Peering and Network Deployment at 10G

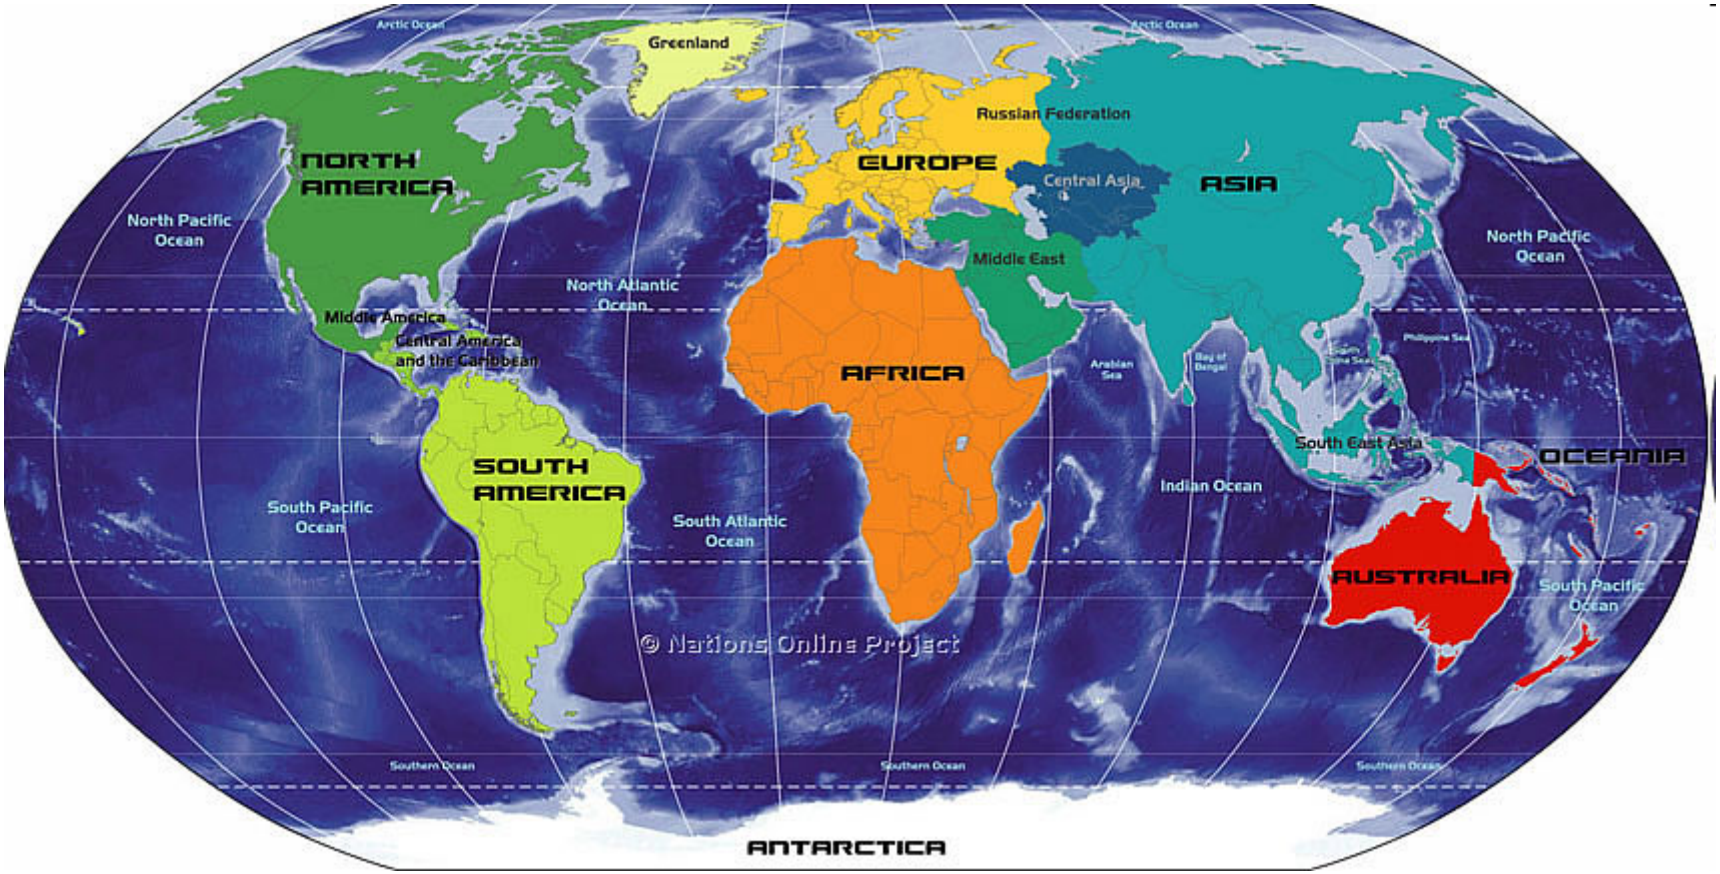

Nigel Titley

Introductions

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

Agenda

- What am I talking about?
- What is peering and why you should do it
- Peering policy – what is it?
- Peering strategy – what is it?
- Exchange points and direct peering
- Building a peering network
- Tools of the trade
- Conclusions
- Questions and Answers

What is Easynet?

- UK and European ISP specialising in Corporate customers
- 1G European network: France, Germany, Italy, Belgium, Netherlands, Spain
- Local Loop unbundling in the UK (over 1000 exchanges undbundled)
- Bought by BSkyB satellite broadcaster (owned by News International)
- Over 1M domestic broadband customers and 40G of traffic
- Traffic mostly eyeballs

What am I talking about?

- Peering and why you should do it
- Why Easynet invested 2.5M in building a 10G network just for peering
- Peering tools of the trade

Definition?

Peering:

- The act of one national Internet backbone provider accepting and passing traffic from another national provider. See NAP.

Peering – what is it?

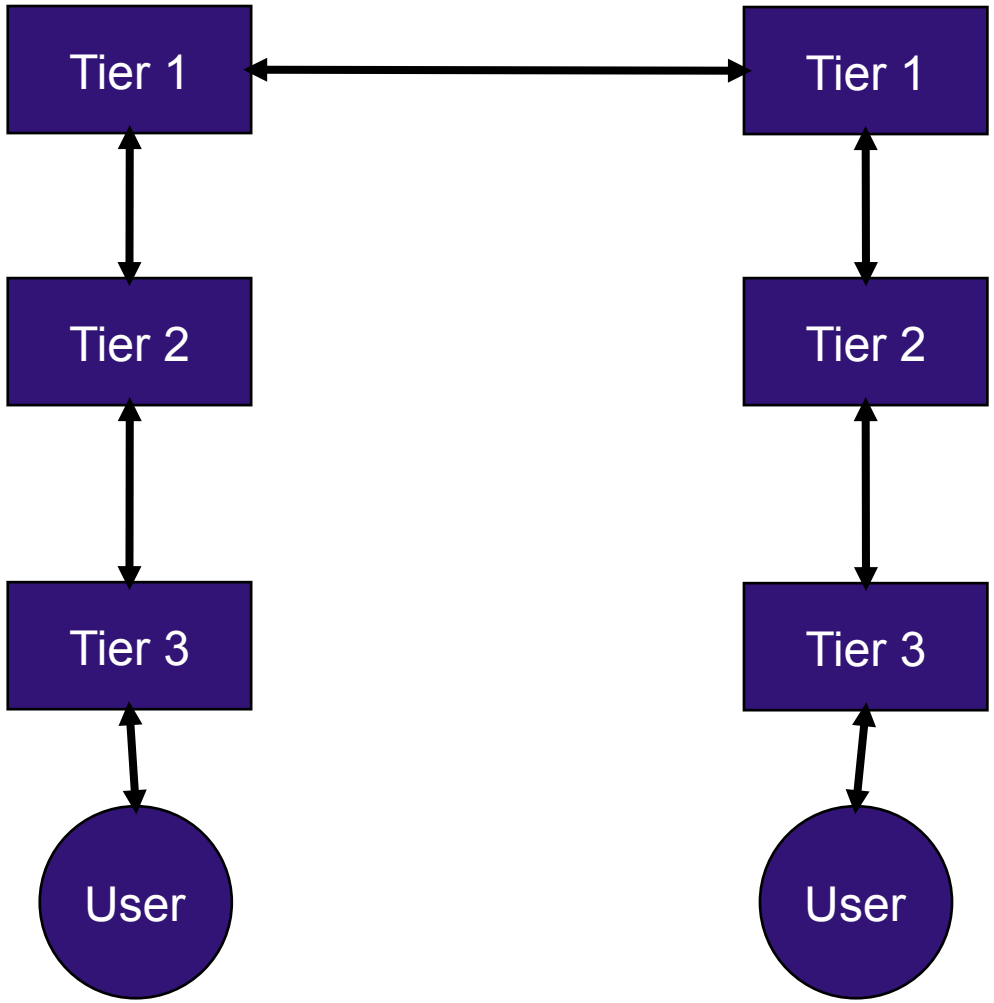

Peering – what is it?

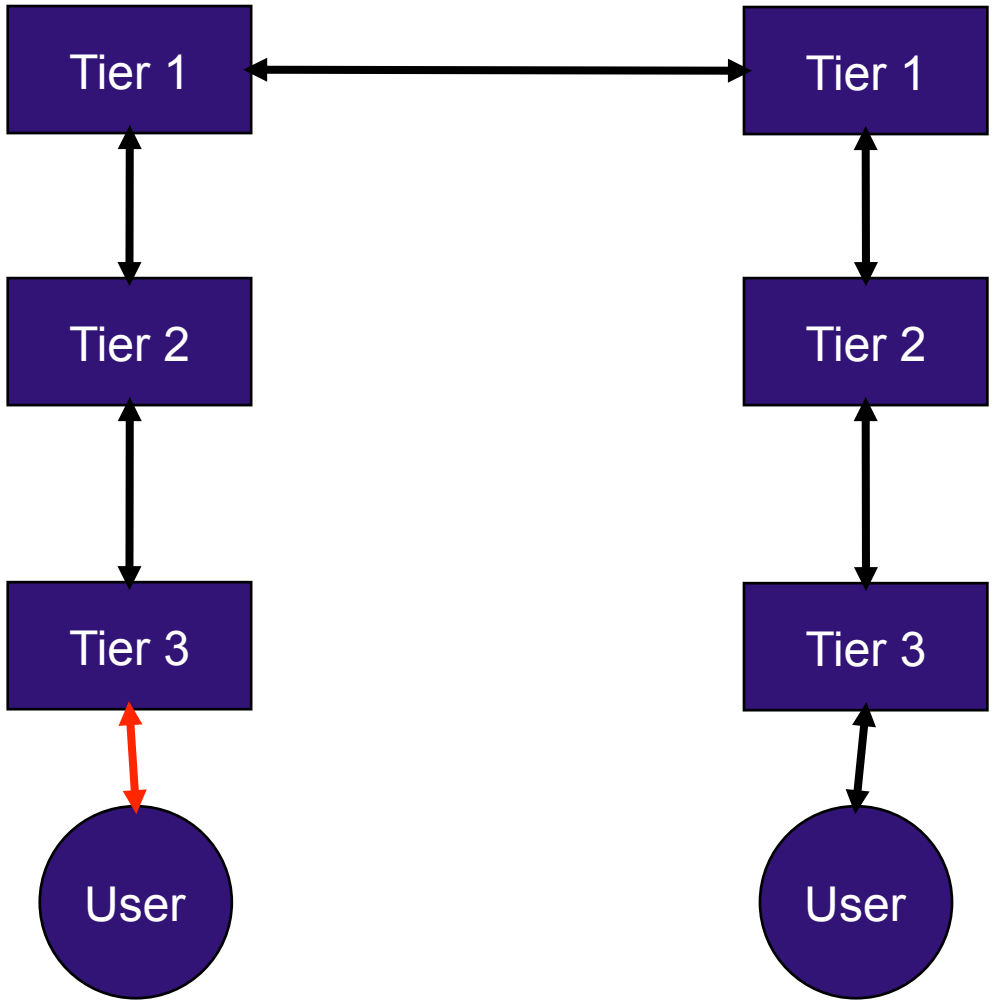

Peering – what is it?

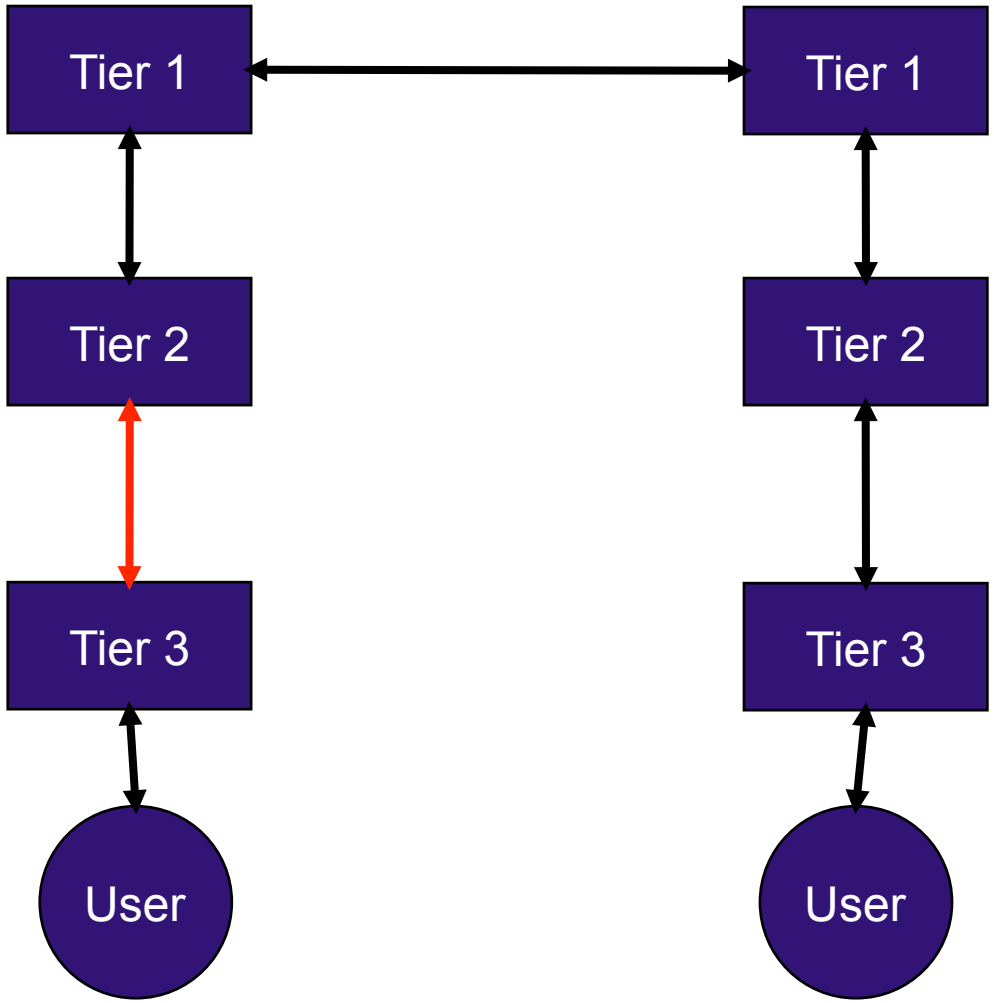

Peering – what is it?

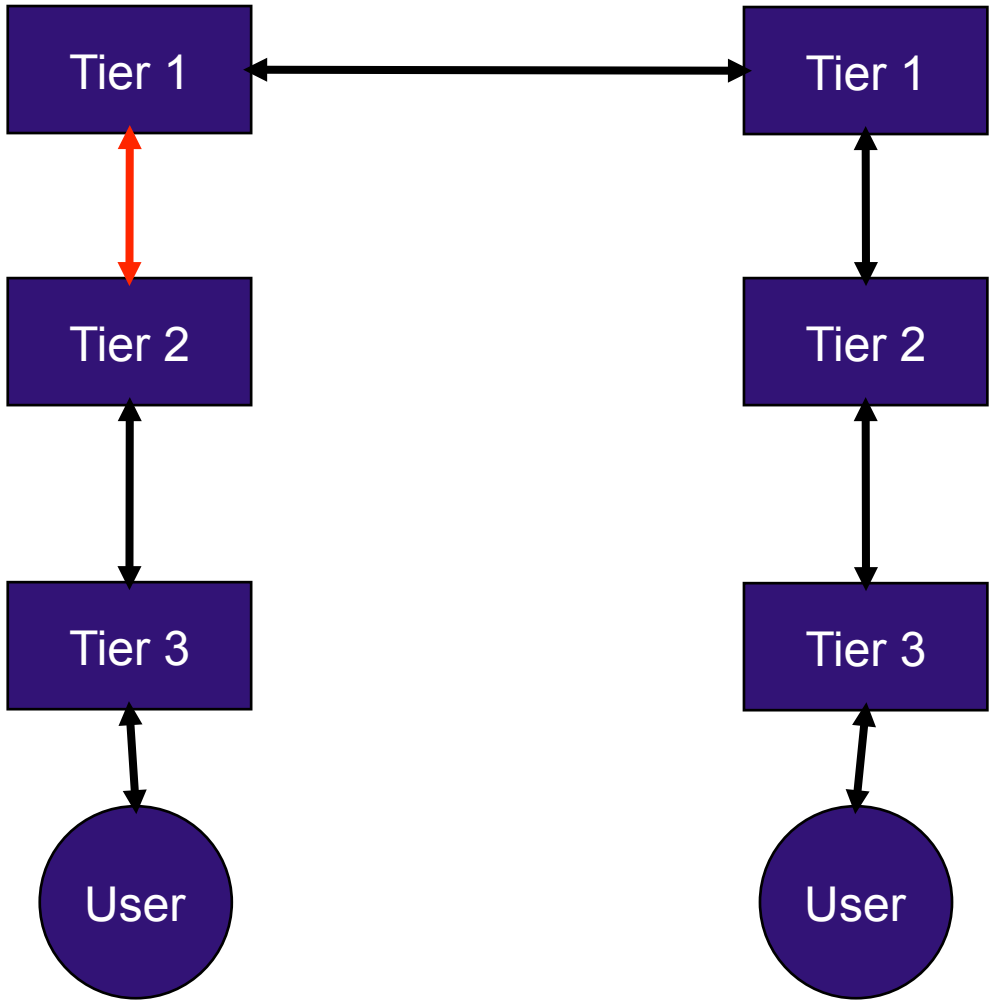

Peering – what is it?

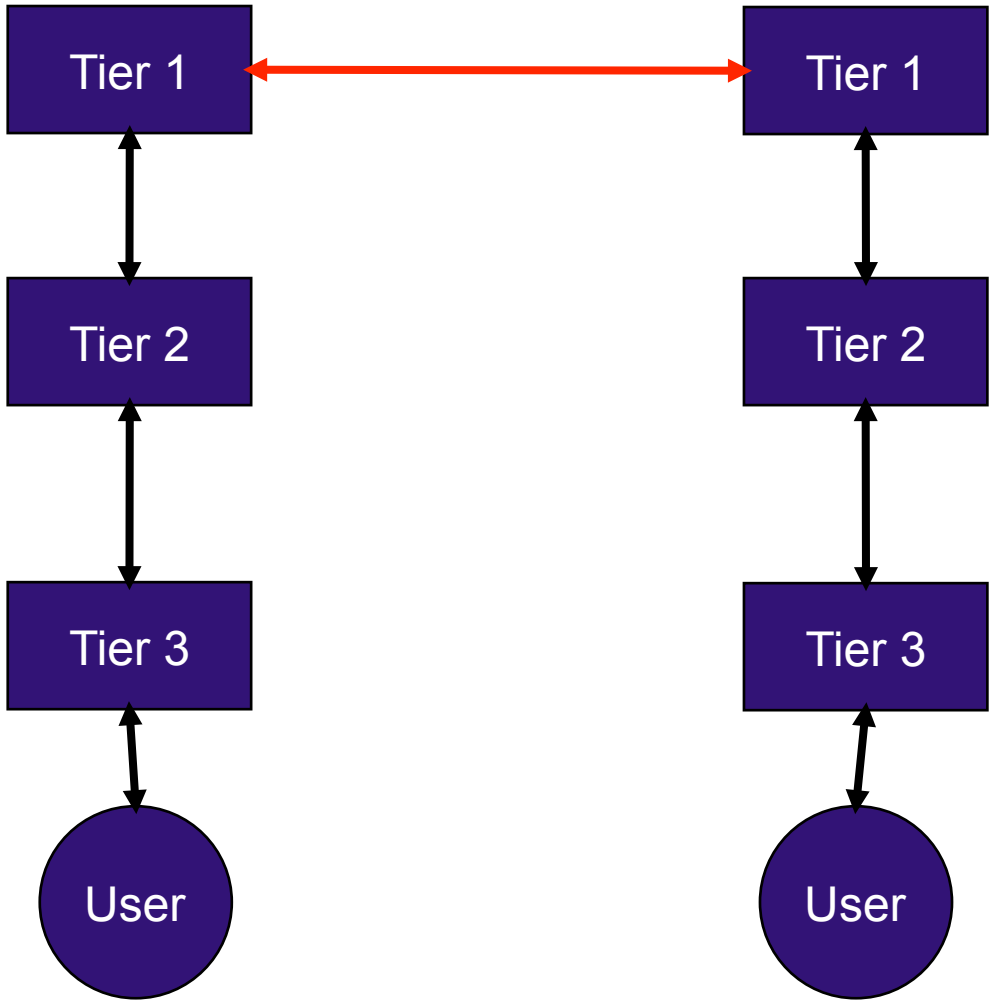

Peering – what is it?

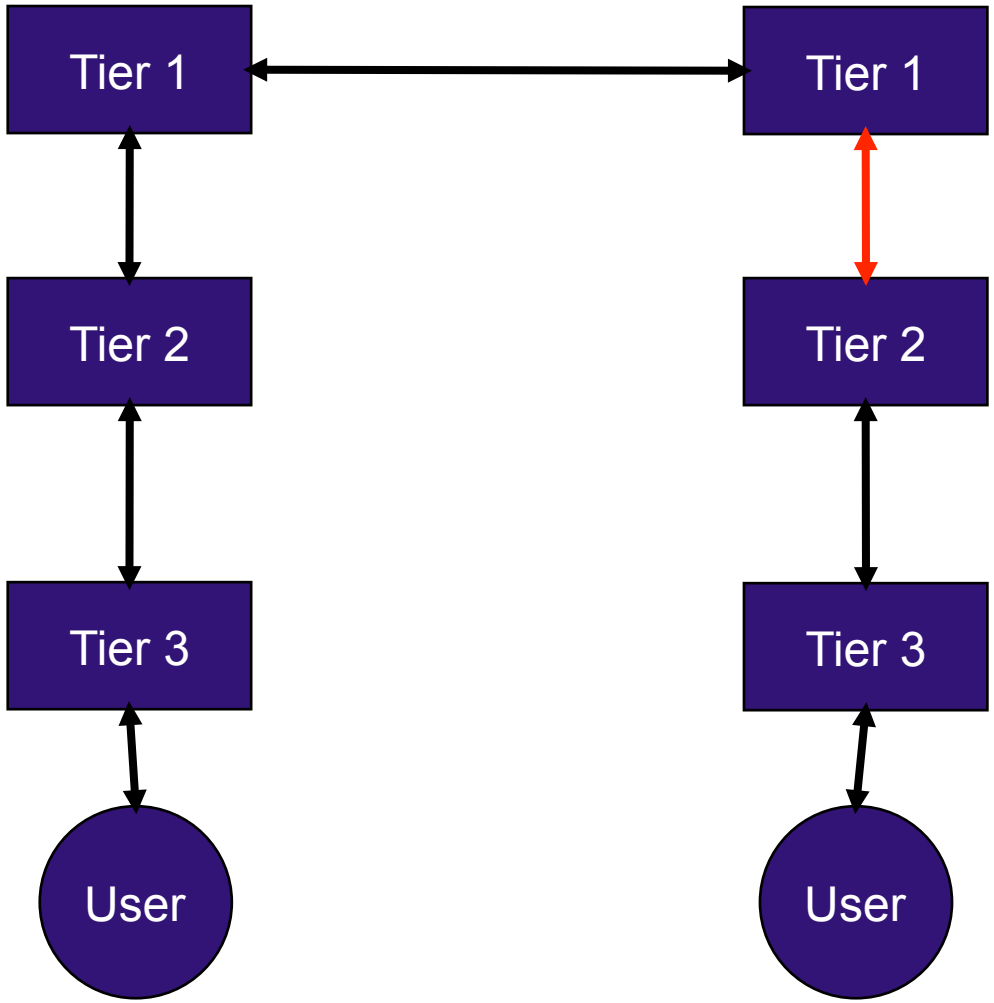

Peering – what is it?

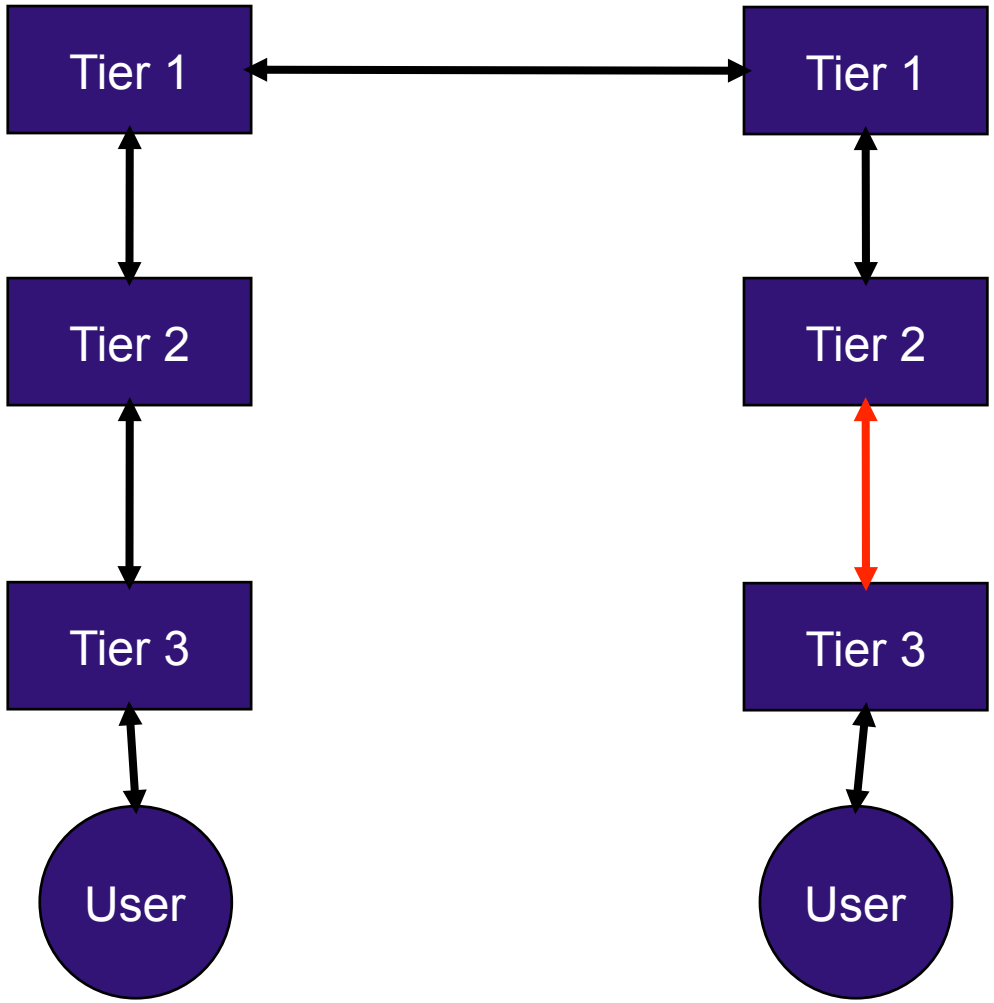

Peering – what is it?

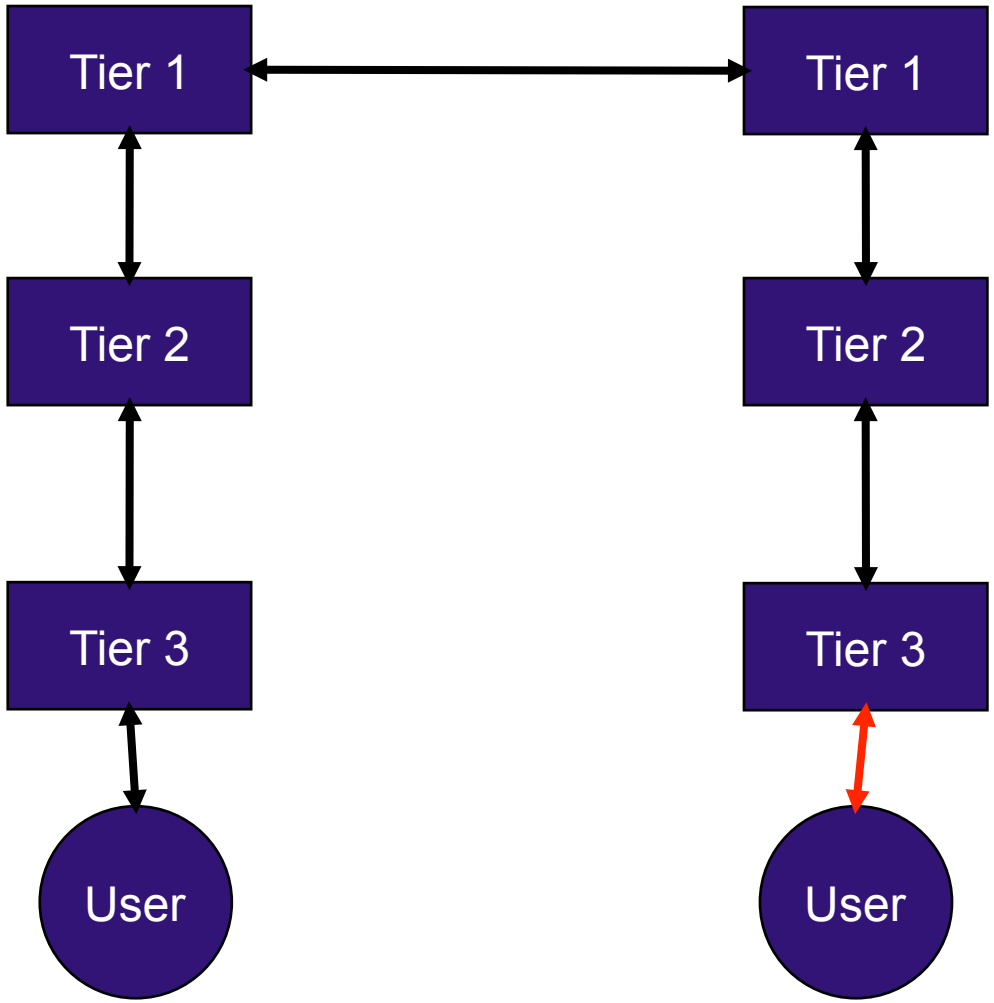

So what is the problem?

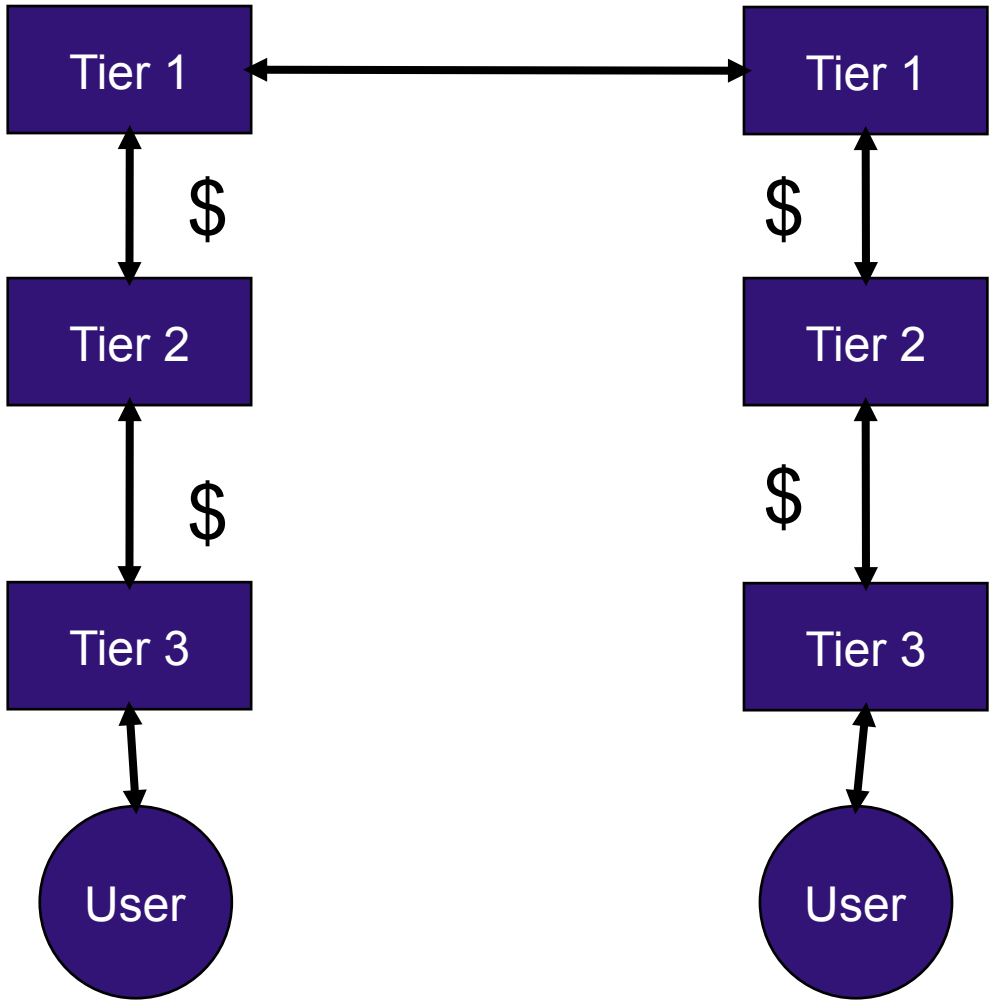

Other Problems

- 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?

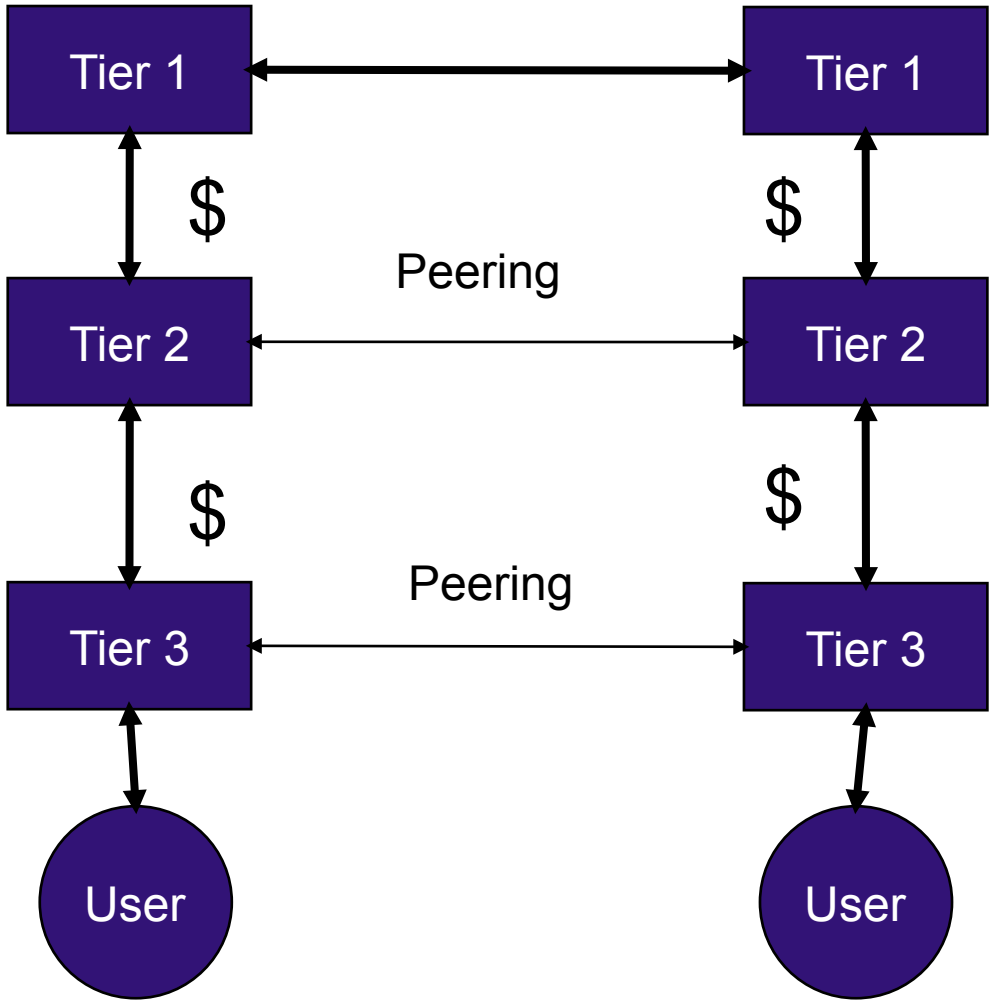

Advantages

- 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?)

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)

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)

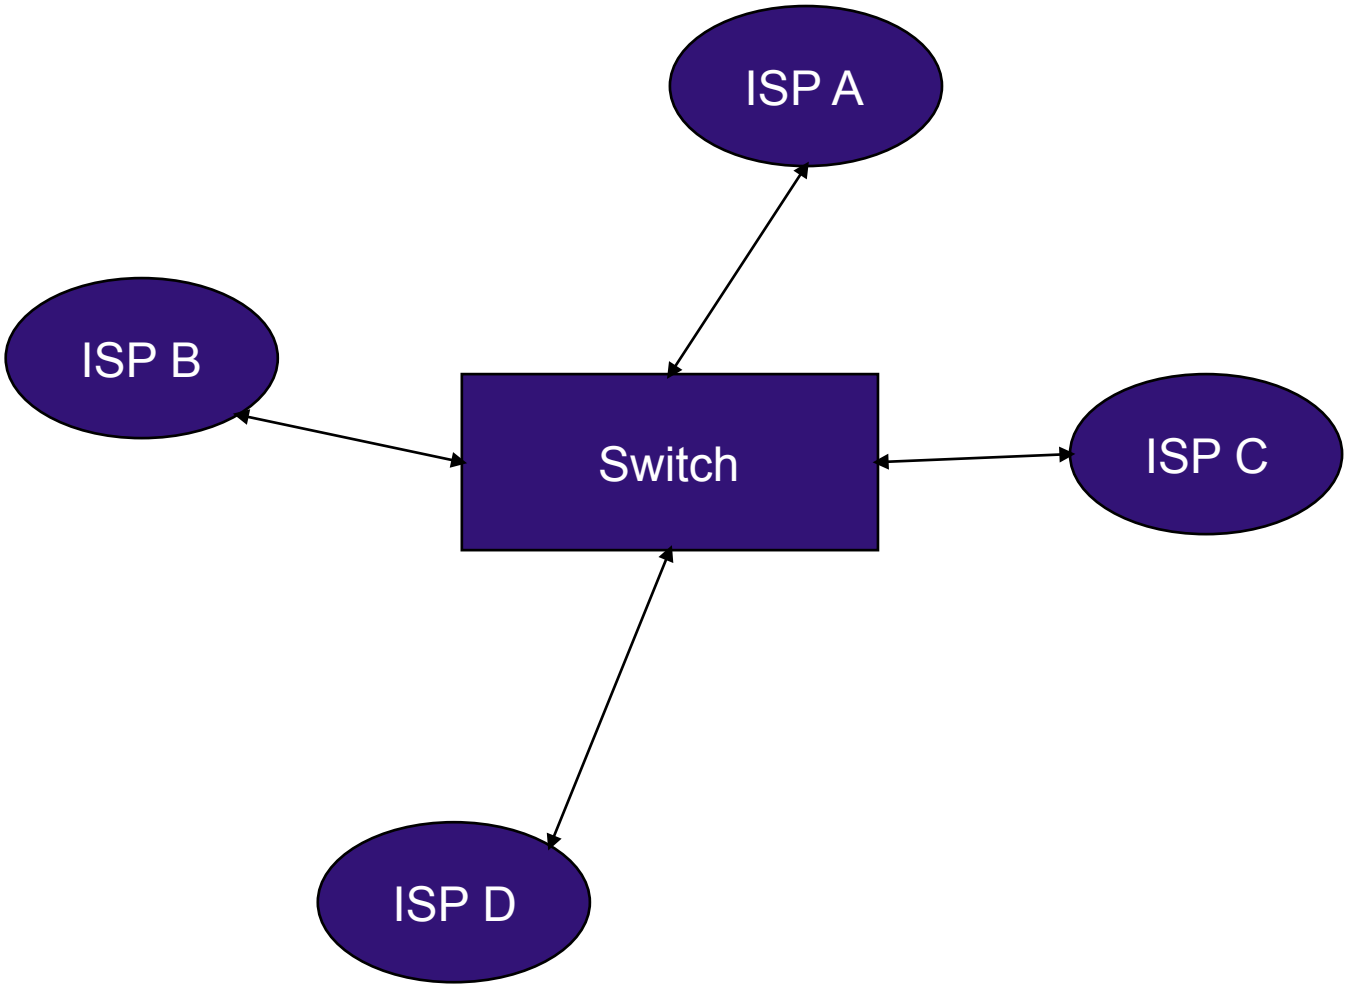

Typical IXP (Logical)

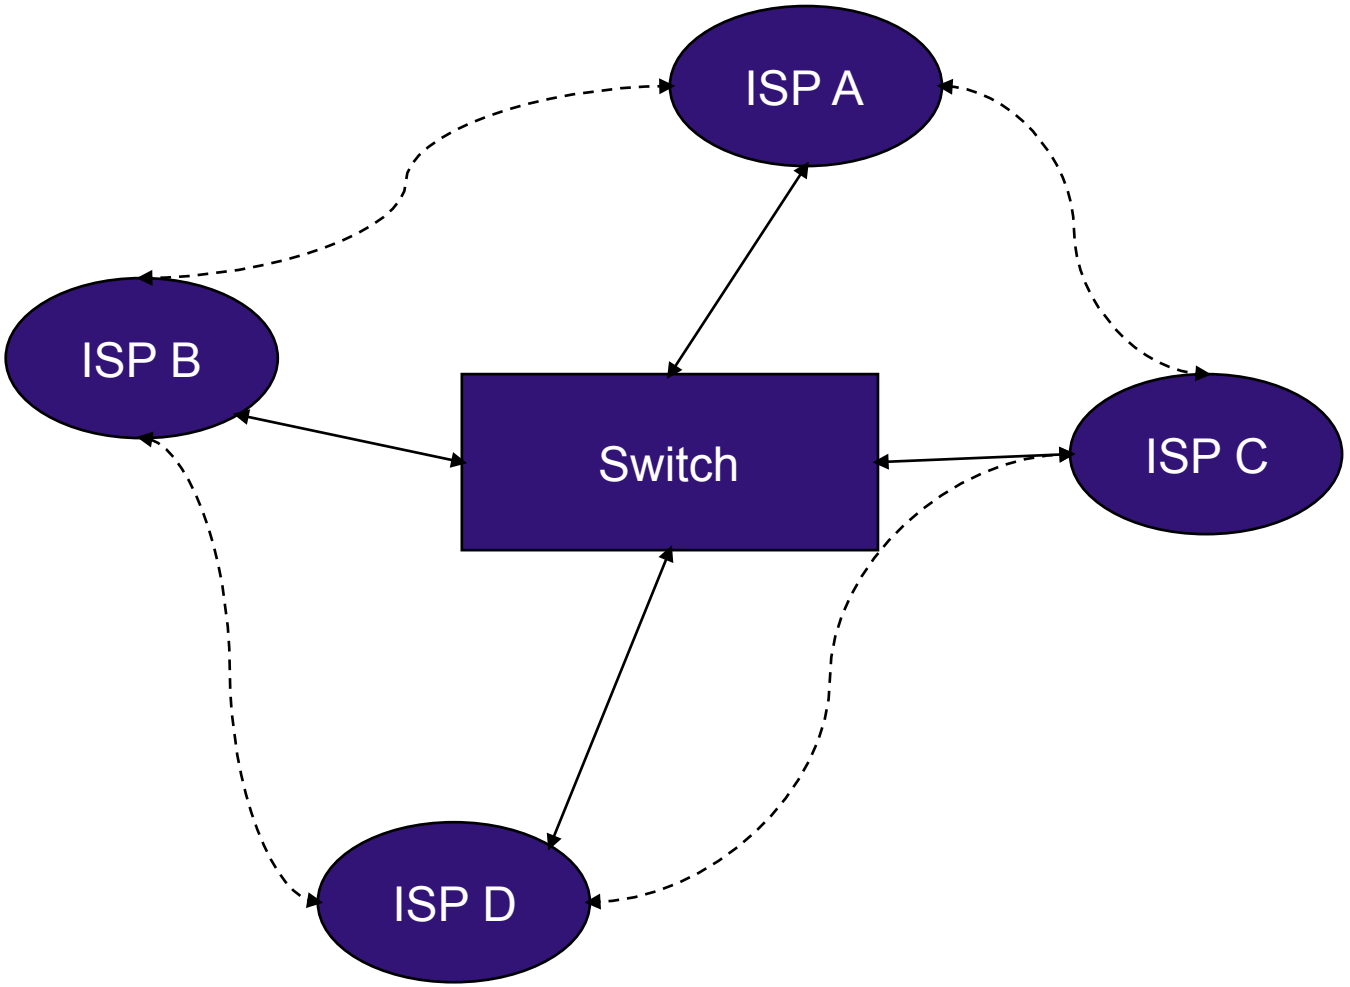

Typical IXP (Physical with route server)

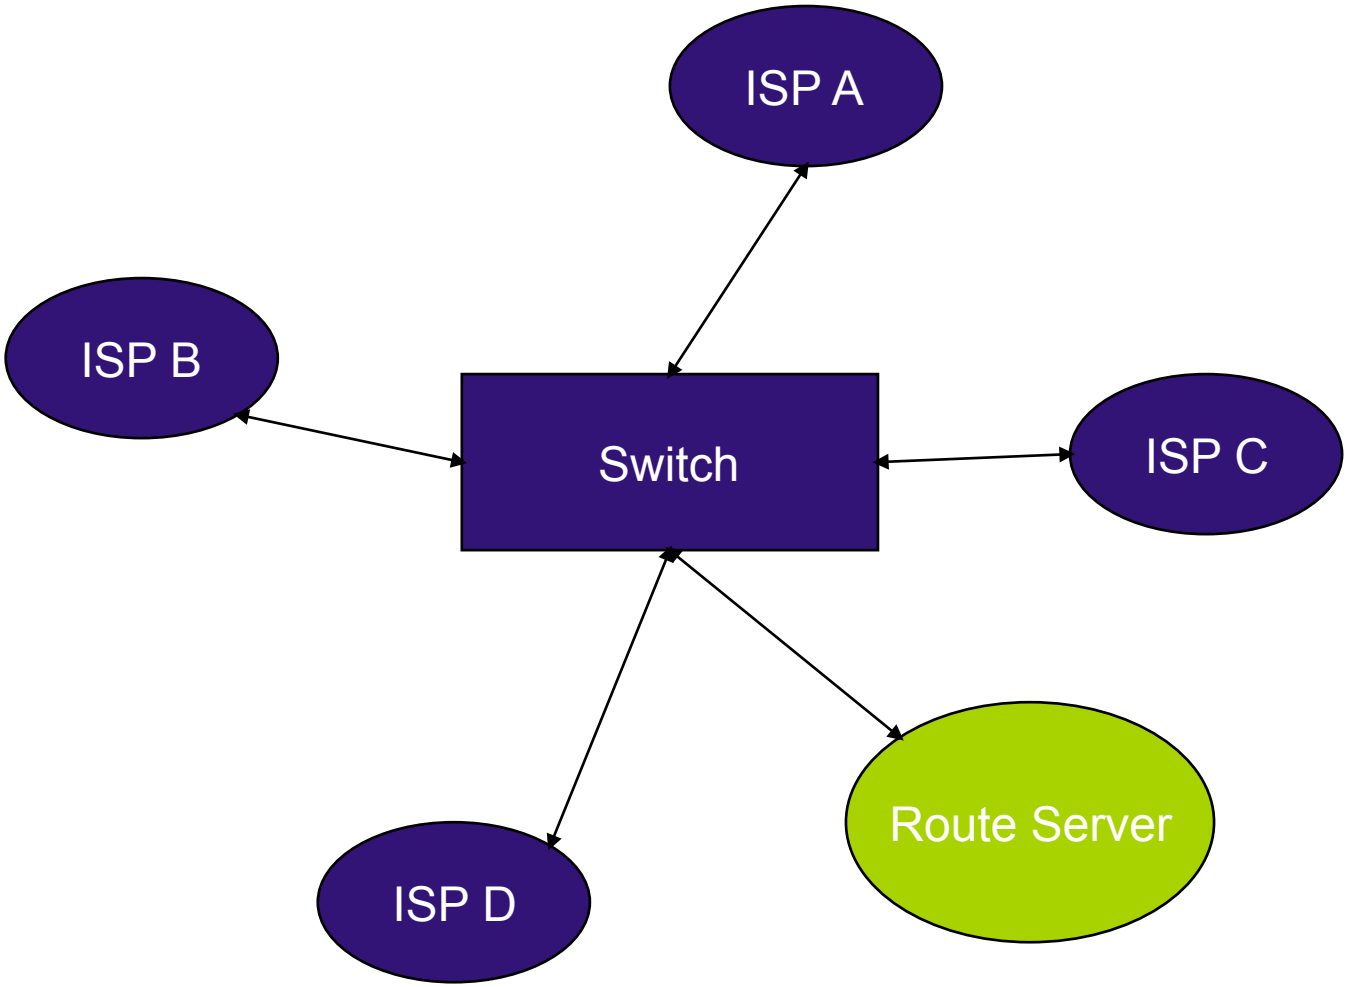

Route Server

- 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)

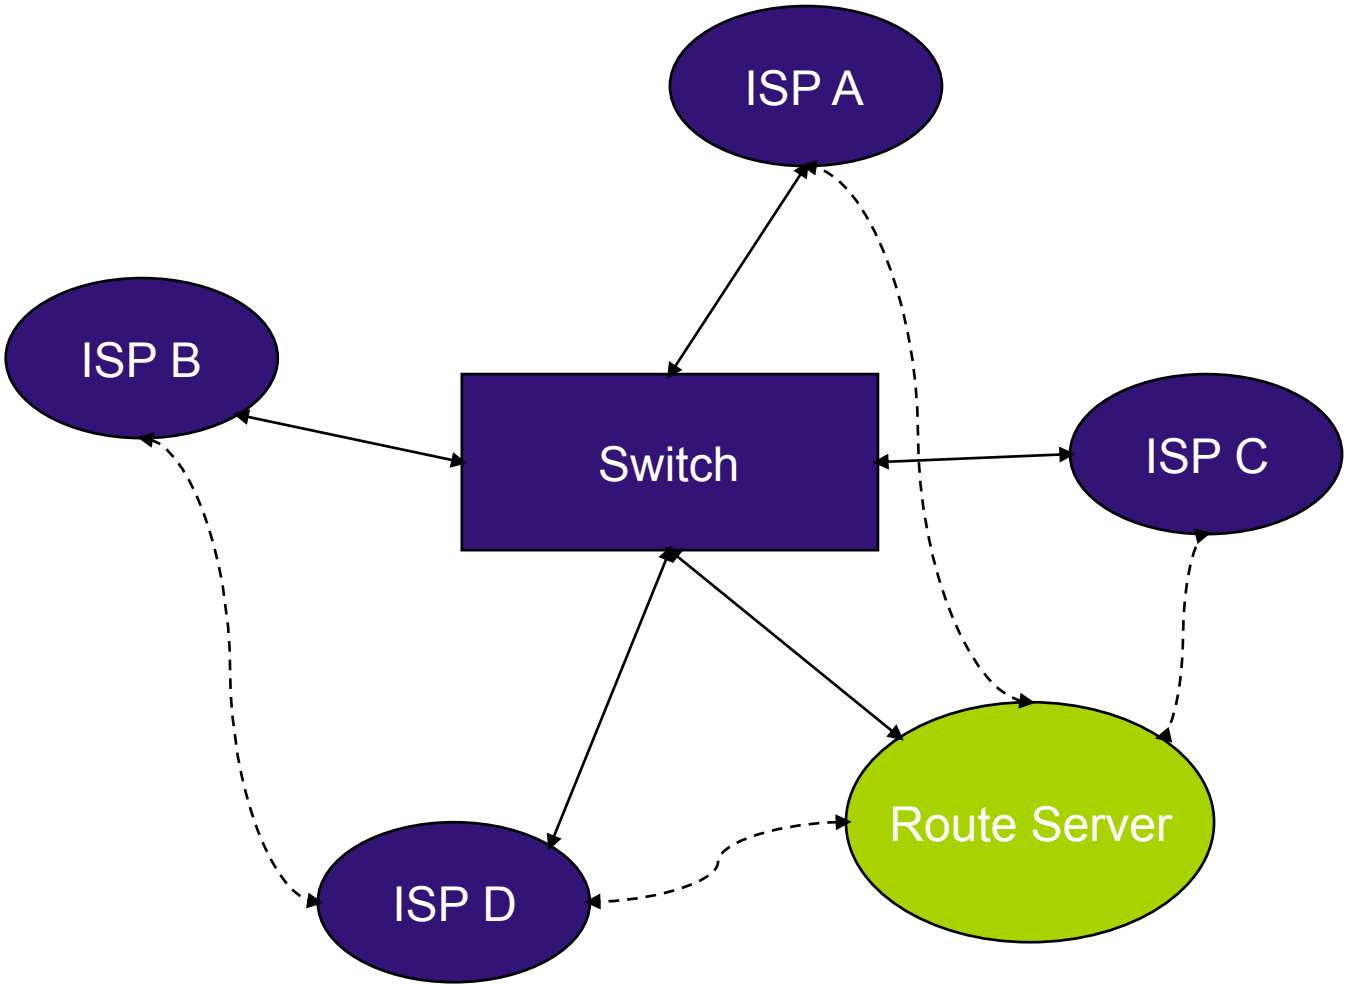

IXP pros and cons

- 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

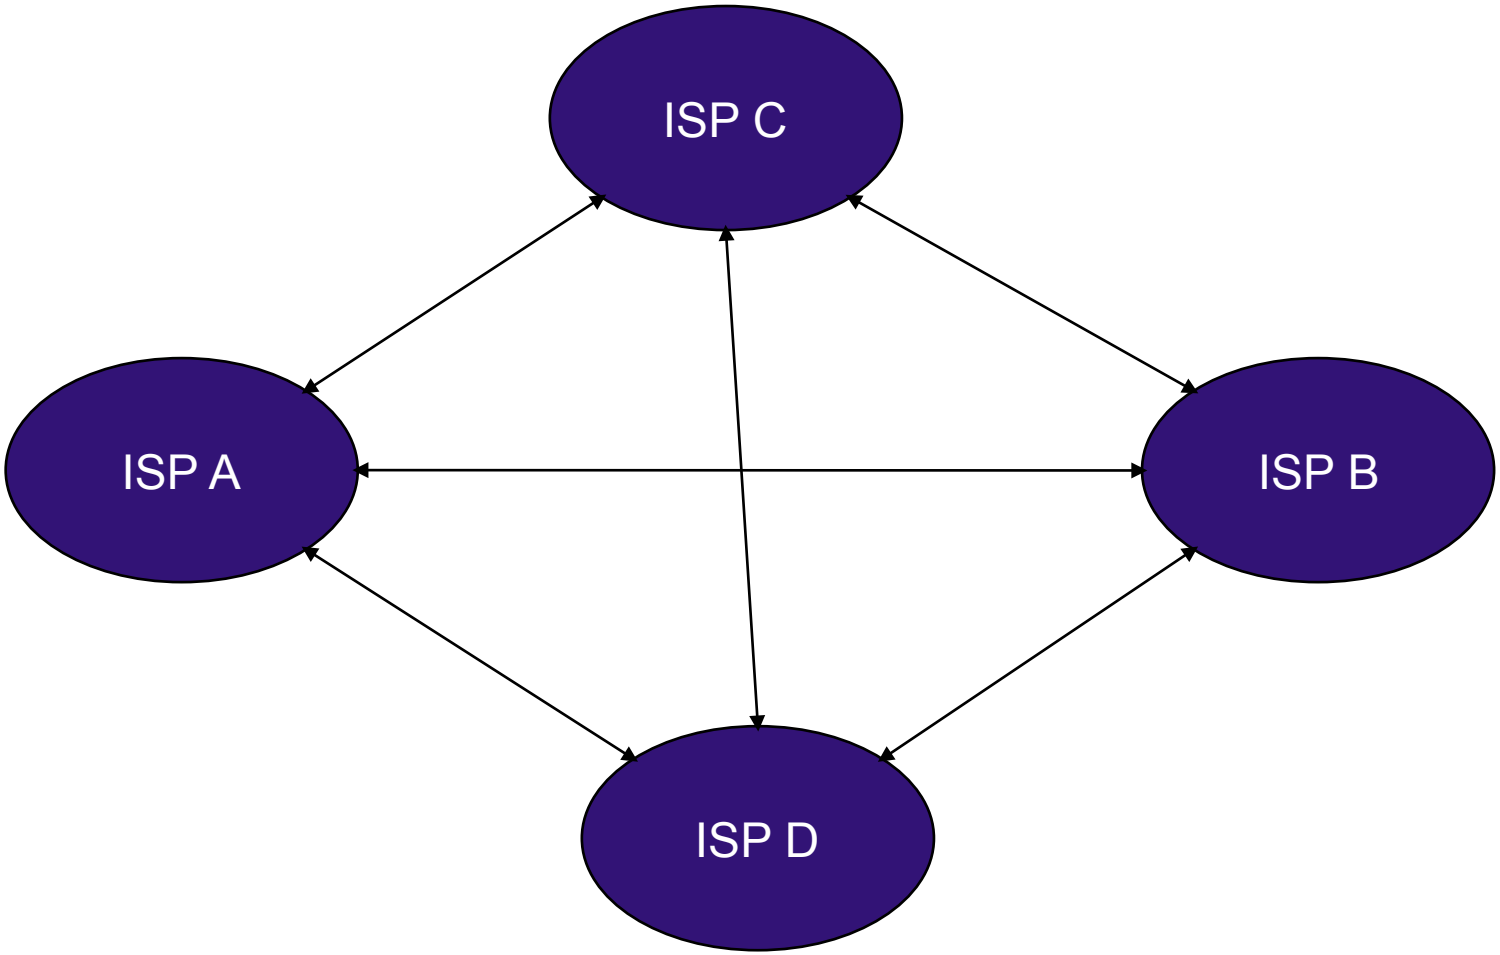

Direct peering pros and cons

- 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

Why build a peering network?

- Cheaper than transit
- More control over traffic
- Traffic flows (especially if your traffic is asymmetrical)
- Ego factor
- Keep me in a job

Planning

- Estimates of amount of peering traffic
  - At least 20G
- Region
  - US
  - Europe
  - Split roughly 50/50
- Target likely peers
  - Content providers
  - Peer to peer (other eyeballs)
- Locations

Network Diagram (Europe)

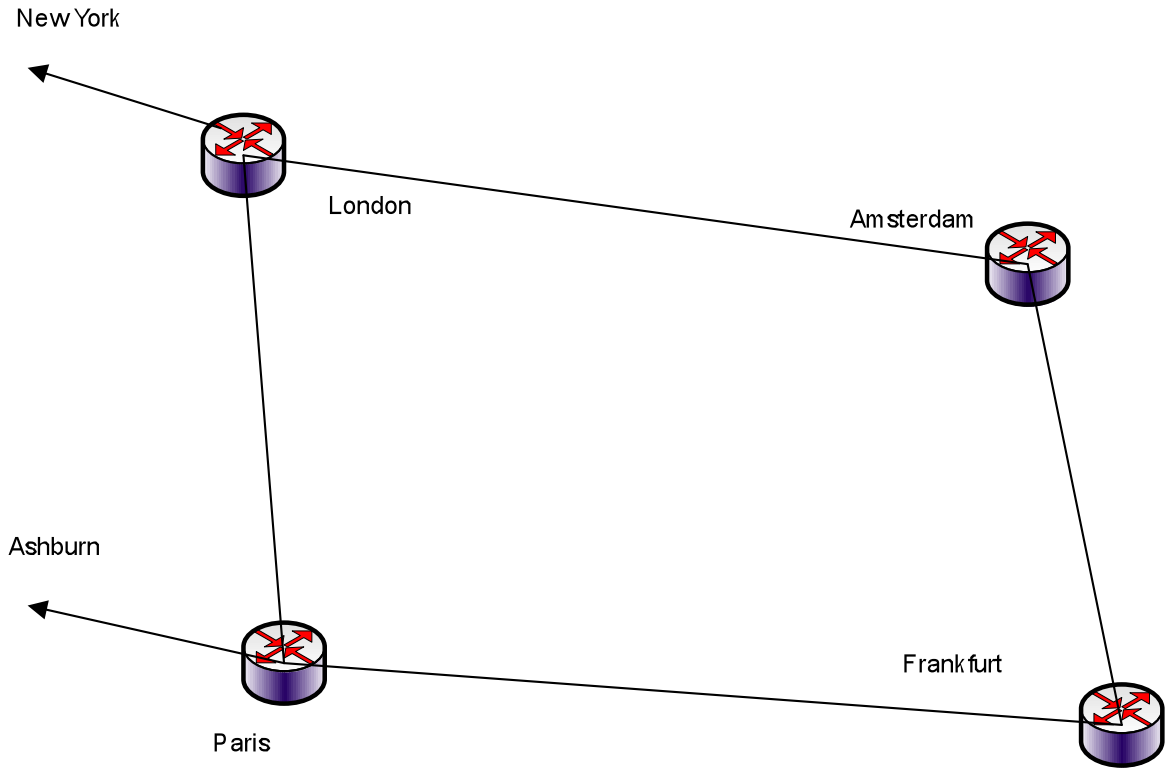

Locations (Europe)

- Fairly Easy as we already had a network in Europe
- Based around IXPs
  - Amsterdam (AMSIX)
  - Frankfurt (DECIX)
  - Paris (SFINX/PANAP/FreeIX)
- Initial Build
  - 10G ring London – Amsterdam – Paris – Frankfurt
  - Cisco CRSes (room to grow)
  - Optics as we need them

Network Diagram (US)

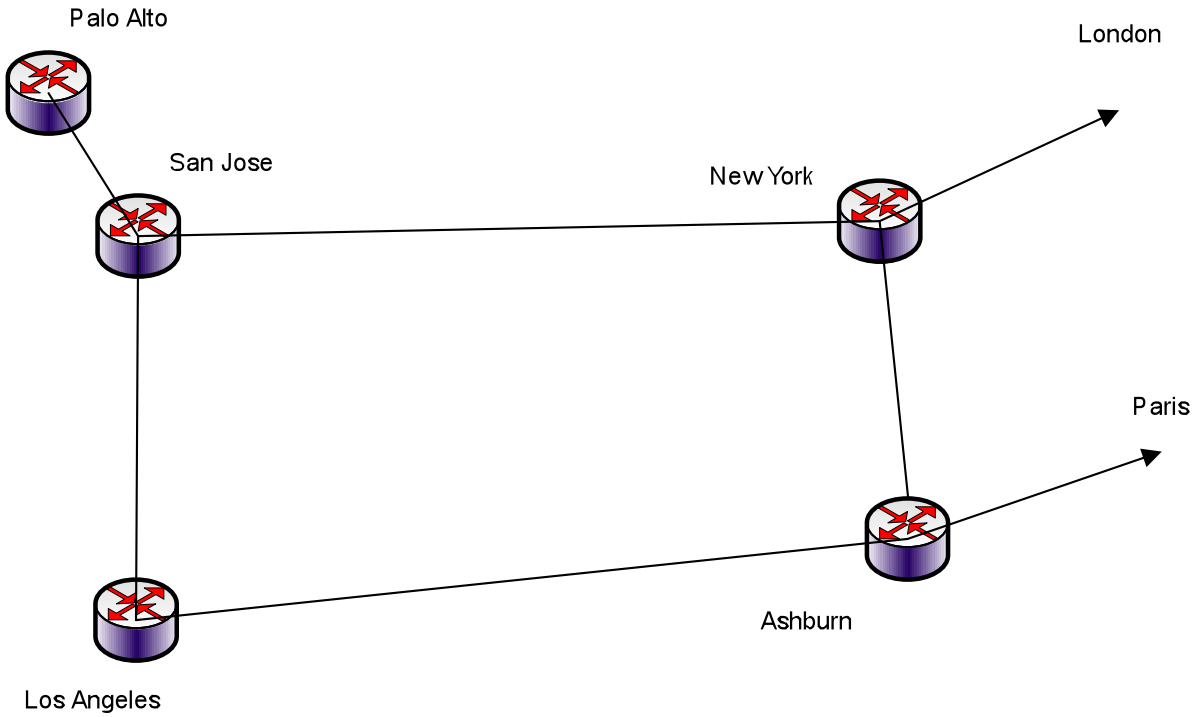

Locations (US)

- Brand New Network
- Based around IXPs
  - New York (NYIIX, PAIX)
  - Ashburn (Equinix)
  - Los Angeles (Any2, LAIIX)
  - Palo Alto (PAIX)
  - San Jose (Equinix)
- Initial Build
  - 10G ring New York – Ashburn – Los Angeles – San Jose
  - 10G spurs to Palo Alto
  - Cisco CRSes (room to grow)
  - Optics as we need them

US traffic growth

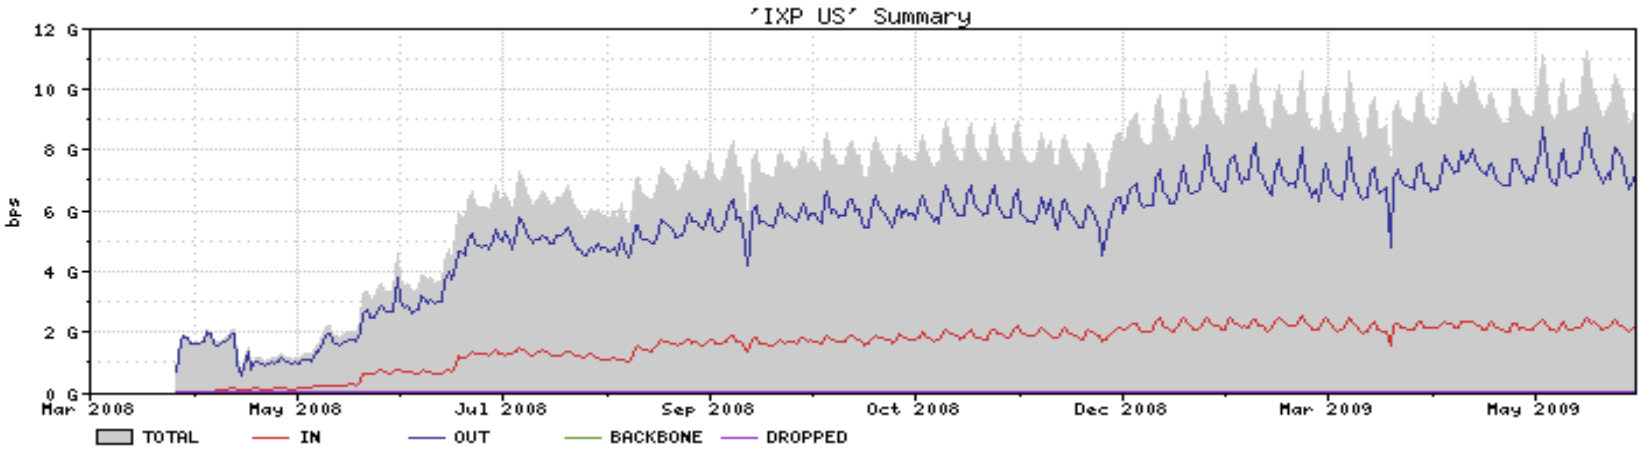

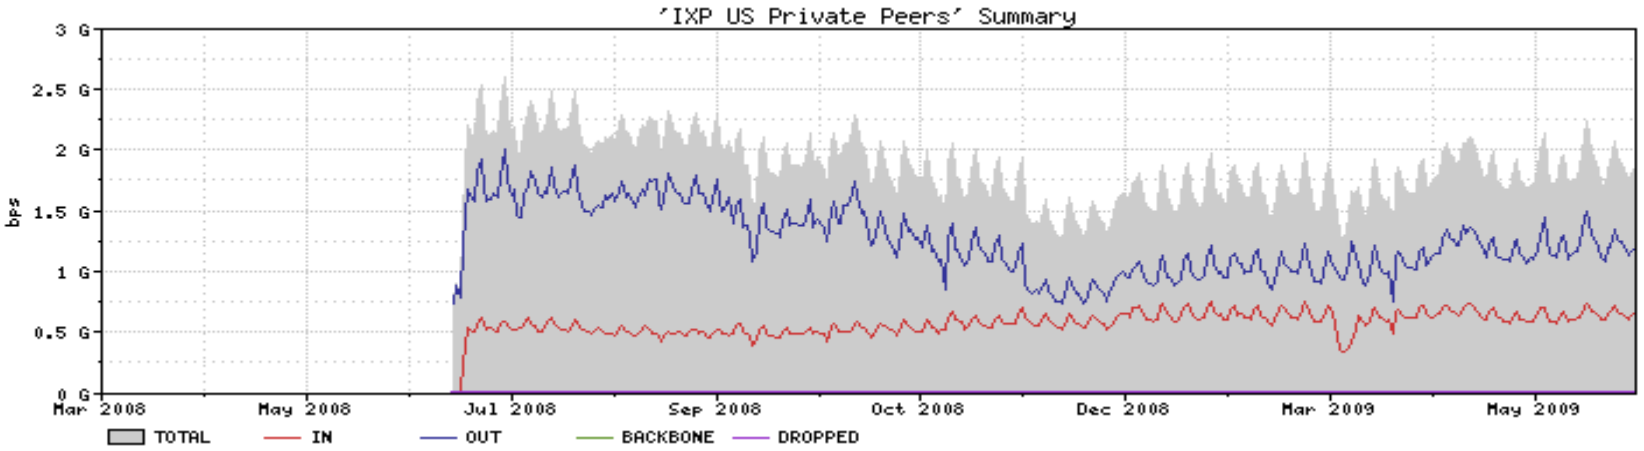

Europe Traffic growth

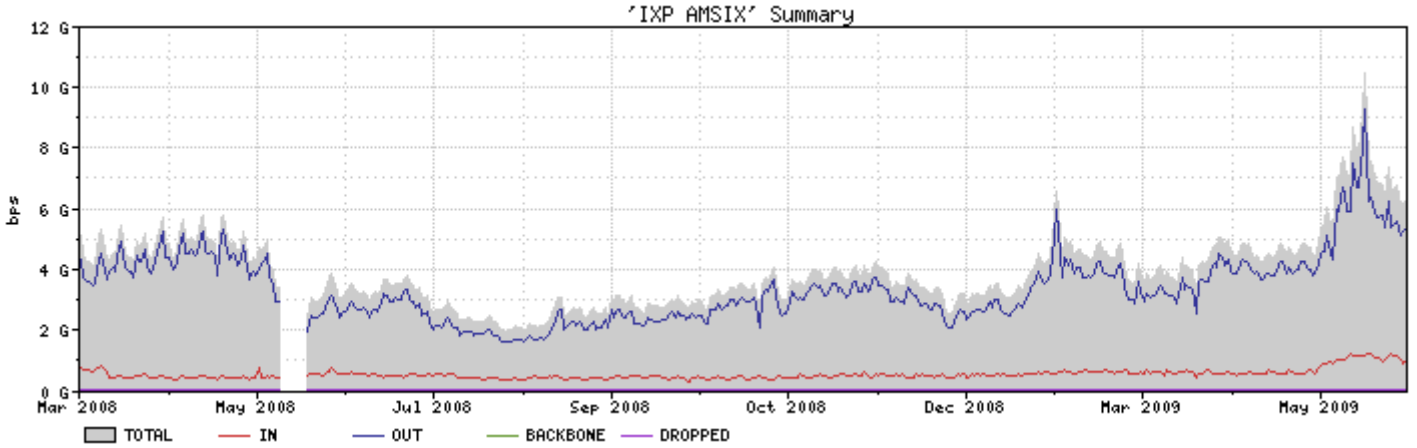

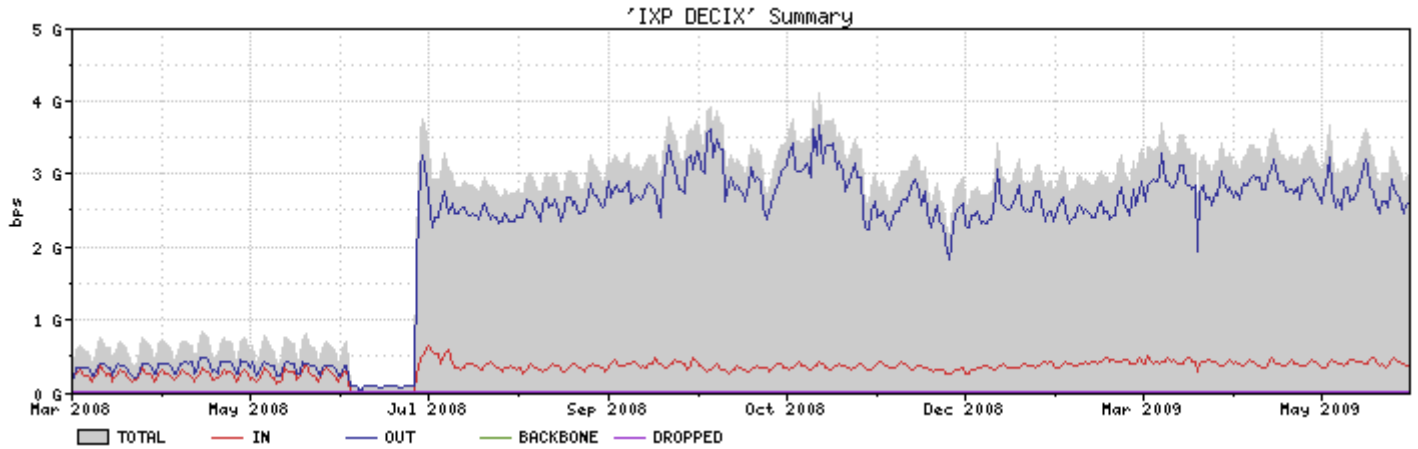

European Traffic growth (2)

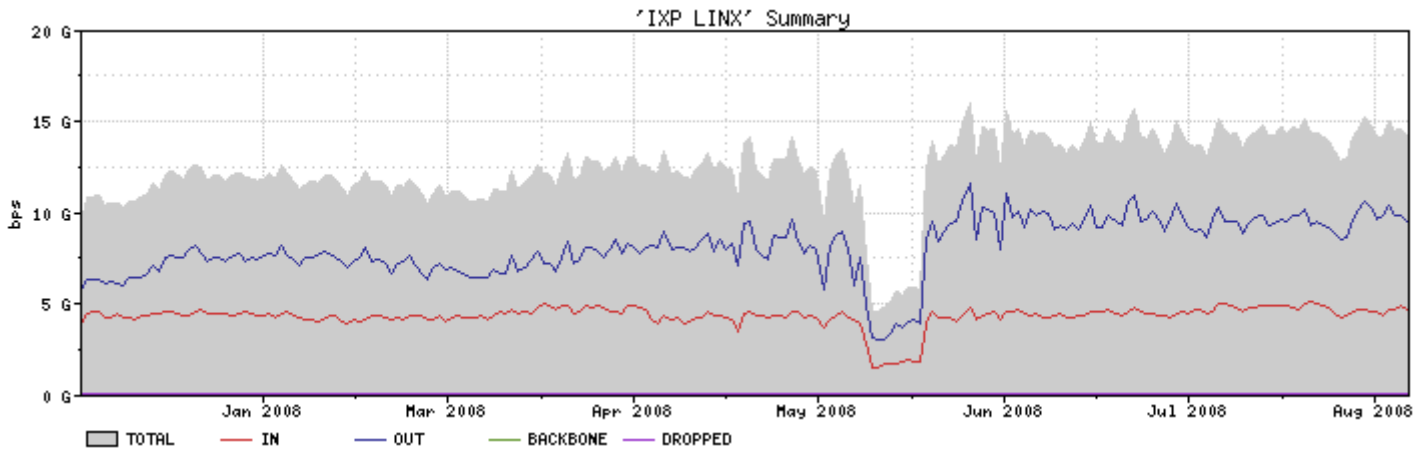

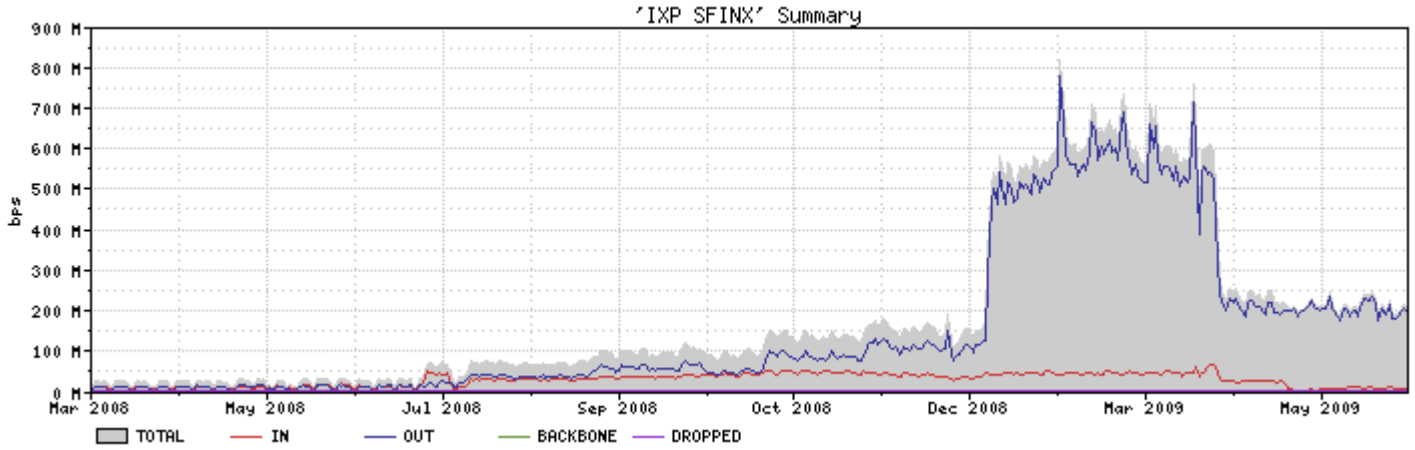

Effect on transit

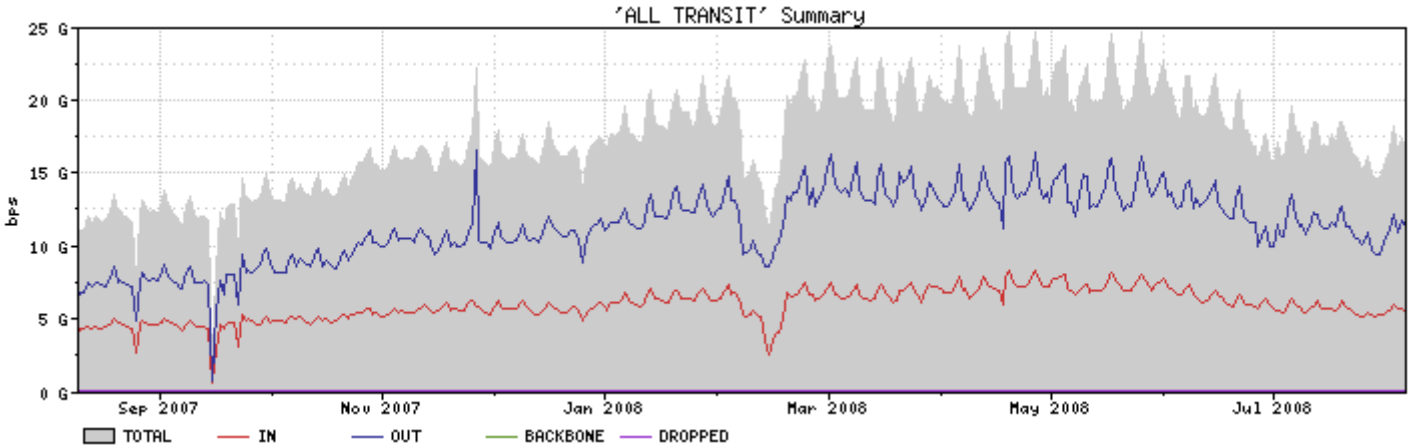

Savings

- Transit – Peering = 30 – 50%
- Total traffic = 140G approx
- Peering ratio approx 70%
- Latency/jitter improved
- Network available for reverse traffic FOC

How to use the New Network

- Analyse traffic (probably already done as part of your planning phase)
- Await requests to peer
- Stalk and hunt desired peers
- Bring up peerings

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

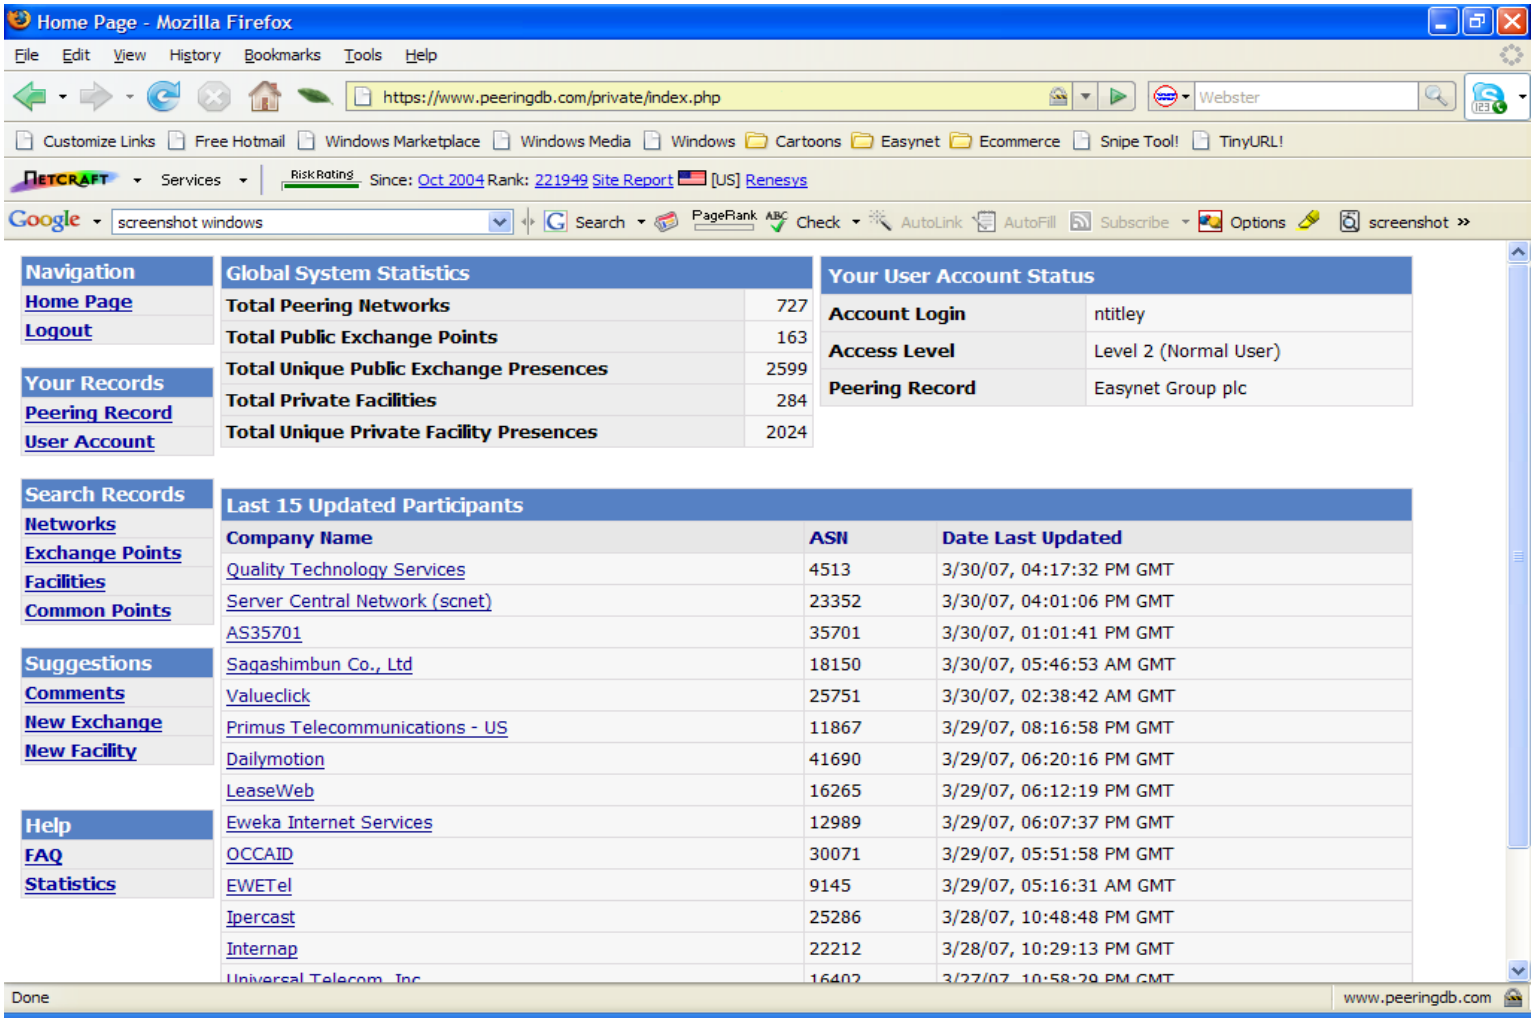
Home Page - Mozilla Firefox
File Edit View History Bookmarks Tools Help
https://www.peeringdb.com/private/index.php
Webster
Customize Links Free Hotmail Windows Marketplace Windows Media Windows Cartoons Easynet Ecommerce Snipe Tool! TinyURL!
NETCRAFT Services Risk Rating Since: Oct 2004 Rank: 221949 Site Report [US] Renesys
Google screenshot windows Search PageRank Check AutoLink AutoFill Subscribe Options screenshot
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
Sagashimbun 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
Done
www.peeringdb.com

PeeringDB (search)

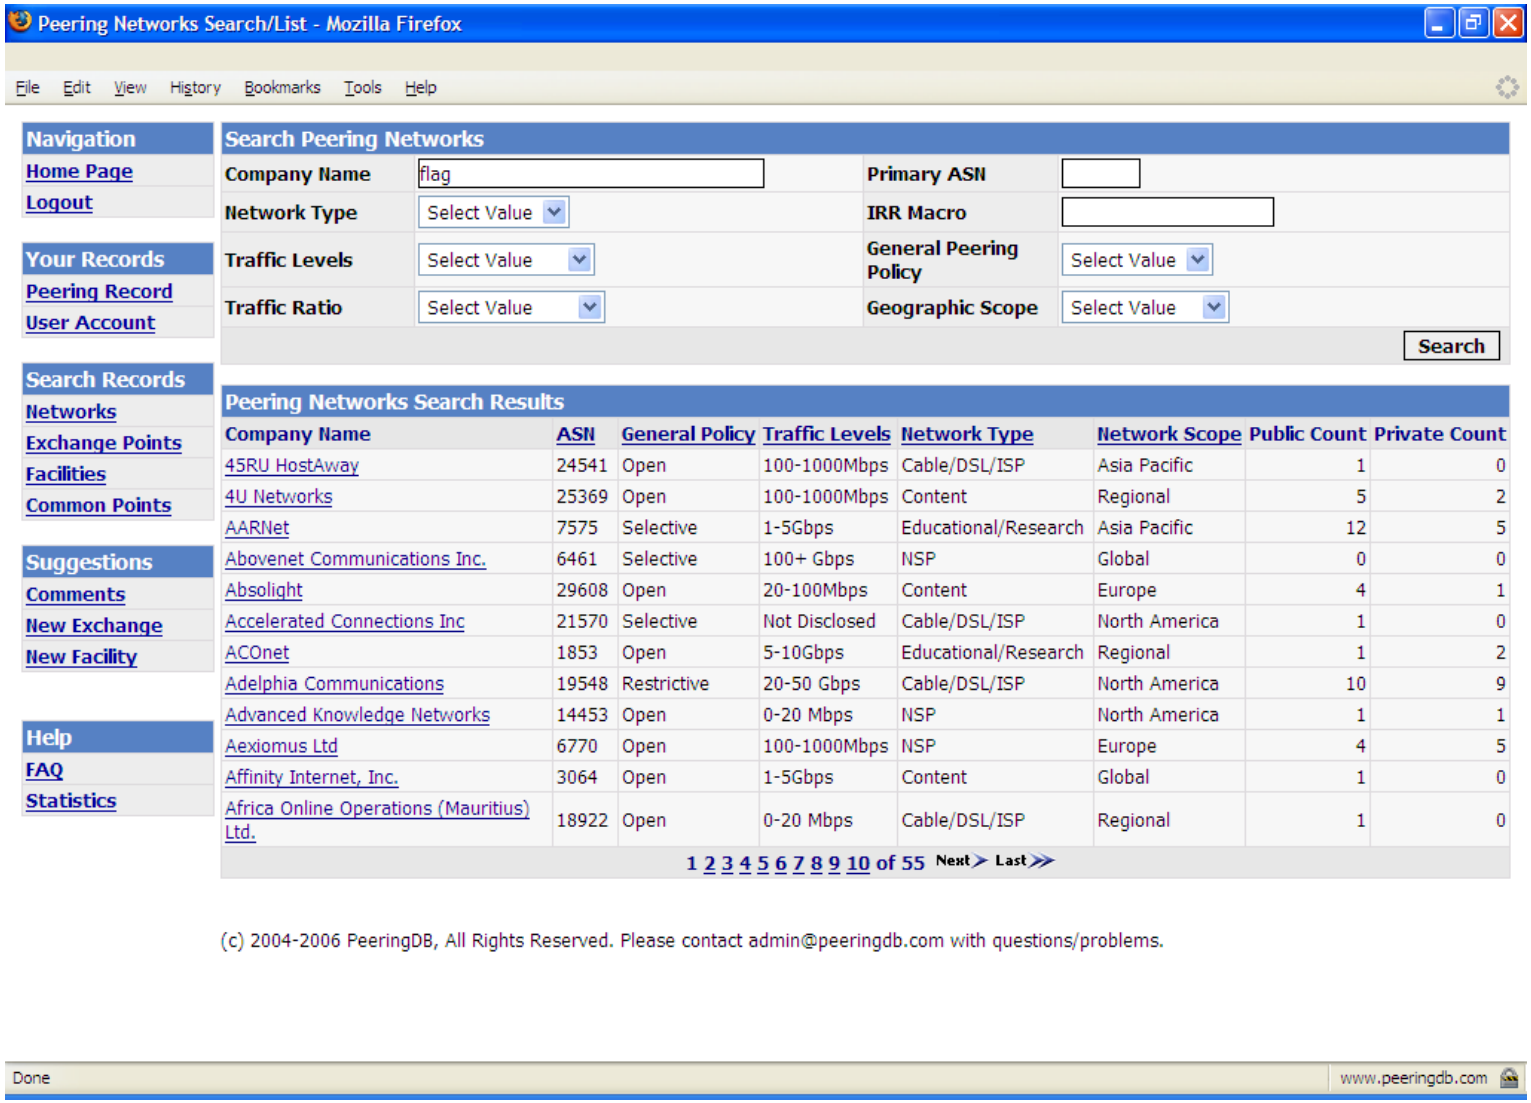

PeeringDB (search results)

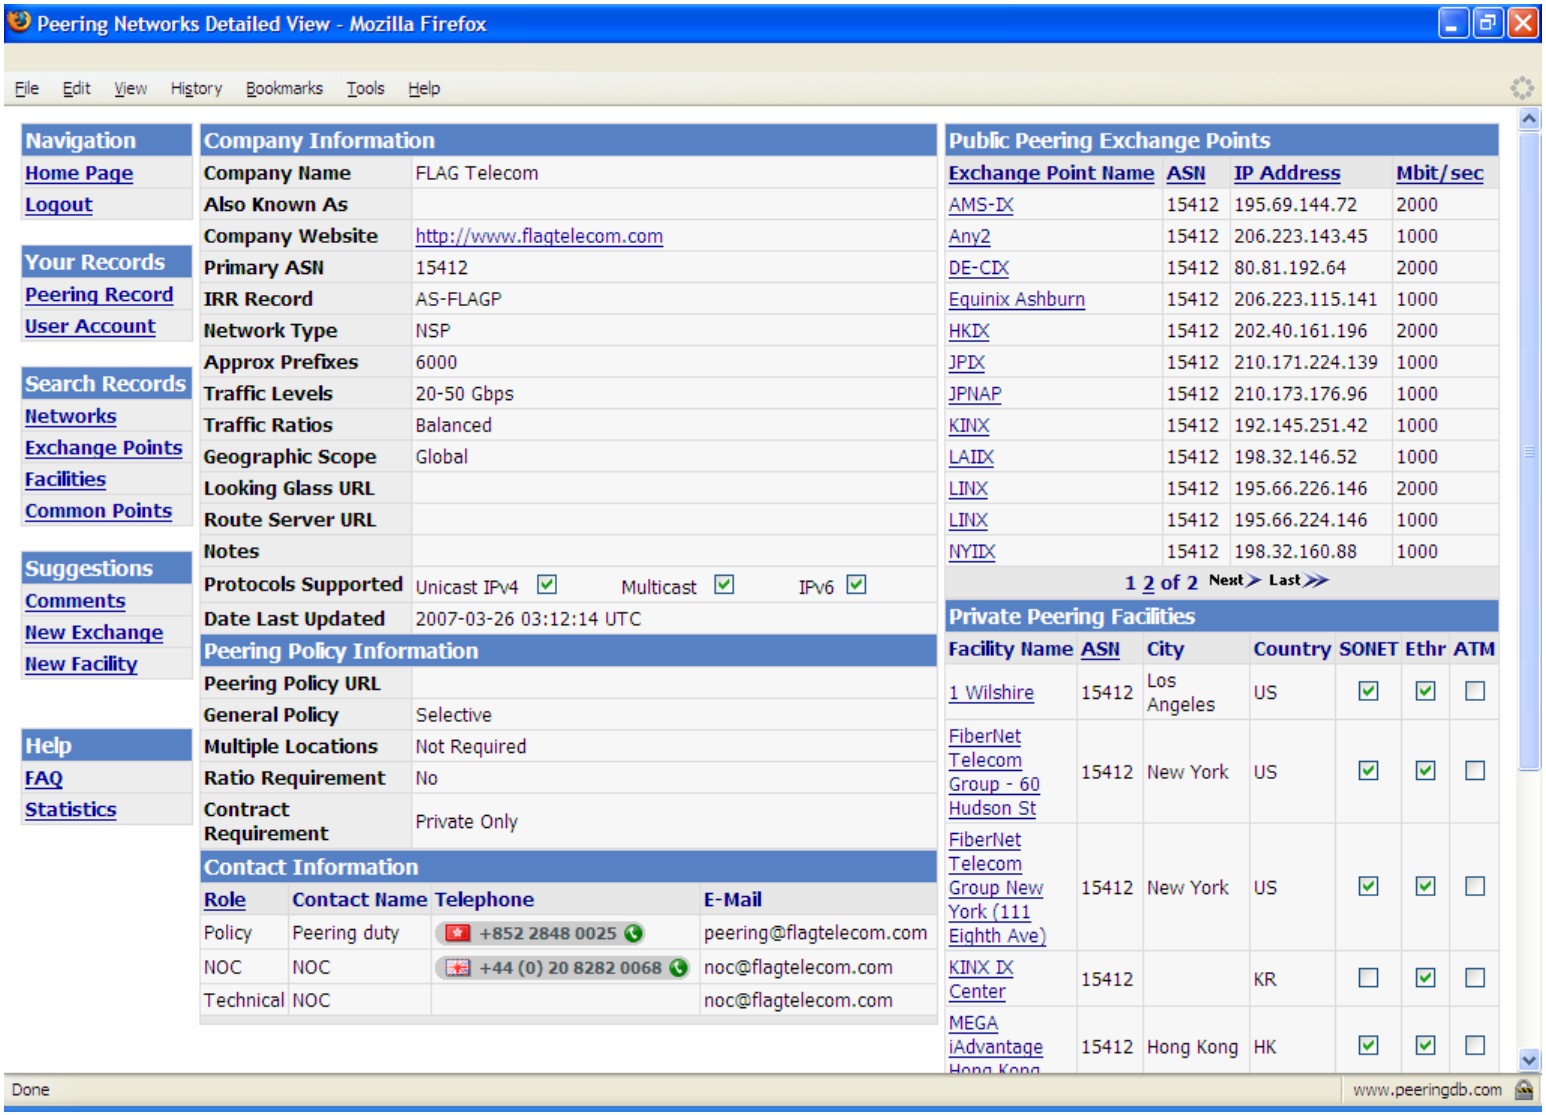

Arbor Peakflow (or similar)

- 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)

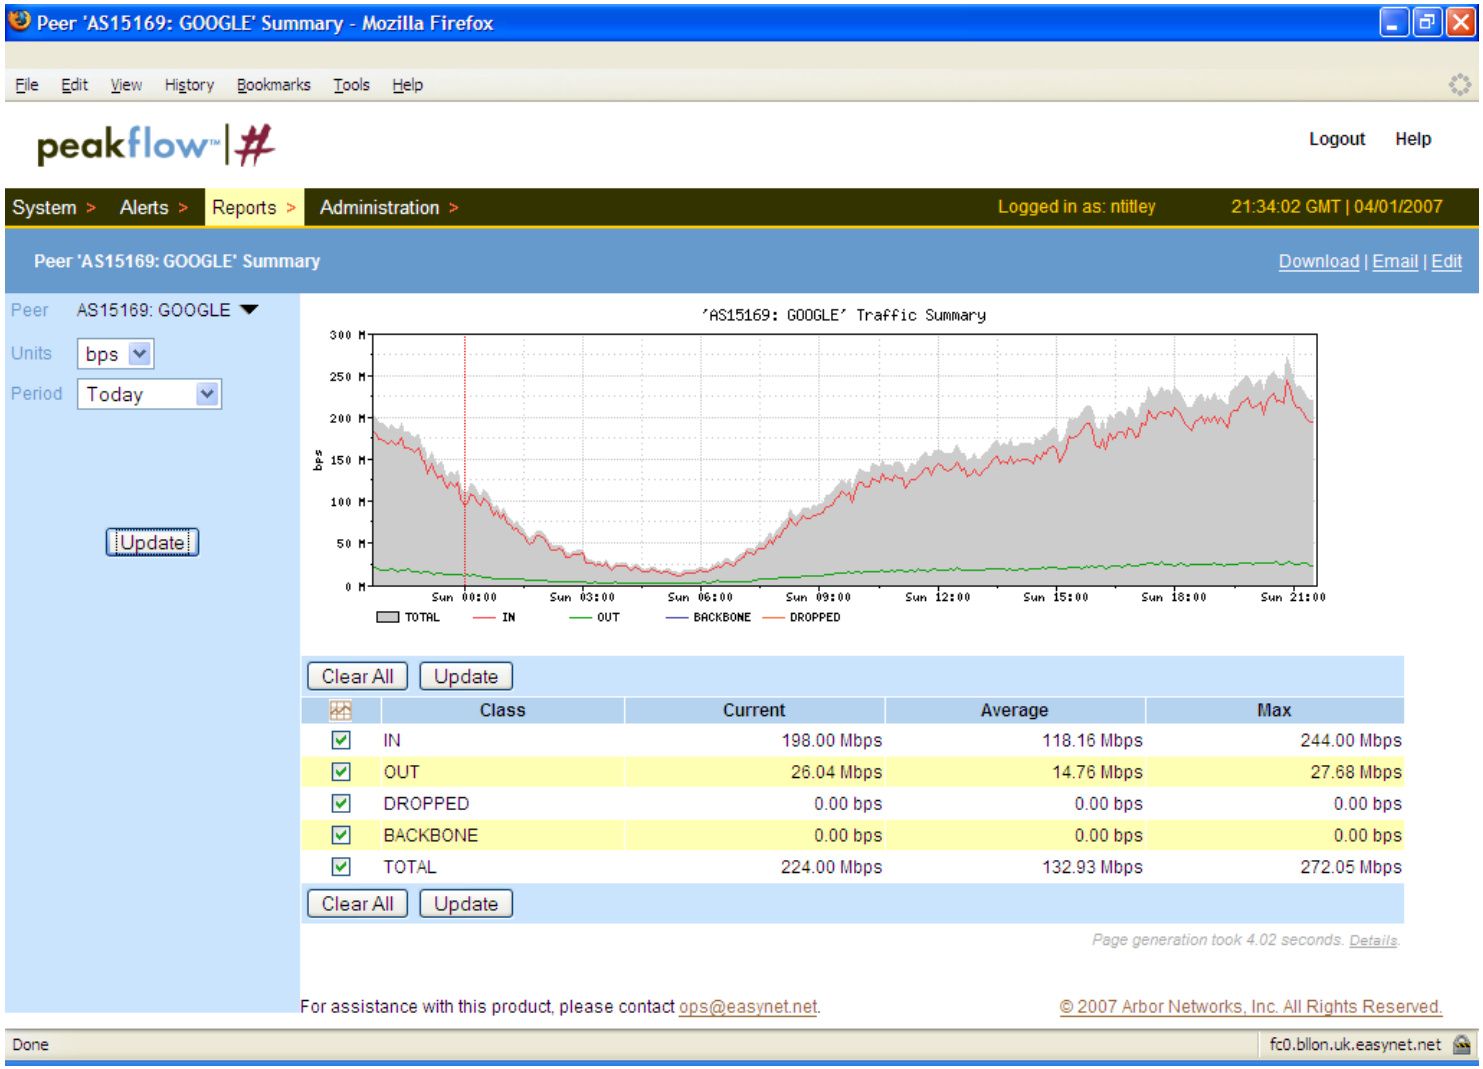

Renesys

- 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
  - Get access to complete history

AS5564 (Scotland Online)

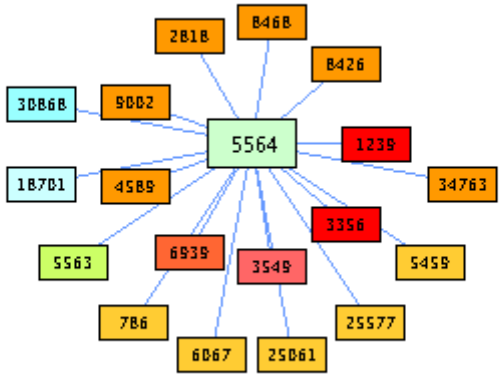

AS4589 (Easynet)

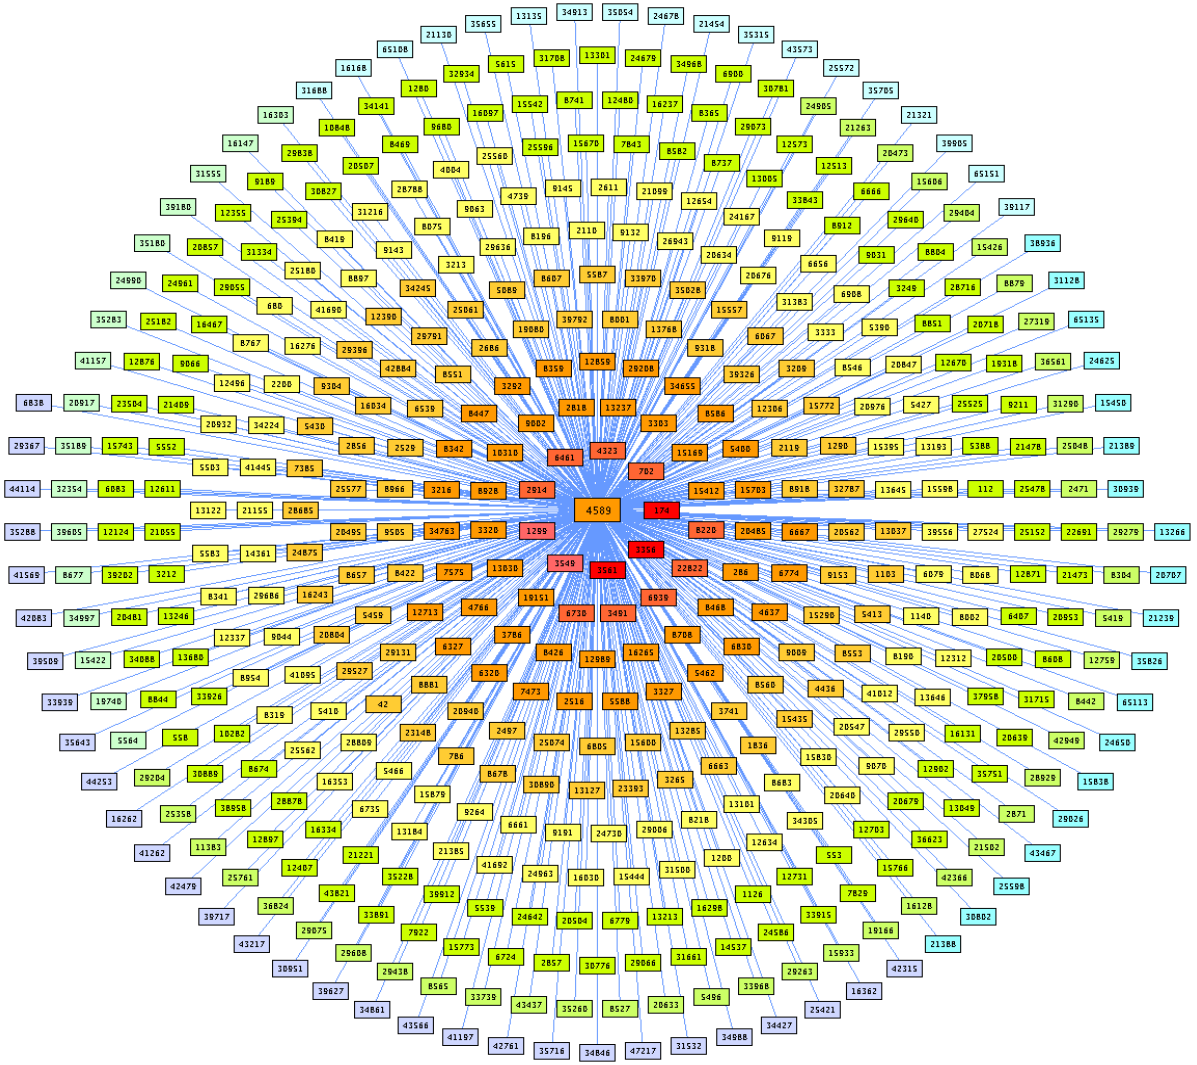

Local Database

- 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

IXP Overview

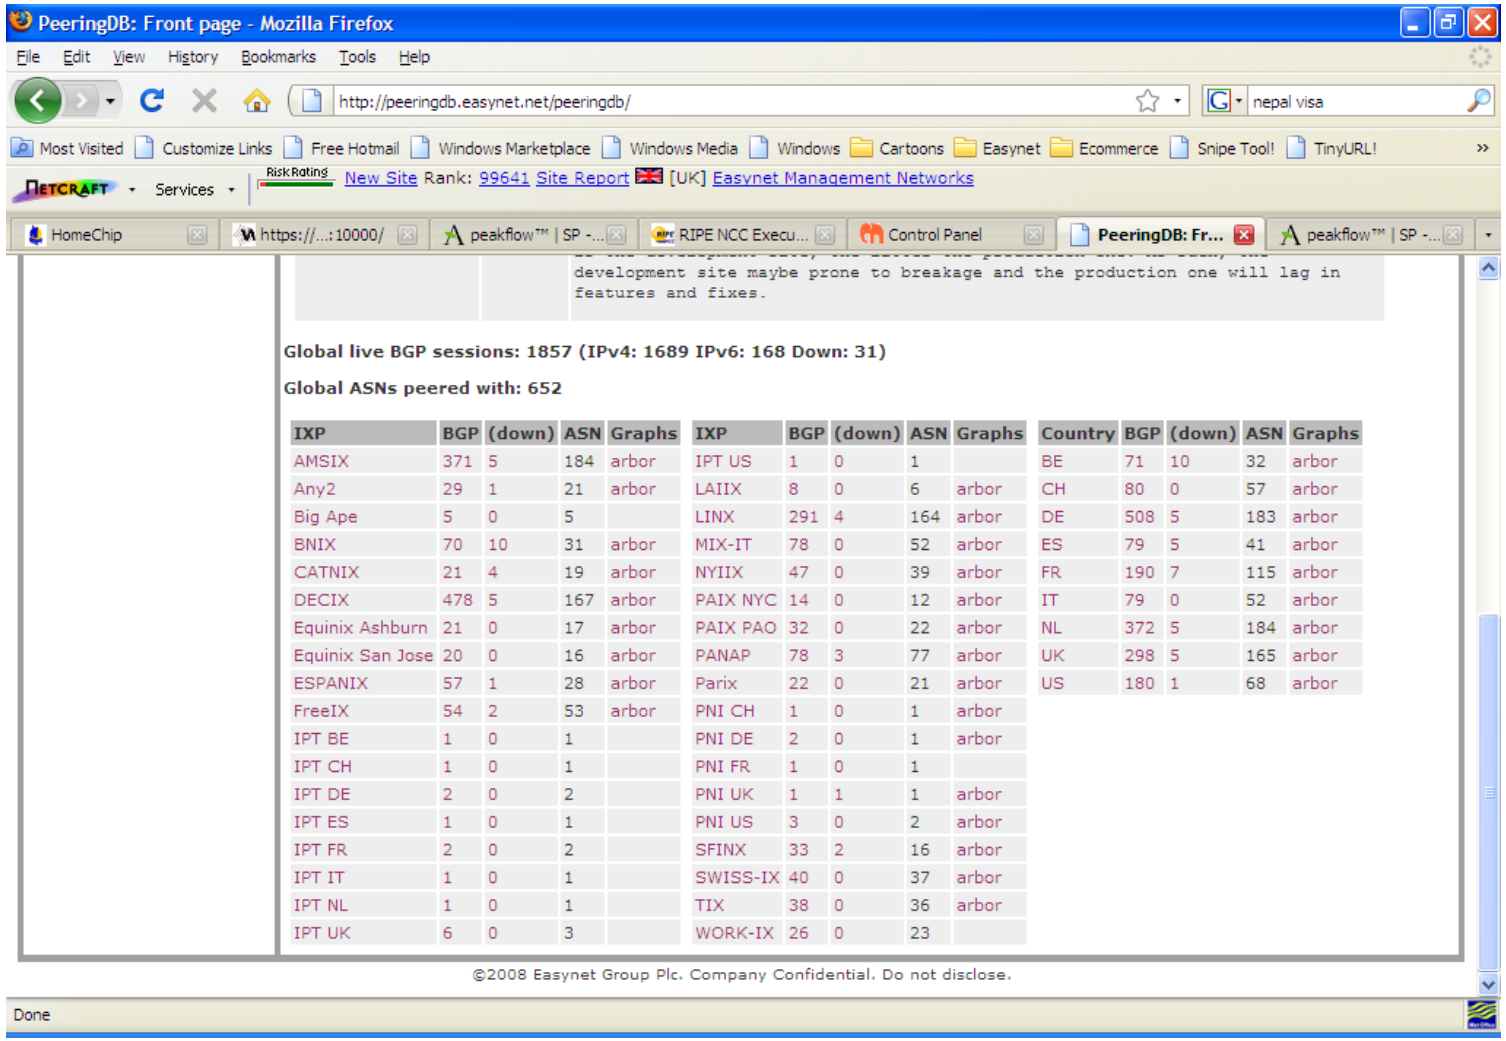

IXP	BGP	(down)	ASN	Graphs
AMSIX	371	5	184	arbor
Any2	29	1	21	arbor
Big Ape	5	0	5	
BNIX	70	10	31	arbor
CATNIX	21	4	19	arbor
DECIX	478	5	167	arbor
Equinix Ashburn	21	0	17	arbor
Equinix San Jose	20	0	16	arbor
ESPANIX	57	1	28	arbor
FreeIX	54	2	53	arbor
IPT BE	1	0	1	
IPT CH	1	0	1	
IPT DE	2	0	2	
IPT ES	1	0	1	
IPT FR	2	0	2	
IPT IT	1	0	1	
IPT NL	1	0	1	
IPT UK	6	0	3	

IXP	BGP	(down)	ASN	Graphs
IPT US	1	0	1	
LAIIX	8	0	6	arbor
LINX	291	4	164	arbor
MIX-IT	78	0	52	arbor
NYIIX	47	0	39	arbor
PAIX NYC	14	0	12	arbor
PAIX PAO	32	0	22	arbor
PANAP	78	3	77	arbor
Parix	22	0	21	arbor
PNI CH	1	0	1	arbor
PNI DE	2	0	1	arbor
PNI FR	1	0	1	
PNI UK	1	1	1	arbor
PNI US	3	0	2	arbor
SFINX	33	2	16	arbor
SWISS-IX	40	0	37	arbor
TIX	38	0	36	arbor
WORK-IX	26	0	23	

Country	BGP	(down)	ASN	Graphs
BE	71	10	32	arbor
CH	80	0	57	arbor
DE	508	5	183	arbor
ES	79	5	41	arbor
FR	190	7	115	arbor
IT	79	0	52	arbor
NL	372	5	184	arbor
UK	298	5	165	arbor
US	180	1	68	arbor

Peer Search

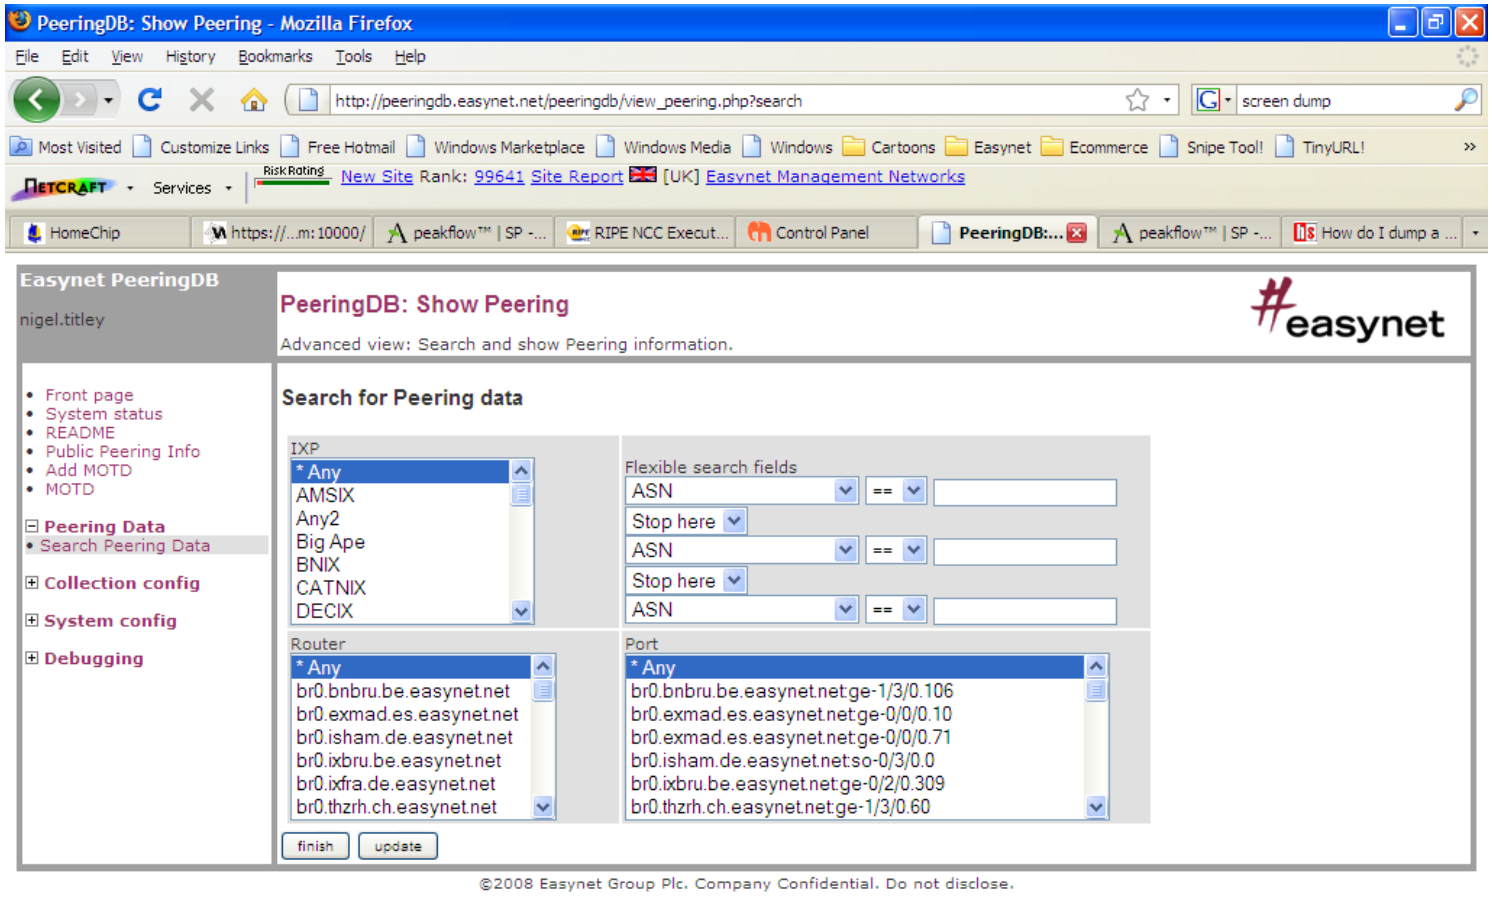
PeeringDB: Show Peering - Mozilla Firefox
File Edit View History Bookmarks Tools Help
http://peeringdb.easynet.net/peeringdb/view_peering.php?search
screen dump
Most Visited Customize Links Free Hotmail Windows Marketplace Windows Media Windows Cartoons Easynet Ecommerce Snipe Tool! TinyURL!
Services New Site Rank: 99641 Site Report [UK] Easynet Management Networks
HomeChip peakflow™ | SP -... RIPE NCC Execut... Control Panel PeeringDB:... peakflow™ | SP -... How do I dump a ...
Easynet PeeringDB
nigel.titley
PeeringDB: Show Peering
Advanced view: Search and show Peering information.
#easynet
Front page
System status
README
Public Peering Info
Add MOTD
MOTD
Peering Data
Search Peering Data
Collection config
System config
Debugging
Search for Peering data
IXP
* Any
AMSIX
Any2
Big Ape
BNIX
CATNIX
DECIX
Flexible search fields
ASN == Stop here
ASN == Stop here
ASN ==
Router
* Any
br0.bnbru.be.easynet.net
br0.exmad.es.easynet.net
br0.isham.de.easynet.net
br0.ixbru.be.easynet.net
br0.ixfra.de.easynet.net
br0.thzrh.ch.easynet.net
Port
* Any
br0.bnbru.be.easynet.net:ge-1/3/0.106
br0.exmad.es.easynet.net:ge-0/0/0.10
br0.exmad.es.easynet.net:ge-0/0/0.71
br0.isham.de.easynet.net:so-0/3/0.0
br0.ixbru.be.easynet.net:ge-0/2/0.309
br0.thzrh.ch.easynet.net:ge-1/3/0.60
finish update
©2008 Easynet Group Plc. Company Confidential. Do not disclose.

Peer details

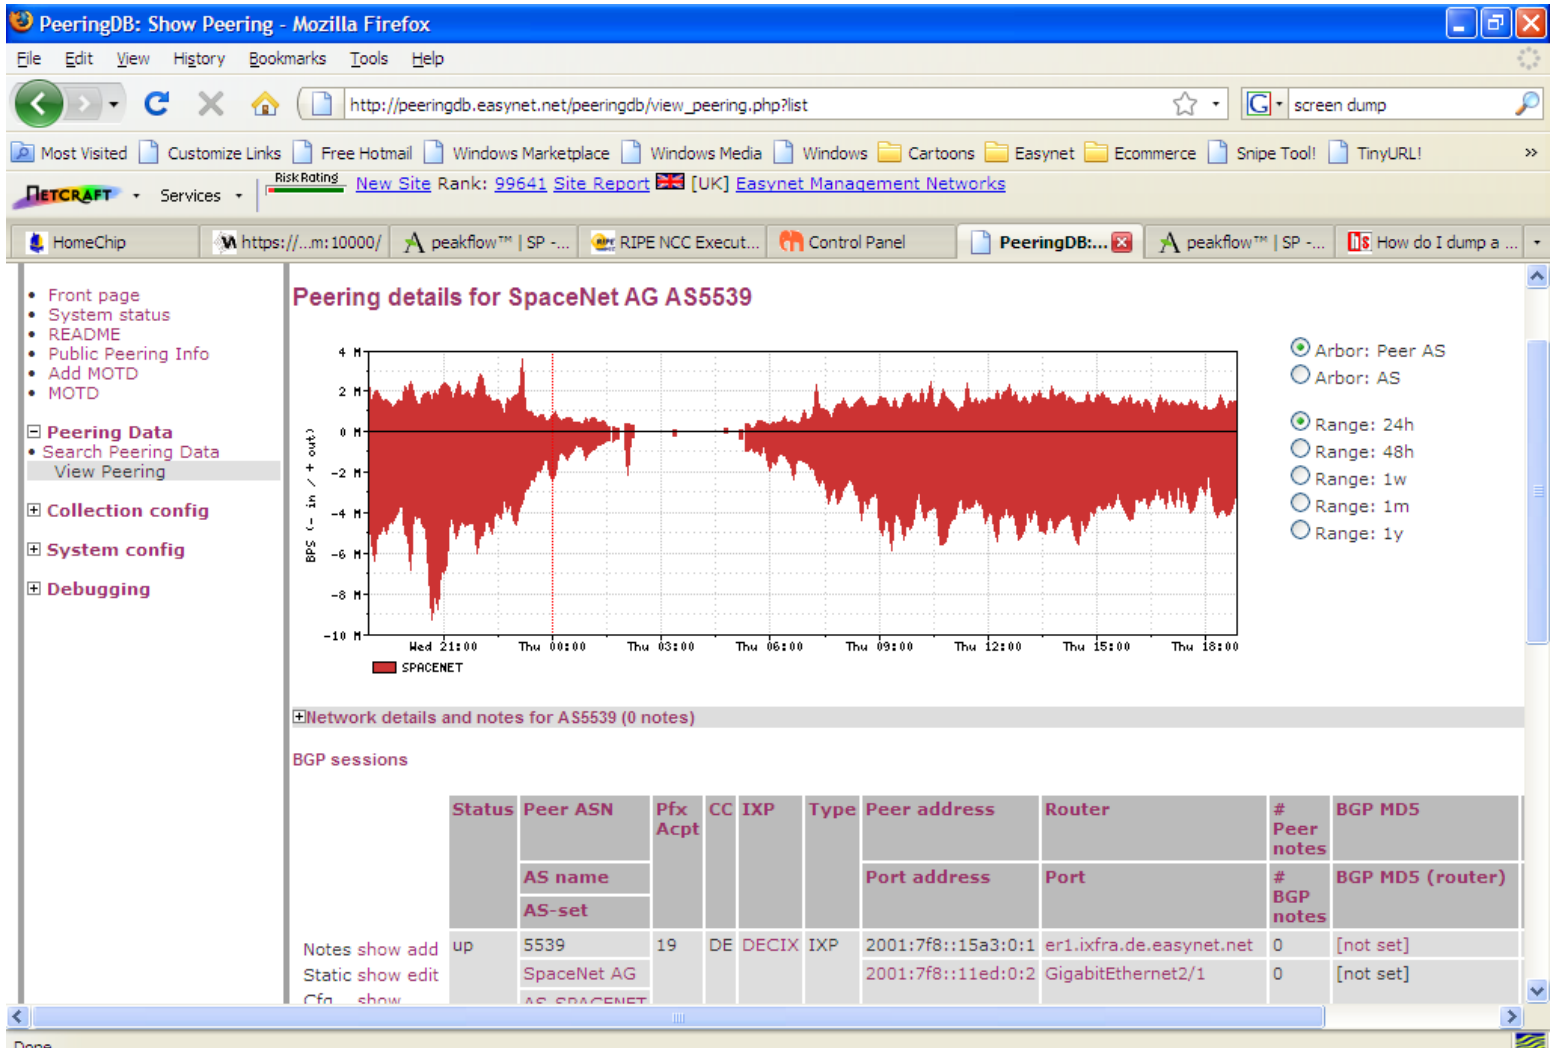

Friendly and Engaging Smile

- It helps to be friendly, but I seem to have managed to get away with it.
- Buy drinks
- Bring tee-shirts
- Go to meetings
  - RIPE
  - NANOG
  - APRICOT
  - SANOG
  - MENOG
  - UKNOF
- Sponsor SANOG meetings

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

Questions?

nigel.titley@uk.easynet.net