



Internet Exchange Points

BGP configuration for IXPs

Definitions

- **Transit** – carrying traffic across a network, usually for a fee
 - traffic and prefixes originating from one AS are carried across an intermediate AS to reach their destination AS
- **Peering** – private interconnect between two ASNs, usually for no fee
- **Internet Exchange Point** – common interconnect location where several ASNs exchange routing information and traffic

IXP Peering Issues

- Only announce your prefixes and your customer prefixes at IXPs
- Only accept the prefixes which your peer is entitled to originate
- Never carry a default route on an IXP (or private) peering router

ISP Transit Issues

Many mistakes are made on the Internet today due to incomplete understanding of how to configure BGP for peering at IXPs



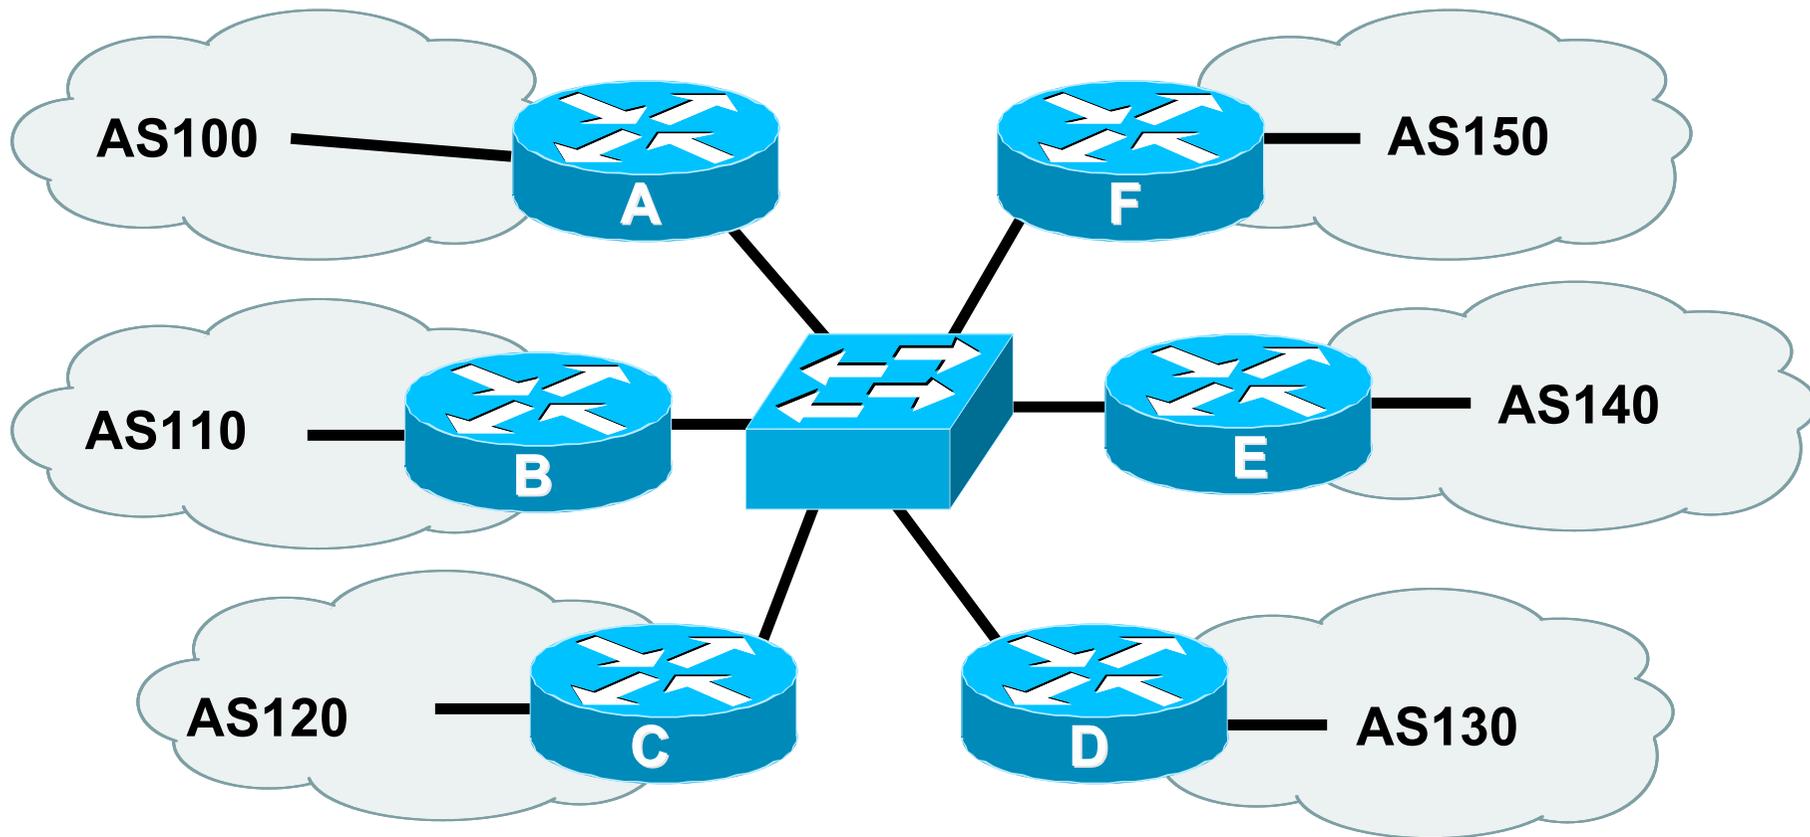
Exchange Points

Simple Example

Exchange Point Example

- Exchange point with 6 ASes present
Layer 2 – ethernet switch
- Each ISP peers with the other
NO transit across the IXP allowed

Exchange Point



- each of these represents a border router in a different autonomous system

Exchange Point Router A configuration

```
interface fastethernet 0/0
  description Exchange Point LAN
  ip address 120.5.10.2 mask 255.255.255.224
  ip verify unicast reverse-path
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
!
router bgp 100
  neighbor ixp-peers peer-group
  neighbor ixp-peers send-community
  neighbor ixp-peers prefix-list myprefixes out
  neighbor ixp-peers route-map set-local-pref in
..next slide
```

Exchange Point

```
neighbor 120.5.10.2 remote-as 110
neighbor 120.5.10.2 peer-group ixp-peers
neighbor 120.5.10.2 prefix-list peer110 in
neighbor 120.5.10.3 remote-as 120
neighbor 120.5.10.3 peer-group ixp-peers
neighbor 120.5.10.3 prefix-list peer120 in
neighbor 120.5.10.4 remote-as 130
neighbor 120.5.10.4 peer-group ixp-peers
neighbor 120.5.10.4 prefix-list peer130 in
neighbor 120.5.10.5 remote-as 140
neighbor 120.5.10.5 peer-group ixp-peers
neighbor 120.5.10.5 prefix-list peer140 in
neighbor 120.5.10.6 remote-as 150
neighbor 120.5.10.6 peer-group ixp-peers
neighbor 120.5.10.6 prefix-list peer150 in
```

Exchange Point

```
ip route 121.10.0.0 255.255.224.0 null0
!
ip prefix-list myprefixes permit 121.10.0.0/19
ip prefix-list peer110 permit 122.0.0.0/19
ip prefix-list peer120 permit 122.30.0.0/19
ip prefix-list peer130 permit 122.12.0.0/19
ip prefix-list peer140 permit 122.18.128.0/19
ip prefix-list peer150 permit 122.1.32.0/19
!
route-map set-local-pref permit 10
  set local-preference 150
!
```

Exchange Point

- Configuration of the other routers in the AS is similar in concept
- Notice inbound and outbound prefix filters
 - outbound announces myprefixes only
 - inbound accepts peer prefixes only
- Notice inbound route-map
 - Set local preference higher than default ensures that local traffic crosses the exchange

Exchange Point

- Ethernet port configuration
 - Use `ip verify unicast reverse-path`
 - Helps prevent “stealing of bandwidth”
- IXP border router must NOT carry prefixes with origin outside local AS and IXP participant ASes
 - Helps prevent “stealing of bandwidth”

Exchange Point

- Issues:

 - AS100 needs to know all the prefixes its peers are announcing

 - New prefixes requires the prefix-lists to be updated

- Alternative solutions

 - Use the Internet Routing Registry to build prefix list

 - Use AS Path filters (could be risky)



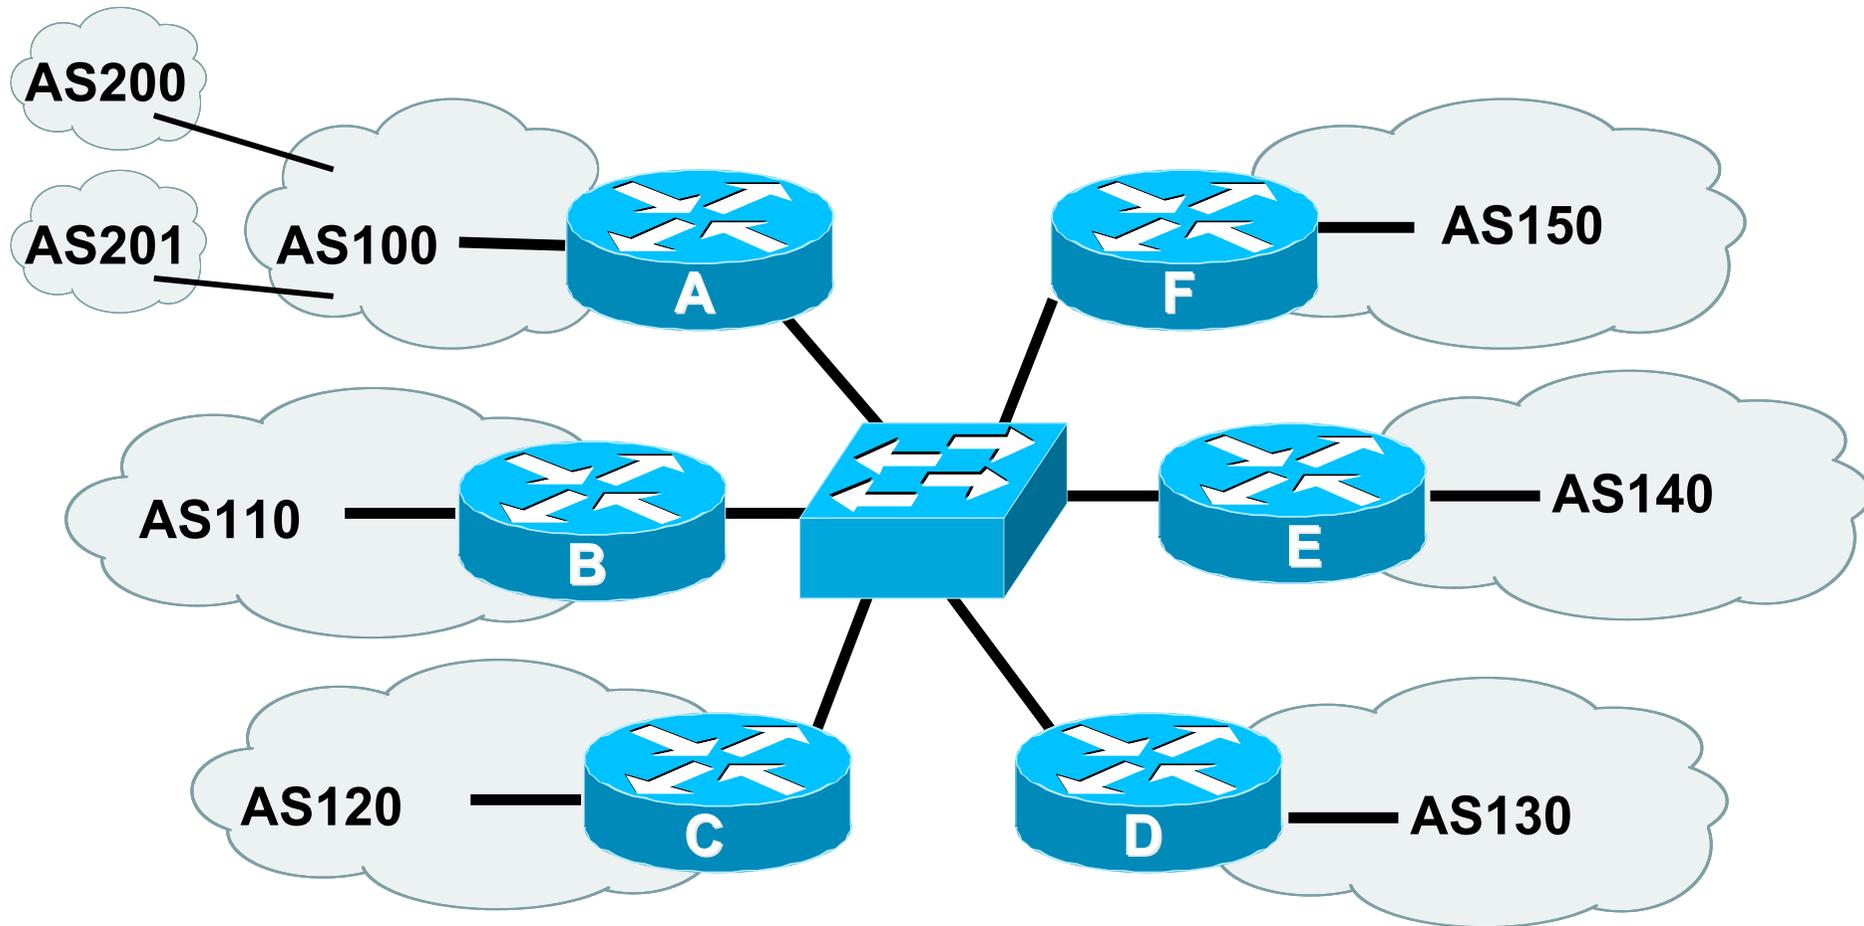
Exchange Points

More Complex Example

Exchange Point Example

- Exchange point with 6 ASes present
 - Layer 2 – ethernet switch
- Each ISP peers with the other
 - NO transit across the IXP allowed
 - ISPs at exchange points provide transit to their customers

Exchange Point



- each of these represents a border router in a different autonomous system

Exchange Point Router A configuration

```
interface fastethernet 0/0
  description Exchange Point LAN
  ip address 120.5.10.2 mask 255.255.255.224
  ip verify unicast reverse-path
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
!
router bgp 100
  neighbor ixp-peers peer-group
  neighbor ixp-peers send-community
  neighbor ixp-peers prefix-list bogons out
  neighbor ixp-peers filter-list 10 out
  neighbor ixp-peers route-map set-local-pref in
..next slide
```

Exchange Point

```
neighbor 120.5.10.2 remote-as 110
neighbor 120.5.10.2 peer-group ixp-peers
neighbor 120.5.10.2 prefix-list peer110 in
neighbor 120.5.10.3 remote-as 120
neighbor 120.5.10.3 peer-group ixp-peers
neighbor 120.5.10.3 prefix-list peer120 in
neighbor 120.5.10.4 remote-as 130
neighbor 120.5.10.4 peer-group ixp-peers
neighbor 120.5.10.4 prefix-list peer130 in
neighbor 120.5.10.5 remote-as 140
neighbor 120.5.10.5 peer-group ixp-peers
neighbor 120.5.10.5 prefix-list peer140 in
neighbor 120.5.10.6 remote-as 150
neighbor 120.5.10.6 peer-group ixp-peers
neighbor 120.5.10.6 prefix-list peer150 in
```

Exchange Point

```
ip route 121.10.0.0 255.255.224.0 null0
!
ip as-path access-list 10 permit ^$
ip as-path access-list 10 permit ^200$
ip as-path access-list 10 permit ^201$
!
ip prefix-list peer110 permit 122.0.0.0/19
ip prefix-list peer120 permit 122.30.0.0/19
ip prefix-list peer130 permit 122.12.0.0/19
ip prefix-list peer140 permit 122.18.128.0/19
ip prefix-list peer150 permit 122.1.32.0/19
!
route-map set-local-pref permit 10
    set local-preference 150
```

Exchange Point

- Notice the change in router A's configuration
 - Filter-list instead of prefix-list permits local and customer ASes out to exchange
 - Prefix-list blocks Special Use Address prefixes – rest get out, could be risky
- Other issues as previously



Internet Exchange Points

BGP configuration for IXPs