

BGP and communities: Fixing bad transit

MBNOG 1 - March 22, 2018

Winnipeg, MB, Canada

by Jonathan Stewart

LES.NET

BGP between ASNs

Autonomous Systems on the internet:

- Run BGP on routers
 - Send routes
 - Receive routes
- Use BGP to select best forwarding path
 - For packets destined outside the AS
- IGP like OSPF is used for packets destined inside the AS

BGP key points

- IPv4 different routes than IPv6
 - Duplicate Everything

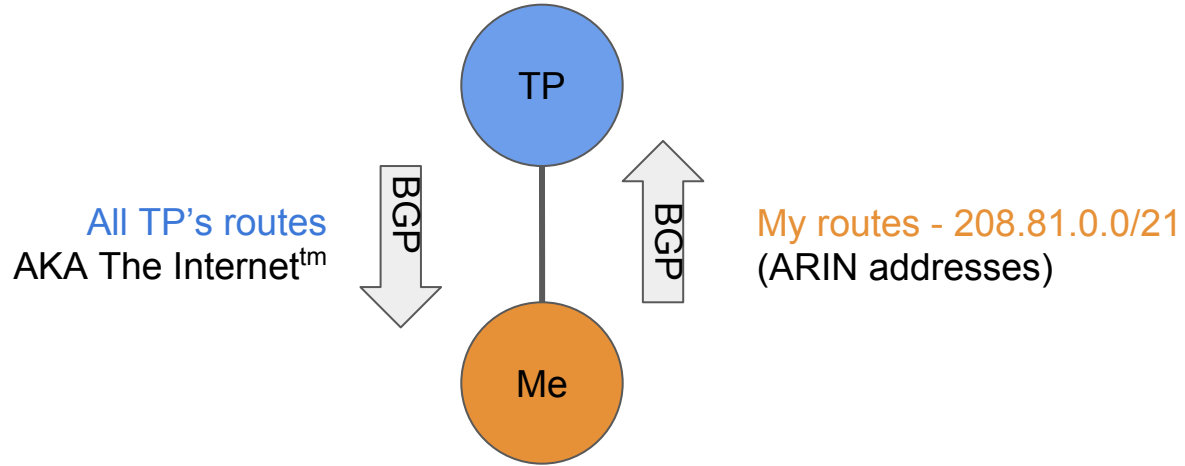
Examples:

- ASN: AS18451 or AS393445
- IPv4 Route: 208.81.0.0/21
- IPv6 Route: 2605:e200::/32

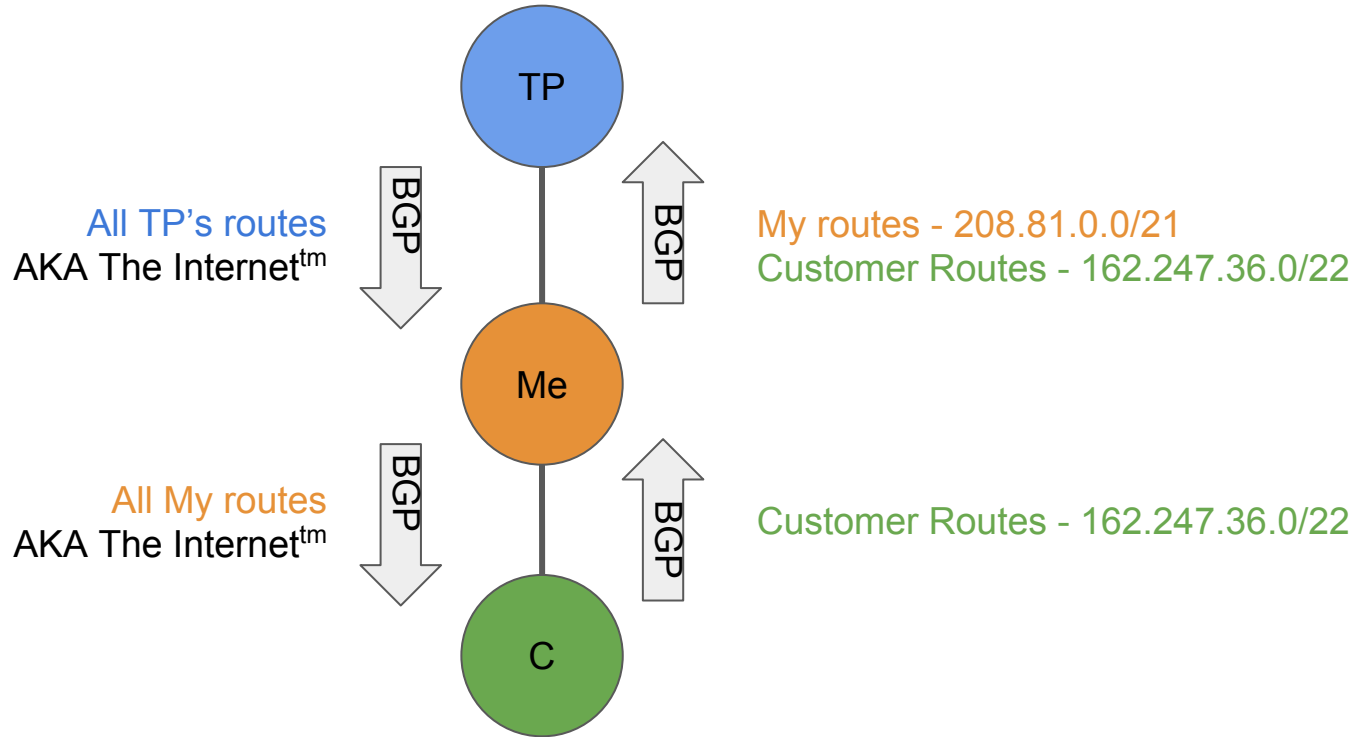
BGP Perspective

- Each AS sees the internet differently
 - Each AS picks its own best paths
-
- For outbound traffic, think of YOUR best path
 - For inbound traffic, think of THEIR best path

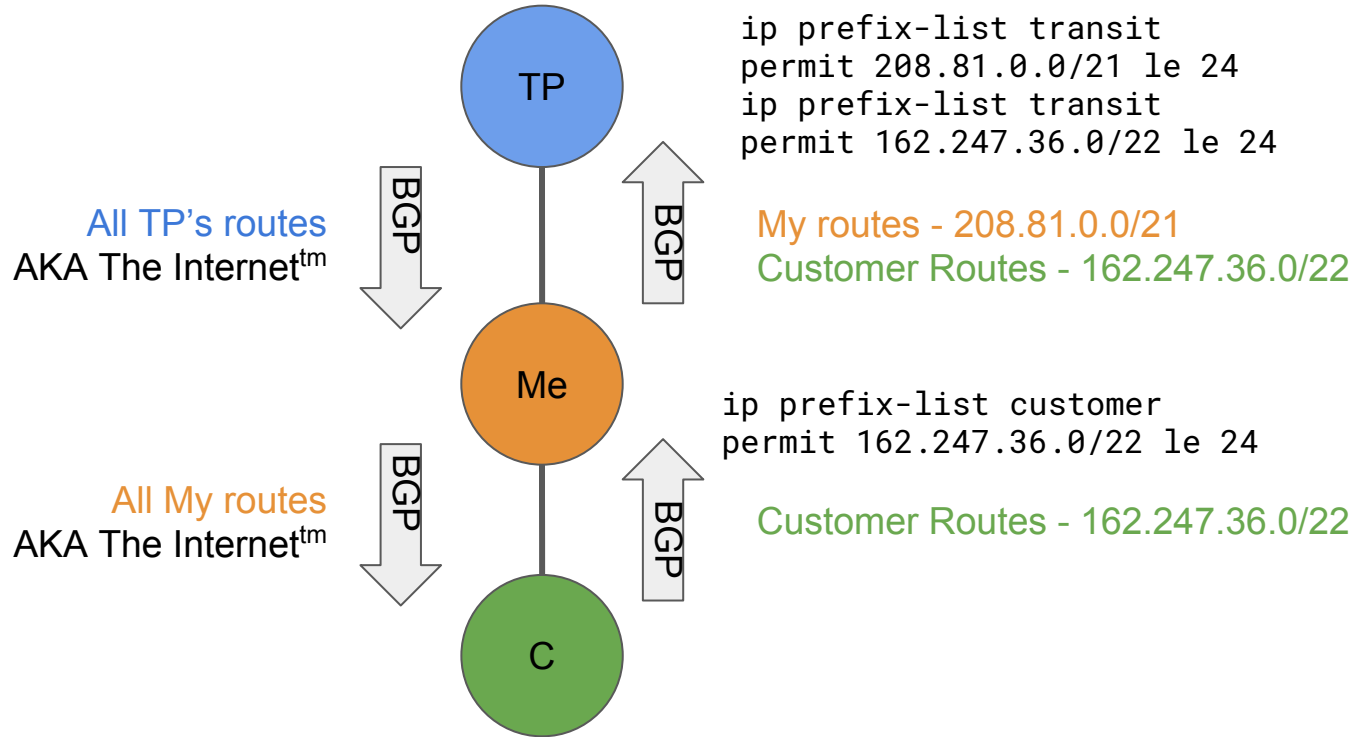
My simplest network



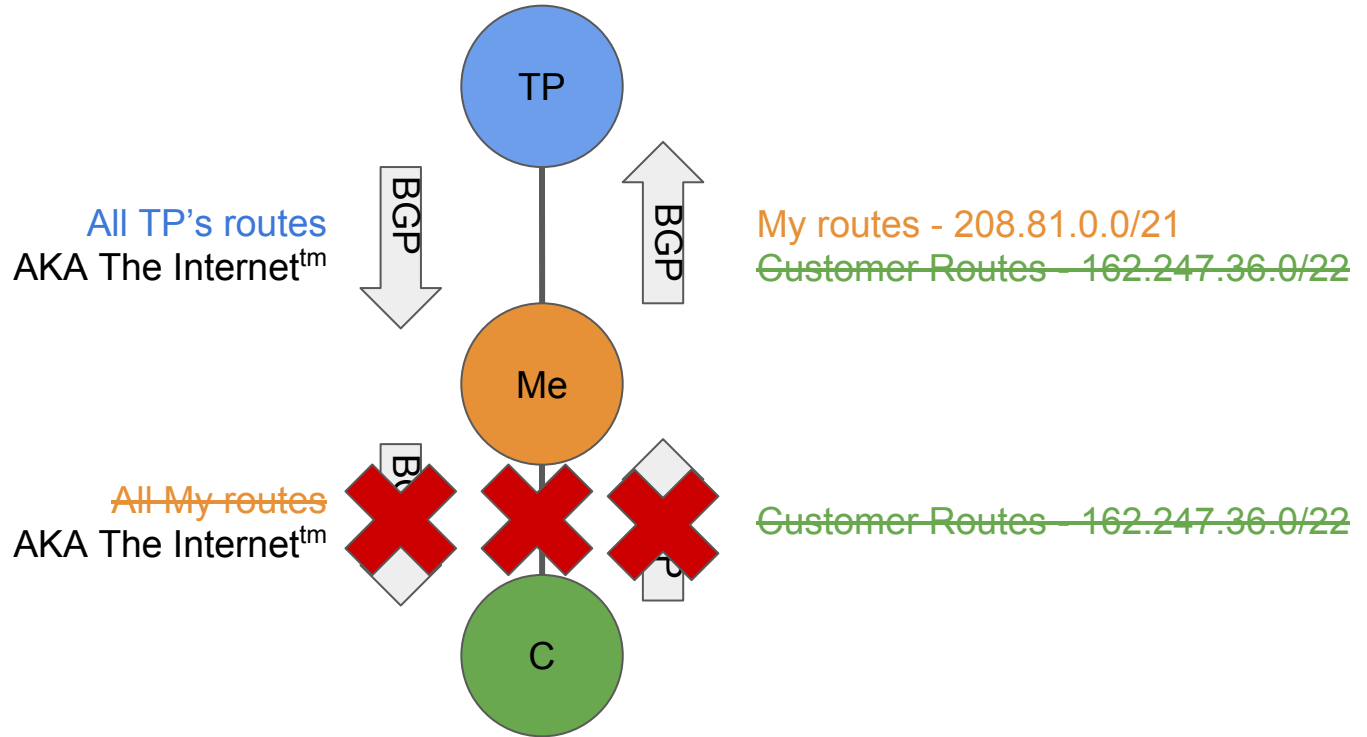
My simple transit network



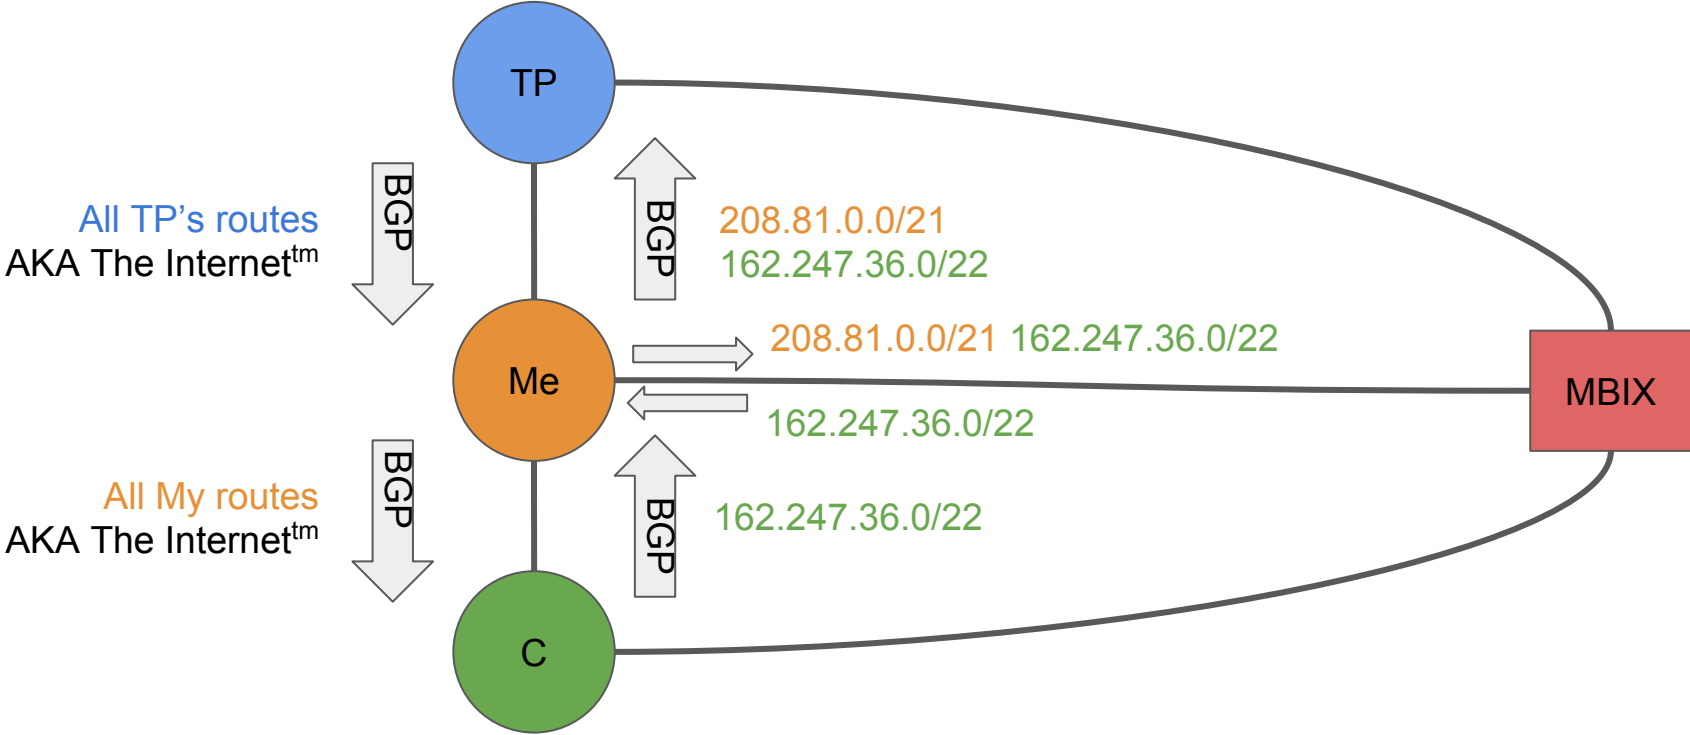
My simple transit network - prefix lists



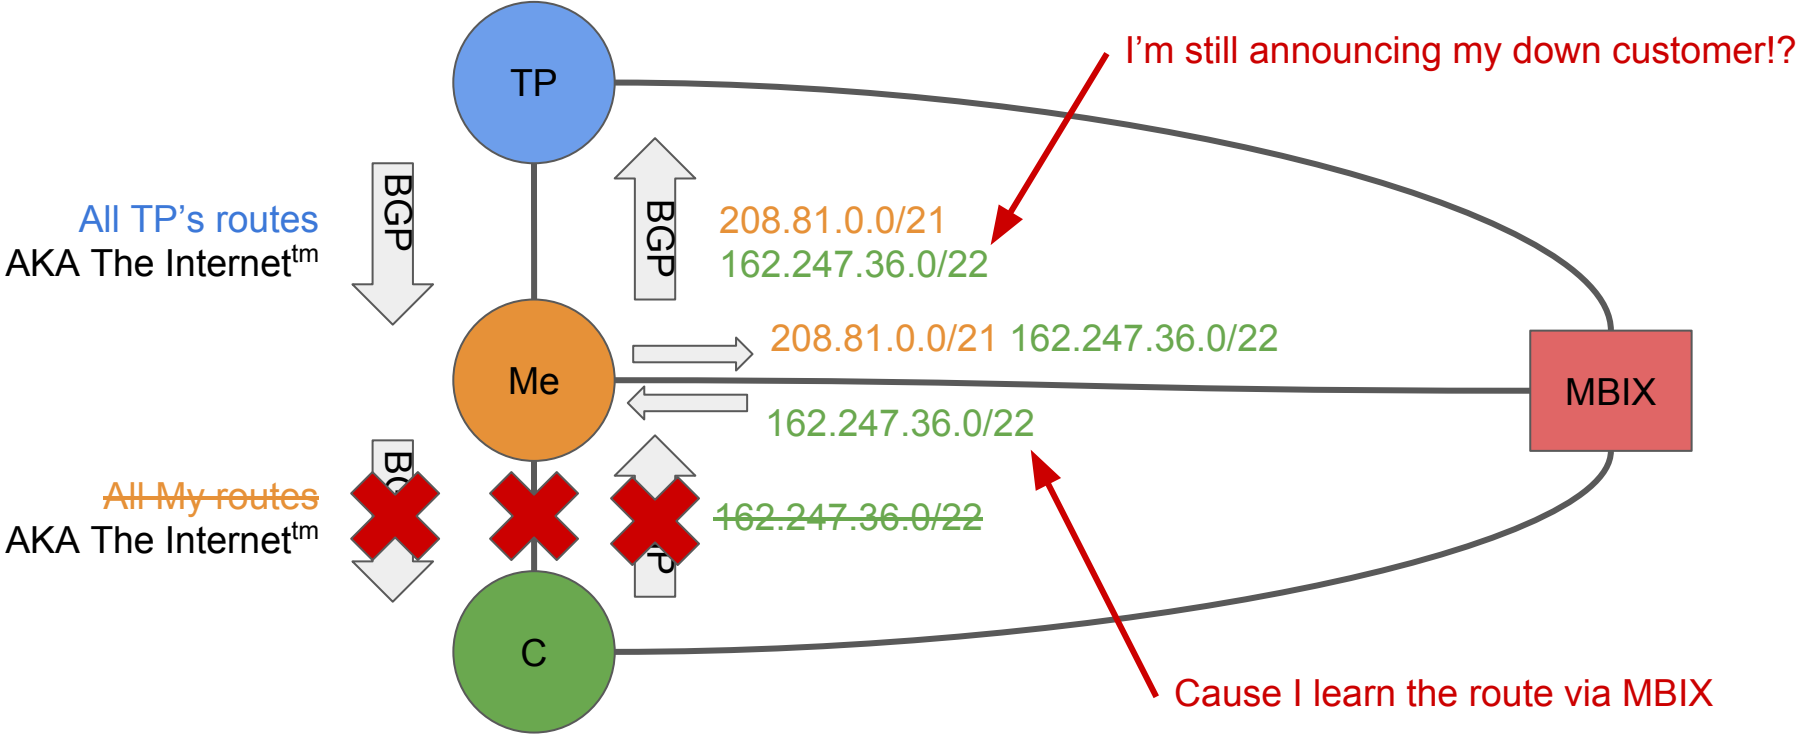
My simple transit network - Down Customer



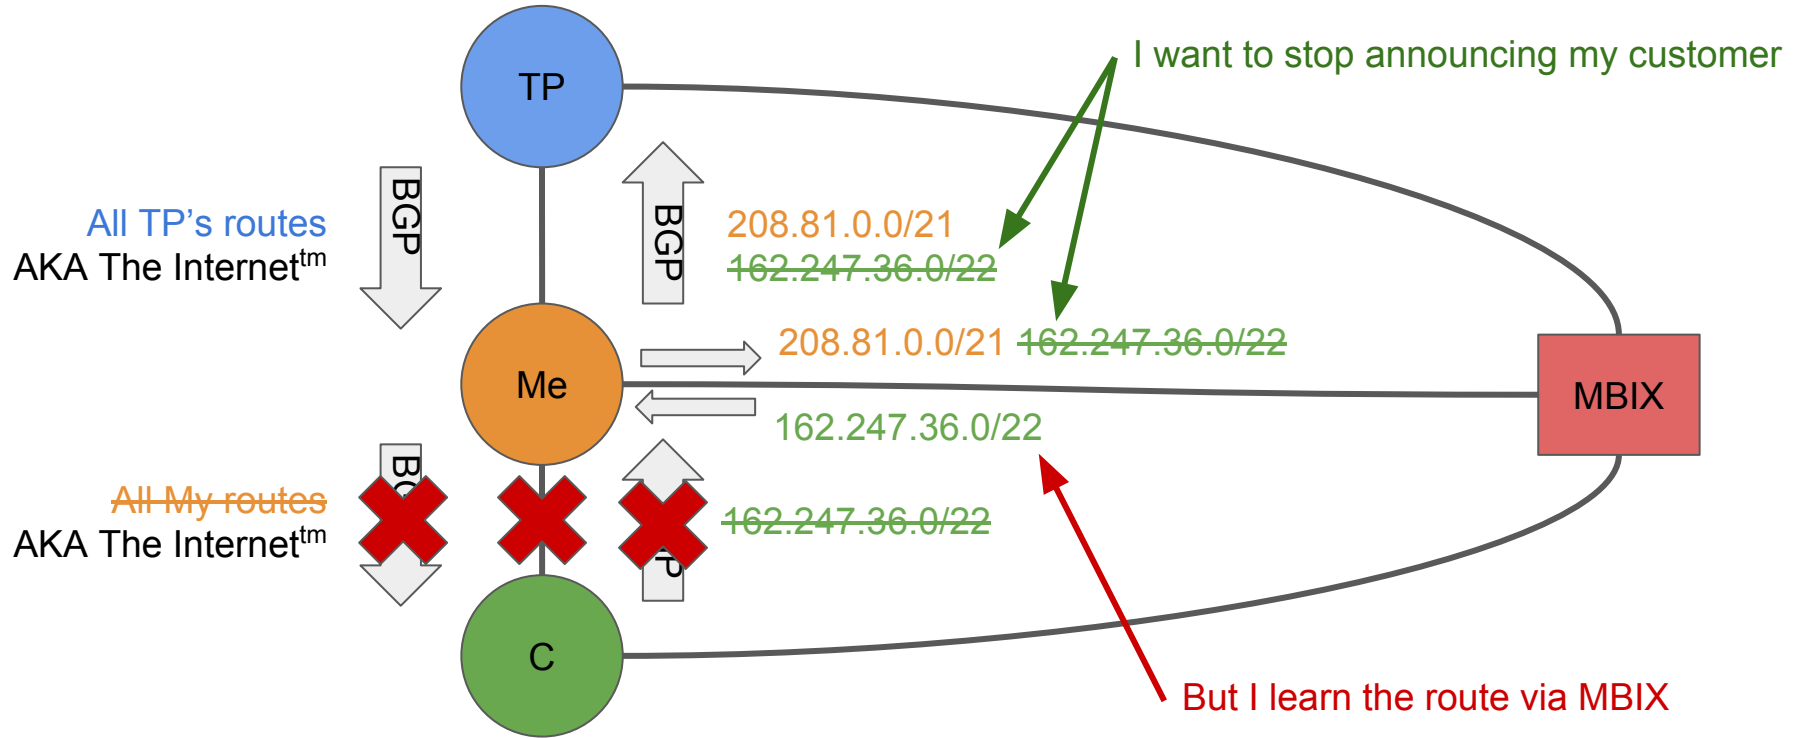
My transit & peering network



My transit & peering network - Down Customer



My transit & peering network - Down Customer Desire



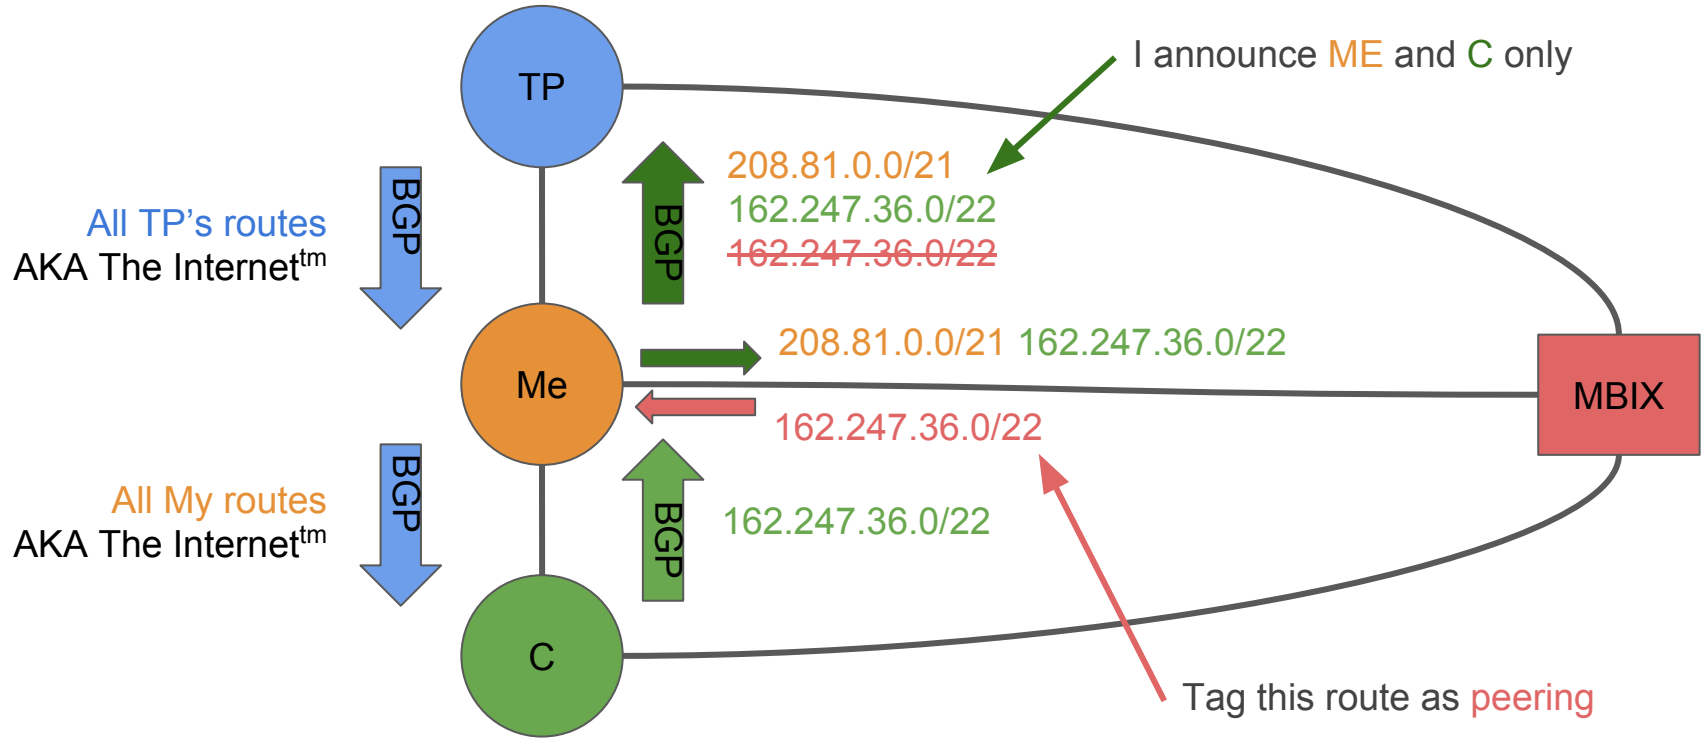
Solution: Tag the routes

- BGP Communities
- Add tag as you learn the route
- Define your own meanings
 - Not that hard

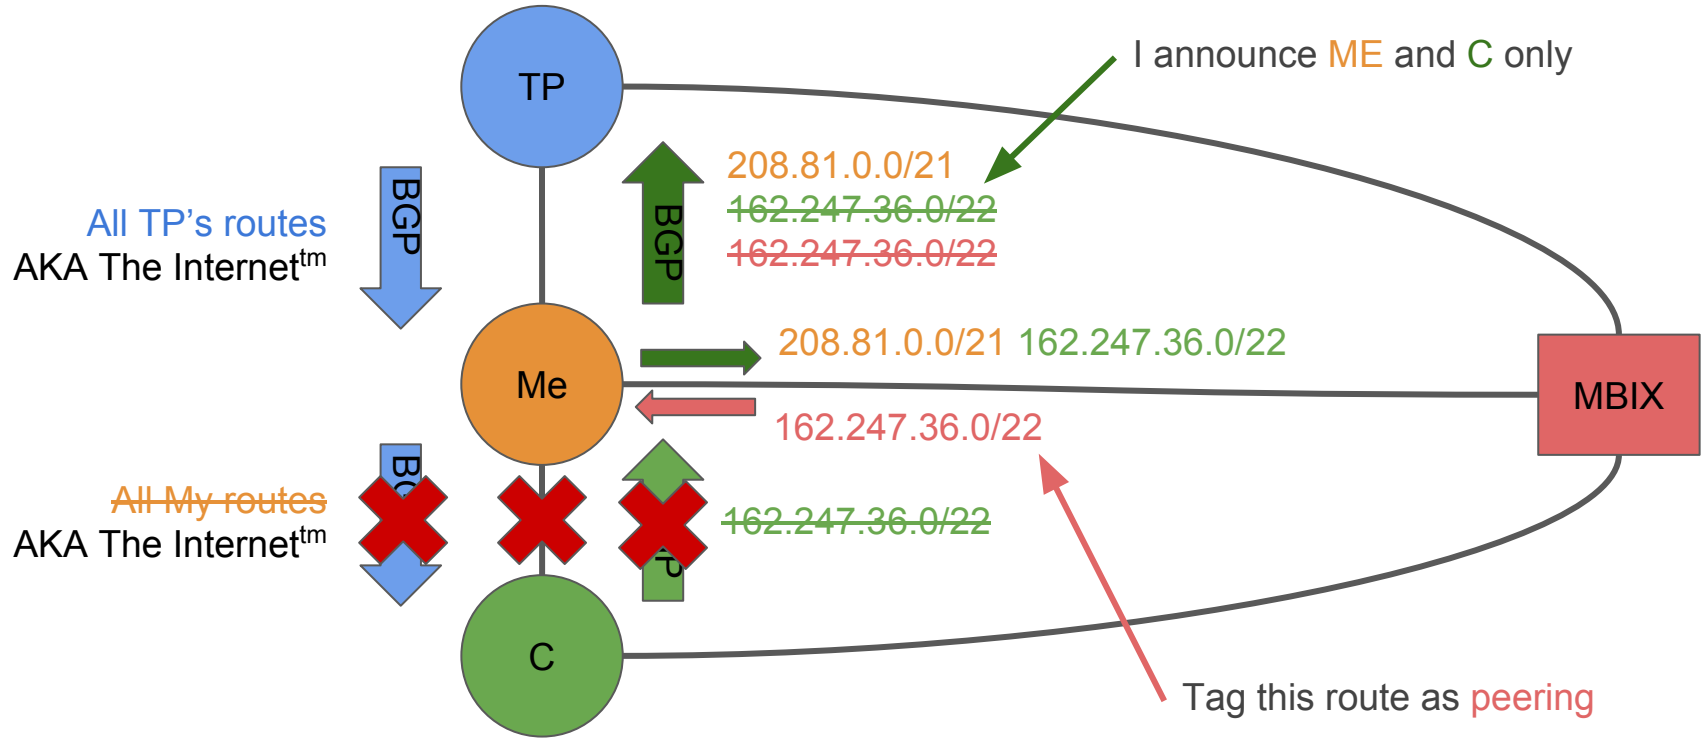
BGP community examples:

- 18451:5
- 6327:3001

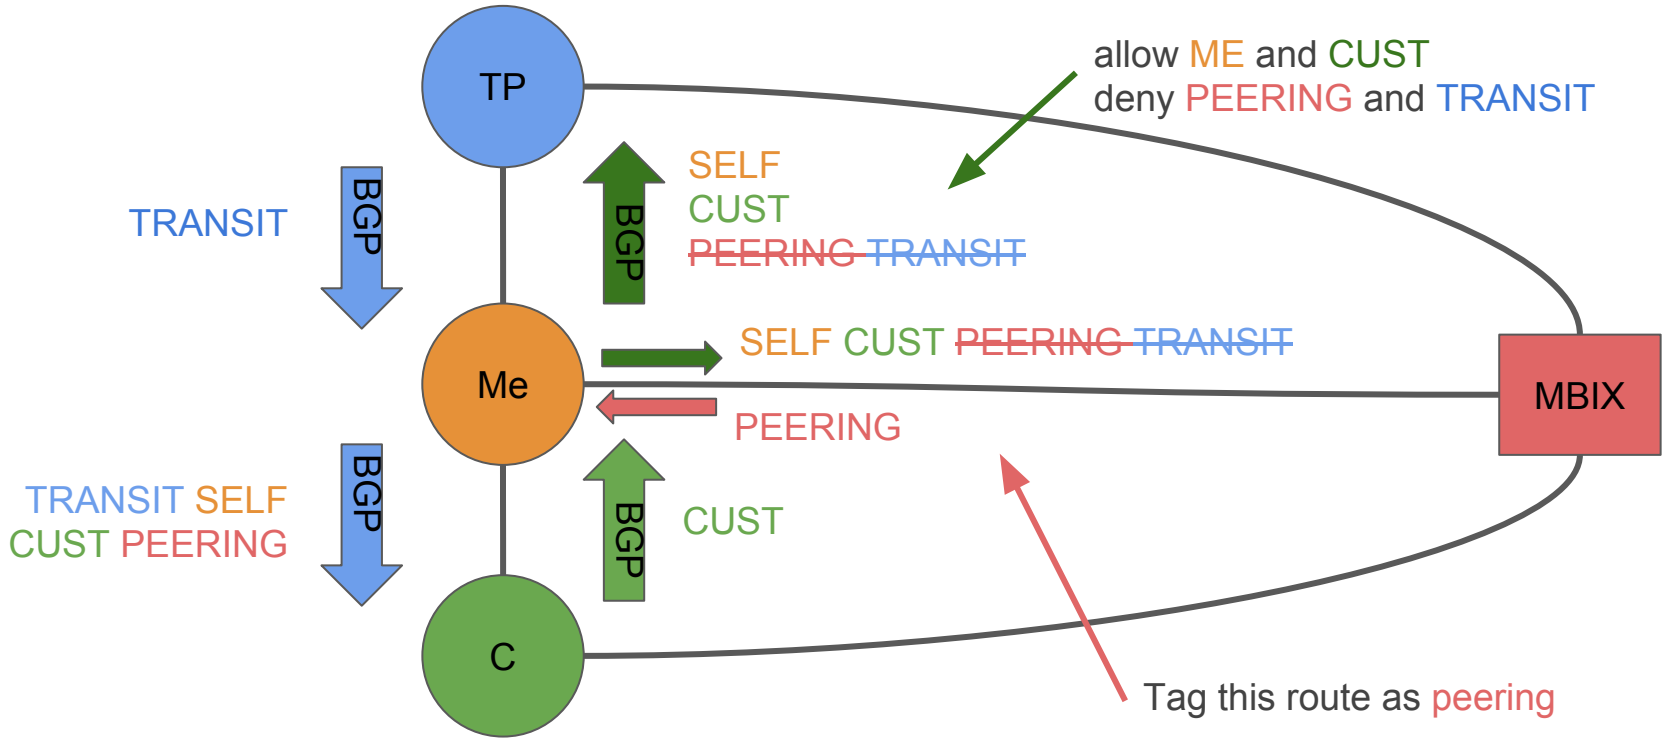
My transit & peering network - Up Customer



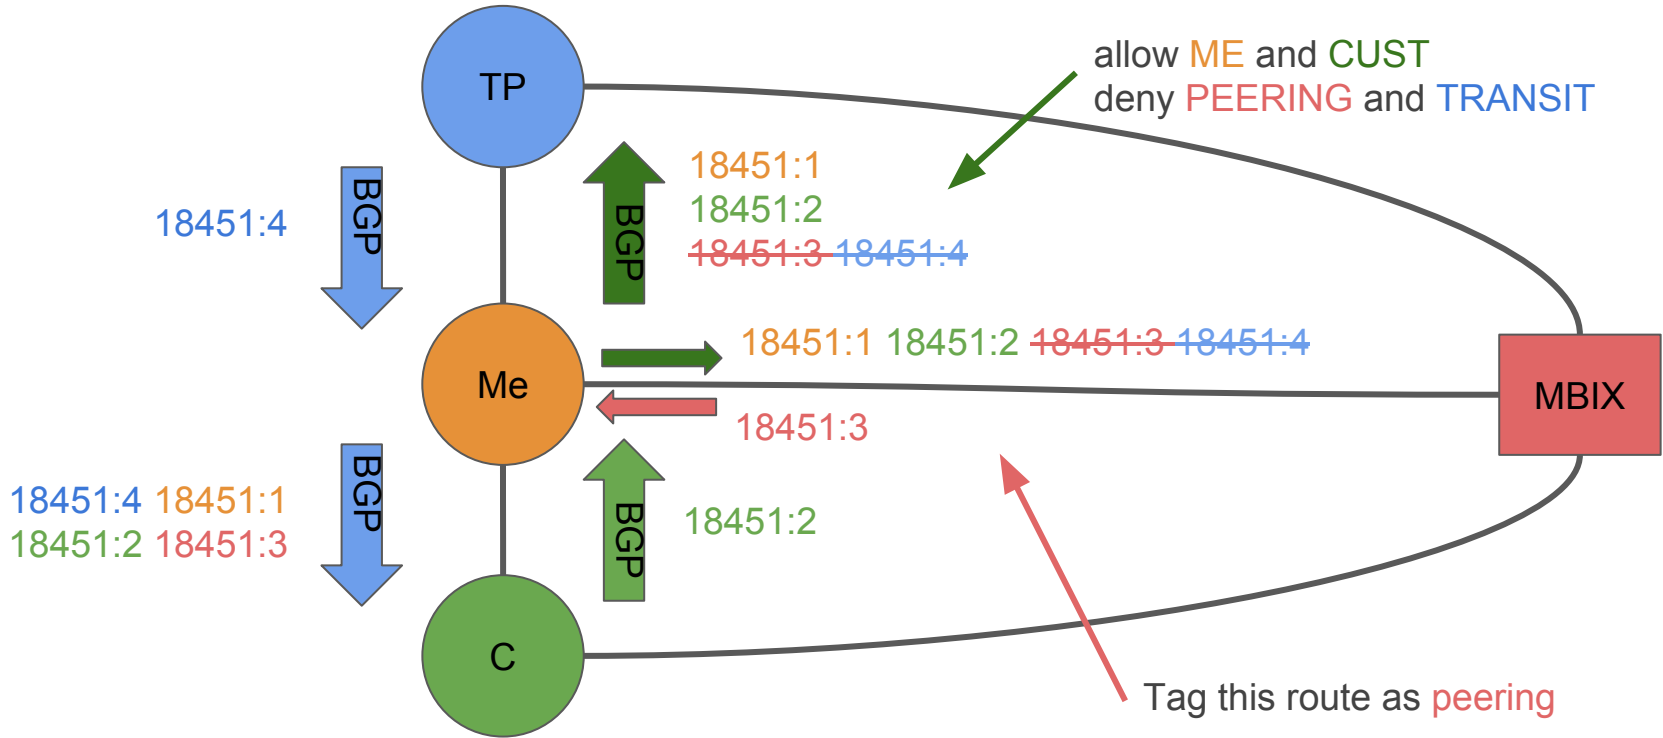
My transit & peering network - Down Customer



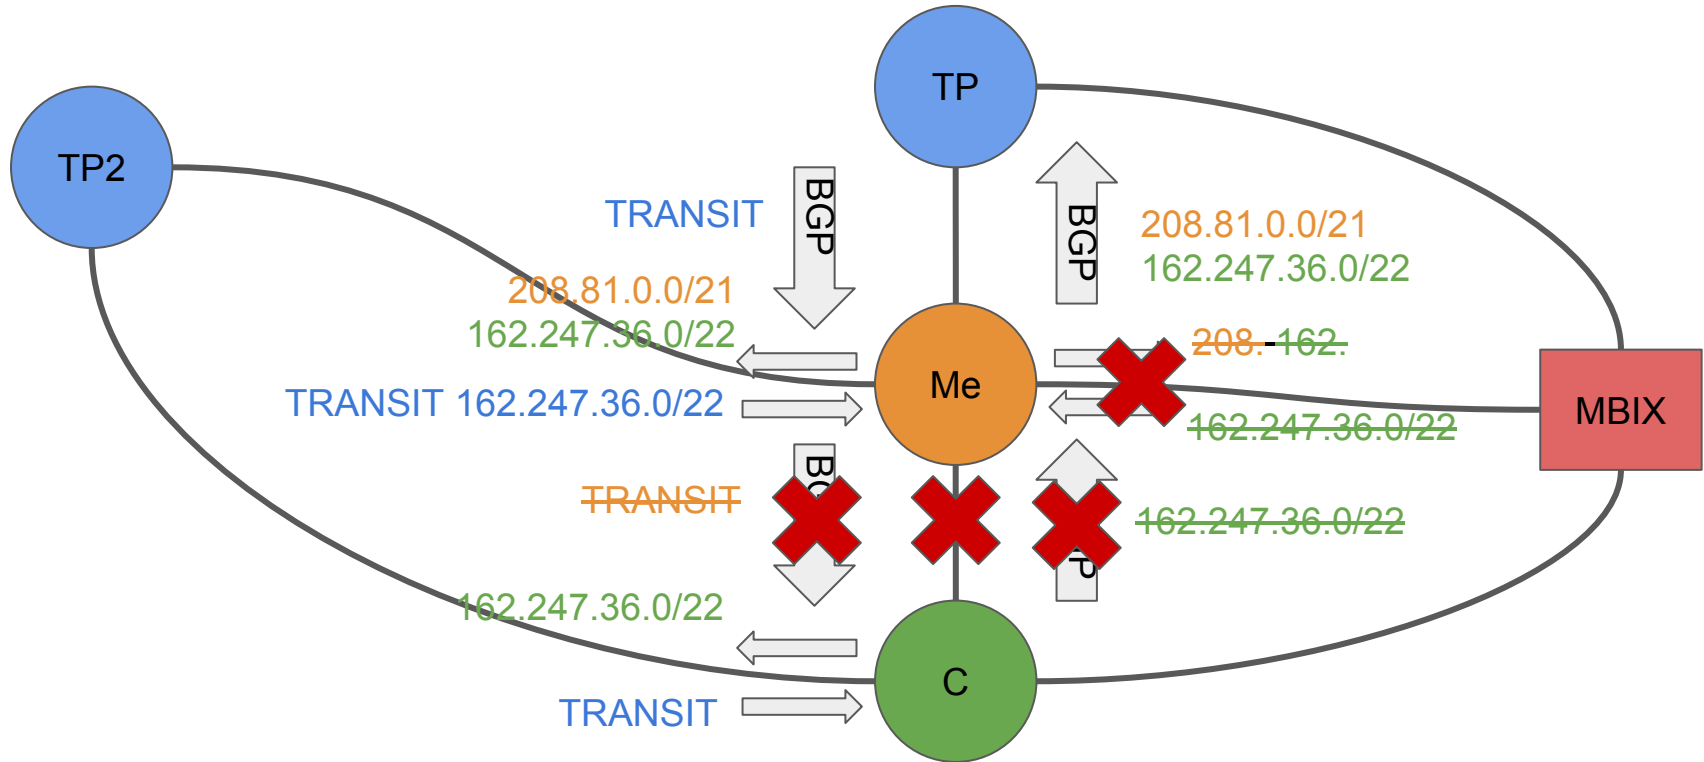
My transit & peering network - Up Customer



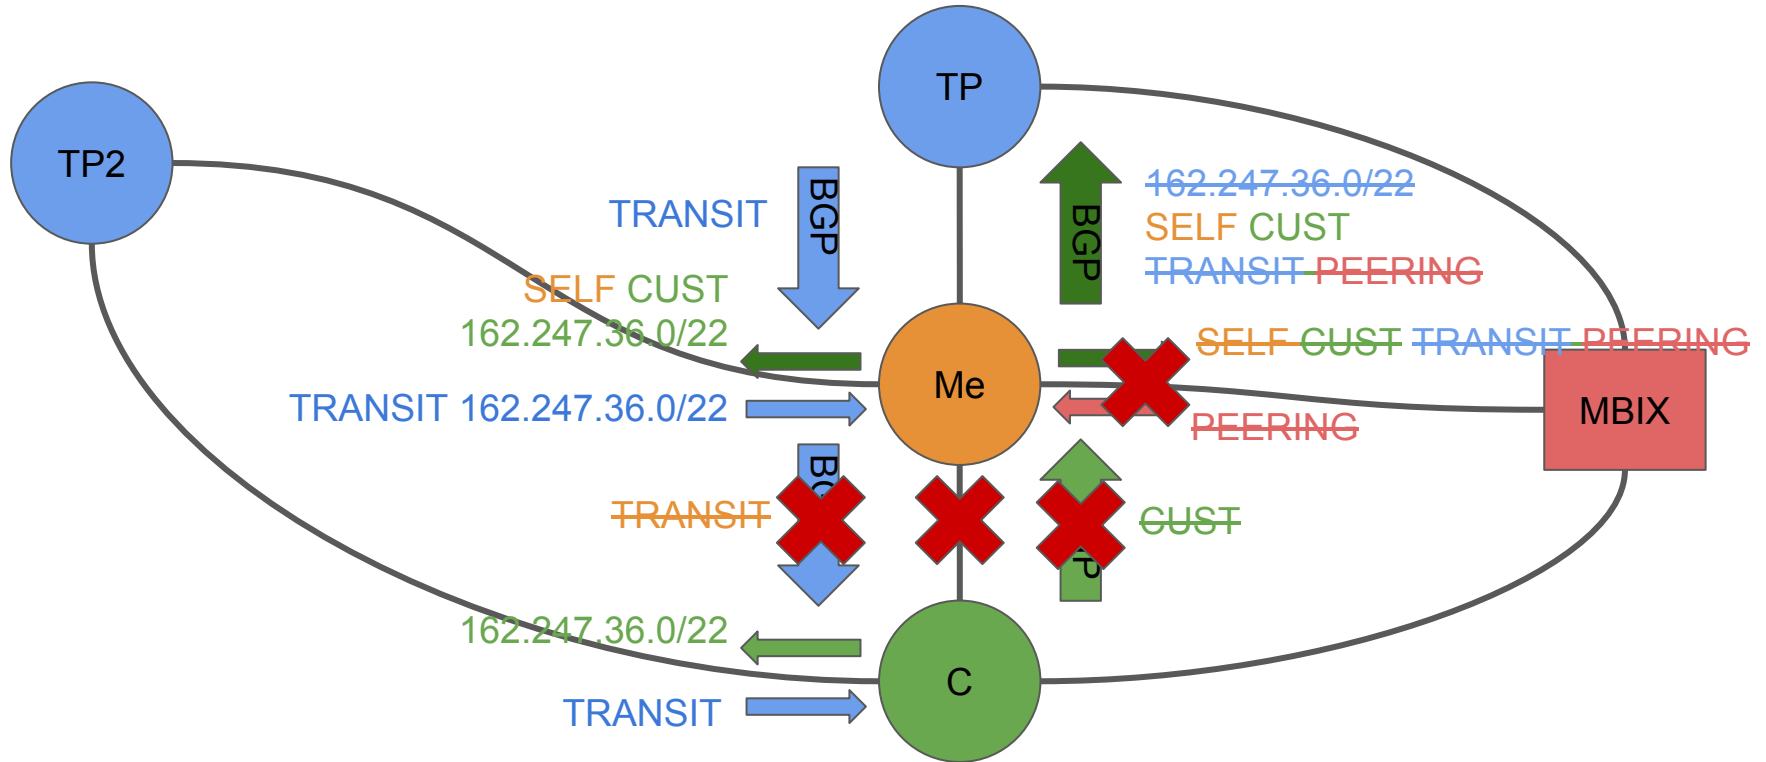
My transit & peering network - Up Customer



My multihomed transit, peering network - The Worst



My multihomed network - The Worst Fixed



Fix: bgp neighbor code

```
router bgp 18451
```

```
address-family ipv4
```

```
neighbor 216.66.74.57 route-map PEER-6939-v4-in in
```

```
neighbor 216.66.74.57 route-map PEER-6939-v4-out out
```

```
address-family ipv6
```

```
neighbor 2001:470:1:40D::1 route-map PEER-6939-v6-in in
```

```
neighbor 2001:470:1:40D::1 route-map PEER-6939-v6-out out
```

Fix: route-map code

```
route-map PEER-6939-v4-in permit 10
  set local-preference 90  ! Bonus--next talk
  set community 18451:5
```

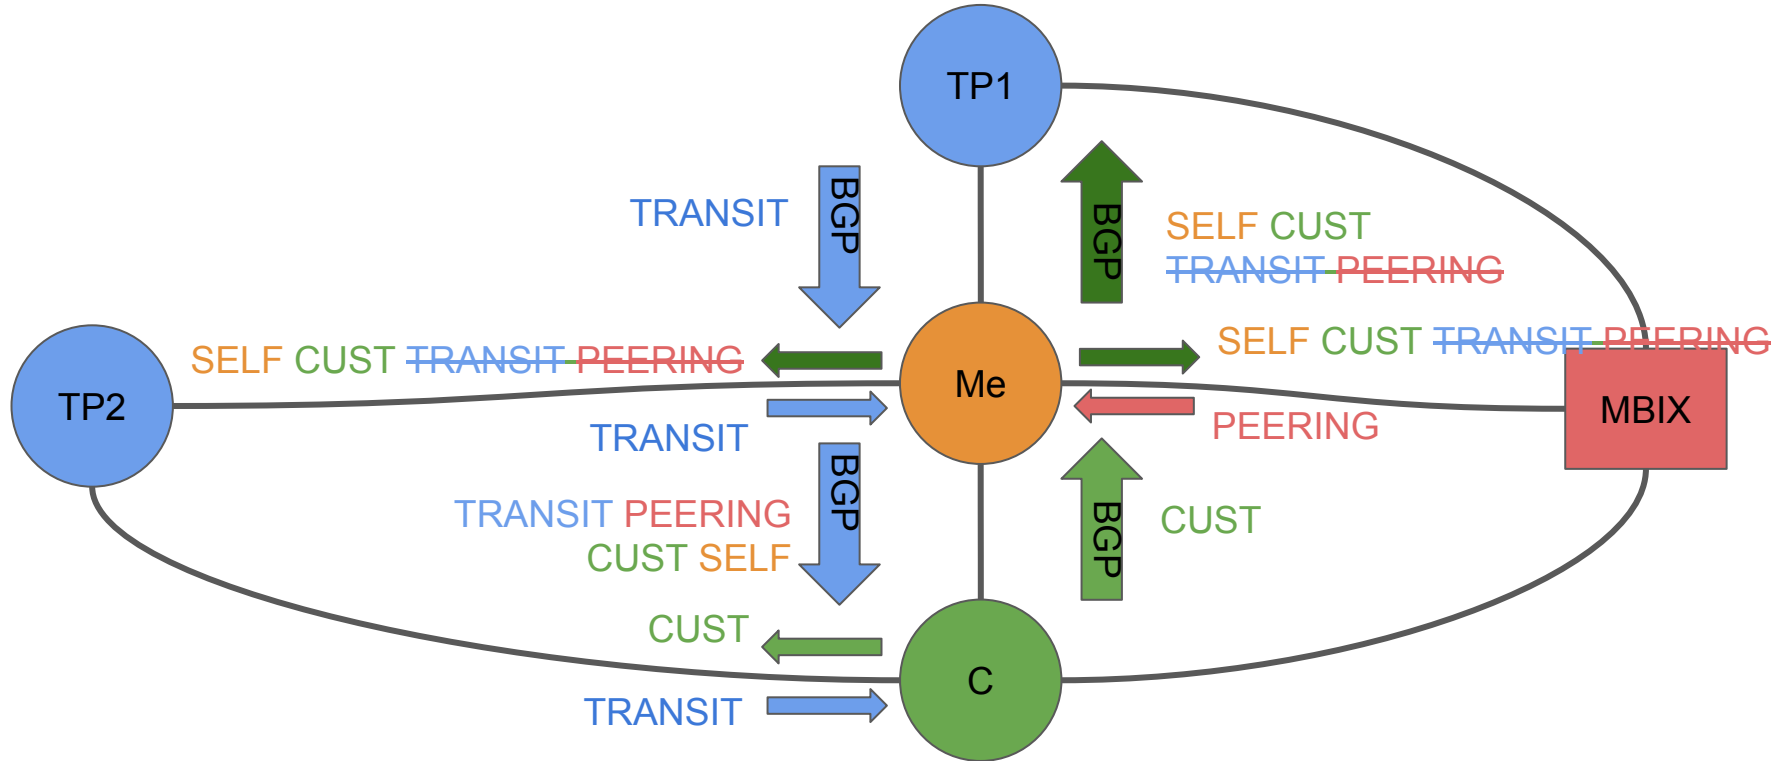
```
route-map PEER-6939-v4-out deny 10
  match community TRANSIT PEER
route-map PEER-6939-v4-out permit 20
  match community CUST SELF
route-map PEER-6939-v4-out deny 999
```

Fix: community-list code

- Make a useful name to match your (numeric) community

```
ip community-list expanded SELF permit 18451:1$
ip community-list expanded SELF permit 18451:1 .*
ip community-list expanded CUST permit 18451:2$
ip community-list expanded CUST permit 18451:2 .*
ip community-list expanded PEER permit 18451:3$
ip community-list expanded PEER permit 18451:3 .*
ip community-list expanded TRANSIT permit 18451:4$
ip community-list expanded TRANSIT permit 18451:4 *
```

My multi-homed transit & peering network- Overview



BGP and communities: Fixing bad transit

Jonathan Stewart

LES.NET Network Ops

jonathan.stewart@gmail.com

Twitter: @fudoreaper