

BGP and the Internet

Co-location practises overseas

Why Co-locate Overseas?

Cisco.com

- **Hard to re-terminate transoceanic circuit in case of “issues” with upstream ISP**
- **No Quality of Service**
- **No Control over infrastructure**
- **No Monitoring of link performance**

Co-location Overseas

Cisco.com

- **Many ISPs outside the US and Western Europe co-locate equipment in the US and/or Western Europe**

install their own router(s) and other hardware (servers, caches,...)

establish peering relationships with regional NSPs and domestic ISPs

buy facilities management services

usually hardware maintenance, installation management

Co-location Overseas

Cisco.com

- **Many ISPs outside the US and Western Europe co-locate equipment in the US and/or Western Europe**

US/EU domestic circuits are “cheap”

Easy to change your upstream

Easy to have multiple upstreams

Easy to implement QoS related features, service differentiation, etc...

Co-location

Simple Example

Co-location

- **Common Scenario:**


AS130 has co-locate space in the US

AS120 and AS100 are transit providers for AS130

AS130 is also present at the local exchange point for regional peers

Co-location

Cisco.com



Co-location

Cisco.com

- **AS130**

Router A is dedicated to peering at local IXP

Router G is dedicated to links with the transit providers

Router H is dedicated to the transoceanic link

Co-location Router A Configuration

Cisco.com

```
interface loopback 0
  description Border Router Loopback
  ip address 221.0.0.1 255.255.255.255
!
interface fastethernet 0/0
  description Exchange Point LAN
  ip address 220.5.10.2 255.255.255.224
  ip verify unicast reverse-path
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
!
..next slide
```

Co-location Router A Configuration

Cisco.com

```
interface fastethernet 1/0
  description Crossover 100Mbps Connection to Router G
  ip address 221.0.10.2 255.255.255.252
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
!
```

```
interface fastethernet 2/0
  description Crossover 100Mbps Connection to Router H
  ip address 221.0.10.6 255.255.255.252
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
```

..next slide

Co-location Router A Configuration

Cisco.com

```
router bgp 130
  neighbor ixp-peers peer-group
  neighbor ixp-peers soft-reconfiguration in
  neighbor ixp-peers prefix-list myprefixes out
  neighbor 221.0.0.2 remote-as 130
  neighbor 221.0.0.2 description Router G - Upstream Peers
  neighbor 221.0.0.2 update-source loopback 0
  neighbor 221.0.0.3 remote-as 130
  neighbor 221.0.0.3 description Router H - transpacific router
  neighbor 221.0.0.3 update-source loopback 0
  neighbor 221.0.0.4 remote-as 130
  neighbor 221.0.0.4 description Router at HQ
  neighbor 221.0.0.4 update-source loopback 0
```

..next slide

Co-location Router A Configuration

Cisco.com

```
neighbor 220.5.10.4 remote-as 110
neighbor 222.5.10.4 peer-group ixp-peers
neighbor 222.5.10.4 prefix-list peer110 in
neighbor 220.5.10.5 remote-as 111
neighbor 222.5.10.5 peer-group ixp-peers
neighbor 222.5.10.5 prefix-list peer111 in
neighbor 220.5.10.6 remote-as 112
neighbor 222.5.10.6 peer-group ixp-peers
neighbor 222.5.10.6 prefix-list peer112 in
!
ip prefix-list myprefixes permit 221.10.0.0/19
ip prefix-list peer110 permit 222.12.0.0/19
ip prefix-list peer111 permit 222.18.128.0/19
ip prefix-list peer112 permit 222.1.32.0/19
```

Co-location Router A Configuration

Cisco.com

- **Router A does NOT originate AS130's prefix block**
if router is disconnected from AS130 either locally or across the ocean, announcement could cause blackhole
- **Prefix-list filtering is the minimum required**
usually include AS path filtering too

Co-location Router G Configuration

Cisco.com

```
interface loopback 0
  description Peering Router Loopback
  ip address 221.0.0.2 255.255.255.255
!
interface fastethernet 0/0
  description Crossover 100Mbps Connection to Router A
  ip address 221.0.10.1 255.255.255.252
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
!
```

..next slide

Co-location Router G Configuration

Cisco.com

```
interface hssi 1/0
  description T3 link to BigISP
  ip address 222.0.0.2 255.255.255.252
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
!
```

```
interface hssi 2/0
  description T3 link to MegaISP
  ip address 218.6.0.2 255.255.255.252
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
```

..next slide

Co-location Router G Configuration

Cisco.com

```
router bgp 130
  neighbor 221.0.0.1 remote-as 130
  neighbor 221.0.0.1 description Router A - US Local IXP
  neighbor 221.0.0.1 update-source loopback 0
  neighbor 221.0.0.1 prefix-list myprefixes out
  neighbor 221.0.0.3 remote-as 130
  neighbor 221.0.0.3 description Router H - transpacific router
  neighbor 221.0.0.3 update-source loopback 0
  neighbor 221.0.0.4 remote-as 130
  neighbor 221.0.0.4 description Router at HQ
  neighbor 221.0.0.4 update-source loopback 0
```

..next slide

Co-location Router G Configuration

Cisco.com

```
neighbor 222.0.0.1 remote-as 120
neighbor 222.0.0.1 prefix-list myprefixes out
neighbor 222.0.0.1 prefix-list rfc1918-sua in
neighbor 218.6.0.1 remote-as 100
neighbor 218.6.0.1 prefix-list myprefixes out
neighbor 218.6.0.1 prefix-list rfc1918-sua in
!
ip prefix-list myprefixes permit 221.10.0.0/19
```

Co-location

Router G Configuration

Cisco.com

- **Router G accepts full BGP prefixes from both AS120 and AS100**
- **Router G announces AS130 prefix to upstreams**
- **Simple Example – policy may also be required for loadsharing etc**

Co-location Router H Configuration

Cisco.com

```
interface loopback 0
  description Peering Router Loopback
  ip address 221.0.0.3 255.255.255.255
!
interface fastethernet 0/0
  description Crossover 100Mbps Connection to Router A
  ip address 221.0.10.5 255.255.255.252
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
!
```

..next slide

Co-location Router H Configuration

Cisco.com

```
interface hssi 1/0
  description T3 link back to home
  ip address 221.1.0.1 255.255.255.252
  rate-limit output access-group 195 ..etc
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
!
```

..next slide

Co-location Router H Configuration

Cisco.com

```
router bgp 130
  neighbor 221.0.0.1 remote-as 130
  neighbor 221.0.0.1 description Router A - US Local IXP
  neighbor 221.0.0.1 update-source loopback 0
  neighbor 221.0.0.2 remote-as 130
  neighbor 221.0.0.2 description Router G - peering router
  neighbor 221.0.0.2 update-source loopback 0
  neighbor 221.0.0.4 remote-as 130
  neighbor 221.0.0.4 description Router at HQ
  neighbor 221.0.0.4 update-source loopback 0
!
```

Co-location Router H Configuration

Cisco.com

- **Router H is dedicated to transoceanic link**
part of ISP core iBGP mesh
- **More complex configuration likely**
CAR, RED, etc
- **More complex links likely**
e.g satellite uplink for low revenue latency insensitive traffic

Co-location

Cisco.com

- **Richer interconnectivity possible**
- **Better redundancy possible**
- **Overall advantage – control!**

BGP and the Internet

Co-location practises overseas