Making BGP filtering an habit: Impact on policies

draft-cardona-filtering-threats-00

JuanCamilo.Cardona@imdea.org
Pierre.Francois@imdea.org

History

- IETF GROW talk in 2010
- RIPE talk in 2010
- GIS paper in 2011
- IETF IDR talk last week
- Should become an IETF GROW WG DOC at next IETF

Agenda

- Local filtering can do harm
- Remotely triggered filtering can do harm

- Still it's needed and used
- Let's be aware and conscious about it

Local filtering as an habit?

Ignoring overlapping prefixes?

- People get serious about filtering
- See INIT7 talk at RIPE63
  - Demo’ing bill reduction through filtering
  - Filter out prefixes at transit to get through peers via a covering prefix
- Requests to vendors for automated filtering features

Filtering of overlapping prefixes...

- *“They make me forward to my transit instead of my peer/customer”,*
- *“I’m loosing money due to their games”*
- *“They violate my policy”*
- It is frustrating to forward traffic with which you could get more ROI, indeed.

Why does it take place?

- What are the reasons for an ISP or a CDN to receive more specific prefixes from providers only, while there is a covering prefix at a peer ?

Reference context 1

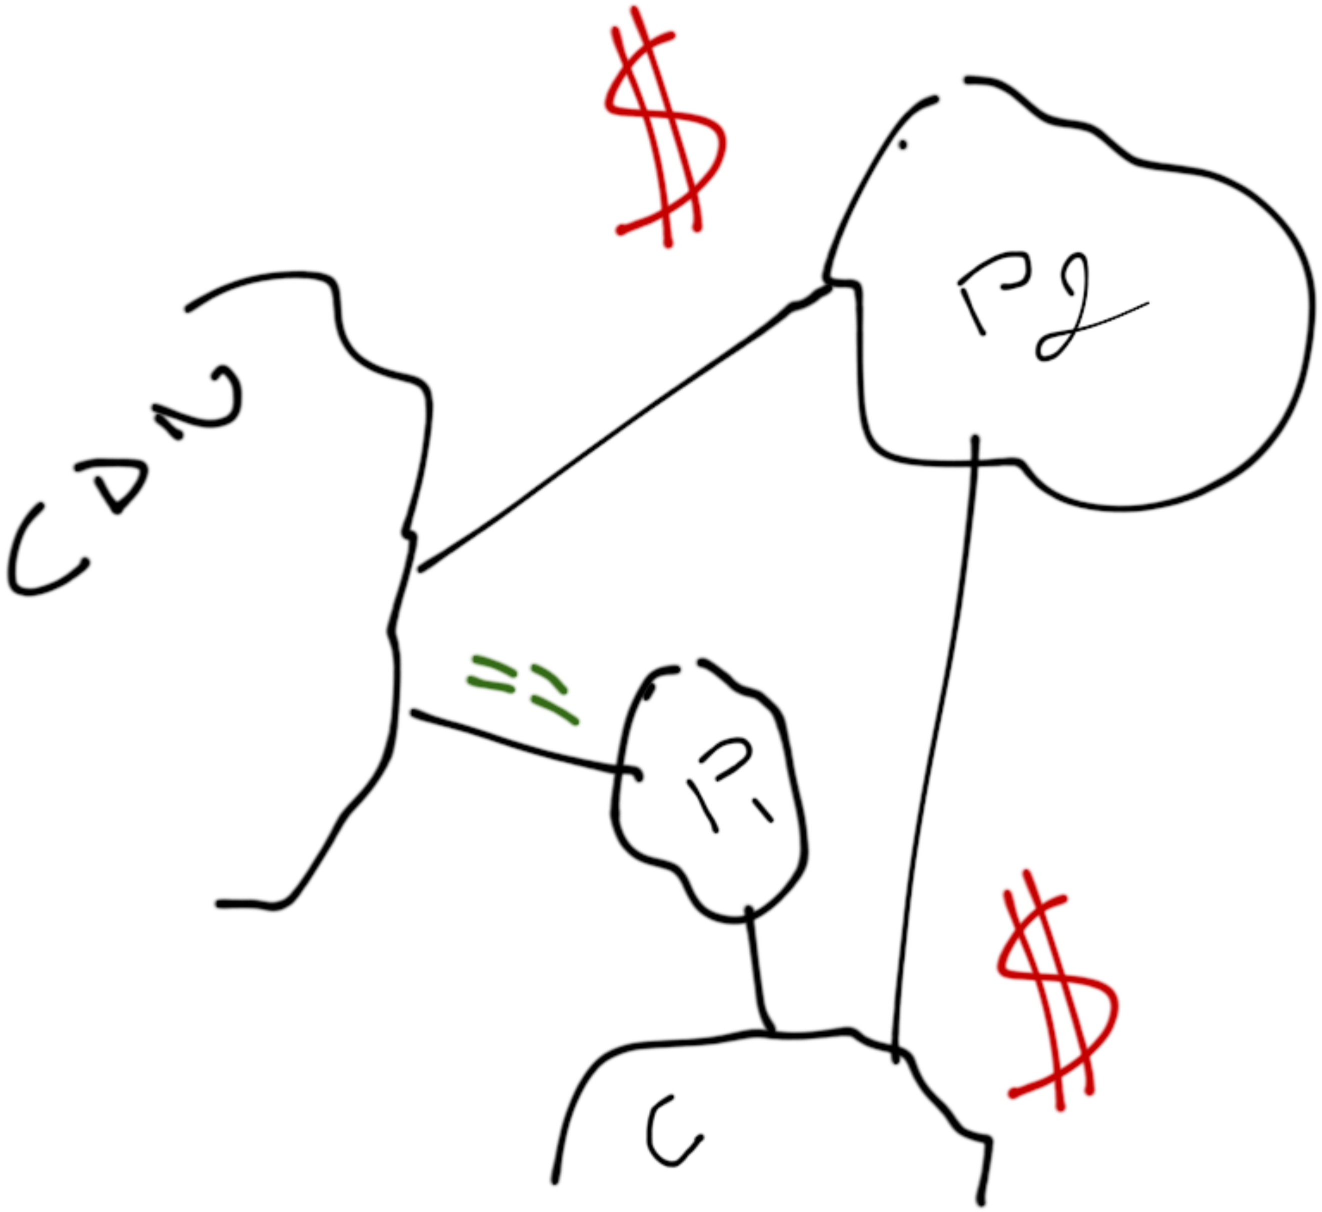

- Destination Eyeball ISP C
- C in customer base of Peer P1
- C in customer base of Provider P2

Case 1
No export

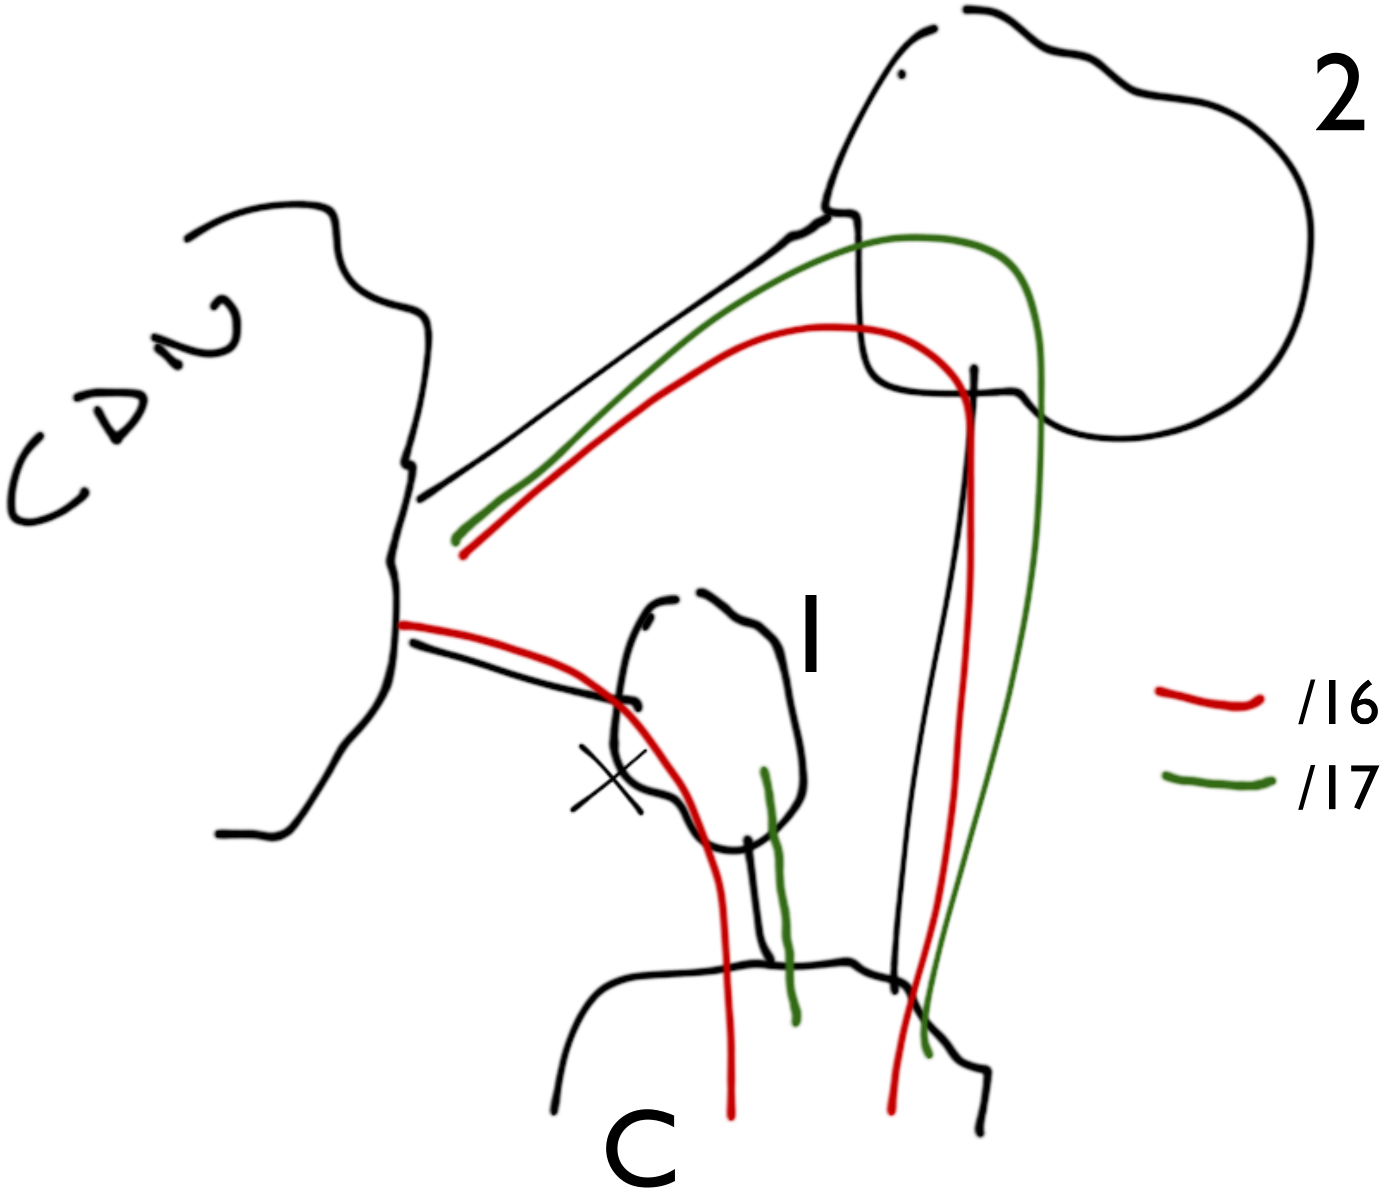

- C tags NO_EXPORT when advertising the more specific to peer P1
    - C does not want the entire incoming traffic shares for the /17 to be delivered by P1
    - C gives traffic shares to P1 only for the single homed customers of P1. **C Expects to receive the rest from P2**
- Can you bypass the TE needs of C?

Case II
Selective advertisement

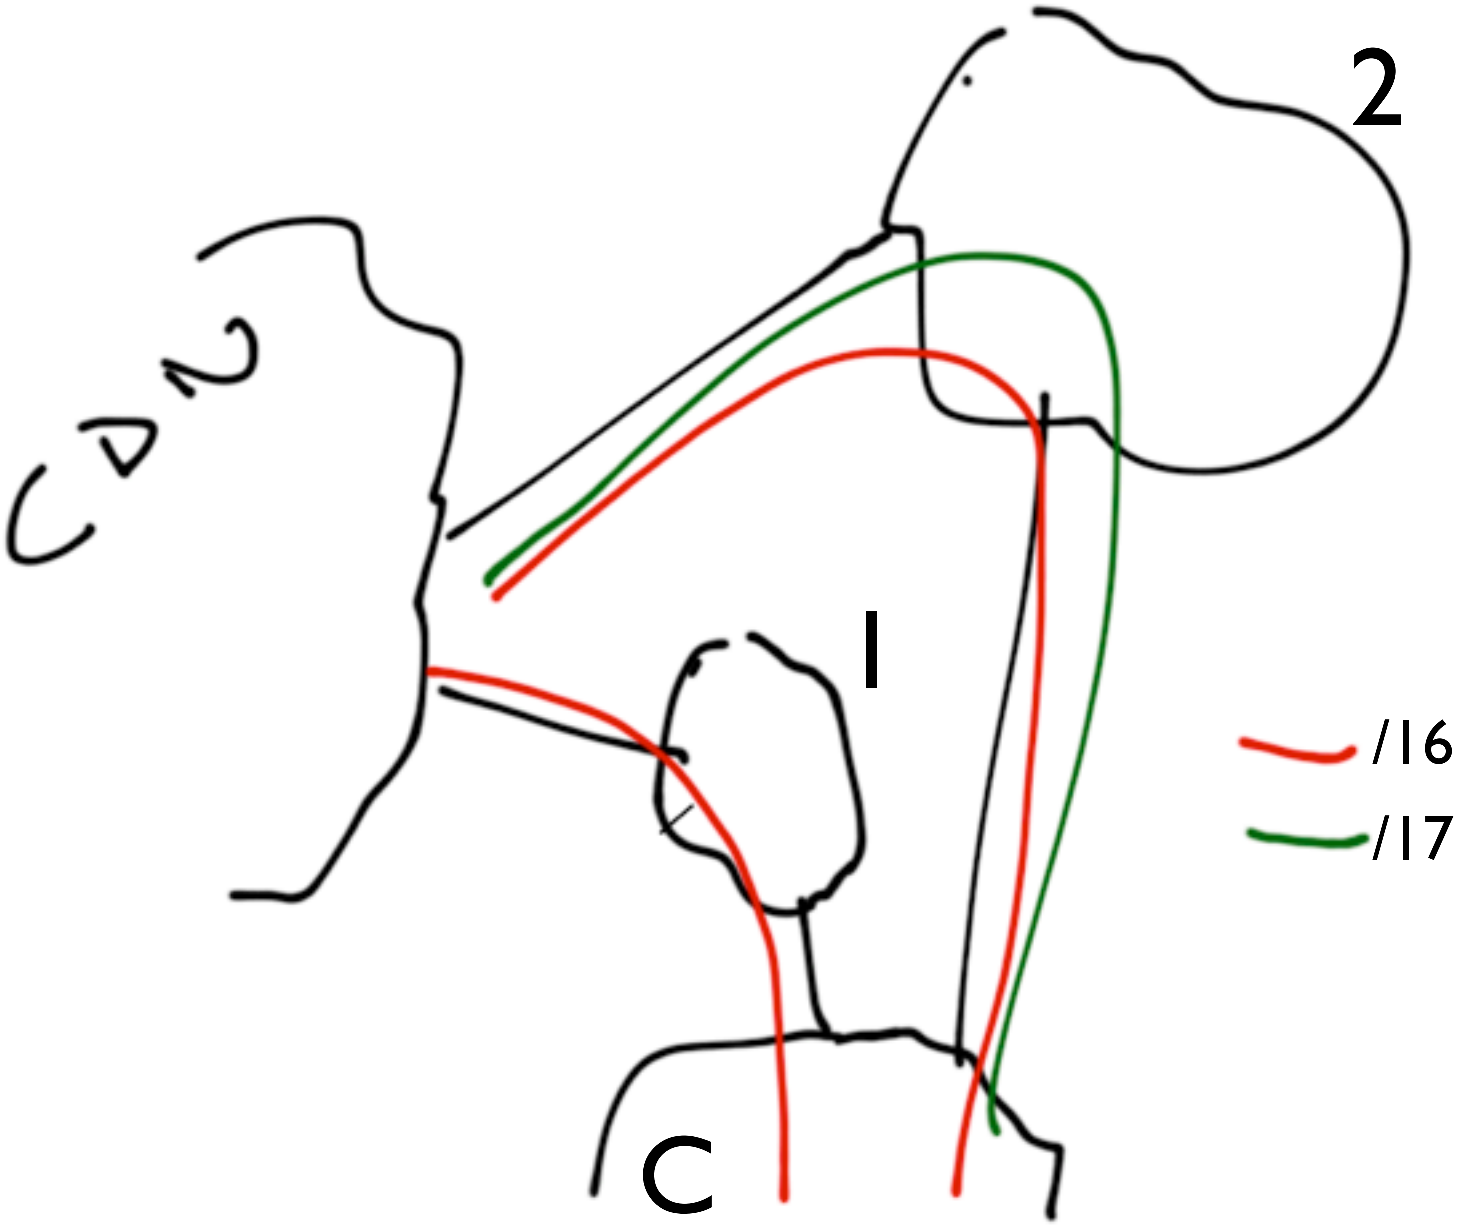

- C does not advertise the /17 to P1
    - C does not want to allow the incoming traffic shares for the /17 to be delivered by P1
    - P1 is only allowed to deliver its own customer traffic to C
- Can you bypass the TE needs of C?

Impact of bypassing more specifics

- Disrespect of your peers' customers traffic engineering requirements/needs
- Up to now, this is a business discussion on who should decide about Internet end-to-end paths...
- The games being played doing so can turn bad for some ISPs

BGP : control plane

- Policy-constrained path selection in BGP...
  Flexible
  Per-prefix granularity

- *"A BGP-router's* ***route processor*** *will pick a path towards a given destination* ***prefix*** *by applying the following rules"*

  Weight
  Local-pref
  As Path Length
  IGP/Med
  ...

Data plane result of BGP

- ... dominated in the data-plane
- A **FIB** will pick a path towards a given destination **address** by applying the following rules

  Longest prefix match to get the prefix

  (
  Best path towards that prefix was picked based on
  Weight
  Local-pref
  As Path Length
  IGP/Med
  ...)

Policy violation at a peer

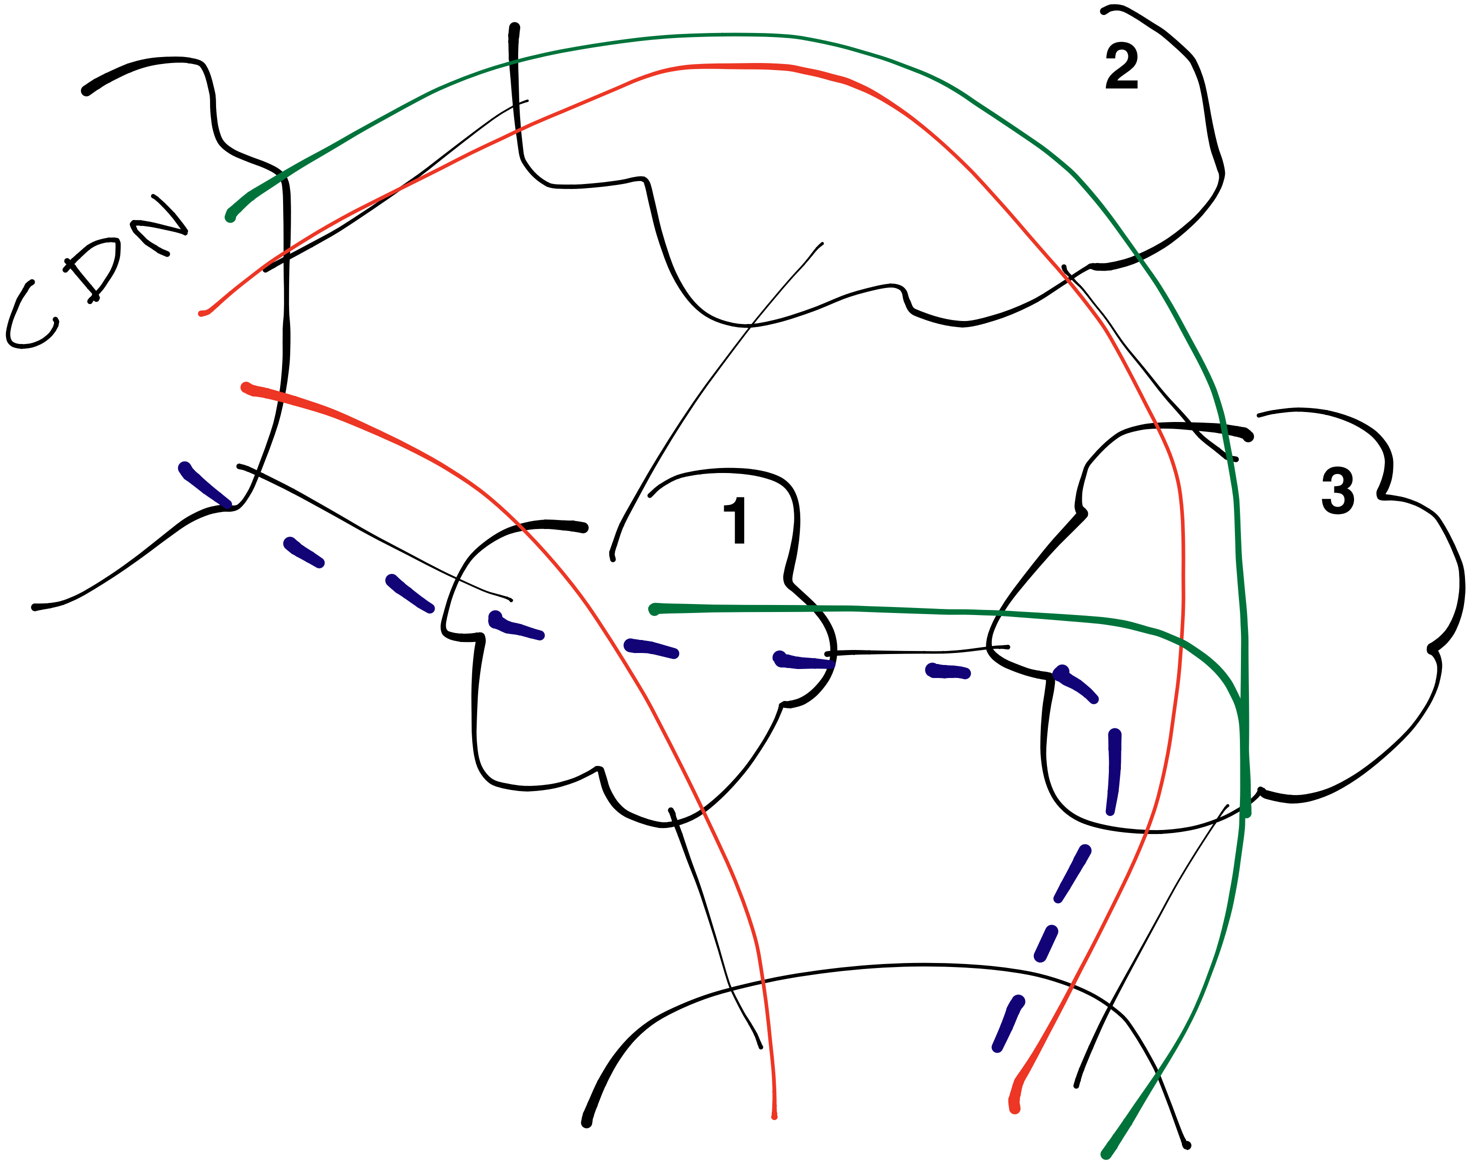

- P3 and P1 are peers
- CDN peers with P1
- C does not advertise the /17 to P1, Only to P3
- If you ignore the transit path, you violate P1's policy doing CDN-P1-P3

Marketing

- You act against your neighborhood
- What is the cost of a public announcement
  “*These CDN guys are the ones making the money, they force me to peer instead of paying me, and now, on top of that, they make moves to get free transit through my network ???*”

Take away

- Ignoring more specifics can do you good
  - vs. your peers, customers, and customers of your peers
  - With a risk of policy violation at your peers
  - Undistinguishable cases without gathering external data
- **Should not be done automatically with simplistic rules**
- **Peering and Customer contracts should accommodate those cases**

Remotely triggered filtering

Remote triggered filtering

- Triggering the same mess from far...
- Example:
  Route propagation control offered by Sprint
  - Have to be a customer of Sprint
  - 65000:XXX : Do not advertise to ASXXX
    can be AOL, NTT, BT, Level3, GBLX, Verizon, AT&T, ...

Powerful complementary means to limit path knowledge towards yourself

- Selective advertisement, performed locally
- Selective propagation, triggered remotely

Control-plane/Data-plane can mismatch

- Paths for **overlapping** prefixes are controlled independently
  - By yourself
  - By your BGP neighborhood
- Forwarding plane dominated by the longest prefix match rule

- What if policies differ for those overlapping prefixes ?

Toy case study

🟥 A BGP advertisement for NLRI P/p

🟩 A BGP advertisement of a prefix more specific than P/p, say P/p+1

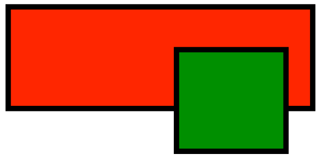

The BGP policy violation trick

- Play with 🟩 and communities
- Make 🟩 reach only a subset of the ASes
  - Some ASes forward 🟩 according to 🟥
  - Until packet reaches an AS knowing 🟩
  - Resulting data-plane not necessarily fitting everyone’s policy...

What can you do with these communities ?

- Turn “don’t advertise to X” values into a only “advertise to Y”
  Just put them all but Y

- or explicit “only advertise to Y” community

Initial routing status

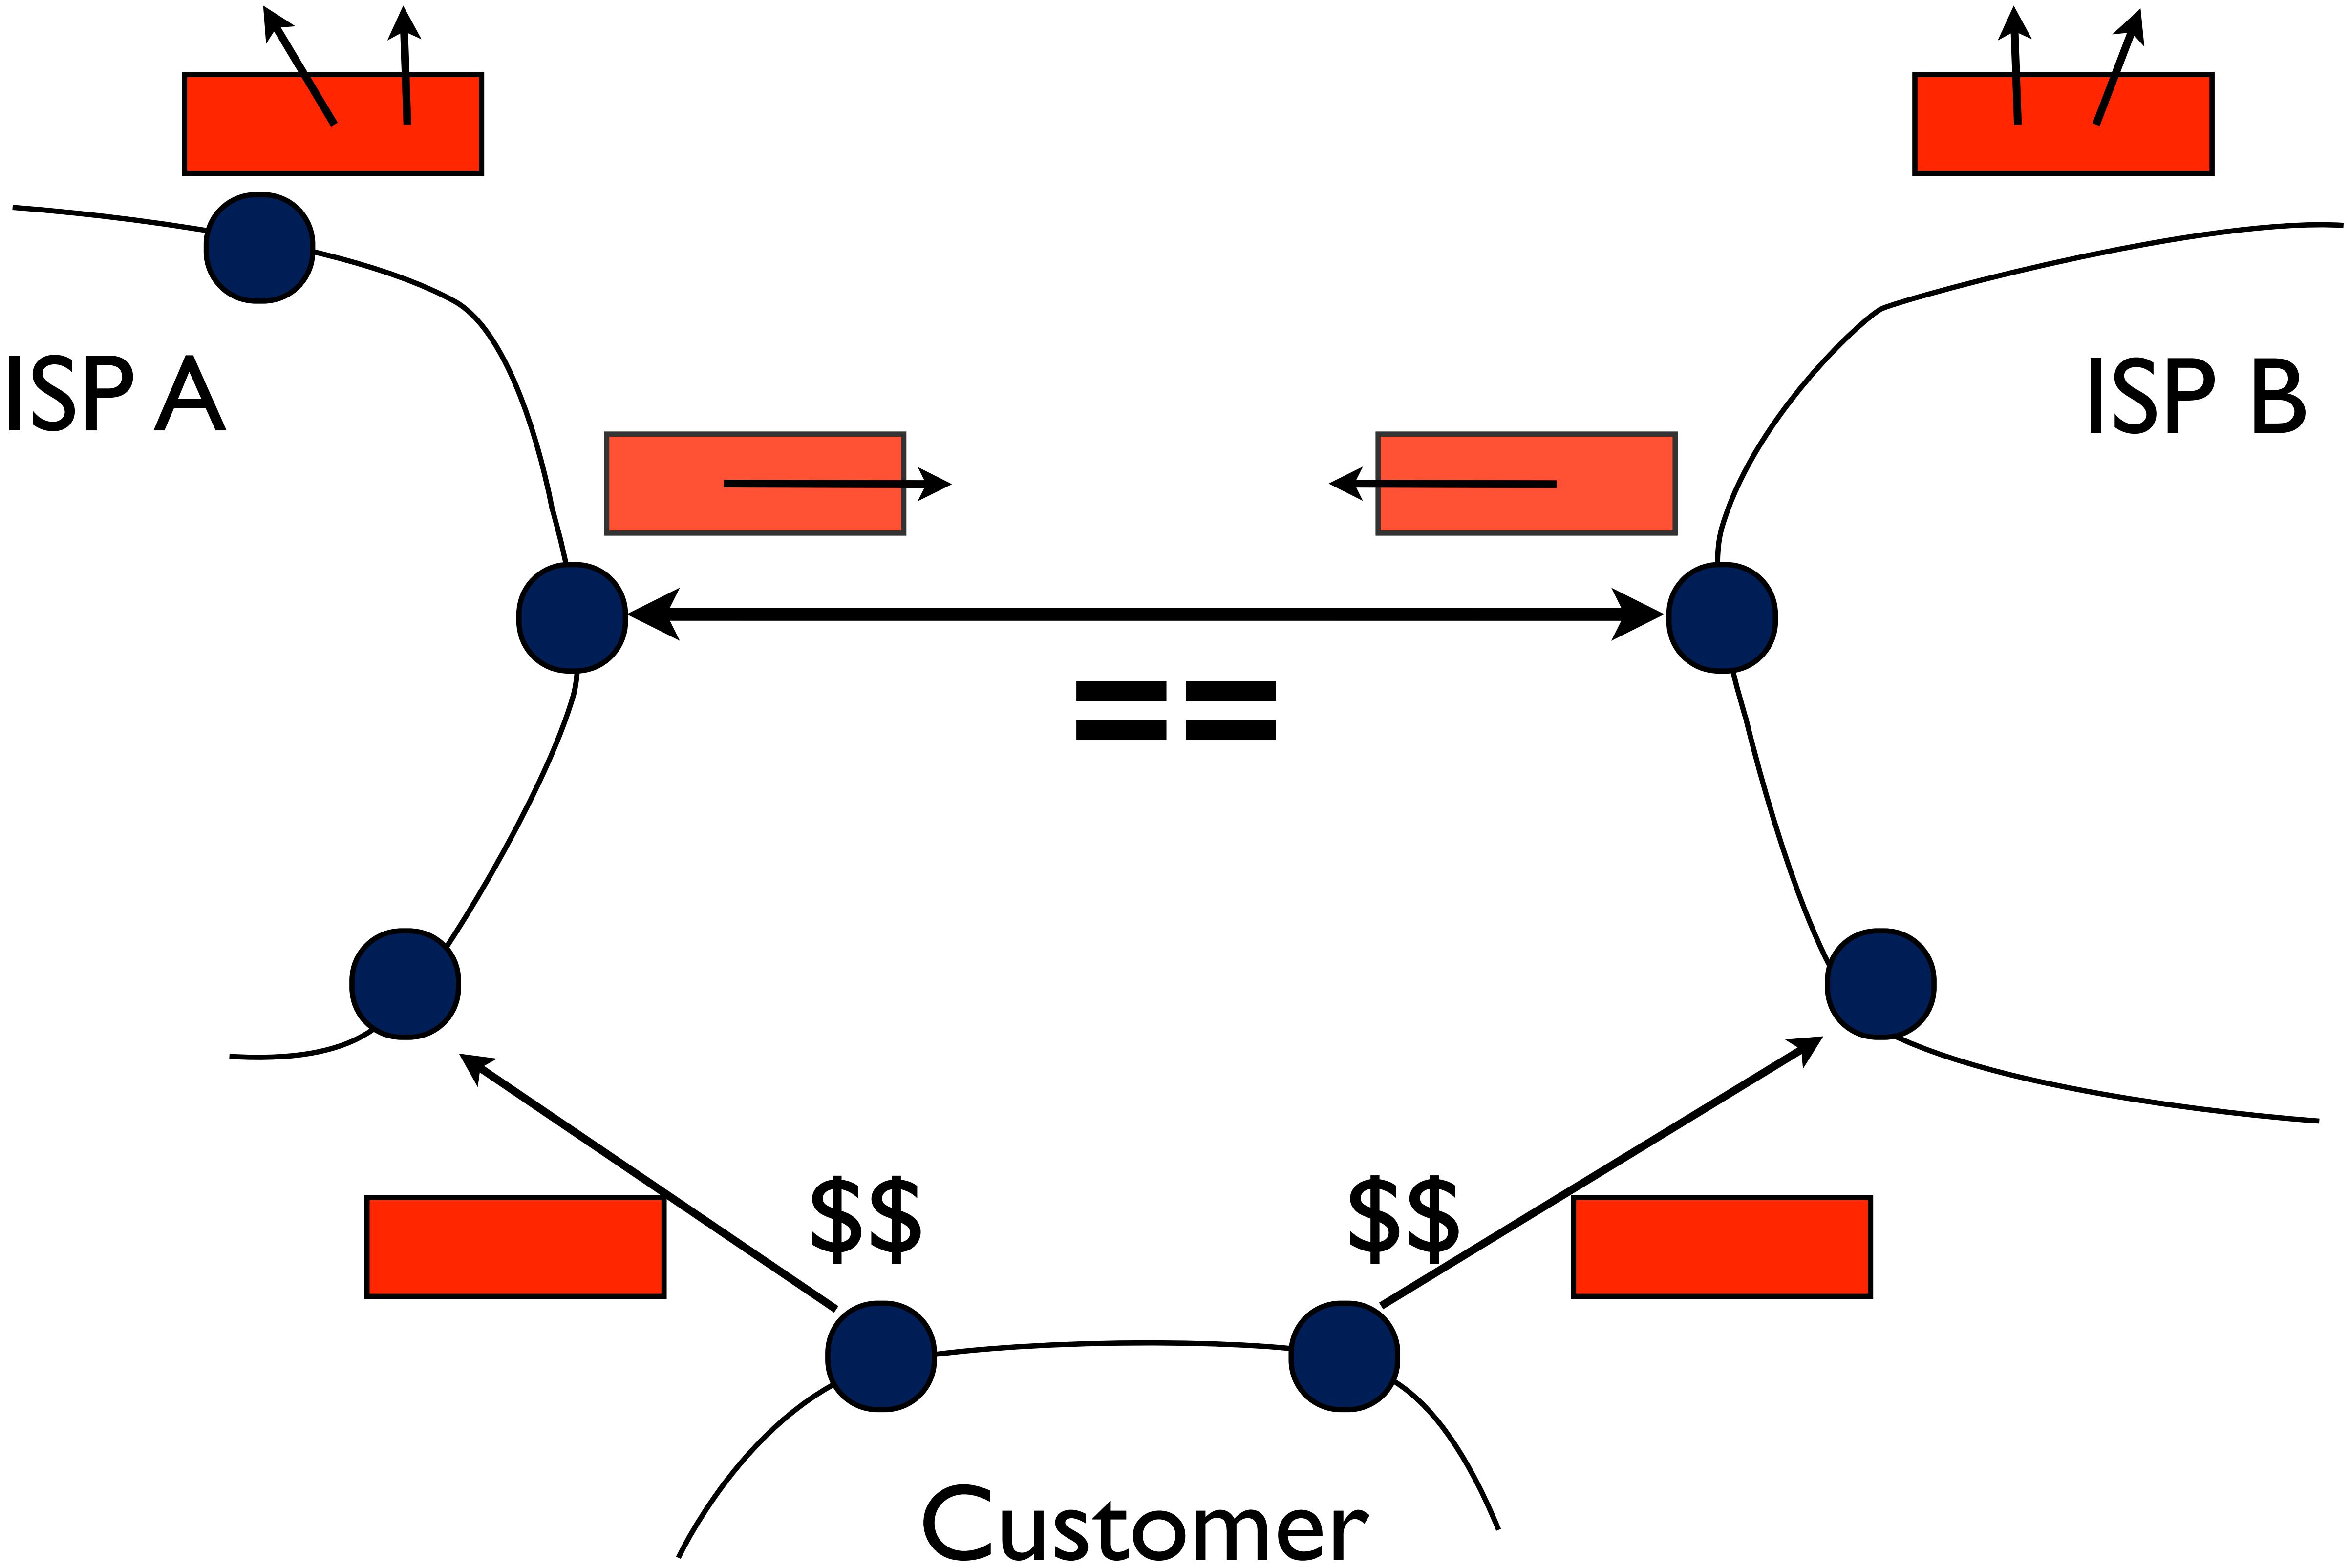

Initial routing status

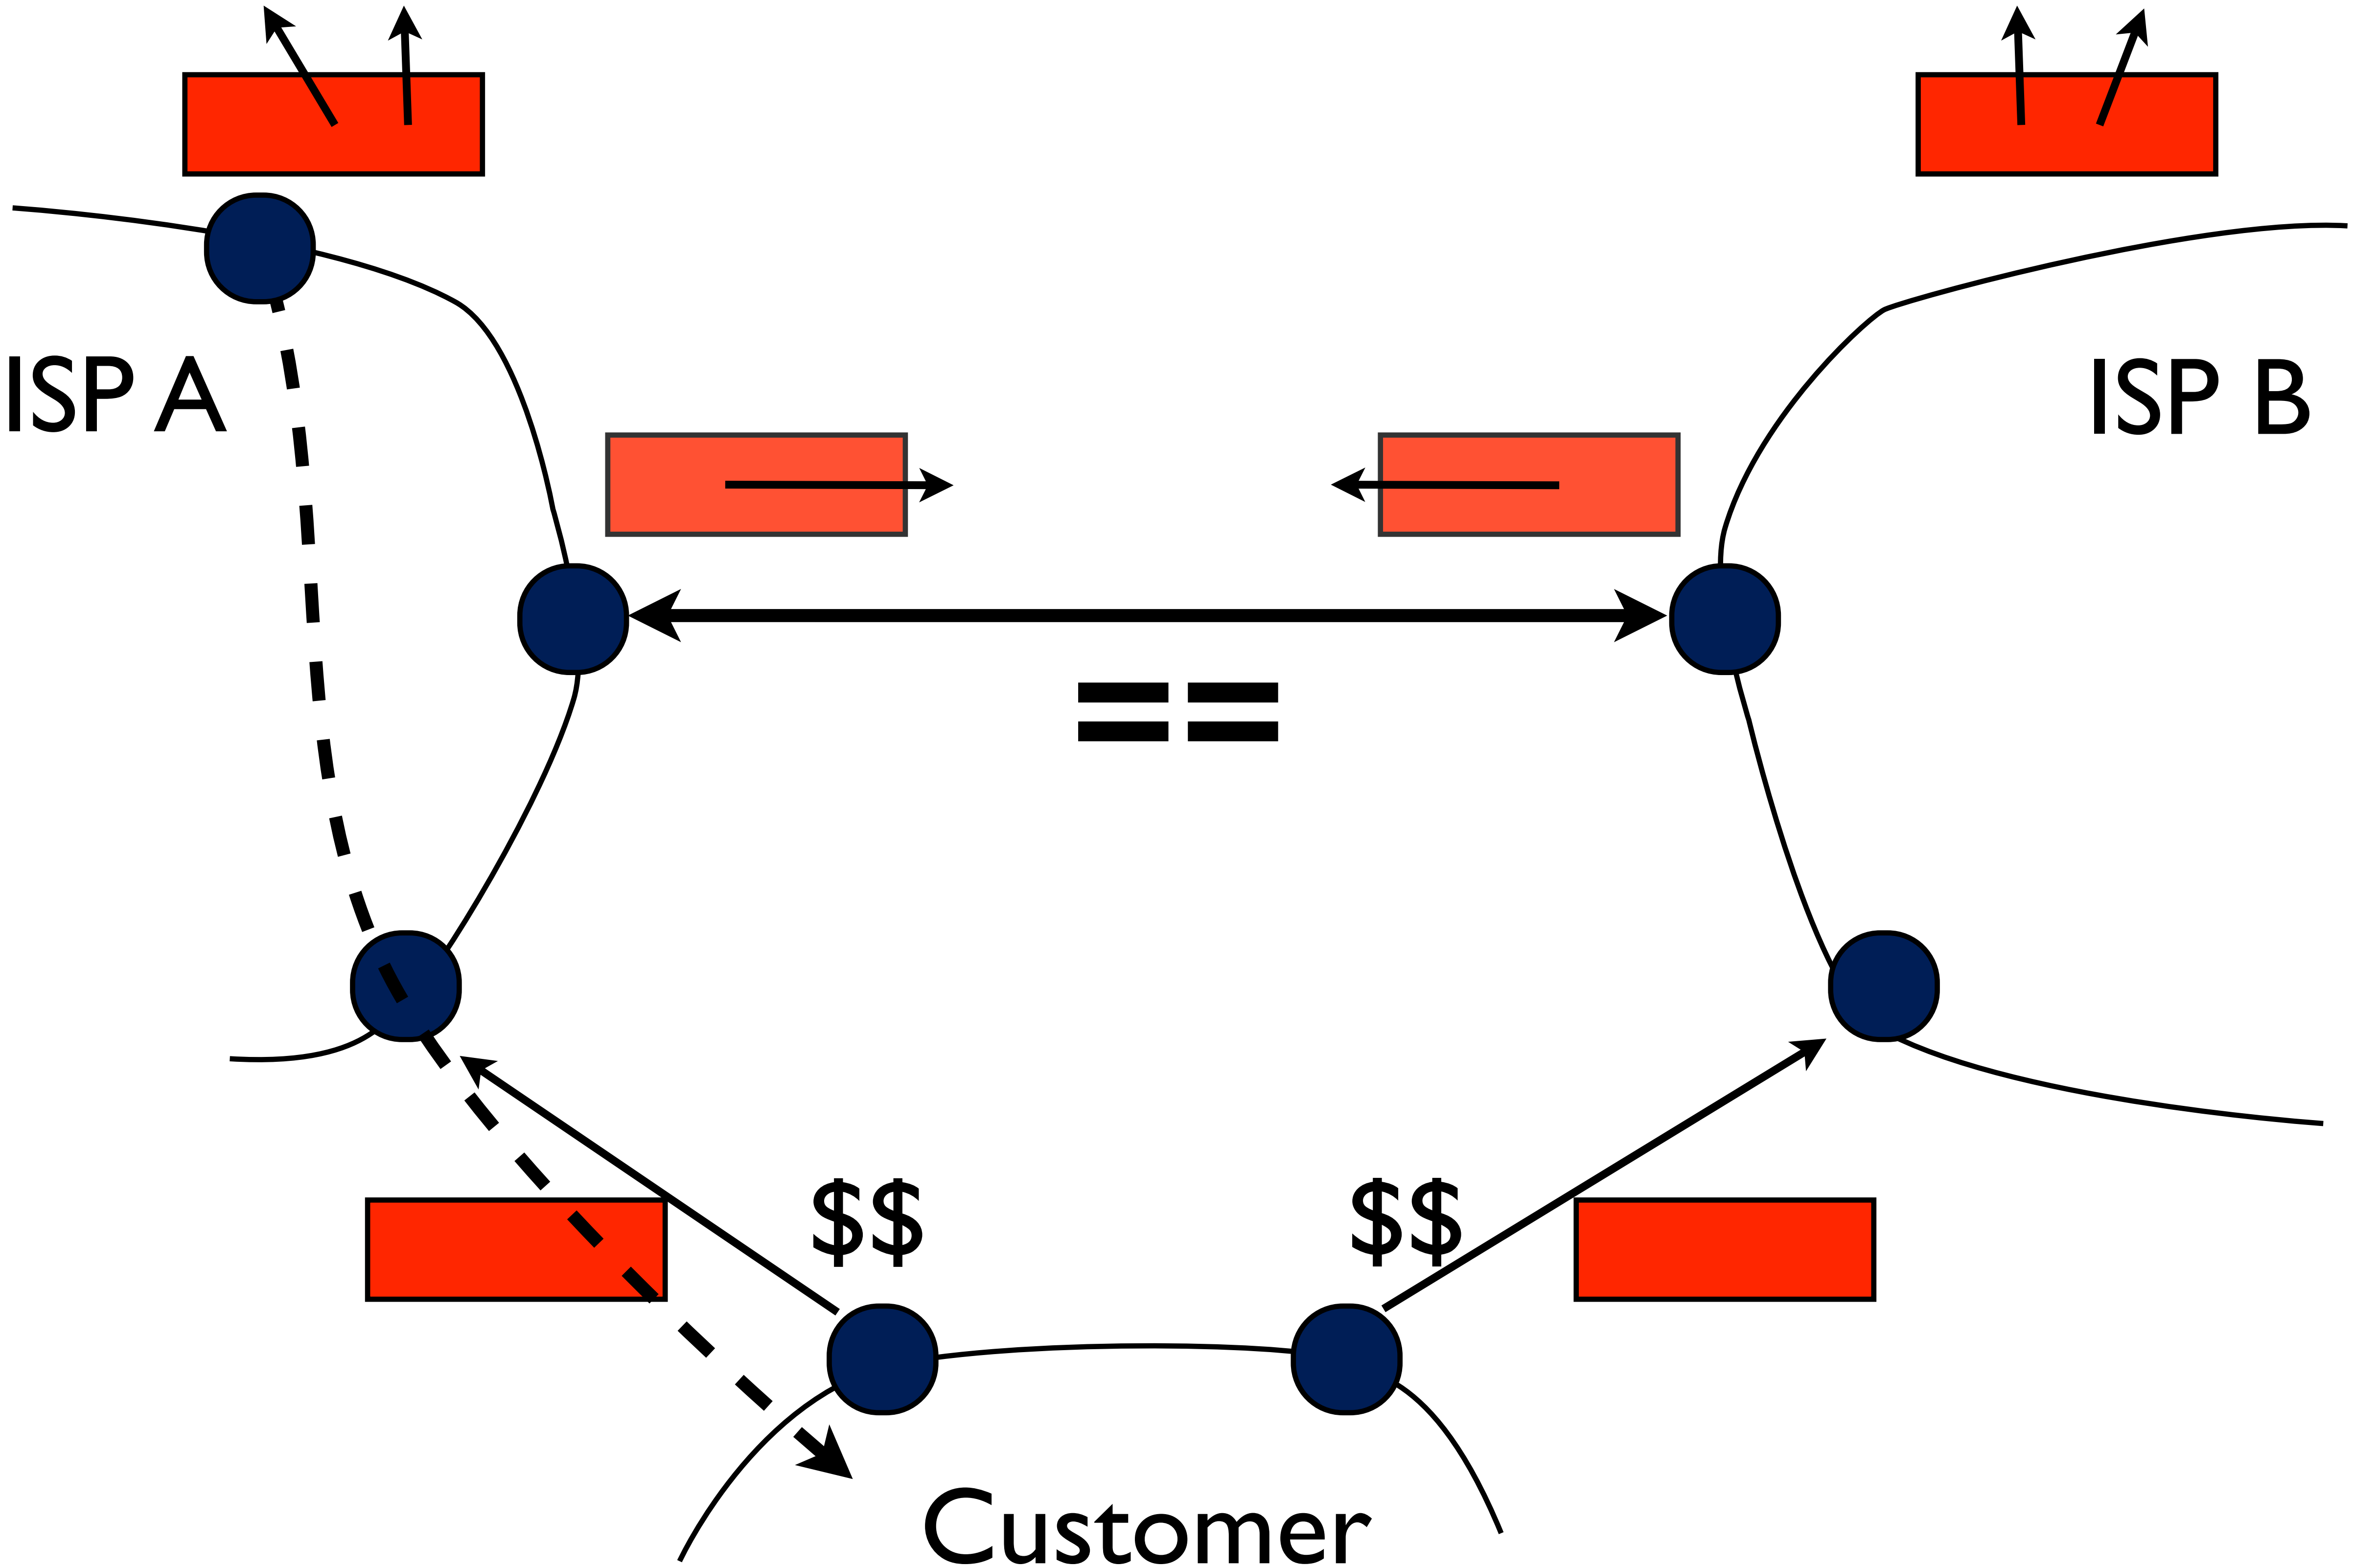

Initial routing status

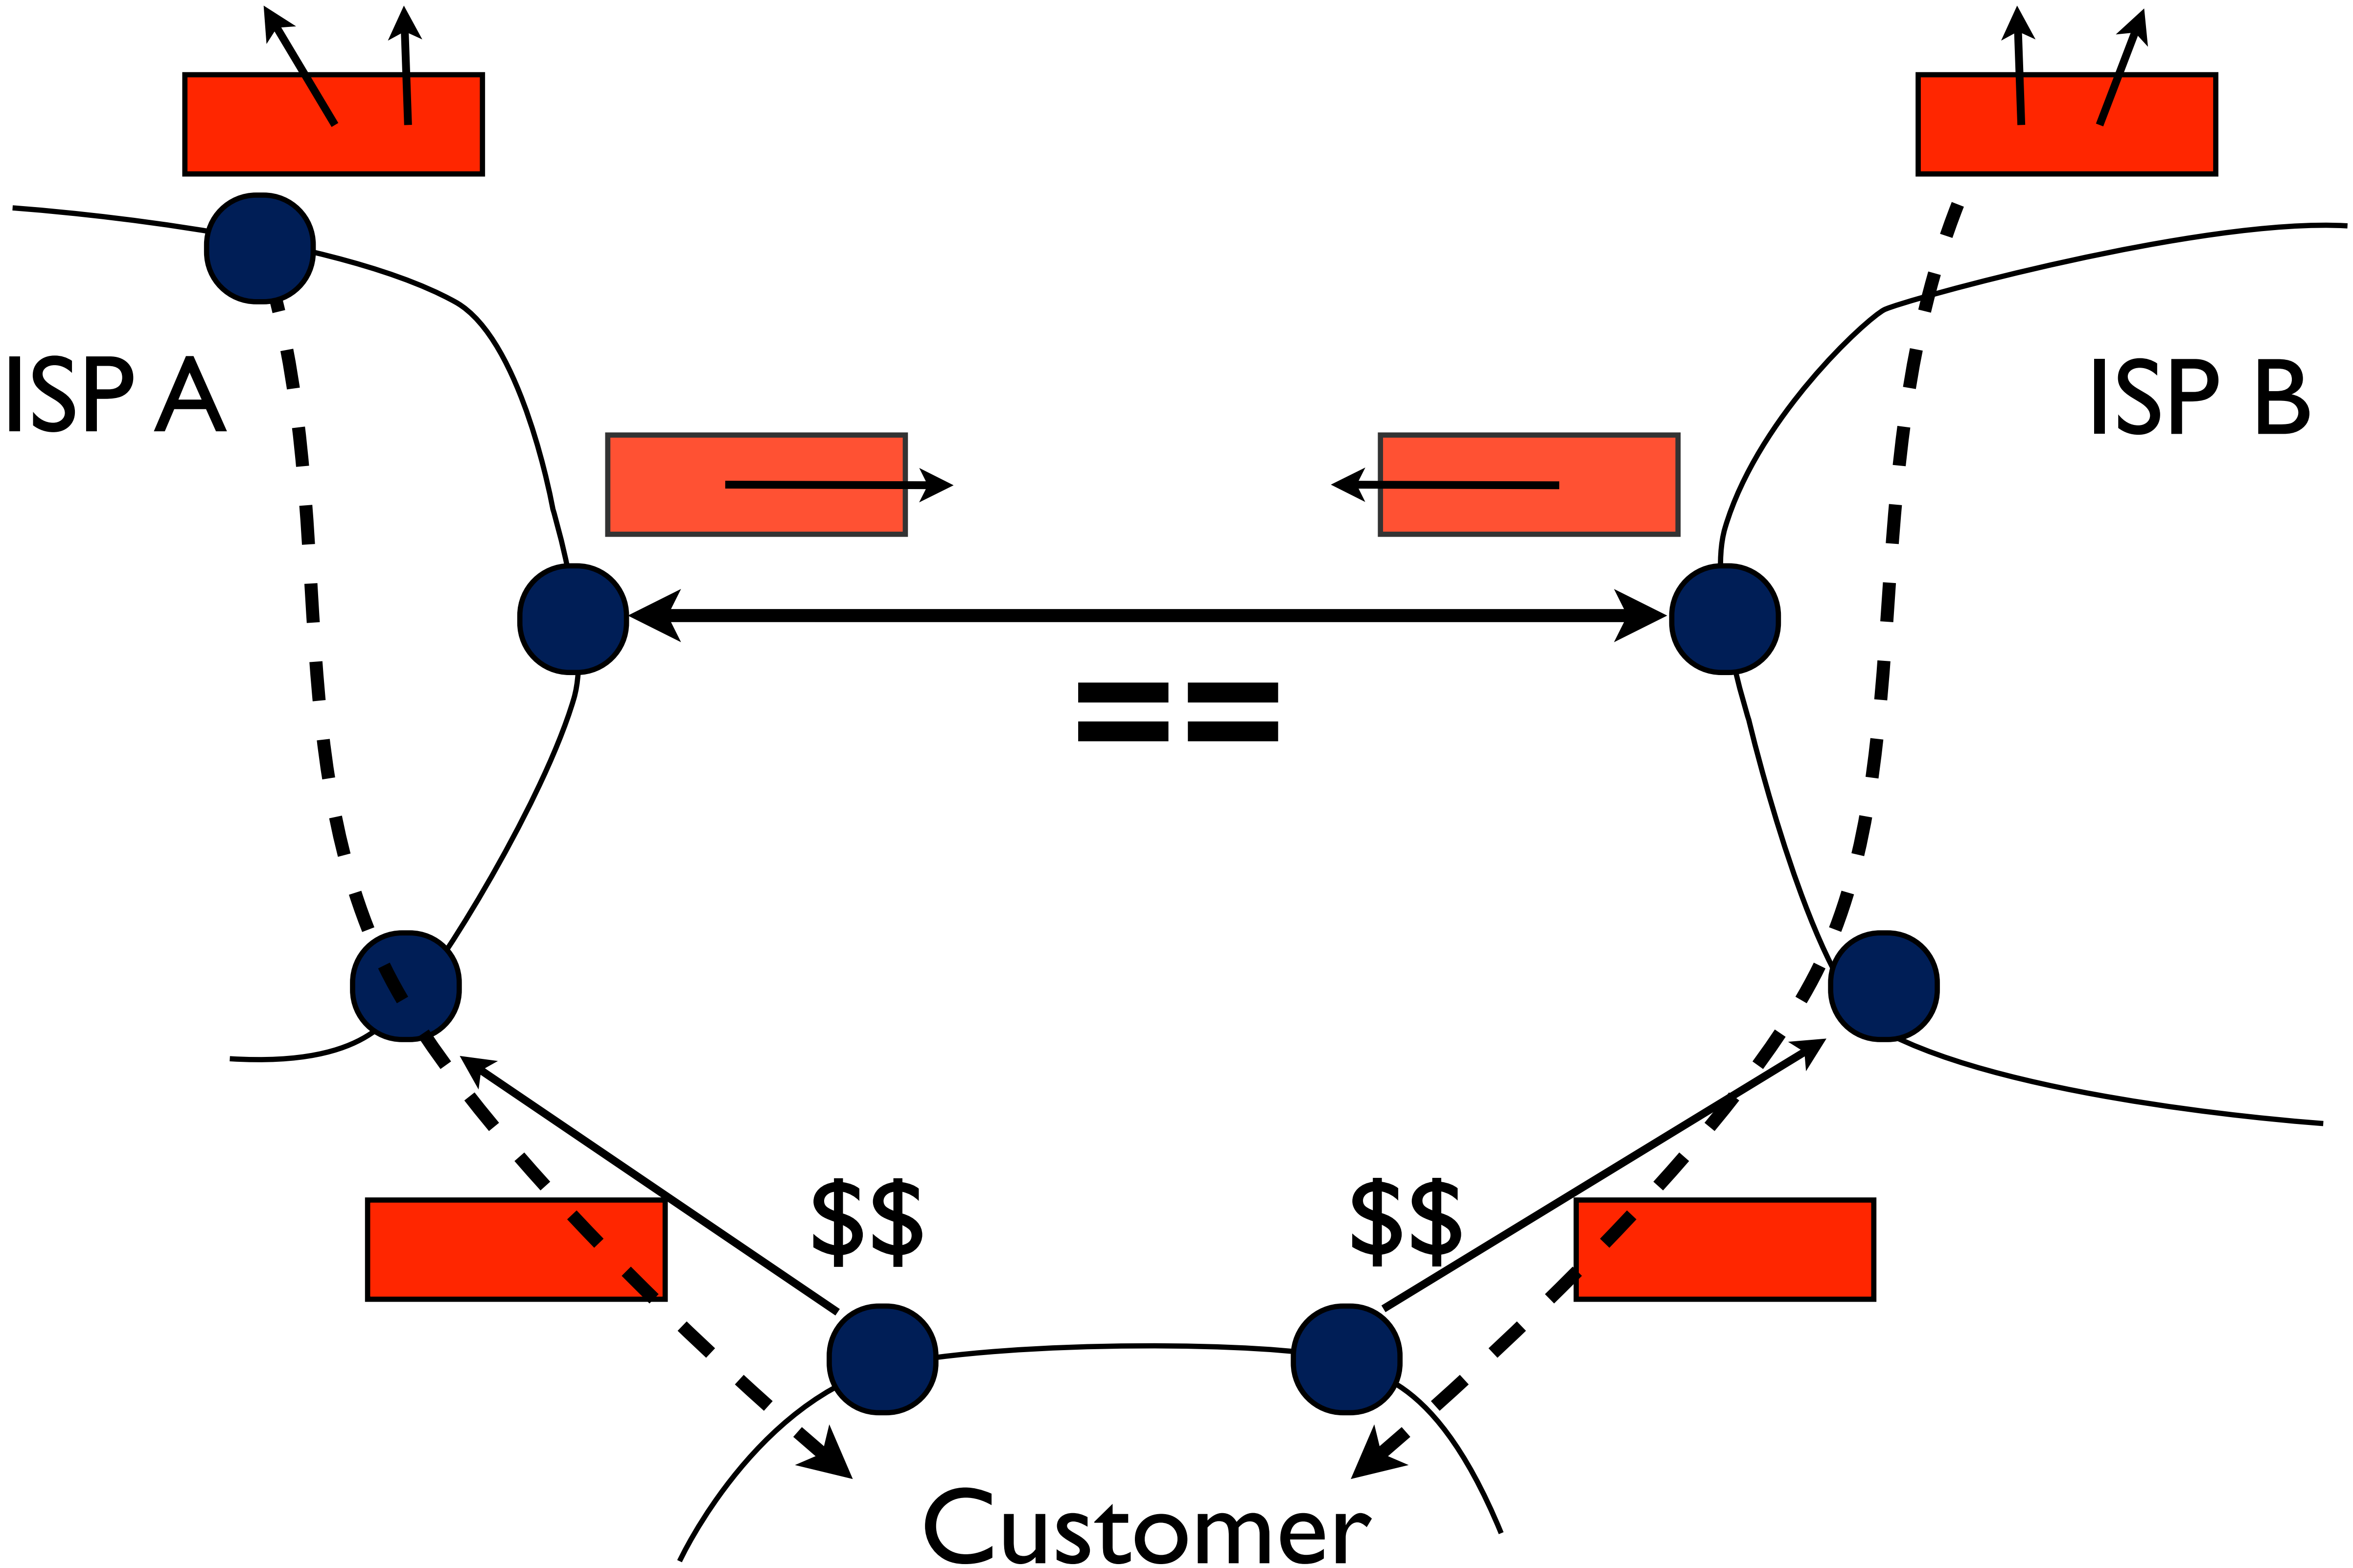

Inbound TE, selective advertisement of a more specific prefix

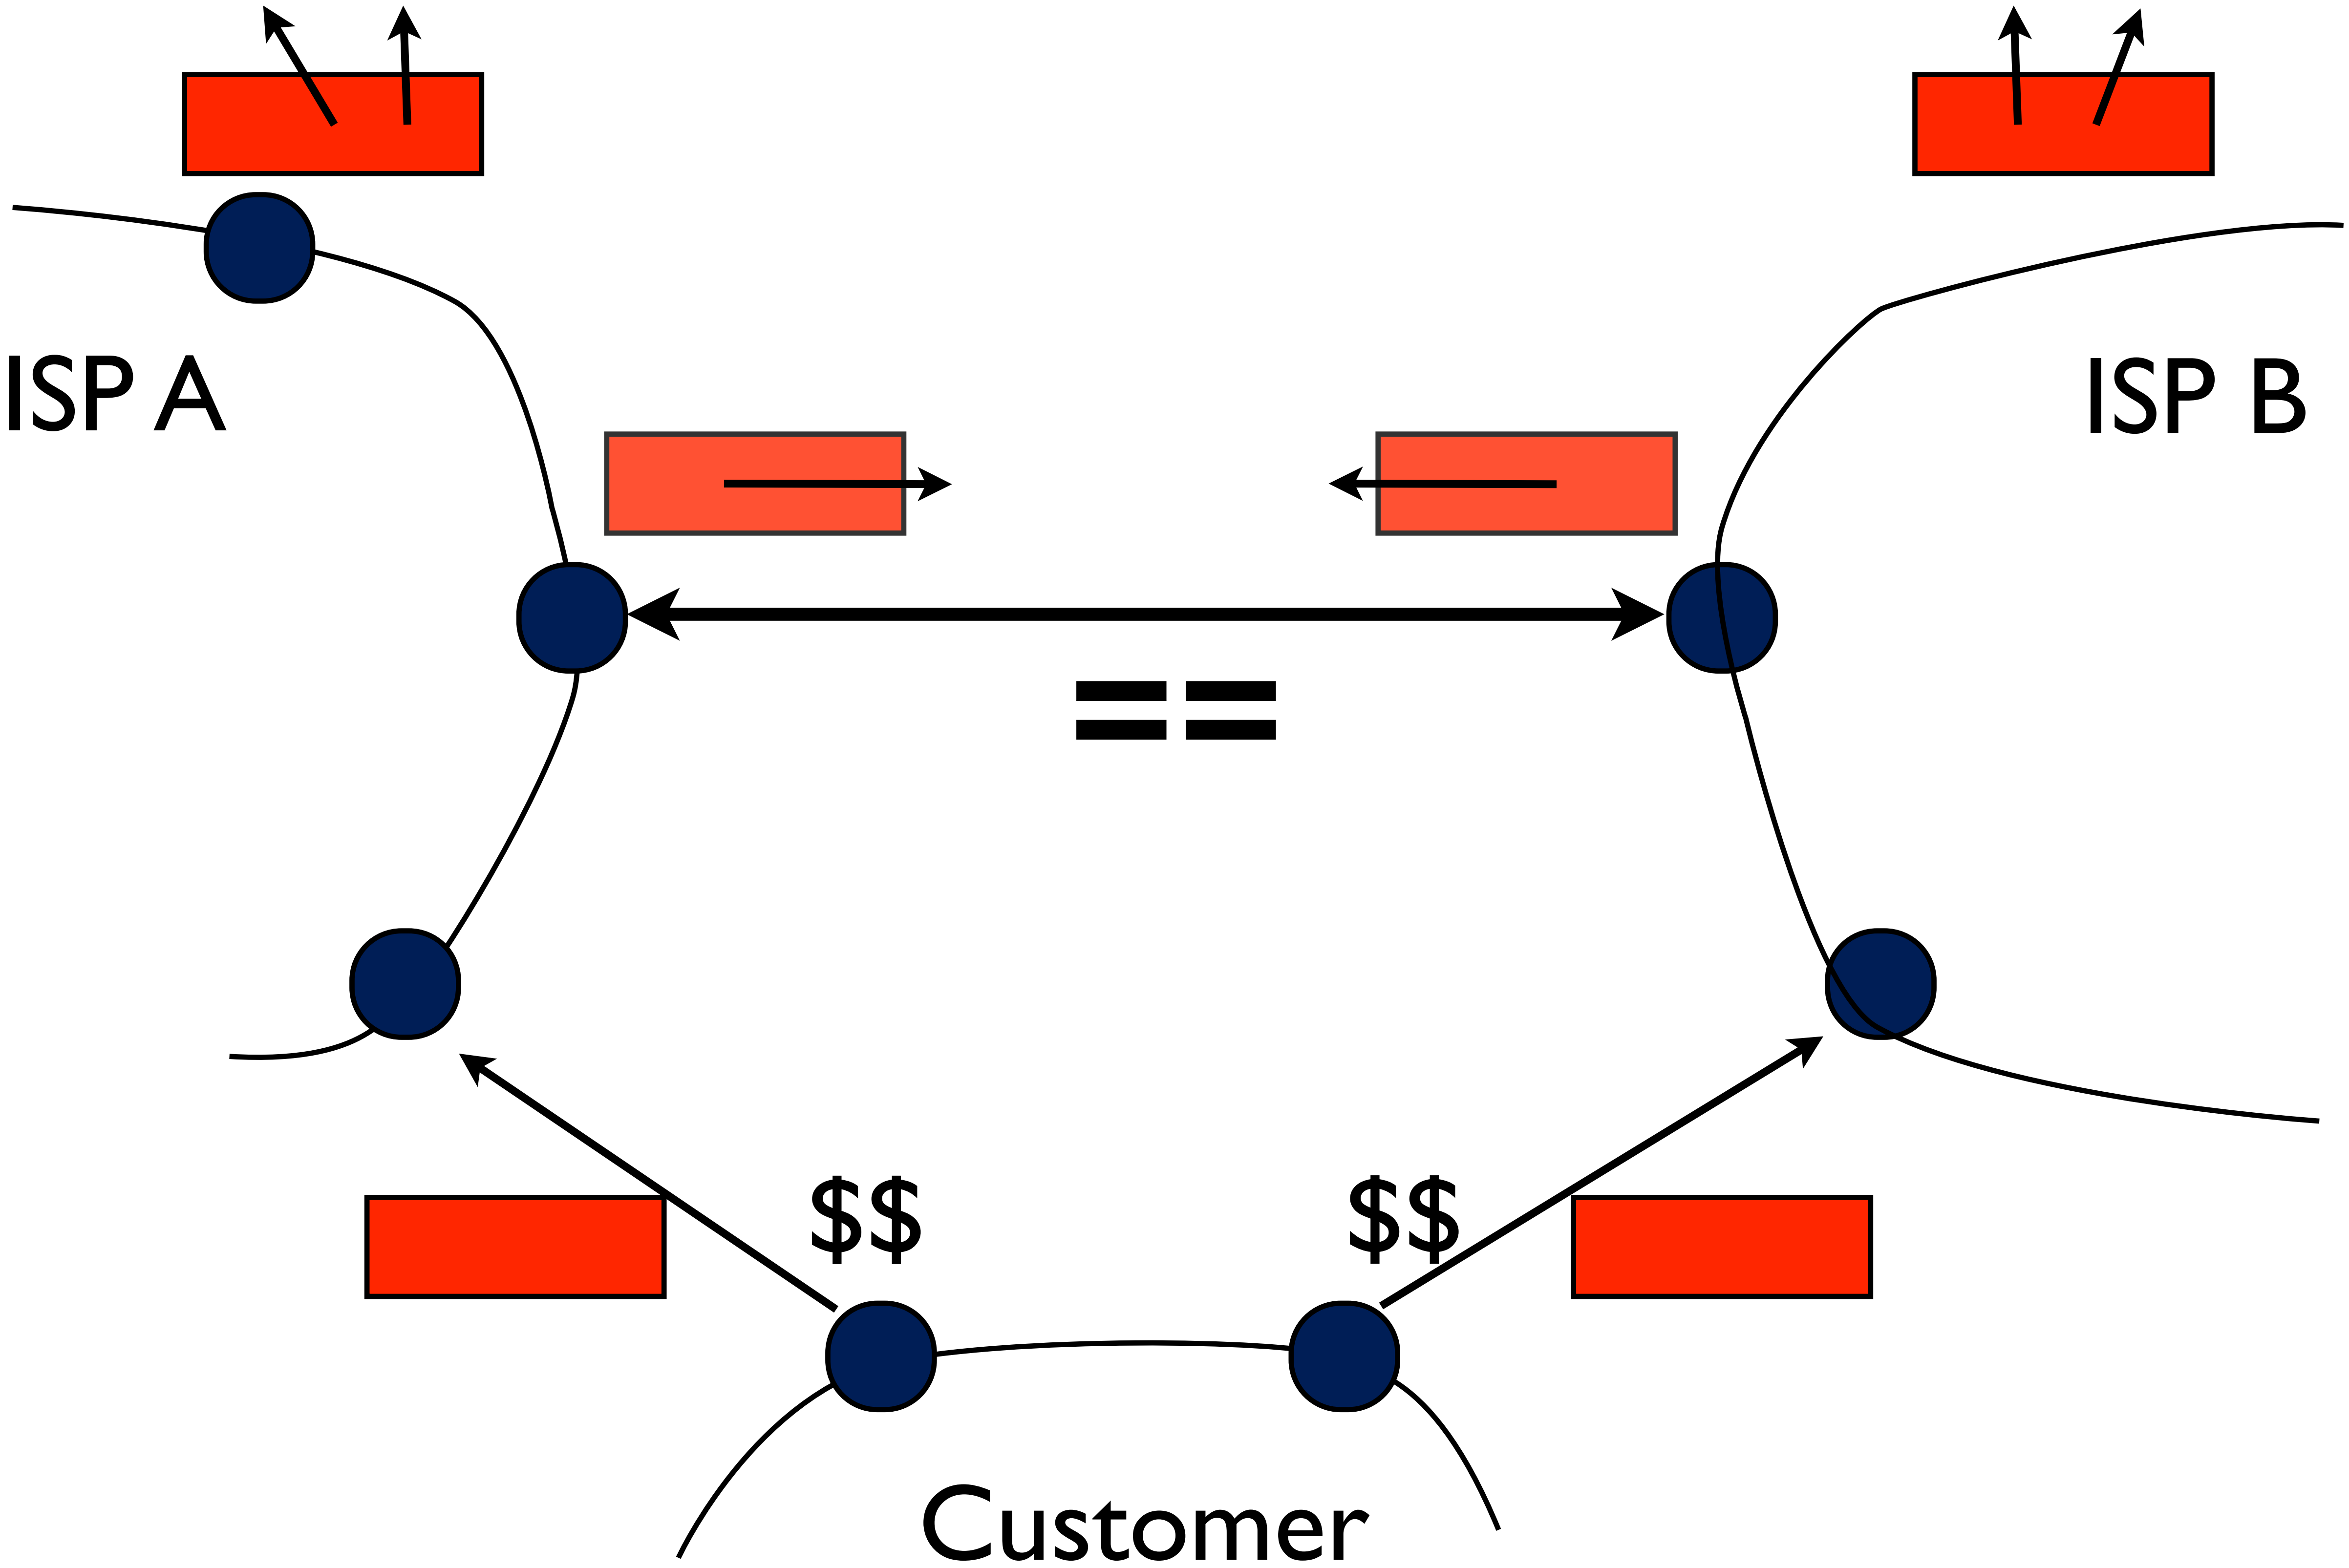

Inbound TE, selective advertisement of a more specific prefix

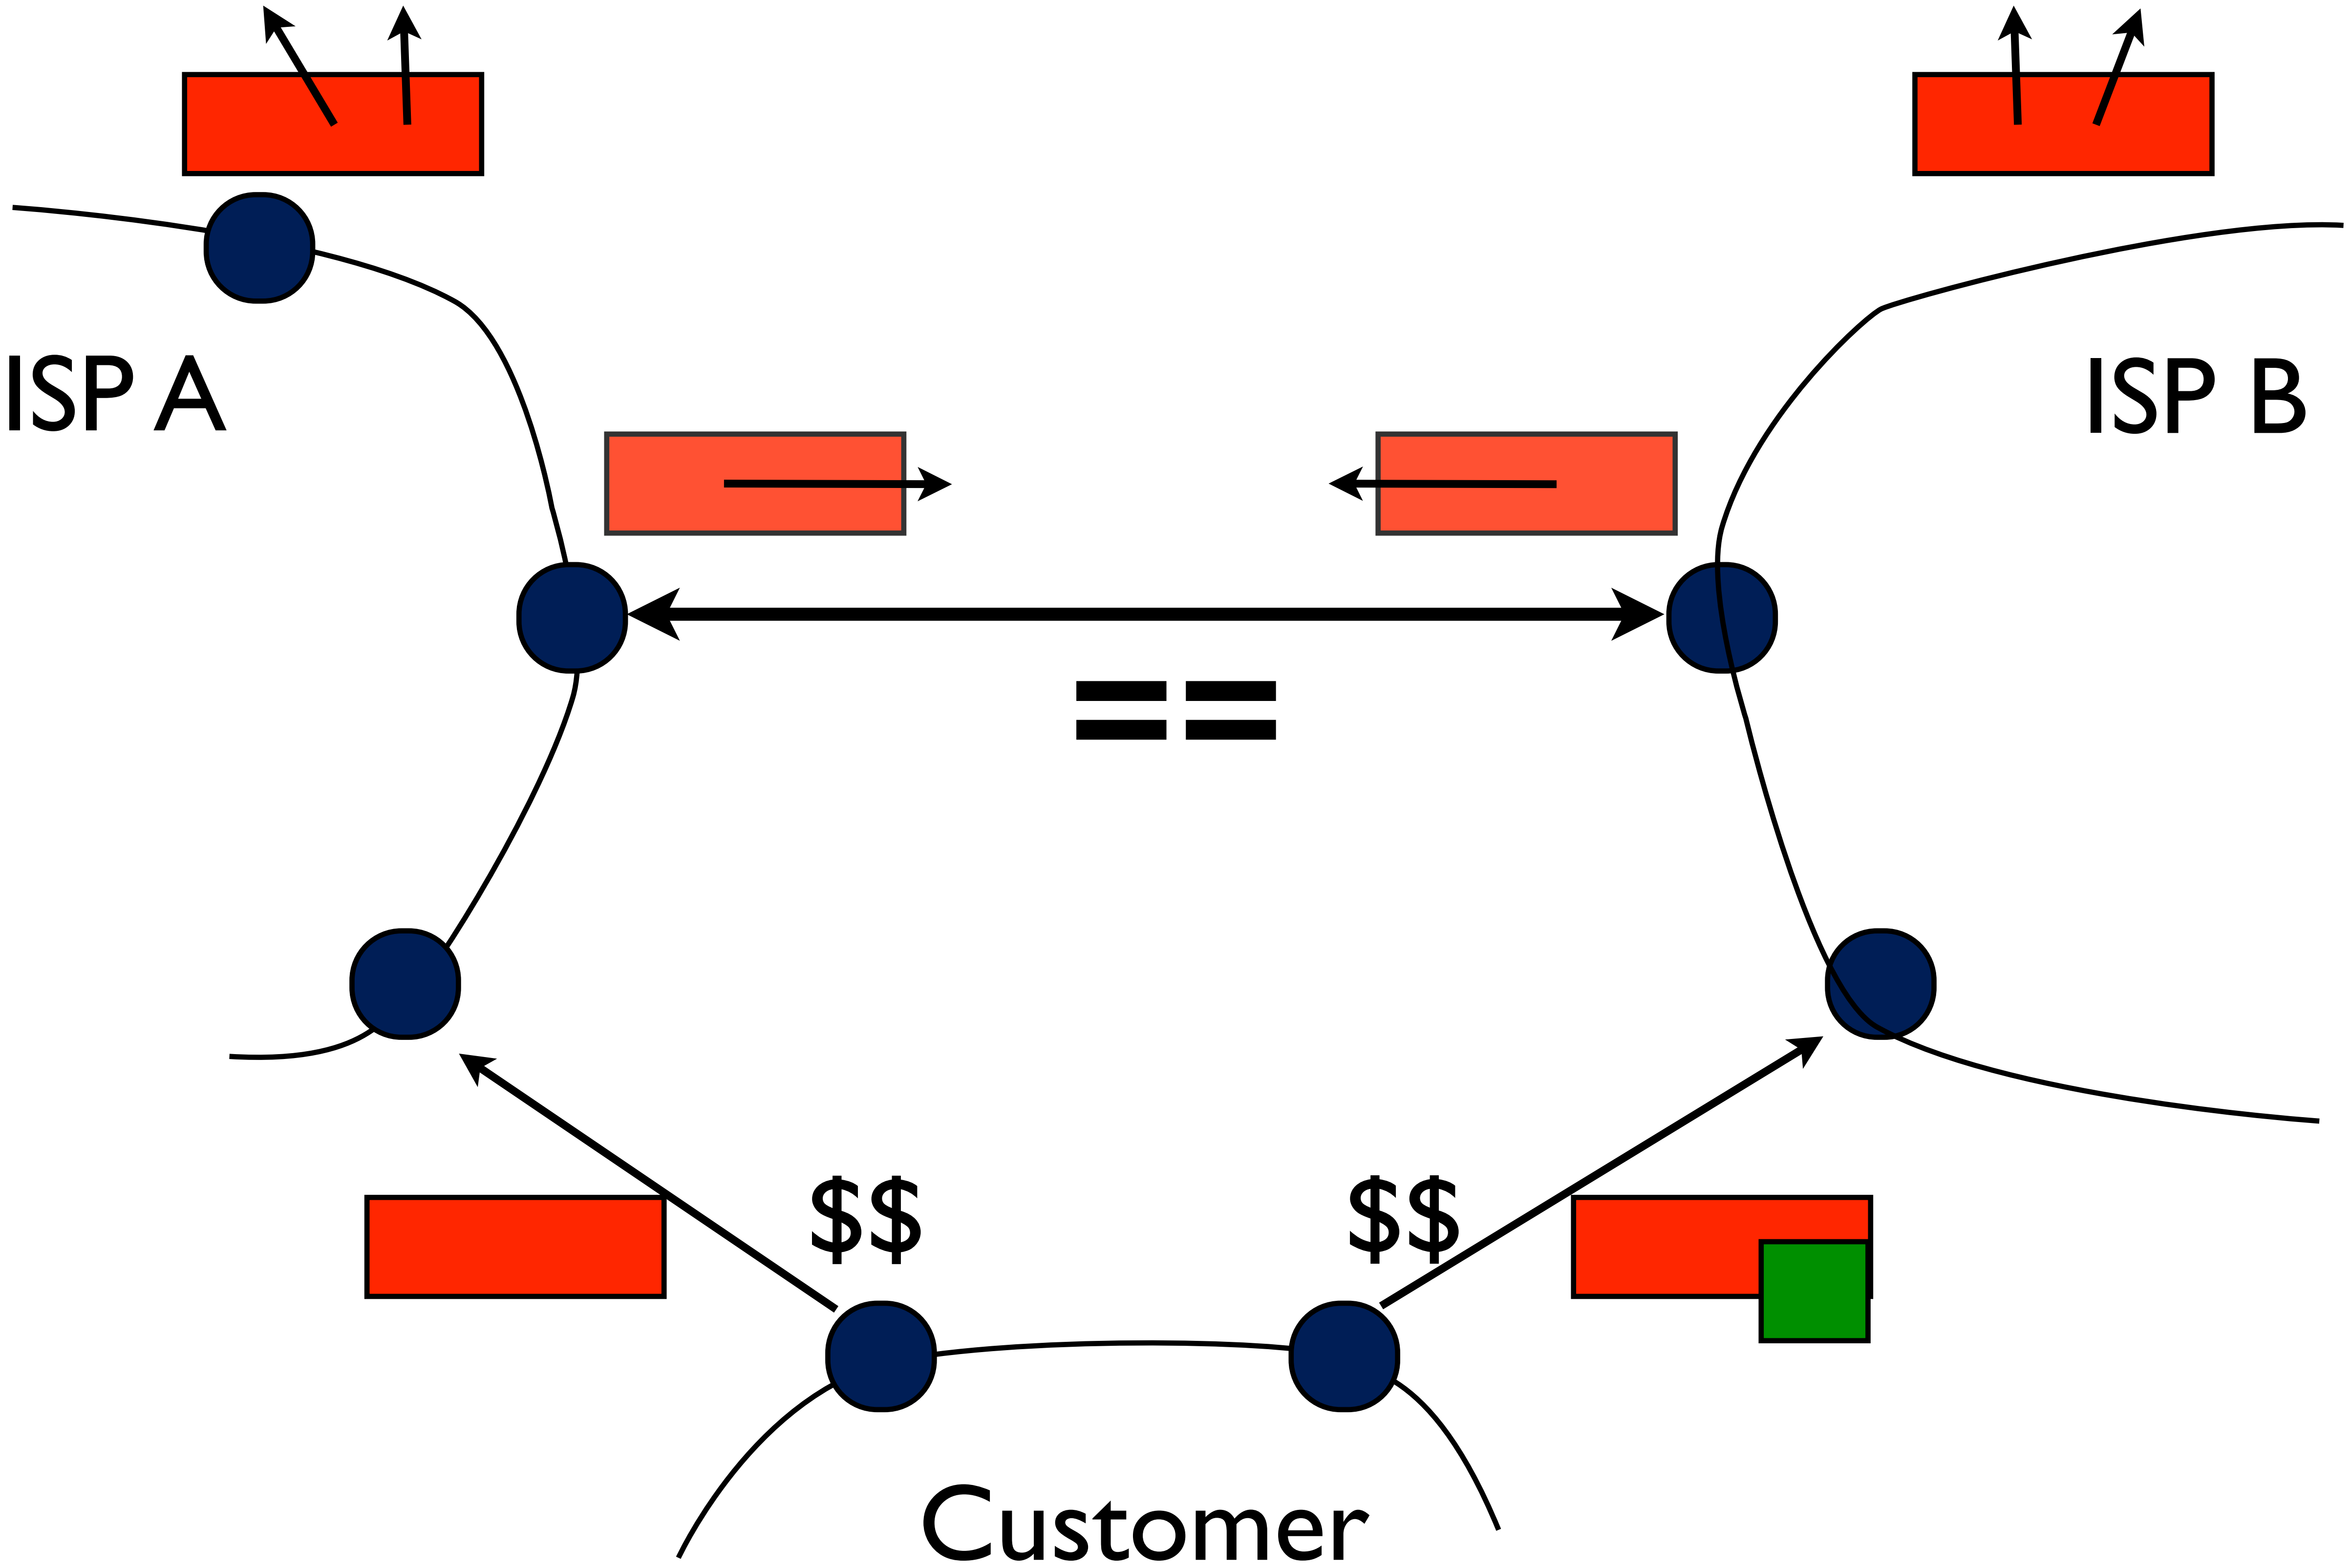

Inbound TE, selective advertisement of a more specific prefix

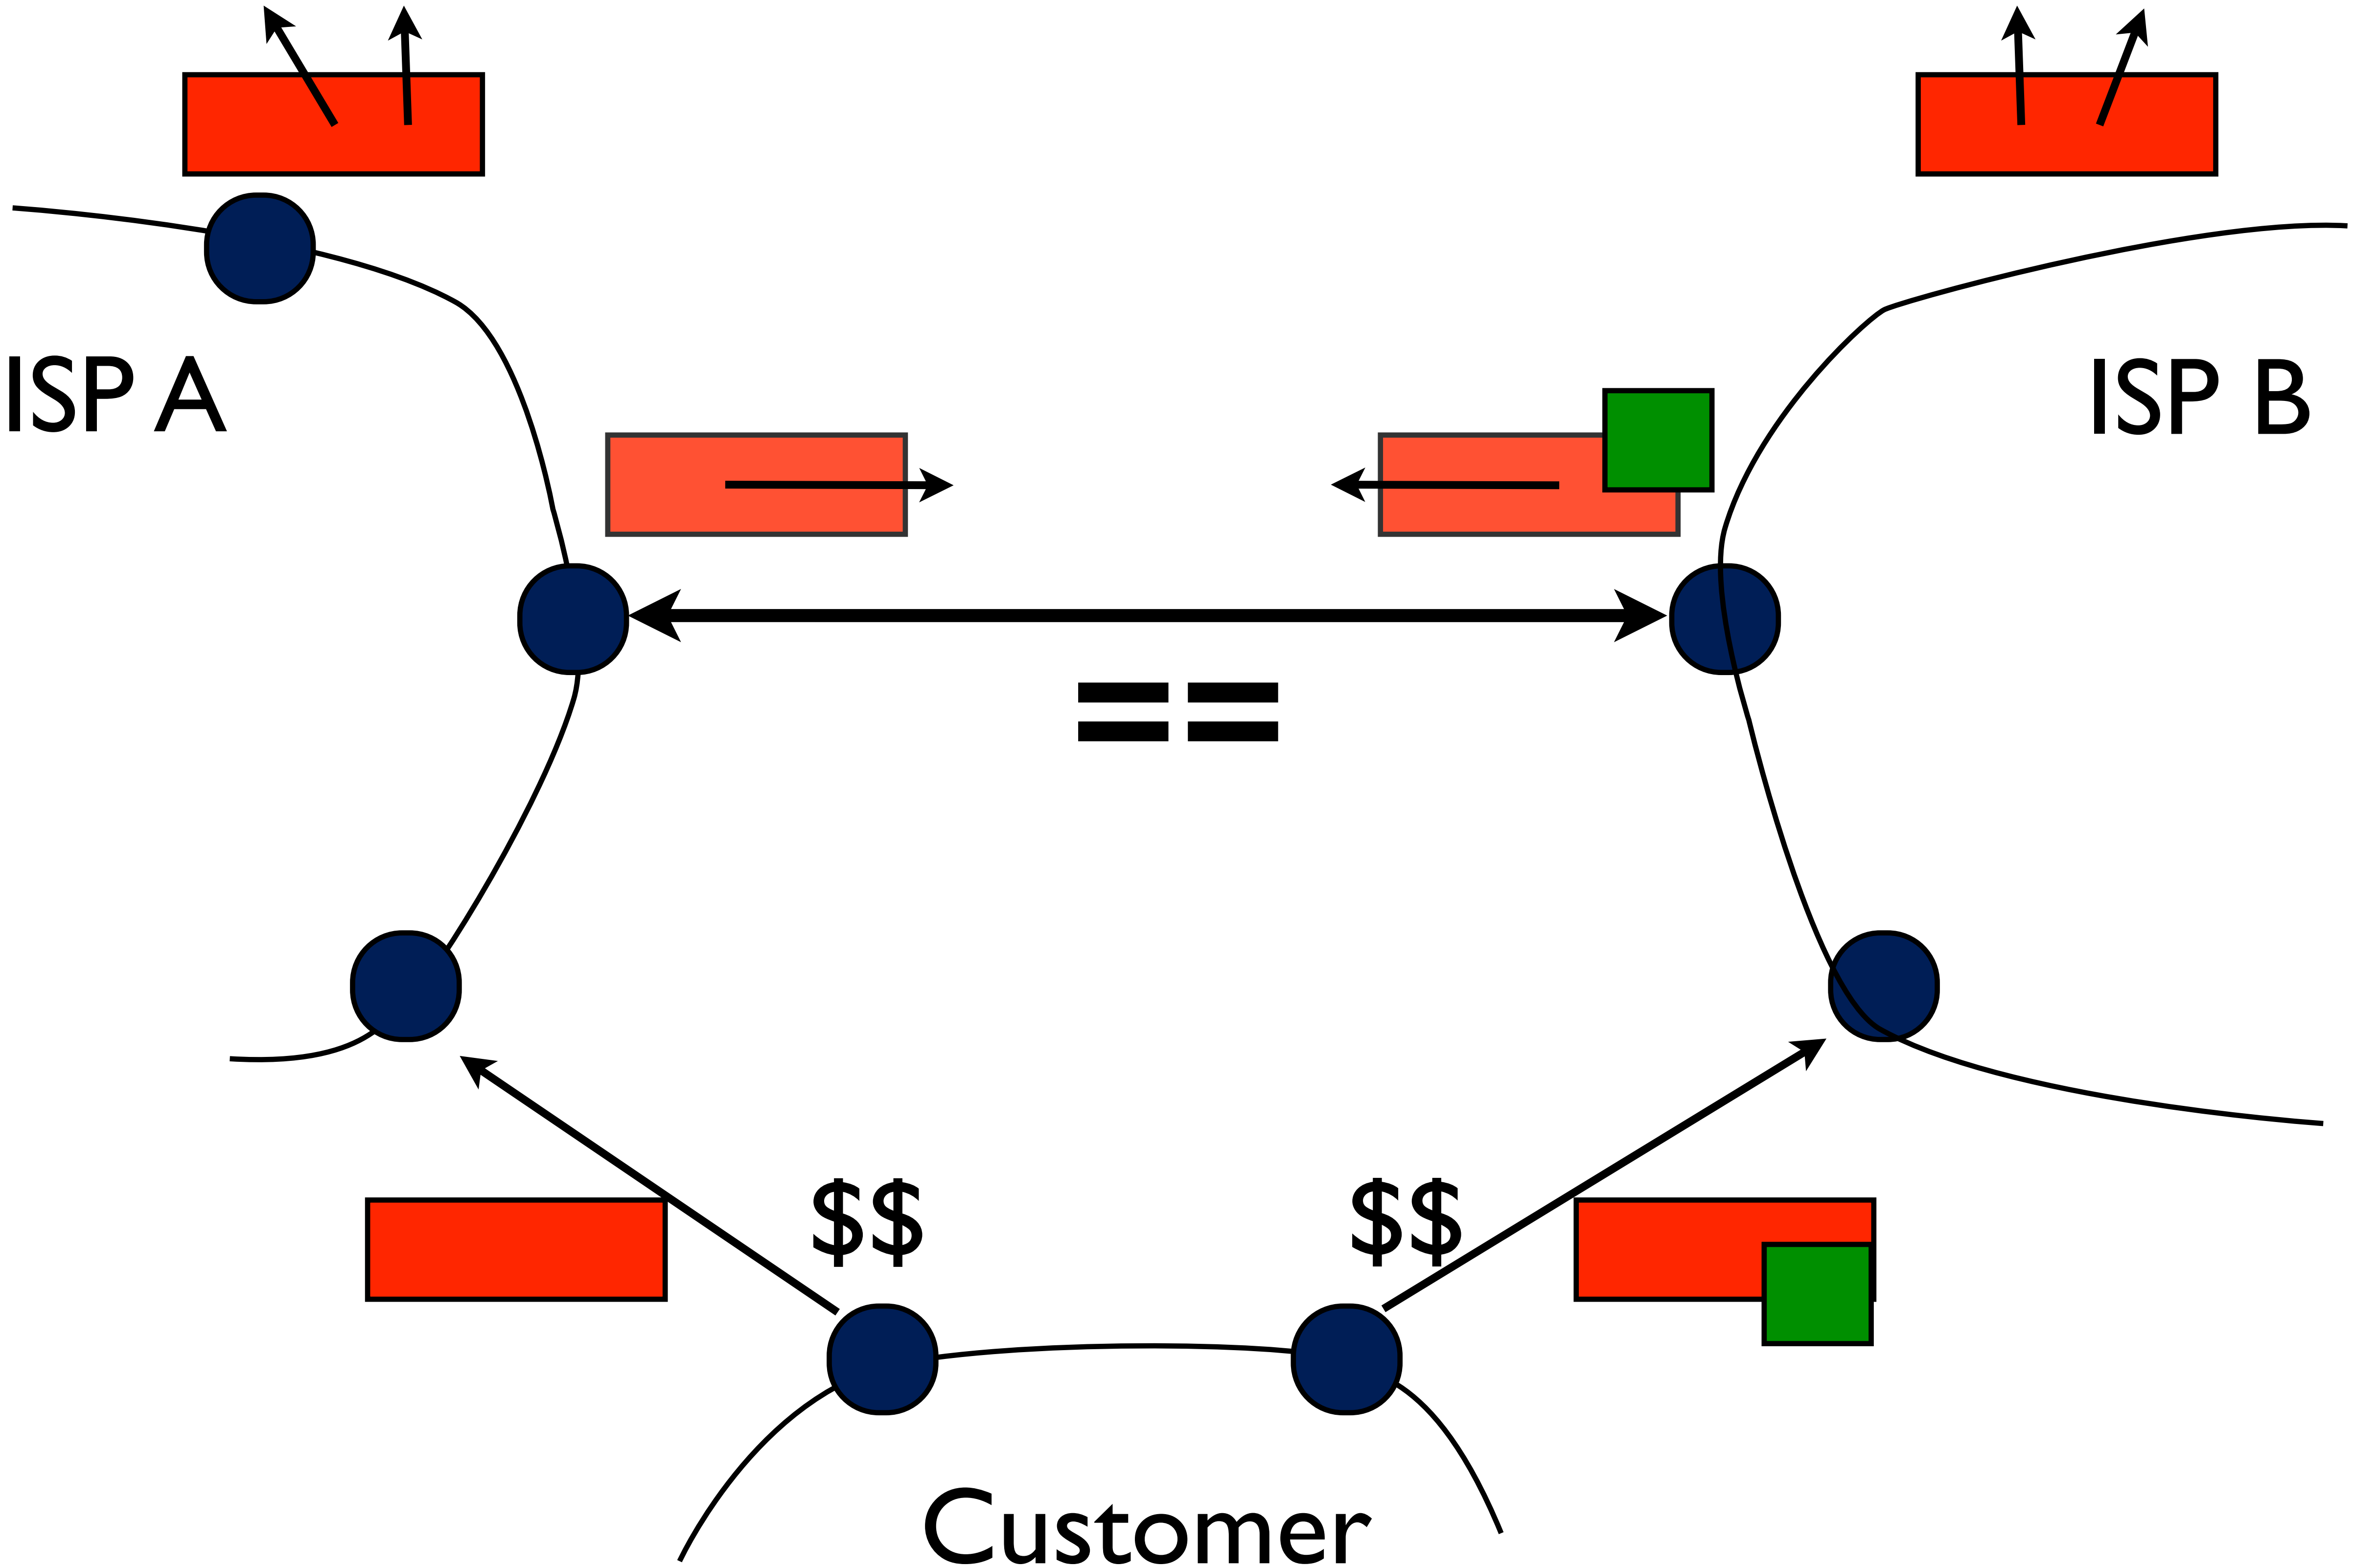

Inbound TE, selective advertisement of a more specific prefix

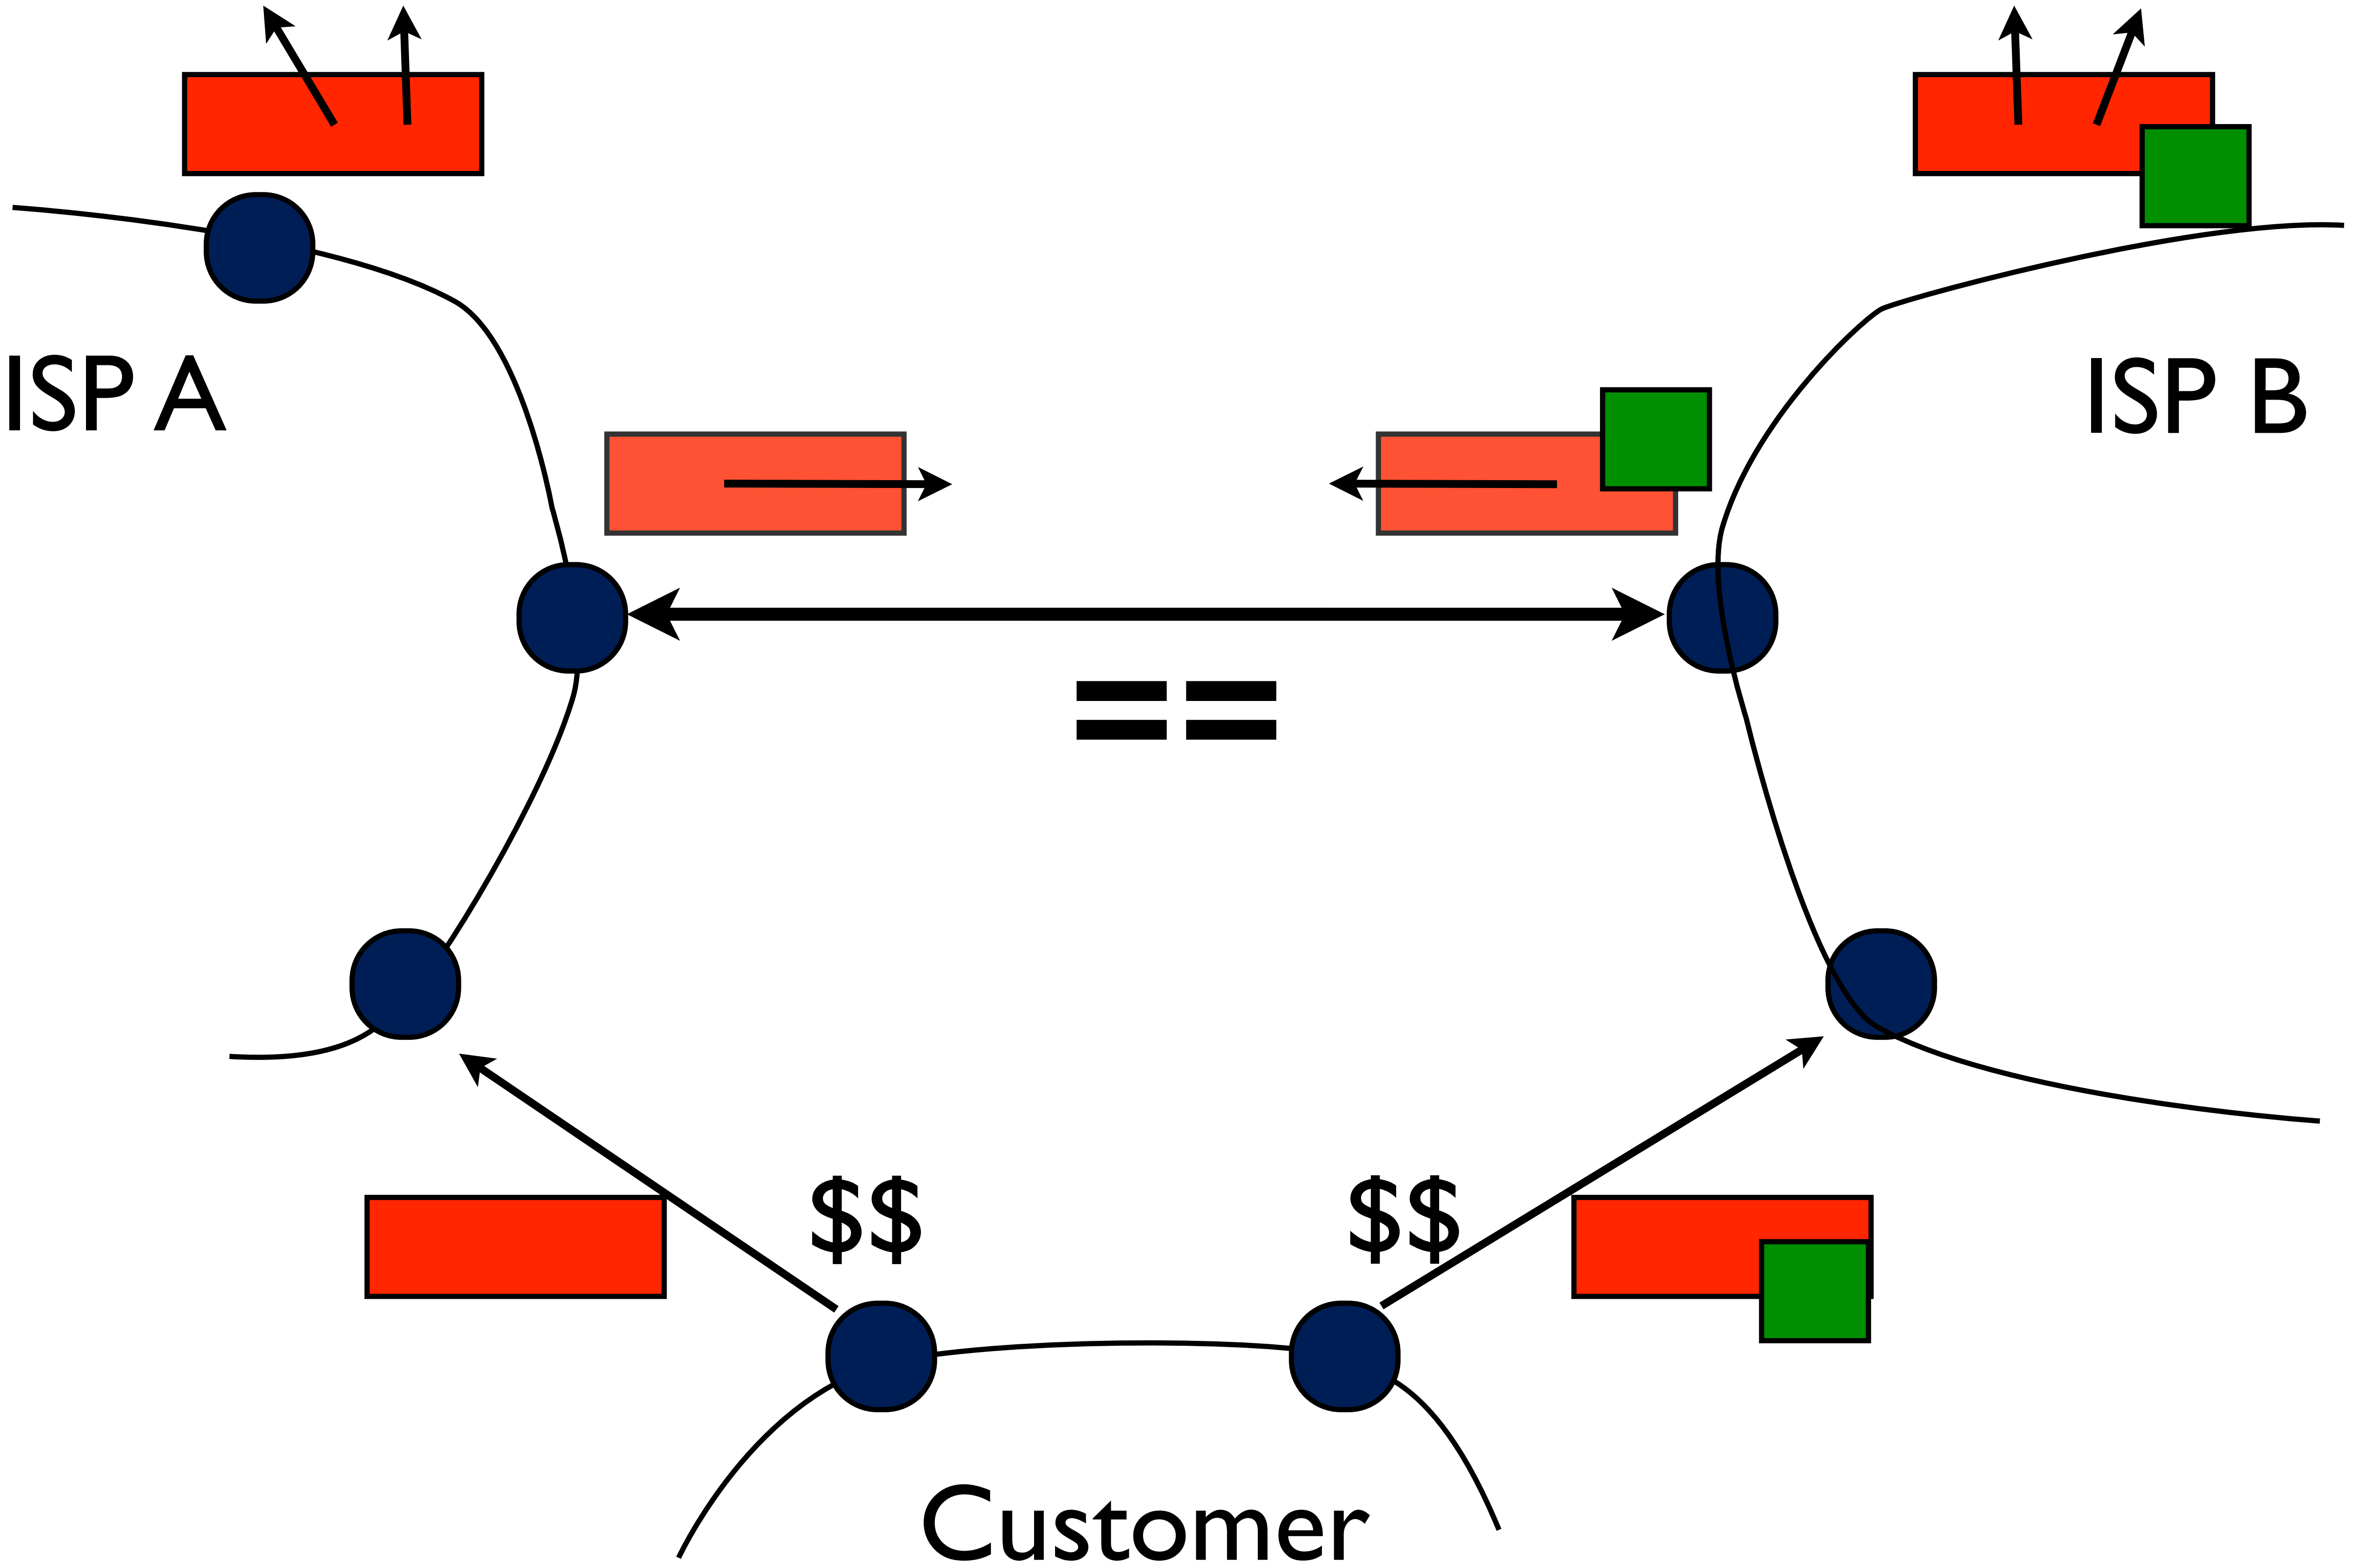

Inbound TE, selective advertisement of a more specific prefix

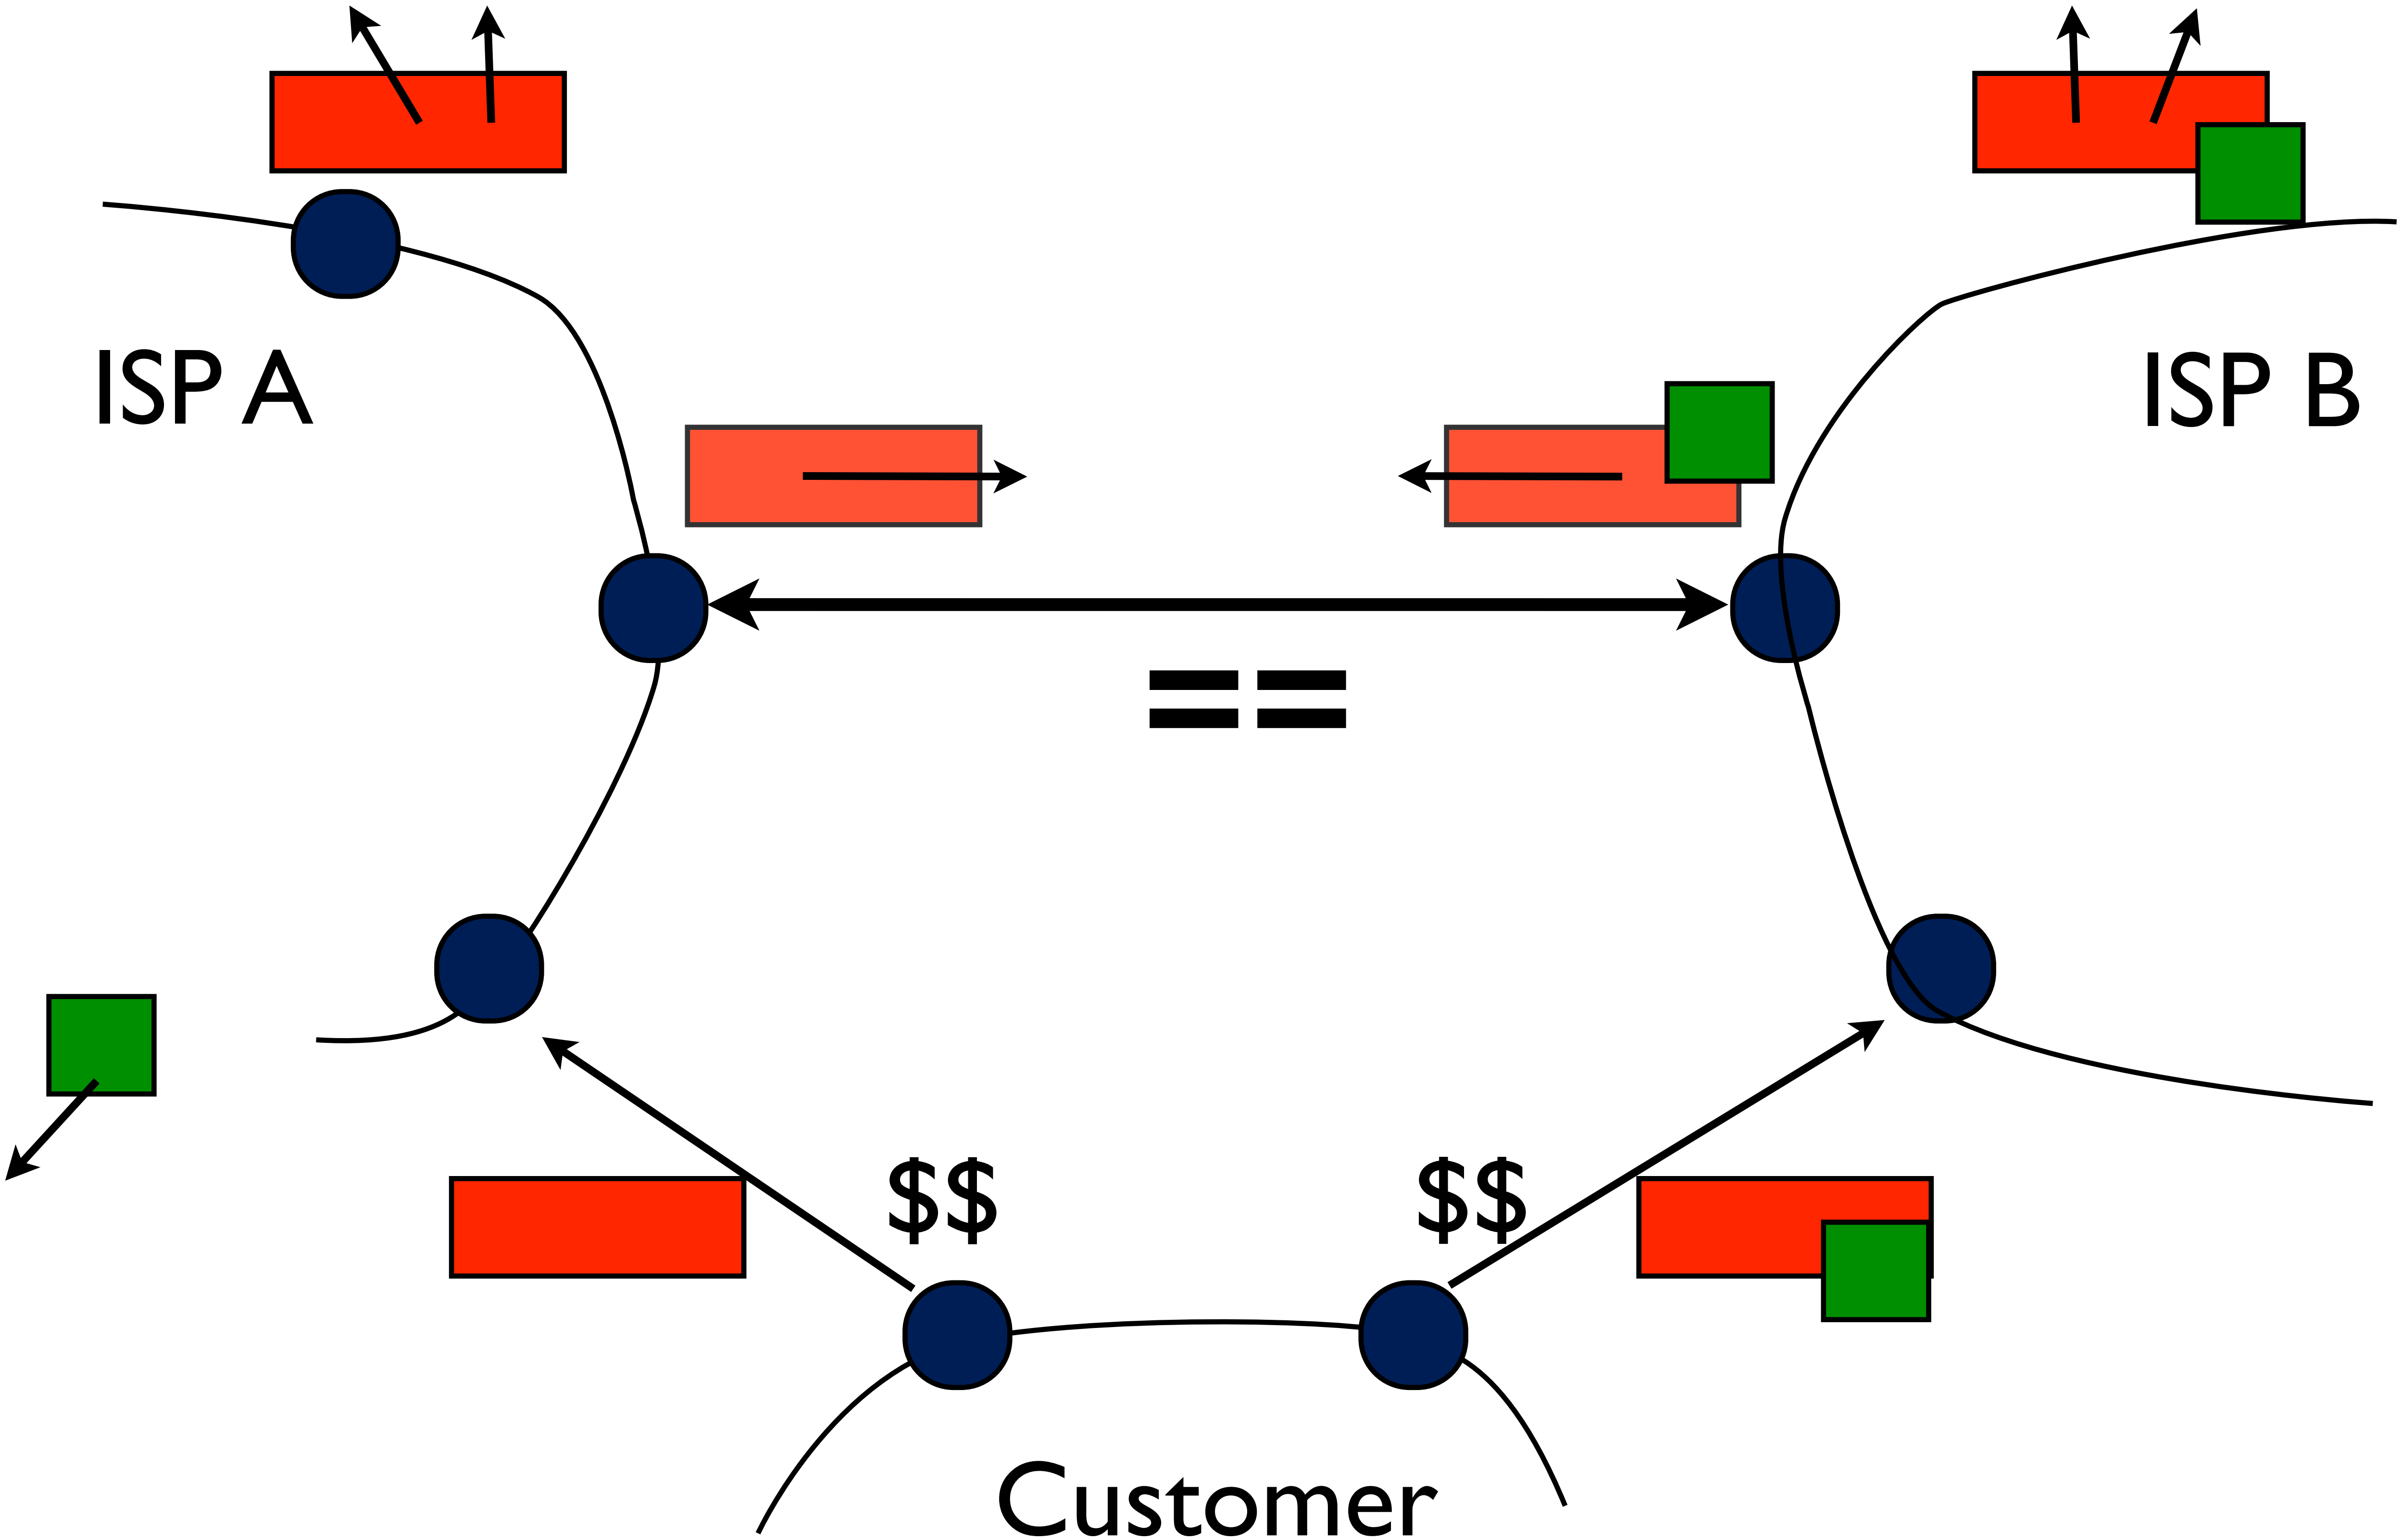

Inbound TE, selective advertisement of a more specific prefix

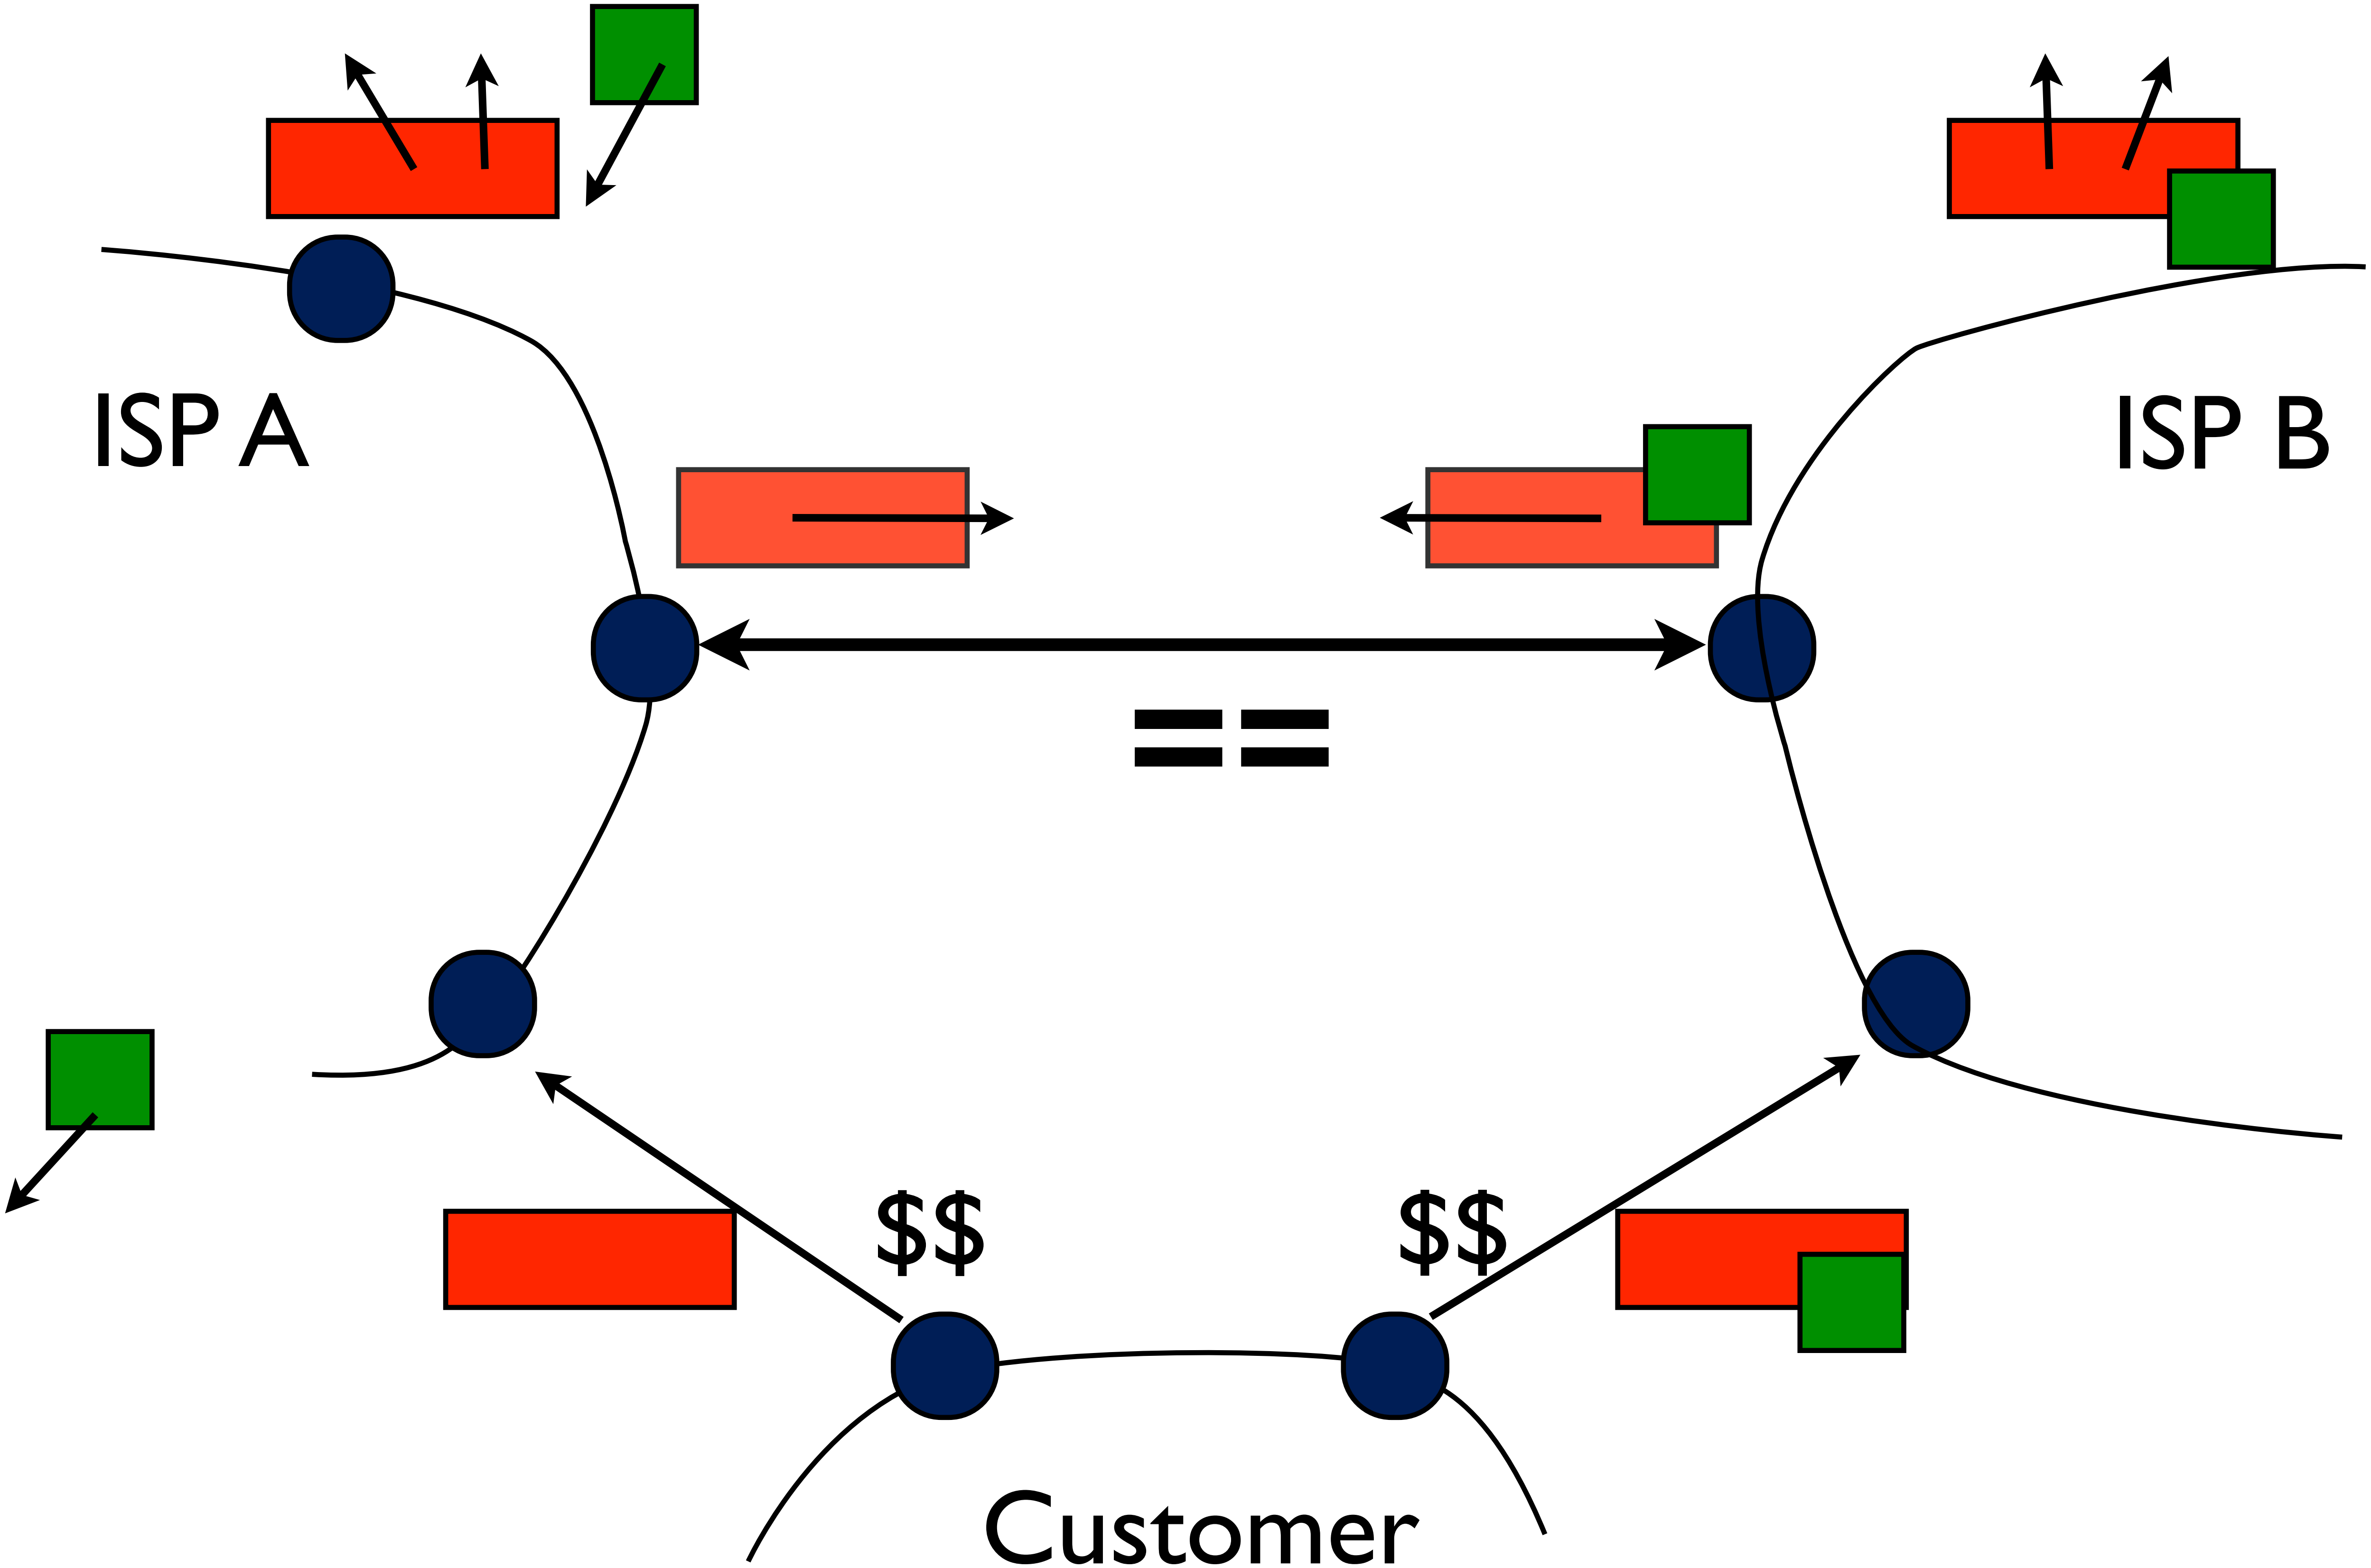

Inbound TE, selective advertisement of a more specific prefix

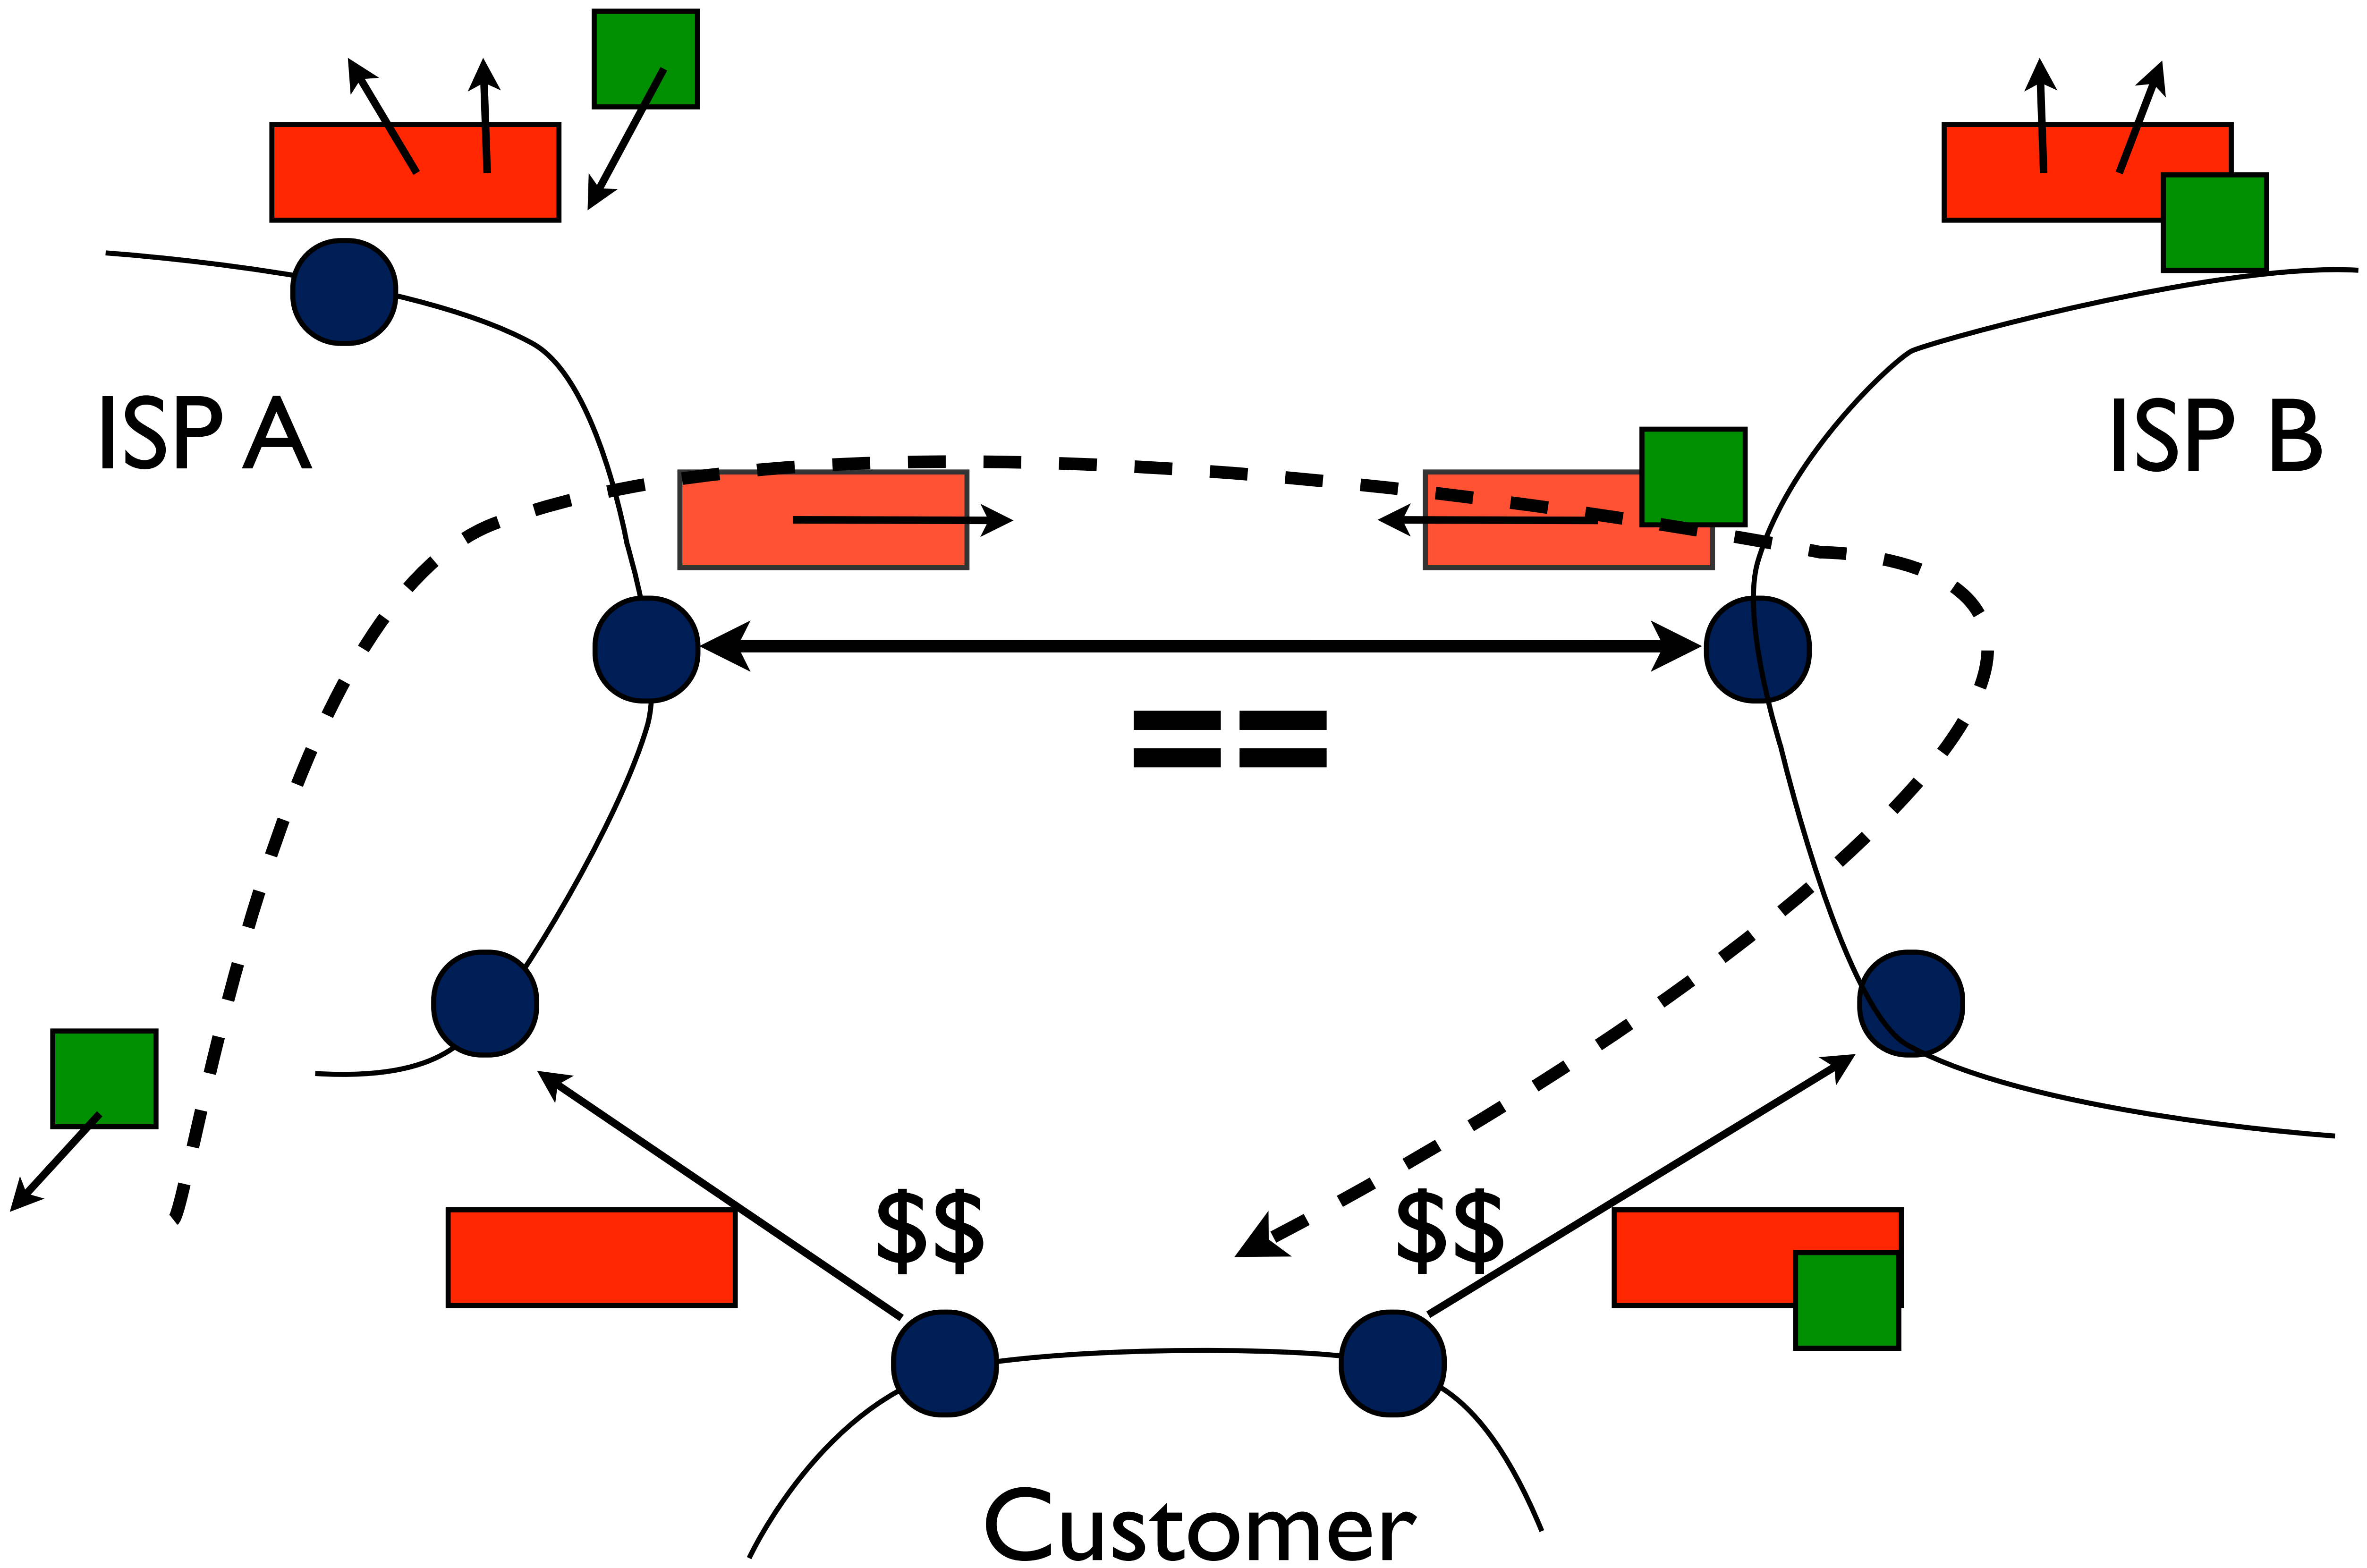

Inbound TE, selective advertisement of a more specific prefix

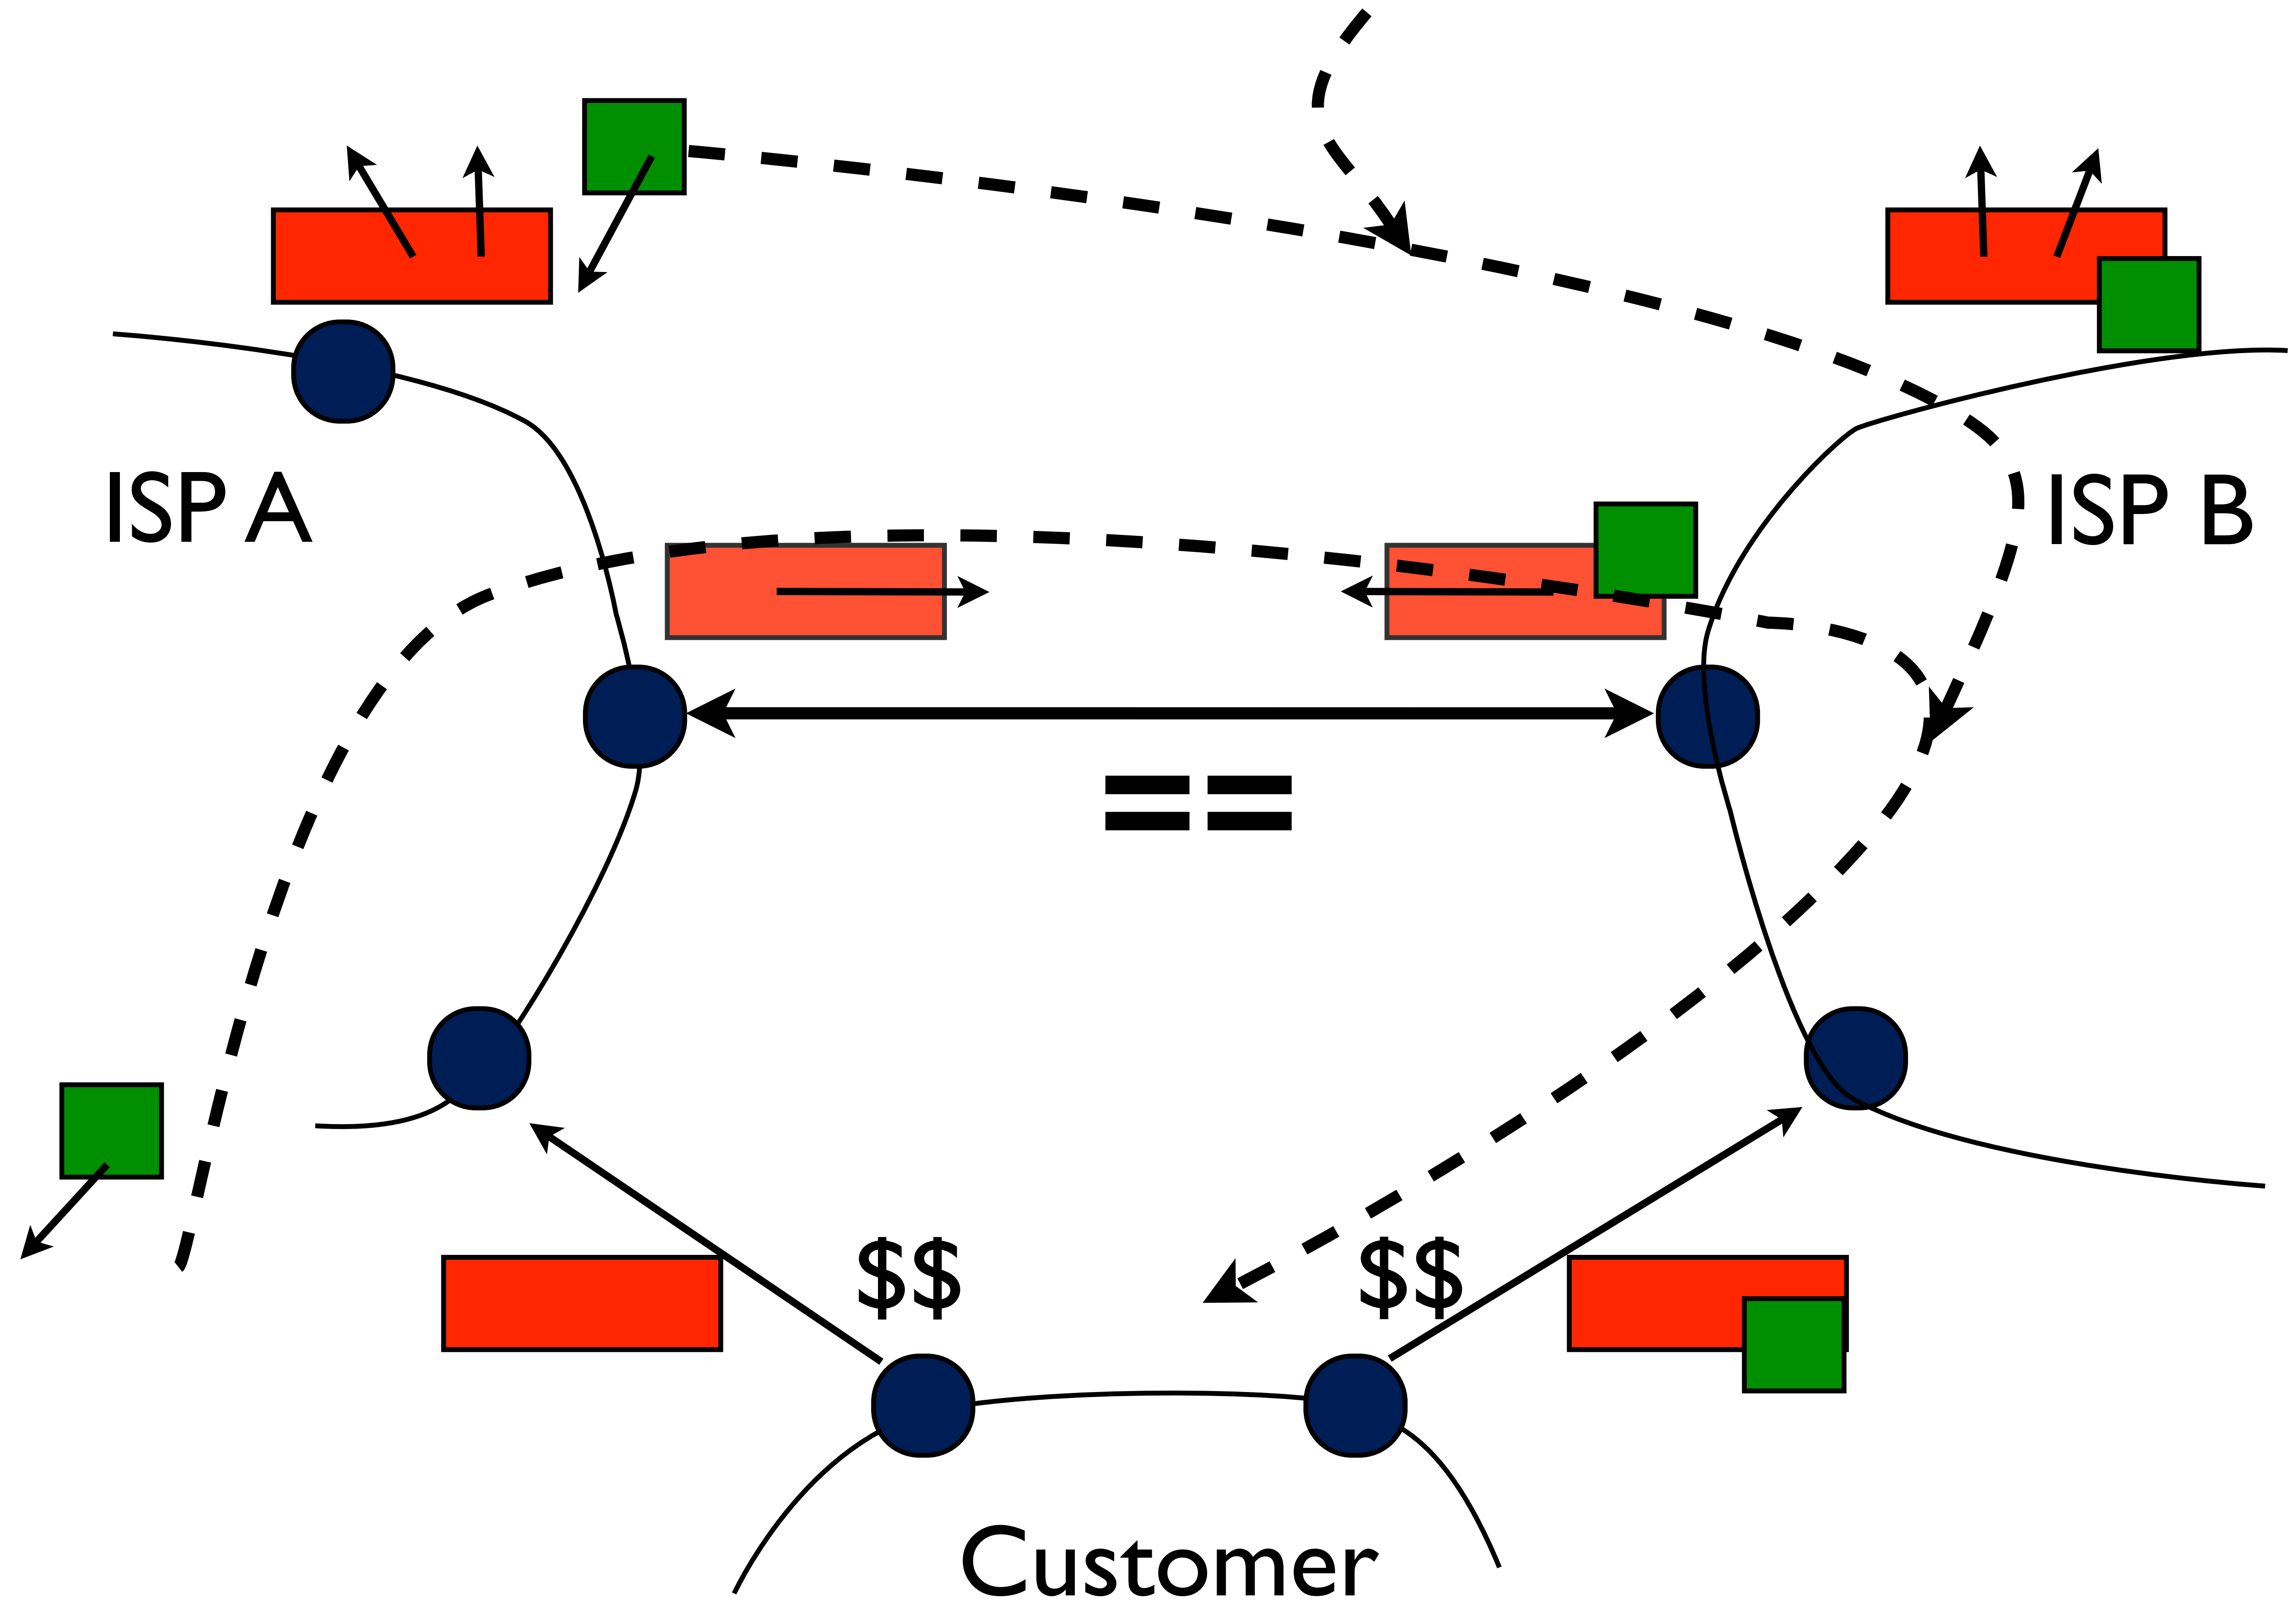

Scope the advertisement of the more specific

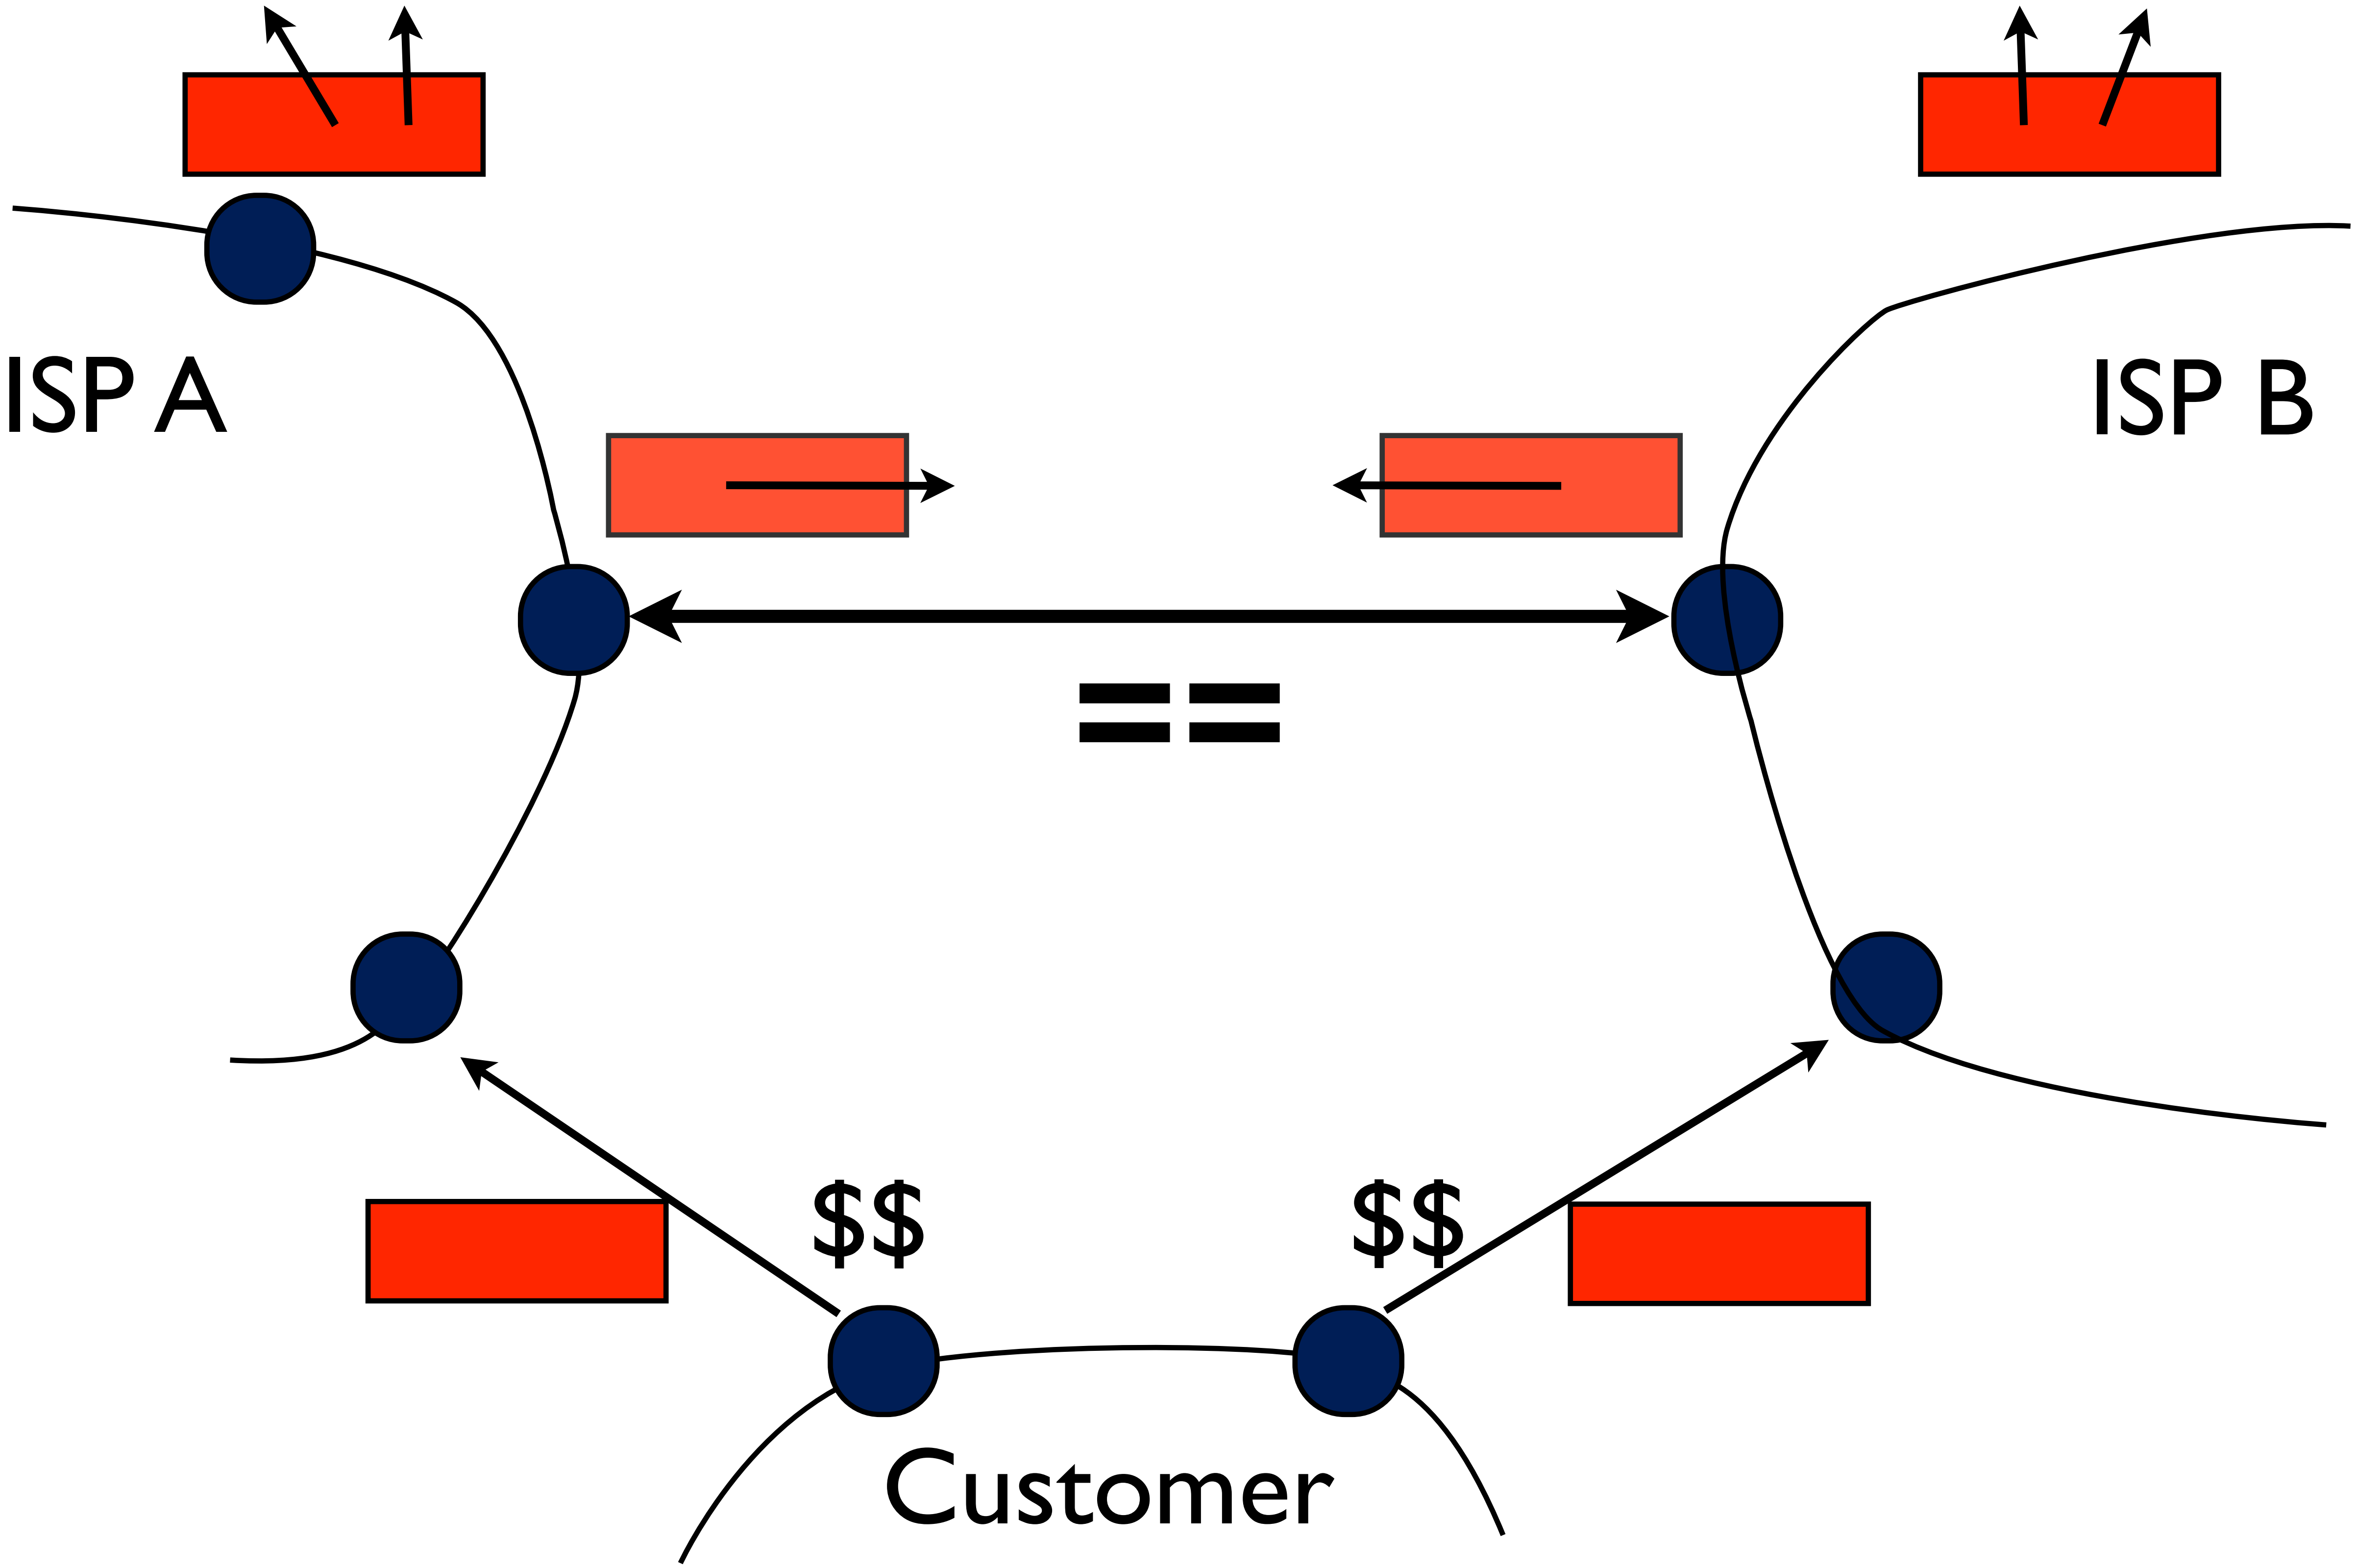

Scope the advertisement of the more specific

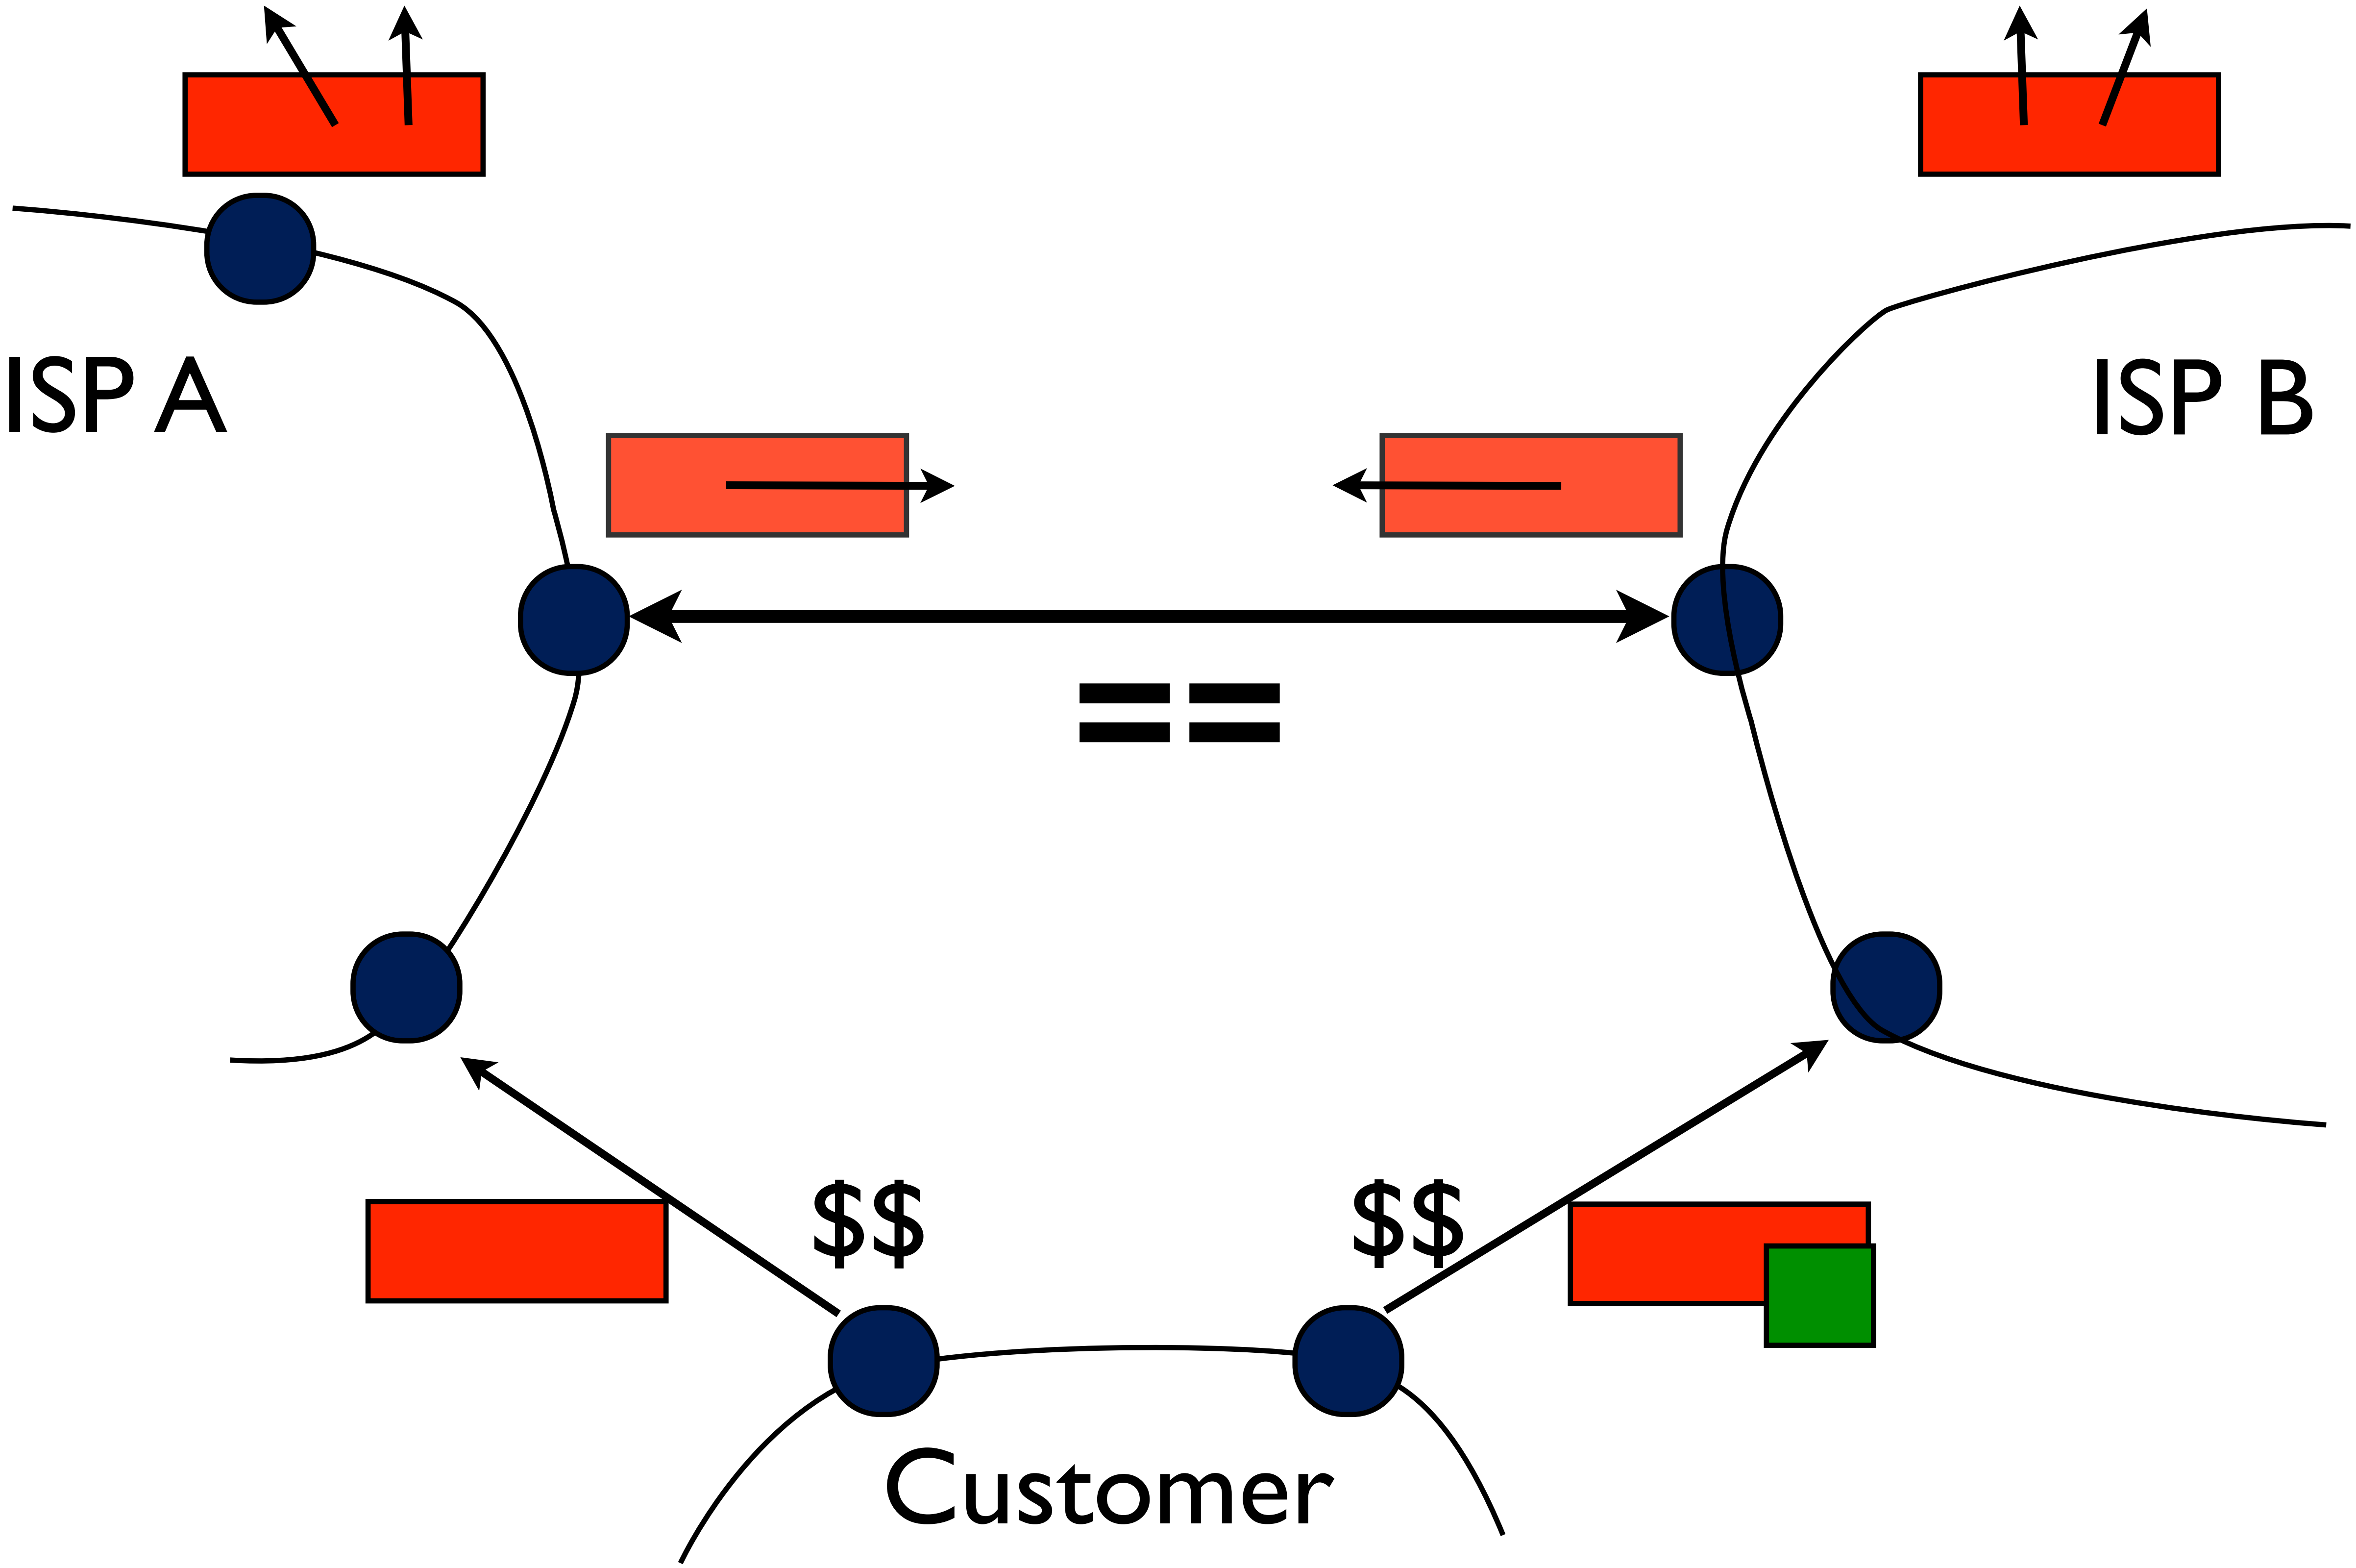

Scope the advertisement of the more specific

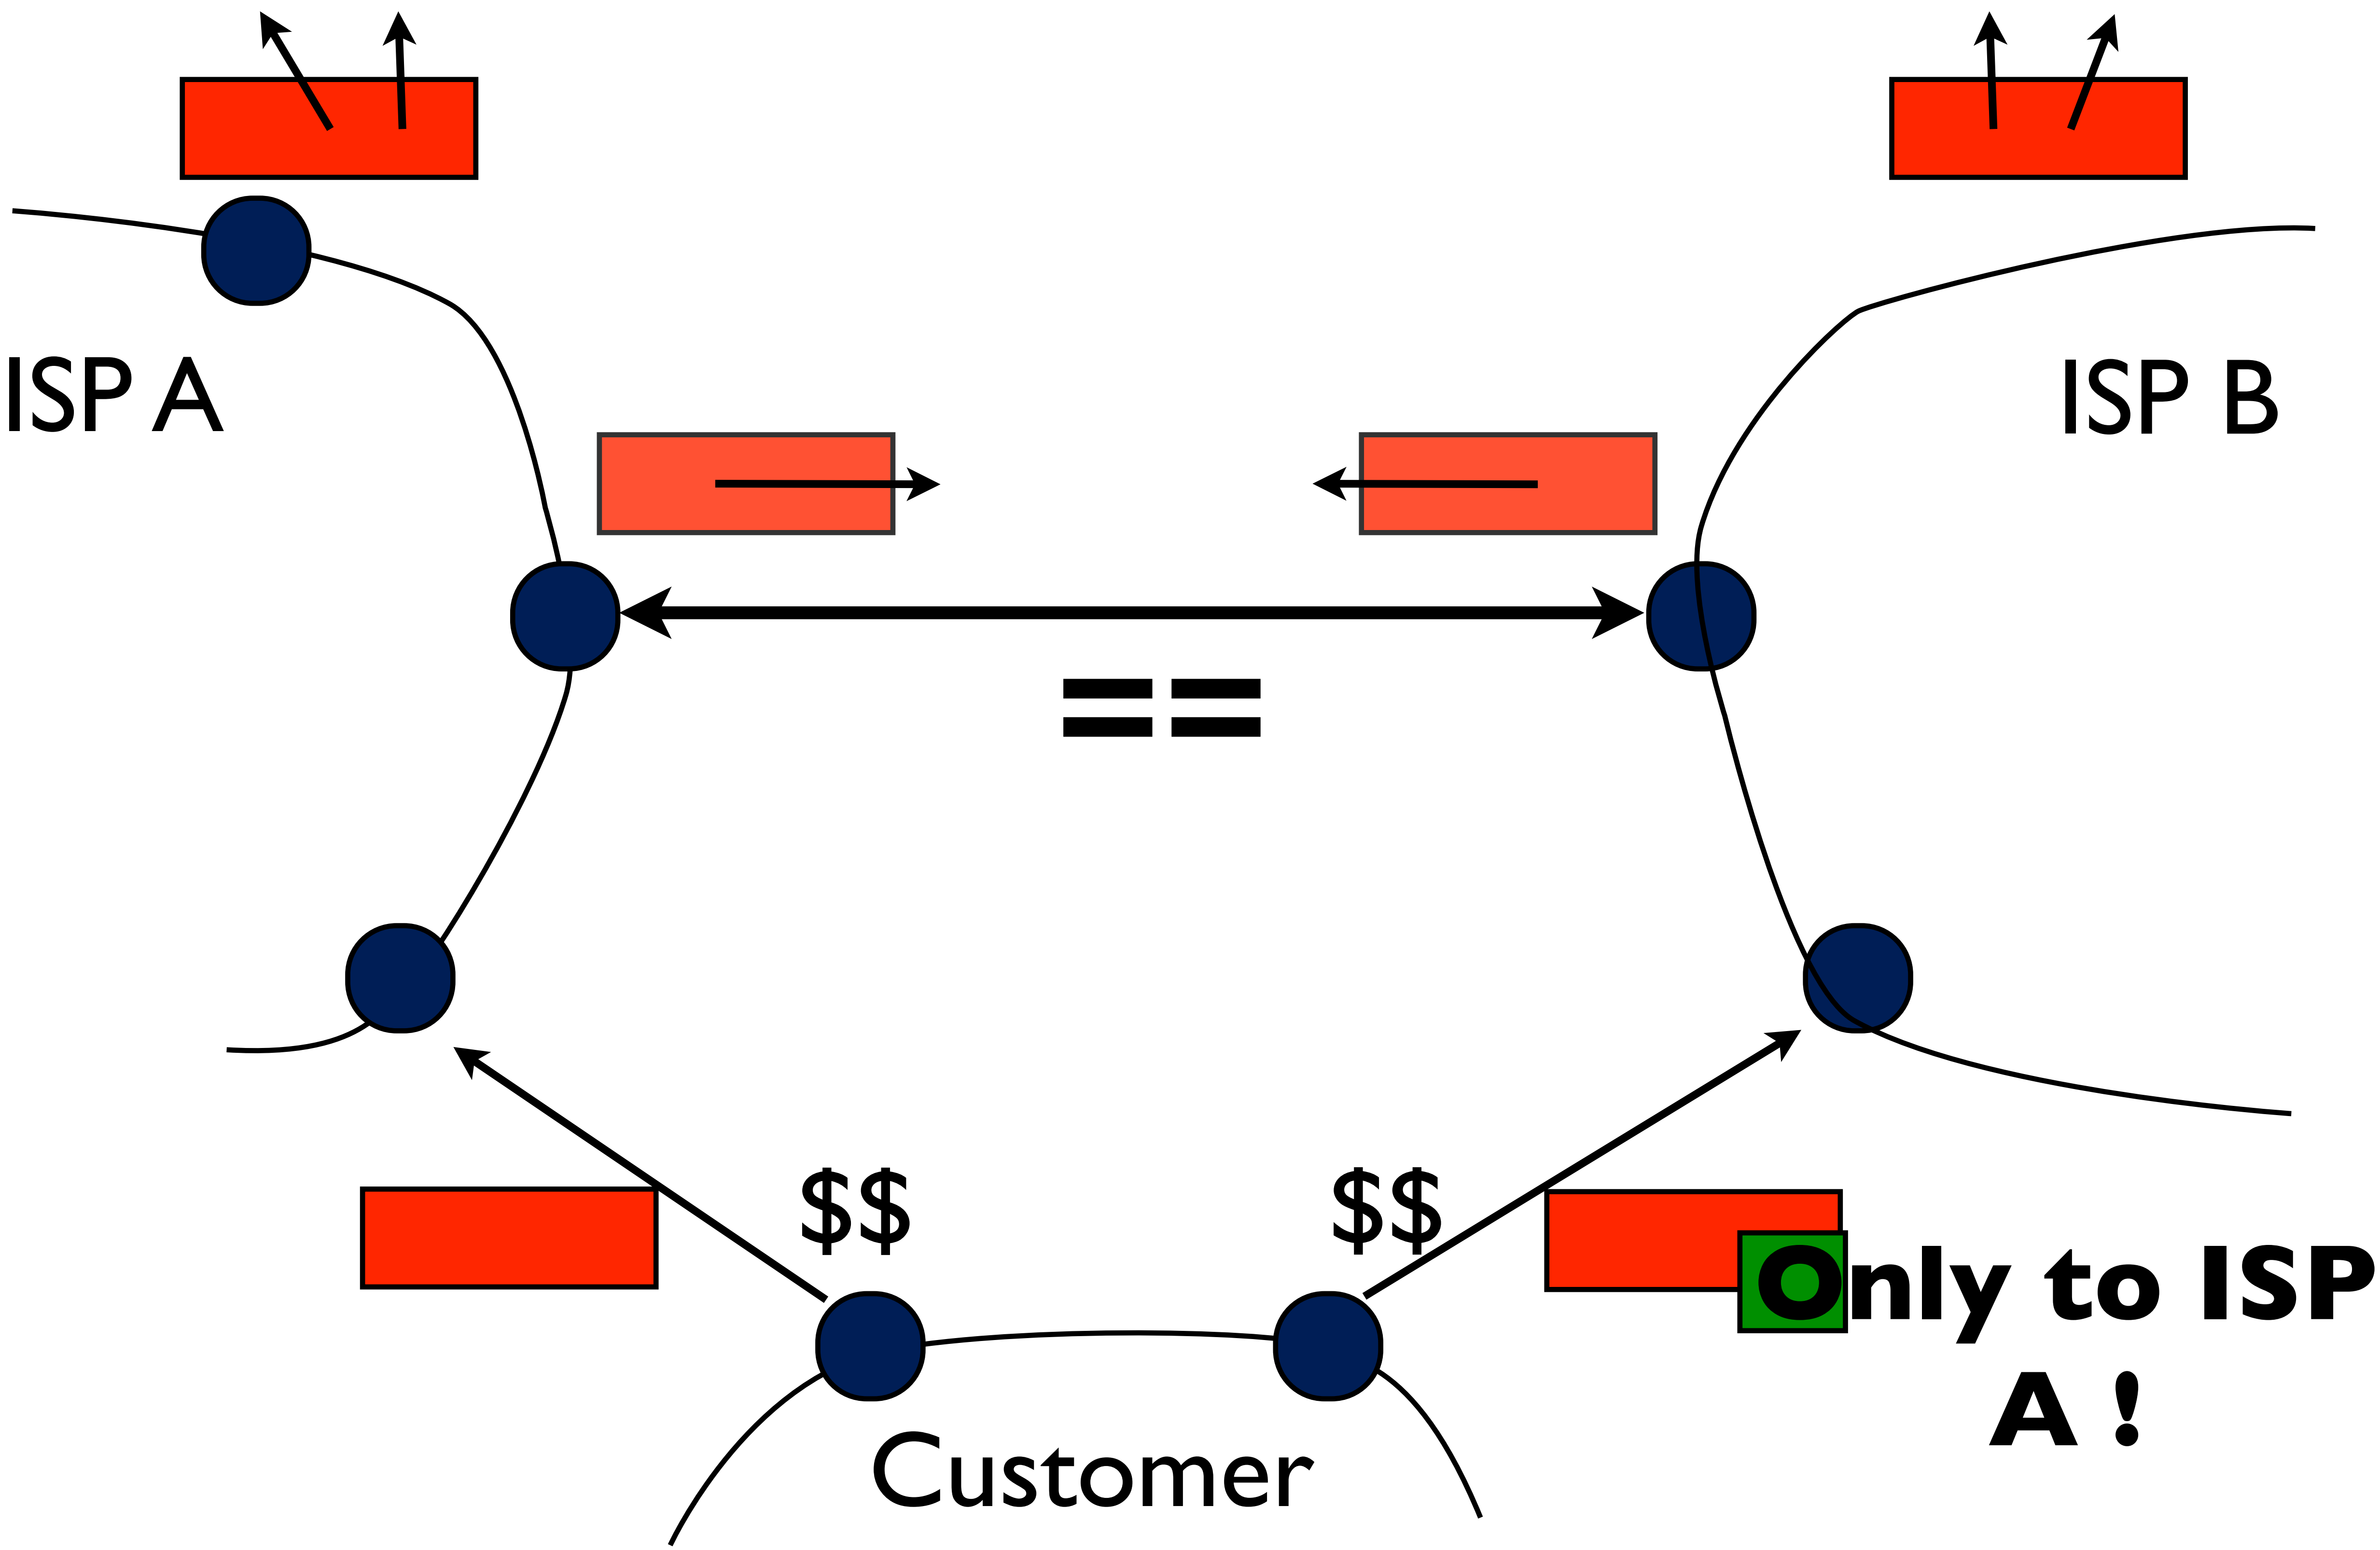

Only to ISP A !

Scope the advertisement of the more specific

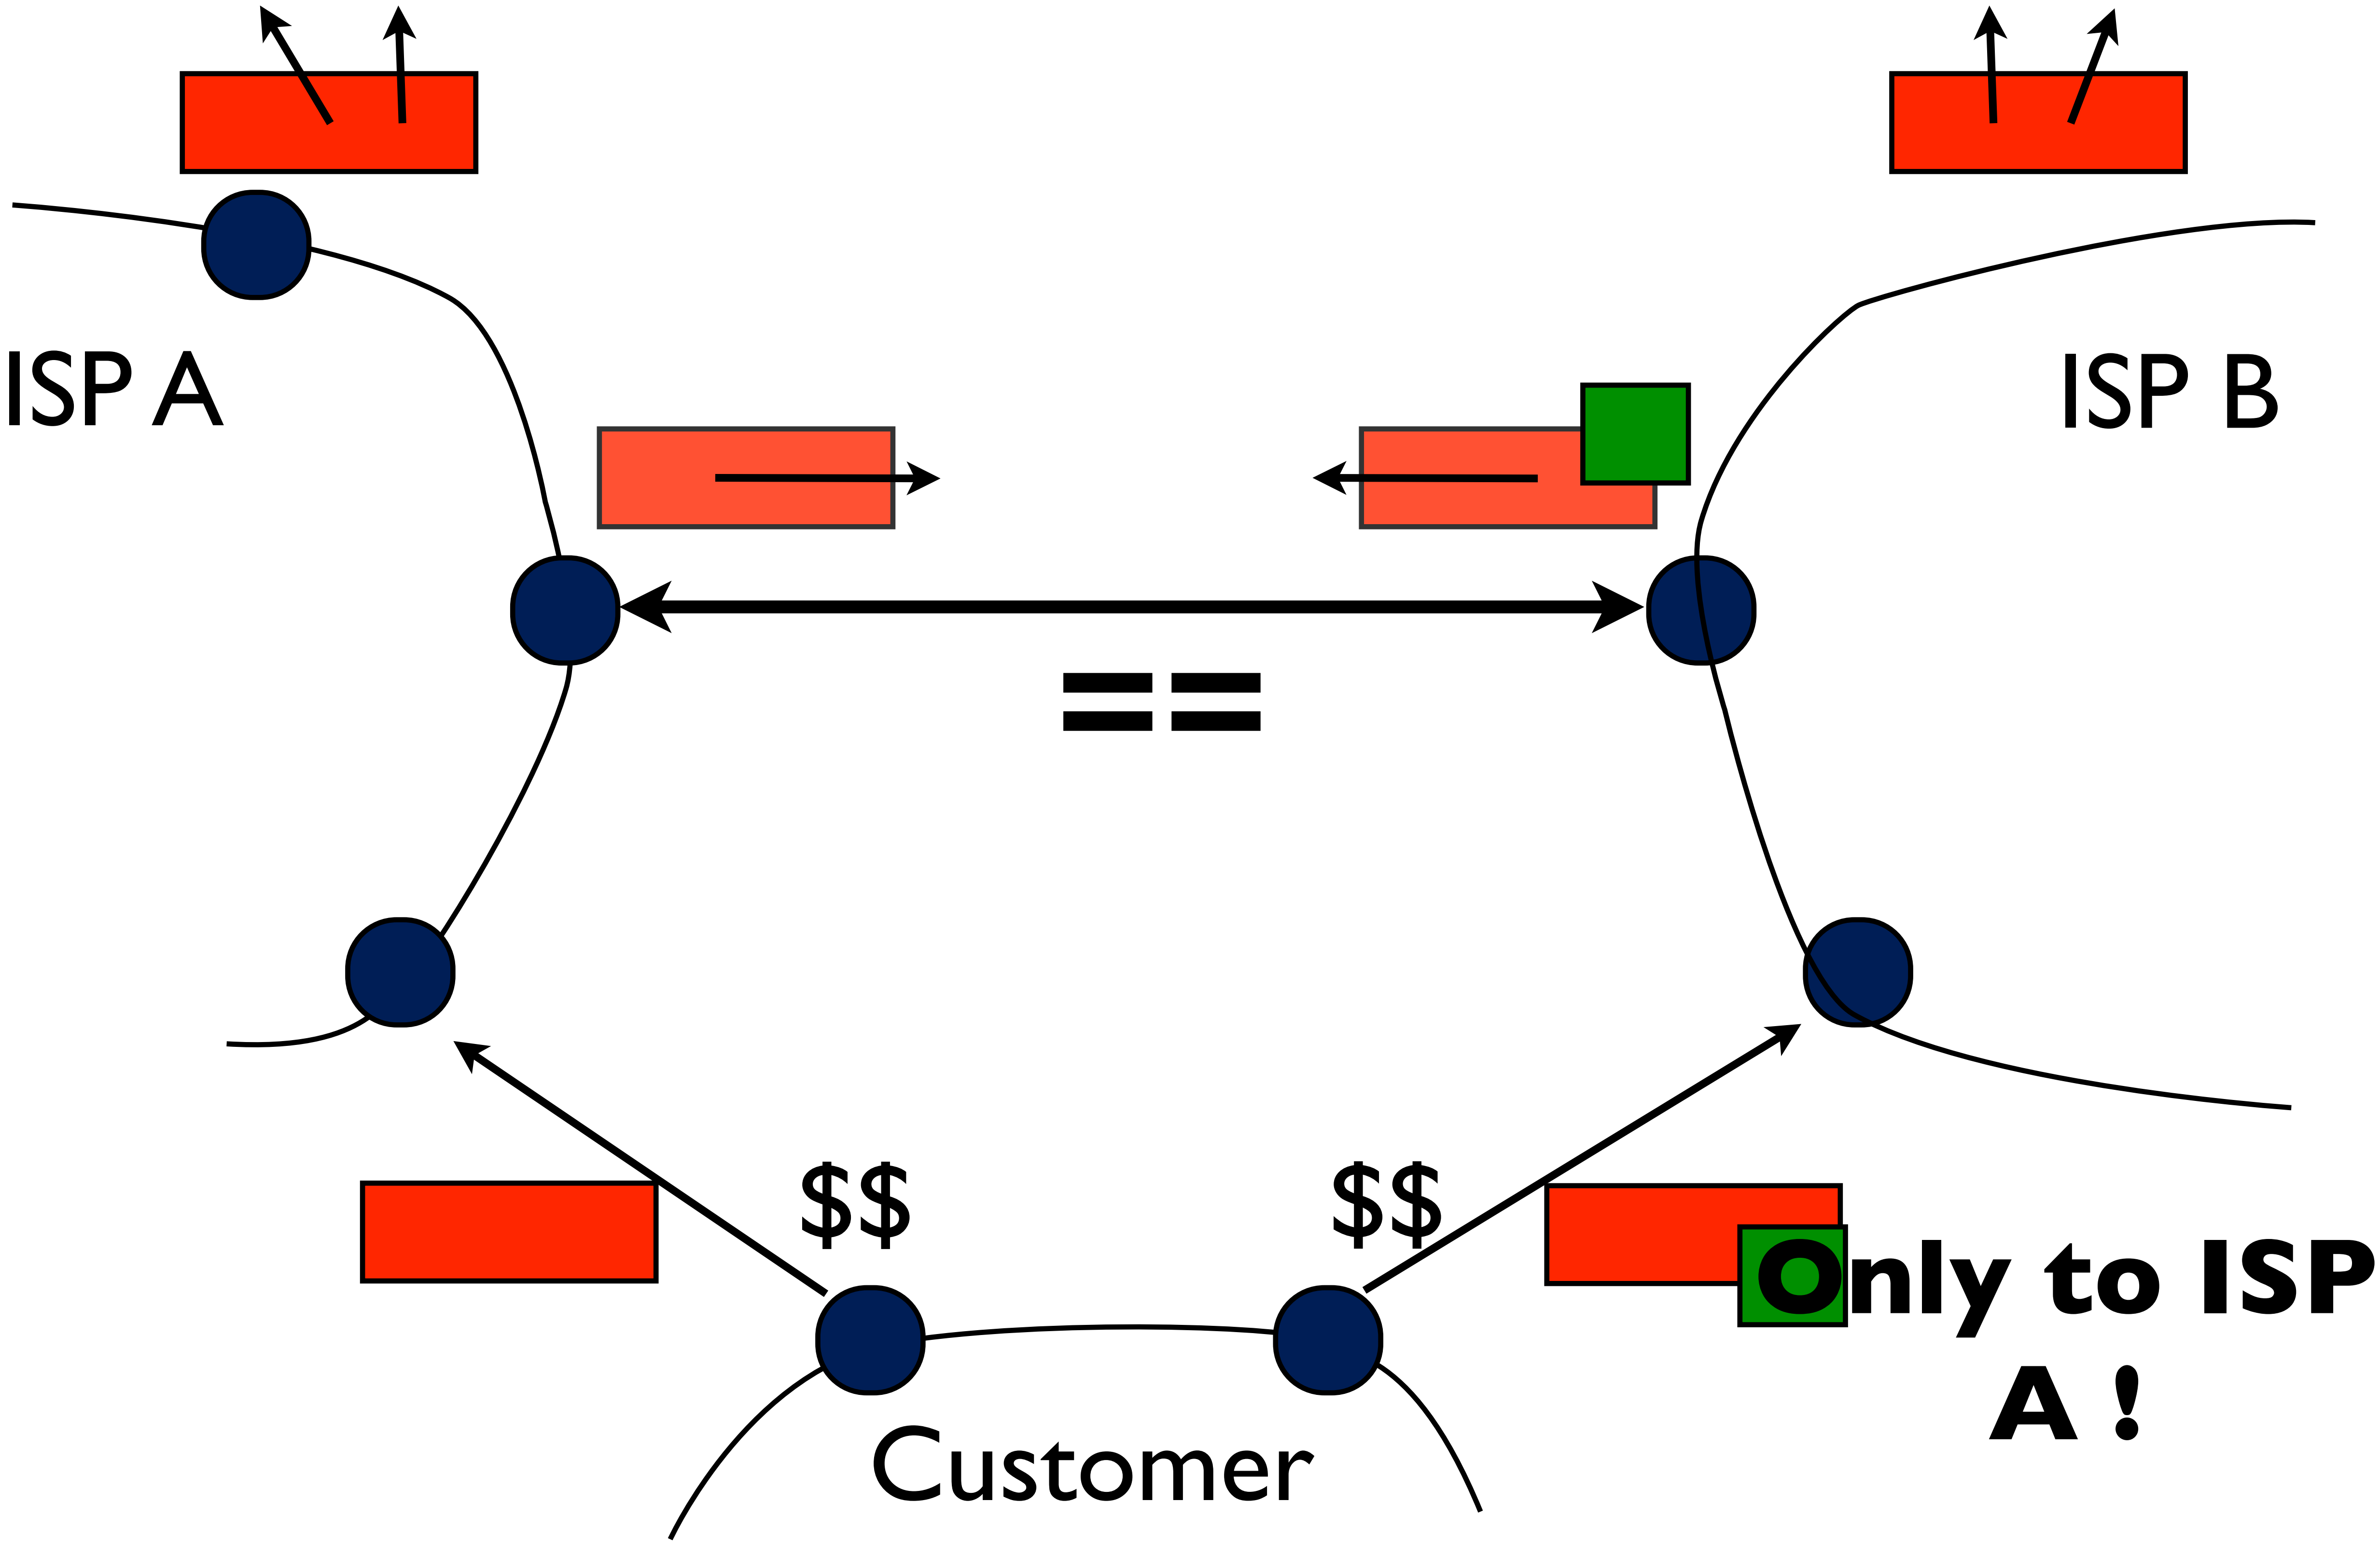

Only to ISP A !

Scope the advertisement of the more specific

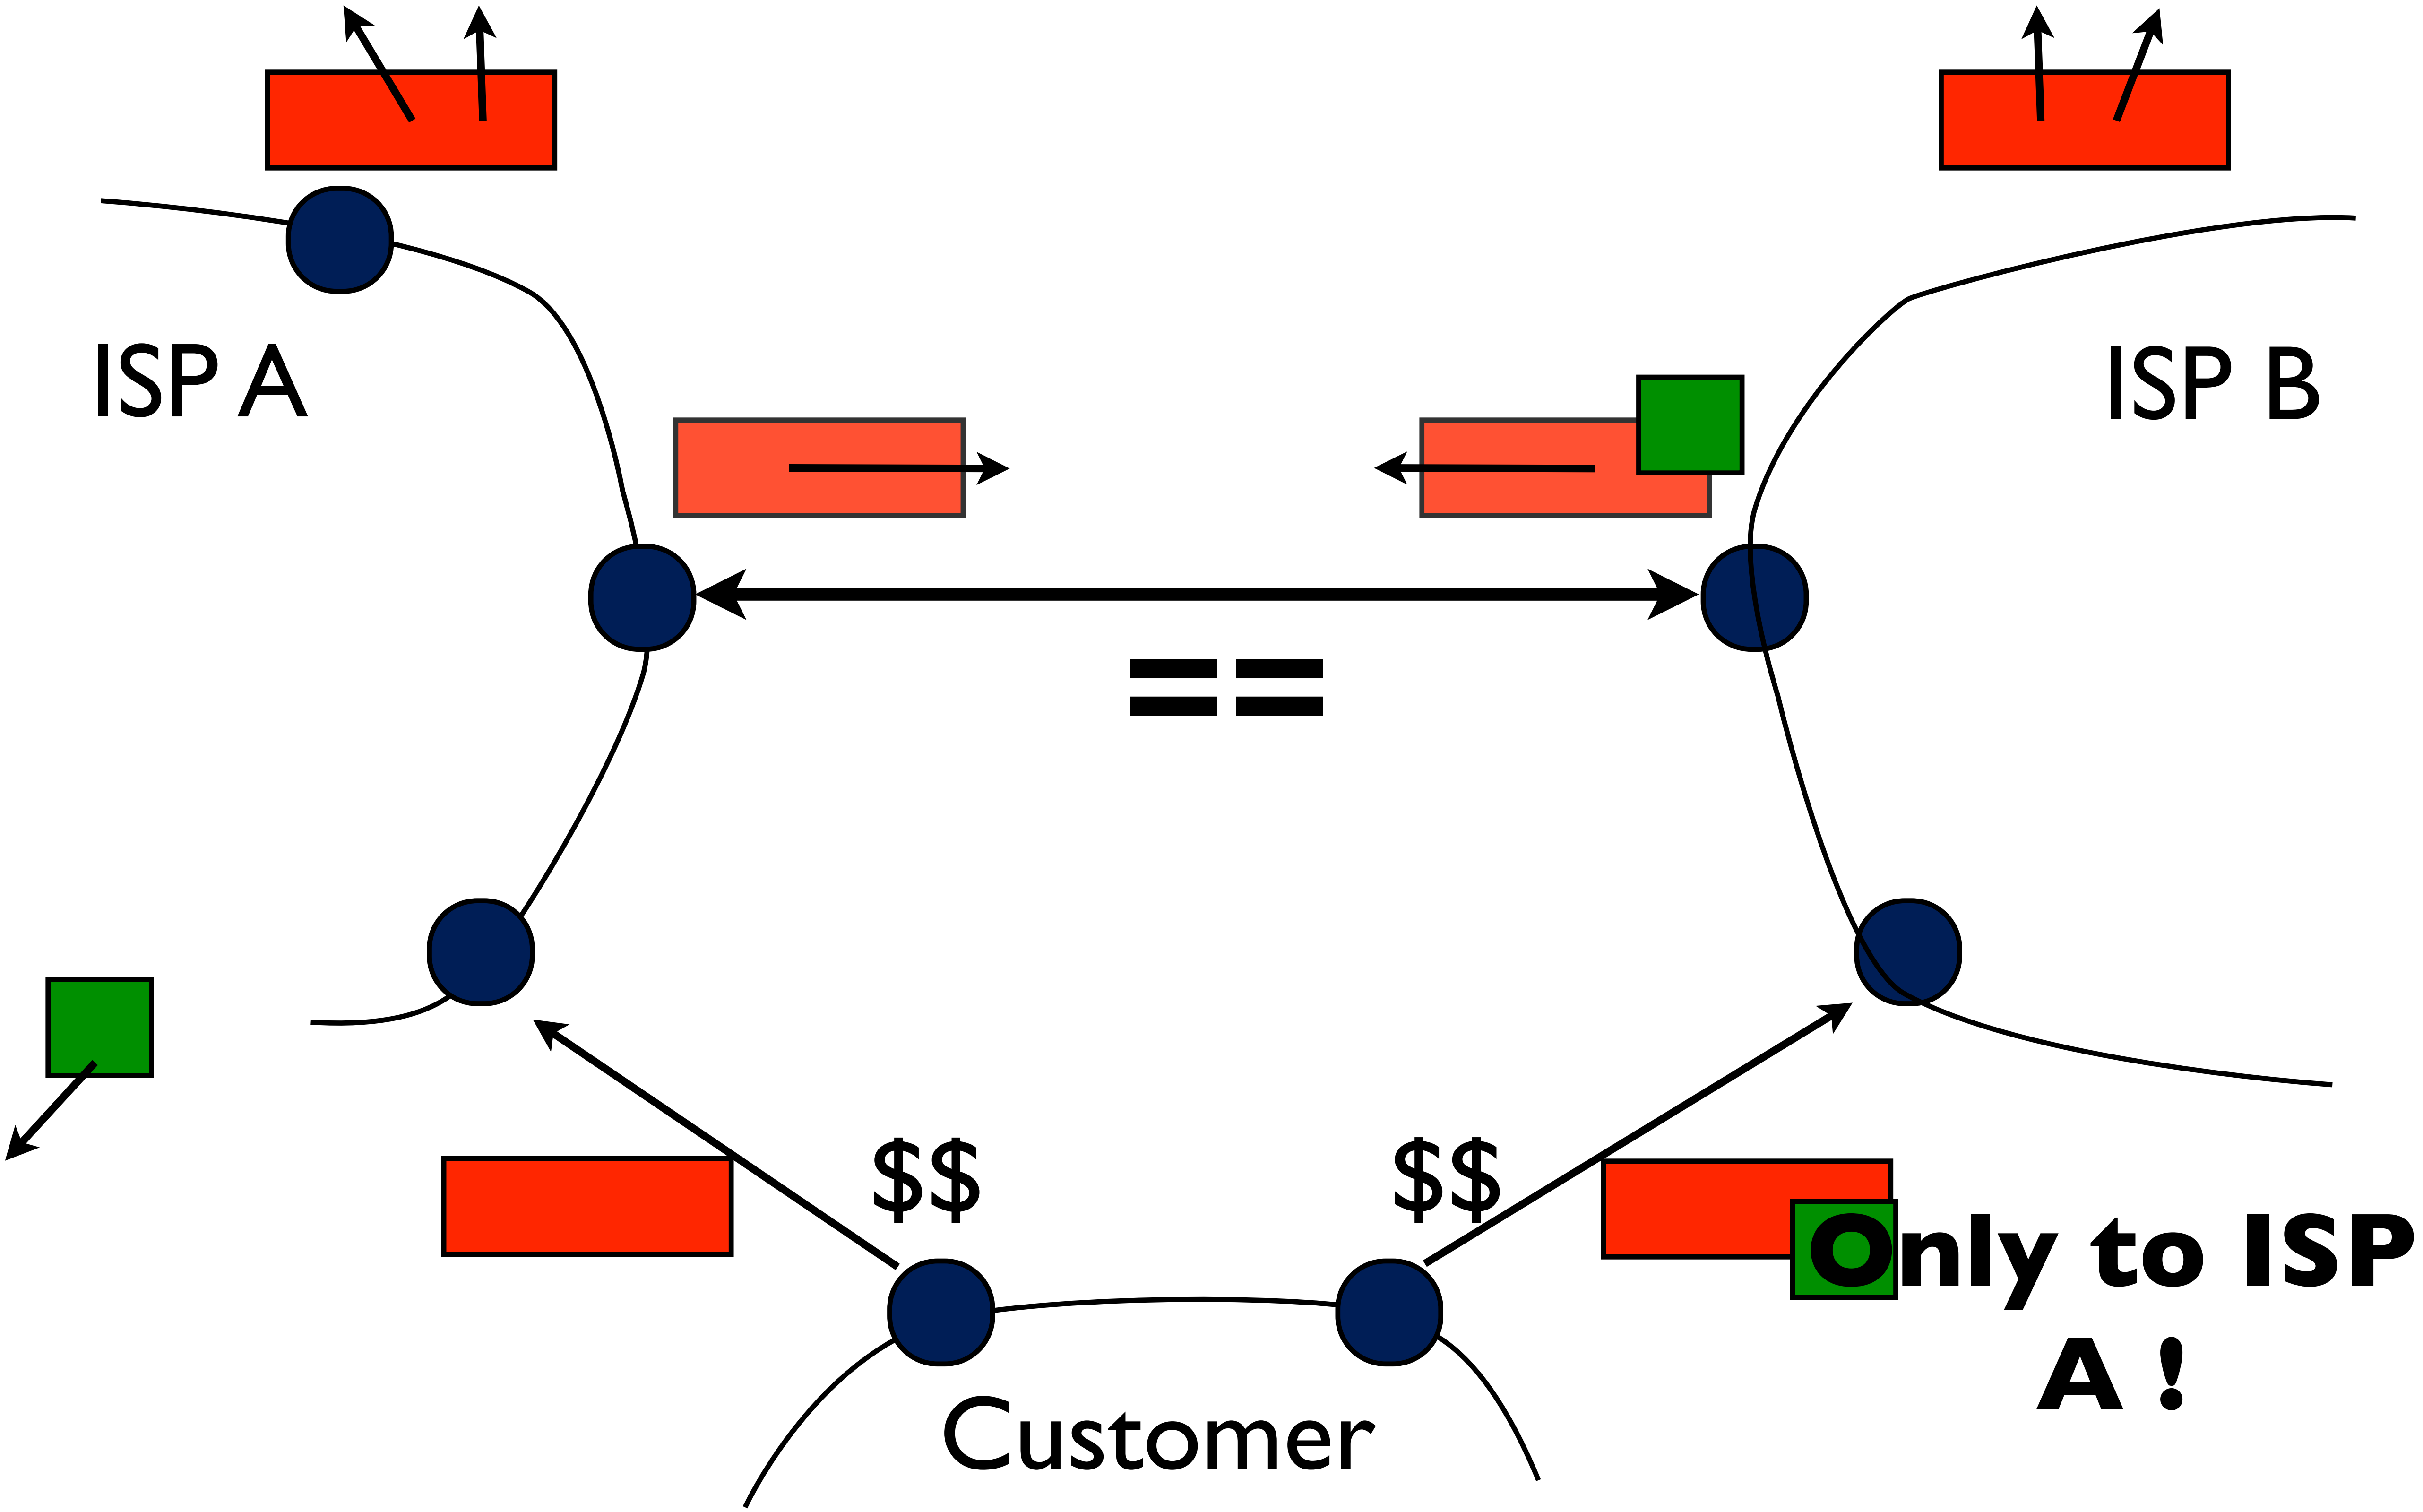

Only to ISP A !

Scope the advertisement of the more specific

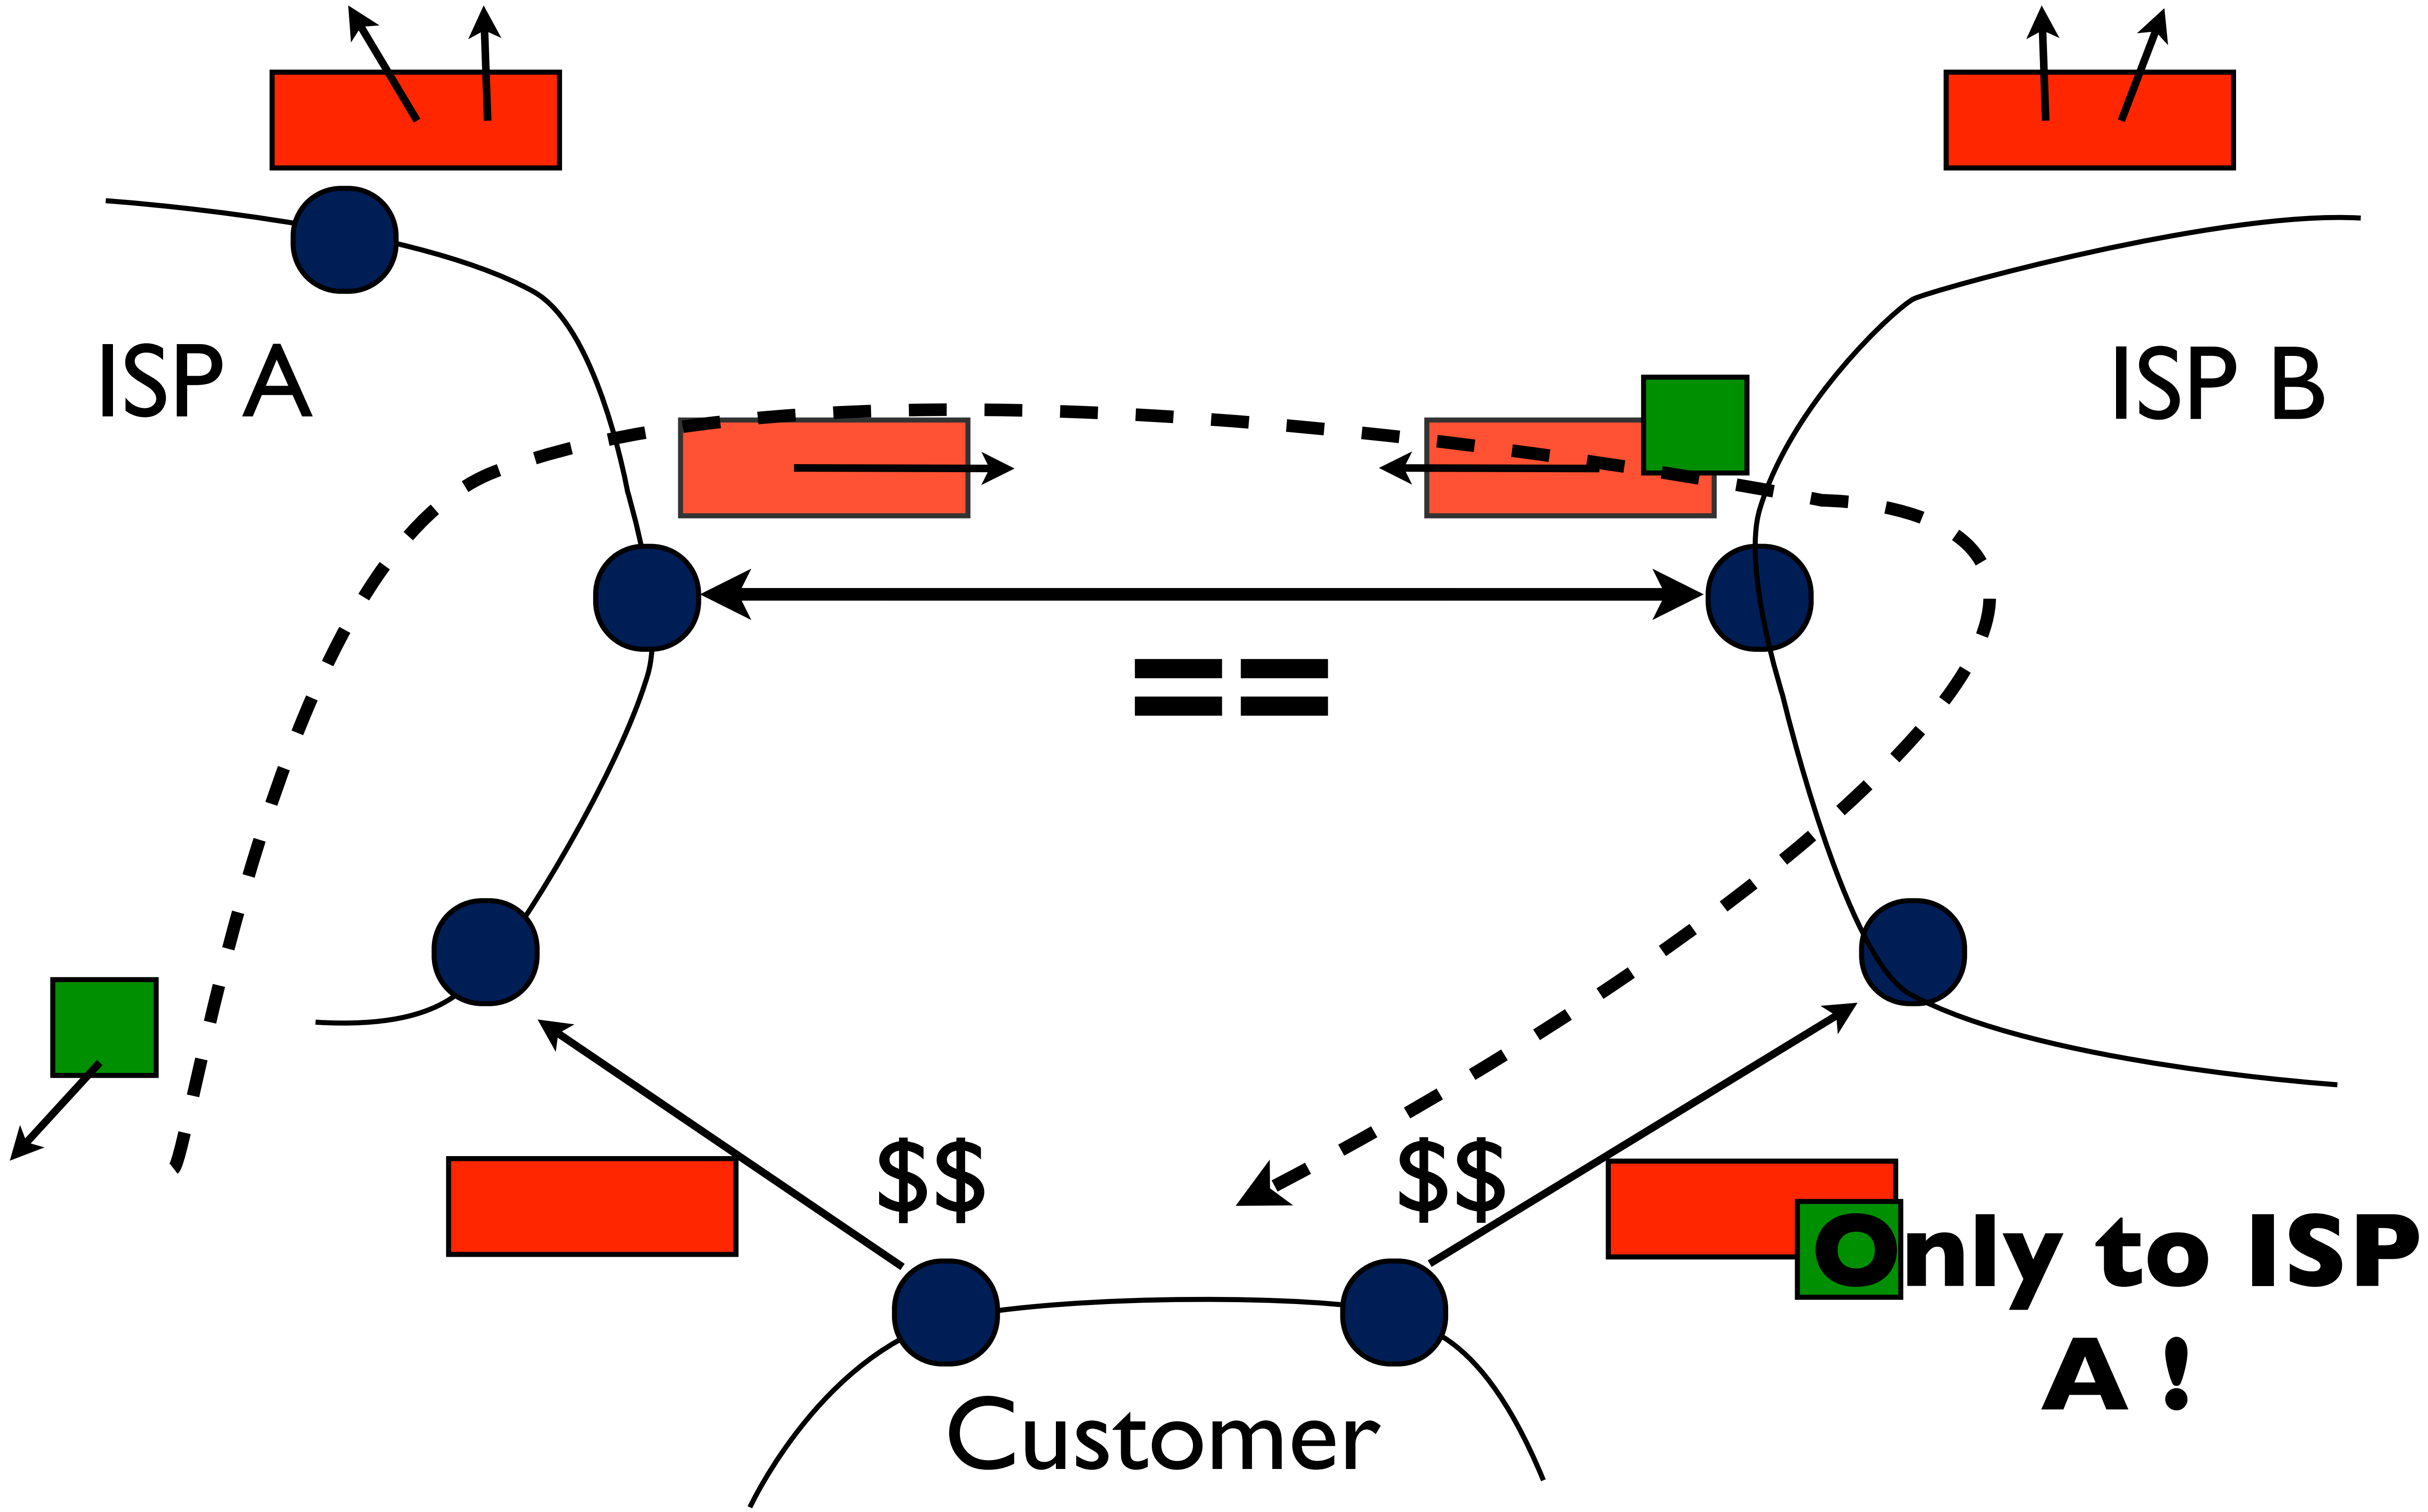

“New paths” through your network

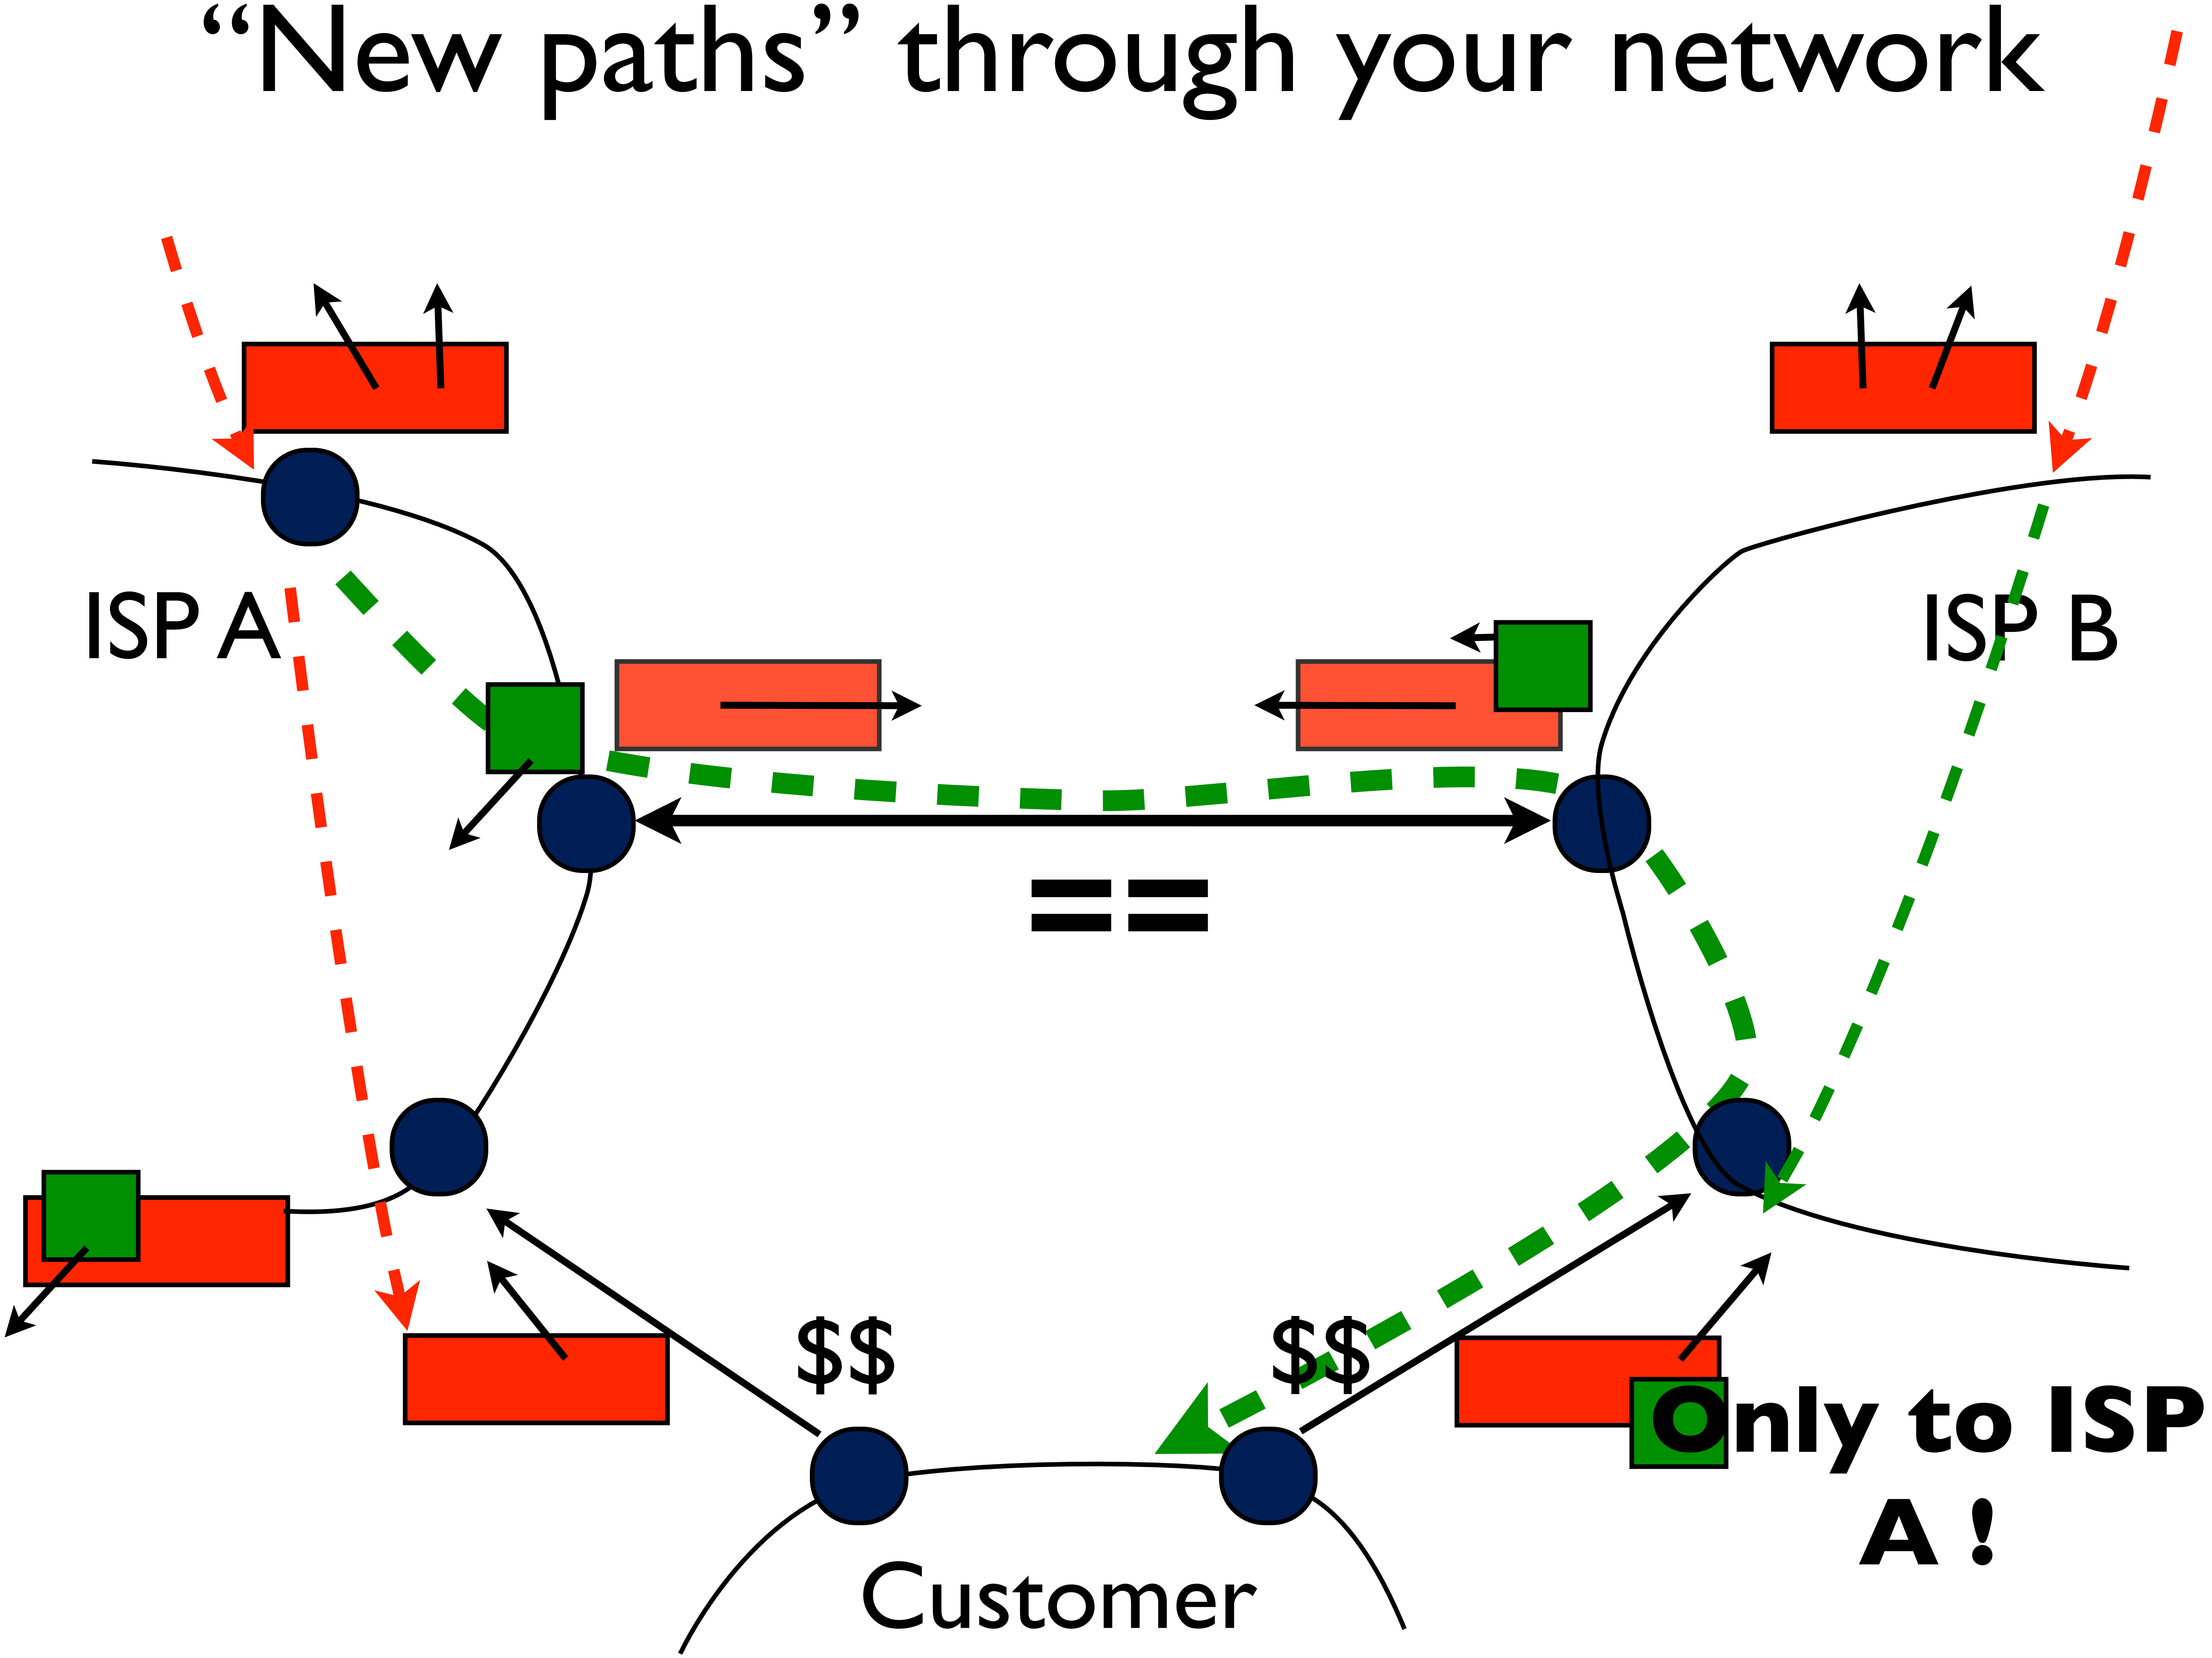

This is annoying

- Policies can be violated, again
- Your flexible routing service can turn **you** into a transit thief when misused by **your** customers
- “Nothing breaks” when the violation takes place
- Ex. : Just consider the Tier-1 clique...

So what can you do ?

- Forward differently
- Filter-out / Drop
- **Monitor !**

Thank you!