BGP Attributes

Value	Code	Reference
1	ORIGIN	[RFC1771]
2	AS_PATH	[RFC1771]
3	NEXT_HOP	[RFC1771]
4	MULTI_EXIT_DISC	[RFC1771]
5	LOCAL_PREF	[RFC1771]
6	ATOMIC_AGGREGATE	[RFC1771]
7	AGGREGATOR	[RFC1771]
8	COMMUNITY	[RFC1997]
9	ORIGINATOR_ID	[RFC2796]
10	CLUSTER_LIST	[RFC2796]
11	DPA	[Chen]
12	ADVERTISER	[RFC1863]
13	RCID_PATH / CLUSTER_ID	[RFC1863]
14	MP_REACH_NLRI	[RFC2283]
15	MP_UNREACH_NLRI	[RFC2283]
16	EXTENDED COMMUNITIES	[Rosen]
...		
255	reserved for development	

From IANA: http://www.iana.org/assignments/bgp-parameters

Most important attributes

Not all attributes need to be present in every announcement

Attributes are Used to Select Best Routes

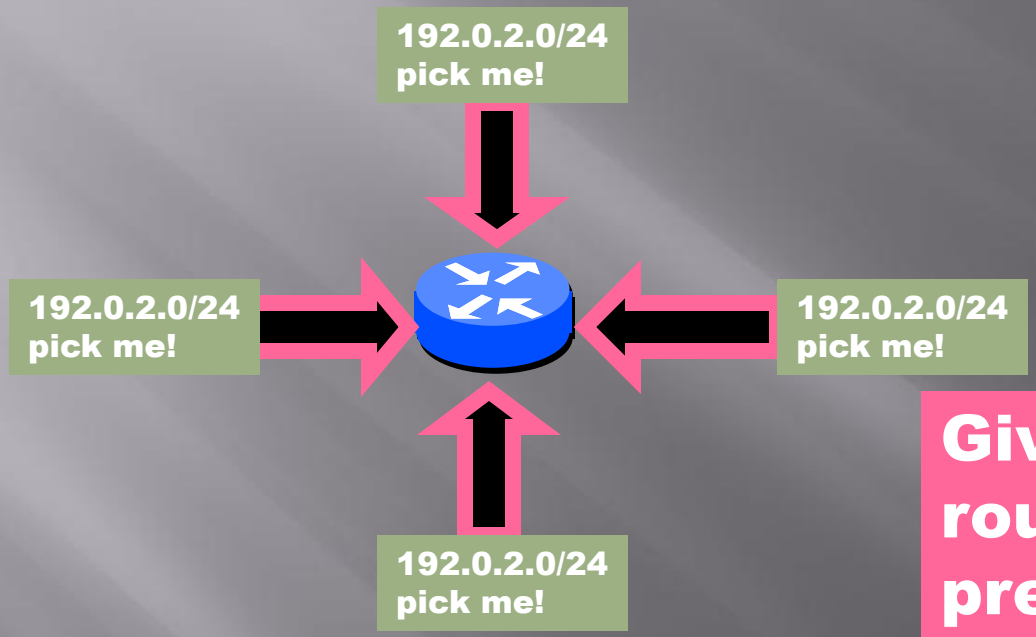

Given multiple routes to the same prefix, a BGP speaker must pick at most <u>one</u> best route (Note: it could reject them all!)

ASPATH Attribute

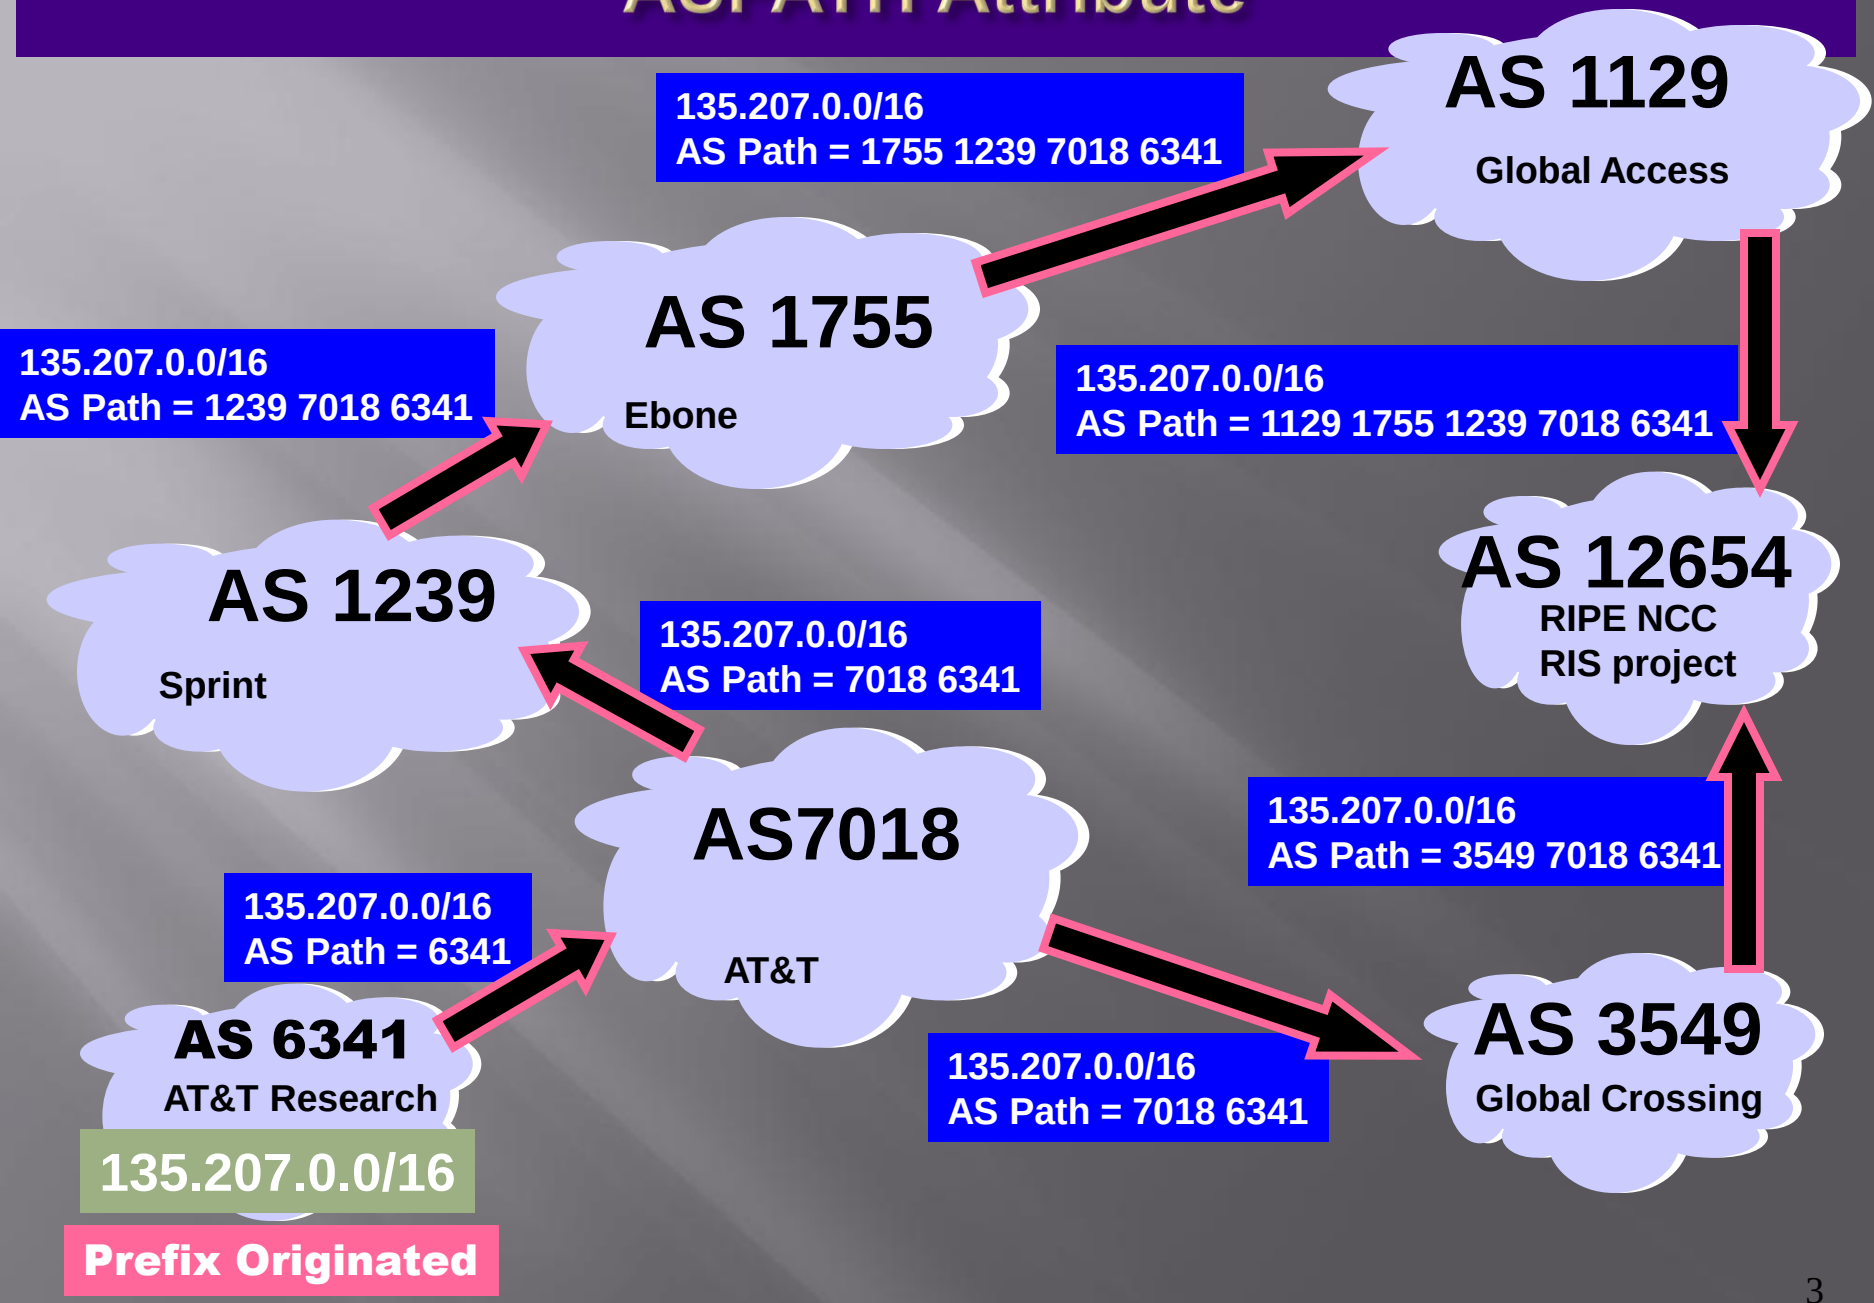

AS Graphs Do Not Show Topology!

BGP was designed to throw away information!

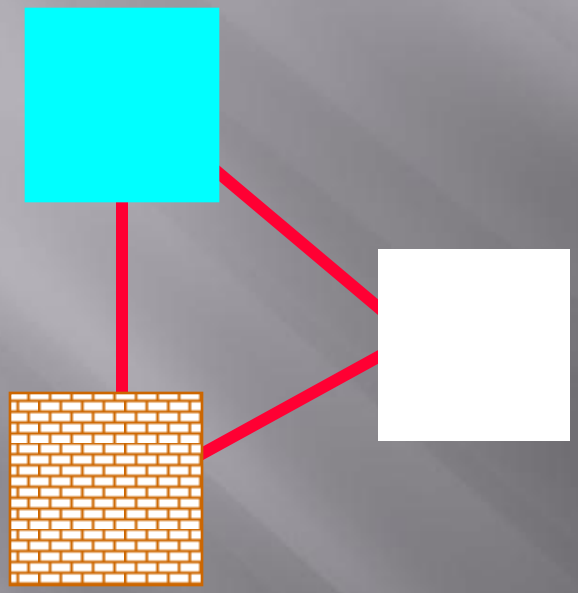

The AS graph may look like this.

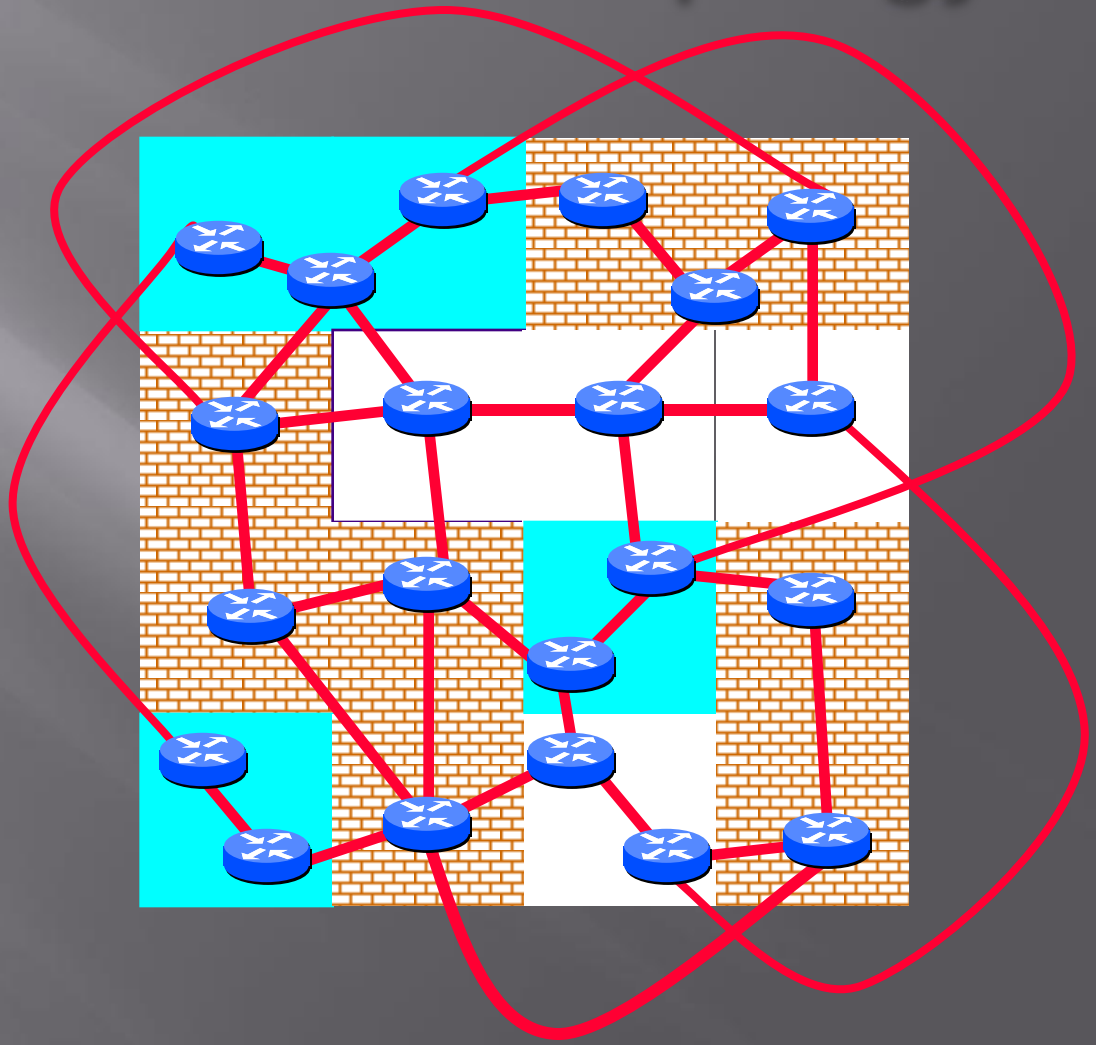

Reality may be closer to this…

AS Graphs Depend on Point of View

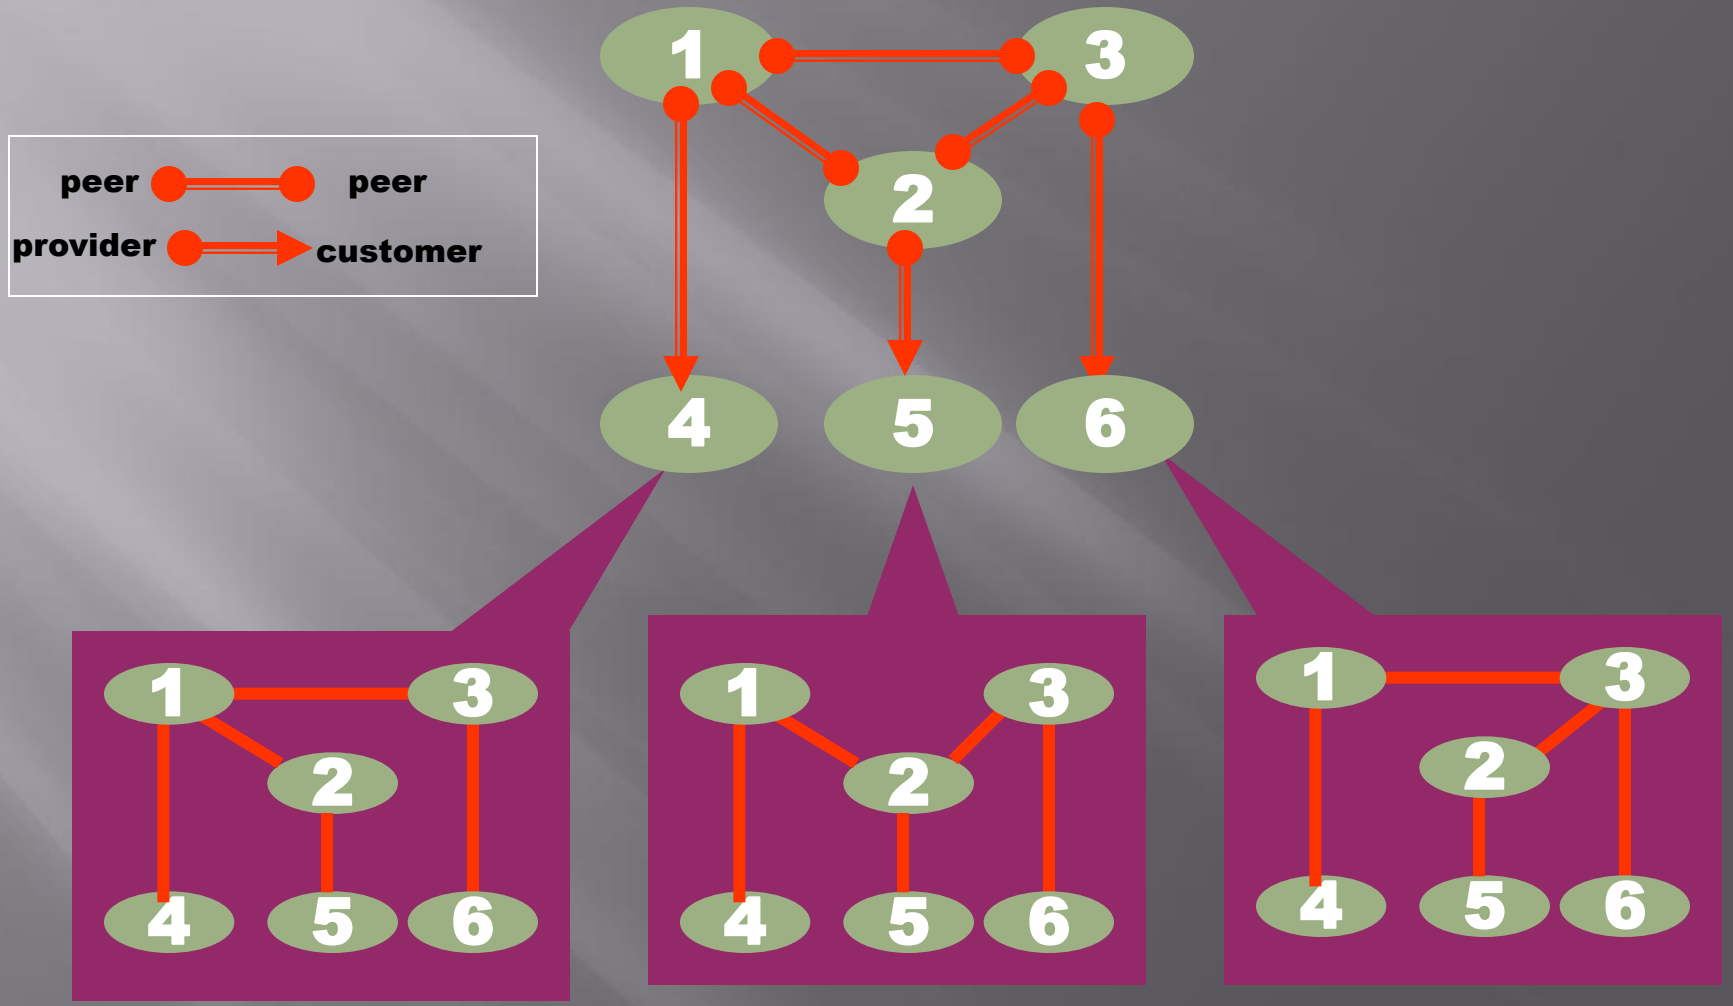

This explains why there is no UUNET (701) Sprint (1239) link on previous slide!

Shorter Doesn't Always Mean Shorter

In fairness: could you do this "right" and still scale? Exporting internal state would dramatically increase global instability and amount of routing state

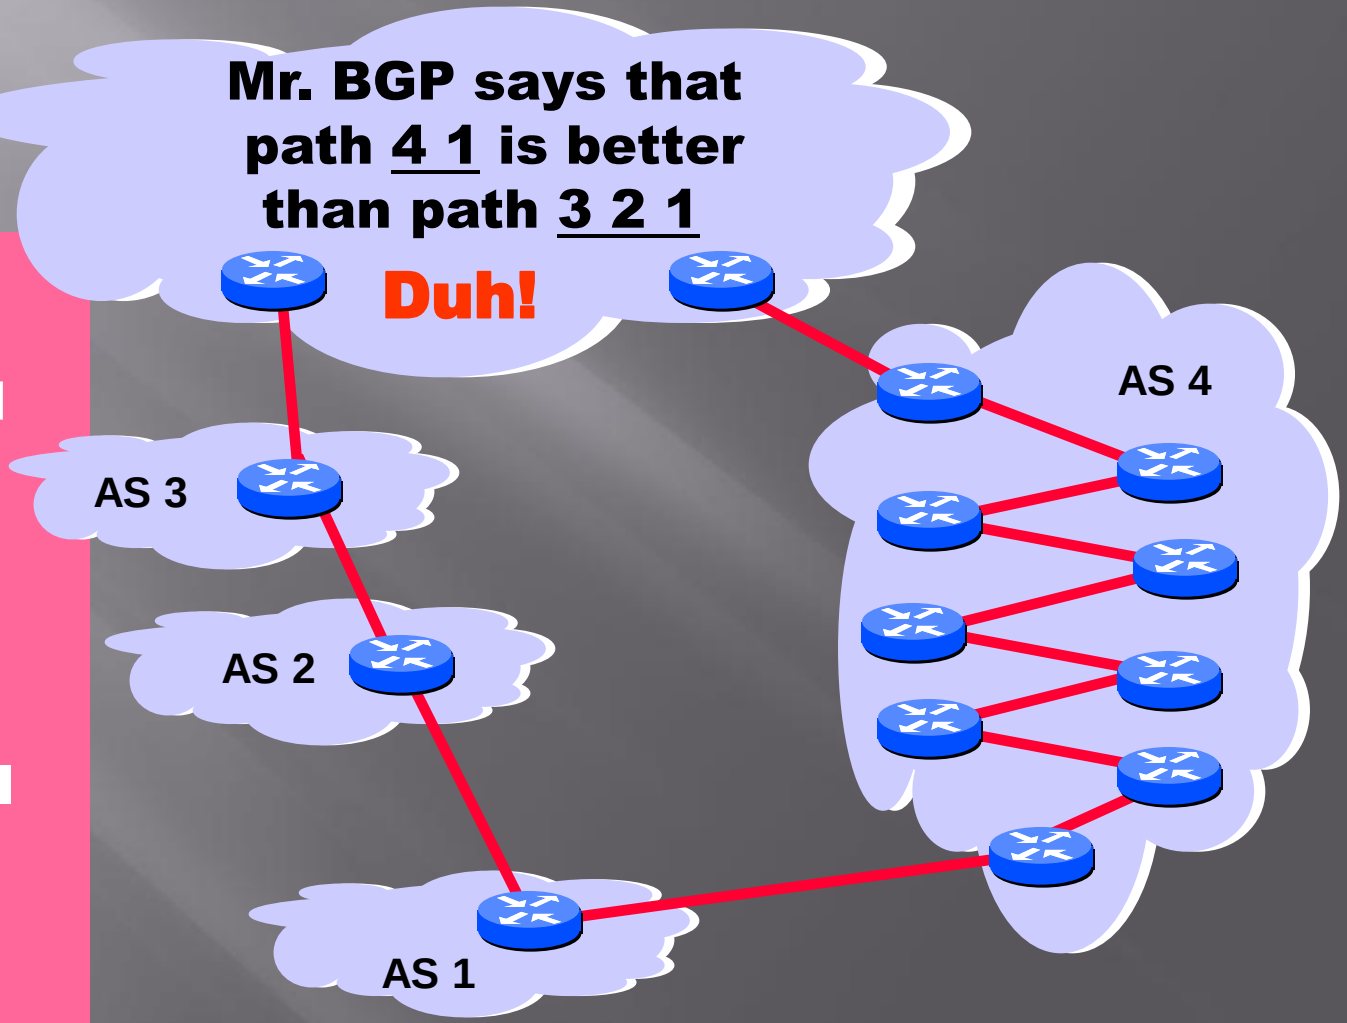

Route Selection Summary

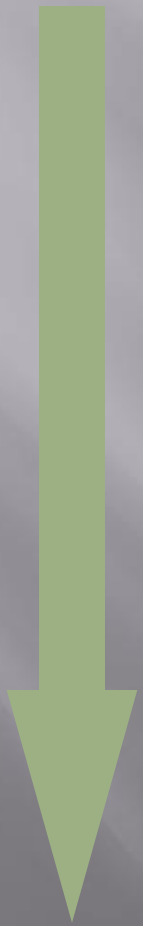

Highest Local Preference	**Enforce relationships**
Shortest ASPATH **Lowest MED** **i-BGP < e-BGP** **Lowest IGP cost to BGP egress**	**traffic engineering**
Lowest router ID	**Throw up hands and break ties**

Implementing Customer/Provider and Peer/Peer relationships

Two parts:

- Enforce transit relationships
  - Outbound route filtering
- Enforce order of route preference
  - provider < peer < customer

Import Routes

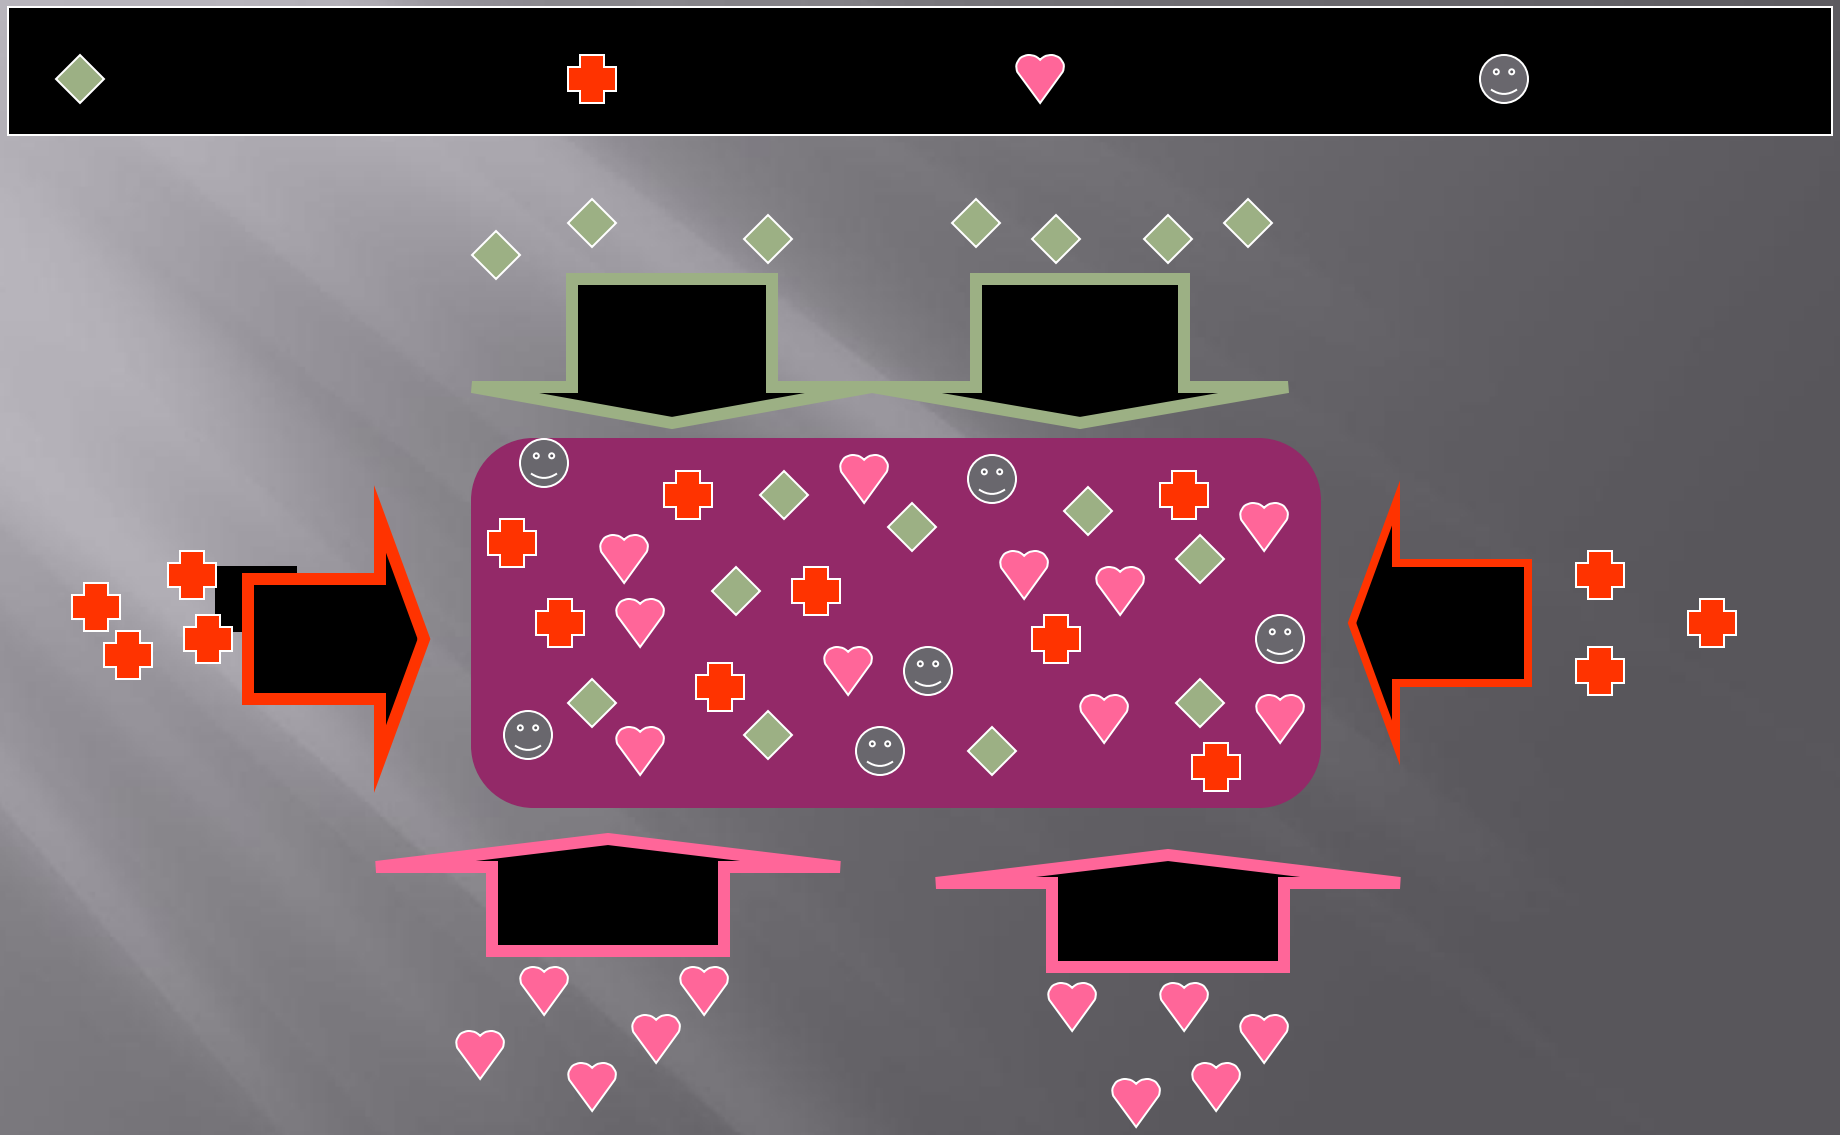

Export Routes