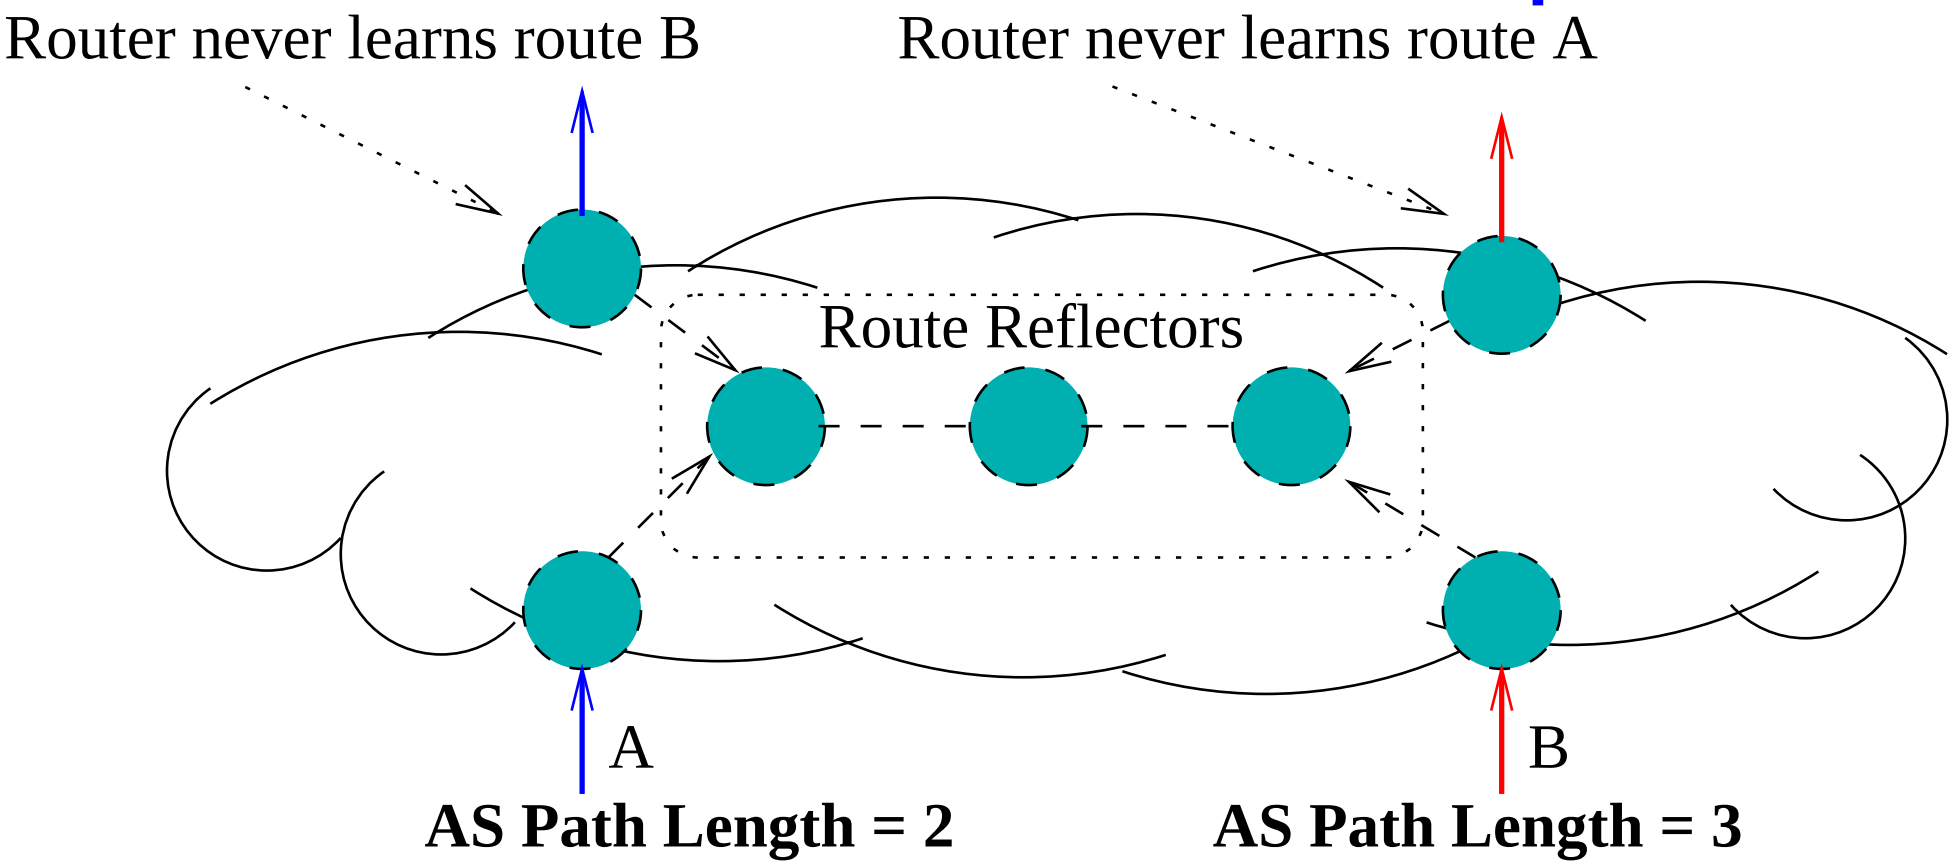

When iBGP is not configured correctly, all routers may not learn "equally good" routes.

(rcc also detects iBGP signaling partitions.)

Inconsistent Advertisements from a Customer

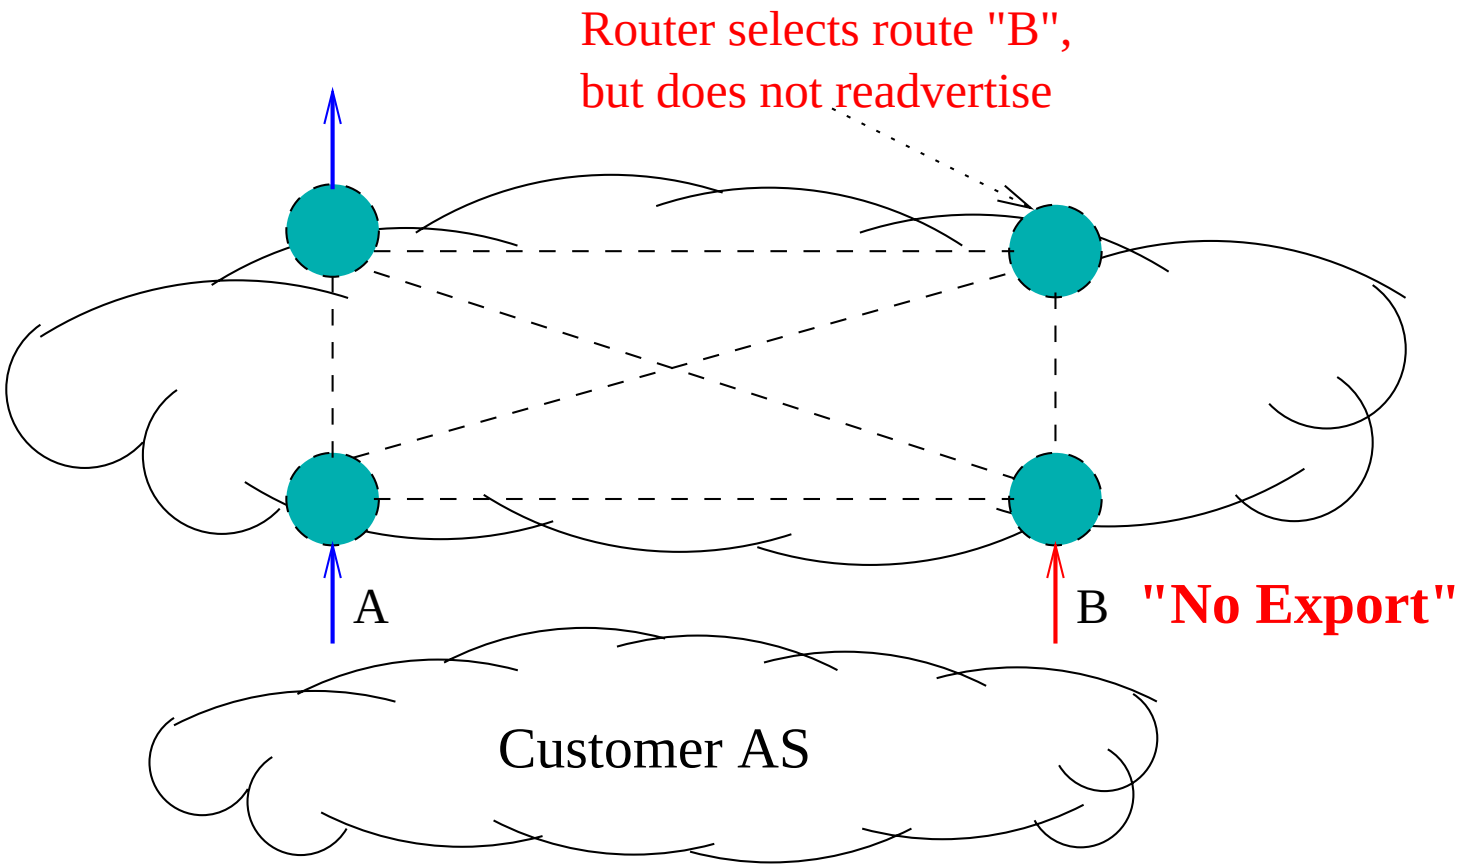

- Some routers may select routes with "no export" community.

Mix of Customer and Peer Routes

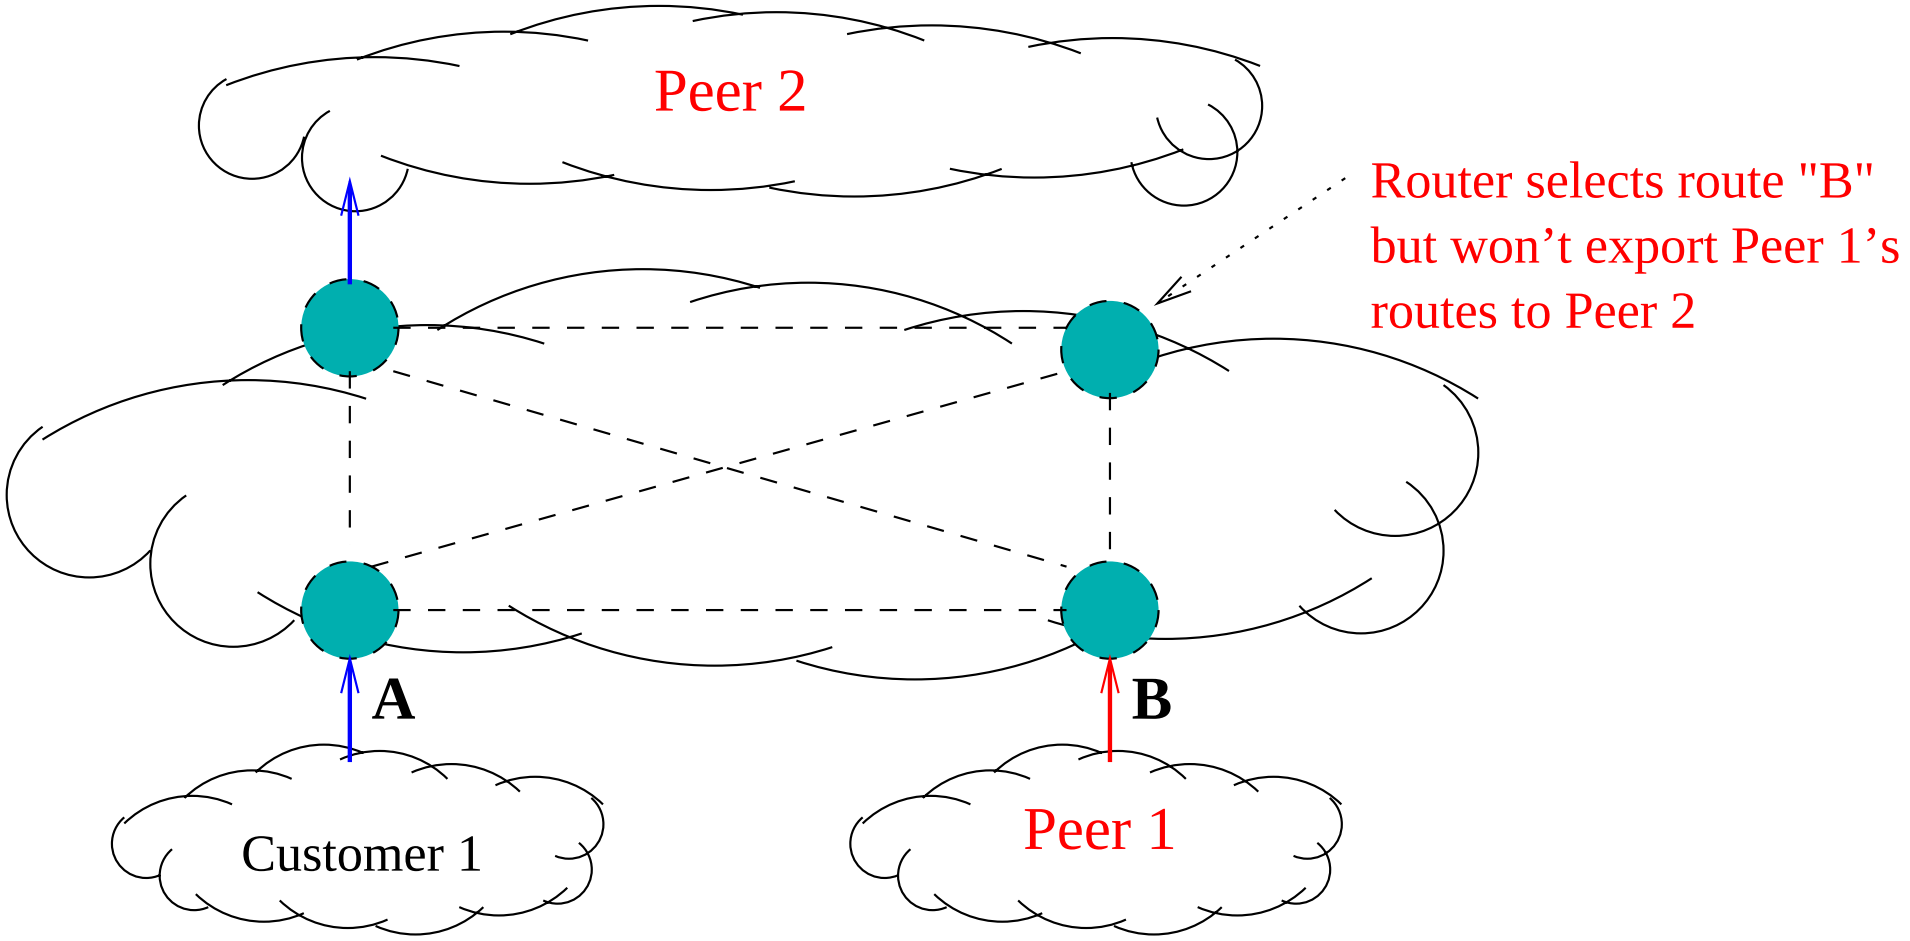

- Some routers may select "customer" routes while others select "peer" routes.

Mix of Customer and Peer Routes

Date: Thu Mar 13 1997 - 02:14:00 EST
From: Randy Bush <randy@psg.com>
Subject: consistent policy != consistent announcements

I have a peer P who connects to multihomed site M. I also have a customer C who connects to M. I will see M as either "P M" or "C M" - again equal length AS paths. If I follow the 'normal' BGP path selection rules (and don't always prefer customers over peers), then in some portions of my network, I will prefer M via P and in others I will prefer M via C.

Therefore, I will not announce any route to M to my peers in some locations, as I don't announce peers to other peers, and in others I will announce "C M". Again, I do not make identical announcements to my peers, yet I have a consistent policy.

Problem Statement:

Detect when a neighboring AS does not announce "equally good" routes on all peering sessions.

Ideally: Compare all eBGP routes from a peer

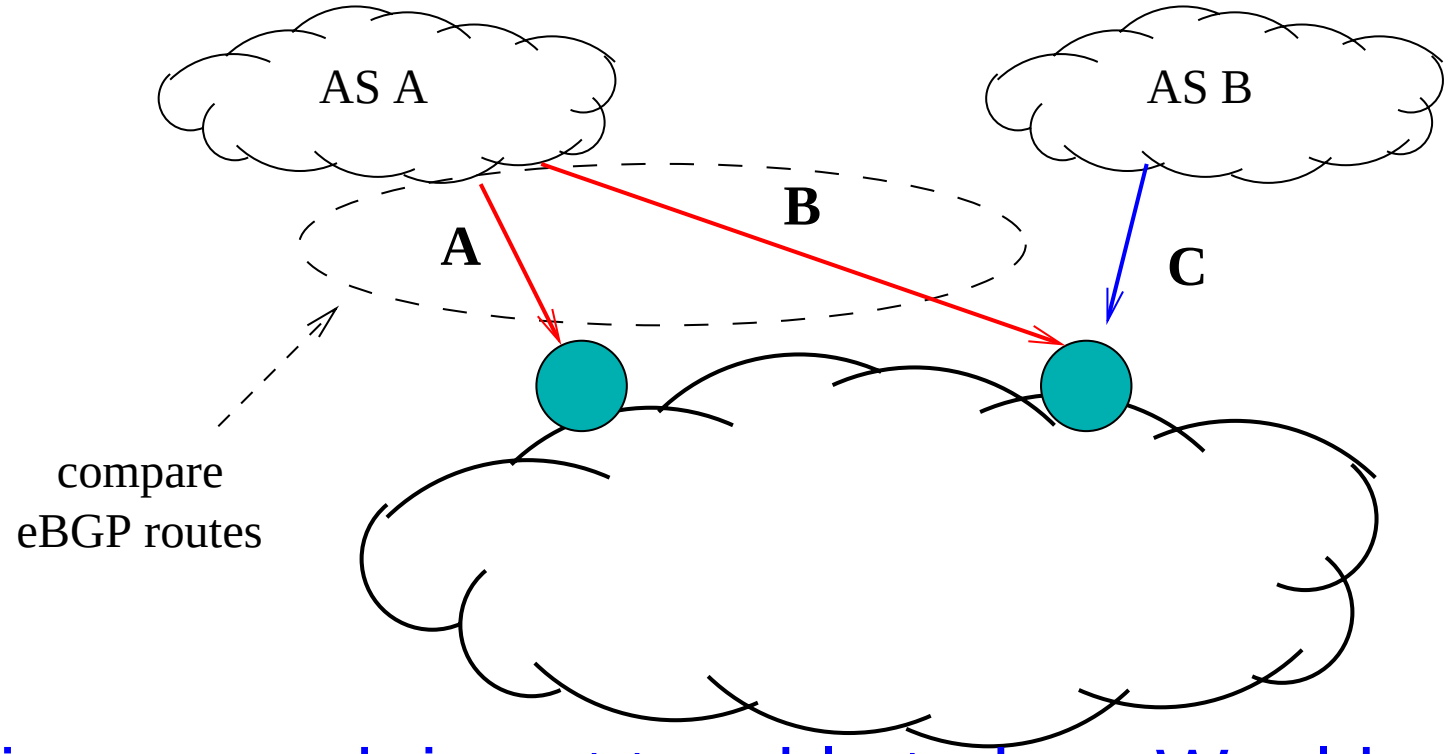

This approach is not tenable today. Would require one of:

- vendor enhancements
- separate BGP feeds from all peers
- packet monitors on every link

Detecting inconsistent advertisements is hard

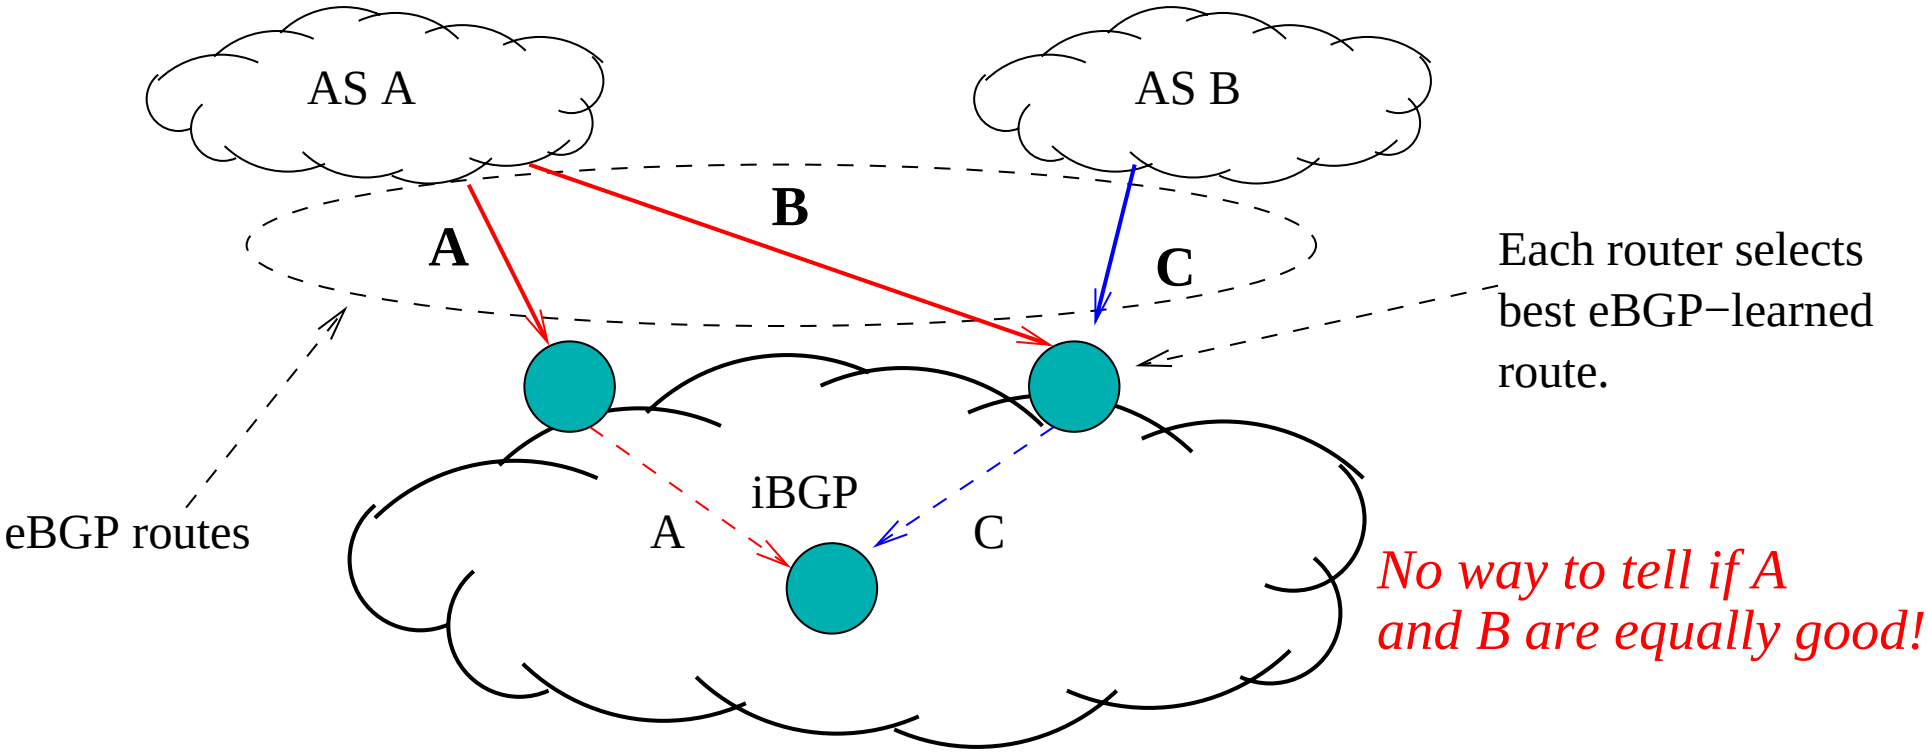

Problem:

- Can’t see all eBGP-learned routes
- Thus, can’t always see all routes from a peer.

Questions

- How to detect inconsistent advertisements with only incomplete monitoring? (i.e., only the best route from each router)

Prevalence of inconsistent advertisements in practice:

- How often do inconsistent route advertisements happen in practice?
- How long do inconsistent route episodes last?
- How "severe" is the resulting cold potato routing?

Making inferences about "invisible" routes

1. Invert the import policies at each router
2. Infer properties of routes we don't see based on the best route

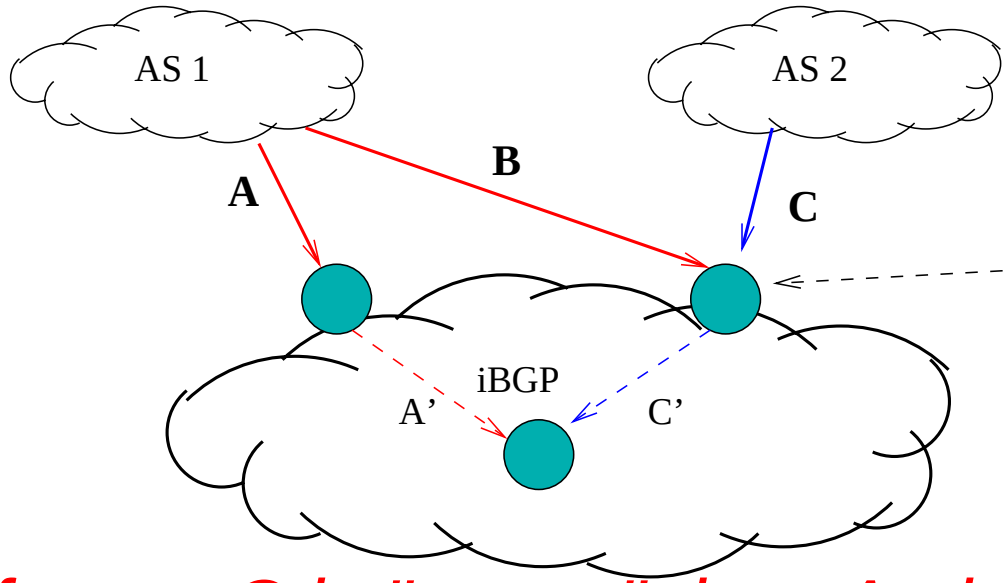

1. Figure out what route C looked like before import policy.
2. Would like to say: "C is at least as good as B."

If route C is "worse" than A, then either:

- B was worse than A

- B was not advertised

When can we do this, given only iBGP routes?

The import policy:

- Must be invertible.
- Must not make consistent routes inconsistent.

Non-Invertible Import Policy

If AS path == "1" then
  prepend "1 1"
If AS path == "1 1" then
  prepend "1"

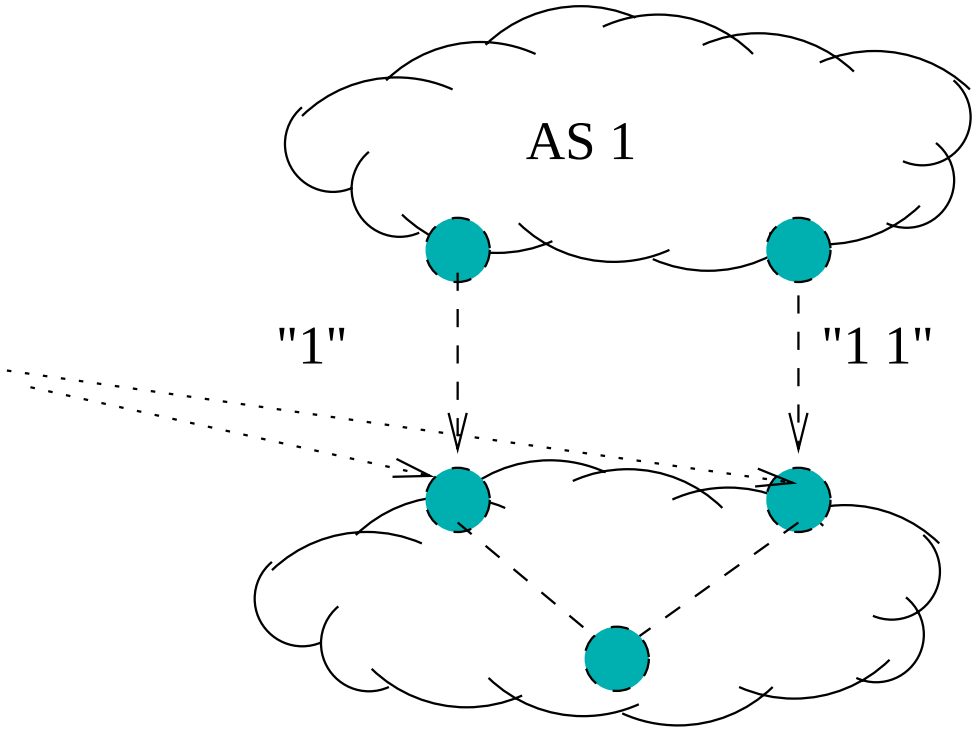

Both routes have AS path "1 1 1"

Impossible to tell from iBGP routes alone whether eBGP routes are consistent.

Import Policy that Creates Inconsistent Routes

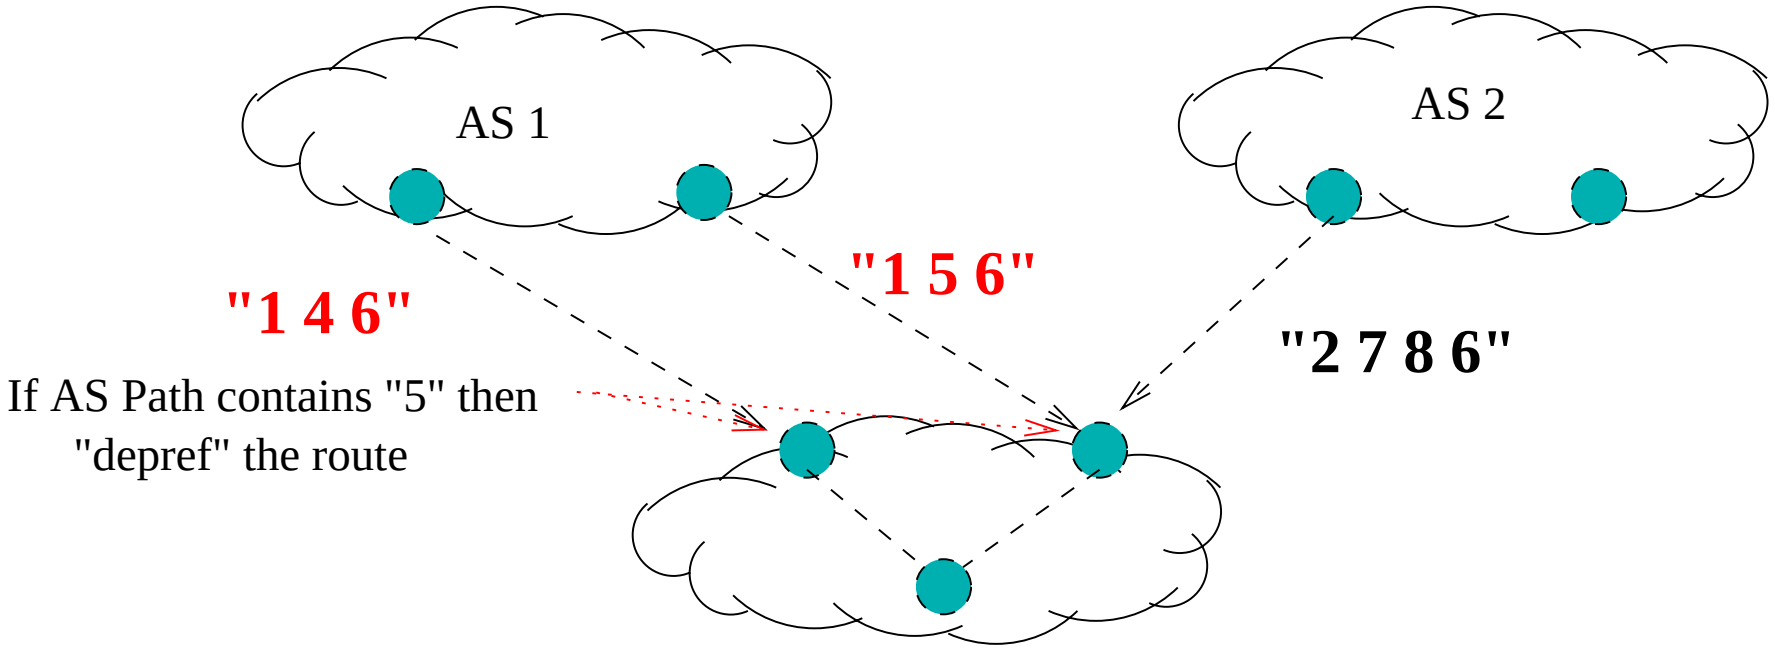

Import policy takes two routes with equal AS path lengths and makes one route look worse.

Impossible to tell from iBGP routes alone whether eBGP routes are consistent.

Questions

Prevalence of inconsistent advertisements in practice:

- How often do inconsistent route advertisements happen in practice?
- How long do inconsistent route episodes last?
- How "severe" is the resulting cold potato routing?

Inconsistent Advertisements Occur Continually

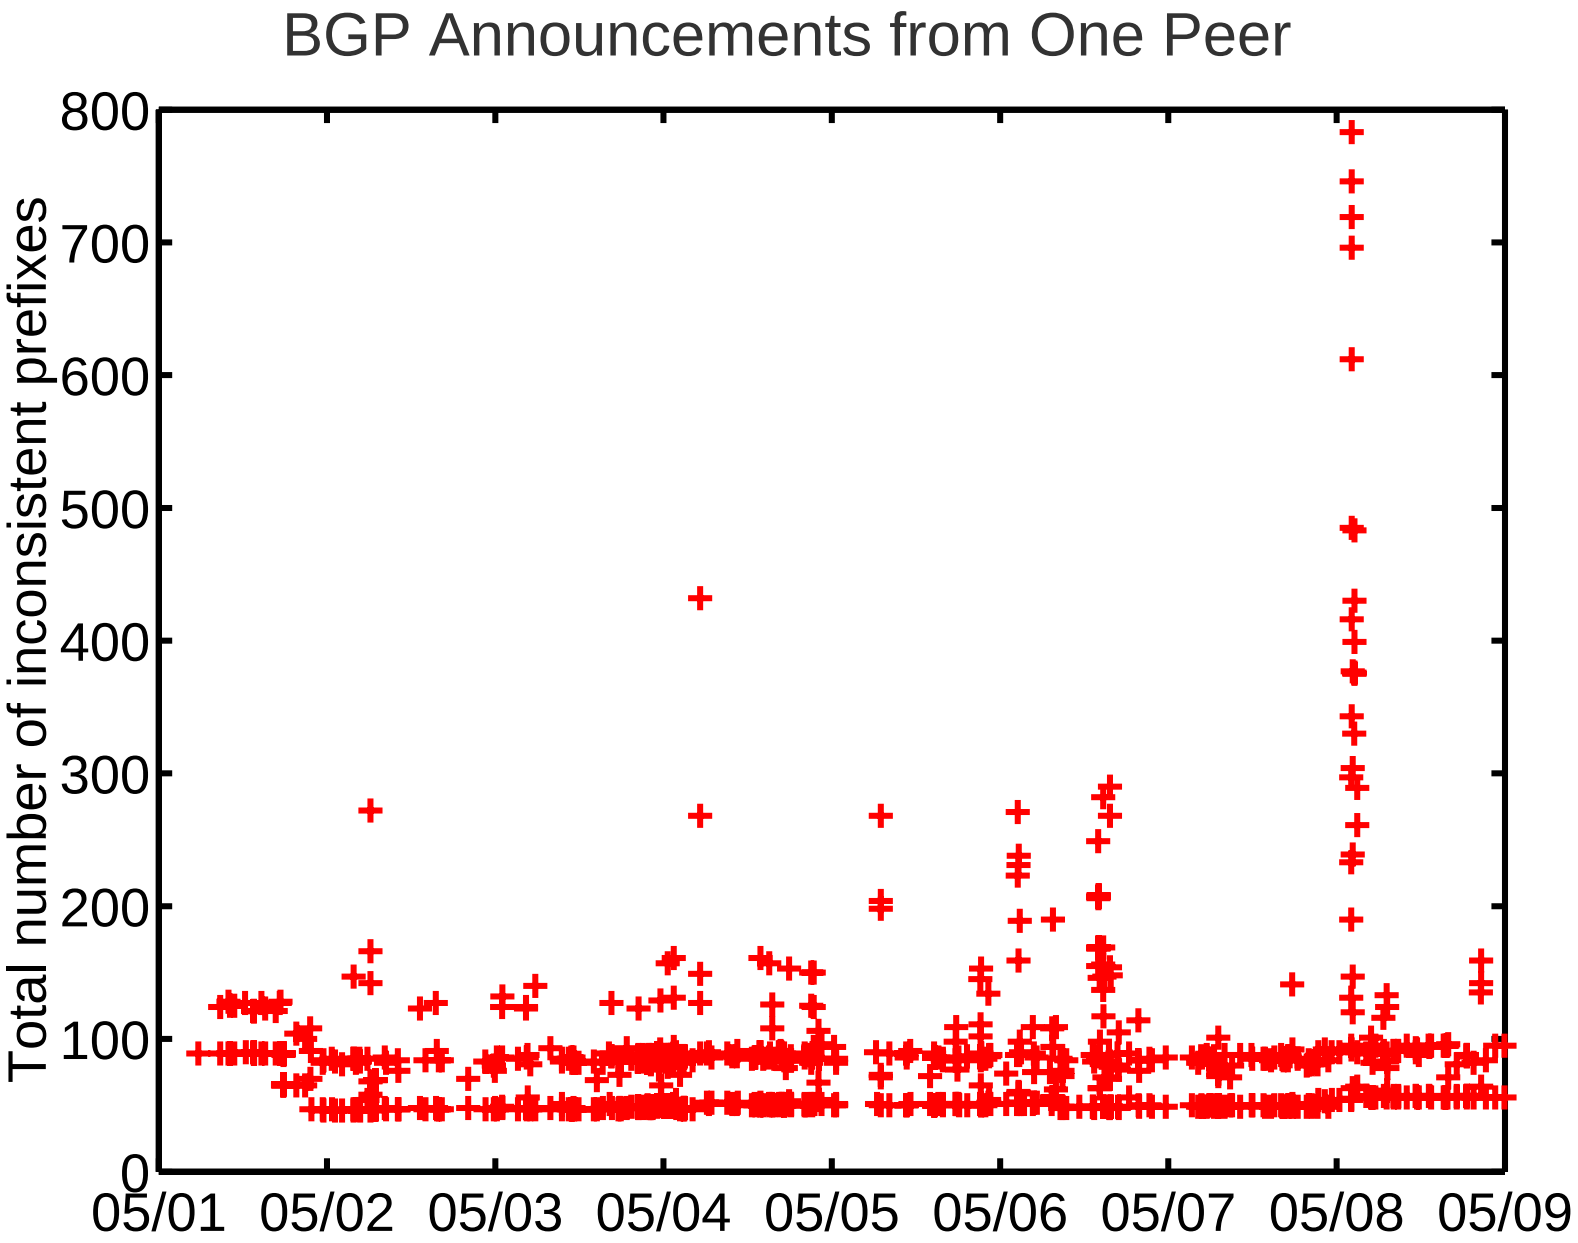

Most Events are Transient, Some Persistent

- Missing Routes:
  - 99% are shorter than 100 seconds.
  - 0.1% of routes are inconsistent for entire week.
- Different AS Path Lengths:
  - 99% are shorter than 100 seconds.
  - All are shorter than 5 minutes.

BGP convergence causes
most, but not all, inconsistencies.

Most Prefixes Consistent Most of the Time

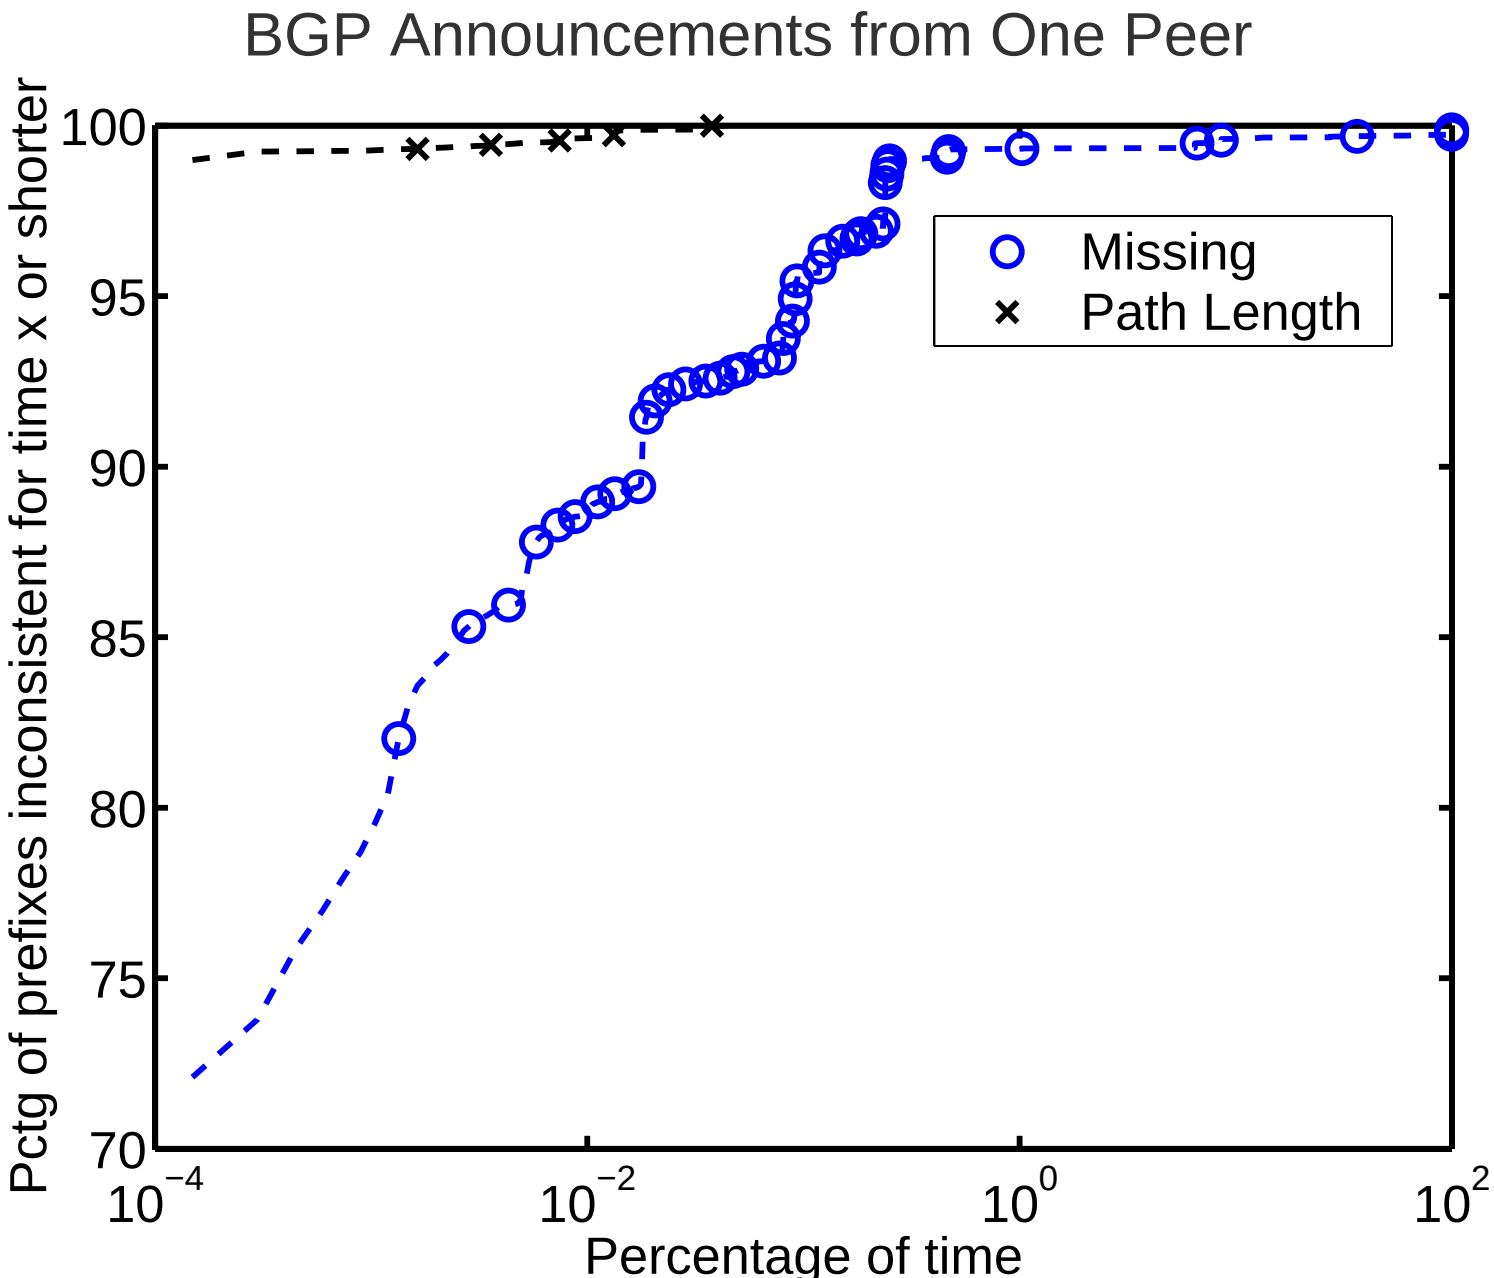

Some Peers are Persistently Inconsistent

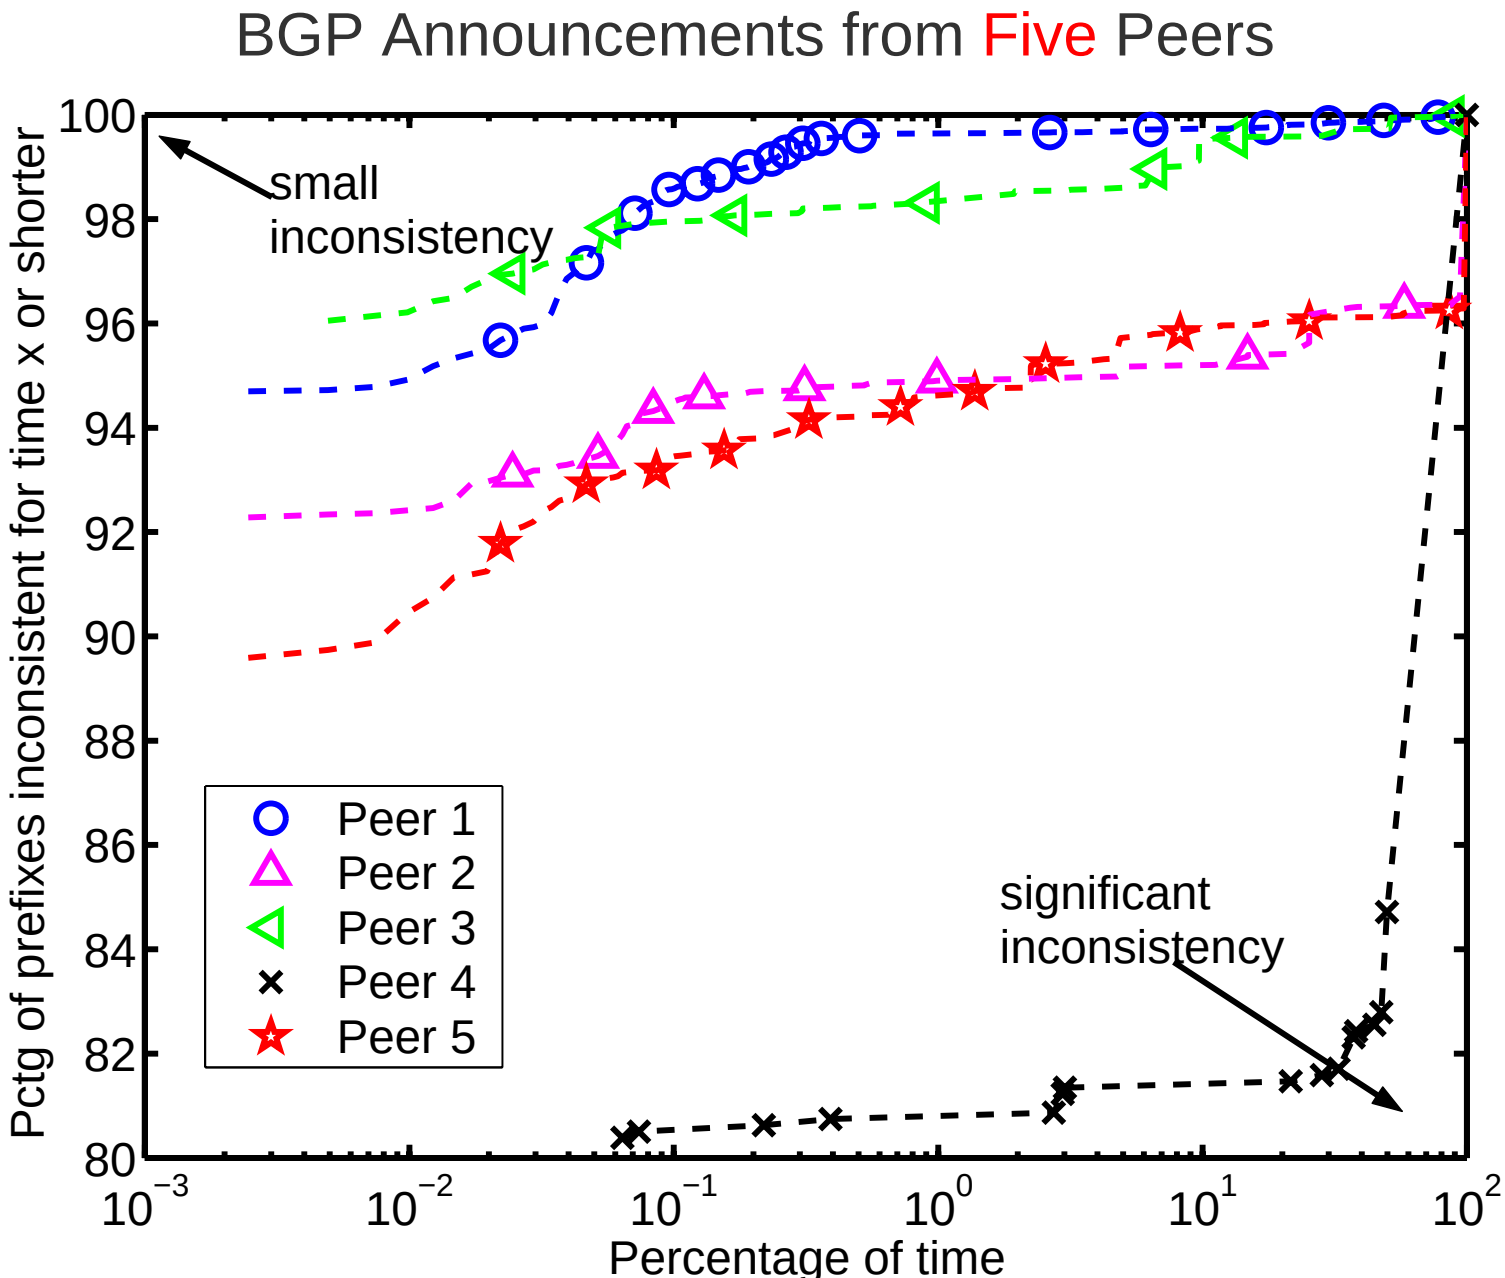

Conclusions

- Settlement-free peering arrangements are violated.
  - Until now, there has been no good way to detect these violations.

A monitoring protocol that dumped all eBGP routes would make this problem easy!

- Thousands of inconsistent advertisements in a week.
  - Most of these are due to convergence.
    (99% of inconsistencies last less than ~2 minutes.)
  - Missing routes are more common that path length differences.

- Most peers advertise consistently in steady state.

- Some peers have consistent inconsistencies.
  - In one case, more than 15% of prefixes were inconsistent for 70% of the trace.

Most Events are Transient, Some Persistent