

The Value of Peering

ISP/IXP Workshops



These materials are licensed under the Creative Commons Attribution-NonCommercial 4.0 International license (<http://creativecommons.org/licenses/by-nc/4.0/>)

Acknowledgements

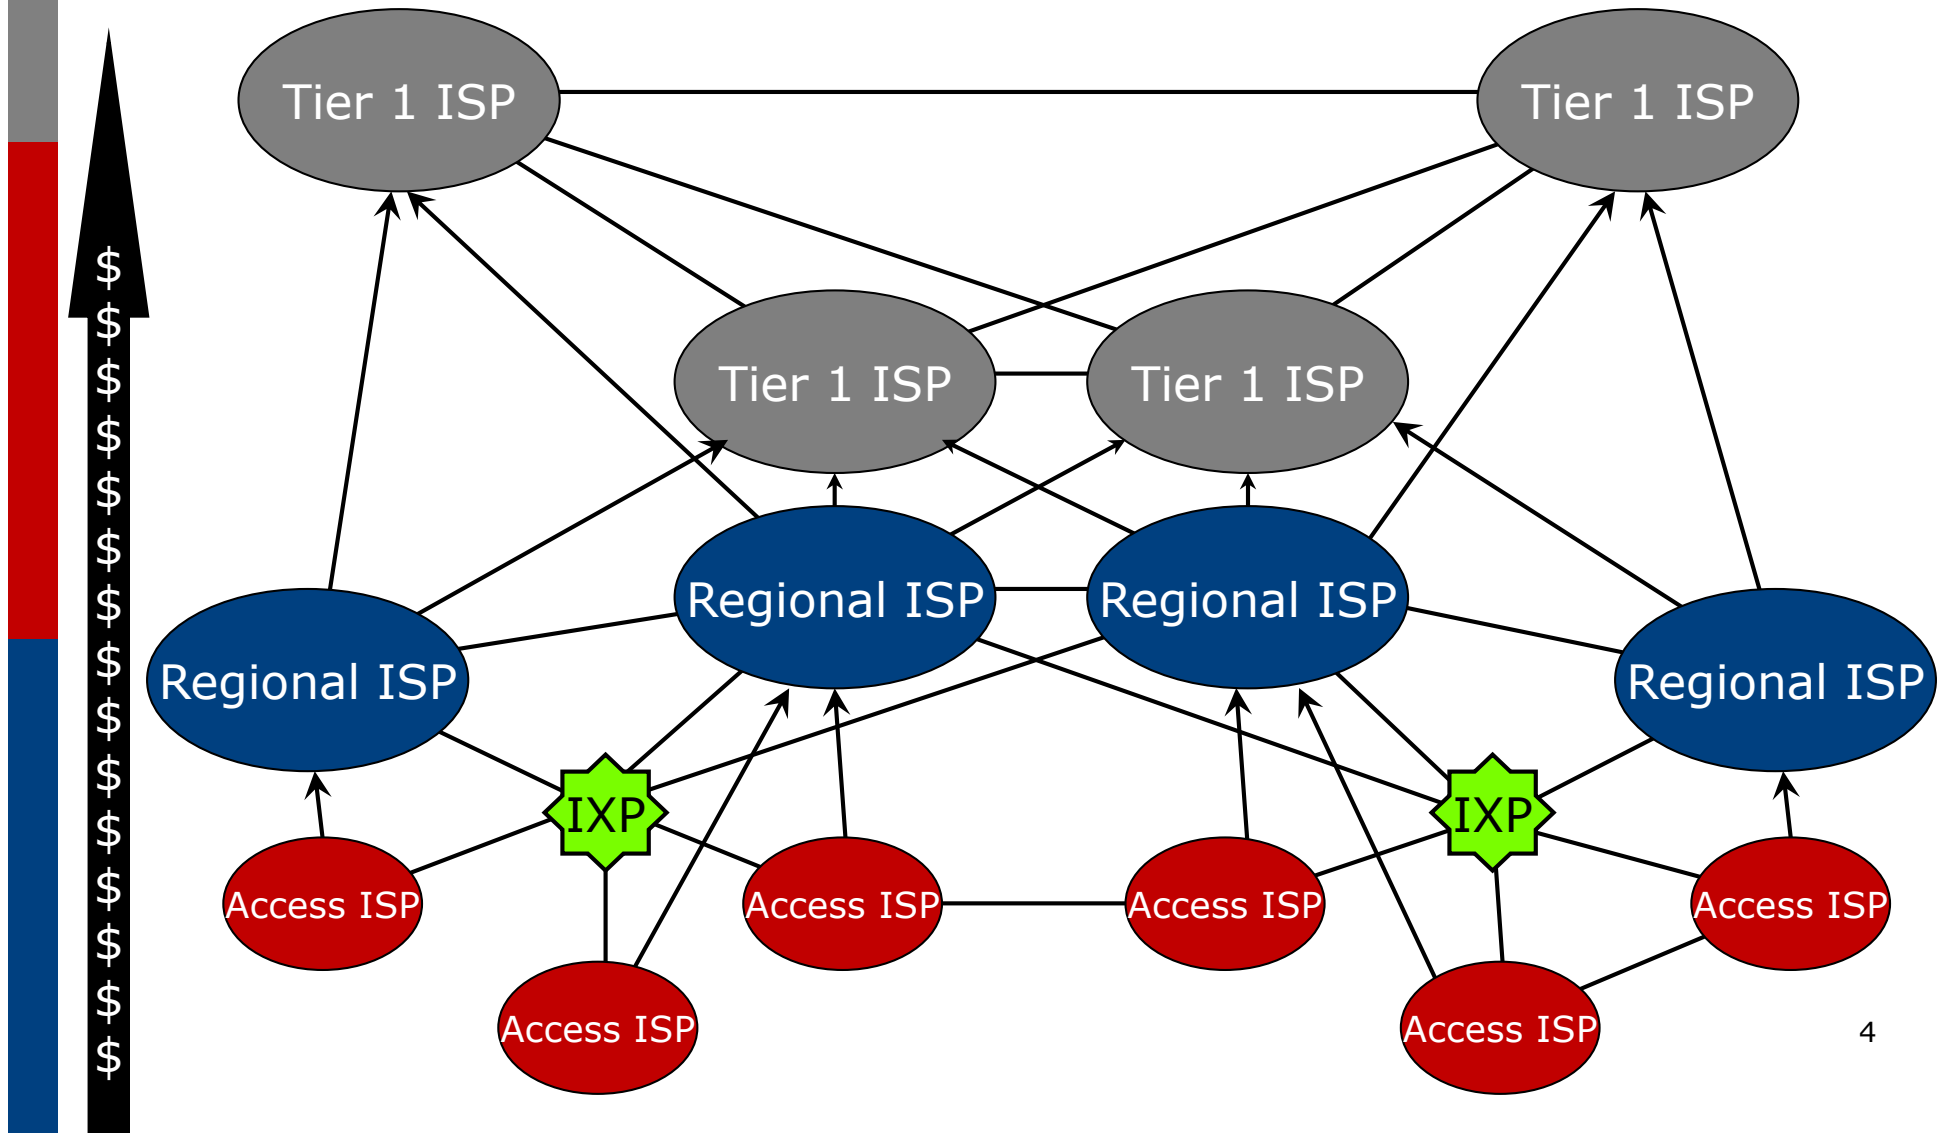
- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene
- Use of these materials is encouraged as long as the source is fully acknowledged and this notice remains in place
- Bug fixes and improvements are welcomed
 - Please email *workshop (at) bgp4all.com*

Philip Smith

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 58100 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
 - Just over 7800 autonomous networks provide transit

Categorising ISPs



Peering and Transit

□ Transit

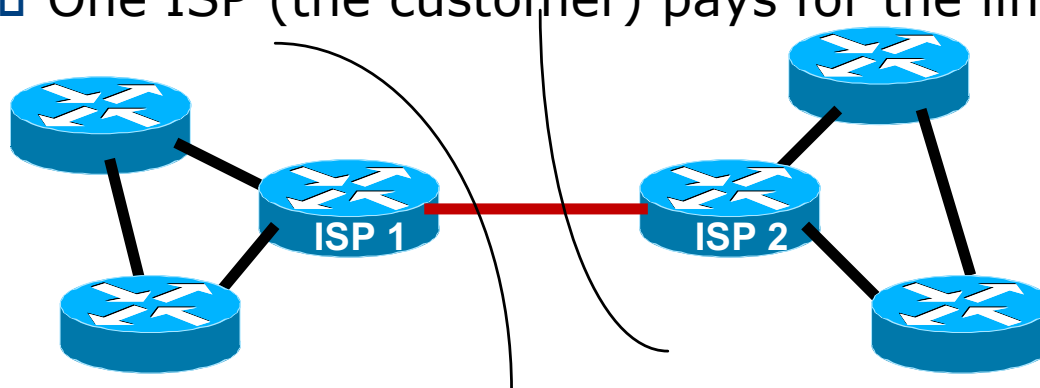
- Carrying traffic across a network
- Usually for a fee
- Example: Access provider connects to a regional provider

□ 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**
 - Private Network Interconnect (PNI)
 - Can be peering arrangement – “Private Peering”
 - ❑ 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 usually takes place at a public peering point (IXP)
- Multilateral Peering
 - Takes place at Internet Exchange Points, where operators all peer with each other via a Route 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

□ Restrictive Peering

- Where an ISP decides who its peering partners are, and is generally not approachable to considering peering opportunities

Types of Peering (4)

- The Peering Database documents ISPs peering policies
 - <https://www.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)
 - Many countries now have their own Peering Fora

Equinix Palo Alto

Organization	Equinix
Long Name	Equinix Internet Exchange Palo Alto
City	Palo Alto
Country	US
Continental Region	North America
Media Type	Ethernet
Protocols Supported	<input checked="" type="checkbox"/> Unicast IPv4 <input checked="" type="checkbox"/> Multicast <input checked="" type="checkbox"/> IPv6

Contact Information

Company Website	https://ix.equinix.com
Traffic Stats Website	
Technical Email	servicesupport@equinix.com
Technical Phone	+1-866-811-8720
Policy Email	servicesupport@equinix.com
Policy Phone	

LAN

MTU	
DOT1Q	<input type="radio"/>
IPv4	198.32.176.0/24
IPv4	198.32.175.0/24
IPv4	198.32.177.0/24
IPv6	2001:504:d::/64

Local Facilities

Facility ▼	Country	City
Digital Realty San Francisco (200 Paul)	United States of America	San Francisco
Equinix Palo Alto (SV8)	United States of	Palo Alto

Peers at this Exchange Point

Peer Name ▼ ASN	IPv4 IPv6	Speed Policy
6connect, Inc. 8038	198.32.176.51 2001:504:d::33	1G Open
AARNet 7575	198.32.176.177 2001:504:d::b1	10G Selective
Academia Sinica Network(ASNet) 9264	198.32.176.174 2001:504:d::ae	2G Open
Advanced Wireless Network Co. Ltd. 45430	198.32.176.129 2001:504:d::4:5430:1	1G Selective
Akamai Prolexic DDoS Mitigation 32787	198.32.176.228 2001:504:d::3:2787:1	10G Selective
Akamai Technologies 20940	198.32.176.127 2001:504:d::2:940:1	60G Open
alibaba 45102	198.32.176.180 None	10G Open
Amazon.com 16509	198.32.176.36 2001:504:d::24	60G Open
Amazon.com 16509	198.32.176.217 2001:504:d::d9	60G Open
Apple Inc 714	198.32.176.237 2001:504:d::714:1	40G Selective
Bell Canada Backbone 577	198.32.176.94 2001:504:d::5e	10G Restrictive
Bharti Airtel Limited 9498	198.32.176.203 2001:504:d::9498:1	20G Selective
Biznet Networks 17451	198.32.176.60 2001:504:d::3c	1G Open
BlinkMind, Inc.	198.32.176.121	1G

Amazon.com

Organization	Amazon.com
Also Known As	
Company Website	http://www.amazon.com
Primary ASN	16509
IRR Record	AS-AMAZON
Route Server URL	
Looking Glass URL	
Network Type	Enterprise
IPv4 Prefixes	2000
IPv6 Prefixes	500
Traffic Levels	Not Disclosed
Traffic Ratios	Balanced
Geographic Scope	Global
Protocols Supported	<input checked="" type="checkbox"/> Unicast IPv4 <input type="checkbox"/> Multicast <input checked="" type="checkbox"/> IPv6
Last Updated	2016-05-23T23:08:16Z

Notes

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

Public Peering Exchange Points

Exchange ▼ ASN	IPv4 IPv6	Speed RS Peer
AMS-IX 16509	80.249.210.100 2001:7f8:1::a501:6509:1	200G <input type="radio"/>
AMS-IX 16509	80.249.210.217 2001:7f8:1::a501:6509:2	200G <input type="radio"/>
AMS-IX Hong Kong 16509	103.247.139.10 2001:df0:296::a501:6509:1	10G <input type="radio"/>
BBIX Osaka 16509	218.100.7.24 2001:de8:c:2:0:1:6509:1	100G <input type="radio"/>
BBIX Tokyo 16509	218.100.6.52 2001:de8:c:1:6509:1	200G <input type="radio"/>
BCIX 16509	193.178.185.95 2001:7f8:19:1::407d:1	100G <input type="radio"/>
CoreSite - Any2 California 16509	206.72.210.146 2001:504:13::146	30G <input type="radio"/>
CoreSite - Any2 California 16509	206.72.211.146 2001:504:13::211:146	30G <input type="radio"/>
DE-CIX Frankfurt <small>Main</small> 16509	80.81.194.152 2001:7f8::407d:0:1	200G <input type="radio"/>
DE-CIX Frankfurt <small>Main</small> 16509	80.81.195.152 2001:7f8::407d:0:2	200G <input type="radio"/>
DE-CIX New York 16509	206.130.10.99 2001:504:36::407d:0:1	40G <input type="radio"/>
Digital Realty Telx Atlanta 16509	198.32.132.95 2001:478:132::95	60G <input type="radio"/>
Digital Realty Telx New York 16509	206.126.115.37 2001:504:17:115::37	10G <input type="radio"/>
ECIX-BER 16509	194.9.117.85 2001:7f8:8:5:0:407d:0:1	100G <input type="radio"/>

Private Peering Facilities

Facility ▼ ASN	Country City
151 Front Street West Toronto 16509	Canada Toronto
365 Data Centers St. Louis (ST1) 16509	United States of America St. Louis

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 than 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 STM-1/OC3 links to 4 different ISPs will almost always cost more than 2x STM-4/OC12 links to 2 different ISPs
 - Yet bandwidth of latter (1.2Gbps) is greater than that of former (620Mbps) 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

The Value of Peering



ISP/IXP Workshops