The Value of Peering

ISP/IXP Workshops

Last updated 23rd March 2015

The Internet

- Internet is made up of ISPs of all shapes and sizes
  - Some have local coverage (access providers)
  - Others can provide regional or per country coverage
  - And others are global in scale
- These ISPs interconnect their businesses
  - They don’t interconnect with every other ISP (over 49500 distinct autonomous networks) – won’t scale
  - They interconnect according to practical and business needs
- Some ISPs provide transit to others
  - They interconnect other ISP networks
  - Around 6300 autonomous networks provide transit

Categorising ISPs

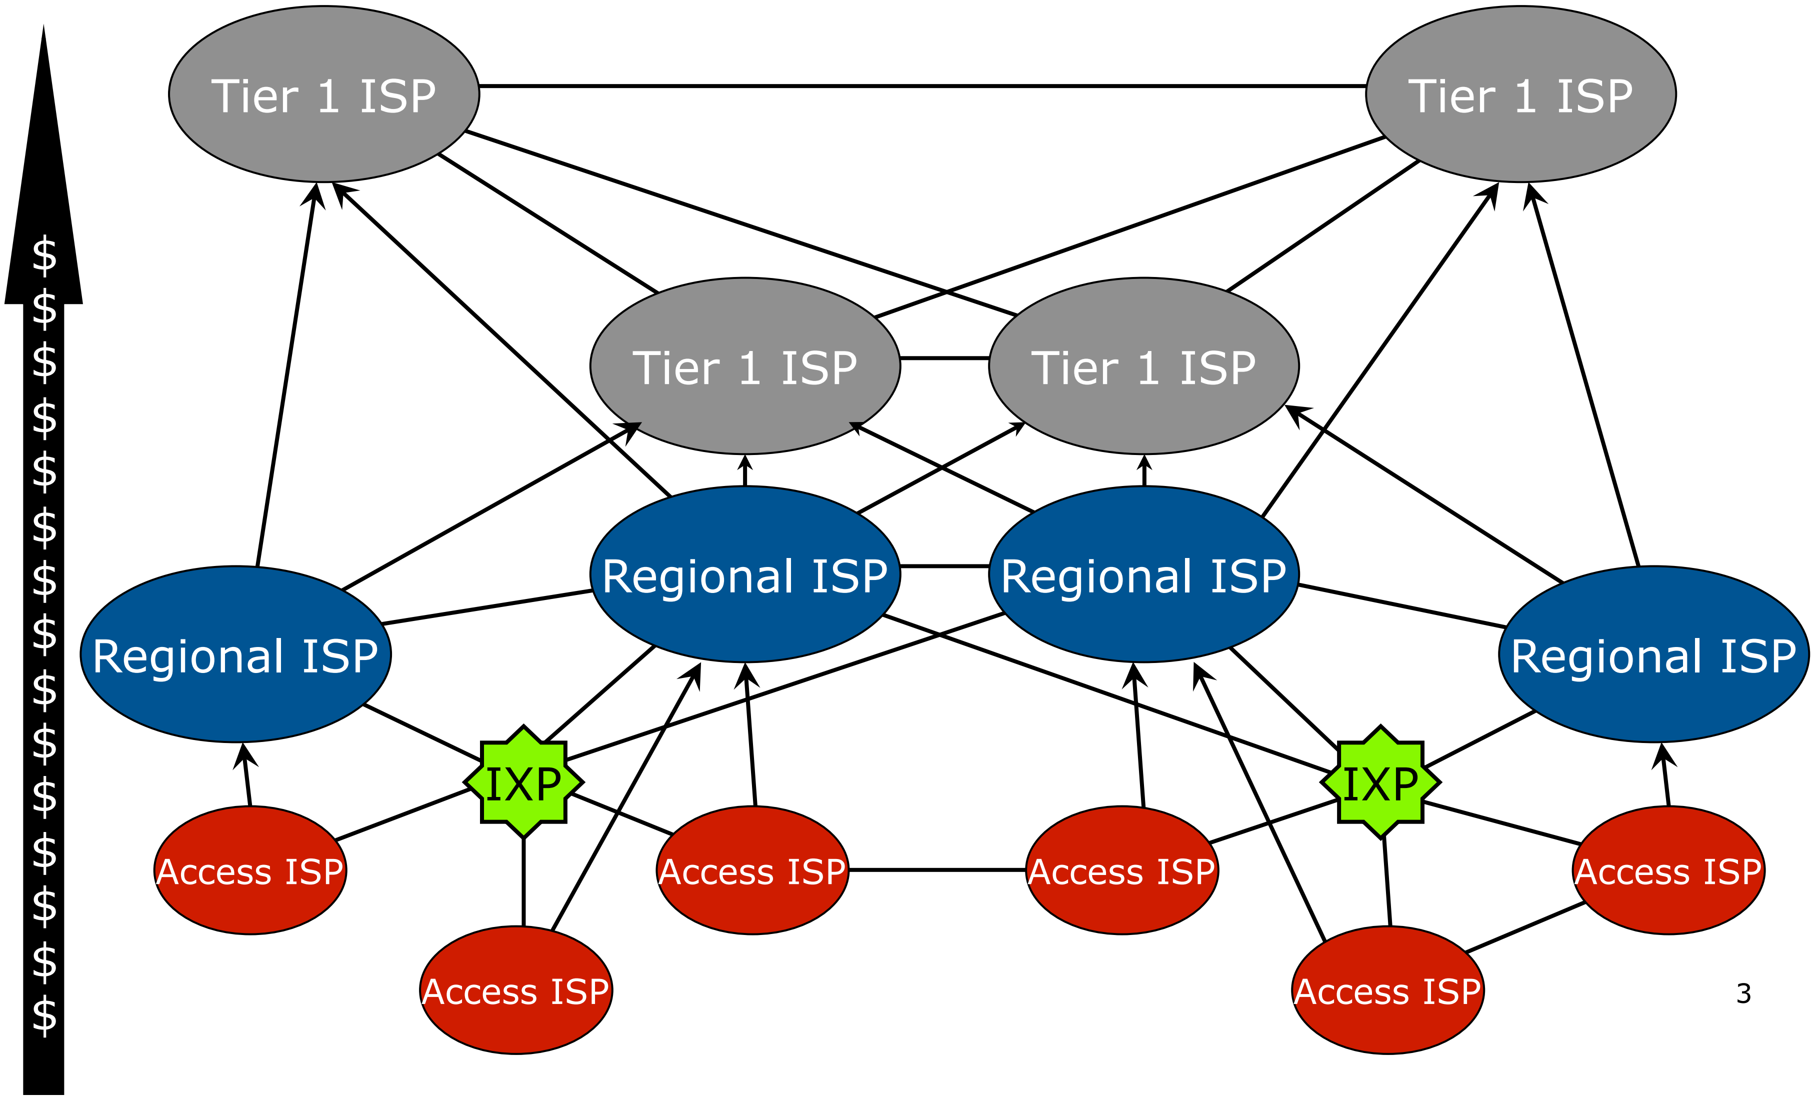

Peering and Transit

- Transit
  - Carrying traffic across a network
  - Usually for a fee
  - Example: Access provider connects to a regional provider
- Peering
  - Exchanging routing information and traffic
  - Usually for no fee
  - Sometimes called settlement free peering
  - Example: Regional provider connects to another regional provider

Private Interconnect

- Two ISPs connect their networks over a private link
  - Can be peering arrangement
    - No charge for traffic
    - Share cost of the link
  - Can be transit arrangement
    - One ISP charges the other for traffic
    - One ISP (the customer) pays for the link

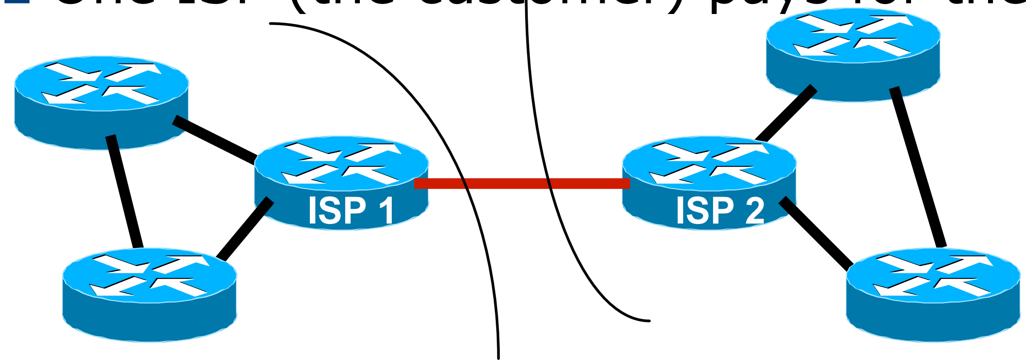

Public Interconnect

- Several ISPs meeting in a common neutral location and interconnect their networks
  - Usually is a peering arrangement between their networks

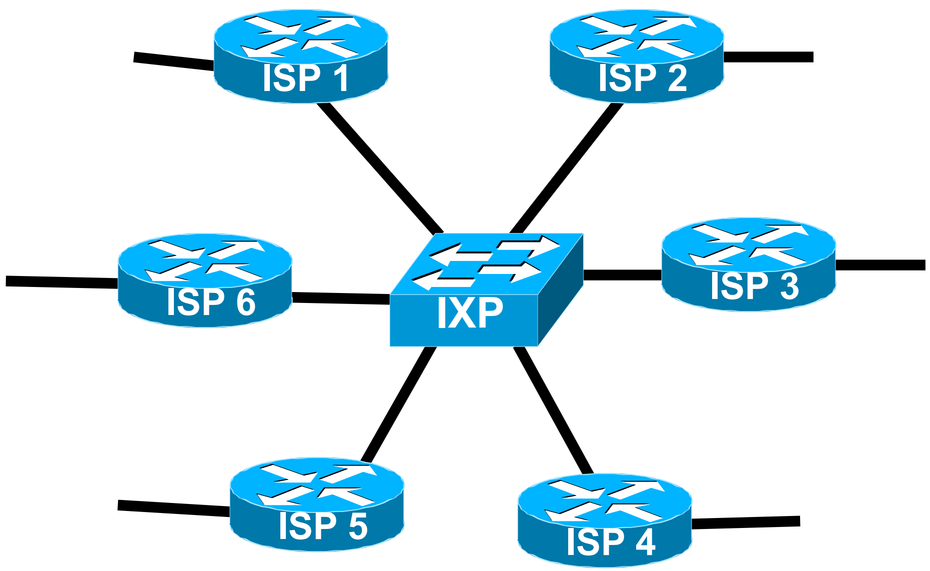

Types of Peering (1)

- Private Peering
  - Where two network operators agree to interconnect their networks, and exchange their respective routes, for the purpose of ensuring their customers can reach each 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 exchange traffic charges for a peering relationship

Types of Peering (2)

- Bi-lateral Peering
  - Very similar to Private Peering, but may take place at a public peering point (IXP)
- Multilateral Peering
  - Takes place at Internet Exchange Points, where operators all peer with each other via a Router Server
- Mandatory Multilateral Peering
  - Where operators are forced to peer with each other as condition of IXP membership
  - **Strongly discouraged: Has no record of success**

Types of Peering (3)

- Open Peering
  - Where an ISP publicly states that they will peer with all parties who approach them for peering
  - Commonly found at IXPs where ISP participates via the Route Server
- Selective Peering
  - Where an ISP's peering policy depends on the nature of the operator who requests peering with them
  - At IXPs, operator will not peer with RS but will only peer bilaterally
- Closed Peering
  - Where an ISP decides who its peering partners are, and is generally not approachable to creating peering opportunities

Types of Peering (4)

- The Peering Database documents ISPs peering policies
  - http://peeringdb.com
- All operators of ASNs should register in the peeringdb
  - All operators who are considering peering or are peering must be in the peeringdb to enhance their peering opportunities
- Participation in peering fora is encouraged too
  - Global Peering Forum (GPF)
  - Regional Peering Fora (European, Middle Eastern, Asian, Caribbean, Latin American)

peeringdb.com

Public Exchanges Point Detailed View

Navigation
Home Page
Logout

Your Records
Peering Record
User Account

Search Records
Networks
Exchange Points
Facilities
Common Points

Suggestions
Comments
New Exchange
New Facility

Help
FAQ
Statistics

Public Exchange Point Detailed View

Common Name	Equinix Palo Alto
Long Name	Equinix Internet Exchange Palo Alto
City	Palo Alto
Country	US
Continental Region	North America
Media Type	Ethernet
Protocols Supported	Unicast IPv4 ✓ Multicast ✓ IPv6 ✓

Contact Information

Website	https://ix.equinix.com
Traffic Statistics Website	
Technical E-Mail	servicesupport@equinix.com
Technical Phone	+1-866-811-8720
Policy E-Mail	servicesupport@equinix.com
Policy Phone	

IP Address Blocks

Type	Address Block	Reverse DNS Scan
IPv4 Unicast	198.32.175.0/24	Link
IPv4 Unicast	198.32.176.0/24	Link
IPv6 Unicast	2001:504:d::/64	Unsupported
IPv4 Multicast	198.32.177.0/24	Link

Local Facilities

Facility Name	City	Country	Participant Count
365 San Jose (formerly Equinix San Jose (SV7))	San Jose	US	4
Digital Realty San Francisco (200 Paul)	San Francisco	US	16
Equinix Palo Alto (SV8)	Palo Alto	US	111
Equinix Sunnyvale (SV6)	Sunnyvale	US	7

List of Peers at this Exchange Point (Total: 115)

Peer Name	Local ASN	IP Address	IPs	Policy
6connect, Inc.	8038	198.32.176.51	2	Open
AARNet	7575	198.32.176.177	2	Selective
Academia Sinica Network(ASNet)	9264	198.32.176.174	2	Open
Akamai Technologies	20940	198.32.176.127	2	Open
Amazon.com	16509	198.32.176.217	2	Open
Apple Inc	714	198.32.176.237	2	Selective
Bell Canada Backbone	577	198.32.176.94	2	Restrictive
Bharti Airtel Limited	9498	198.32.176.203	2	Selective
Biznet Networks	17451	198.32.176.60	2	Open
Black Oak Computers Inc	22781	198.32.176.153	2	Open
BlinkMind, Inc.	40739	198.32.176.121	1	Open
BroadbandONE (formerly WV Fiber)	19151	198.32.176.164	2	Selective
CDNetworks Inc.	36408	198.32.176.221	2	Open
CENIC / CalREN	2152	198.32.176.33	2	Selective
Chunghwa Telecom	9505	198.32.176.160	2	Open
CNS-KBT	9416	198.32.176.212	1	Open
CRITEO (USA)	19750	198.32.176.110	2	Selective
Dailymotion	41690	198.32.176.151	1	Open
DBolical Pty Ltd	55651		1	Open
Docomo Pacific	3605	198.32.176.100	2	Selective
Dropbox	19679	198.32.176.200	4	Open
Dynamic Network Services, Inc.	33517	198.32.176.56	2	Selective
Electronic Arts	22220	198.32.176.23	1	Open

1 2 3 4 5 of 5 Next > Last >>

peeringdb.com

Peering Networks Detailed View

Navigation
Home Page
Logout

Your Records
Peering Record
User Account

Search Records
Networks
Exchange Points
Facilities
Common Points

Suggestions
Comments
New Exchange
New Facility

Help
FAQ
Statistics

Company Information

Company Name	Amazon.com
Also Known As	
Company Website	http://www.amazon.com
Primary ASN	16509
IRR Record	AS-AMAZON
Network Type	Enterprise
Approx Prefixes	1000
Traffic Levels	Not Disclosed
Traffic Ratios	Balanced
Geographic Scope	Global
Looking Glass URL	
Route Server URL	

The following Amazon US locations and associated IX's carry routes/traffic specific only to the services with infrastructure in that metro. For example, Jacksonville is CloudFront only, whereas Ashburn is CloudFront, EC2, S3, etc.)

- Seattle
- Palo Alto
- San Jose
- Los Angeles
- Dallas
- St Louis
- South Bend
- Jacksonville
- Miami
- Ashburn
- Vienna
- Newark
- New York

The following locations and associated IX's are part of Amazon's European Backbone, carrying routes/traffic for other AWS services (e.g. EC2, S3), Amazon retail, as well as local CloudFront caching. For each of these locations, Amazon will provide local routes/traffic for all services within that locality unless peers are able to meet in at least two diverse locations within the region, in which case, routes/traffic for all services within

Public Peering Exchange Points

Exchange Point Name	ASN	IP Address	Mbit/sec
AMS-IX	16509	80.249.210.217	60000
AMS-IX	16509	2001:7f8:1::a501:6509:2	60000
AMS-IX	16509	80.249.210.100	60000
AMS-IX	16509	2001:7f8:1::a501:6509:1	60000
AMS-IX Hong Kong	16509	103.247.139.10	10000
BBIX Tokyo	16509	218.100.6.52	10000
CoreSite - Any2 California	16509	2001:504:13::146	30000
CoreSite - Any2 California	16509	206.72.210.146	30000
DE-CIX Frankfurt	16509	80.81.195.152	80000
DE-CIX Frankfurt	16509	2001:7f8::407d:0:1	80000
DE-CIX Frankfurt	16509	80.81.194.152	80000
DE-CIX Frankfurt	16509	2001:7f8::407d:0:2	80000

1 2 3 4 5 of 9 Next > Last >>

Private Peering Facilities

Facility Name	ASN	City	Country	SONET	Ethr	ATM
365 St. Louis (formerly Equinix St Louis)	16509	St. Louis	US	☐	☑	☐
421 West Church Street	16509	Jacksonville	US	☐	☑	☐
CoreSite - LA1 - One Wilshire	16509	Los Angeles	US	☐	☑	☐
CoreSite - NY1	16509	New York	US	☐	☑	☐
Equinix Ashburn (DC1-DC11)	16509	Ashburn	US	☐	☑	☐
Equinix Dallas (DA3)	16509	Dallas	US	☐	☑	☐
Equinix Frankfurt KleyerStrasse (FR5)	16509	Frankfurt	DE	☐	☑	☐
Equinix Hong Kong (HK1)	16509	Hong Kong	HK	☐	☑	☐
Equinix Los Angeles (LA1)	16509	Los Angeles	US	☐	☑	☐
Equinix New York (111 8th)	16509	New York	US	☐	☑	☐
Equinix Newark (NY1)	16509	Newark	US	☐	☑	☐
Equinix Palo Alto (SV8)	16509	Palo Alto	US	☐	☑	☐

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 dial up, 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

The IXP's role

- Private peering makes sense when there are very few equivalent players
  - Connecting to one other ISP costs X
  - Connecting to two other ISPs costs 2 times X
  - Connecting to three other ISPs costs 3 times X
  - Etc… (where X is half the circuit cost plus a port cost)
- The more private peers, the greater the cost
- IXP is a more scalable solution to this problem

The IXP's role

- Connecting to an IXP
  - ISP costs: one router port, one circuit, and one router to locate at the IXP
- Some IXPs charge annual "maintenance fees"
  - The maintenance fee has potential to significantly influence the cost balance for an ISP
- Generally connecting to an IXP and peering there becomes cost effective when there are at least three other peers
  - The real $ amount varies from region to region, IXP to IXP

Who peers at an IXP?

- Access Providers
  - Don't have to pay their regional provider transit fees for local traffic
  - Keeps latency and costs for local traffic low
  - 'Unlimited' bandwidth through the IXP (compared with costly and limited bandwidth through transit provider)
- Regional Providers
  - Don't have to pay their global provider transit fees for local and regional traffic
  - Keeps latency and costs for local and regional traffic low
  - 'Unlimited' bandwidth through the IXP (compared with costly and limited bandwidth through global provider)

Who peers at an IXP?

- Content Providers & Content Distribution Services
  - Don't have to pay their regional provider transit fees for local traffic
  - Keeps latency and costs for local traffic low
  - 'Unlimited' bandwidth through the IXP (compared with costly and limited bandwidth through transit provider)
- Root, ccTLD and gTLD operators
  - Adds to the resiliency of the global DNS system
  - Keeps latency and response time for local resolver traffic very low

The IXP's role

- Global Providers can be located close to IXPs
  - Attracted by the potential transit business available
- Advantageous for access & regional providers
  - They can peer with other similar providers at the IXP
  - And in the same facility pay for transit to their regional or global provider
  - (Not across the IXP fabric, but a separate connection)

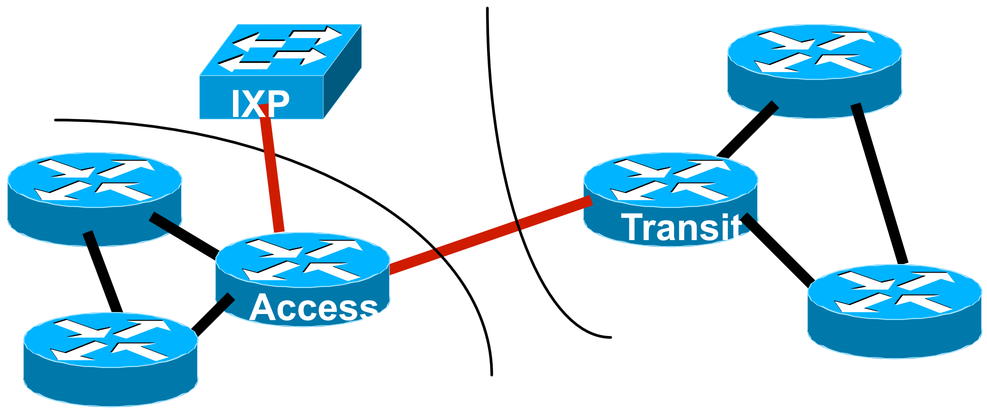

Connectivity Decisions

- Transit
  - Almost every ISP needs transit to reach rest of Internet
  - One provider = no redundancy
  - Two providers: ideal for traffic engineering as well as redundancy
  - Three providers = better redundancy, traffic engineering gets harder
  - More then three = diminishing returns, rapidly escalating costs and complexity
- Peering
  - Means low (or zero) cost access to another network
  - Private or Public Peering (or both)

Transit Goals

1. Minimise number of transit providers
   - But maintain redundancy
   - 2 is ideal, 4 or more is hard
2. Aggregate capacity to transit providers
   - More aggregated capacity means better value
     - Lower cost per Mbps
   - 4x 45Mbps circuits to 4 different ISPs will almost always cost more than 2x 155Mbps circuits to 2 different ISPs
     - Yet bandwidth of latter (310Mbps) is greater than that of former (180Mbps) and is much easier to operate

Peering or Transit?

- How to choose?
- Or do both?
- It comes down to cost of going to an IXP
  - Free peering
  - Paying for transit from an ISP co-located in same facility, or perhaps close by
- Or not going to an IXP and paying for the cost of transit directly to an upstream provider
  - There is no right or wrong answer, someone has to do the arithmetic

Private or Public Peering

- Private peering
  - Scaling issue, with costs, number of providers, and infrastructure provisioning
- Public peering
  - Makes sense the more potential peers there are (more is usually greater than "two")
- Which public peering point?
  - Local Internet Exchange Point: great for local traffic and local peers
  - Regional Internet Exchange Point: great for meeting peers outside the locality, might be cheaper than paying transit to reach the same consumer base

Local Internet Exchange Point

- Defined as a public peering point serving the local Internet industry
- Local means where it becomes cheaper to interconnect with other ISPs at a common location than it is to pay transit to another ISP to reach the same consumer base
  - Local can mean different things in different regions!

Regional Internet Exchange Point

- These are also "local" Internet Exchange Points
- But also attract regional ISPs and ISPs from outside the locality
  - Regional ISPs peer with each other
  - And show up at several of these Regional IXPs
- Local ISPs peer with ISPs from outside the locality
  - They don't compete in each other's markets
  - Local ISPs don't have to pay transit costs
  - ISPs from outside the locality don't have to pay transit costs
  - Quite often ISPs of disparate sizes and influences will happily peer – to defray transit costs

Which IXP?

- How many routes are available?
  - What is traffic to & from these destinations, and by how much will it reduce cost of transit?
- What is the cost of co-lo space?
  - If prohibitive or space not available, pointless choosing this IXP
- What is the cost of running a circuit to the location?
  - If prohibitive or competitive with transit costs, pointless choosing this IXP
- What is the cost of remote hands/assistance?
  - If no remote hands, doing maintenance is challenging and potentially costly with a serious outage

Example: South Asian ISP @ LINX

- Time: 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

- Time: 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
- Servers at remote co-lo
  - Benefits: mail filtering, content caching, etc
  - 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)
  - Some consumers on dial (1000 on V.34 type speeds)
- 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 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; or $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: $9k
  - Somewhat cheaper than $33k
  - Business case for local peering is straightforward - $24k 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

Worked Example (6)

- Costs:
  - Router for Regional IXP (e.g. 2801) at $3k over three years
  - Co-lo space at Regional IXP venue at $3k per year
  - Best price for transit at the Regional IXP venue by competitive tender is $30 per Mbps per month, plus $1k port charge
  - 30% of traffic offloads to IXP, leaving 70% of 16Mbps to transit provider = $330 per month, or $5k per annum
  - Total with this model is $9k per year, plus the cost of the circuit (still $100k)
  - Compare this with paying $50 per Mbps per month to the transit provider = $10k per annum (plus cost of the circuit)

Worked Example (7)

- Result:
  - ISP co-locates at Regional IXP
  - Pays reduced transit charges to transit provider (competitive tender)
  - Pays no charges for traffic across Regional IXP
- Bonuses:
  - Rate limits on router at Regional IXP Co-lo
    - Can prioritise congestion dependent on customer demands
  - Install servers at Regional IXP co-lo facility
    - Filters e-mail (spam and viruses) – relieves some capacity on link
    - Caches content – relieves a little more capacity on link

Conclusion

- Within the original costs of having one international transit provider:
  - ISP has turned up at the local IXP and offloaded local traffic for free
  - ISP has turned up at a major regional IXP and offloaded traffic, avoiding paying transit charges to transit provider
  - ISP has reduced remaining transit charges by competitive tender at the regional IXP co-location facility
- Caveat
  - These numbers are typical of the Internet today
  - As ever, your mileage may vary – but do the financial calculations first and in the context of potential technical advantages too

The Value of Peering

ISP/IXP Workshops