BGP Configuration for International Co-location

ISP Workshops

Agenda

- Simplistic Transit Model
- Why place equipment in remote transit points?
- And how to configure BGP to handle this

Simplistic Transit Model (1)

- Conventional wisdom is that an ISP buys a circuit to and service from their transit provider
  - This circuit 'plugs into' the transit provider router
  - The ISP pays transit provider for all Internet traffic
- Unfortunately this is a very naïve view
  - And long term very expensive for the ISP

Simplistic Transit Model (2)

- ISP locked into their transit provider for:
  - Service
  - Support
  - Performance
  - Reliability
  - Internet access costs
- No redundancy should there be a problem with their transit provider's network

Simplistic Transit Model (3)

- Hard to re-terminate international leased circuit in case of "issues" with transit ISP
  - Takes time (days, weeks,…)
  - Means service disruption
- No Quality of Service
  - Not possible to differentiate services
- No Control over infrastructure
  - Traffic that you may not want traverses your most expensive link
- No Monitoring of link performance
  - View of one end of the link only

Why Invest in International Colo?

- "International" means outside the local or regional Internet presence
- Bandwidth saving at the transit edge
  - Content filtering & caching
  - Security filters for common misconfigurations
  - Email washing (anti-spam, anti-virus)
- Security at the transit edge
  - Border filters
  - DDOS attack protection before impacting international link

Why Invest in International Colo?

- Reliability & performance
  - Choice of transit providers & service quality
  - Migration between transit providers without breaking service
- Cost reduction
  - Opportunity participate at IXPs rather than paying transit costs
  - Opportunity to peer privately
  - Opportunity to seek most cost-effective transit provider

International Co-location

- Many ISPs invest in international co-location facilities
  - They install equipment at major co-lo's including:
    - London
    - Amsterdam
    - Frankfurt
    - New York/Washington/Miami
    - Seattle/San Francisco/Los Angeles
    - Hong Kong
    - Singapore

International Co-location

- Installations include:
  - Their own router(s)
  - Other hardware (servers, caches,…)
  - Buying transit at domestic rates from transit providers
  - Establishing peering relationships with regional NSPs and domestic ISPs
    - Privately
    - At Internet Exchange Points
  - Buy facilities management services, usually hardware maintenance, installation management, etc

International Co-location

- Benefits include:
  - US/EU domestic circuits are “cheap”
  - Easy to change transit provider
  - Easy to have multiple transits
  - Major cost reduction through peering rather than paying for transit
    - Over 60% of traffic can be obtained via peering
  - Easy to implement traffic and content filtering, security and 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
- Equipment Provision
  - One router for link back home
  - One router for IXP
  - One router for paid Transit
  - (Other servers for web caching, mail washing, bandwidth management, &c)

Co-location

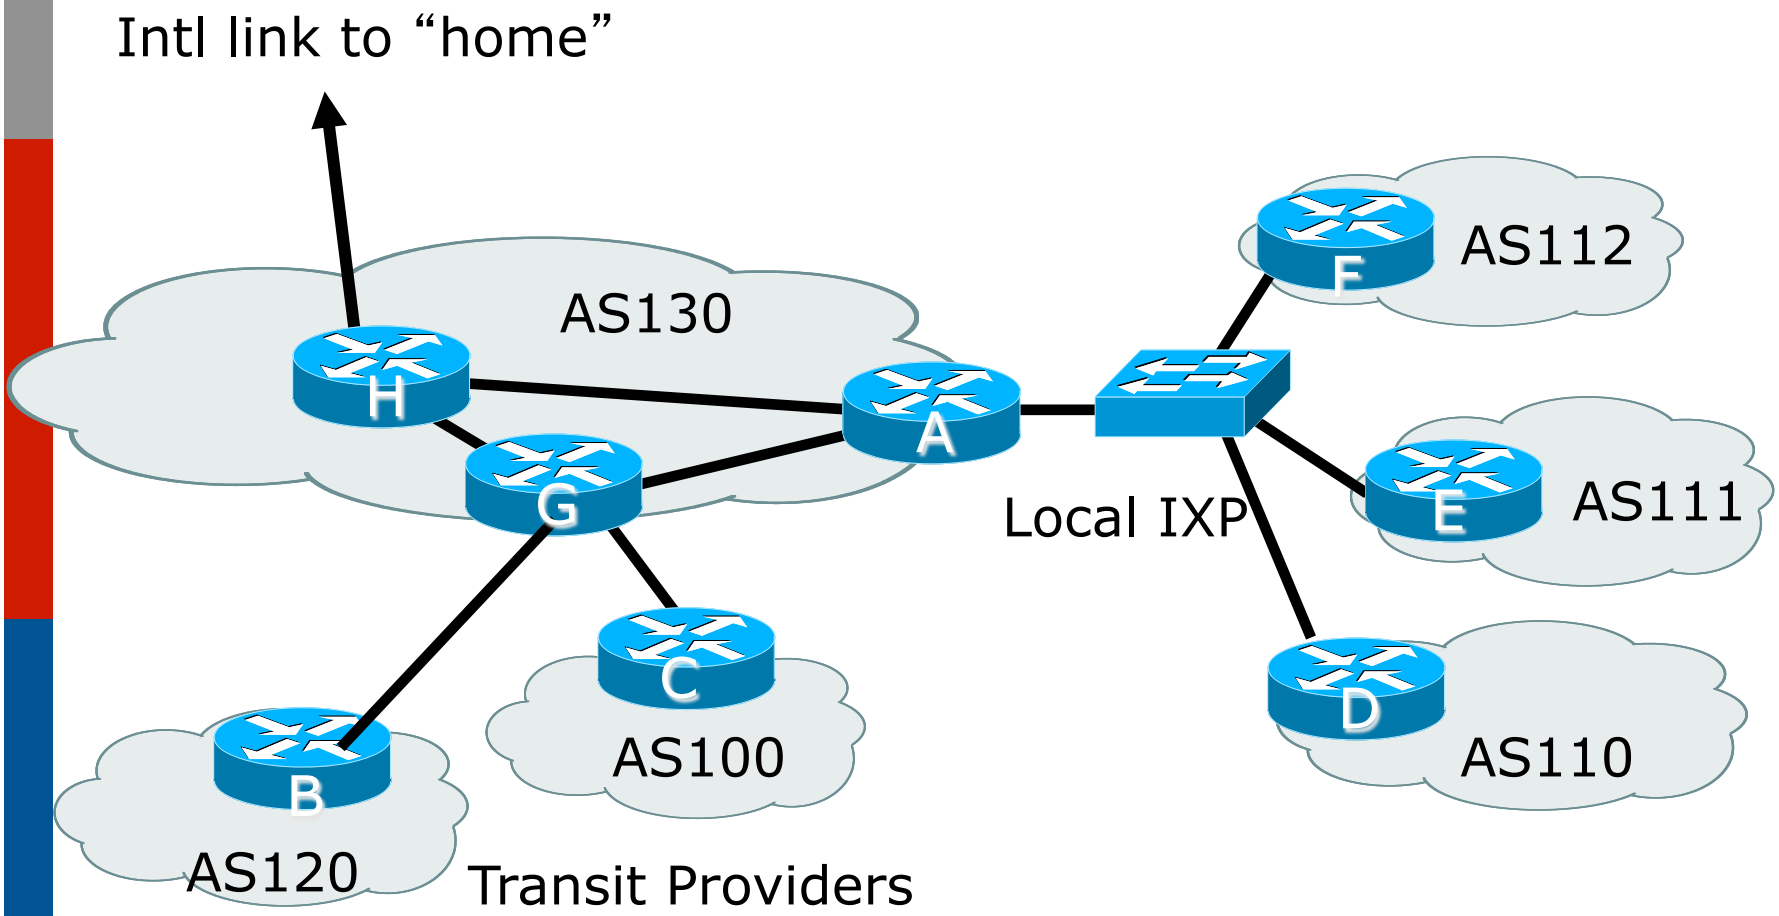
Intl link to “home”
AS130
H
A
G
F
AS112
Local IXP
E
AS111
C
AS100
D
AS110
B
AS120
Transit Providers

Co-location

- 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
  - Is route reflector for Router A and G
  - Is RR client off a route reflector back home
- Three routers means that in the event of failure of one, the other two can provide temporary backup until repairs are completed

Co-location
Router A Configuration

```
interface loopback 0
 description Border Router Loopback
 ip address 221.0.0.1 255.255.255.255
!
interface gigabitethernet 0/0
 description Exchange Point LAN
 ip address 220.5.10.2 255.255.255.224
 no ip directed-broadcast
 no ip proxy-arp
 no ip redirects
!
```

...next slide

Co-location
Router A Configuration

```
interface gigabitethernet 1/0
 description Crossover 1Gbps 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 gigabitethernet 2/0
 description Crossover 1Gbps 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

```
router bgp 130
 bgp deterministic-med
 neighbor ixp-peers peer-group
 neighbor ixp-peers prefix-list myprefixes out
 neighbor rr peer-group
 neighbor rr remote-as 130
 neighbor rr update-source loopback 0
 neighbor rr send-community
 neighbor 221.0.0.3 peer-group rr
 neighbor 221.0.0.3 description Router H - Intl Link
 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
```

...next slide

Co-location
Router A Configuration

```
 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
!
ip route 221.10.0.0 255.255.224.0 null0 250
```

Co-location
Router A Configuration

- Router A does NOT originate AS130's prefix block
  - If router is disconnected from AS130 either locally or across the international link, the announcement will blackhole AS130's entire network
  - Static route to null0 for AS130's address block performs integrity function
- Prefix-list filtering is the minimum required
  - Usually include AS path filtering too

Co-location
Router G Configuration

```
interface loopback 0
 description Peering Router Loopback
 ip address 221.0.0.2 255.255.255.255
!
interface gigabitethernet 0/0
 description Crossover 1Gbps 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

```
interface POS 1/0
 description STM-1 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 POS 2/0
 description STM-1 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

```
router bgp 130
 neighbor rr peer-group
 neighbor rr remote-as 130
 neighbor rr update-source loopback 0
 neighbor rr send-community
 neighbor 221.0.0.3 peer-group rr
 neighbor 221.0.0.3 description Router H - Intl Link
```

..next slide

Co-location
Router G Configuration

```
 neighbor 222.0.0.1 remote-as 120
 neighbor 222.0.0.1 prefix-list myprefixes out
 neighbor 222.0.0.1 prefix-list bogons 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 bogons in
!
ip prefix-list myprefixes permit 221.10.0.0/19
!
ip route 221.10.0.0 255.255.224.0 null0 250
```

Co-location
Router G Configuration

- Multihoming to upstreams:
  - 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
  - Obviously this can and should be refined, as per multihoming recommendations covered earlier

Co-location
Router H Configuration

```
interface loopback 0
 description Peering Router Loopback
 ip address 221.0.0.3 255.255.255.255
!
interface gigabitethernet 0/0
 description Crossover 1Gbps 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

```
interface POS 1/0
 description STM-1 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

```
router bgp 130
 neighbor rr-client peer-group
 neighbor rr-client remote-as 130
 neighbor rr-client update-source loopback 0
 neighbor rr-client send-community
 neighbor 221.0.0.1 peer-group rr-client
 neighbor 221.0.0.1 description Router A - US IXP
 neighbor 221.0.0.2 peer-group rr-client
 neighbor 221.0.0.2 description Router G - US transit
 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

- Router H is dedicated to transoceanic link
  - Router reflector for the other routers in the overseas PoP
  - Client of route reflector in the ISP's domestic backbone
- More complex configuration likely
  - CAR, RED, etc
- More complex links likely
  - e.g satellite uplink for low revenue latency insensitive traffic

Co-location

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

BGP Configuration for International Co-location

ISP Workshops