



Peering for Pleasure and Profit

Nigel Titley

Introductions

- Thanks very much for inviting me to speak to you
- Currently Head of Capacity Planning and Peering and Transit strategy at Easynet/Sky
- RIPE NCC Board member
- Peering Coordinator at British Telecom, Level 3 (Europe and Asia), PacketExchange, and Flag



Agenda

- What am I talking about?
- What is peering and why you should do it
- Exchange points and direct peering
- Tools of the trade
- Peering policy – what is it
- Peering Strategy – what is it
- Sample Strategies
 - Small ISP or content provider
 - National ISP
 - Regional ISP
 - Global with Tier 1 pretensions
- Conclusions
- Questions and Answers



Why the title?

Pleasure: The state or feeling of being pleased or gratified.

1. A source of enjoyment or delight: *The graceful skaters were a pleasure to watch.*
2. Amusement, diversion, or worldly enjoyment: *“Pleasure . . . is a safer guide than either right or duty”* (Samuel Butler).
3. Sensual gratification or indulgence.
4. One's preference or wish: *What is your pleasure?*



Why the title?

Profit:

- An advantageous gain or return; benefit.
- The return received on a business undertaking after all operating expenses have been met.




Why the title?

Peering:


- The act of one national Internet backbone provider accepting and passing traffic from another national provider. See [NAP](#).



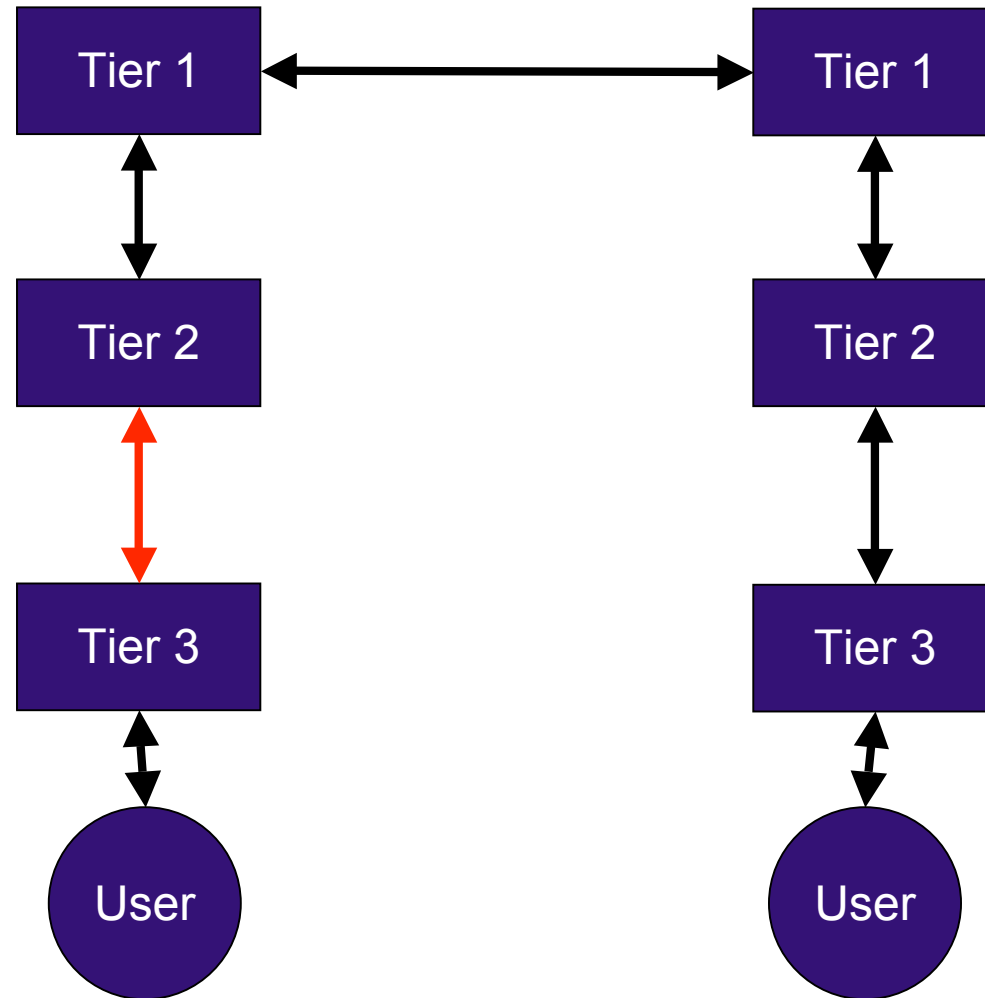
Peering – what is it?




Peering – what is it?




Peering – what is it?




Peering – what is it?




Peering – what is it?




Peering – what is it?




Peering – what is it?



Peering – what is it?



So what is the problem?

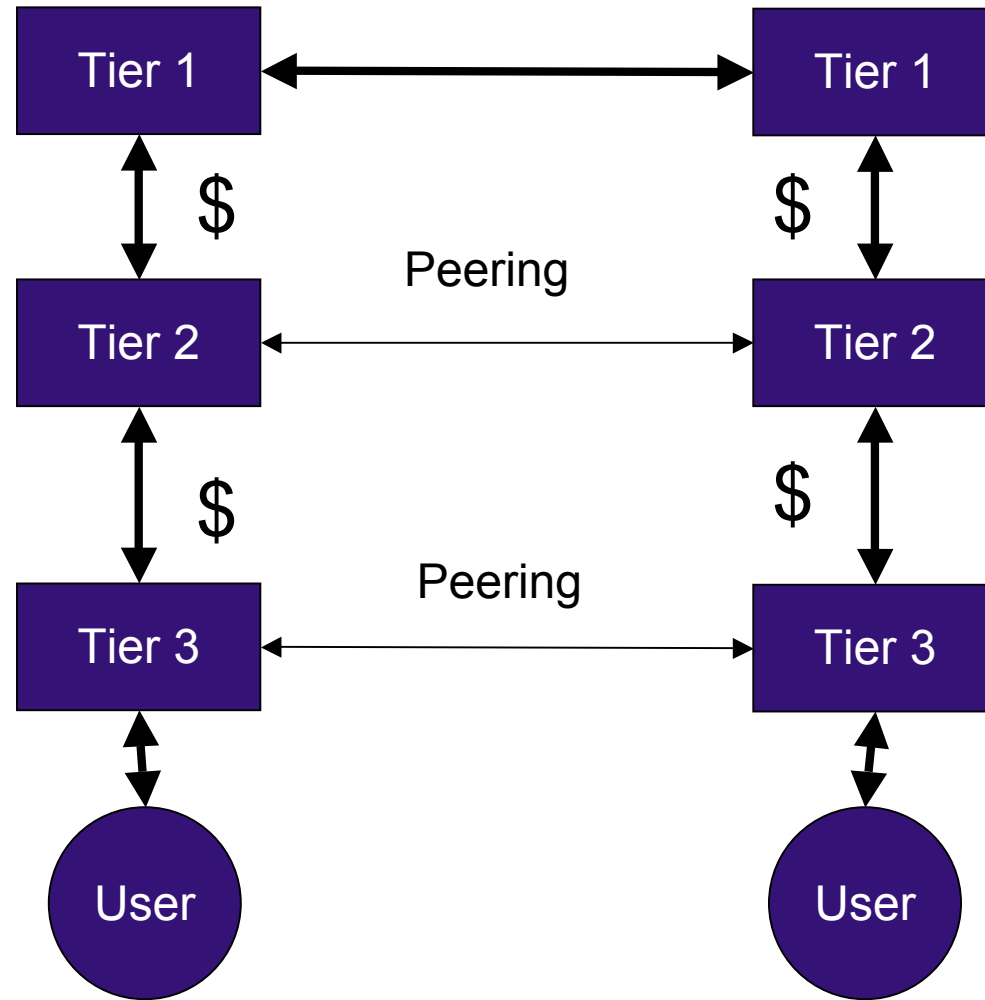


Other Problems

- Latency
 - Traffic may dog-leg via the US or Europe
 - Especially important for peer to peer traffic or gaming
- Congestion
 - Expensive international links
 - Makes everything unresponsive
- Jitter (or delay variance)
 - A combination of the above
 - Makes VOIP and video difficult to use or unusable



What do we do to fix it?



Advantages

- Tier 2 and 3 pay less transit charges (Hurrah!)
- Local traffic stays local
 - Lower latency
 - Less jitter
 - Less chance of congestion
 - Less dependency on external factors (like undersea cable breaks)
- Cooperation between ISPs
 - Overall better service
 - Possible moves towards a trade association



Disadvantages

- Tier 1 sees less revenue (but who cares)
- Tier 2 may see less revenue (but is paying less to Tier 1)
- Management may see peering as cooperation with potential competitors (but we all know how to manage our management don't we?)




How do we do it?

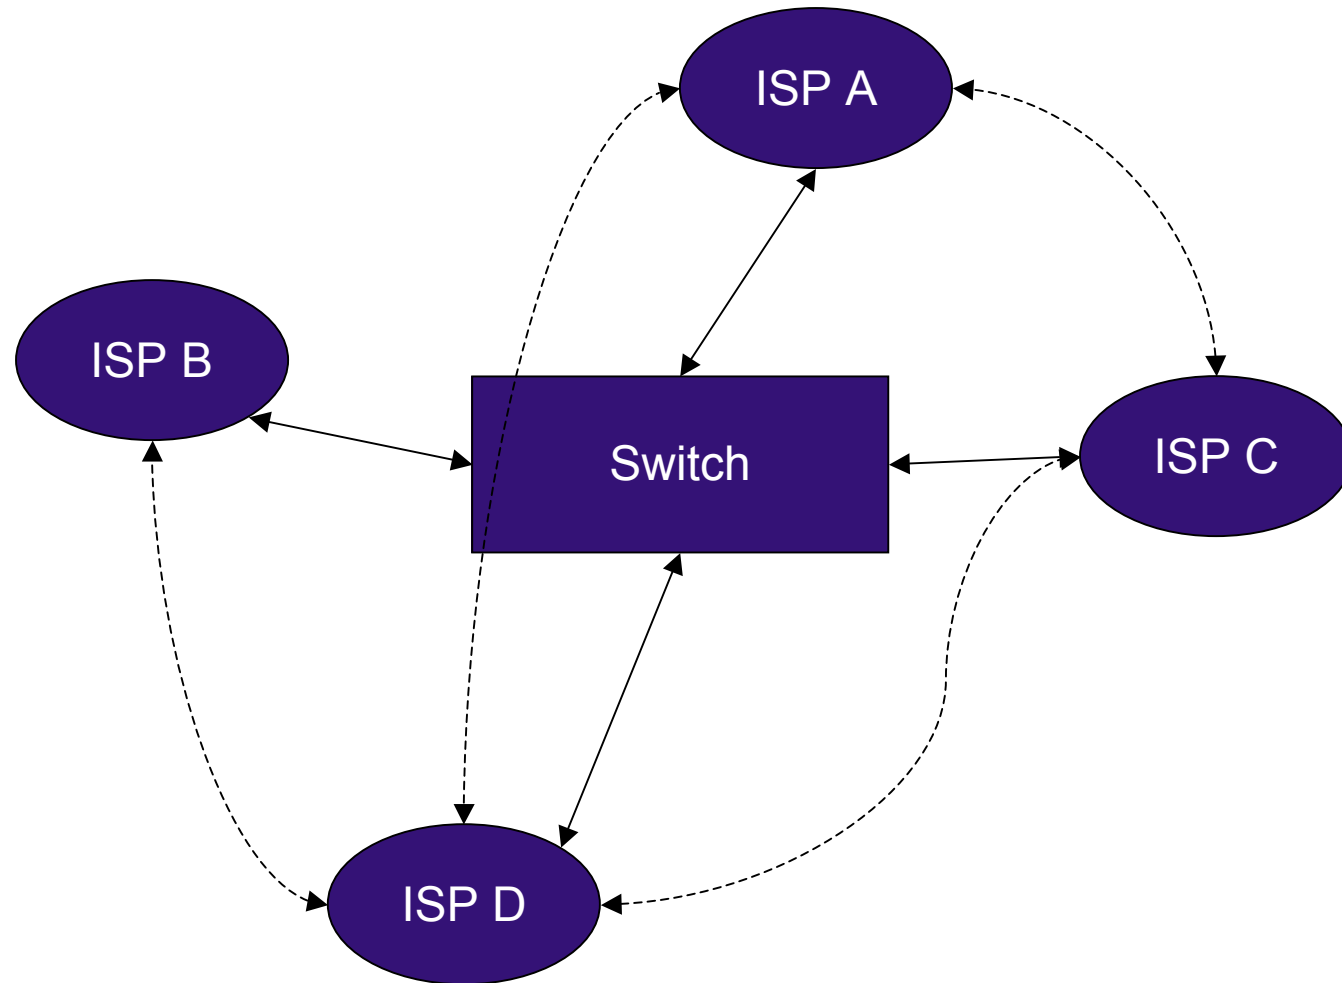
- Via an Internet Exchange Point
 - A neutrally managed layer 2 switch
- Via direct peering
 - A direct connection between two ASes



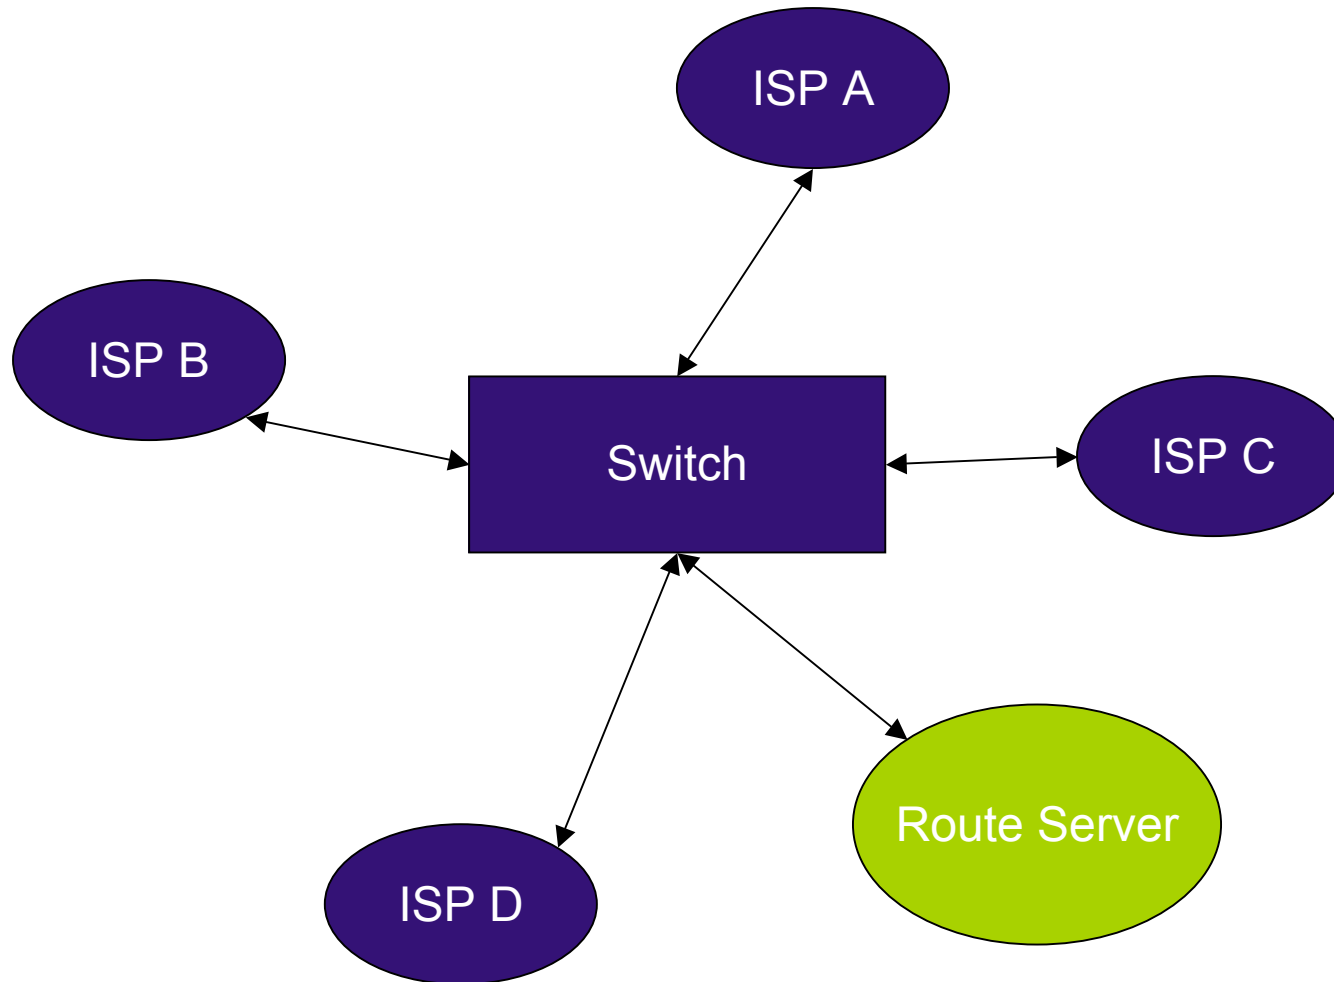
Typical IXP (Physical)



Typical IXP (Logical)



Typical IXP (Physical with route server)

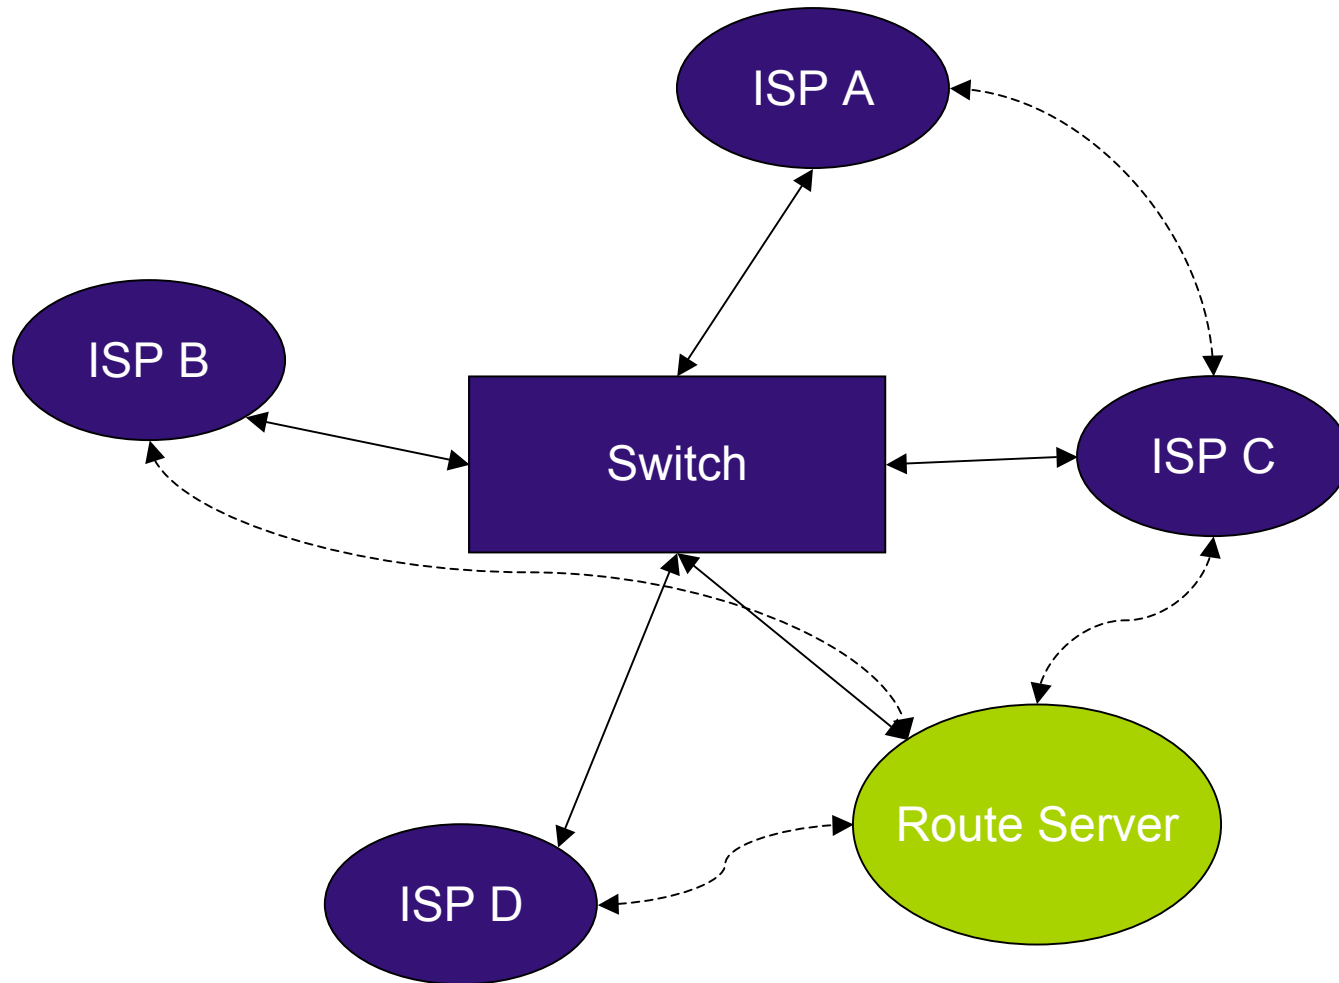


Route Server

- Typically a PC running UNIX/Linux
- Zebra or Quagga
- Sets up BGP sessions with IXP members
- Distributes routes (not traffic)
- May be mandatory or optional



Typical IXP (Logical with route server)




IXP pros and cons

- Pros
 - Only one connection needed
 - Can be very cheap
 - All potential peers immediately available (especially with route server)
 - Can form the basis for cooperative ventures such as trade associations
- Cons
 - Infrastructure may congest (unlikely)
 - Single point of failure
 - Bad traffic (broadcast storms) may disrupt peering
 - Lack of flexibility (with route server)
 - May be difficult to measure traffic to individual peers
 - There may not be an IXP available (so create one, see PCH)



Direct Peering



Direct peering pros and cons

- Pros
 - Easy to see how much traffic is flowing to your peer
 - No single point of failure
 - No interference between peering session
- Cons
 - Port required for each peer (expensive)
 - Bringing up session needs physical installation (so tends to inhibit peering)
 - Cannot share bandwidth between several peers



Compromises

- Start peers on a shared infrastructure (IXP)
- Measure peering flows
- Migrate onto direct peerings when economical to do so
- Best of both worlds
- Assumes that you have the means to measure traffic flows



Tools of the Trade

- Good business card file
- PeeringDB
- Arbor (or similar)
- Renesys
- Local database
- Friendly and engaging smile



Business card file

- Contact details
- Phone number
- Email address
- Write things on the back (like AS number)
- May get extensive
- Review regularly (but don't throw away cards, peering coordinators move between companies but keep the same job)



PeeringDB

- Free!
- Incredibly useful resource
- <http://www.peeringdb.com>
- Guest access (guest/guest)
- Register your own account
(<https://www.peeringdb.com/registration/register.php>)
- Enter your own network's details
- Easily search for details of potential peers



Peering DB

The screenshot shows the Peering DB website interface. The browser window title is "Home Page - Mozilla Firefox" and the address bar shows "https://www.peeringdb.com/private/index.php". The page is divided into several sections:

- 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

Global System Statistics

Total Peering Networks	727
Total Public Exchange Points	163
Total Unique Public Exchange Presences	2599
Total Private Facilities	284
Total Unique Private Facility Presences	2024

Your User Account Status

Account Login	ntitley
Access Level	Level 2 (Normal User)
Peering Record	Easynet Group plc

Last 15 Updated Participants

Company Name	ASN	Date Last Updated
Quality Technology Services	4513	3/30/07, 04:17:32 PM GMT
Server Central Network (scnet)	23352	3/30/07, 04:01:06 PM GMT
AS35701	35701	3/30/07, 01:01:41 PM GMT
Saqashimbun Co., Ltd	18150	3/30/07, 05:46:53 AM GMT
Valueclick	25751	3/30/07, 02:38:42 AM GMT
Primus Telecommunications - US	11867	3/29/07, 08:16:58 PM GMT
Dailymotion	41690	3/29/07, 06:20:16 PM GMT
LeaseWeb	16265	3/29/07, 06:12:19 PM GMT
Eweka Internet Services	12989	3/29/07, 06:07:37 PM GMT
OCCAID	30071	3/29/07, 05:51:58 PM GMT
EWETel	9145	3/29/07, 05:16:31 AM GMT
Ipercast	25286	3/28/07, 10:48:48 PM GMT
Internap	22212	3/28/07, 10:29:13 PM GMT
Universal Telecom, Inc	16402	3/27/07, 10:58:20 PM GMT



PeeringDB (search)

Peering Networks Search/List - Mozilla Firefox

File Edit View History Bookmarks Tools Help

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](#)

Search Peering Networks

Company Name Primary ASN

Network Type IRR Macro

Traffic Levels General Peering Policy

Traffic Ratio Geographic Scope

Peering Networks Search Results

Company Name	ASN	General Policy	Traffic Levels	Network Type	Network Scope	Public Count	Private Count
45RU HostAway	24541	Open	100-1000Mbps	Cable/DSL/ISP	Asia Pacific	1	0
4U Networks	25369	Open	100-1000Mbps	Content	Regional	5	2
AARNet	7575	Selective	1-5Gbps	Educational/Research	Asia Pacific	12	5
Abovenet Communications Inc.	6461	Selective	100+ Gbps	NSP	Global	0	0
Absolight	29608	Open	20-100Mbps	Content	Europe	4	1
Accelerated Connections Inc	21570	Selective	Not Disclosed	Cable/DSL/ISP	North America	1	0
ACONet	1853	Open	5-10Gbps	Educational/Research	Regional	1	2
Adelphia Communications	19548	Restrictive	20-50 Gbps	Cable/DSL/ISP	North America	10	9
Advanced Knowledge Networks	14453	Open	0-20 Mbps	NSP	North America	1	1
Axiomus Ltd	6770	Open	100-1000Mbps	NSP	Europe	4	5
Affinity Internet, Inc.	3064	Open	1-5Gbps	Content	Global	1	0
Africa Online Operations (Mauritius) Ltd.	18922	Open	0-20 Mbps	Cable/DSL/ISP	Regional	1	0

1 2 3 4 5 6 7 8 9 10 of 55 Next > Last >>

(c) 2004-2006 PeeringDB, All Rights Reserved. Please contact admin@peeringdb.com with questions/problems.



PeeringDB (search results)

Peering Networks Detailed View - Mozilla Firefox

File Edit View History Bookmarks Tools Help

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			Public Peering Exchange Points					
	Company Name FLAG Telecom	Exchange Point Name		ASN	IP Address	Mbit/sec			
	Also Known As	AMS-IX	15412	195.69.144.72	2000				
	Company Website http://www.flagtelecom.com	Any2	15412	206.223.143.45	1000				
	Primary ASN 15412	DE-CIX	15412	80.81.192.64	2000				
	IRR Record AS-FLAGP	Equinix Ashburn	15412	206.223.115.141	1000				
	Network Type NSP	HKIX	15412	202.40.161.196	2000				
	Approx Prefixes 6000	JPIX	15412	210.171.224.139	1000				
	Traffic Levels 20-50 Gbps	JPNAP	15412	210.173.176.96	1000				
	Traffic Ratios Balanced	KINX	15412	192.145.251.42	1000				
Geographic Scope Global	LAIX	15412	198.32.146.52	1000					
Looking Glass URL	LINX	15412	195.66.226.146	2000					
Route Server URL	LINX	15412	195.66.224.146	1000					
Notes	NYIX	15412	198.32.160.88	1000					
Protocols Supported Unicast IPv4 <input checked="" type="checkbox"/> Multicast <input checked="" type="checkbox"/> IPv6 <input checked="" type="checkbox"/>	1 2 of 2 Next > Last >>								
Date Last Updated 2007-03-26 03:12:14 UTC	Private Peering Facilities								
Peering Policy Information			Facility Name	ASN	City	Country	SONET	Ethr	ATM
Peering Policy URL	1 Wilshire	15412	Los Angeles	US	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
General Policy Selective	FiberNet Telecom Group - 60 Hudson St	15412	New York	US	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
Multiple Locations Not Required	FiberNet Telecom Group New York (111 Eighth Ave)	15412	New York	US	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
Ratio Requirement No	KINX IX Center	15412		KR	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
Contract Requirement Private Only	MEGA iAdvantage Hong Kong	15412	Hong Kong	HK	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
Contact Information									
Role	Contact Name	Telephone	E-Mail						
Policy	Peering duty	+852 2848 0025	peering@flagtelecom.com						
NOC	NOC	+44 (0) 20 8282 0068	noc@flagtelecom.com						
Technical	NOC		noc@flagtelecom.com						

Done www.peeringdb.com



Arbor Peakflow (or similar)

- For each router in network
 - Takes netflow data
 - Takes an iBGP feed
 - Takes SNMP feed
- Used to determine traffic to and from a peer (or potential peer)
- For a potential peer will show the current path
- Also shows
 - Traffic breakdown (by destination port)
 - Traffic flowing from a peer to a peer (useful for detecting misconfigurations)
 - Peering adviser mode, shows peers to aim for
 - Traffic flows within network
 - Historical data
- Disadvantage
 - Cost (licensed per router)
 - May lie to you



Arbor Peakflow (example)

Peer 'AS15169: GOOGLE' Summary - Mozilla Firefox

File Edit View History Bookmarks Tools Help

peakflow™ # Logout Help

System > Alerts > Reports > Administration > Logged in as: nttley 21:34:02 GMT | 04/01/2007

Peer 'AS15169: GOOGLE' Summary Download | Email | Edit

Peer AS15169: GOOGLE ▾

Units bps ▾

Period Today ▾

Update

'AS15169: GOOGLE' Traffic Summary

Legend: TOTAL (grey area), IN (red line), OUT (green line), BACKBONE (blue line), DROPPED (orange line)

Clear All Update

<input type="checkbox"/>	Class	Current	Average	Max
<input checked="" type="checkbox"/>	IN	198.00 Mbps	118.16 Mbps	244.00 Mbps
<input checked="" type="checkbox"/>	OUT	26.04 Mbps	14.76 Mbps	27.68 Mbps
<input checked="" type="checkbox"/>	DROPPED	0.00 bps	0.00 bps	0.00 bps
<input checked="" type="checkbox"/>	BACKBONE	0.00 bps	0.00 bps	0.00 bps
<input checked="" type="checkbox"/>	TOTAL	224.00 Mbps	132.93 Mbps	272.05 Mbps

Clear All Update

Page generation took 4.02 seconds. [Details](#).

For assistance with this product, please contact ops@easynet.net. © 2007 Arbor Networks, Inc. All Rights Reserved.

Done fc0.blon.uk.easynet.net

Renesys

- <http://www.renesys.com>
- Historical routing topology data
- Very useful
 - Debugging routing problems
 - Working out routing topologies
- Peering advisor tool
- Free
 - Give them a peering
 - Get access to basic tools
- Pay
 - Get automatic advisories of routing topology problems to feed to your NOC



Local Database

- Where do I peer
- Who do I peer with?
- How do I contact them?
- Who *don't* I peer with?
- Missing peering points
- Anything else you might want to record



Friendly and Engaging Smile

- It helps to be friendly (but see later caveats)
- Buy drinks
- Bring tee-shirts
- Go to meetings
 - RIPE
 - NANOG
 - APRICOT
 - SANOG
 - MENOOG
 - UKNOF
- Sponsor MENOOG meetings



Policies and Strategies

- Peering Strategy
 - How do I plan to achieve: reduced transit costs, increased profits, better customer experience, world domination etc
 - Private
- Peering Policy
 - What do I tell people who want to peer with me?
 - Should be publicly available (on your web site)
 - Includes contact info
 - Referred to on your PeeringDB page
 - If you are a large player or are very selective then should be objective in order to avoid problems with regulators



Types of Peering Policies

- Open (we peer with anyone)
- Selective (we are a bit fussy about who we peer with)
- Restrictive (we actively discourage people from peering with us)
- Closed (we won't peer with anyone)



Example policy: Small ISP or Content Provider

- Open policy
- Primary aim to reduce costs
- Peering with anyone
- Probably only able to attract peering from other similar peers
- Bring up as many peers as possible while trying to increase market share



Example Policy: National ISP

- Selective Policy
- Aim to exchange traffic with similar sized ISPs
- Will probably involve connecting to foreign peering points
- Avoid peering with customers of existing peers or customers of short term target peers
- Possibly depeer smaller existing peers if this will increase traffic to larger peers
- Analyse traffic and monitor target peers policies
- Acquire more customers (and the right sort of customers)
- Watch the markets and the technical web sites
- Gradually tighten peering policy
- May take several years



Example Policy: Regional with pretensions

- Policy selective verging on restrictive
- Depeer smaller peers if this moves traffic to larger peers or targets
- Adjust BGP policies to concentrate traffic on larger peers, especially if their peering policy changes
- Attempt to negotiate transit contracts with peering escape clauses
- Expand network further
- Acquire more customers
- Study markets and peering flows
- Monitor target peers peering policies and strike immediately you satisfy them



Example Policy: Global tier 1

- Restrictive peering policy
- Stop paying transit charges
- Continually monitor traffic and other tier 1 peers' peering policies
- All peerings probably direct by now
- Meet with your peers peering coordinators whenever possible
- Collect your bonus



Conclusions

- Pleasure
 - It can be a lot of fun being a peering coordinator
 - Cooperation
 - Meeting like minded people
 - Making the internet a better place (without peering there would be no internet)
- Profit
 - Peering is your markup
 - Difference between the price you pay for transit and the price you charge your customers
 - More peering generally means more profit and happier customers



The image shows the word "sly" in a stylized, 3D font. The letters are light blue with a metallic sheen and a dark blue outline. The 's' is a simple, rounded shape. The 'l' is a tall, thin vertical bar. The 'y' is a complex, angular shape with a sharp peak and a tail that curves to the right.

Questions?

nigel.titley@uk.easynet.net