Interconnectivity in South Asia

APRICOT 2018 - Kathmandu, Nepal

Countries under consideration

- Afghanistan
- Bangladesh
- Bhutan
- India
- Maldives
- Nepal
- Pakistan
- Sri Lanka

How South Asia is connected the world?

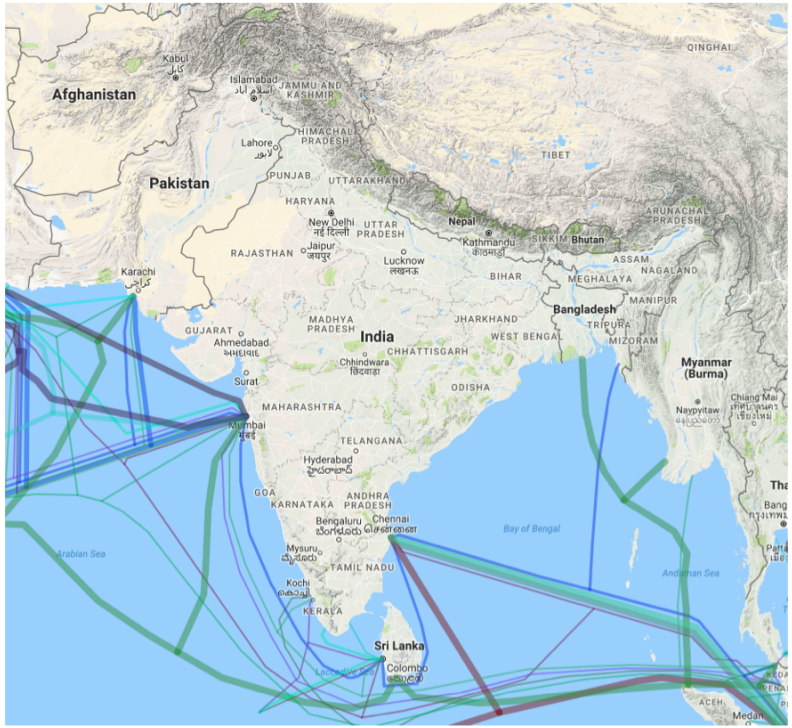

* Source: CableMap.info

Submarine cable landing in South Asia

1. Karachi, Pakistan (6 cables) - SEA-ME-WE-4, SEA-ME-WE 3, I-ME-WE, Transworld (TWA-1), SEA-ME-WE-5, AAE-1
2. Mumbai, India (12 cables) - SEACOM, SEA-ME-WE-4, SEA-ME-WE 3, EIG, I-ME-WE, FLAG Europe Asia (FEA), FLAG Alcatel-Lucent Optical Network (FALCON), Gulf Bridge International, BBG, SEA-ME-WE-5, MENA, AAE-1
3. Kochi, India (2 cables) - SEA-ME-WE 3, WARF
4. Thiruvananthapuram, India (1 cable) - FALCON
5. Tuticorin, India (1 cable) - Bharat Lanka Cable
6. Chennai, India (5 cables) - i2i, SEA-ME-WE-4, Tata Indicom (TIISCS), BRICS Cable, BBG
7. Kuatka, Bangladesh (1 cable) - SEA-ME-WE-5
8. Cox Bazaar, Bangladesh (1 cable) - SEA-ME-WE-4
9. Ngwesaung, Mayanmar (1 cable) - SEA-ME-WE-5

Submarine cable landing in South Asia (cont.)

1. Ngwesaung, Mayanmar (1 cable) - SEA-ME-WE-5
2. Colombo, Sri Lanka (5 cables) - SEA-ME-WE-4, SLT-Dhiraagu, FALCON, WARF, BBG
3. Mount Lavinia, Sri Lanka (1 cable) - SEA-ME-WE 3, Bharat Lanka Cable System
4. Male, Maldives (3 cables) - SLT-Dhiraagu, FALCON, WARF

But how well is L3 connectivity?

Ways to analyse L3 connectivity

1. Traceroutes from NLNOG ring nodes in the region

2. Traceroutes from active RIPE Atlas in the region

NLNOG Ring - Connectivity test

From	To	IPv4 Path	IPv6 Path
digitalocean05.ring.nlnog.net	dhiraagu01.ring.nlnog.net	Local / Direct	Goes via Europe!
amazon12.ring.nlnog.net	dhiraagu01.ring.nlnog.net	Goes from outside region	Goes from outside region
anuragbhatia01.ring.nlnog.net	dhiraagu01.ring.nlnog.net	Goes via US!	Local / Direct
hostvirtual01.ring.nlnog.net	dhiraagu01.ring.nlnog.net	Goes via US!	Goes via Europe

NLNOG Ring - Connectivity test (cont.)

From	To	IPv4 Path	IPv6 Path
dhiraagu01.ring.nlnog.net	digitalocean05.ring.nlnog.net	Local / Direct	Goes via Europe!
dhiraagu01.ring.nlnog.net	amazon12.ring.nlnog.net	Goes via Europe	Goes from outside region
dhiraagu01.ring.nlnog.net	anuragbhatia01.ring.nlnog.net	Goes via Europe	Local / Direct
dhiraagu01.ring.nlnog.net	hostvirtual01.ring.nlnog.net	Local / Direct	Local / Direct

Connectivity test towards Atlas anchors

From/To	Afghanistan	Bangladesh	Bhutan	India	Maldives	Nepal	Pakistan	Sri Lanka
Afghanistan	-	-	-	-	-	-	-	-
Bangladesh	-	-	Indirect	Mostly Direct	Indirect	Mostly Direct	Indirect	-
Bhutan	-	Indirect	-	Mostly Indirect	Indirect	Mostly Indirect	Indirect	-
India	-	Mostly Direct	Mostly Direct	-	Indirect	Mostly Direct	Indirect	-
Maldives	-	Indirect	Indirect	Indirect	-	Indirect	Indirect	-
Nepal	-	Mostly Direct	Mostly Direct	Indirect	Indirect	-	Indirect	-
Pakistan	-	Indirect	Indirect	Indirect	Indirect	Indirect	-	-
Sri Lanka	-	Indirect	Indirect	Mostly Indirect	Mostly Indirect	Indirect	Indirect	-

Note: No probes & anchors in Afghanistan and no anchors in Sri Lanka

Interesting Patterns / Observations

1. Data shows indirect routing between Nepal to India while India to Nepal routing is mostly direct. This shows inconsistent more specific announcements by Indian networks.

2. While India geographically located between the South Asian countries, very few routes seems to be transiting via India on layer 3. Possibly a lot on layer 1 circuits. Handoff mostly happening in Singapore instead of South Asian region.

3. Interconnection of Pakistan with all other South Asian countries seems to be happen either in Europe or Singapore.

Examples of poor IP connectivity - Case 1

Trace from Lucknow, India to Kathmandu, Nepal

Traceroute to 202.52.0.25 (202.52.0.25), 48 byte packets

1 180.233.120.1 AS45942.180.233.120.1.lucknow.sikkanet.com AS45942 9.052ms 2.656ms 2.662ms
2 182.19.13.6 AS38266 13.594ms 13.447ms 13.428ms
3 182.19.106.103 40.918ms 73.03ms 39.429ms
4 195.89.115.197 ae28-100-xcr1.mar.cw.net AS1273 139.05ms 140.113ms 138.929ms
5 62.115.153.190 AS1299 139.113ms 140.013ms 138.955ms
6 62.115.133.176 prs-bb4-link.telia.net AS1299 176.217ms 188.051ms 175.083ms
7 62.115.114.228 ldn-bb2-link.telia.net AS1299 175.417ms 175.351ms 217.761ms
8 213.155.137.119 ldn-b7-link.telia.net AS1299 154.767ms 154.612ms 62.115.138.155 ldn-b7-link.telia.net AS1299 175.903ms
9 62.115.154.231 worldlink-ic-323350-ldn-b7.c.telia.net AS1299 175.857ms 176.615ms 175.718ms
10 103.225.212.146 AS133372 256.02ms 257.749ms 257.768ms
11 103.225.212.178 AS133372 258.403ms 258.29ms 263.401ms
12 202.79.41.210 ge-210.41.nren-rtr.wlink.com.np AS17501 253.034ms 252.822ms 254.69ms
13 202.52.0.25 atlas-anchor.nren.net.np AS45170 258.316ms 258.159ms 258.212ms

Examples of poor IP connectivity - Case 2

Trace from Amritsar, Punjab to Islamabad, Pakistan

Traceroute to 210.56.11.202 (210.56.11.202), 48 byte packets

1 192.168.1.2 0.838ms 0.416ms 0.377ms
2 100.66.160.3 5.05ms 5ms 4.898ms
3 192.168.241.41 5.138ms 5ms 4.89ms
4 192.168.249.133 4.959ms * 5.185ms
5 59.144.34.125 aes-static-125.34.144.59.airtel.in AS9498 14.031ms 14.396ms 14.096ms
6 182.79.189.102 140.76ms 140.396ms 141.056ms
7 149.14.124.17 te0-4-0-19.ccr21.mrs01.atlas.cogentco.com AS174 141.621ms 140.376ms 141.949ms
8 149.14.125.250 ptcl.demarc.cogentco.com AS174 237.079ms 238.111ms 238.304ms
9 202.125.128.173 static-10GE-KHI275-P01-SwA.pie.net.pk AS17557 239.407ms 236.305ms 233.951ms
10 221.120.254.13 rwp44.pie.net.pk AS17557 261.468ms 257.14ms 260.167ms
11 221.120.253.34 rwp44.pie.net.pk AS17557 257.078ms 278.971ms 263.372ms
12 202.125.149.58 Touchstone-4.rwp44d1.pie.net.pk AS9557 257.244ms 257.097ms 257.291ms
13 210.56.8.180 AS7590 257.321ms 257.294ms 257.462ms
14 210.56.11.202 AS7590 258.452ms 256.389ms 256.468ms

Examples of poor IP connectivity - Case 3

Trace from Mumbai India to Bangladesh

Traceroute to 59.152.99.17 (59.152.99.17), 48 byte packets

1 180.149.246.1 AS33480 0.658ms 0.637ms 0.36ms
2 223.31.163.17 AS9583 1.355ms 1.343ms 1.067ms
3 100.65.218.170 132.27ms 133.156ms 132.234ms
4 149.6.148.217 te0-2-0-2-2.ccr21.lon01.atlas.cogentco.com AS174 132.529ms 132.249ms 132.039ms
5 154.54.57.153 be2868.ccr41.lon13.atlas.cogentco.com AS174 132.525ms 132.32ms 132.294ms
6 154.54.56.130 be12497.ccr41.par01.atlas.cogentco.com AS174 139.782ms 139.864ms 139.507ms
7 130.117.49.154 be3092.ccr21.mrs01.atlas.cogentco.com AS174 150.334ms 150.374ms 150.598ms
8 149.14.126.122 bsccl.demarc.cogentco.com AS174 202.056ms 202.195ms 202.234ms
9 103.16.152.30 103-16-152-30-noc.bsccl.com AS132602 202.275ms 202.024ms 202.156ms
10 103.16.155.106 103-16-155-106-noc.bsccl.com AS132602 201.744ms 208.567ms 201.879ms
11 103.21.41.110 AS58715 201.864ms 201.776ms 201.537ms
12 59.152.96.38 AS63969 202.027ms 202.054ms 201.821ms
13 59.152.99.17 AS63969 203.095ms 202.035ms 202.161ms

Reasons for bad routing

1. Commercial reasons
2. Routing design challenges
3. Regulatory Challenges
4. Others

Commercial Reasons

Commercial Reasons for bad routing

- Circuits to major hubs tend to be cheaper and this becomes Chicken egg issue. E.g circuit between Maldives & Sri Lanka is way more expensive than circuit between Maldives & Singapore. (It's cheaper to send traffic from Maldives to Sri Lanka via Singapore!)
- Operators tend to peer outside & not in their home market which leads to more peering at hubs far off from the location
- Larger hubs tend to pull more content & CDN players resulting in better utilisation of larger circuits
- Often regional players tend to sell bandwidth at a commercially expensive price, thus it becomes easy to reach them from larger International players from major hub via a less optimised path
- Price of Internet transit is as low as $0.15/Mbps/month in major connected markets & hence for less latency specific traffic, operators always have option to pump traffic via less optimised path

Routing design challenges

Routing design challenges

- Peering in selective regions and not all tends to push traffic via those peering points
- Many incumbent telcos tend to buy long haul circuits (IPLCs) to far off locations without connecting to networks in nearby regions. This leads to traffic flow towards them from outside region
- Often missing of dual stacking on IPv6 BGP sessions

Regulatory Challenges

Regulatory Challenges

- Some countries (like India) in the region tend to have additional regulations for connecting outside the country. For instance it’s much harder for Bangladeshi or Nepalese ISPs to have PoP in India & peer with content networks in India Vs having a PoP in Singapore or Hong Kong

- The requirement of LIM (Lawful Interception) by various countries tend to further complicate network design & capacity planning. In the world of end to end encrypted text messages, SSL enabled websites, Letsencrypt and more, the benefits of LIM need to be revisited by regulators

- Regulatory burden of declaration of “International PoPs” leads limited large L3 connection locations in the region even when fiber is available on better optimised paths

Other reasons...

Other misc reasons

- Political tensions - There seems no visible circuits between Pakistan & India for public Internet traffic despite of direct submarine cable
- Traffic is getting more & more concentrated. This has moved focus of eyeball networks away from optimising regional connectivity
- Nepal recently activated it’s first link via China & is expected to push more traffic via hubs in East Asian region over that link
- Lack of peering across incumbents leads to exchange of their respective customer traffic far outside the region

Conclusions

Conclusions

- Routing within South Asian region appears to be poor overall. There are limited cases of direct symmetric routing.
- Countries like Nepal & Bangladesh seem to have direct connectivity to India to some extent since many operators there are customer of Indian operators.
- There are almost Zero International peers at IX’es across India (NIXI), Nepal (NPIX), Bangladesh (BDIX) etc. Most of traffic exchange seem to be transit-downstream relationship based on visibility of announcement.
- Pakistan, Maldives, and Sri Lanka to a great extent seem to be exchanging traffic only at larger International peering hubs.
- There is a extremely high dependency on majority submarine cables. Any natural disaster can lead to very poor connectivity between South Asian countries

Misc Points / Methodology used

1. It is assumed to be the case of bad routing if 75% or more probes show bad routing data (and vice versa) since purpose of study was to look at routing in general and not specific to network.

2. All the available RIPE Atlas probes were used to monitor routing towards anchors & reverse trace from anchors to probes was also done to verify cases of asymmetric bad routing.

3. Relationship between two networks is assumed to be transit relationship if one is announcing routes further & route is visible across Tier 1 network operators due to that announcement.

References

1. Countries in South Asia - https://en.wikipedia.org/wiki/South_Asia

2. NLNOG Ring - https://ring.nlnog.net

3. RIPE Atlas Project - https://atlas.ripe.net

Questions?

Drop a mail - anurag@he.net
Web: https://he.net