

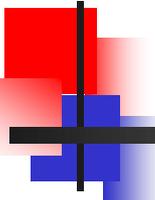
APRICOT 2013
Singapore

19 February - 1 March 2013



ISP and NSP Security Workshop

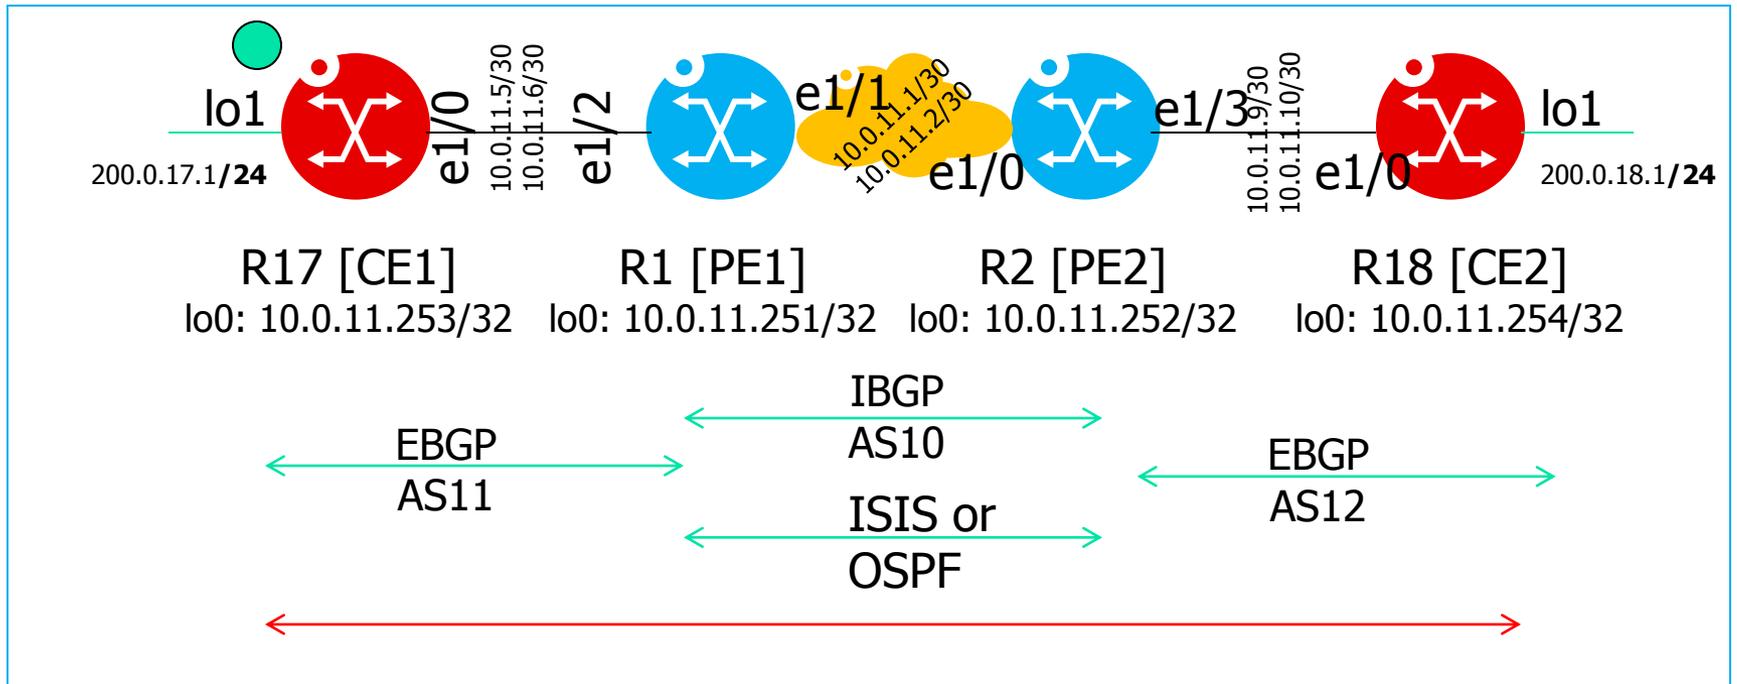
APRICOT 2013
Prefix hijacking

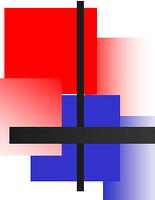


Prefix hijacking and more specifics

- Prefix hijacking can occur in both accidental and malicious scenarios
 - Accidental: YouTube Pakistan Telecom hijack
 - Malicious: Someone wearing a black hat wants to impact your network / steal your reputation
 - Think spammers
- In this lab, we'll take a look at a simple prefix hijack through *more specific* routes, and mitigations for same

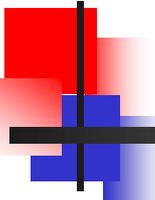
Lab starting point





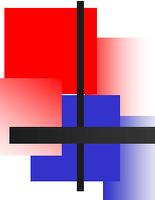
Lab starting check point

- IBGP and OSPF operating PE1-PE2 as per day 1
- EBGP operating between CE1-PE1, CE2-PE2
- Remove IPSEC



Lab

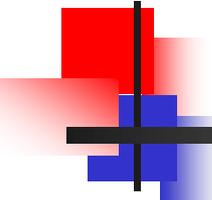
- Configure 200.0.XX.1/**24** as a BGP network on the CE
 - Our existing networks are **/32** networks
 - Remove the /32 networks
- Check that the /24 prefix is received at the PE routers



Lab

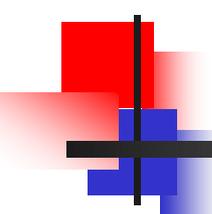
- Put on your black hat!
- On CE2, configure a new loopback interface (loopback555) with 200.0.XX.1 – the IP address of CE1!
- On CE2, inject 200.0.XX.0/**25** into BGP

- **What do you see on PE1 and PE2 for BGP routes?**



Lab

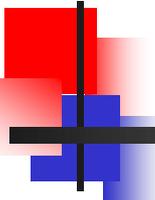
- Pinging from PE1 to CE1's loopback1 IP address will now go to CE2!
- **This is a simple example of prefix hijacking**
- What are the options to mitigate this?



Lab

- Configure a prefix-list on the PE2-CE2 BGP session that will permit only CE2's **legitimate prefixes**

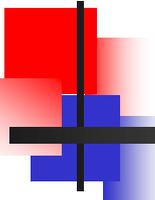
- Are there other options for protecting the network from this attack?



CE1 configuration

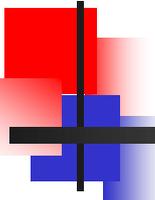
```
interface ethernet X/X
  description "To PE1"
  ip address 10.0.X.X 255.255.255.252
  no ip redirects
  no ip proxy-arp
!
interface loopback0
  ip address 10.0.X.X 255.255.255.255
!
interface loopback1
  ip address 200.0.XX.1 255.255.255.255
!
router bgp X1
  no synchronization
  no auto-summary
  network 200.0.XX.1 255.255.255.0
  network 10.0.XX.X 255.255.255.255
  neighbor 10.0.X.X remote-as XX
  neighbor 10.0.X.X password <PASSWORD>
  neighbor 10.0.X.X send-community
!
ip route 200.0.XX.0 255.255.255.0 null0
```

We don't need the RTBH
route-maps anymore, so
remove those if present



PE1 configuration

```
interface ethernet X/X
  description "To PE2"
  ip address 10.0.X.X 255.255.255.252
  no ip redirects
  no ip proxy-arp
!
interface loopback0
  ip address 10.0.X.X 255.255.255.255
!
router bgp X0
  no synchronization
  no auto-summary
  neighbor 10.0.X.X remote-as X1
  neighbor 10.0.X.X password <PASSWORD>
  neighbor 10.0.X.X send-community
  neighbor 10.0.X.X ebgp-multihop 2
  neighbor 10.0.X.Y
  neighbor 10.0.X.Y send-community
  neighbor 10.0.X.Y next-hop-self
  neighbor 10.0.X.Y password <PASSWORD>
```



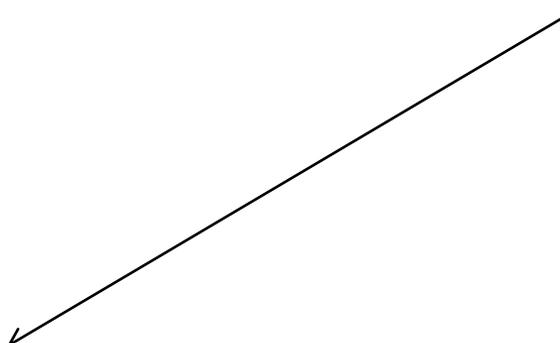
PE2 configuration

```
interface ethernet X/X
  description "To PE2"
  ip address 10.0.X.X 255.255.255.252
  no ip redirects
  no ip proxy-arp
!
interface loopback0
  ip address 10.0.X.X 255.255.255.255
!
router bgp X0
  no synchronization
  no auto-summary
  neighbor 10.0.X.X remote-as X1
  neighbor 10.0.X.X password <PASSWORD>
  neighbor 10.0.X.X send-community
  neighbor 10.0.X.X ebgp-multihop 2
  neighbor 10.0.X.Y
  neighbor 10.0.X.Y send-community
  neighbor 10.0.X.Y next-hop-self
  neighbor 10.0.X.Y password <PASSWORD>
```

CE2 configuration

```
interface ethernet X/X
  description "To PE2"
  ip address 10.0.X.X 255.255.255.252
  no ip redirects
  no ip proxy-arp
!
interface loopback0
  ip address 10.0.X.X 255.255.255.255
!
interface loopback1
  ip address 200.0.XX.1 255.255.255.255
!
interface loopback555
  ip address 200.0.YY.1 255.255.255.255
!
router bgp X2
  no synchronization
  no auto-summary
  network 200.0.XX.1 255.255.255.255
  network 10.0.XX.X 255.255.255.255
  network 200.0.YY.1 255.255.255.128
  neighbor 10.0.X.X remote-as XX
  neighbor 10.0.X.X password <PASSWORD>
  neighbor 10.0.X.X send-community
!
ip route 200.0.YY.1 255.255.255.128 null0
```

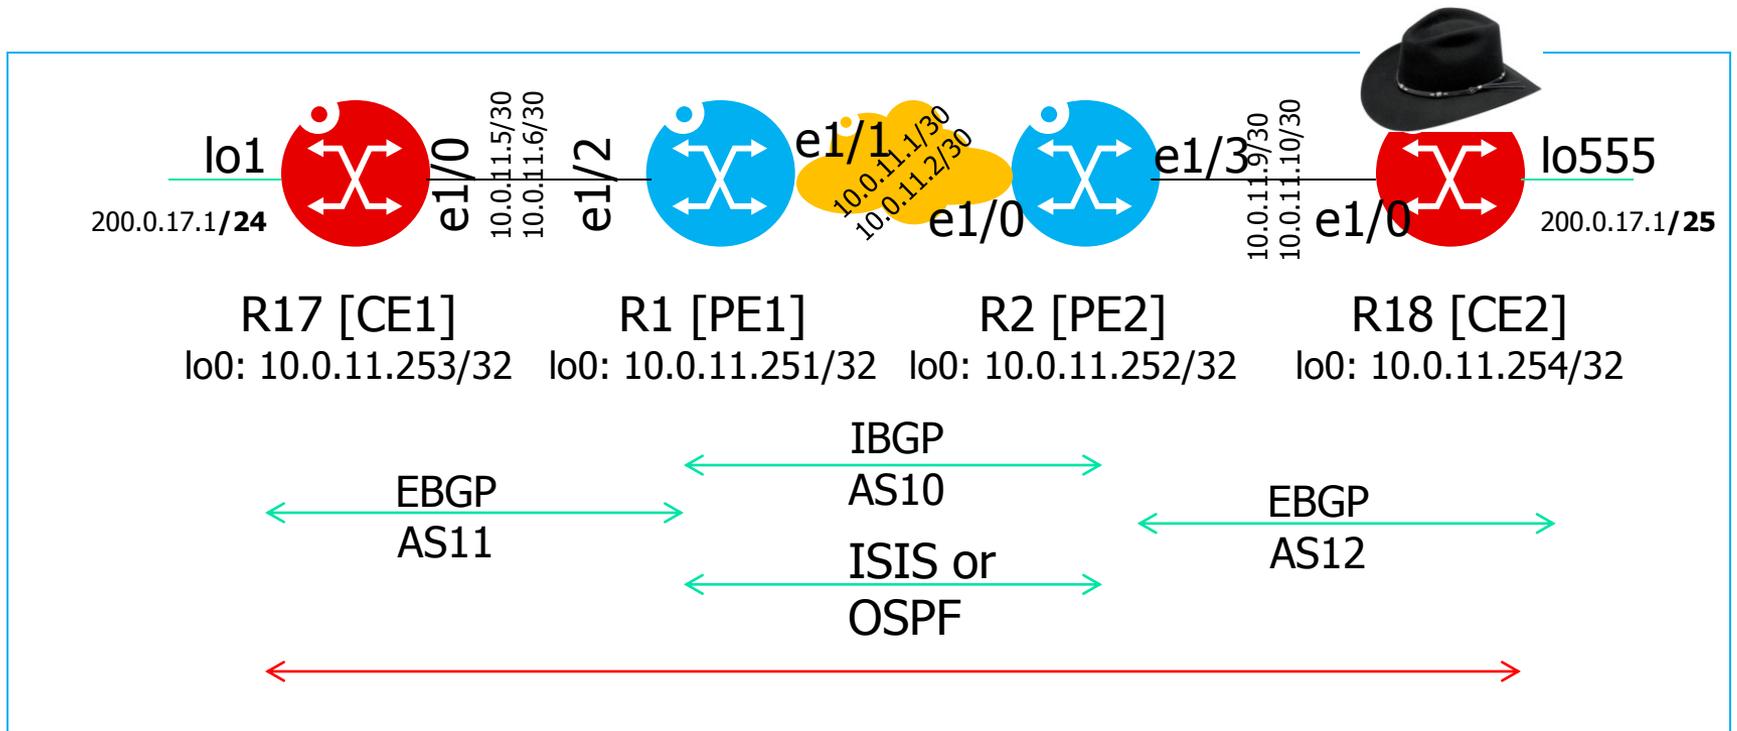
This is CE1's IP

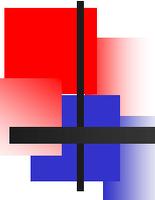


We don't need the RTBH
route-maps anymore, so
remove those if present



Lab black-hat mode





Lab

- Because CE2 announced 200.0.17.0/25 to PE2, and CE1 announced 200.0.17.0/24 to CE1, the longer prefix won
- Traffic for 200.0.17.1 ended up at CE2!
- Prefix-list on PE2-CE2 BGP session will prevent PE2 from announcing CE1 addresses