

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/>)

Last updated 8th December 2021

Acknowledgements

- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene
 - Appreciative thanks for the contributions from Kurt Erik Lindqvist, Mark Tinka, Nigel Titley, and other members of the peering community

- 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

BGP Videos

- NSRC has made a video recording of this presentation, as part of a library of BGP videos for the whole community to use:
 - https://learn.nsrc.org/bgp#value_of_peering

The screenshot displays the NSRC (Network Startup Resource Center) website. The navigation bar includes links for Home, About, BGP for All (highlighted), perfSONAR, ScienceDMZ, FedIdM, and Contact Us, along with a search bar. The main content area is divided into three columns:

- BGP for All:** A text-based introduction to BGP, explaining its role as the primary routing protocol for the Internet and autonomous systems. It also mentions that understanding routing options can lead to efficiencies and collaboration opportunities.
- Introduction to Routing:** A list of video topics including Internet Routing, Routing Protocols, Introduction to IS-IS (UPDATED), IS-IS Levels, IS-IS Adjacencies, Best Configuration Practices for IS-IS on Cisco IOS, IS-IS Authentication, Default Routes and IPv6, Introduction to OSPF, OSPF Areas, OSPF Adjacencies, Best Configuration Practices for OSPF on Cisco IOS, OSPF Authentication, Default Routes and IPv6, Comparing OSPF and IS-IS, Choosing between OSPF and IS-IS, Migrating from OSPF to IS-IS, Migration Plan, and Finalizing Migration.
- Introduction to BGP:** A list of video topics including Introduction to Border Gateway Protocol, Transit and Peering, Autonomous Systems (UPDATED), How BGP works, Supporting Multiple Protocols, IBGP versus EBGP, Setting up EBGP, and Setting up IBGP.

On the right side, there is a video player for "BGP for All" with a play button and a "Watch on YouTube" button. Below the video player, there are sections for "BGP Case Studies" and "Communities".

BGP Case Studies:

- Peering Priorities ^{NEW}
- Transit Provider Peering at an IXP ^{NEW}
- Customer Multihomed between two IXP members ^{NEW}
- Traffic Engineering for an ISP connected to two IXes ^{NEW}
- Traffic Engineering for an ISP with two interfaces on one IX LAN ^{NEW}
- Traffic Engineering and CDNs ^{NEW}

Communities:

- Communities: RFC 1998 Traffic Engineering
- Communities: Simplifying Traffic Engineering
- How to Apply Communities to Originated Routes
- How to Use Communities for Service Identification



Network Operator Goals?

- Today, the vast majority of content and resources consumed by end-users is available by peering:
 - The multi-national content providers (Google, Facebook, etc)
 - The multi-national “cloud” providers
 - Private cross connects
 - Internet Exchange Points
- A network operator’s goal is to obtain as much peering as possible
- Transit is for the last resort, for any content not available by peering



Network Operator Goals?

□ Peering

- Locally by direct cross-connect with other providers
- Locally at an Internet Exchange Point
- Getting to the nearest IXP or other interconnect

□ Transit

- Relying on another network operator to get the rest of the Internet
- Considered a last resort now



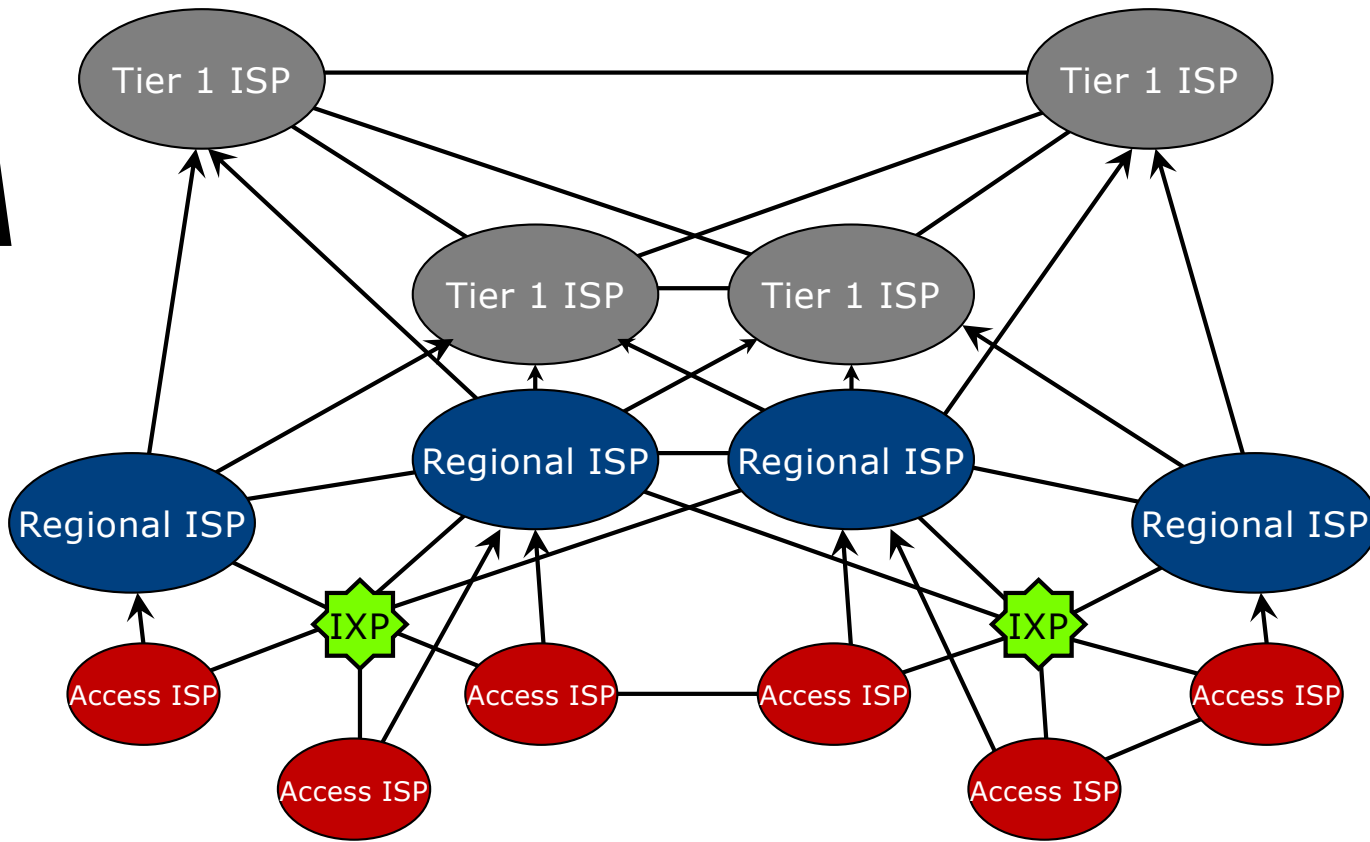
Other concerns for Network Operators

- ❑ More economies and regions worried about “data sovereignty”
- ❑ Data held at multi-national cloud operators could be anywhere on Earth
- ❑ Local peering and local interconnects mean:
 - Less “domestic” traffic going crossing national boundaries
 - Greater opportunity for domestic cloud/hosting providers
 - More responsive hosting and “cloud” services
 - More assurance about “data sovereignty”
 - Greater opportunity for creating a vibrant local Internet economy

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 72400 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 10200 autonomous networks provide transit to another AS

Categorising ISPs



Peering and Transit

□ Transit


- Carrying traffic across a network
- Usually have to pay traffic charges
- Example: Access provider connects to a regional provider

□ Peering

- Exchanging routing information and traffic
- No traffic charges
- Normally 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

- Private Peering
- Public Peering (at Internet Exchange Points)
 - Bi-Lateral
 - Multi-Lateral
 - Mandatory Multi-Lateral

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
- ❑ Once operators interconnect:
 - 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
 - ❑ (Marketing name for Local Transit? 🤔)

Public Peering

- *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**
 - (But some are still determined to prove 30 years of industry experience wrong 😬)

Peering Business Policies

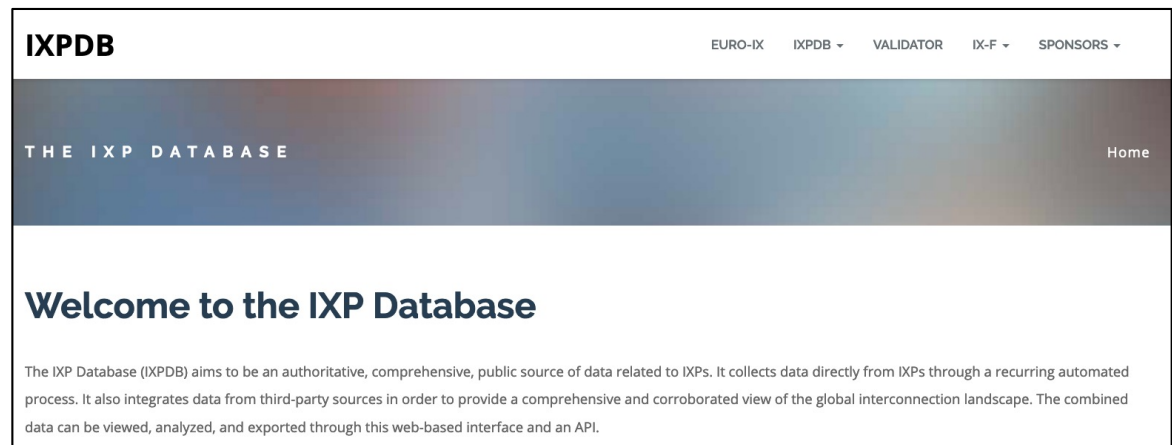
- *Open Peering*
 - Where a network operator publicly states that they will peer with all parties who approach them for peering
 - Commonly found at IXPs where the member participates via the Route Server
- *Selective Peering*
 - Where a network operator's peering policy depends on the nature of the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

The Peering Database

- The Peering Database documents network operator peering policies
 - <https://www.peeringdb.com>
- All operators with AS numbers 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) – for North American operators
 - Regional Peering Fora (European, Middle Eastern, Asian, Caribbean, Latin American)
 - Many countries now have their own Peering Fora


The IXP Database

- ❑ The IXPDB documents IXPs and their participants around the world
 - <https://ixpdb.euro-ix.net/en/>
- ❑ All Internet Exchange Point operators are recommended to register their IXP in the database
 - IXPs using IXP Manager (<https://www.ixpmanager.org>) will have this happen as part of the IXP Manager set up
 - Provides the LAN IP addresses of each member to facilitate automation



The screenshot shows the homepage of the IXP Database (IXPDB). The header includes the IXPDB logo and navigation links for EURO-IX, IXPDB, VALIDATOR, IX-F, and SPONSORS. The main content area features the text 'THE IXP DATABASE' and a 'Home' link. Below this is a large heading 'Welcome to the IXP Database' followed by a paragraph describing the database's mission: 'The IXP Database (IXPDB) aims to be an authoritative, comprehensive, public source of data related to IXPs. It collects data directly from IXPs through a recurring automated process. It also integrates data from third-party sources in order to provide a comprehensive and corroborated view of the global interconnection landscape. The combined data can be viewed, analyzed, and exported through this web-based interface and an API.'

HKIX

Organization	Hong Kong Internet eXchange Limited
Long Name	Hong Kong Internet Exchange
City	Hong Kong
Country	HK
Continental Region	Asia Pacific
Media Type	Ethernet
Protocols Supported	<input checked="" type="radio"/> Unicast IPv4 <input type="radio"/> Multicast <input checked="" type="radio"/> IPv6
Notes 	


Contact Information

Company Website	https://www.hkix.net/
Traffic Stats Website	https://www.hkix.net/hkix/stat/aggt/hkix-aggregate.html
Technical Email	noc@hkix.net
Technical Phone	+85239439900
Policy Email	info@hkix.net
Policy Phone	+85239438800


LAN

MTU	1500
DOT1Q	<input type="radio"/>
IPv6	2001:7fa:0:1::/64
IPv4	123.255.88.0/21

Local Facilities

Facility 	Country	City
CUHK	Hong Kong	Hong Kong
MEGA Two (iAdvantage Hong Kong)	Hong Kong	Hong Kong
MEGA-i (iAdvantage Hong Kong)	Hong Kong	Hong Kong

Peers at this Exchange Point

Peer Name  ASN	IPv4 IPv6	Speed Policy
ASGCNET HKIX Peering LAN 24167	123.255.91.53 2001:7fa:0:1::ca28:a135	10G Open
Asia Pacific Telecom HKIX Peering LAN 17709	123.255.91.86 2001:7fa:0:1::ca28:a156	10G Open
ASLINE HKIX Peering LAN 18013	123.255.92.13 2001:7fa:0:1::ca28:a20d	10G Open
AT&T AP - AS2687 HKIX Peering LAN 2687	123.255.91.46 2001:7fa:0:1::ca28:a12e	10G Selective
Automattic HKIX Peering LAN 2635	123.255.90.71 2001:7fa:0:1::ca28:a047	10G Open
Badoo Ltd HKIX Peering LAN 12678	123.255.90.220 None	2G Open
Baidu HKIX Peering LAN 55967	123.255.90.131 2001:7fa:0:1::ca28:a083	10G Open
Baidu HKIX Peering LAN 55967	123.255.91.61 2001:7fa:0:1::ca28:a13d	10G Open
Bayan Telecommunications Inc. HKIX Peering LAN 6648	123.255.91.45 2001:7fa:0:1::ca28:a12d	3G Open
BGP Network Limited HKIX Peering LAN 64050	123.255.91.177 2001:7fa:0:1::ca28:a1b1	100G Open
BIGHUB-ISP HKIX Peering LAN 137989	123.255.90.207 2001:7fa:0:1::ca28:a0cf	1G Open
BIGHUB-ISP HKIX Peerina LAN	123.255.91.98	10G

Amazon.com Diamond Sponsor

Organization	Amazon.com
Also Known As	Amazon Web Services
Company Website	http://www.amazon.com
Primary ASN	16509
IRR as-set/route-set ?	AS-AMAZON
Route Server URL	
Looking Glass URL	
Network Type	Enterprise
IPv4 Prefixes ?	5000
IPv6 Prefixes ?	2000
Traffic Levels	Not Disclosed
Traffic Ratios	Balanced
Geographic Scope	Global
Protocols Supported	<input checked="" type="radio"/> Unicast IPv4 <input type="radio"/> Multicast <input checked="" type="radio"/> IPv6 <input type="radio"/> Never via route servers
Last Updated	2019-12-29T14:56:38Z
Notes ?	<p>If you have a connectivity issue to Amazon then please visit:</p> <ul style="list-style-type: none"> • IPv4: http://ec2-reachability.amazonaws.com/ • IPv6: http://ipv6.ec2-reachability.amazonaws.com/ <p>And include detail on prefixes you think you have a problem with if you contact our Ops alias. This will reduce time with troubleshooting.</p> <p>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.)</p> <ul style="list-style-type: none"> • Seattle • Palo Alto • San Jose • Los Angeles • Dallas

Public Peering Exchange Points

Exchange ▼ ASN	IPv4 IPv6	Speed RS Peer
AMS-IX 16509	80.249.210.100 2001:7f8:1::a501:6509:1	400G <input type="radio"/>
AMS-IX 16509	80.249.210.217 2001:7f8:1::a501:6509:2	400G <input type="radio"/>
AMS-IX Chicago 16509	206.108.115.36 2001:504:38:1:0:a501:6509:1	100G <input type="radio"/>
AMS-IX Hong Kong 16509	103.247.139.10 2001:df0:296::a501:6509:1	100G <input type="radio"/>
AMS-IX India 16509	223.31.200.29 2001:e48:44:100b:0:a501:6509:2	10G <input type="radio"/>
AMS-IX India 16509	223.31.200.30 2001:e48:44:100b:0:a501:6509:1	10G <input type="radio"/>
BBIX Osaka 16509	218.100.9.24 2001:de8:c:2:0:1:6509:1	40G <input type="radio"/>
BBIX Tokyo 16509	218.100.6.52 2001:de8:c::1:6509:1	200G <input type="radio"/>
BBIX Tokyo 16509	218.100.6.207 2001:de8:c::1:6509:2	200G <input type="radio"/>
BCIX BCIX Peering LAN 16509	193.178.185.95 2001:7f8:19:1::407d:1	200G <input type="radio"/>
BIX.BG Main 16509	193.169.198.87 2001:7f8:58::407d:0:1	100G <input type="radio"/>
RNIX	194.53.172.122	100G

Private Peering Facilities

Facility ▼ ASN	Country City
151 Front Street West Toronto 16509	Canada Toronto
25 Lake Street / 250 Front Street West	Canada

Telia Carrier

Organization	Telia Group
Also Known As	TeliaSonera, Telia, TSIC
Company Website	http://www.teliacarrier.com/
Primary ASN	1299
IRR as-set/route-set ?	RIPE::AS-TELIANET RIPE::AS-TELIANET-V6
Route Server URL	
Looking Glass URL	https://lg.telia.net/
Network Type	NSP
IPv4 Prefixes ?	426000
IPv6 Prefixes ?	40000
Traffic Levels	1 Tbps+
Traffic Ratios	Balanced
Geographic Scope	Global
Protocols Supported	<input checked="" type="radio"/> Unicast IPv4 <input type="radio"/> Multicast <input checked="" type="radio"/> IPv6 <input type="radio"/> Never via route servers
Last Updated	2020-02-05T11:43:25Z
Notes ?	<p>IPv4 + IPv6 Prefixes above would be actuals, not proposed max- prefix values.</p> <p>AS1299 is matching RPKI validation state and reject invalid prefixes from peers and customers. Our looking-glass marks validation state for all prefixes. Please review your registered ROAs to reduce number of invalid prefixes.</p> <p>All trouble ticket requests or support related emails should be sent to carrier-csc@teliacompany.com.</p>

Peering Policy Information

Peering Policy	https://www.teliacarrier.com/dam/jcr:d1e83942-3db1-4334-a5f8-431578633d26/Telia_Carrier_Global_Peering_Policy.pdf
General Policy	Restrictive

Public Peering Exchange Points

Exchange ▼ ASN	IPv4 IPv6	Speed RS Peer
-------------------	--------------	------------------

No filter matches.
You may filter by **Exchange**, **ASN** or **Speed**.

Private Peering Facilities

Facility ▼ ASN	Country City
365 Data Centers Buffalo (BU1) 1299	United States of America Buffalo
365 Data Centers Detroit (DT1) 1299	United States of America Southfield
365 Data Centers Nashville (NA1) 1299	United States of America Nashville
365 Data Centers Tampa (TA1) 1299	United States of America Tampa
3U Rechenzentrum Berlin 1299	Germany Berlin
Altus IT 1299	Croatia Zagreb
Borovaya 57 1299	Russia St. Petersburg
CE Colo Prague 1299	Czechia Prague
CINECA - DC NaMeX 1299	Italy Roma
COD BM-18 1299	Russia St.Petersburg
Caldera21 1299	Italy Milan
CarrierColo Berlin Luetzow (I/P/B/ site B) 1299	Germany Berlin
Cologix MTL3 1299	Canada Montreal
Cologix TOR1 1299	Canada Toronto

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 \$ cost 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

Transits Chosen	Configuration Complexity
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 providers	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 1Gbps links to 4 different ISPs will usually cost more than having 2x 10Gbps links to 2 different ISPs
 - ▣ Yet bandwidth of latter (20Gbps) is greater than that of former (4Gbps) 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 work out what makes financial & operational sense

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

- Regional Internet Exchanges are NOT built from scratch
 - Even today, there are too many “well meaning” attempts to build a Regional IXP
 - There have been several attempts since the 1990s, all have failed
 - Yet there are still entities determined to repeat the same mistakes thinking they can get a better outcome

- Definition: A Regional IXP is a Local IXP that has become so successful that it attracts members from outside its normal service area

Regional Internet Exchange Point

- Regional IXPs are also local IXPs:
 - Regional ISPs join and 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



Value propositions

- Peering at a local IXP
 - Reduces latency & transit costs for local traffic
 - Improves Internet quality perception
 - Encourages local Internet economy (content, hosting, “cloud” services)
- 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