

BorderGuard: Detecting Cold Potatoes from Peers


Nick Feamster, M.I.T.

Z. Morley Mao, University of Michigan

Jennifer Rexford, AT&T Labs -- Research

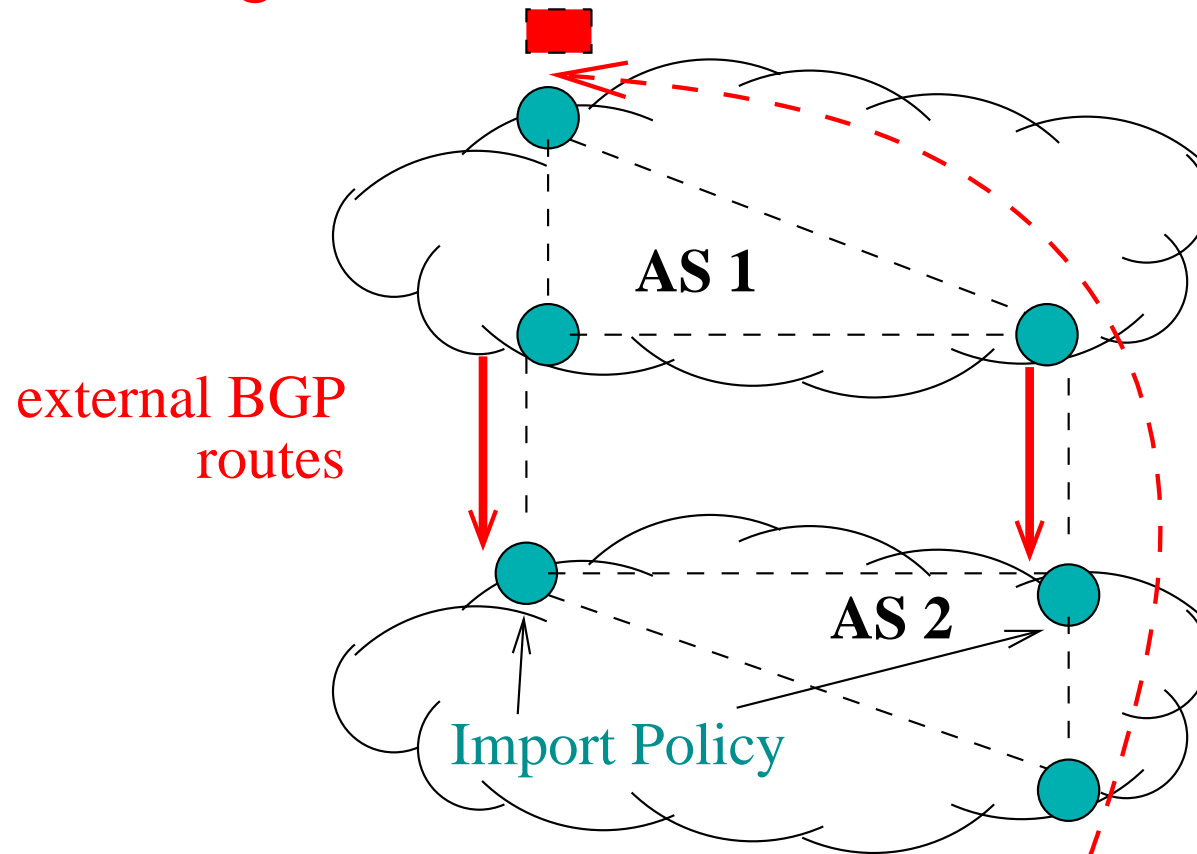
Networks Exchange Traffic in Multiple Places

Hot-Potato Routing:



Networks Exchange Traffic in Multiple Places

Hot-Potato Routing:




Peers Typically Exchange Consistent Routes

Rules of "settlement-free peering":


1. Advertise routes to its peer at all peering points.
2. The routes must have equal AS path length.

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!



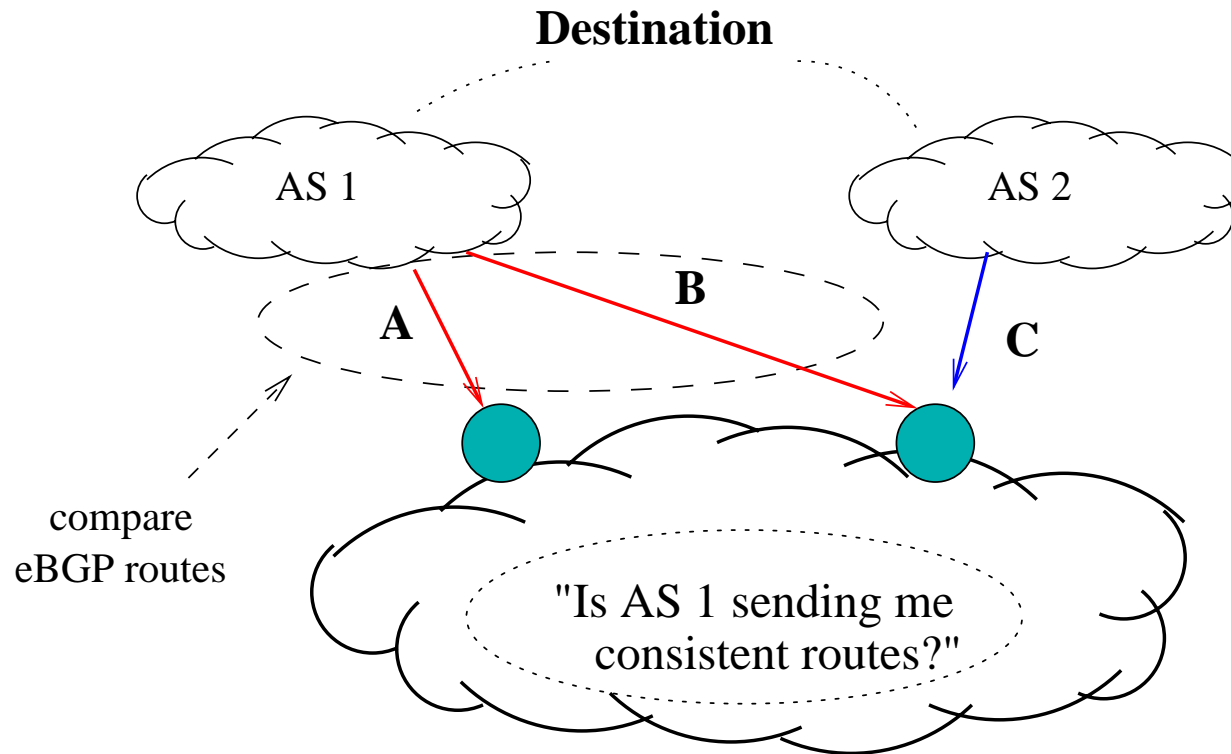
Why do inconsistent advertisements happen?

- Intentional Cheating
- Convergence
- Misconfiguration
 - ▶ Inconsistent export policies to a peer
 - ▶ iBGP partitions
 - ▶ Inconsistent advertisements from customer
- Fundamental Limitations
 - ▶ Feature interaction between multiple goals
 - ◆ Don't export routes from one "peer" to another
 - ◆ Selecting routes with the shortest AS path length

Questions

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

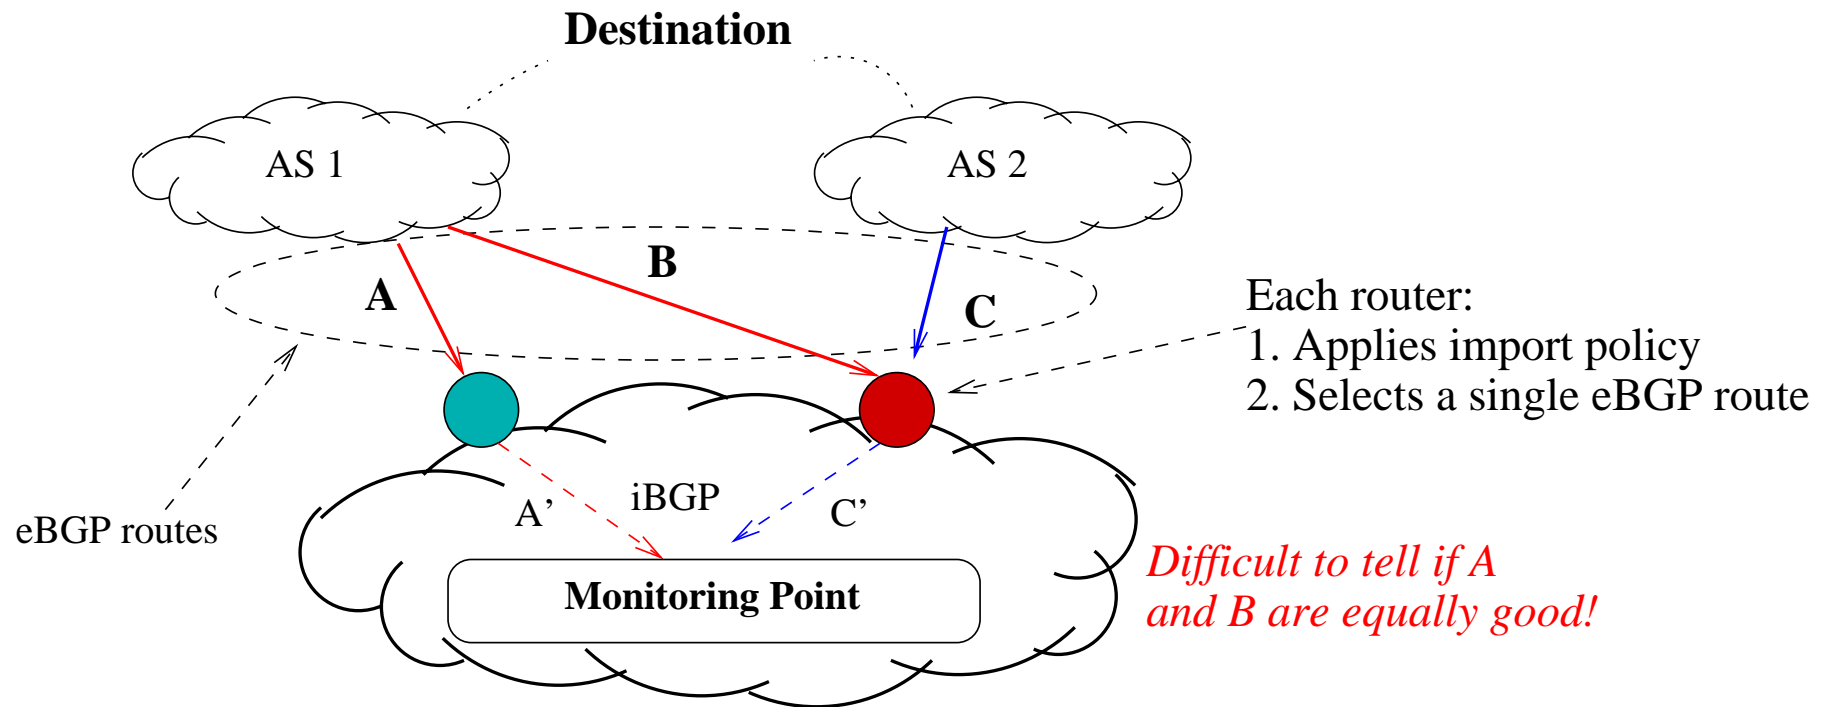
Ideally: Compare all eBGP routes from a peer



Unfortunately, we don't see every eBGP route. Need:

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

Reality: Only Each Router's Best BGP Route



Problems:

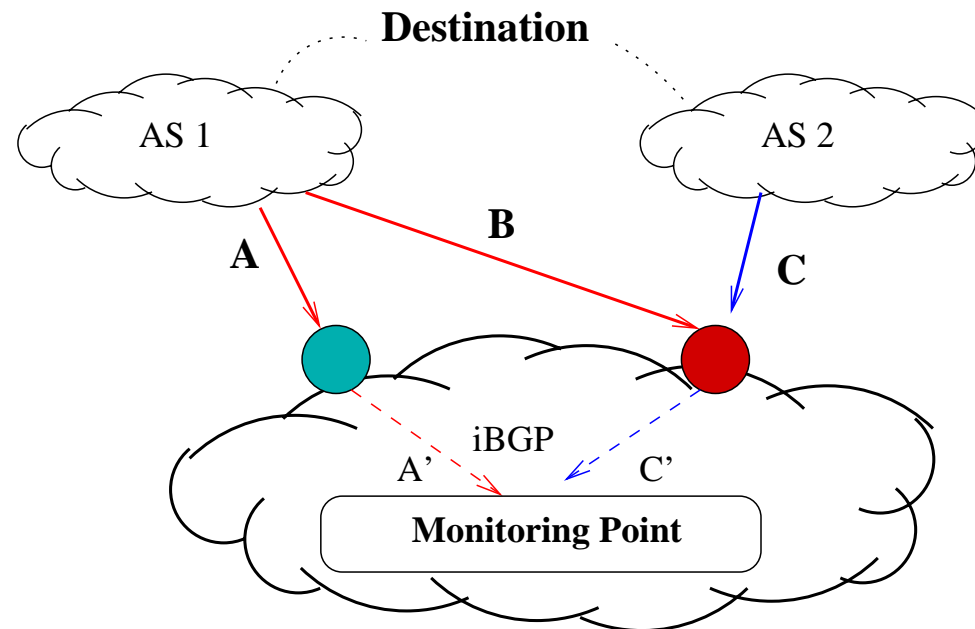
- Can't see all eBGP-learned routes
- Only see iBGP routes after import policy

We do know:

The red router picked C as its best route.

BorderGuard Algorithm


1. Invert the import policies at each router. ($A' \rightarrow A$)
2. Infer properties of routes we don't see based on the best routes.
(B should have the same AS path length as A.)



- If C' is "worse" than B' would have been, then either:
- B has a longer AS path length than A
 - B was not advertised

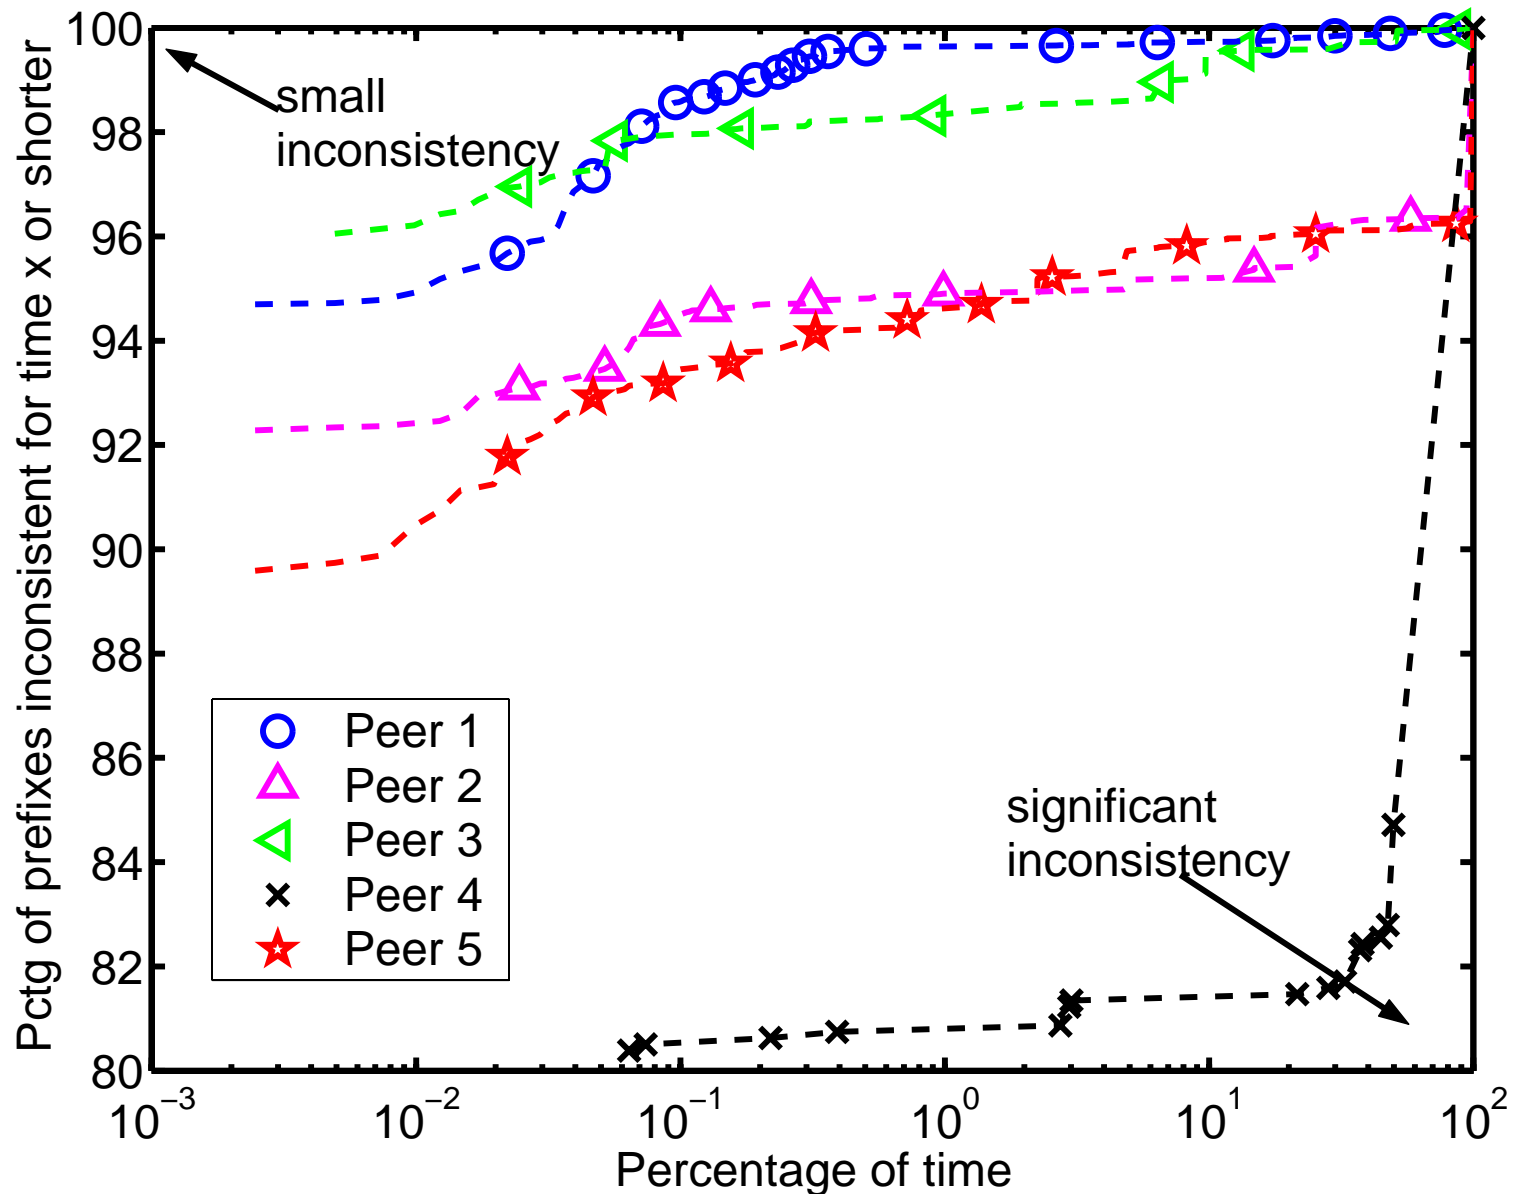
Inconsistent Advertisements are Continual

BGP Announcements from One Peer

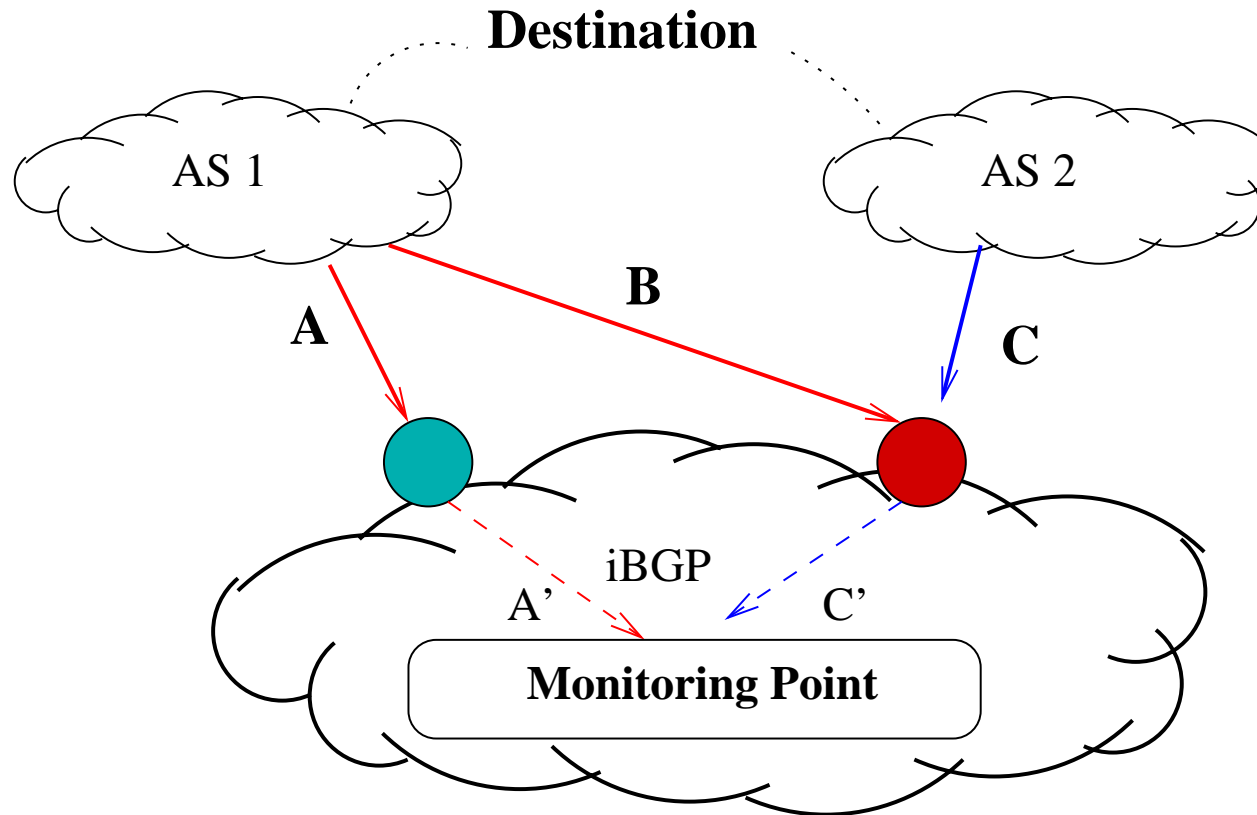


Most Inconsistencies are Short; A Few Long...

BGP Announcements from **Five** Peers



Can we always use this algorithm?



The import policy:

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

Non-Invertible Import Policy


Import Policy:

If AS path == "1" then

prepend "1 1"

If AS path == "1 1" then

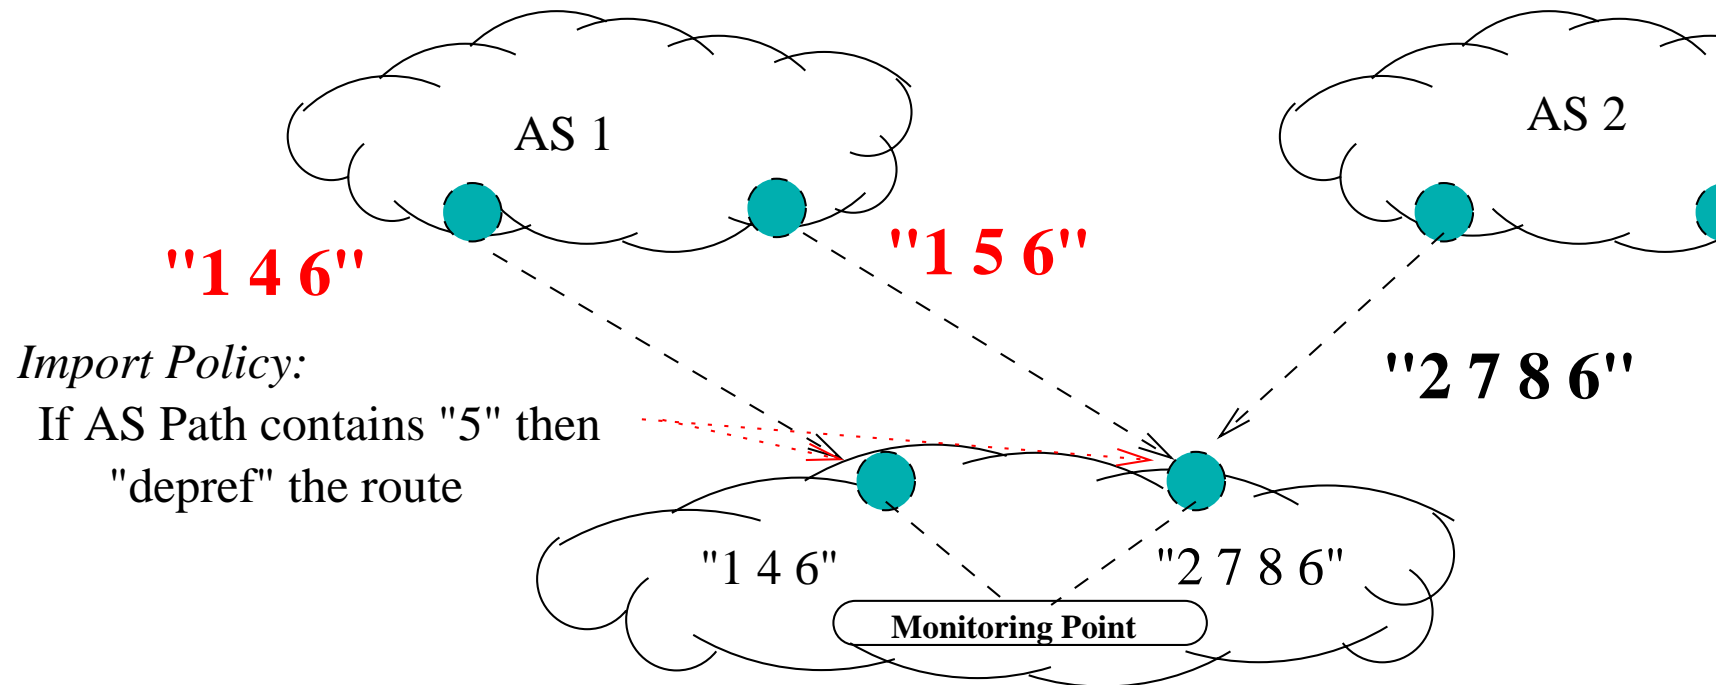
prepend "1"



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.

Conclusion

- Detecting inconsistent routes from peer ASes.
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.)
- Settlement-free peering arrangements are violated.
 - ▶ In one case, more than 15% of prefixes were inconsistent for 70% of the trace.
- Inconsistencies have many causes.
 - ▶ Cheating, misconfiguration, feature interaction



Questions

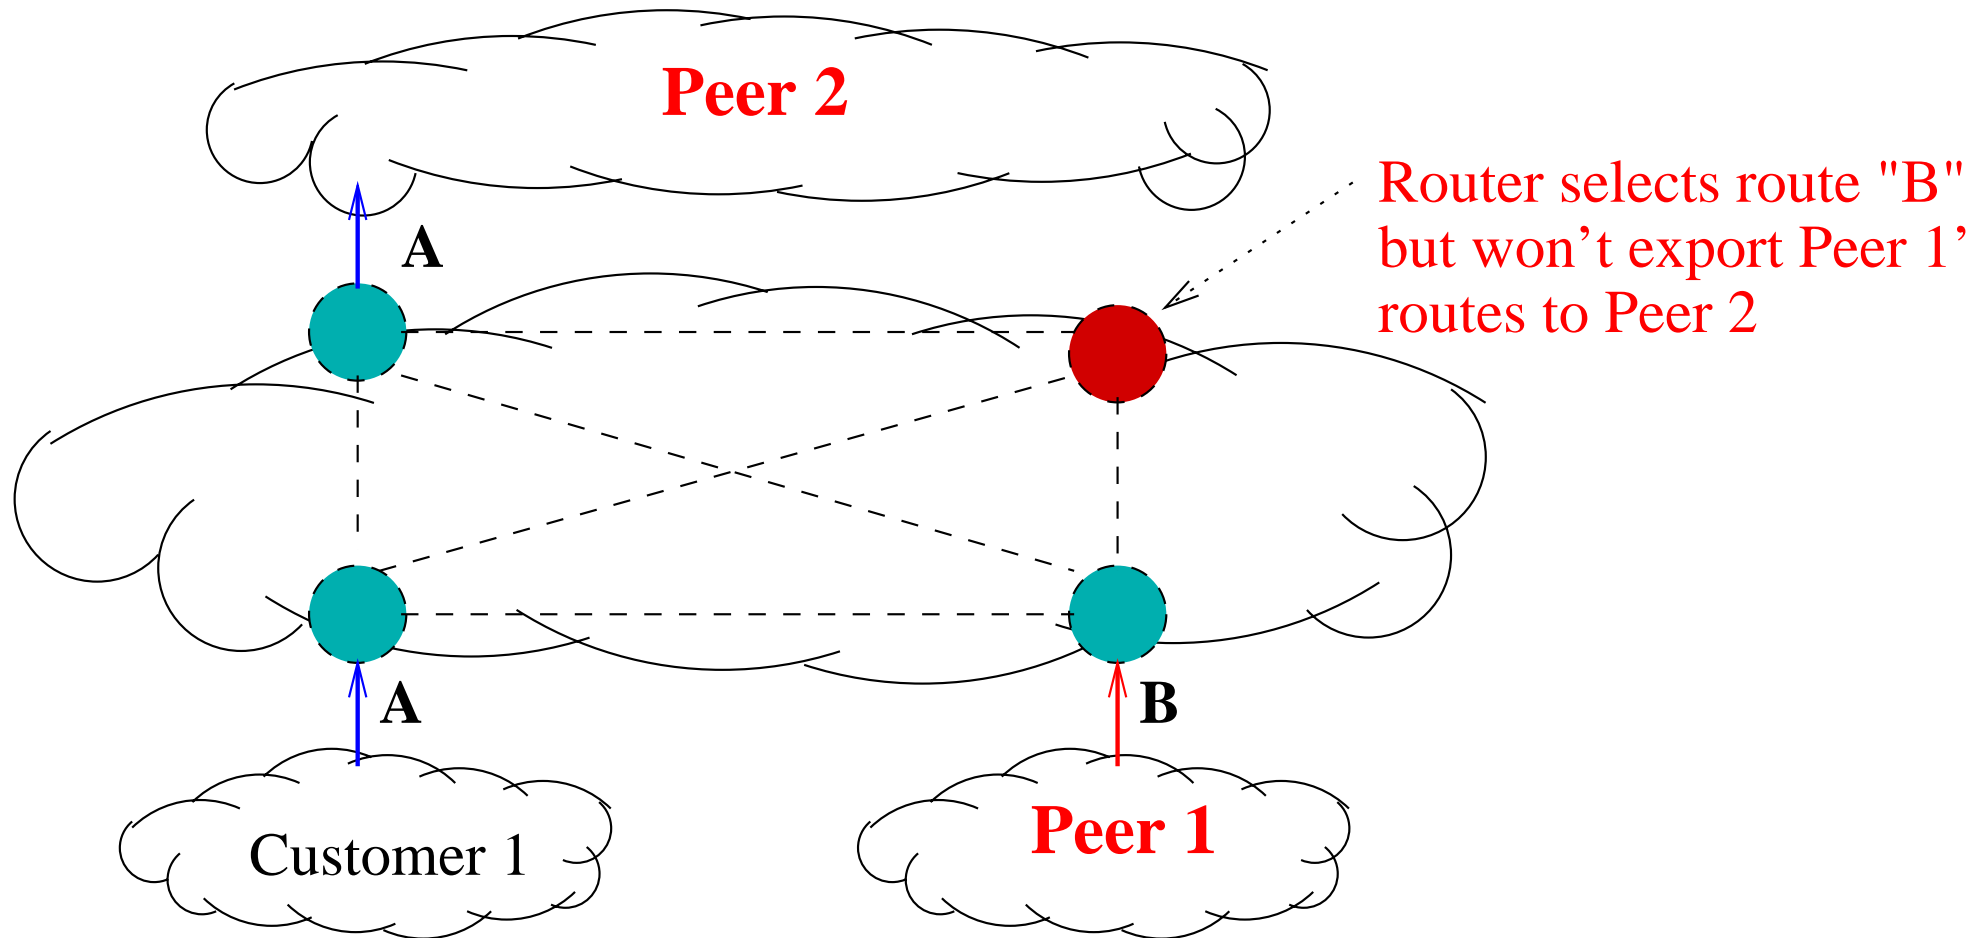
- How to detect inconsistent advertisements with only incomplete monitoring? (i.e., only the best route from each router)

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?

Mix of Customer and Peer Routes

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



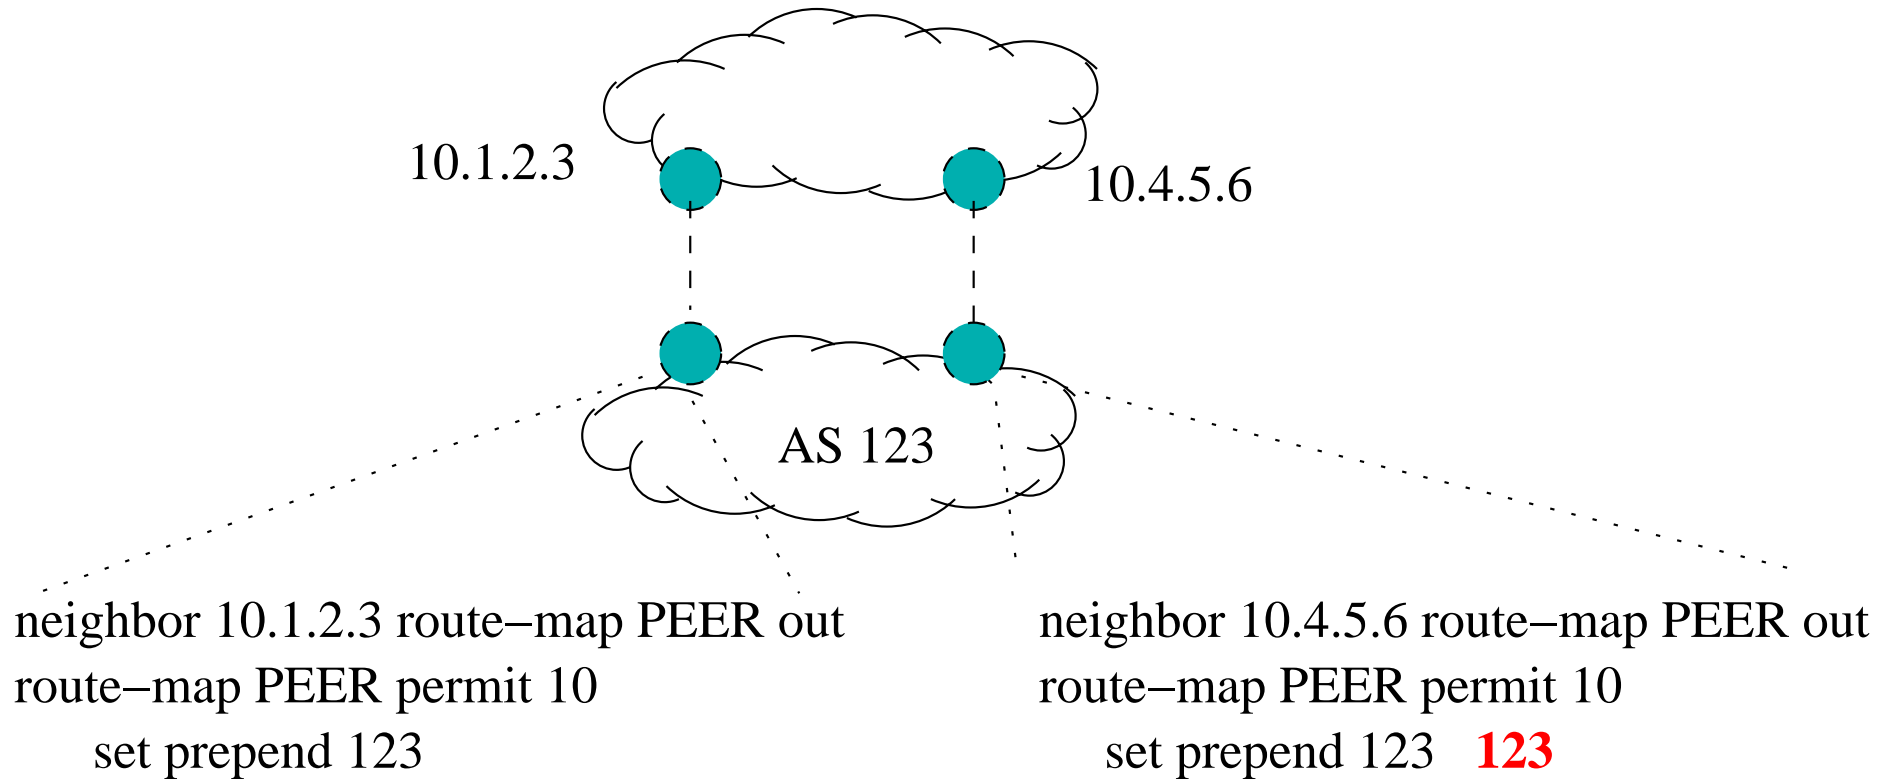
Most Events are Transient, Some Persistent

- 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.

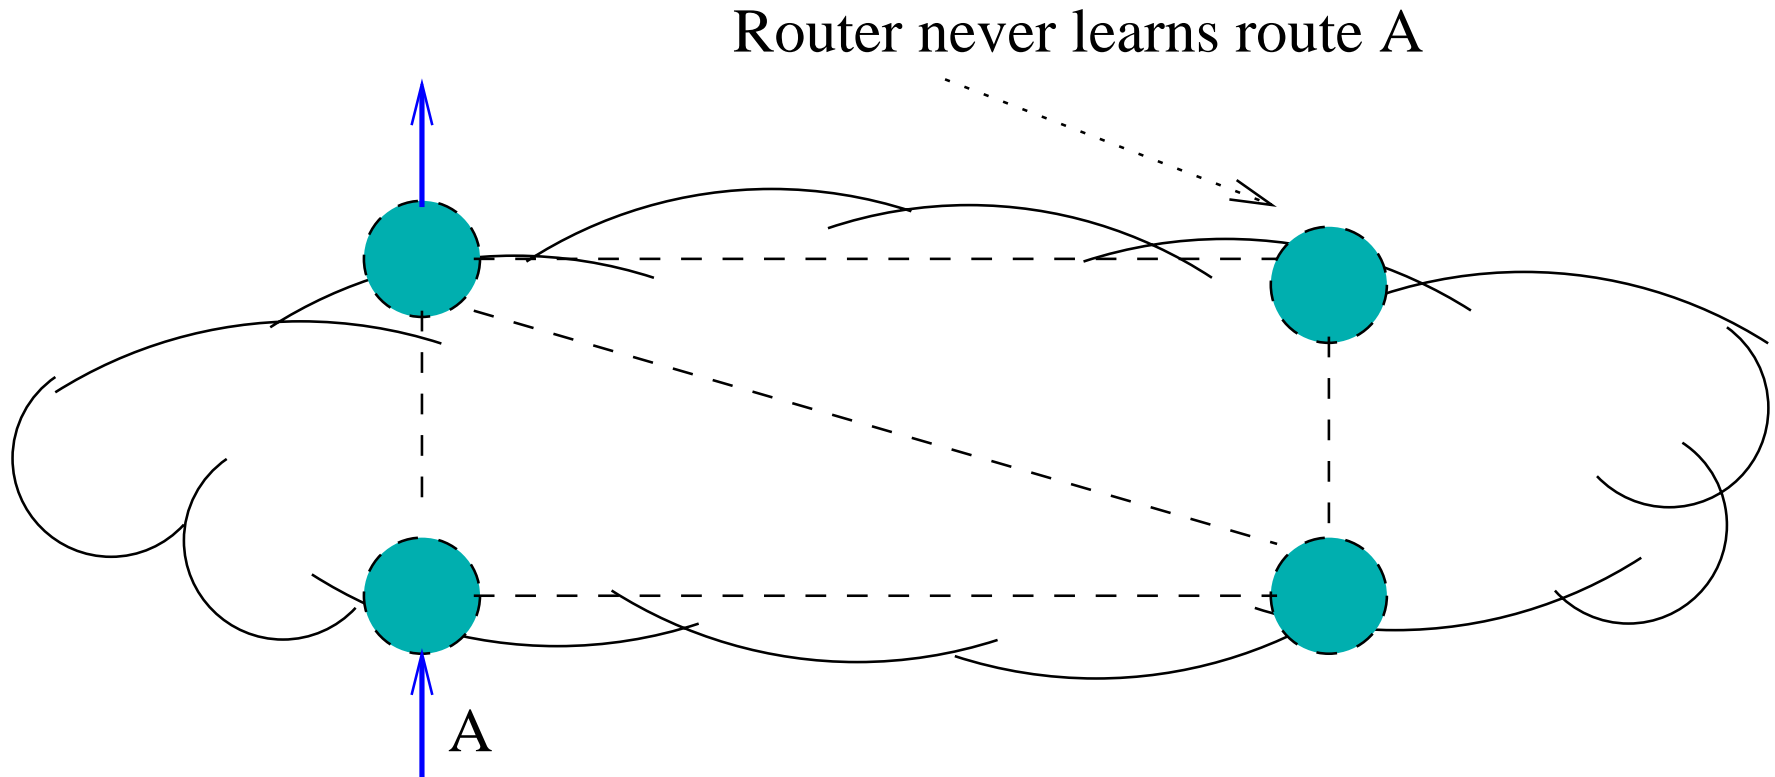
Inconsistent Export Policies

Different routers may apply different export policies.



iBGP Signaling Partitions

iBGP misconfiguration can prevent some routers from learning a route.




iBGP Signaling Partitions

iBGP misconfiguration can prevent some routers from learning the best route.

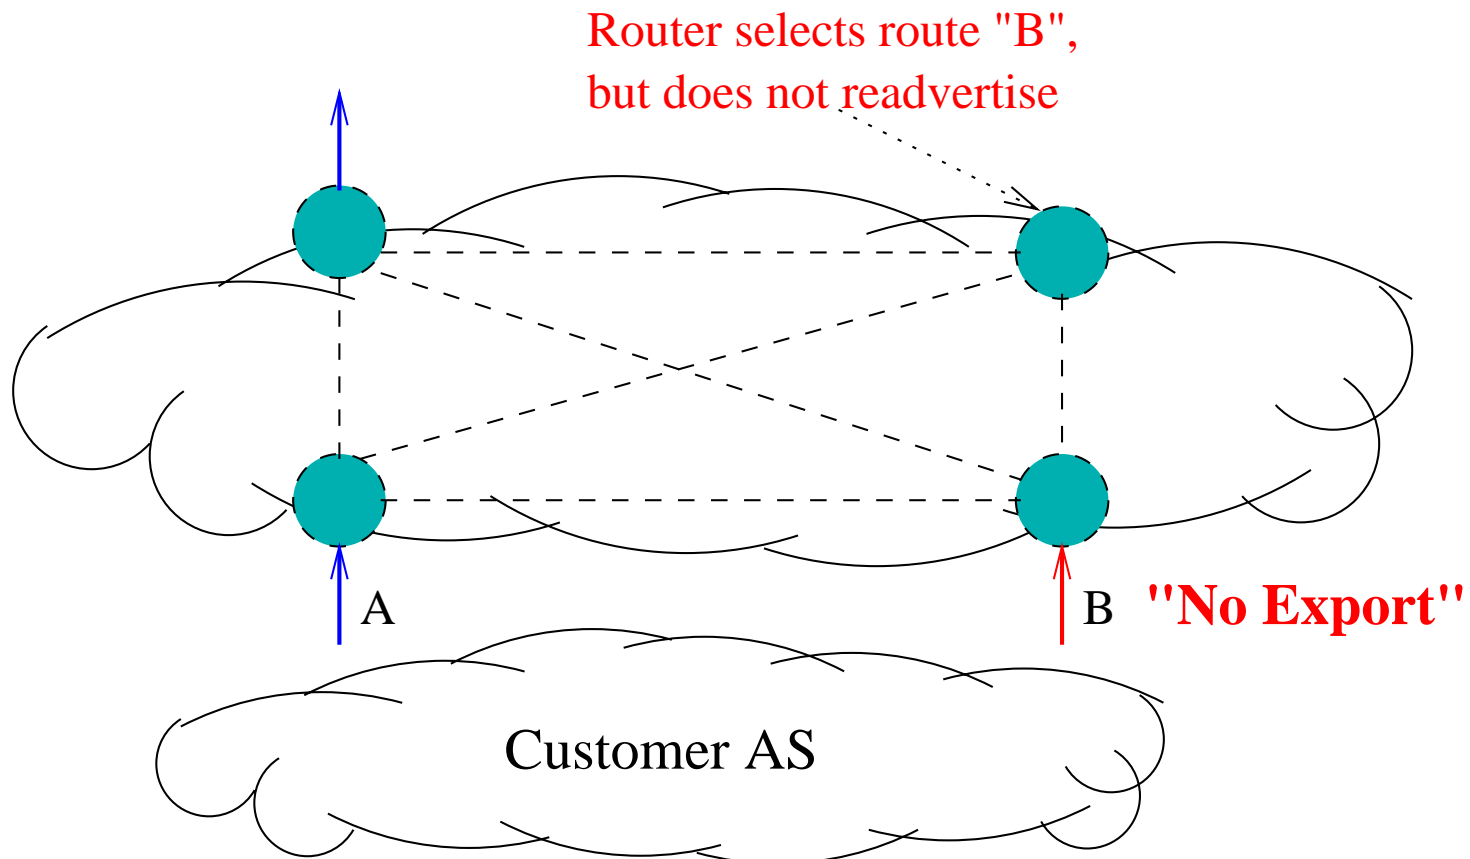
Router never learns route B

Router never learns route A

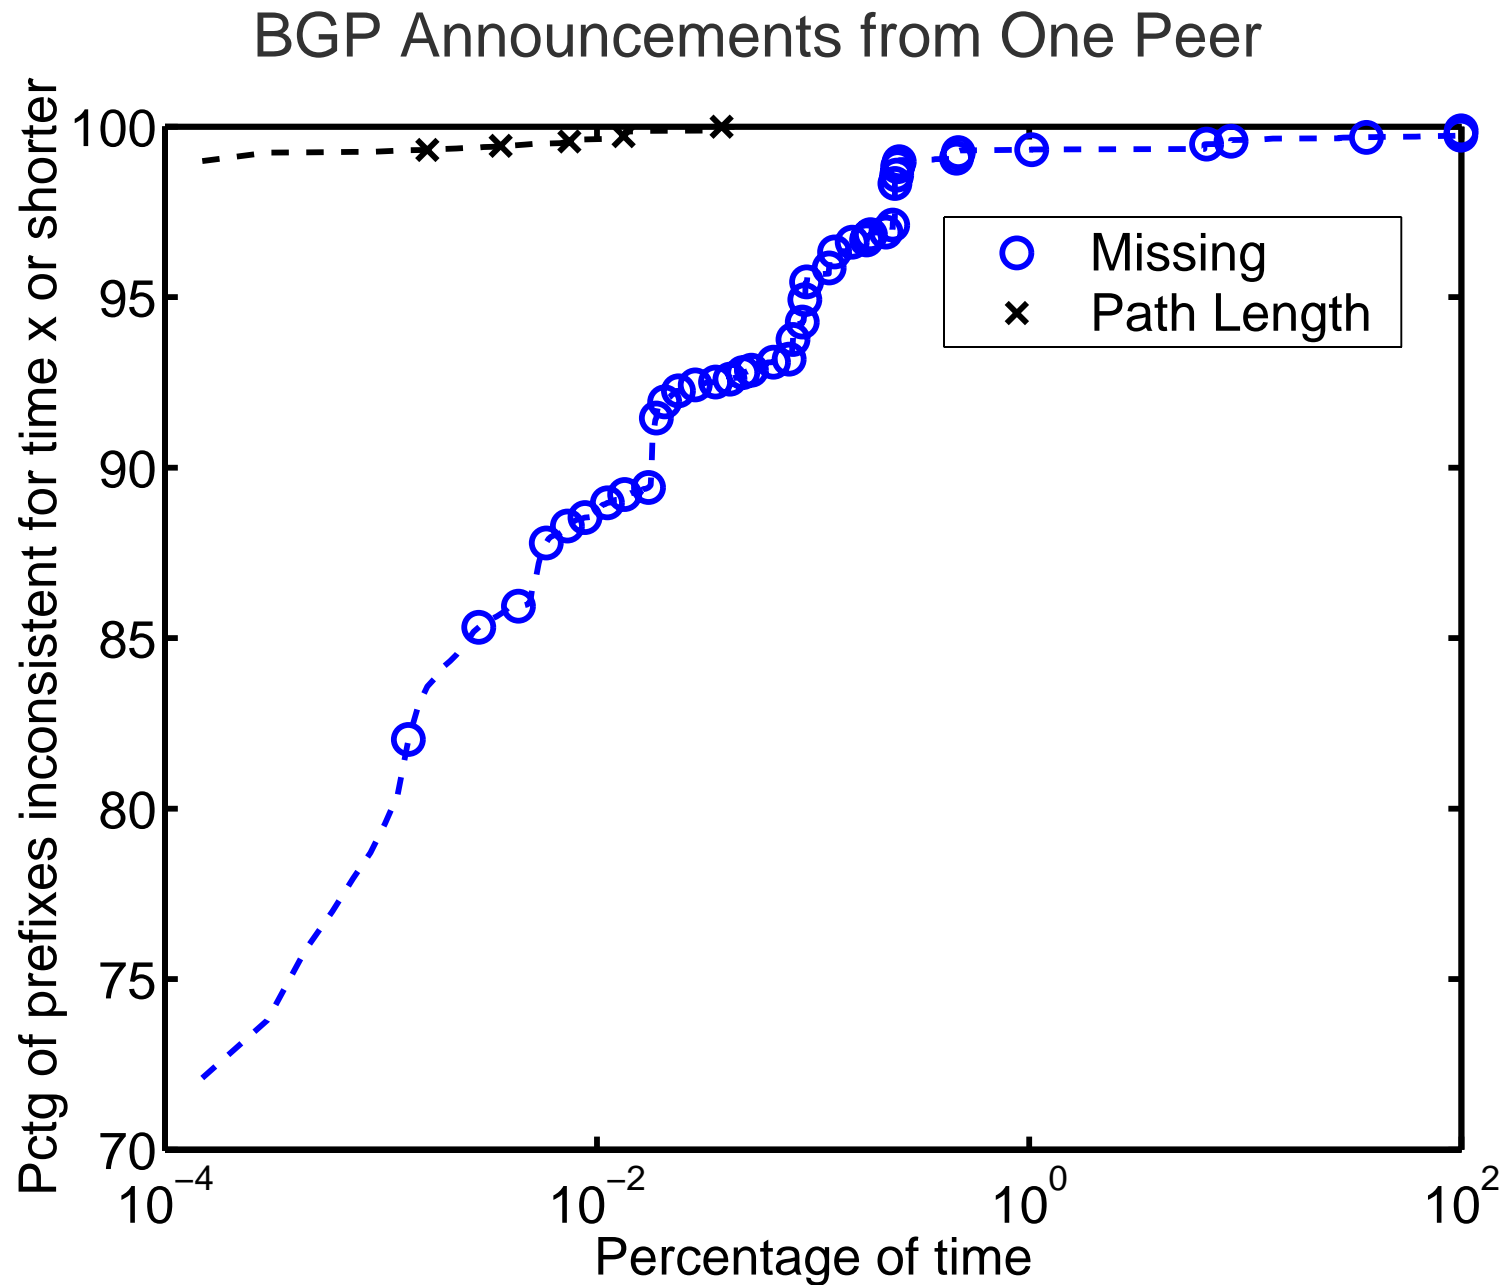


Inconsistent Advertisements from a Customer

Some routers select routes with "no export" community.




Most Prefixes Consistent Most of the Time



Most Events are Transient, Some Persistent

BGP Announcements from One Peer



Some Peers Have Long-Lived Inconsistencies

