What is an IXP?

What is an Internet eXchange Point (IXP)?

- The Internet is an interconnection of networks
  - Each controlled by separate entities
    - Generally called Internet Service Providers (ISPs)
  - Grouped by Autonomous System (AS) numbers
- Transit
  - Where ISP pays to send/receive traffic
    - Downstream ISP will pay upstream ISP for transit service
- Peering
  - ISPs will (generally) not pay each other to exchange traffic
  - Works well if win win for both
  - Reduce cost on expensive transit link

Peering and Transit example

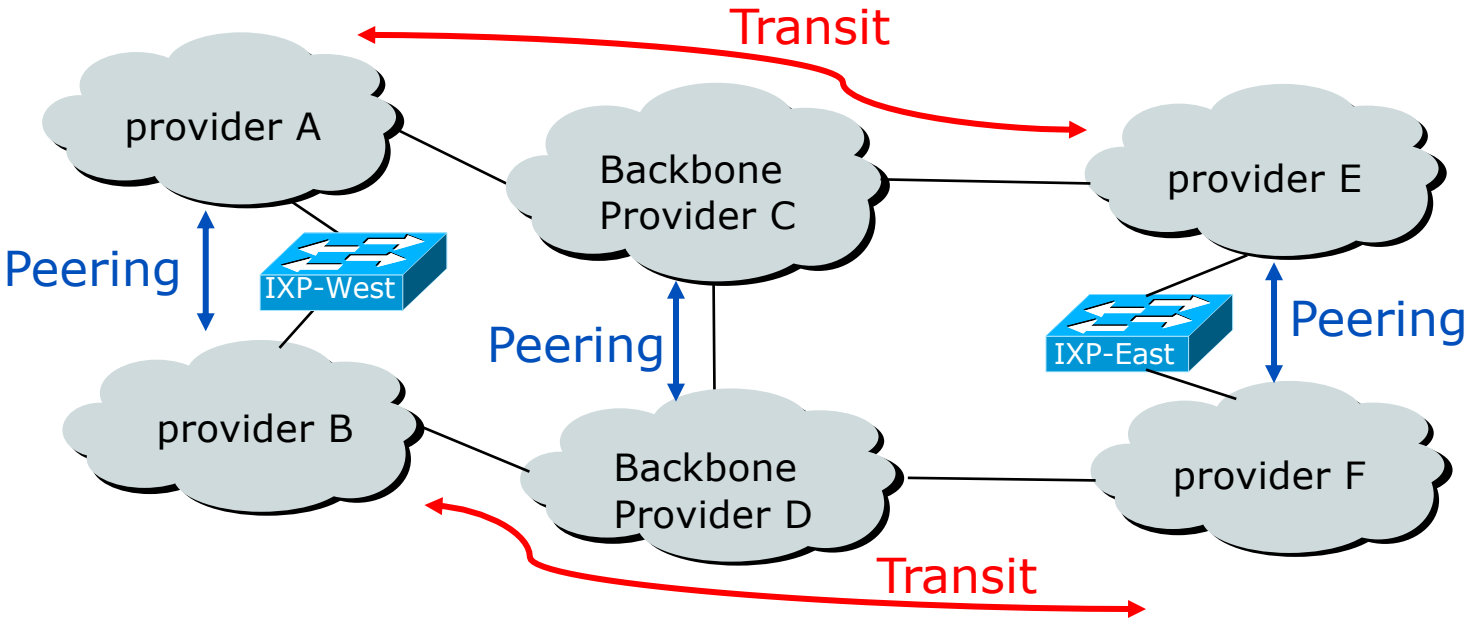

A and B peer for free, but need transit arrangements with C and D to get packets to/from E and F

Private Interconnect

- Two ISPs connect their networks over a private link
  - Could be a peering arrangement
    - No charges for traffic
    - Share the cost of the link
  - Could be transit
    - One ISP charges the other for traffic (and also for the link)

Private Interconnect

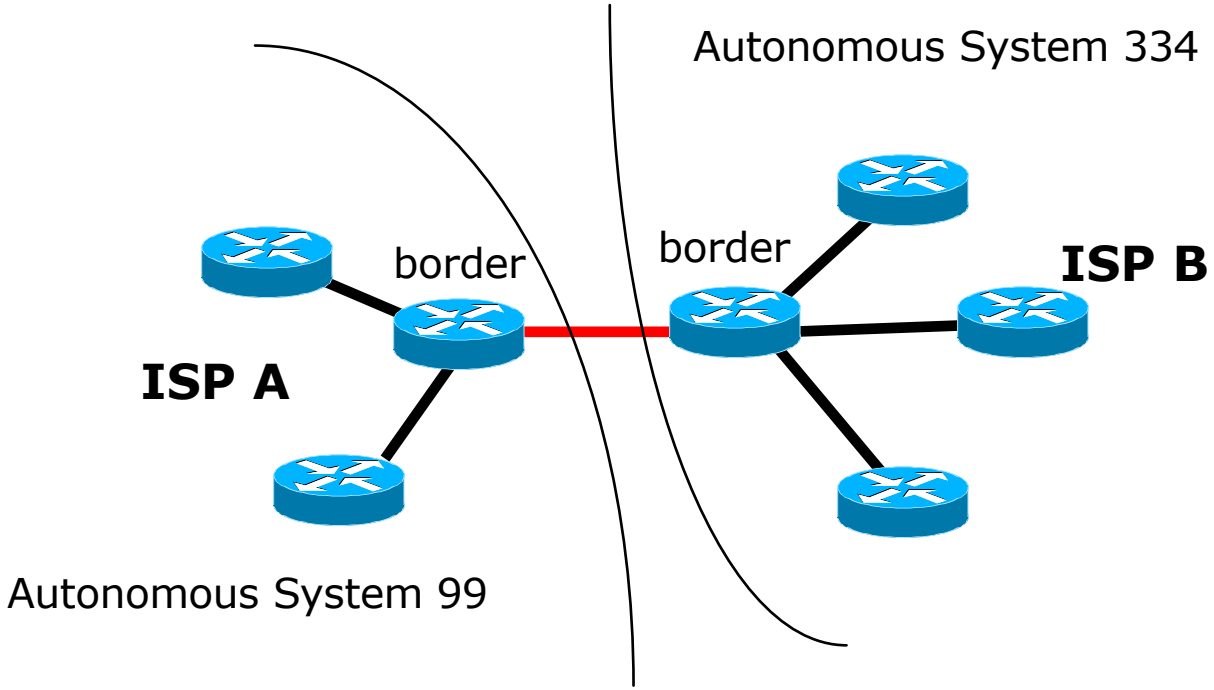

Public Interconnect

- A location or facility where several ISPs are present and connect to each other over a common shared media
  - Ex: Ethernet
- Why?
  - To save money, reduce latency, improve performance
- IXP – Internet eXchange Point
- Each provider establishes **peering** relationships with other providers at the IXP

Public Interconnect

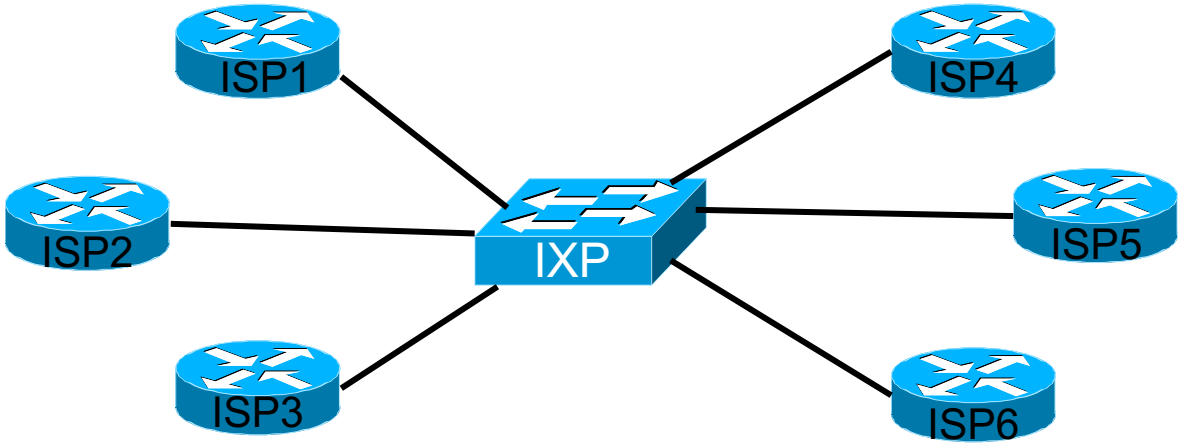

- Border routers in different Autonomous Systems

Global Internet: High Level View

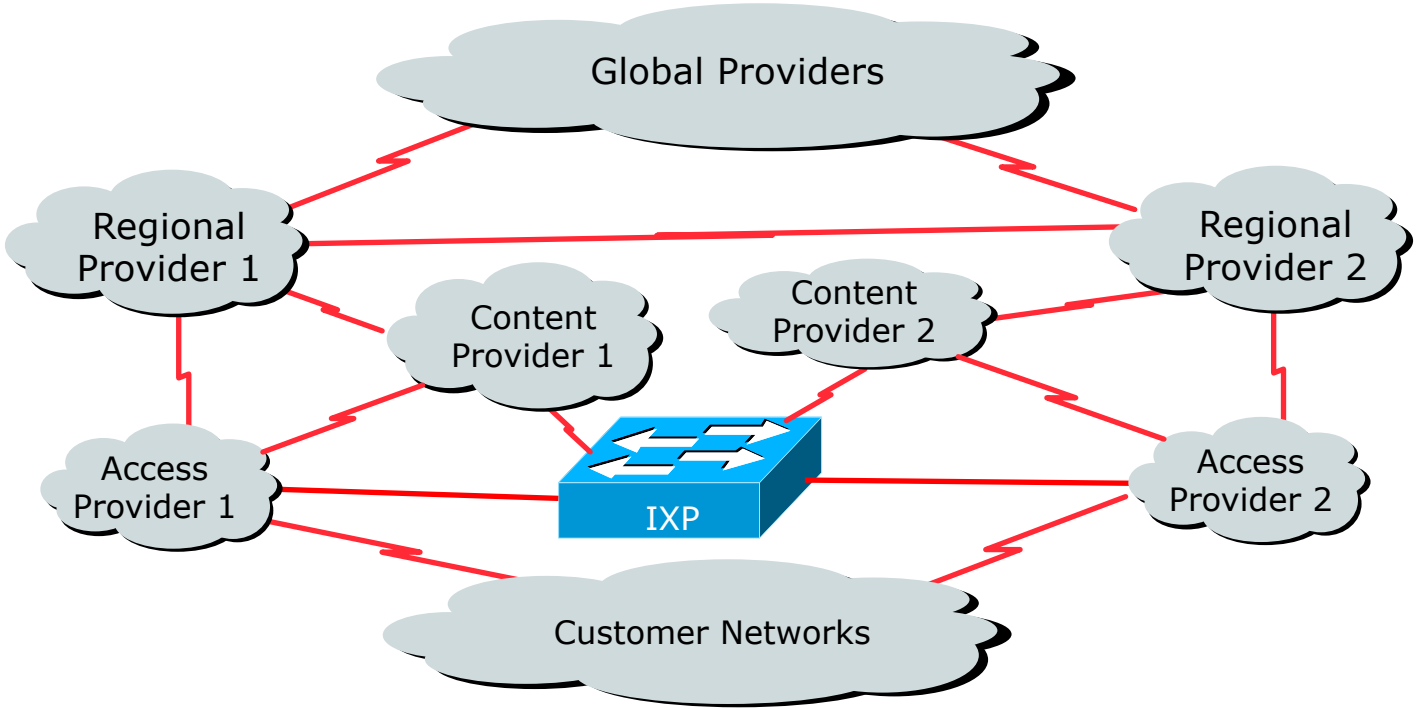

Detailed View: Global Internet

- Global Transit Providers
  - Connect to each other
  - Provide connectivity to Regional Transit Providers
- Regional Transit Providers
  - Connect to each other
  - Provide connectivity to Content Providers
  - Provide connectivity to Access Providers
- Content Providers
  - Cross-connect with Access Providers
  - Peer at IXPs (free traffic to Access Providers)
- Access Providers
  - Connect to each other across IXPs (free peering)
  - Provide access to the end user

Categorizing Network Operators

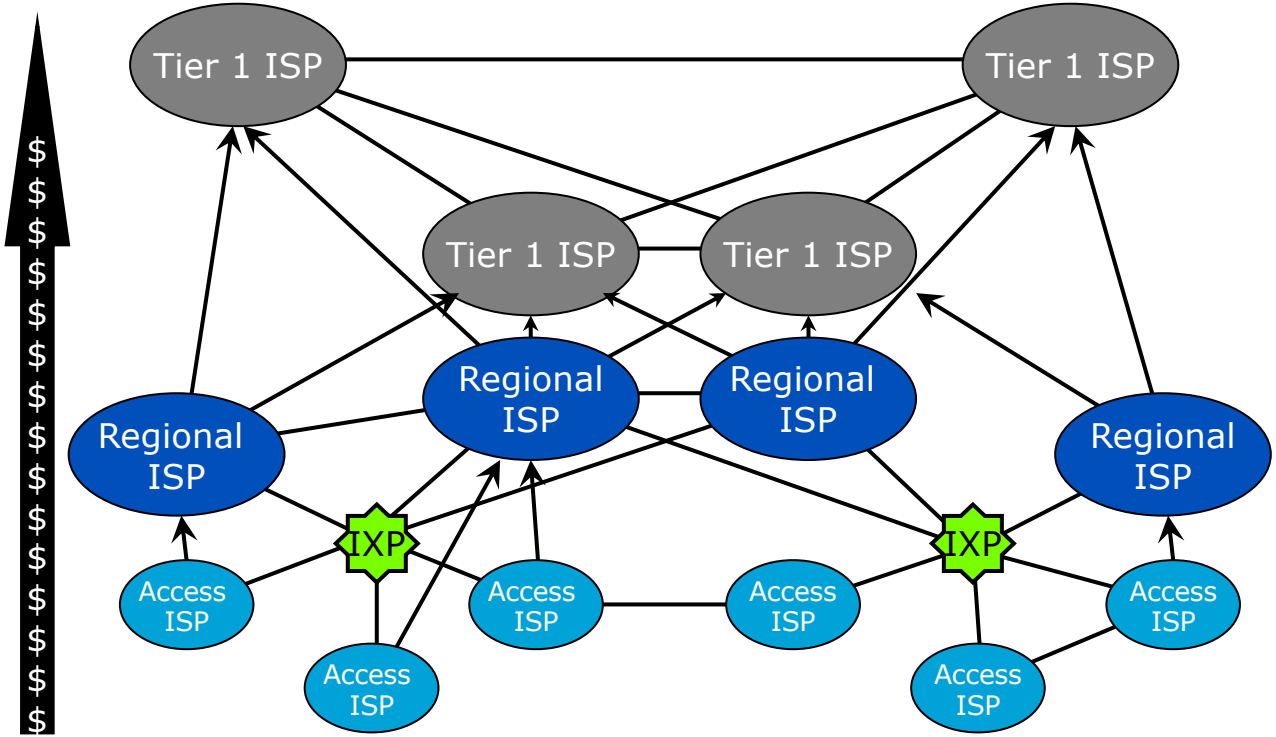

Categorising Network Operators

- Tier-1 Providers
  - A provider that peers with other Tier-1s and does NOT pay for transit
  - Caution:
    - Many ISPs market themselves as Tier-1 even though they may be paying for transit themselves to reach some parts of the Internet
- Regional Providers often have the same reach as Tier-1s but still rely on one or two Tier-1s to reach the whole Internet
  - Often provide access too (in-country access networks)
- Access Providers provide connectivity in their locale

Inter-provider relationships

- Peering between equivalent sizes of service providers (e.g. Regional to Regional)
  - Shared cost private interconnection, equal traffic flows
  - No cost peering
- Peering across eXchange points
  - If convenient, of mutual benefit, technically feasible
- Fee based peering
  - Unequal traffic flows, “market position”

Default Free Zone

- NOT related to where an ISP is in the hierarchy!

- *Made up of Internet routers which have explicit routing information about the rest of the Internet, and hence, do not need to use a default route!*

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 overall performance (end-user experience)

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 QoS,…
  - More international bandwidth for expensive international traffic
  - Everyone is happy

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
- All agree to peer with each other
  - To save money
  - To keep local traffic local
  - To improve network performance, QoS,…

Internet eXchange Point - Why peer?

- Private 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 fabric

- 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
  - QoS considerations for local traffic is not an issue
  - RTTs are typically sub 10ms
  - Customers enjoy the Internet experience
  - Local Internet economy grows rapidly

Summary - Objectives of IXP

- Main objective of an IXP is to keep local traffic local
- It also helps bypass 3rd-party network infrastructure for easy interconnections and direct traffic exchange among participating networks
  - Reduced cost
  - Enhanced network performance
  - Reduced latency
- Every participant is benefited
  - Some may gain more, some may gain less but all will gain

Peering Aside

- About 55-60K network operators globally
- Only ~7.5K are transit providers
  - Only about 5 are global (Tier-1s) providers
- What does this say?
  - Peering is what makes the Internet work!

Internet eXchange Points

- Layer 2 exchange point
  - Ethernet (100Gbps/10Gbps/1Gbps/100Mbps)

Layer 2 Exchange

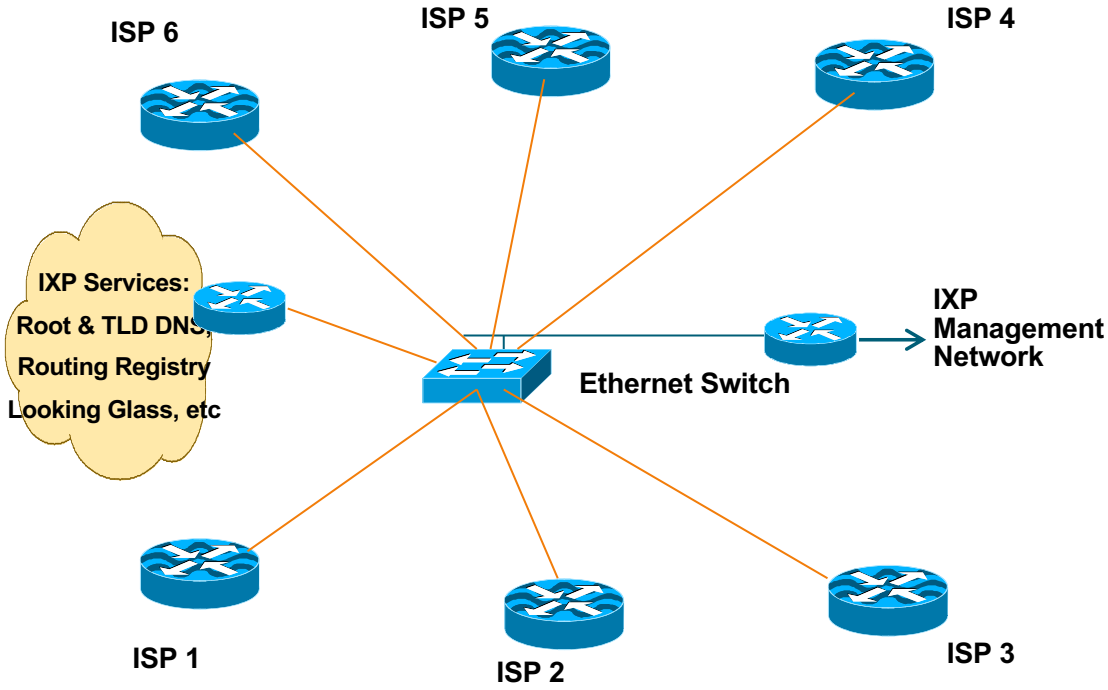

Layer 2 Exchange

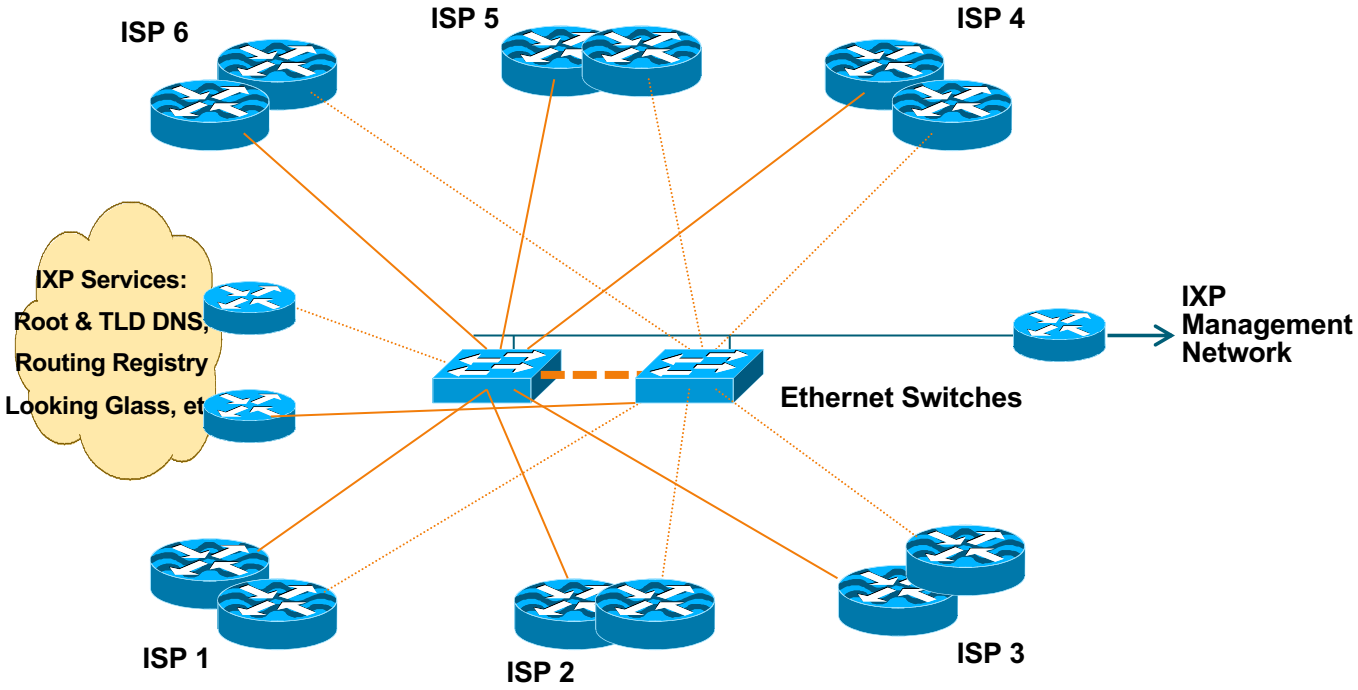

Layer 2 Exchange

- Two switches for redundancy
- ISPs use dual routers for redundancy or loadsharing
- Offer services for the “common good”
  - Internet portals and search engines
  - DNS Root & TLDs, NTP servers
  - Routing Registry and Looking Glass

Layer 2 Exchange

- Requires neutral IXP management
  - Usually funded equally by IXP participants
  - 24x7 cover, support, value add services
- Secure and neutral location
- Configuration
  - Private address space if non-transit and no value add services
  - Otherwise public IPv4 (/24) and IPv6 (/48, /56, /64)
  - ISPs require AS, basic IXP does not

Layer 2 Exchange

- Network Security Considerations
  - LAN switch needs to be securely configured
  - Management routers require TACACS+ authentication, vty security
  - IXP services must be behind router(s) with strong filters

The Value of Peering

Types of Peering

- Private Peering
  - Two operators agree to interconnect and exchange their respective routes to ensure their customers can reach other directly over the peering link
- Settlement Free Peering
  - No traffic charges
  - The most common form of peering!
- Paid Peering
  - Where two operators agree to peer and charges for carrying traffic

Types of Peering

- Bilateral Peering
  - Similar to private peering but may take place at a public peering point (IXP)
- Multilateral Peering
  - Takes place at an IXP, where operators peer with each other via a route server
- Mandatory Multilateral Peering
  - Where all operators are forced to peer with each other (as a condition of IXP membership)
  - Strongly discouraged: no record of success!

Types of Peering

- Open Peering
  - An ISP publicly states that they will peer with anyone who approaches them for peering
  - Commonly found at IXPs, where ISPs participate via the Route Server
- Selective Peering
  - An ISP's peering policy depends on who requests peering with them
  - At an IXP, the operator will not peer with the RS but will only peer bilaterally
- Closed Peering
  - Where an ISP decides who its peering partners are
  - Generally not approachable to creating peering opportunities

ISP Goals

- Minimise the cost of operating the business
- Transit
  - ISP has to pay for circuit (international or domestic)
  - ISP has to pay for data (usually per Mbps)
  - Repeat for each transit provider
  - Significant cost of being a service provider
- Peering
  - ISP shares circuit cost with peer (private) or runs circuit to public peering point (one off cost)
  - No need to pay for data
  - Reduces transit data volume, therefore reducing cost

Transit – How it works

- Small access provider provides Internet access for a city's population
  - Mixture of wireless and fixed broadband
  - Possibly some business customers
  - Possibly also some Internet cafes
- How do their customers get access to the rest of the Internet?
- ISP buys access from one, two or more larger ISPs who already have visibility of the rest of the Internet
  - This is transit – they pay for the physical connection to the upstream and for the traffic volume on the link

Peering – How it works

- If two ISPs are of equivalent sizes, they have:
  - Equivalent network infrastructure coverage
  - Equivalent customer size
  - Similar content volumes to be shared with the Internet
  - Potentially similar traffic flows to each other's networks

- This makes them good peering partners

- If they don't peer
  - They both have to pay an upstream provider for access to each other's network/customers/content
  - Upstream benefits from this arrangement, the two ISPs both have to fund the transit costs

Which IXP?

- How many routes are available?
  - How many other operators/providers are at the IX?
  - What is the traffic to and from these destinations, and how much will it reduce the transit cost?
- What is the cost of co-lo space?
  - If prohibitive or space not available, pointless to choose this IXP
- What is the cost of a circuit to the location?
  - If similar to transit costs, pointless
- What is the cost of remote-hands?
  - For maintenance purposes to avoid serious outages

Example: South Asian ISP @ LINX

- Date: May 2013

- Data:
  - Route Server plus bilateral peering offers 70k prefixes
  - IXP traffic averages 247Mbps/45Mbps
  - Transit traffic averages 44Mbps/4Mbps

- Analysis:
  - 85% of inbound traffic comes from 70k prefixes available by peering
  - 15% of inbound traffic comes from remaining 380k prefixes from transit provider

Example: South Asian ISP @ HKIX

- Date: May 2013

- Data:
  - Route Server plus bilateral peering offers 67k prefixes
  - IXP traffic is 159Mbps/20Mbps
  - Transit traffic is 108Mbps/50Mbps

- Analysis:
  - 60% of inbound traffic comes from 67k prefixes available by peering
  - 40% of inbound traffic comes from remaining 383k prefixes from transit provider

Example: South Asian ISP

- Summary:
  - Traffic by Peering: 406Mbps/65Mbps
  - Traffic by Transit: 152Mbps/54Mbps

  - 73% of incoming traffic is by peering
  - 55% of outbound traffic is by peering

Example: South Asian ISP

- Router at remote co-lo
  - Benefits: can select peers, easy to swap transit providers
  - Costs: co-lo space and remote hands

- Overall advantage:
  - Can control what goes on the expensive connectivity “back to home”

Value propositions

- Peering at a local IXP
  - Reduces latency & transit costs for local traffic
  - Improves Internet quality perception
- Participating at a Regional IXP
  - A means of offsetting transit costs
- Managing connection back to home network
- Improving Internet Quality perception for customers

Summary

- Benefits of peering
  - Private
  - Internet Exchange Points

- Local versus Regional IXPs
  - Local services local traffic
  - Regional helps defray transit costs

Worked Example

Single International Transit

Versus

Local IXP + Regional IXP + Transit

Worked Example

- ISP A is local access provider
  - Some business customers (around 200 fixed links)
  - Some co-located content provision (datacentre with 100 servers)
  - Some consumers on broadband (5000 DSL/Cable/Wireless)
- They have a single transit provider
  - Connect with a 16Mbps international leased link to their transit's PoP
  - Transit link is highly congested

Worked Example (2)

- There are two other ISPs serving the same locality
  - There is no interconnection between any of the three ISPs
  - Local traffic (between all 3 ISPs) is traversing International connections
- Course of action for our ISP:
  - Work to establish local IXP
  - Establish presence at overseas co-location
- **First Step**
  - Assess local versus international traffic ratio
  - Use NetFlow on border router connecting to transit provider

Worked Example (3)

- Local/Non-local traffic ratio
  - Local = traffic going to other two ISPs
  - Non-local = traffic going elsewhere
- Example: balance is 30:70
  - Of 16Mbps, that means 5Mbps could stay in country and not congest International circuit
  - 16Mbps transit costs $50 per Mbps per month
    - local traffic charges = $250 per month, or $3000 per year for local traffic
  - Circuit costs $100k per year => $30k is spent on local traffic
- Total is $33k per year for local traffic

Worked Example (4)

- IXP cost:
  - Simple 8 port 10/100 managed switch plus co-lo space over 3 years could be around US$30k total => $3k per year per ISP
  - One router to handle 5Mbps (e.g. 2801) would be around $3k (good for 3 years)
  - One local 10Mbps circuit from ISP location to IXP location would be around $5k per year, no traffic charges
  - Per ISP total: $11k
  - Somewhat cheaper than $33k
  - Business case for local peering is straightforward - $22k saving per annum

Worked Example (5)

- After IXP establishment
  - 5Mbps removed from International link
  - Leaving 5Mbps for more International traffic – and that fills the link within weeks of the local traffic being removed
- Next step is to assess transit charges and optimise costs
  - ISPs visits several major regional IXPs
  - Assess routes available
  - Compares routes available with traffic generated by those routes from its NetFlow data
  - Discovers that 30% of traffic would transfer to one IXP via peering

How to setup an IXP?

How to setup an IXP?

- The IXP core is an Ethernet switch
  - Managed switch with reasonable security features
- Has superseded all other types of network devices for an IXP
  - From the cheapest and smallest 12 or 24 port 100M/1G switch
  - To the largest switches now handling 10GE, 40GE, 100GE interfaces

How to setup an IXP?

- Each ISP participating in the IXP brings a router to the IXP location
  - Note: ISPs may connect directly to the IXP (availability of fiber connection) instead of a dedicated router at the IXP
- Router needs:
  - One Ethernet port to connect to IXP switch
  - One WAN port to connect to the WAN media leading back to the ISP backbone
  - To be able to run BGP

How to setup an IXP?

- IXP switch located in one equipment rack dedicated to IXP
  - Also includes other IXP operational equipment
- Routers from participant ISPs located in adjacent rack(s)
- Copper (UTP) connections made for 10/100Mbps or 1Gbps connections
- Fibre used for 10Gbps and 40Gbps

Peering

- Each participant needs to run BGP
  - They need their own AS number
  - **Public** ASN, **NOT** private ASN
- Each participant configures external BGP directly with the other participants in the IXP
  - Peering with all participants

    or
  - Peering with a subset of participants

Routing Advice

- ISP border routers at the IXP should NOT advertise default route or the full Internet routing table
  - Carrying default or full table means that this router and the ISP network is open to abuse by non-peering IXP members
  - Correct configuration is only to carry routes offered to IXP peers on the IXP peering router
- Note: Some ISPs offer transit across IX fabrics
  - They do so at their own risk – see above

Routing (more)

- ISP border routers at the IXP should not be configured to carry the IXP LAN network within the IGP or iBGP
  - Use next-hop-self BGP concept
- Don't generate ISP prefix aggregates on IXP peering router
  - If connection from backbone to IXP router goes down, normal BGP failover will then be successful

Address Space

- Some IXPs use private addresses for the IX LAN
  - Public address space means IXP network could be leaked to Internet which may be undesirable
  - Because most ISPs filter RFC1918 address space, this avoids the problem
- Some IXPs use public addresses for the IX LAN
  - Address space available from the RIRs
  - IXP terms of participation often forbid the IX LAN to be carried in the ISP member backbone

APNIC Policy on IXP Address Space

- The End-User Assignments policy caters for IXP's Public Address space under IXP Address Assignment
  - https://www.apnic.net/get-ip/faqs/ixp-address-assignment/

- It requires that IXP have minimum 3 ISPs connected and have clear and open policy for joining
- The minimum IXP Assignment is /24 of IPv4 and /48 for IPv6

Hardware

- Ethernet switch needs to be managed
  - Unmanaged switch means an unmanaged IXP

- Insist that IXP participants bring their own router
  - moves buffering problem off the IXP
  - Avoid spanning tree and other L2 security issues
  - Run port-security (MAC filtering) to protect the IX
  - security of the ISP connection is responsibility of the ISP, not the IXP

How to set up an IXP?

- The hard part with establishing an IXP is NOT the technical part, but for relevant stakeholders to come together to build a creditable governance structure for the IXP with which everyone is happy (TRUST)