Advanced Computer Networks

More on BGP

Jianping Pan
Summer 2007

Review: BGP

- Border Gateway Protocol
  - path vector routing
    - prefix: AS-path
  - policy-based routing
    - import/export filters, local preference, MED
  - goal: stable and scalable
    - MRAI
    - route flap dampening
- Reality check

BGP entries

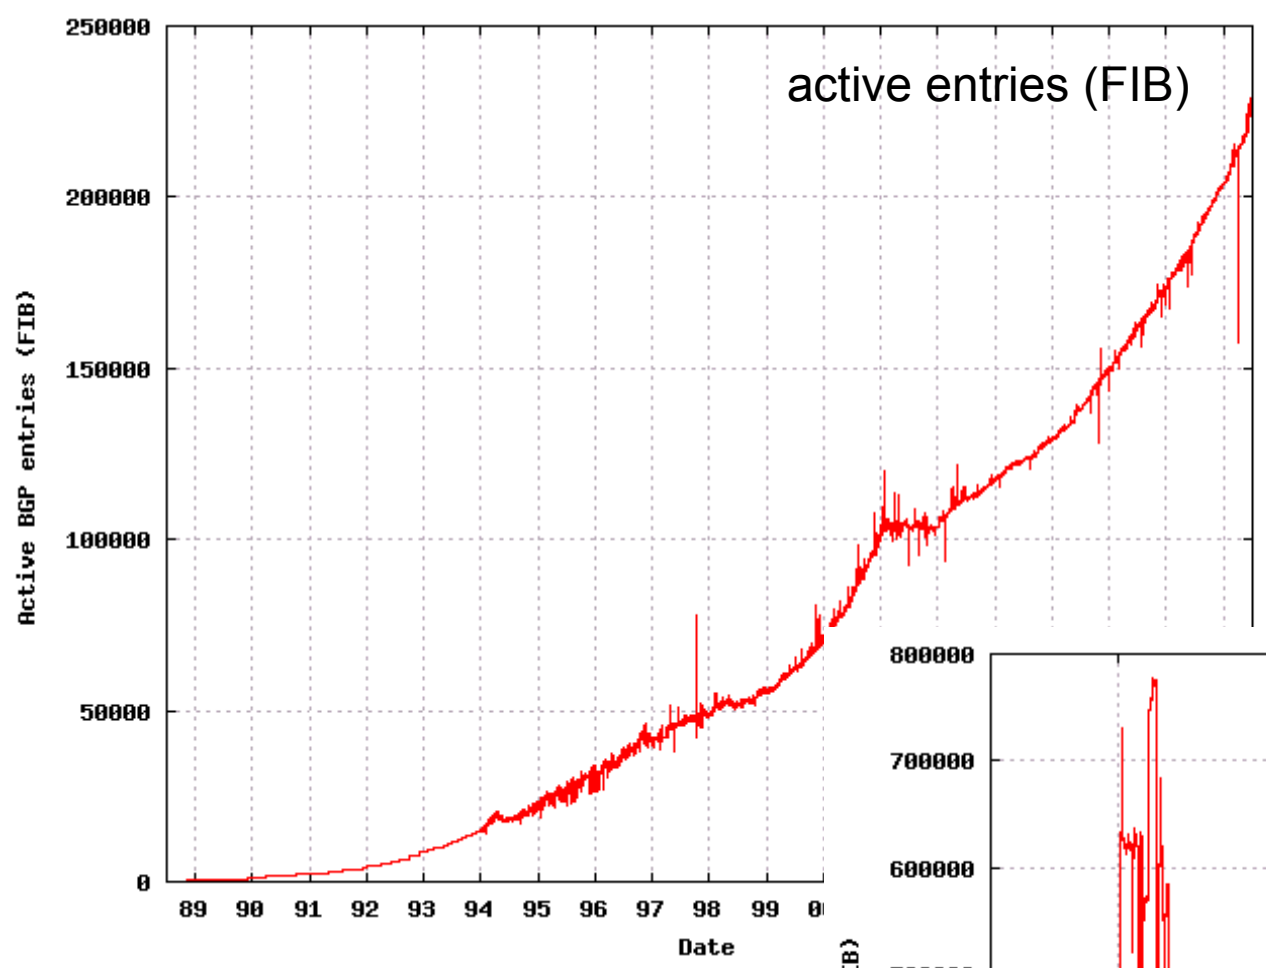

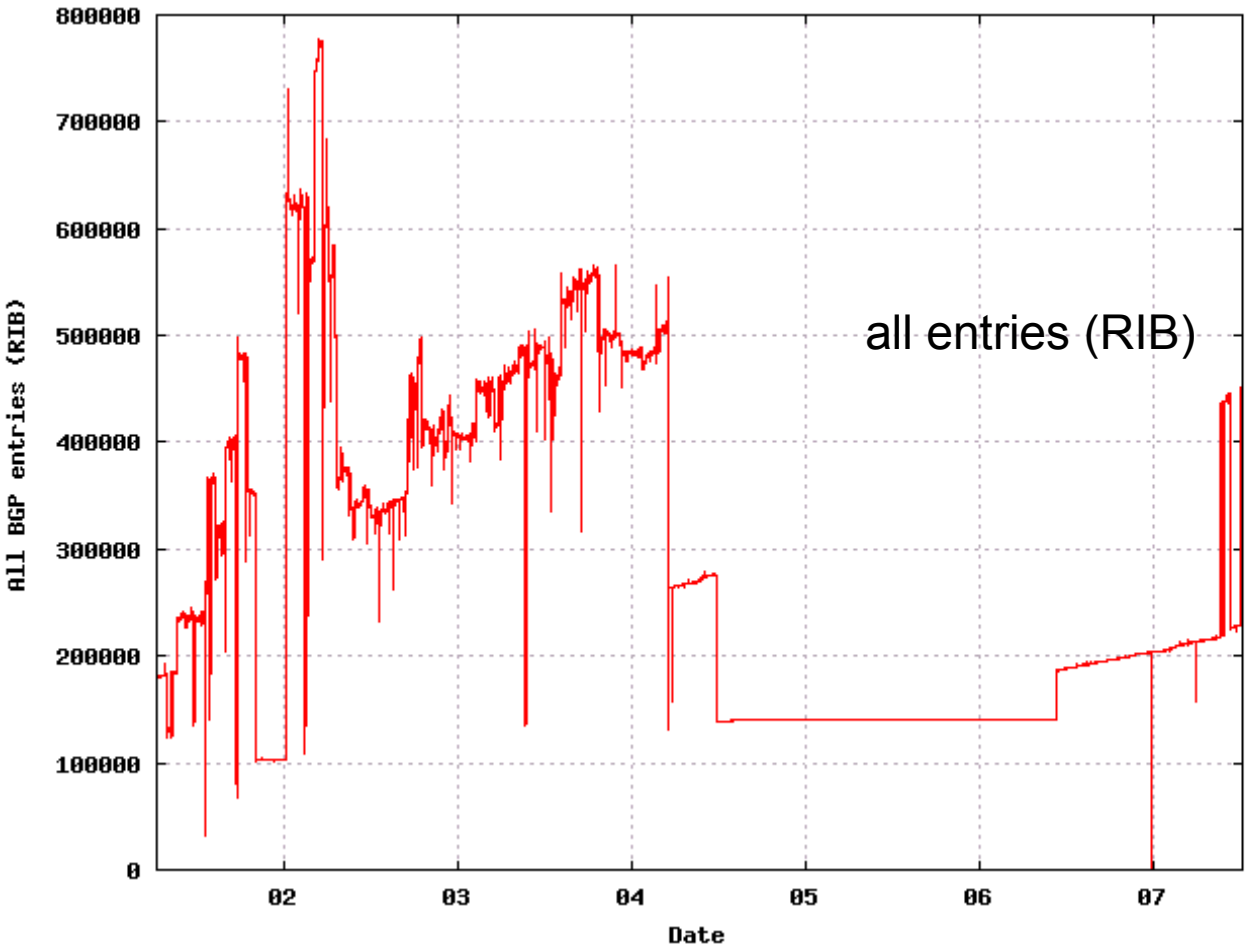

http://bgp.potaroo.net/

Customer-provider relationship

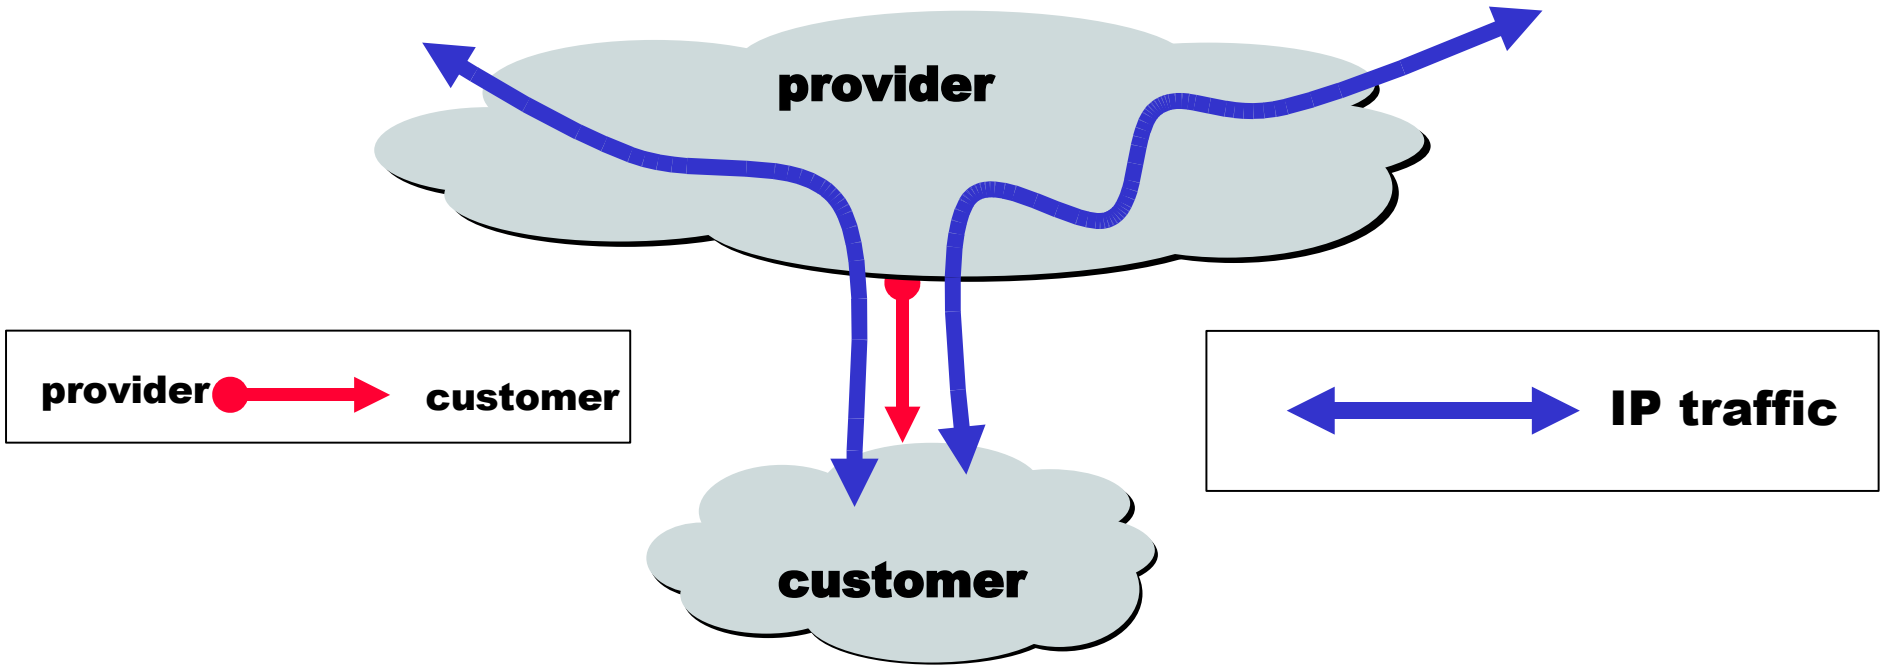

Customer pays provider for access to the Internet

http://www.cl.cam.ac.uk/~tgg22/talks/BGP_TUTORIAL_ICNP_2002.ppt

Peer-to-peer relationship

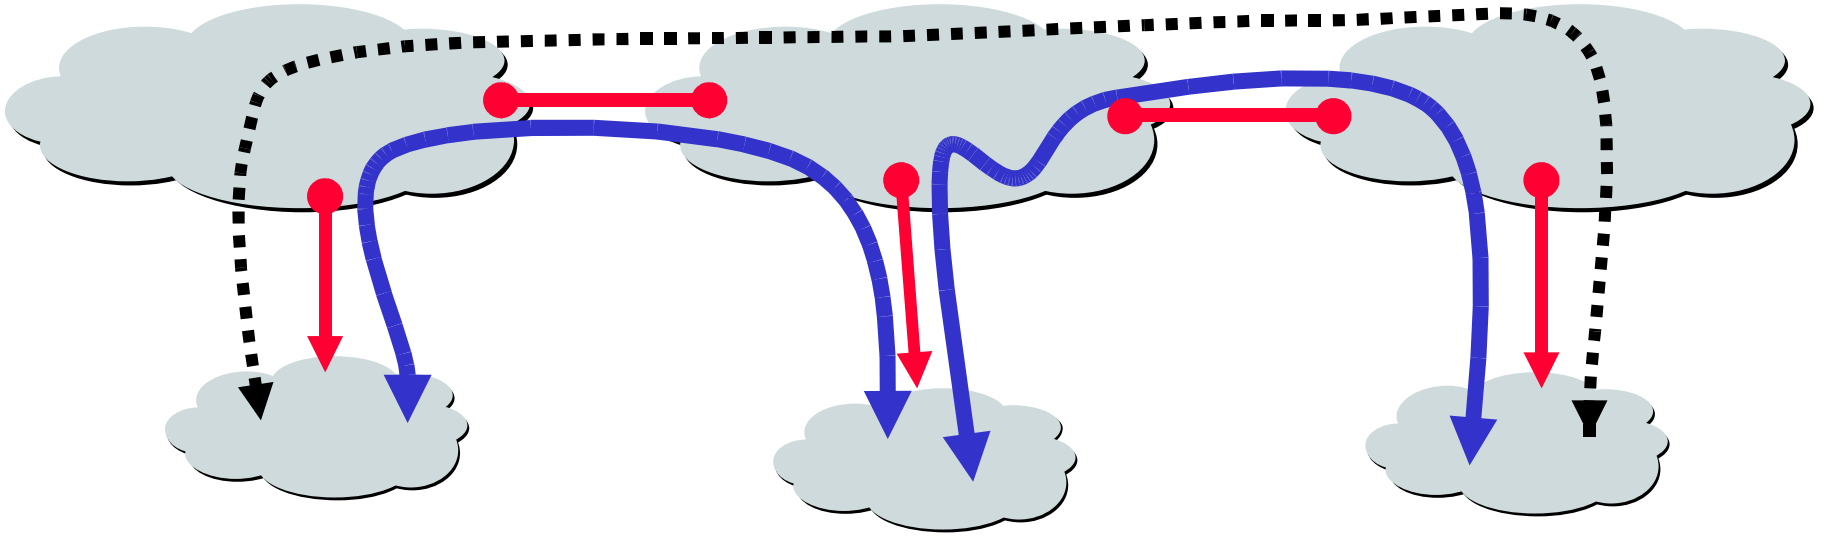

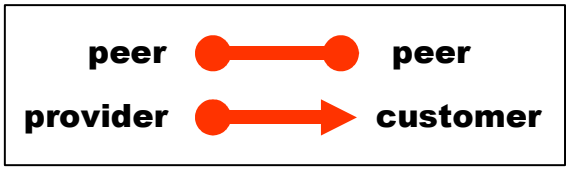

traffic allowed

traffic NOT allowed

Peers provide transit between their respective customers

Peers do not provide transit between peers

Peers (often) do not exchange $$$

Tiered structures

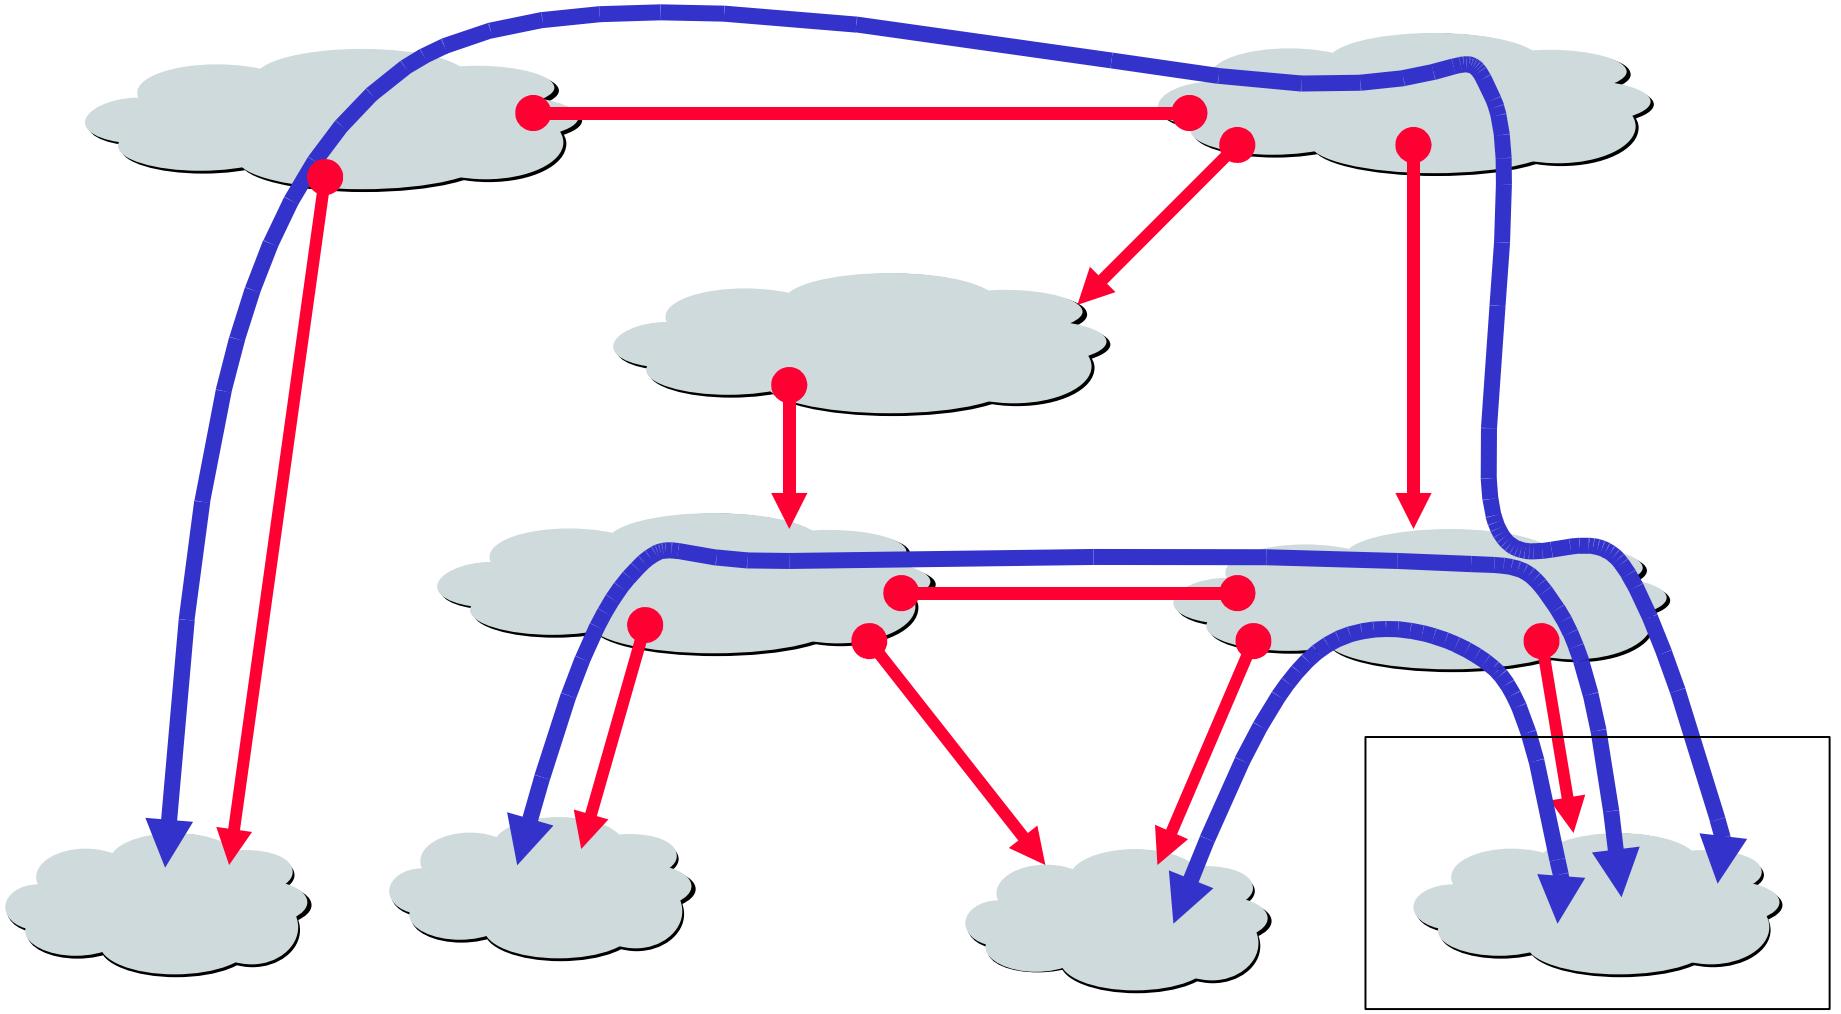

Peering also allows connectivity between the customers of "Tier 1" providers.

eBGP and iBGP

- **External Neighbor (eBGP) in a different Autonomous Systems**
- **Internal Neighbor (iBGP) in the same Autonomous System**

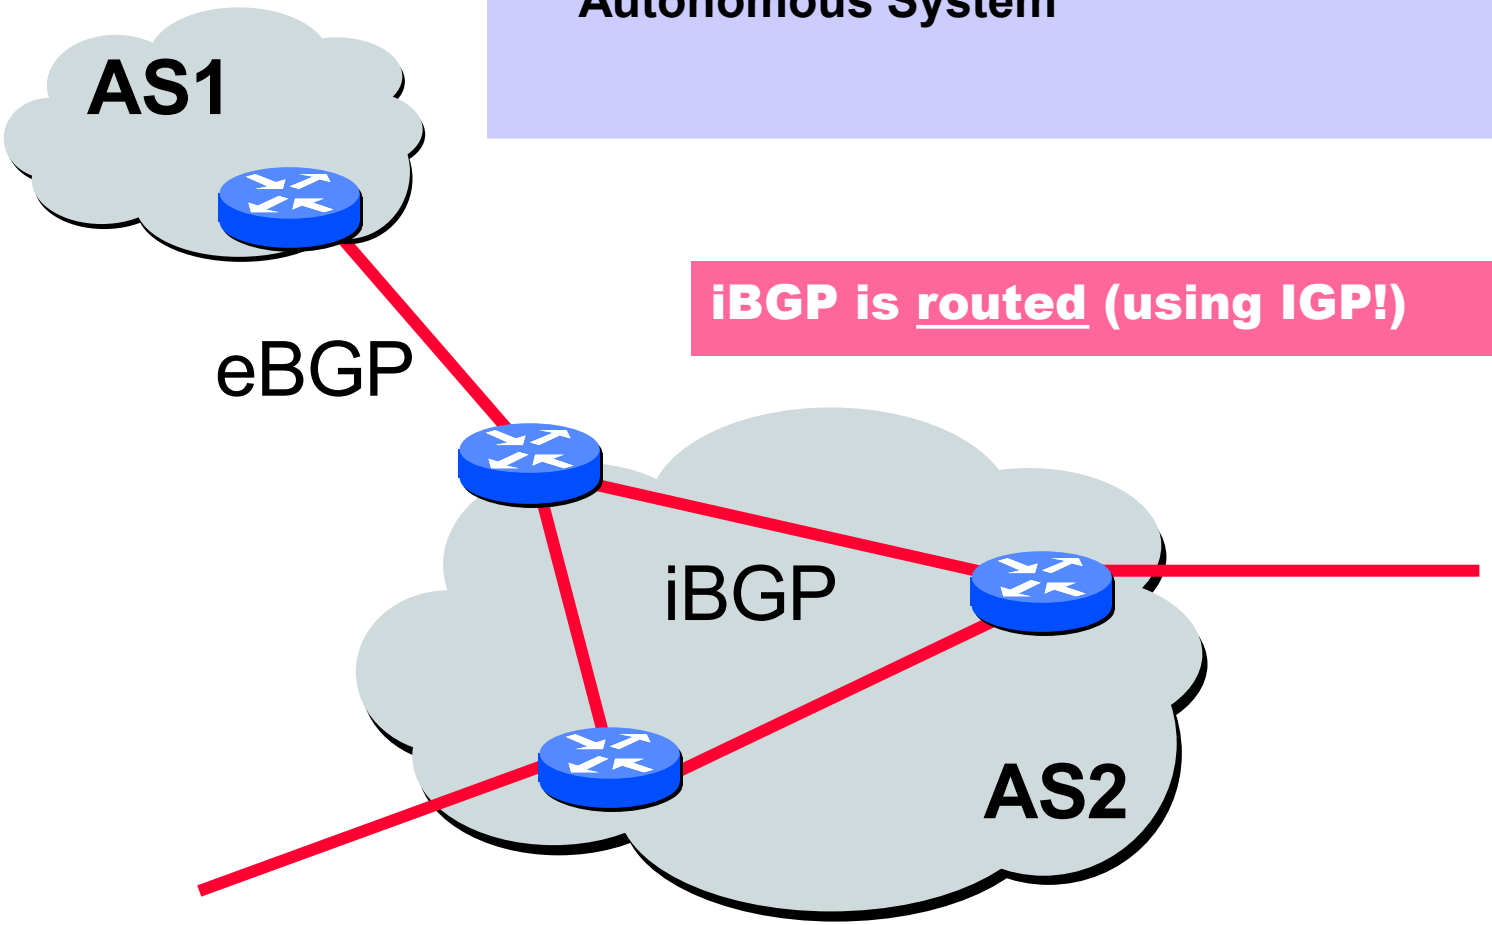

iBGP is routed (using IGP!)

iBGP mesh

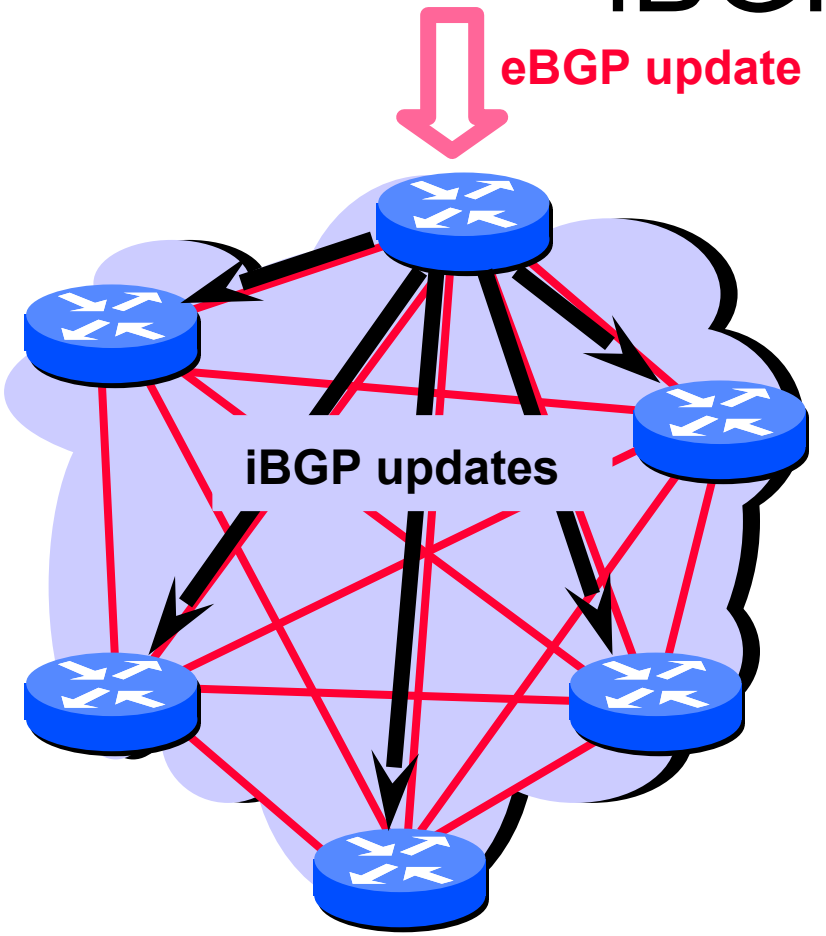

- N border routers means N(N-1)/2 peering sessions
- Each router must have N-1 iBGP sessions configured
- The addition a single iBGP speaker requires configuration changes to all other iBGP speakers
- Size of iBGP routing table can be order N larger than number of best routes (remember alternate routes!)
- Each router has to listen to update noise from each neighbor

Currently four solutions:
(0) Buy bigger routers!
(3) Break AS into smaller ASes
(4) BGP Route reflectors
(5) BGP confederations

Route reflector

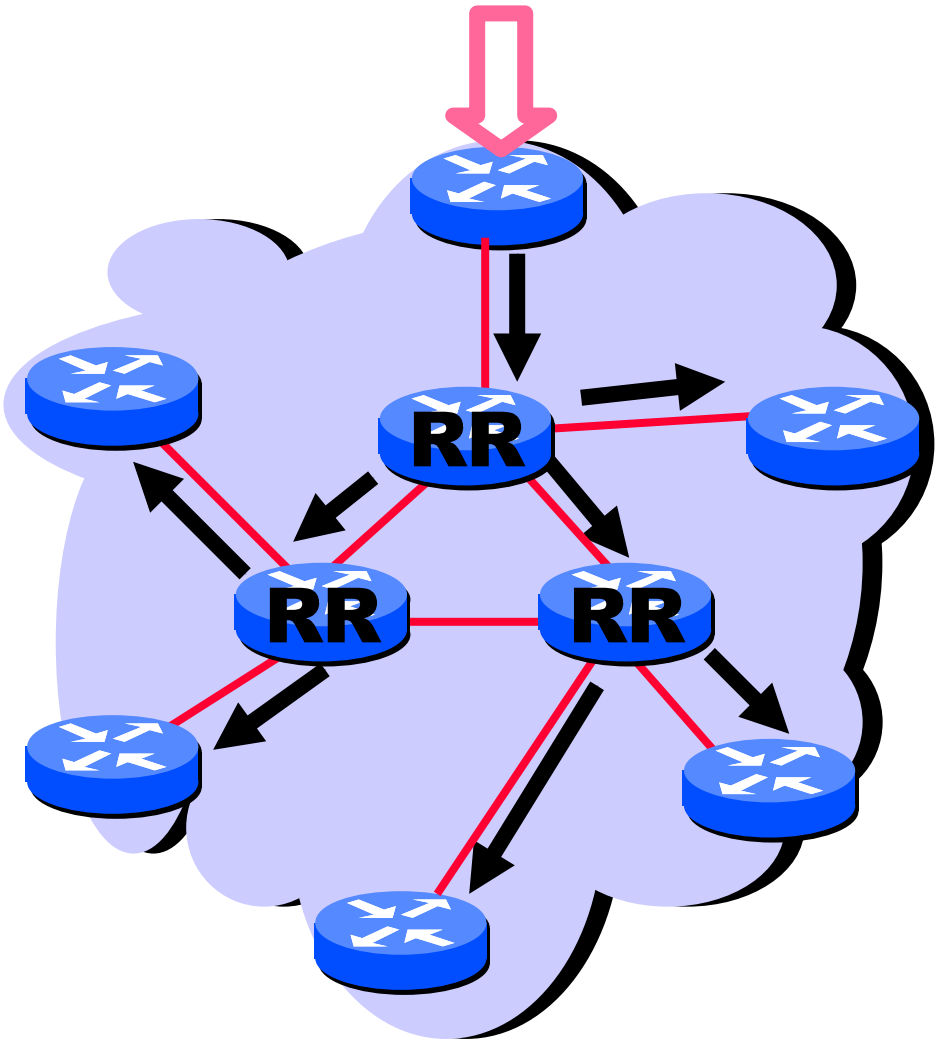

- **Route reflectors can pass on iBGP updates to clients**
- **Each RR passes along ONLY best routes**
- **ORIGINATOR_ID and CLUSTER_LIST attributes are needed to avoid loops**

AS confederation

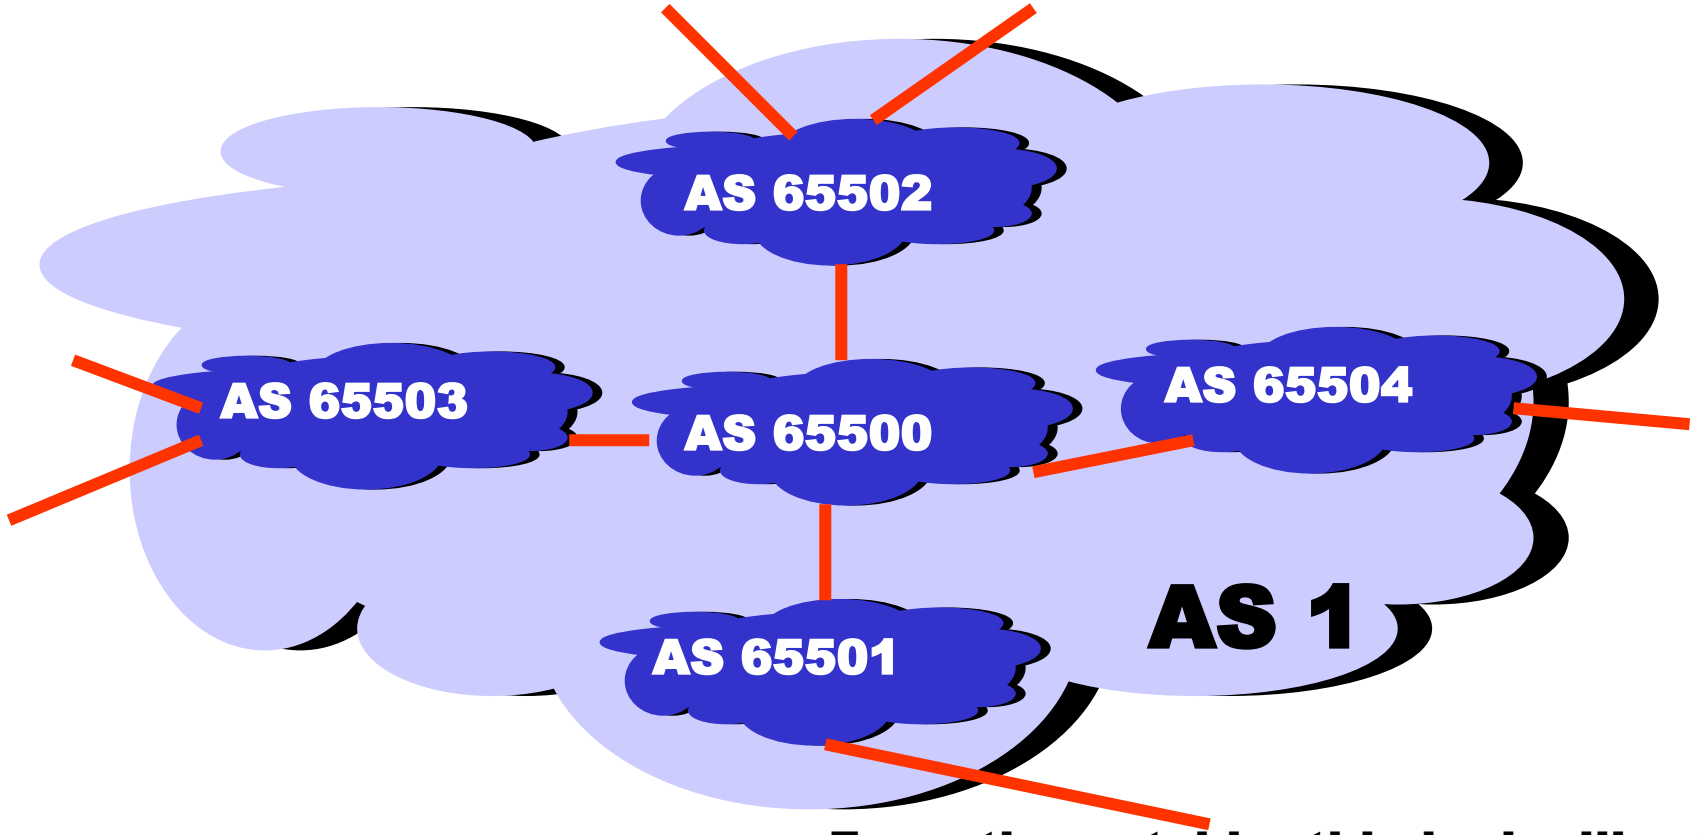

From the outside, this looks like AS 1

Confederation eBGP (between member ASes) preserves LOCAL_PREF, MED, and BGP NEXTHOP.

Path selection

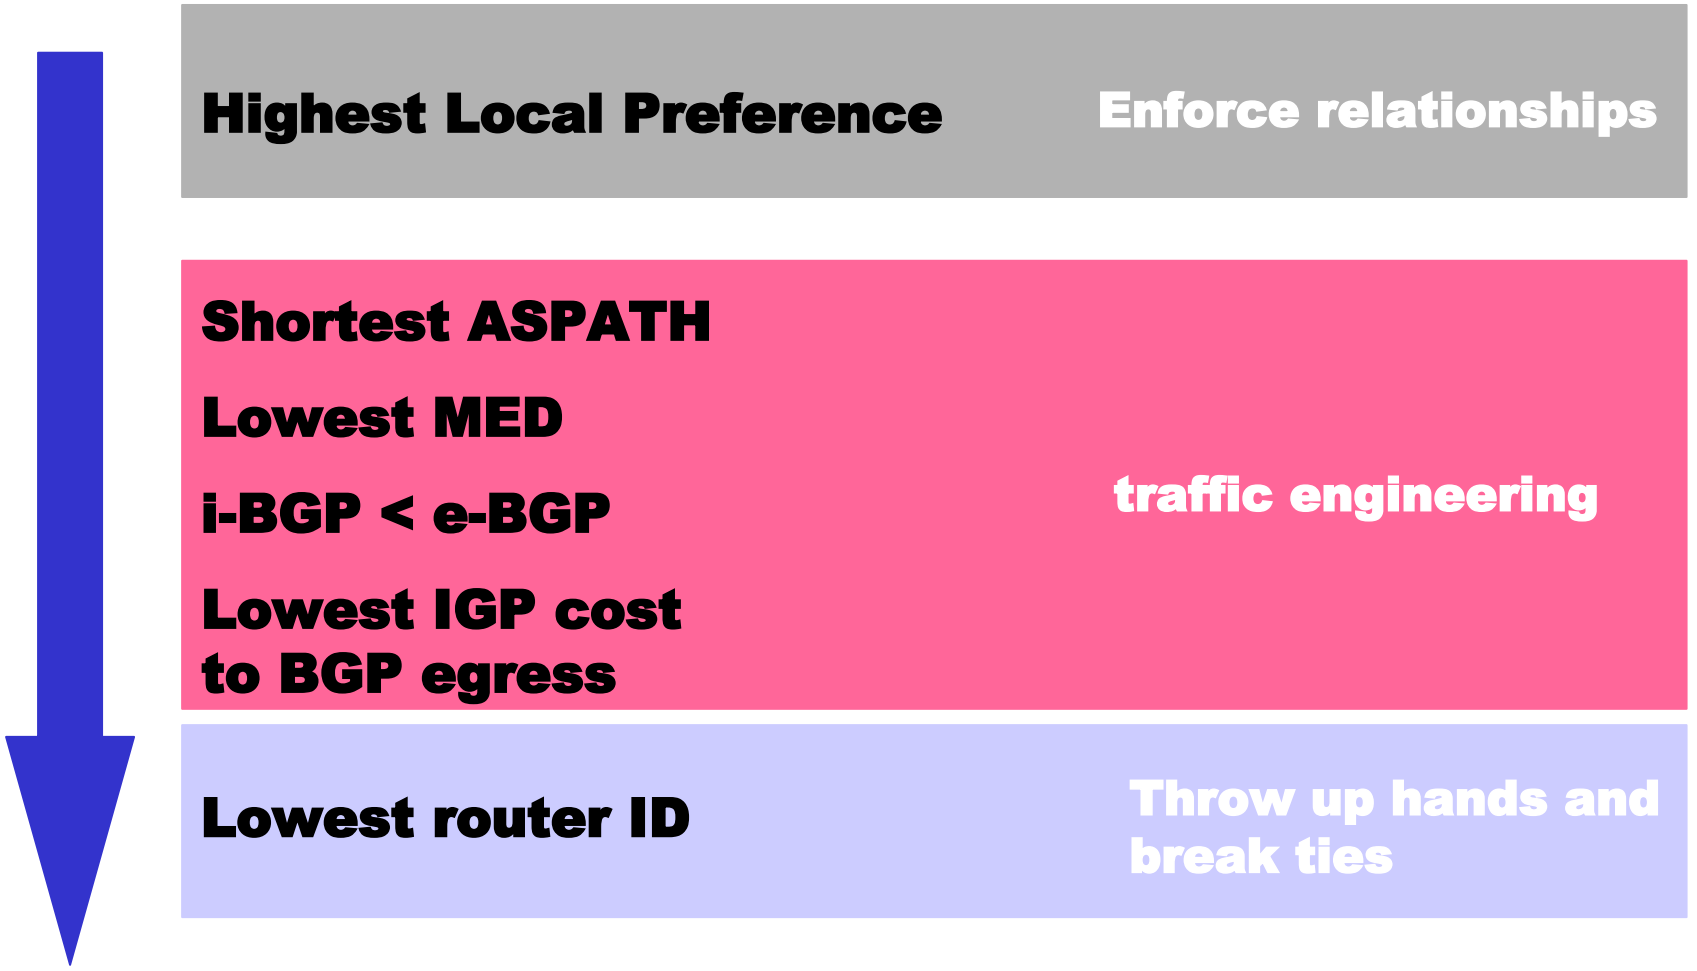

Local preference

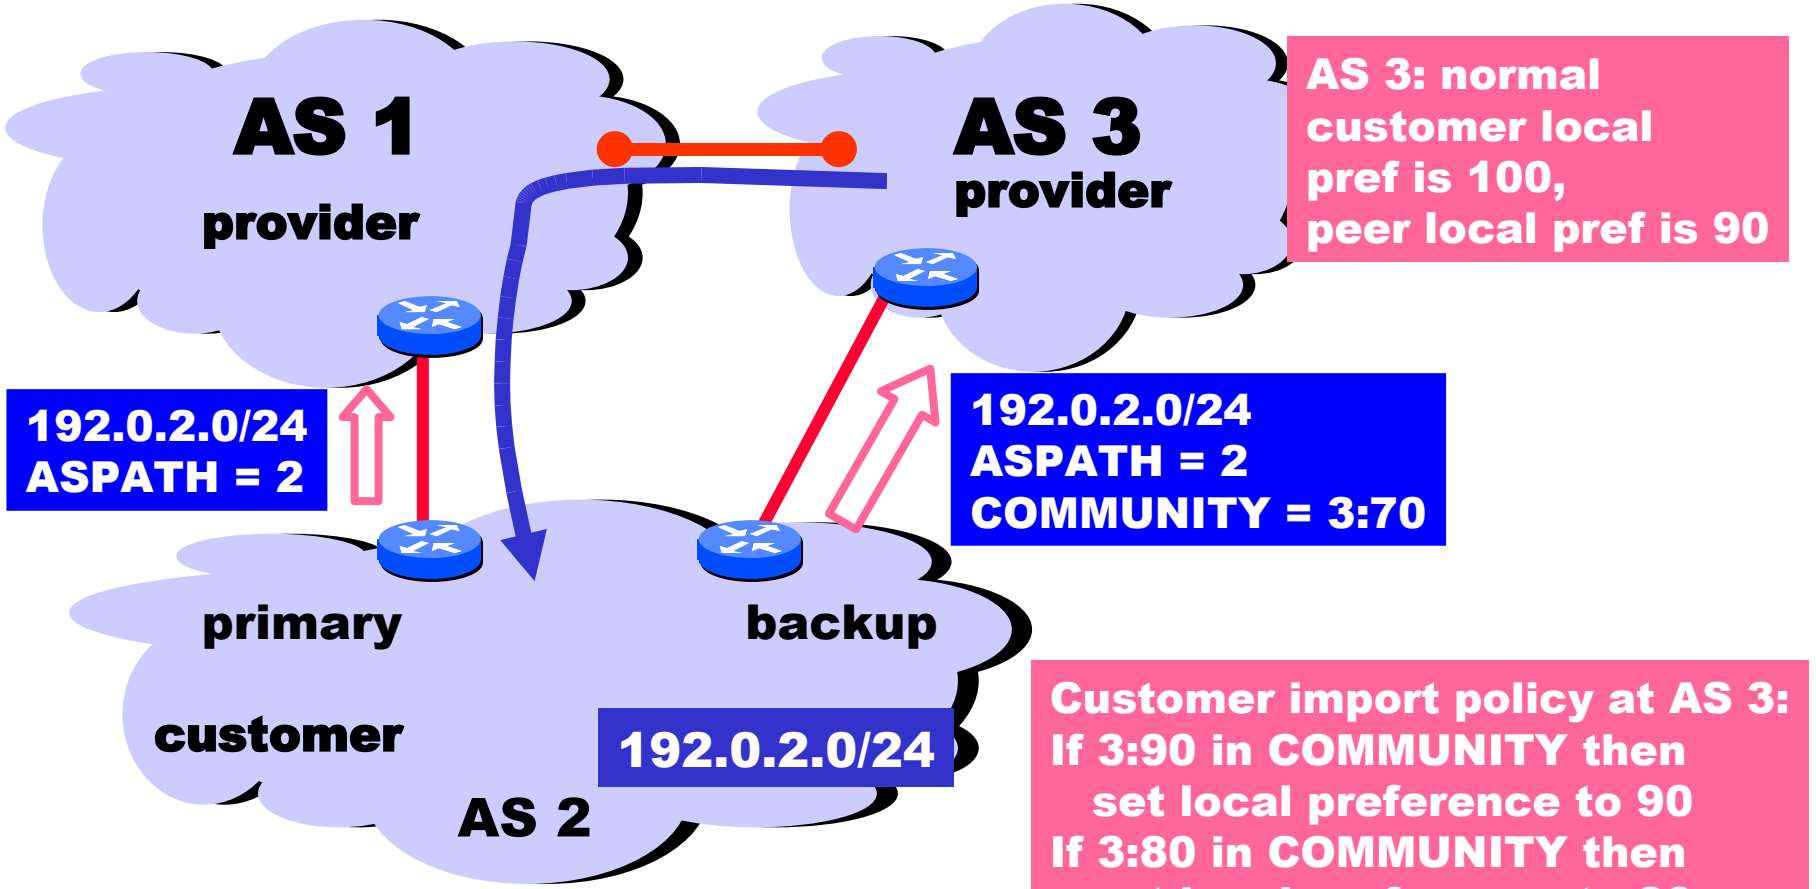

Multi-exit discriminator

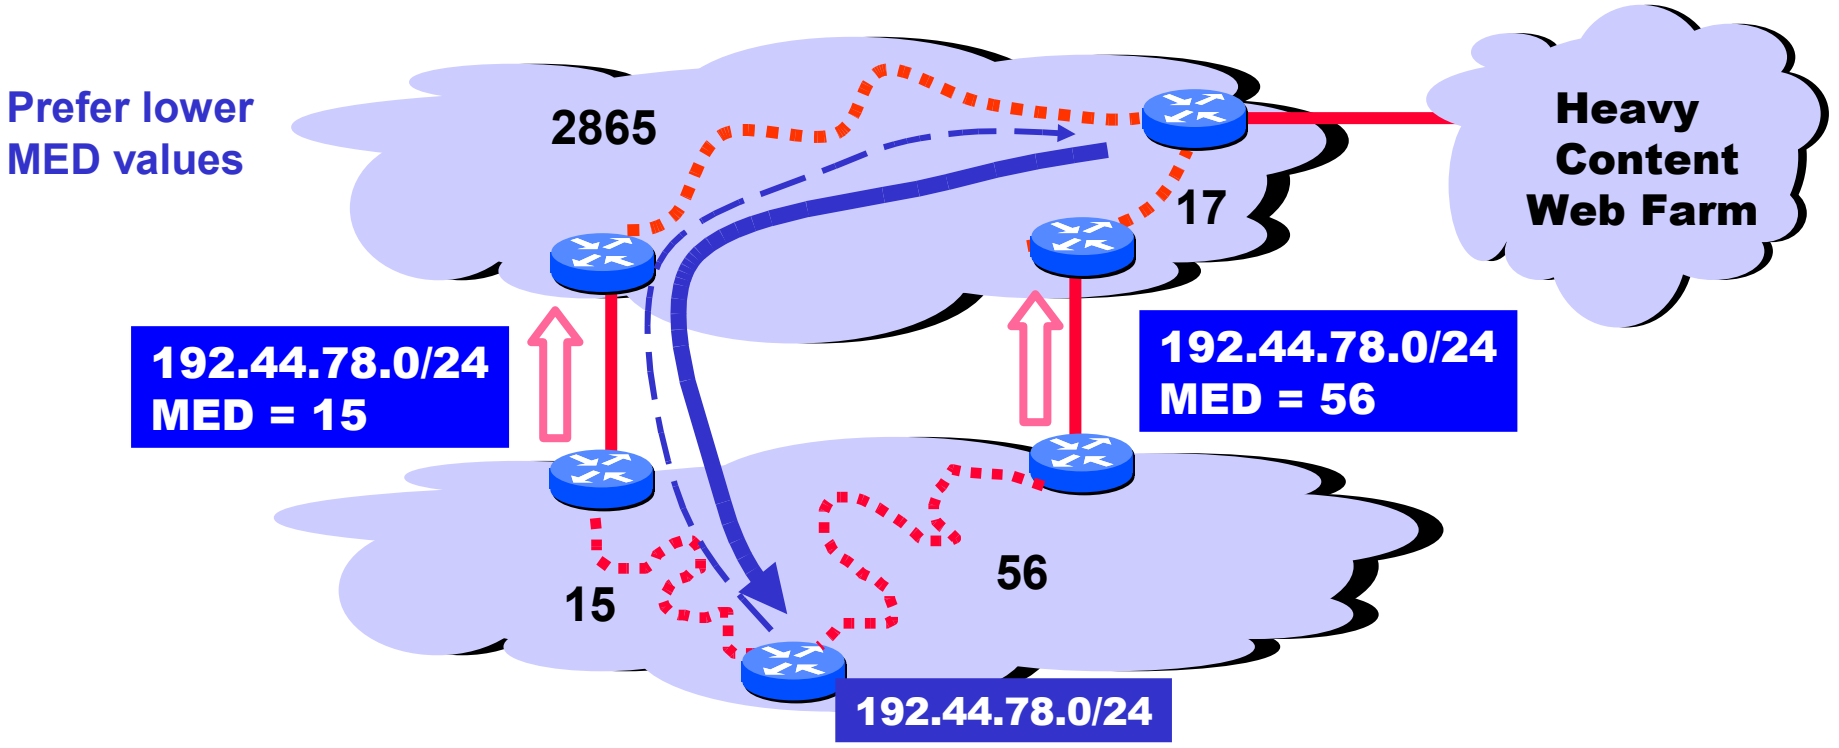

Import routes

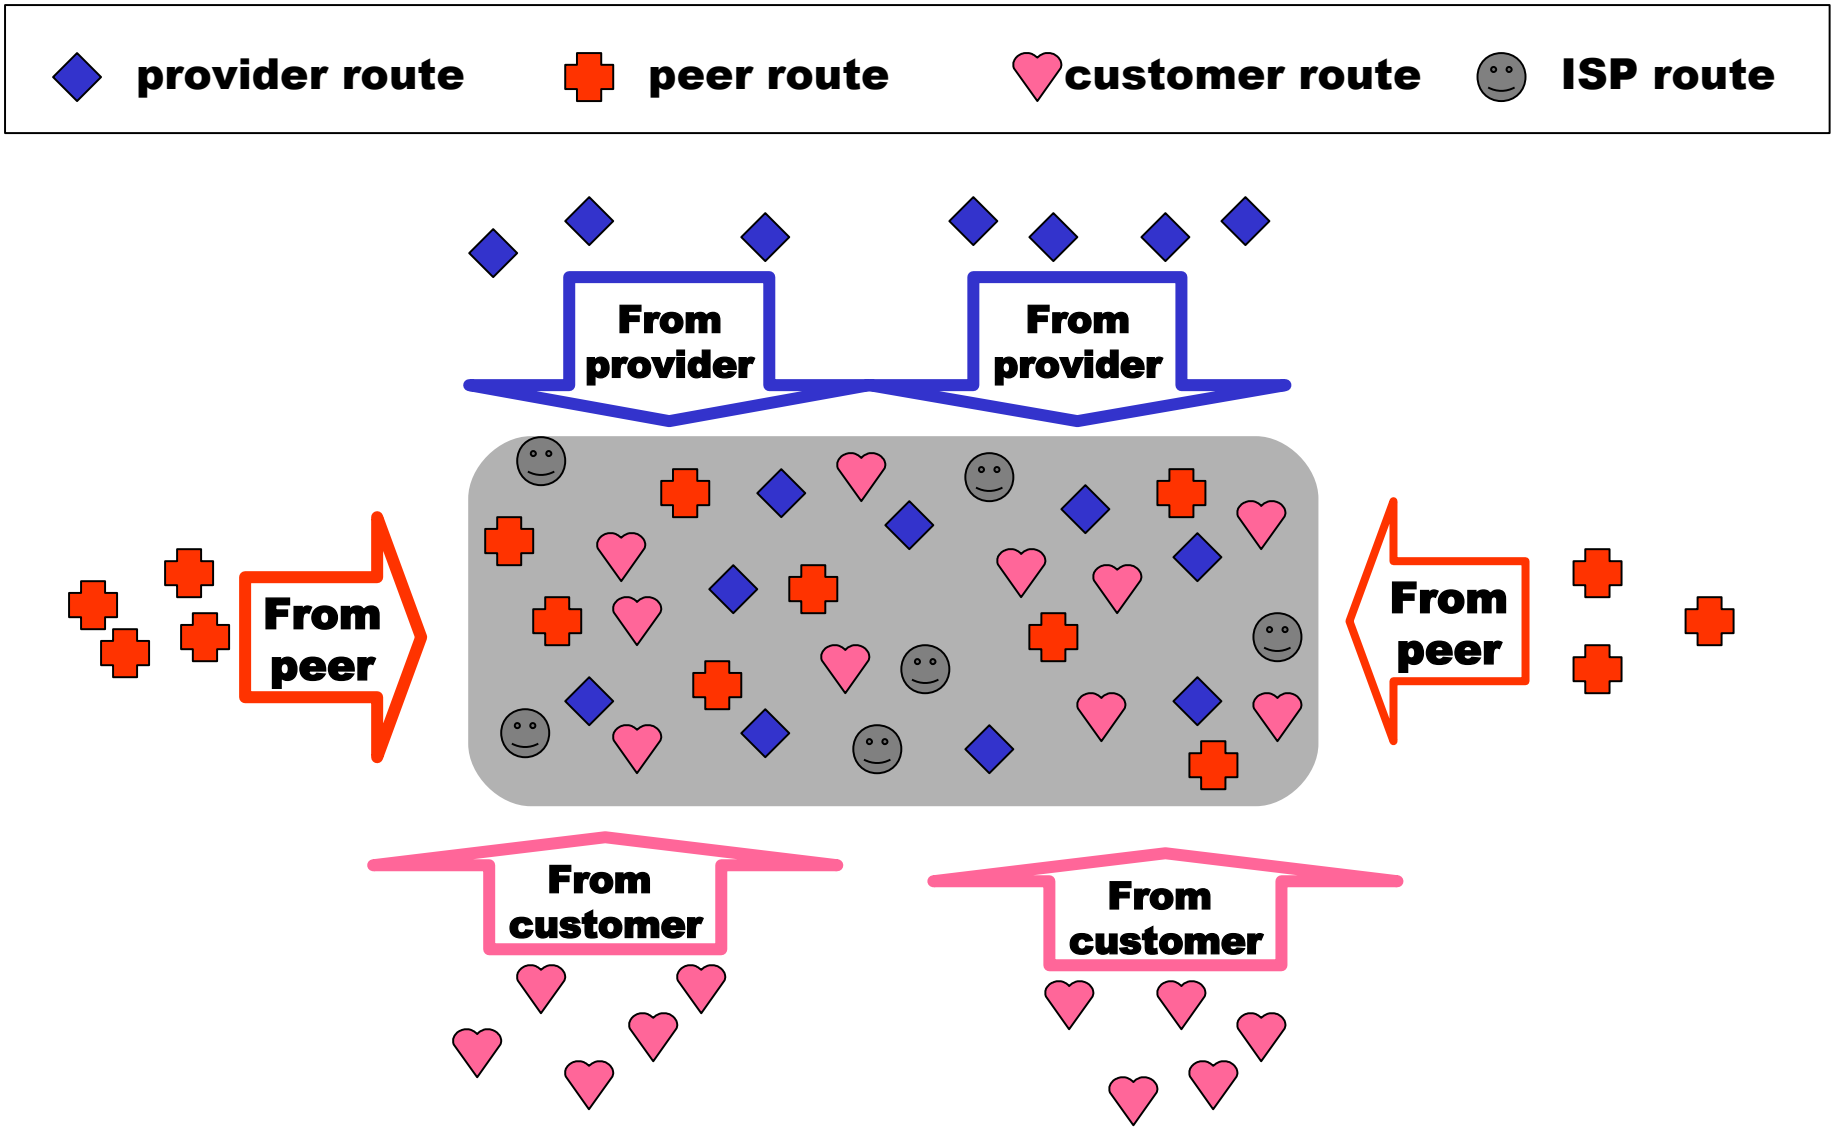

Export routes

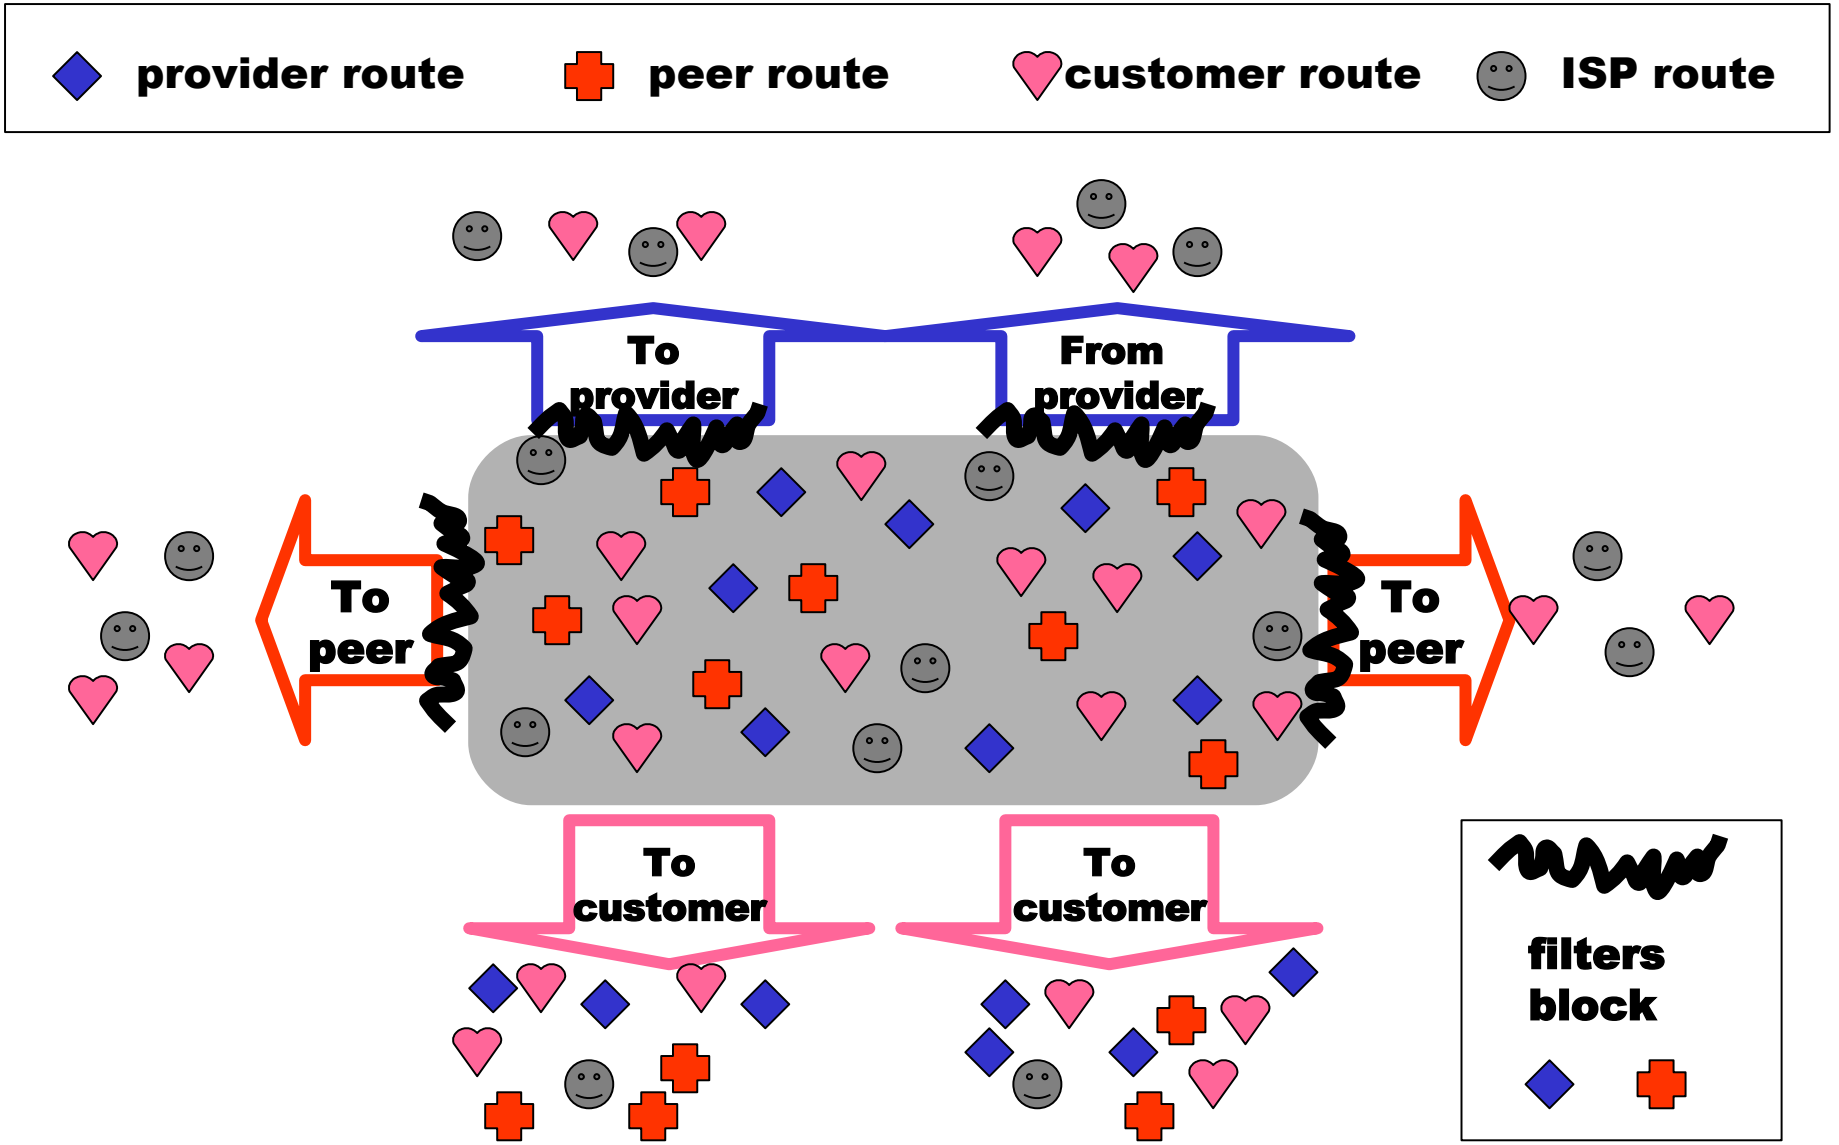

Route preferences

- Prefer customer's routes first
- Then peer's routes
- and finally provider's routes
- Why?

Student presentation

- Justyn Bussey
  - [GR00] Lixin Gao and Jennifer Rexford, "Stable Internet Routing Without Global Coordination". In Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems. 2000.

Further discussion

- Stable path problem

This lecture

- More on BGP
  - stability and scalability in practice
- Explore further
  - http://bgp.potaroo.net/
  - http://www.routeviews.org/
  - http://www.cl.cam.ac.uk/~tgg22/interdomain/

Next lectures

- Traffic management
  - [CSZ92] D. Clark and S. Shenker and L. Zhang, "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanism". In Proceedings of SIGCOMM '92, Baltimore, Maryland, Aug, 1992, pp 14-26. [IntServ]
  - [SSZ98] I. Stoica , S. Shenker , and H. Zhang , "Core -Stateless Fair Queueing: Achieving Approximately Fair Allocations in High Speed Networks" , Proc. ACM SIGCOMM , Vancouver, Canada, September 1998. [CSFQ]