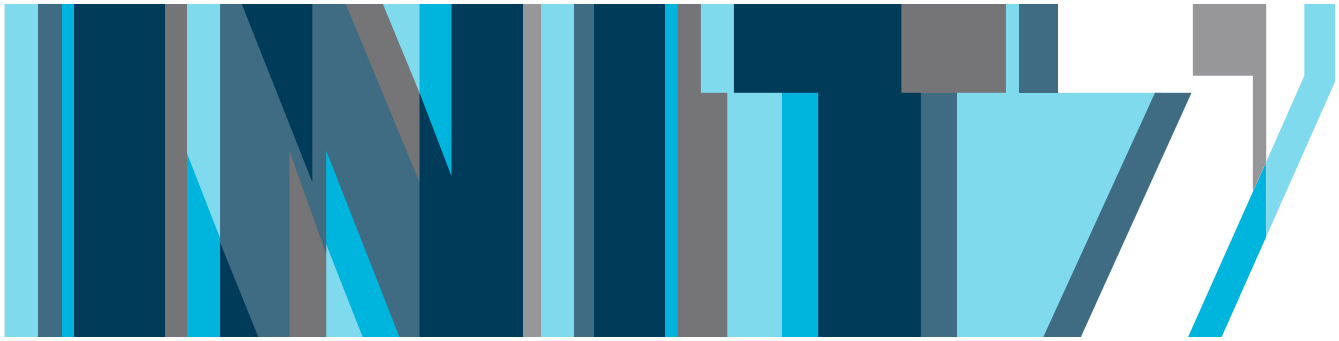

Traffic Engineering with Excel: Commercial Aspects of IP Capacity

MENOG #13 – Kuwait City – September 22/23, 2013.

DISCLAIMER

These slides show experience examples of the Init7 / AS13030 backbone over various years. They may work or may not work for you. Please use the methods described with care and at your own risk. Init7 or the author cannot be held responsible for any damage occurred by using the methods described here.

Your Speaker Today...

Fredy Künzler

kuenzler@init7.net
www.init7.net
www.blogg.ch
www.bgp-and-beyond.com
AS13030

Twitter: @kuenzler / @init7

Init7 (Switzerland) AG
St.-Georgen-Strasse 70
CH-8400 Winterthur
Switzerland

phone: +41 44 315 44 00

The Company Init7

Init7 operates its own backbone with the AS number 13030

Facts

- Init7 operates an international n*10Gbit IP backbone with AS number 13030 (autonomous system)
- The AS13030 is located at around 20 Internet exchange points, where nearly 1000 further networks are so-called peering partners
- This means that direct interconnections exist with all these other networks, enabling direct contact with approx. 60% of the global routing table
- As a result, we can secure optimal connectivity, latency, capacity and availability
- The remaining approx 40% of the targets are connected via globally distributed upstreams

Advantages

- Full control over the quality of IP Transit Services
- Autonomy from suppliers

Init7 Backbone Europe

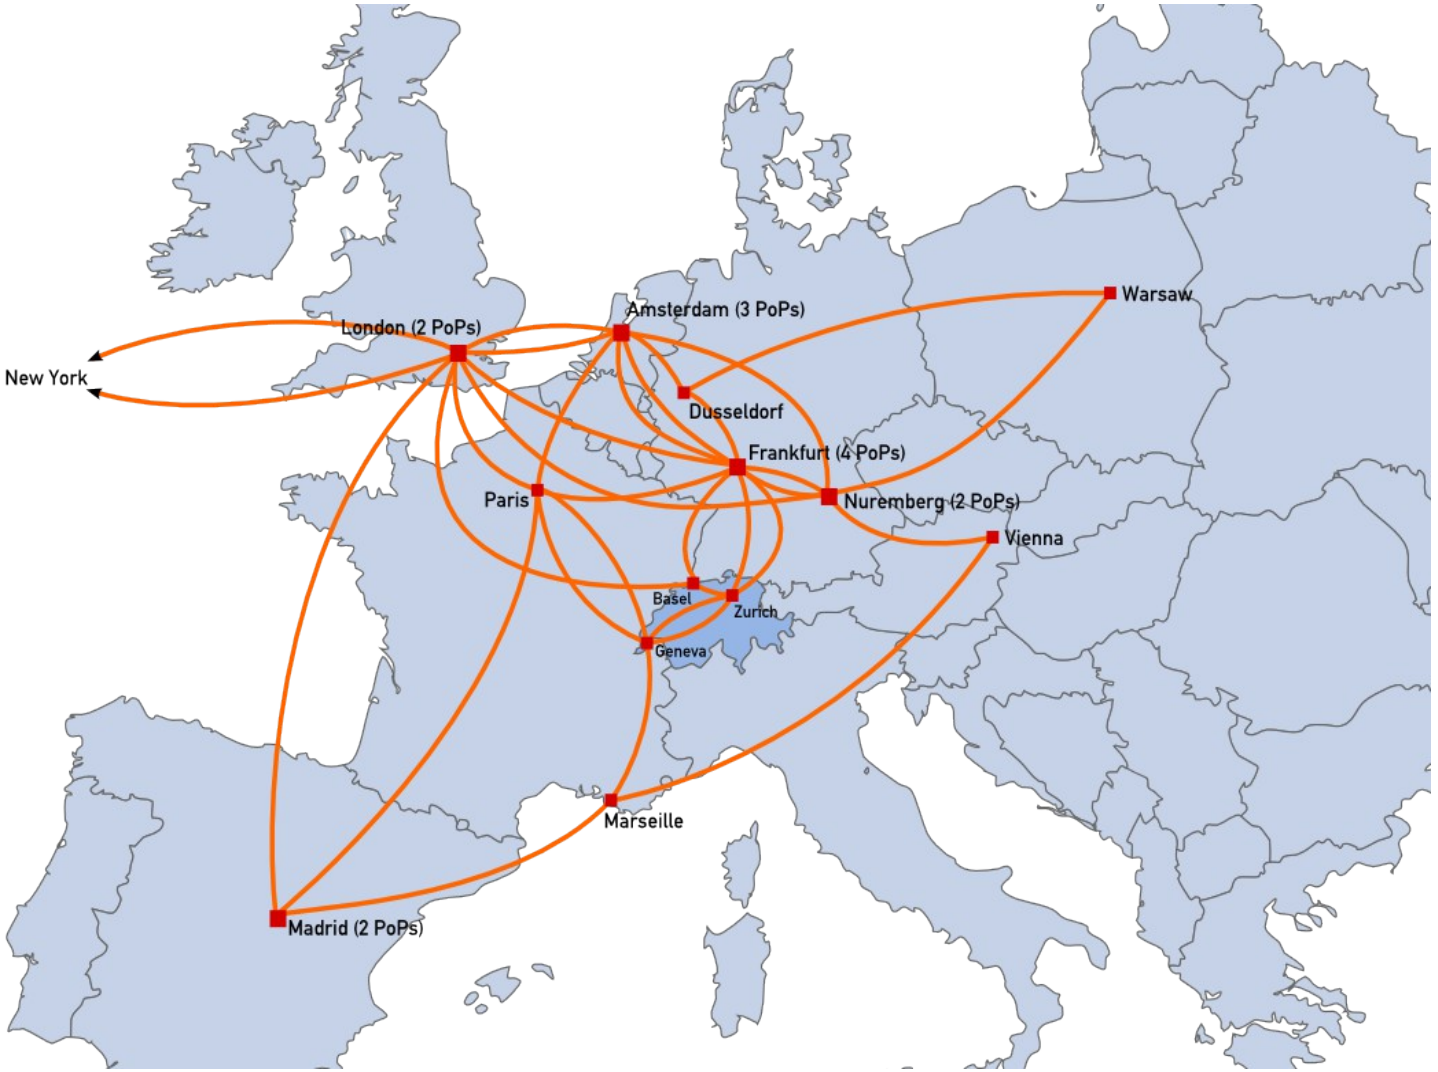

IP capacity buyers just choose by...

→PRICE!

Myth #1

- There is nothing but Transit and Public Peering

… no:
- Paid Peering
- Content Connect
- Partial Transit
- European Transit

Myth #2

- a TIER-1 is the preferred supplier. TIER-2 is second class.

  … it depends what you want to achieve.

  - usually a combination of TIER-1 and TIER-2 shows the best result
  - TIER-2 have a strategic advantage over TIER-1 and can produce IP capacity more efficient, which usually results in better pricing

Myth #3

- Public Peering reduces Transit cost
  (if you believe the marketing bla-bla of the big exchanges)

  … it depends...
  - on the amount of traffic (scales only from 10Gbps+)
  - on the skills of the Network engineering crew (to maintain 100+ peers is hard work)
  - transit is so cheap that peering often doesn't pay off anymore

The goal: sufficient capacity for end users

- **Packet loss (people / packets fall off the bus)**
- **Journey is more comfortable on less packed buses**

… but we do need to maintain the budget

- We could build the best user experience with not congested links but it would cost a fortune
- Required capacity is constantly growing, and the CFO wants us to produce more bandwidth with the same budget
- The average usage per user is rising (but the ARPU remains the same!):
  - 200kBit/s last year
  - 400kBit/s this year
  - 1MBit/s in two, three years

The Example Challenge #1

- Average Middle East ISP (Eyeball Provider) Traffic Ratio 8:1

- Year 2013:
  - 100000 Broadband Customers
  - average bandwidth per user: 400kbps
- total traffic: 40Gbps
  - purchasing from 4 TIER-1 carrier, each carrier supplies 10gig capacity in Europe
  - subsequently 4 STM-64 are purchased from various submarine cable operators (for redundancy)

Excel Calculation 2013

- Average Middle East ISP (Eyeball Provider)

Amount	Item	Cost	Total
4	STM-64 Submarine	100000$	400000$
10Gbps	IP Transit Vendor C	10000$	10000$
10Gbps	IP Transit Vendor G	12000$	12000$
10Gpbs	IP Transit Vendor L	15000$	15000$
10Gbps	IP Transit Vendor T	13000$	13000$
		Total	450000$
	Broadband Subscribers	100000	
	Cost per User (average)	(400kBps)	4.50$

The Example Challenge #2

- Average Middle East ISP (Eyeball Provider) Traffic Ratio 7:1 (note the change – outbound increases, but not significantly)

- Year 2014:
  - 150000 Broadband Customers
  - average bandwidth per user: 600kbps
  - total traffic: 90Gbps

- What should we purchase to supply the bandwidth demand? Expanding capacity the same way would be 225% of today's cost, while revenue is only 150% with the same ARPU.

Excel Calculation 2014

- User Growth: 50% / Usage Growth: 50%

Amount	Item	Cost	Total
9	STM-64 Submarine	80000$ -20%	720000$
30Gbps	IP Transit Vendor C	8000$ -20%	24000$
20Gbps	IP Transit Vendor G	10000$ -17%	20000$
20Gpbs	IP Transit Vendor L	12000$ -20%	24000$
20Gbps	IP Transit Vendor T	11000$ -15%	22000$
		Total	810000$
	Broadband Subscribers	150000	
	Cost per User (average)	(600kBps)	5.40$

Excel Calculation 2014

- User Growth: 50% / Usage Growth: 50%

Amount	Item	Cost	Total
9	STM-64 Submarine	80000$	720000$
30Gbps	IP Transit Vendor C	8000$	24000$
20Gbps	IP Transit Vendor G	10000$	20000$
20Gpbs	IP Transit Vendor L	12000$	24000$
20Gbps	IP Transit Vendor T	11000$	22000$
		Total	810000$
	Broadband Subscribers	150000	
	Cost per User (average)	(600kBps)	5.40$

NOT APPROVED!

… the budget is not approved!

- ARPU remains stable
- Average usage per User rises
- Average cost per user were $4.50 in 2013, but $5.40 in 2014 – no go!

… do the homework #1

- Buy cheaper? …we already talked down the IP transit suppliers –~ 20% rebate on renewals and upgrades
- STM-64 Submarine Capacity gets cheaper too, but demand increases faster –~ 20% rebate (sometimes IRU is an issue, too)
- Local peering? Yes, but it would substitute only 5 to 10% of the transit
- Content Cache Servers locally? Yes, definitely!

… do the homework #2

- Rethink buying strategy!
- So far, you just bought „IP Transit“ (plus Submarine Transport)... but do you really need it?
- If you want to fly to London, you don't buy a ticket to New York and jump off the plane in London, right?
- A ticket to London is cheaper than a ticket to New York – so buy the right [ticket] | [product]!

… do the homework #3

- Airline Industry:
  Know where to go before buying a ticket!

- Internet Industry:
  Know your backbone traffic before purchasing IP Transit and Submarine Transport! This is, in fact, much more important than just shopping around for the best price...

Buying transit by price per Mbps might not end in a good result:

You need telemetry data #1

- To know the top-20 traffic sources of your backbone by percentage could make a huge difference in your capacity purchasing bill.
- Router gear of every major vendor supports Netflow or Sflow, which easily produces raw traffic statistics (so called flow stats).
- These flow stats can be aggregated and produce nice and easy readable traffic graphs by source ASN (Autonomous Systems).

You need telemetry data #2

AS-Stats – (free) Open Source tool to aggregate and display traffic flows by ASN

https://neon1.net/as-stats/

Analysis of Traffic #1

- Top 5 Traffic Sources / ASN are responsible for ~40% of the traffic
- Top 6 - 20 responsible for another ~40%
- Top 21 – 100 responsible for another ~10%
- Long tail... thousands of ASNs are just making noise and are irrelevant for the calculation

Analysis of Traffic #2

Some of the top 10 Traffic Sources (likely):

- Google/Youtube: AS15169
- Akamai: AS20940 (built capacity: ~15TBit)
- Microsoft: AS8075
- Leaseweb: AS16265 (built capacity: ~2TBit)
- Limelight: AS22822

Analysis of Traffic #3

Traffic sources by geographic scope (assumed):

- Middle East: 5%
- Europe: 65%
- United States: 20%
- Rest of the world: 10%

Analysis of Traffic #4

Surprised that the majority of traffic sources in Europe?

TeleGeography (Sept 18, 2013):
Europe Emerges as Global Internet Hub

„Why Europe? The continent benefits from IP transit prices that are among the lowest in the world, geographic proximity, and is home to many large carriers and major Internet exchanges that provide rich peering opportunities. Europe also benefits from a multitude of submarine cable landings, and new cable builds in Africa and the Middle East, specifically [...]"

http://www.telegeography.com/press/marketing-emails/2013/09/18/europe-emerges-as-global-internet-hub/index.html

Conclusion #1

We need to focus on the “fat boys“. How can we handle their traffic more efficient?

Two of the top 5 are rather easy:

- Google AS15169: deploy Google Cache Servers within your network in your own data center
- Akamai AS20940: AANP – similar to Google Cache, build local cache servers within your network

Conclusion #2

- Buying Full IP Transit from 4 vendors is a waste of money. Less than 30% really require full transit (20% US + 10% RoW) – remember: No-one is buying an airline ticket and jumps off the plane halfway.

- Ask vendors for quote for European Transit or Partial Transit... (this is not Paid Peering!).

Conclusion #3

- Peering locally / in the middle East: could possibly offload 5% of the transit traffic

- Peering in Europe: could be very effective, but has to be planned well, and don't underestimate the required long-haul capacity. Buying European Transit or Partial Transit from a well-peered network could be hassle free and result in the same effect for just a little more.

Excel Calculation 2014 - Revised

- User Growth: 50% / Usage Growth: 50%

Amount	Item	Cost	Total
7	STM-64 Submarine	80000$ -20%	560000$
20Gbps	IP Transit Vendor C	8000$ -20%	16000$
20Gbps	IP Transit Vendor G	10000$ -17%	20000$
30Gpbs	Partial European Transit	5000$	15000$
40Gbps	Akamai / Google Cache	4000$	4000$
10Gbps	Middle East Peering	3000$	3000$
=120Gbps!		Total	618000$
	Broadband Subscribers	150000	
	Cost per User (average)	(600kBps)	4.12$

Conclusion #4

- Despite that the average usage per user went up 50% (from 400kbps to 600kbps), the average cost per user went down from $4.50 to $4.12. The CFO loves you for this!
- Capacity is now over-provisioned – while the calculated demand is 90Gbps, the built capacity is 120Gbps – which results in happier end user!
- CAPEX / Labor cost are not considered in this examples.

Contact

Fredy Künzler

kuenzler@init7.net
www.init7.net
www.blogg.ch
www.bgp-and-beyond.com
AS13030

Twitter: @init7

Init7 (Switzerland) AG
St.-Georgenstrasse 70
CH-8400 Winterthur
Switzerland

phone: +41 44 315 44 00

COPYRIGHT