Peering, Transit and IXPs

Philip Smith

MENOG 11

Amman

30th September – 9th October 2012

Background

- Presentation discusses:
  - Transit
  - Peering
  - Internet Exchange Points
- Tracking the growth of an Internet access provider
  - Transit
  - Adding Peering
  - Participating in an IXP

Transit

- A network operator which provides access to other parts of the Internet
  - Local/Regional
  - More usually "The Whole Internet"
- Transit providers must be chosen wisely:
  - Only one = no redundancy
  - Too many:
    - No economy of scale
    - Traffic engineering is very difficult
    - Service quality very difficult to provide
- Recommendation:
  - **At least two, not more than three**

Common Mistakes

- Signing up with too many transit providers
  - Lots of small circuits
    - These cost more per Mbps than larger ones
    - Transit rates per Mbps reduce with increasing transit bandwidth purchased
  - Hard to implement reliable traffic engineering
    - High operational overhead fine tuning peering arrangements
    - Serious service quality concerns due to more frequent path changes and “Internet” disruptions

Common Mistakes

- No diversity for chosen transit providers
  - They are not diversely connected on local operator's network backbone
  - All reached over same satellite or same submarine cable
  - All connect to the same upstream
  - All have poor onward transit and peering arrangements

Peer

- A peer is another autonomous system with which the local network has agreed to exchange locally sourced routes and traffic
- Private peer
  - Private link between two providers for the purpose of interconnecting
- Public peer
  - Internet Exchange Point, where providers meet and **freely** decide who they will interconnect with
- **Recommendation: peer as much as possible!**

Peering Advice

- Analyse traffic sources and use that knowledge to determine peers
- Self-list in the Peering Database
  - www.peeringdb.com
- Participate in the various regional and Global peering fora
  - www.peeringforum.com
- Work as hard as possible to get as much peering as possible
  - No traffic costs (usually)
  - Consider transit as last resort – it costs money!

Common Mistakes

- Mistaking a transit provider's "Exchange" business for a no-cost public peering point
- Being physically close to a public peering point but not participating
- Ignoring/avoiding competitors because they are competition
  - Even though potentially valuable peering partner to give customers a better experience

Types of Interconnect

- Private Interconnect
  - Where two network operators agree to share costs of a direct interconnection
  - Exchange their local routes/traffic
  - No traffic costs
- Public Interconnect
  - Where a network operator participates at an Internet Exchange Point, interconnecting with other network operators
  - Exchange routes/traffic with other peers
  - No traffic costs

Types of IXP peering

- Bi-lateral peering
  - Like private peering, two operators agree to interconnect their networks, but over the IXP fabric
- Multi-lateral peering
  - Operator peers with the IXP route server
  - Route server sends all routes it knows to the operator
  - Route server sends operator's routes to all other operators peering with route server
- Most IXPs provide the opportunity for participants to use both bi-lateral and multi-lateral peering

Why an Internet Exchange Point?

Saving money, improving service quality, encouraging a local Internet economy

Internet Exchange Point
Why peer?

- Consider a region with one ISP
  - They provide internet connectivity to their customers
  - They have one or two international connections
- Internet grows, another ISP sets up in competition
  - They provide internet connectivity to their customers
  - They have one or two international connections
- How does traffic from customer of one ISP get to customer of the other ISP?
  - Via the international connections

Internet Exchange Point
Why peer?

- Yes, International Connections…
  - If satellite, RTT is around 550ms per hop
  - So local traffic takes over 1s round trip
- International bandwidth
  - Costs significantly more than domestic bandwidth
  - Congested with local traffic
  - Wastes money, harms performance

Internet Exchange Point
Why peer?

- Solution:
  - Two competing ISPs peer with each other
- Result:
  - Both save money
  - Local traffic stays local
  - Better network performance, better service quality,…
  - More international bandwidth for expensive international traffic
  - Everyone is happier

Internet Exchange Point
Why peer?

- A third ISP enters the equation
  - Becomes a significant player in the region
  - Local and international traffic goes over their international connections
- They agree to peer with the two other ISPs
  - To save money
  - To keep local traffic local
  - To improve network performance, service quality,…

Internet Exchange Point
Why peer?

- Peering means that the three ISPs have to buy circuits between each other
  - Works for three ISPs, but adding a fourth or a fifth means this does not scale
- Solution:
  - Internet Exchange Point

Internet Exchange Point

- Every participant has to buy just one whole circuit
  - From their premises to the IXP
- Rather than N-1 half circuits to connect to the N-1 other ISPs
  - 5 ISPs have to buy 4 half circuits = 2 whole circuits → already twice the cost of the IXP connection

Internet Exchange Point

- Solution
  - Every ISP participates in the IXP
  - Cost is minimal – one local circuit covers all domestic traffic
  - International circuits are used for just international traffic – and backing up domestic links in case the IXP fails
- Result:
  - Local traffic stays local
  - Service quality considerations for local traffic is not an issue
  - RTTs are typically sub 10ms
  - Customers enjoy the Internet experience
  - Local Internet economy grows rapidly

How to start?

- It needs three network operators to agree:
  - To interconnect their networks
  - A common neutral location for the IX
  - To share costs:
    - Infrastructure (data centre, rack, switch, power, a/c)
    - Operational (data centre, switch management)
  - Basic behavioural rules (MoU)
- And that's really all there is to it

How to scale?

- Start up model works well for a few participants (<10)
- After that, need to consider:
  - Cost recovery model of the IXP
  - Data centre value
  - Permanent staffing arrangement
  - Ethernet switch & other network equipment
  - Scaling the peering arrangements
  - Governance: i.e. consortium/management board

Other Opportunities

- IXP is primarily about facilitating local peering
- But other entities are interested in IXPs too:
  - Content providers
    - Lower transit costs, fast delivery, better end-user experience
  - Root nameserver operators
    - Local instance of F, I, K, L, etc
  - ccTLD and gTLD operators
    - Domestic ccTLD is priority

Other Services

- Other services can be provided:
  - Time synchronisation (ntp)
  - Route Collector
    - Marketing tool for IXP
    - Troubleshooting tool for ISPs and global Internet
  - Route Server
    - Scales BGP peering at IXP
- Services should avoid competing with the membership

Adding more participants?

- With an established IX:
  - Content providers connected
  - Root nameserver operator present
  - Existing participants have superior domestic internet performance
- Non-participants miss out on benefits
  - Motivated to join
  - Customer word of mouth is powerful
  - Especially when local content delivery is superior via IXP connected participants ISPs

Scaling further?

- IXP becomes "critical infrastructure" for local Internet traffic
- How to scale:
  - ISPs bring second router (for redundancy)
  - Second switch (for redundancy)
  - Second site (for redundancy)

Other issues

- Obtaining unanimity in the local industry before setting up the IX is usually impossible
  - Three network operators are all that are needed to start an IX
- Technically the IX is very simple to set up
  - Ethernet switch, one router per ISP, and eBGP
- Politically the IX could be complicated to set up
  - Participants try and gain advantage over others
  - Government or Regulator may want to operate it
  - Incumbent telco is usually last to participate

Advice on IX construction

- Establish local peering before being forced by Government to do so
- Avoid:
  - Complex rules and stifling bureaucracy
  - Complex cost models and barriers to entry
- Obtain minimum critical mass
- Get the IXP established technically (easy!)
- Lobby content providers, root nameserver operators and the local ccTLD to participate

Conclusion

- Tracked the growth of an Internet access provider
  - Transit gives global Internet connectivity – traffic costs
  - Peering – no traffic costs, reduced dependency on Transit
  - IXP – scalable Peering (no traffic costs), essential for a growing Internet economy