Martijn Schmidt & Bob Stevens

# Cables under the sea & other nautical nonsense

# i3D.net

**40+ metro areas, 60+ PoPs**

**6 continents**

**Backbone mostly using leased waves**

**Capacity on 20+ different subsea cables**

**100+ IXPs, 9200+ unique ASNs adjacent**

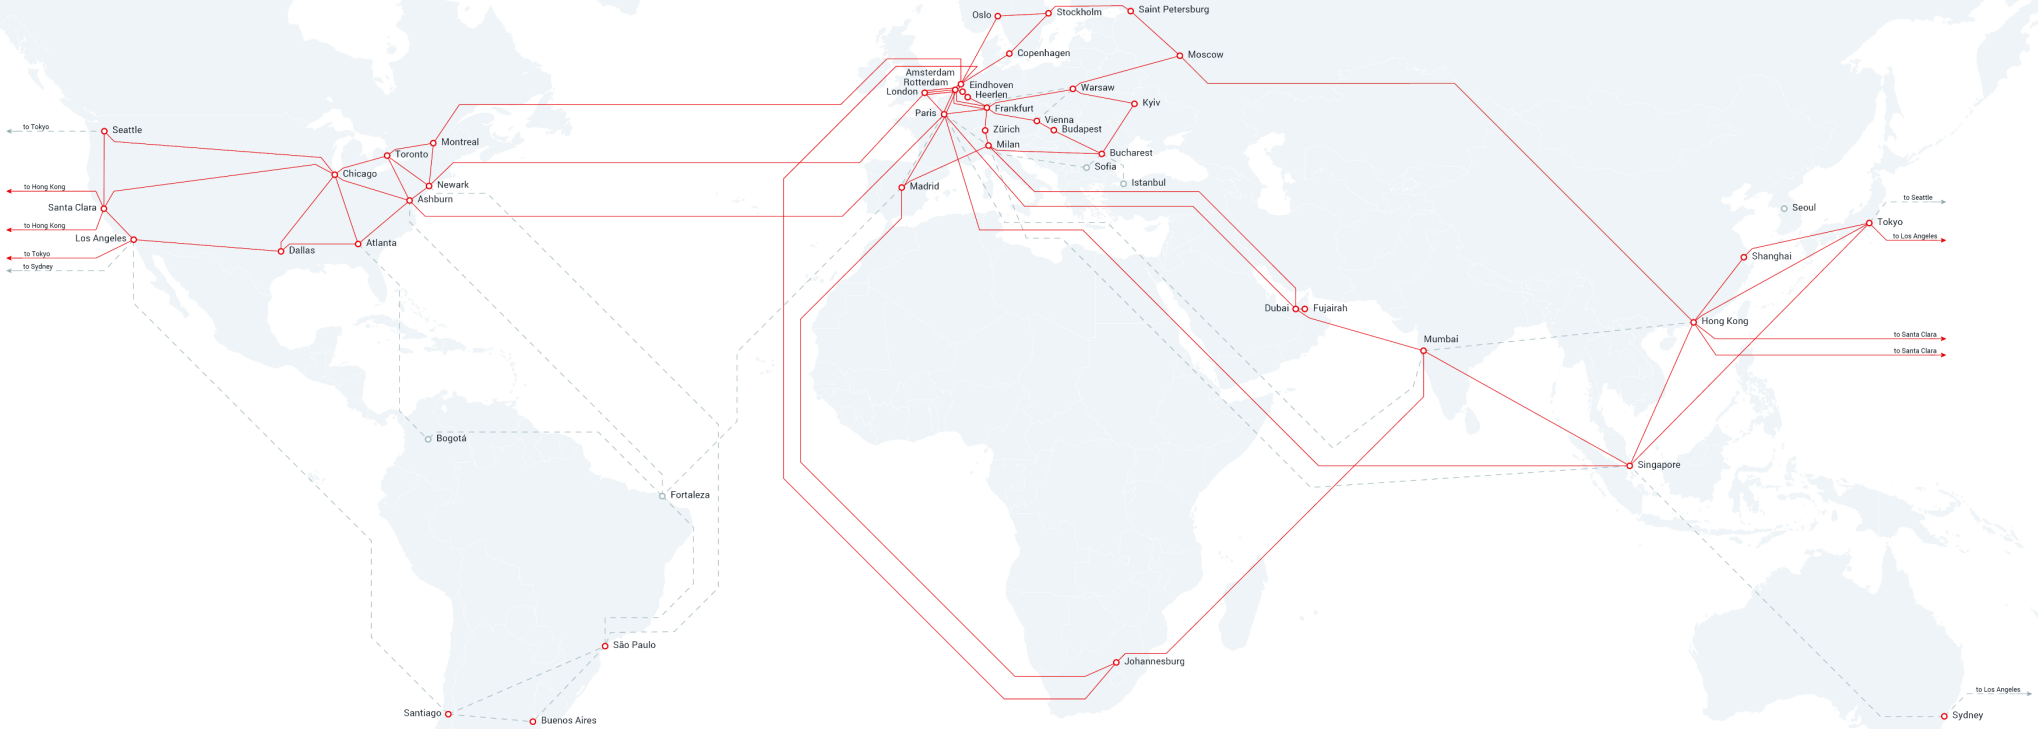


i3D.net

# Why do we go through all this effort?

- **IP transit is generally cheaper per Mbps than subsea**
  - You are paying your €/Mbps for a blended traffic consumption
  - Transit providers include contractual clauses to cover themselves like:
    - "no more than X% traffic to other continents"
    - "no more than X% traffic to any one neighbor"
  - "Backplane transit" is practically free, subsidizes expensive long-haul

- **Cost efficiency possible with i3D.net's specific use case**
  - We originate traffic & peer equally in all continents
  - As such, our subsea capacity can be used bidirectionally
  - Note: pie chart also includes non-subsea long-haul backbone

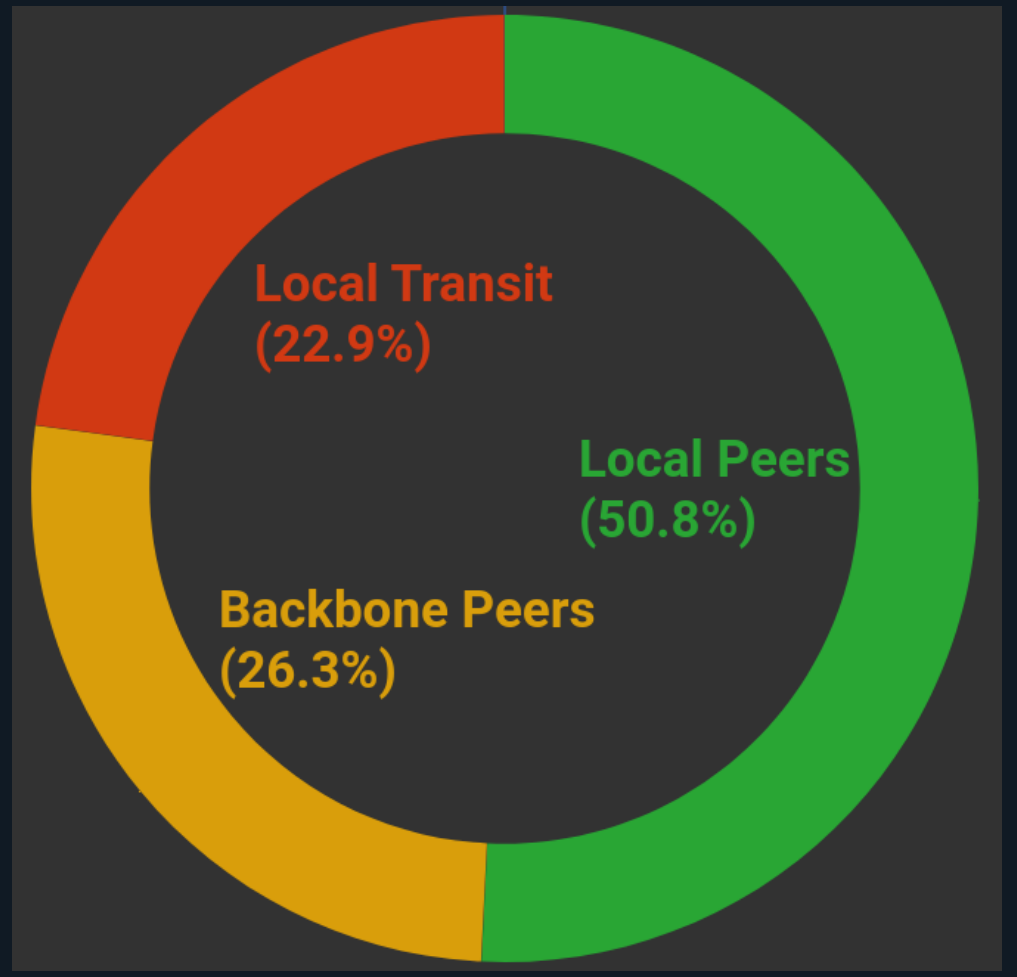


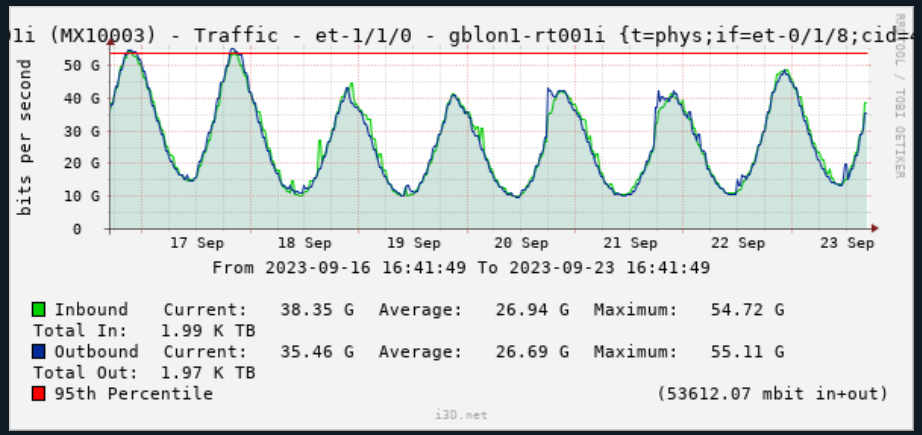


i3D.net

# On the technical side of the equation..

- **Control, stability, visibility**
  - Long-distance performance is important for games, but not all NSPs (and/or their customers) need it / do it well
  - Blast zone size reduction, insight in true capacity & self-mitigate (temporary) saturation events
  - Visibility on the Internet, look for example to Arjan's BCP38 talk

- **Be our own DDoS mitigation**
  - Tier-2 upstreams in complex regions would melt under our DDoS attacks (graphed 1 week of DDoS + volumes in Gbps)
  - Scrub at ingress, backhaul just the clean traffic over expensive subsea capacity

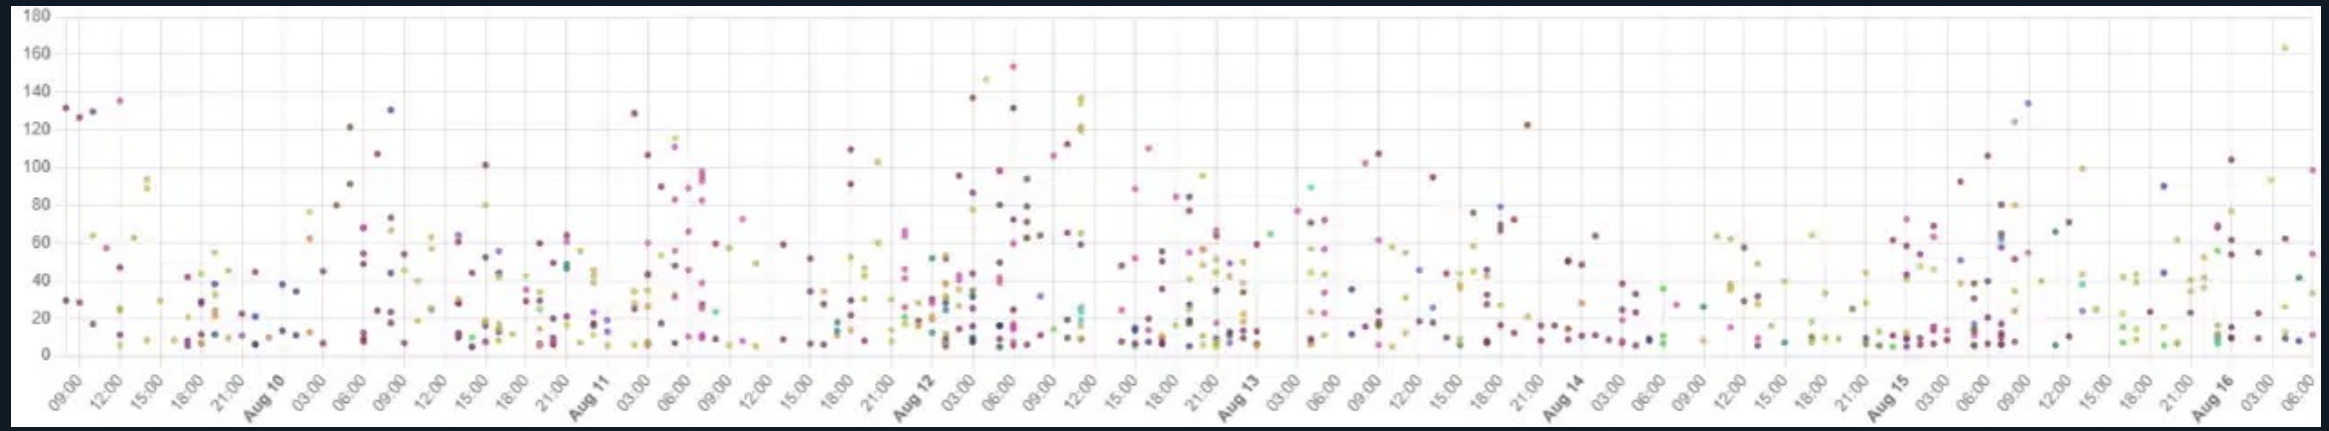

# Important milestones

- 1858 – Transatlantic Telegraph Cable
- 1956 – TAT-1 – First transatlantic telephone cable (coax)
- 1988 – TAT-8 – First fiber optic sea cable
- 1997 – FLAG Europe-Asia – first non-consortium cable
- 2000 - SEA-ME-WE 3 – Longest cable (until 2Africa is finished)
- 2001 - TAT-14 – First 10G transatlantic sea cable

# Designs

- **Point to point**
  - Direct routes between point A and point B
- **Fish-bone**
  - More complex design with "Express routes" and less optimised paths via
  - Branching units that are used to break up some of the cable's fibre-pairs
  - Multiple PFEs needed to power up the whole system
- **Ring System**
  - Mostly deployed in SDH/SONET era to offer customers path protection
  - Protected path has higher latency
- **Festoon**
  - Runs along the coastline from CLS to CLS, usually unrepeatered

ZDnet UK

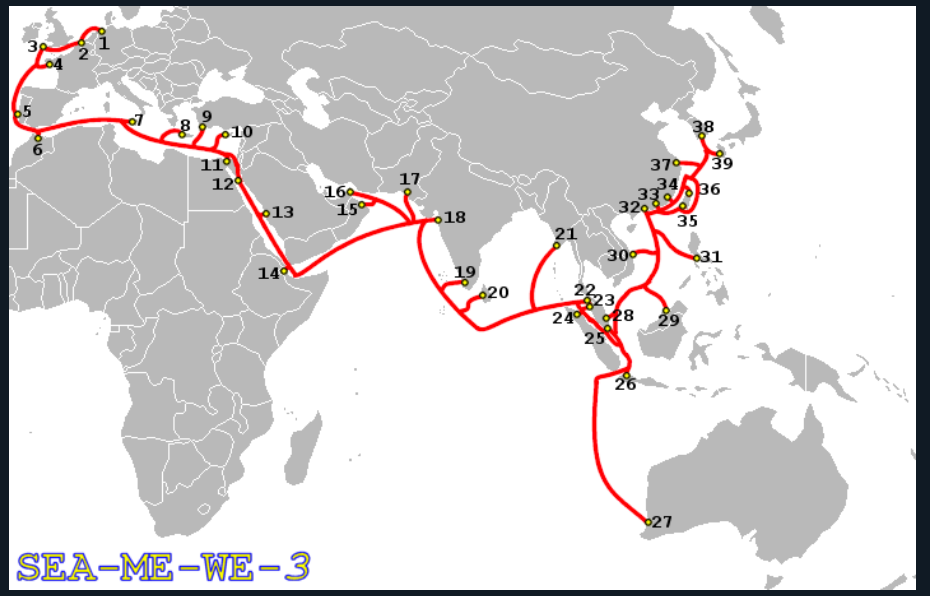


i3D.net

# Cable Landing Stations

- CLS or Cable Landing stations are usually close to the shore, where high-voltage power is fed into the cable to power up subsea amplifiers (if used).
- Resilient design just like datacenters to guarantee uptime & stability
- “Wet segment” connects to the terrestrial backhaul via "Beachjoint" near cable landingstation
- Back in the day used to drop connections at the CLS, now mostly ”Datacenter to Datacenter”
- Older CLS have "colocation rooms" for all providers to house terrestrial network
- CLS can charge crazy "landing fees" for cross connects from the Seacable towards 3rd party providers
- Landing stations of consortium cables are usually operated by local incumbents

i3D.net

# Two types of cable systems

- **Unrepeatered cable systems**
  - Use Laser class 4 - high power RAMAN Amplifiers
  - Extra safety measures need to be applied by personnel working
  - Max 400km distance can be crossed.
  - Very commonly used for shorter intra-European distances
  - Much more fiber-pairs possible since no limitation on electricity
  - Economic life dictated by # of repairs, not as much age of the system


i3D.net

# Two types of cable systems

- **Repeatered cable systems**
  - Every 60 – 80 kilometers are repeater is placed
  - Long point to point cables can have 150+ repeaters
  - Over-engineered for life span of 25 years
  - Repeaters are powered from shore ends
  - Repeaters are “managed” via supervisory channels to retrieve status & performance info
  - On old cable systems floppy disks are still part of the day-to-day work


Alcatel Twitter

# Types of cable

- Different kind of sea cables are used for different segments
- Cables can become exposed due to movement of underwater dunes
- Higher risk areas use more armored cables due to:
  - Fishing
  - Anchors
- Lightweight and Lightweight Protected cable is used for deep sea segments

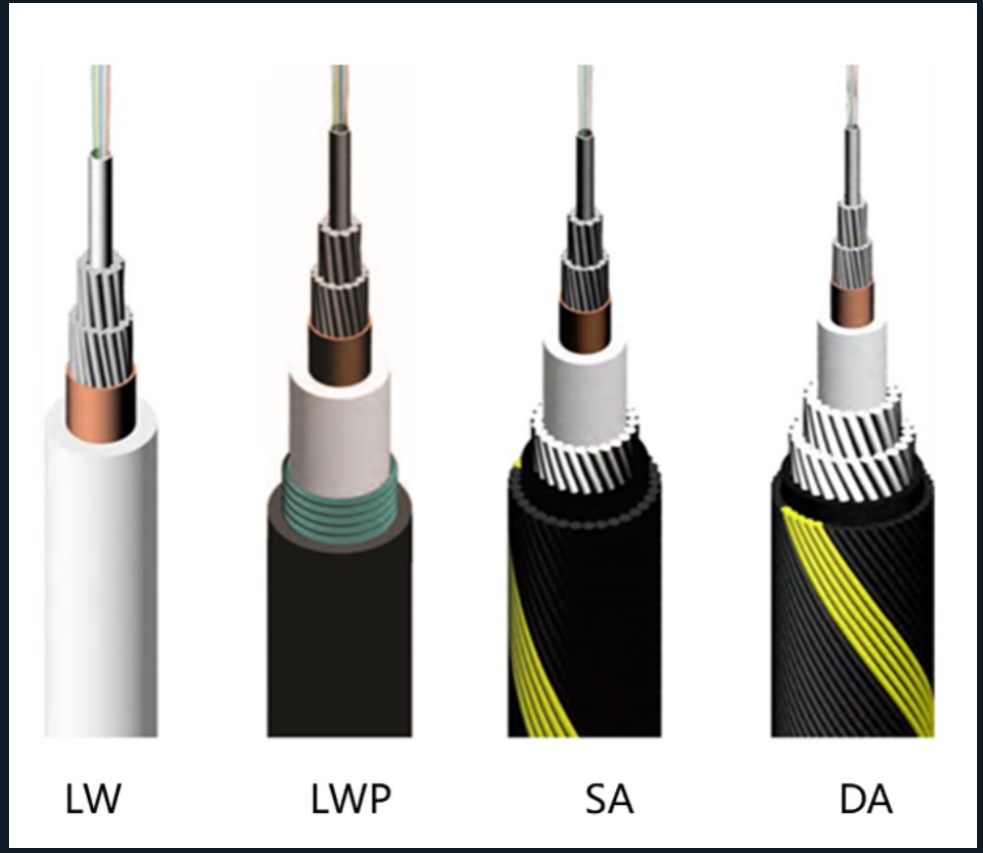


Source: hmntechnologies.com

# PFE - Power Feed Equipment

- Repeaters are powered with DC from both Landing stations.
- Cable is fed by negative and positive current
- Multiple rectifiers are in place to provide high power 6000 volt for longer path’s
- Typically half of a cable is powered by the PFEs
- Shorter cables can be powered from one landing station so dual feed improves reliability
- Cabinets, buttons/switches are locked by keys and have safety measures
- Includes a tone generator that the cable ship uses to track/find cables during repairs

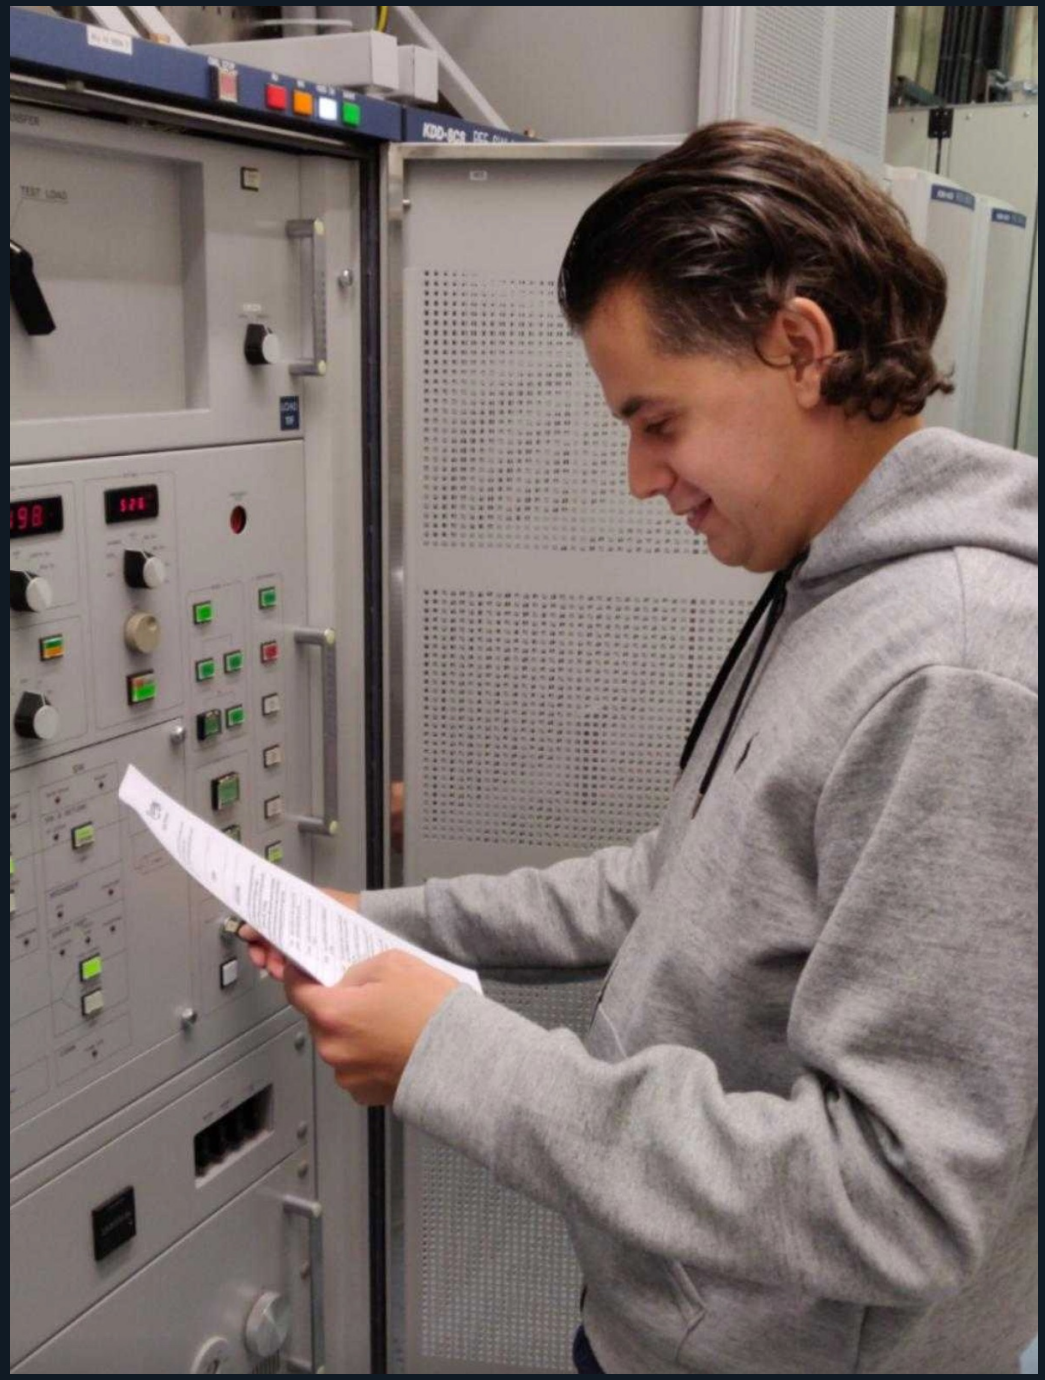

i3D.net

# SLTE - Submarine Line Terminal Equipment

- SLTE is where the client signals from the terrestrial backhaul are connected to the active equipment of the wet segment
- Older sea cables were designed with SDH connecting to the DWDM equipment in CLS
- Optimized for long-haul with enhanced FEC capabilities
- During the lifetime of a seacable SLTE is upgraded to support next-gen speeds and improve capacity

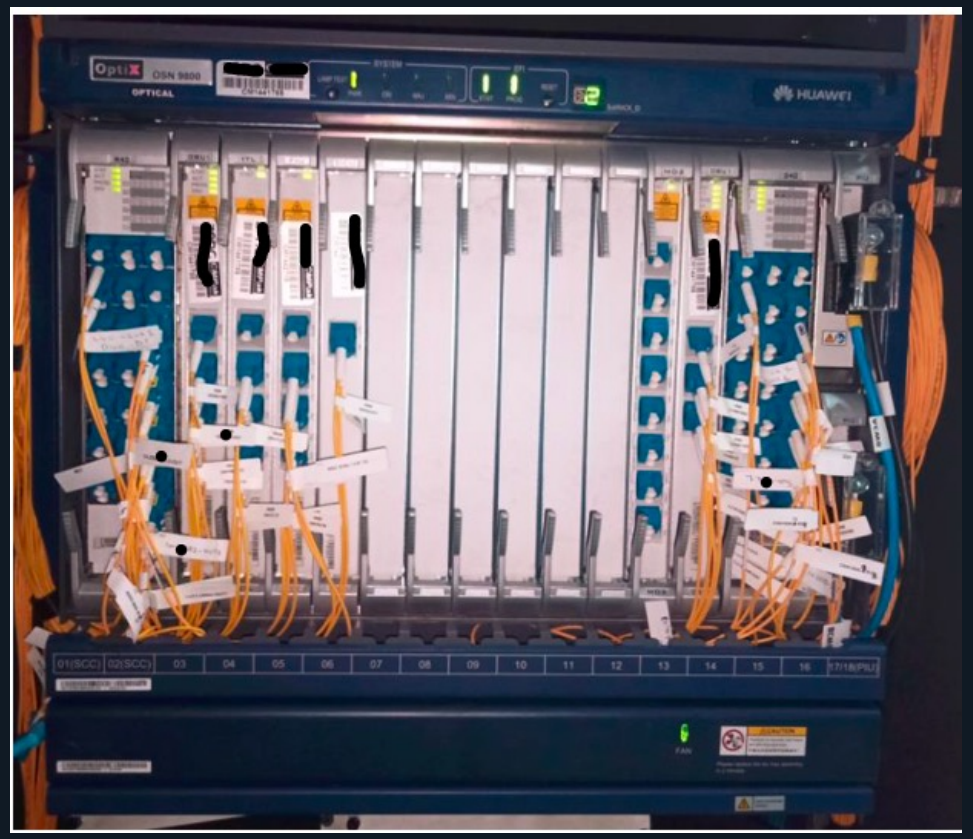


forum.huawei.com

# Usual faults impacting Cable systems

- Different type of faults impact subsea cable systems:
  - Complete outage of power & fibers
  - Shunt fault
    - Insulation is damaged causing short circuit
    - Hard to detect / measure fault location
  - Repeater & branching unit failure
  - Partial failures where one branch is impacted

# Cable repairs

- Repairs are performed by cable ships
  - Used to be "converted" ships but now purpose-built for subsea cable repairs and laying
  - Special equipment to recover and repair cable
  - Ships work closely with sea cable station for different phases of the repair

- Cable systems usually join non-profit cable maintenance agreements like the Atlantic Cable Maintenance Agreement
  - Multiple ships on standby for emergency call-outs
  - Ships load spares at depot, then sail to the repair grounds

- We have a premium account at MarineTraffic to see movement of the ships if one of our paths is impacted.


Orange Marine

i3D.net

# Challenges and planning for faults

- Bob talked about when things go according to the plan/design on L1/L2/L3 of the OSI model.
- I'm spending my time more and more at L8/L9 these days.
- Things can go (badly) wrong, it's a game of playing the odds:
  - Frequency of outages
  - Duration of outages
  - Impact of outages

- RFC1925 applies:
  - (1)  It Has To Work.
  - (8) It is more complicated than you think.
  - (9a) Every networking problem always takes longer to solve than it seems like it should.

# ..simple, right? Call the CMA and go

- Not all subsea cables are operated by (groups of) companies that subscribe to a CMA: back of the queue
- When there are more cable faults than repair ships in a given CMA zone, how do you determine the repair sequence?
  - First in, First Out?
  - Least-connected countries first?
  - Highest capacity cables first?
  - ..no nuclear fallout at repair grounds? (Fukushima 2011)
- Still a multi-week repair job in the very best case, ships simply don't move as fast as terrestrial repair crews
- The repair work may pause for various reasons:
  - Bad weather at sea, typhoons/cyclones/hurricanes, etc
  - Live-fire naval training exercise near repair grounds

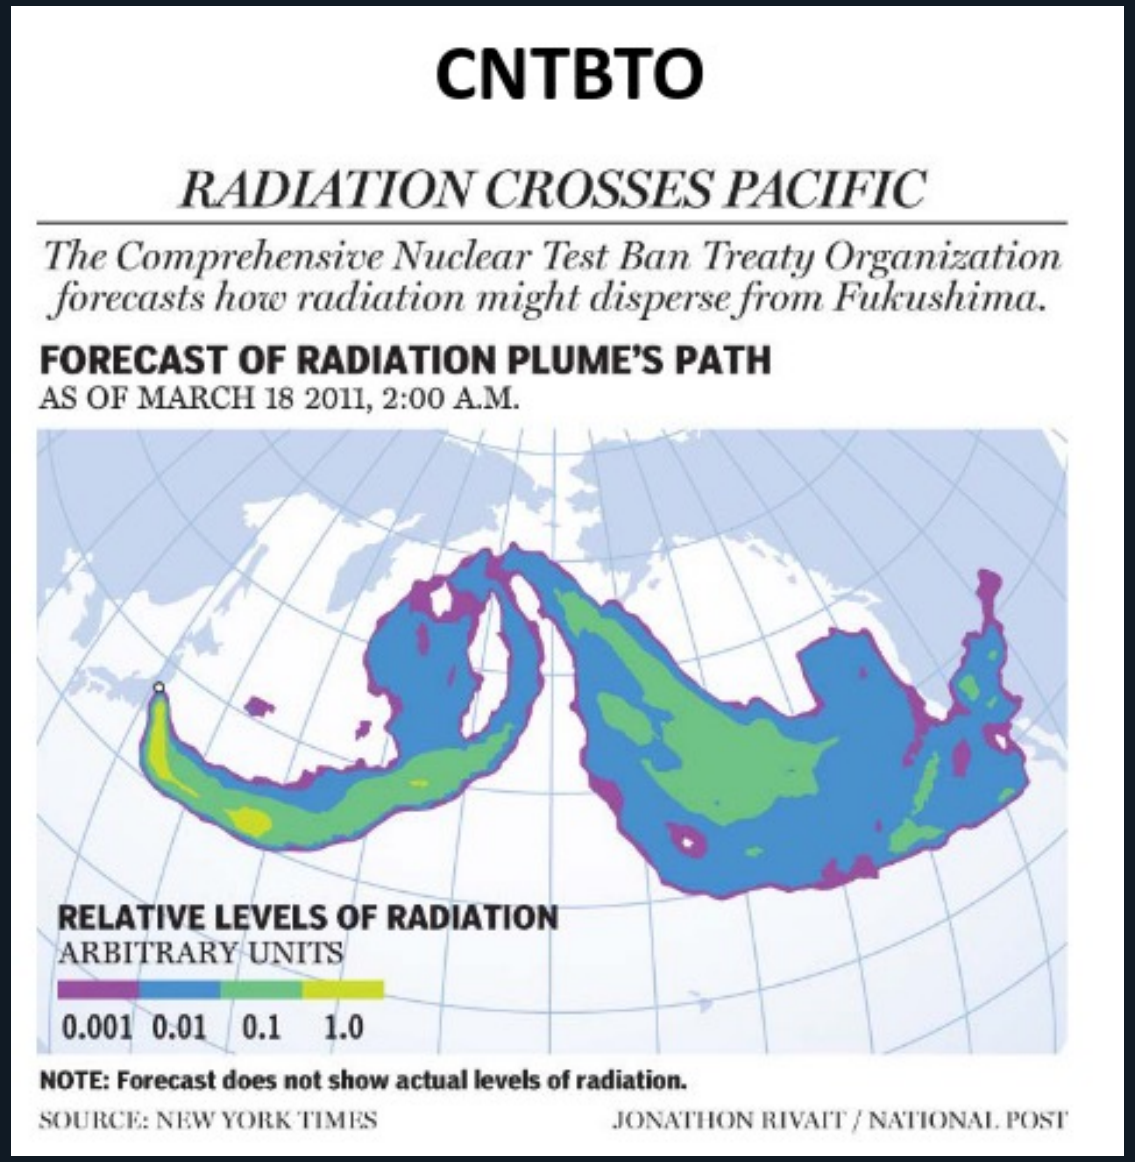

# Did you think you could just sail out?

- Governments also follow RFC1925, specifically (6a): It is always possible to add another level of indirection.
- You need a permit to operate the repair ship in a country's territorial waters. No big deal in most Western countries.
  - Can take a long time to approve repair in certain regions even if $country is itself a victim of the subsea outage.
  - What if multiple countries claim the same territorial waters – sometimes up to 4 concurrent? Get permits from everyone to err on the safe side, or skip and risk offense?
- Some countries have cabotage policies:
  - Several weeks for domestic bids on the repair contract
  - But $country has no companies operating cable ships
  - Or if they do, $cable_owner may have policies to avoid $country's services for geopolitical reasons

# Fake network diversity

- Even if you use multiple subsea cables, you may have fake redundancy – especially coastal waters are the danger zone:
  - Cables landing at the same few km of beach may be damaged simultaneously by anchors or fishing nets
  - Underwater landslides – the Taiwan 2006 earthquake is infamous, Congo River mudslides in 2021 and 2023
- Hurricane Sandy flooded NYC in 2012, affected multiple DCs:
  - Mains grid outages
  - Fuel trucks couldn't always reach their destination
  - Sometimes even generator & fuel storage in basements
- When all forced to use same incumbent's fibers in $country
- Terrestrial backhaul overlap due to provisioning mistakes despite KMZ design efforts, "same path" vs "same exact path" between ROADM locations

# Planning for true network diversity

- We built EU <> NA connectivity on triple paths with landings hundreds of kilometers apart after being taught a lesson by CyberSquirrel1 & "same exact path":
  - London <> NY/NJ, the "default route" for many NSPs
  - Amsterdam <> Montréal, though relatively longer path to big interconnection hubs may require IGP fake metrics
    - Boston can reach US big telco PNIs, but no i3D.net PoP
    - Investigating potential Dublin breakout now
  - Paris <> Ashburn, but southern route means higher latency due to less benefit from curvature of the earth (Mercator projection distortion)
  - Adding a 4th path London <> NY/NJ, traffic naturally gravitates to that bundle, further improves diversity with discussed caveats

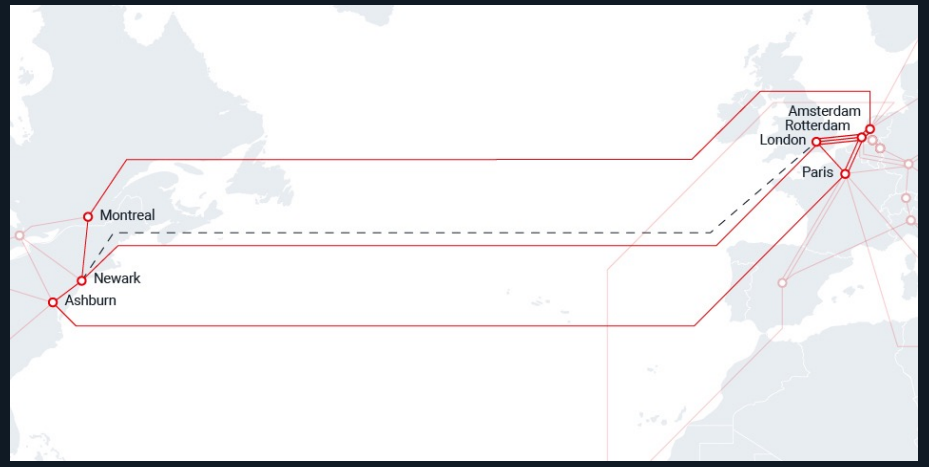


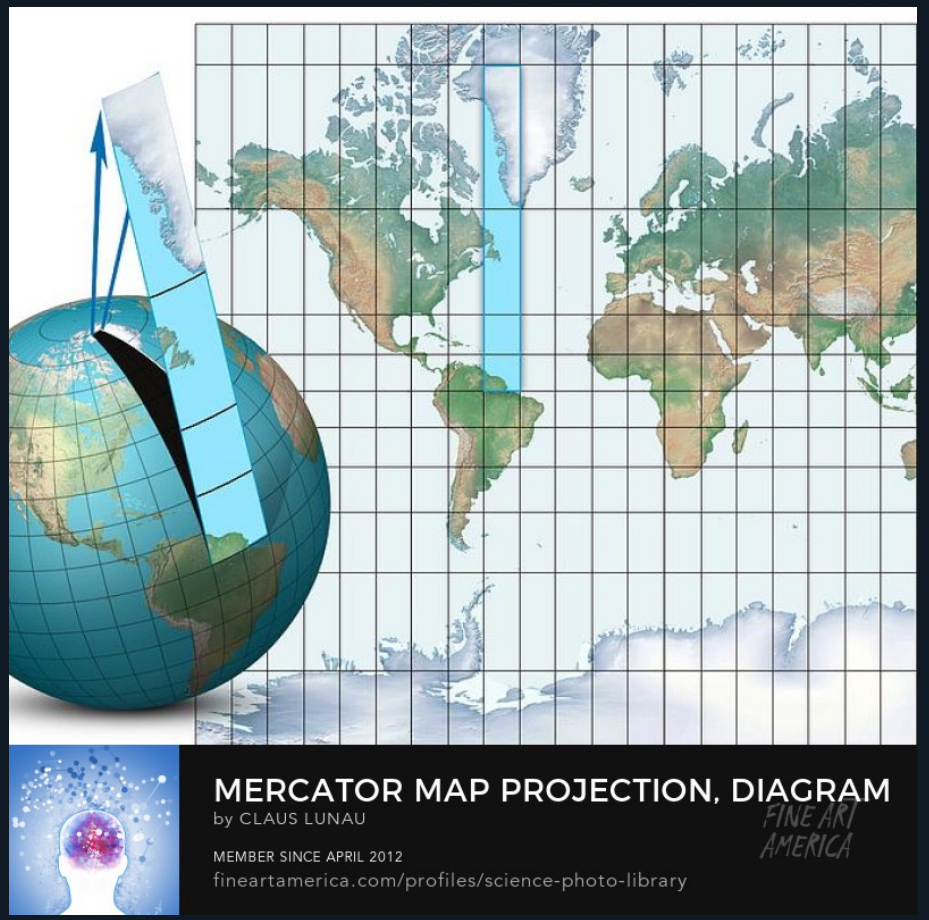


i3D.net

# Capacity crunch: political edition

## Building a new cable is a multi-year endeavour

- Finding investors, performing marine surveys
- Contracting with manufacturers, cable layers
- Building CLS if not re-using existing ones

## US-China rivalry affected new trans-Pacific cables

- PLCN already laid, but denied activation of US landing unless Chinese consortium members fully divested & HK branch gone
- BtoBE reconfigured to CAP-1 without Chinese consortium members & HK/SG/MY branch replaced by PH
- Several other cables cancelled, now being replaced by the likes of Apricot, Bifrost, Echo, TPU but it resets the planning progress

**Team Telecom Recommends that the FCC Deny Pacific Light Cable Network System's Hong Kong Undersea Cable Connection to the United States**

Wednesday, June 17, 2020

Share

**For Immediate Release**

Office of Public Affairs

**Committee Recommends FCC Conditionally Grant Portions of the Application Requesting Connections between the United States, Taiwan, and the Philippines Subject to Google and Facebook Entering into Mitigation Agreements with U.S. Government Agencies**

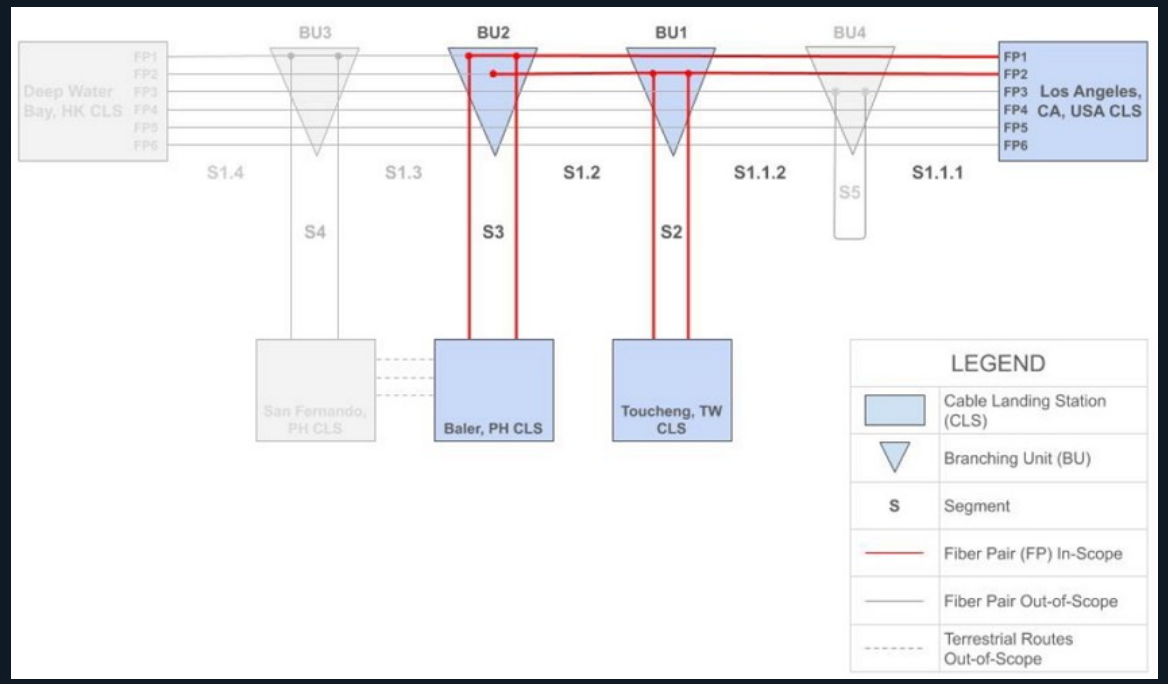


i3D.net

# Missing subsea diversity & blank spots

## (Fake) redundancy issues between some regions

- EllaLink (100 Tbps) is the only direct EU-LATAM cable for now, Atlantis-2 (160 Gbps) shut down due to end of economic life
- Nearly all EU-MiddleEast-Asia cables pass through Egypt, small capacity bypasses are possible but expensive. Will be partially tackled with the launch of Blue-Raman via Israel/Jordan
- Asia to Sub-Saharan Africa is only served by the ancient SAFE cable with a design capacity of 440 Gbps – a single OTN superchannel carries more capacity nowadays

## Completely missing paths (see TeleGeography)

- Oceania-Africa
- Oceania-LATAM
- Europe-West Asia

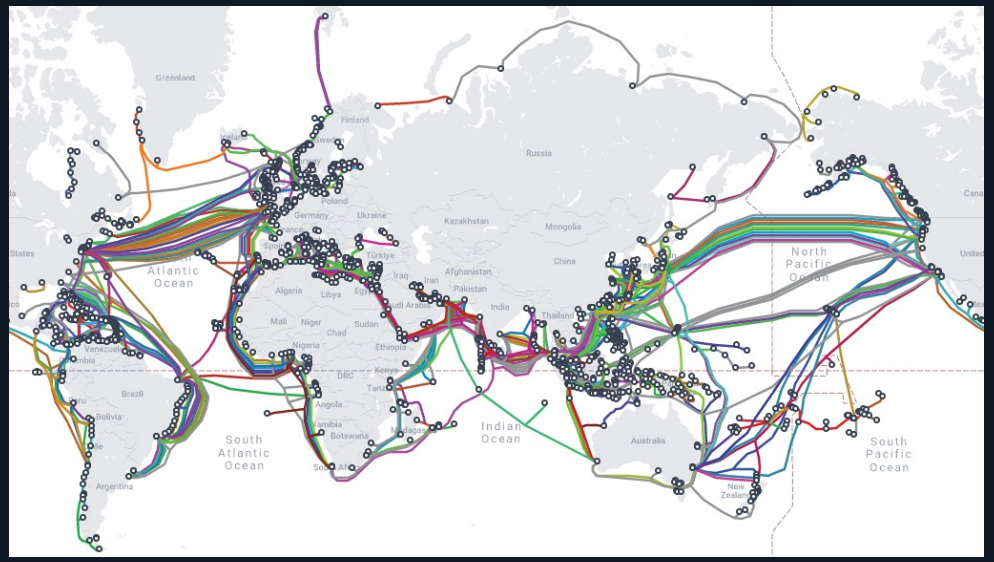


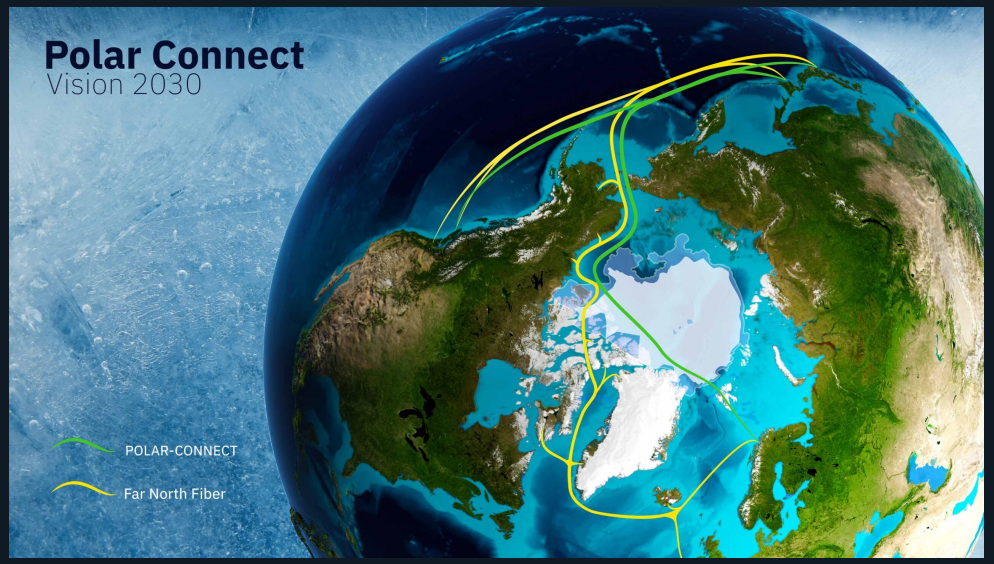


i3D.net

# Looking forward

## Dutch subsea cable coalition

- Various government entities and market players (builders, operators, users, etc) pushing for new cable landings in NL
- Erasmus cable from Rotterdam to London, to avoid the Zandvoort/Lowestoft area landing the most recent cables
- Polar Connect, Far North Fiber, connecting Europe to Japan

## Possible latency gains

- RFC1925 (2): No matter how hard you push and no matter what the priority, you can't increase the speed of light.
  - Light in a fiber-optical cable travels at ~2/3rd of speed of light in vacuum due to optical refraction: v=c/n
  - New medium with a refractive index closer to 1?
  - Space lasers with meshed LEO satellites, but low capacity and only efficient at very long distances due to earth-orbit-earth penalty (~25ms according to Starlink).

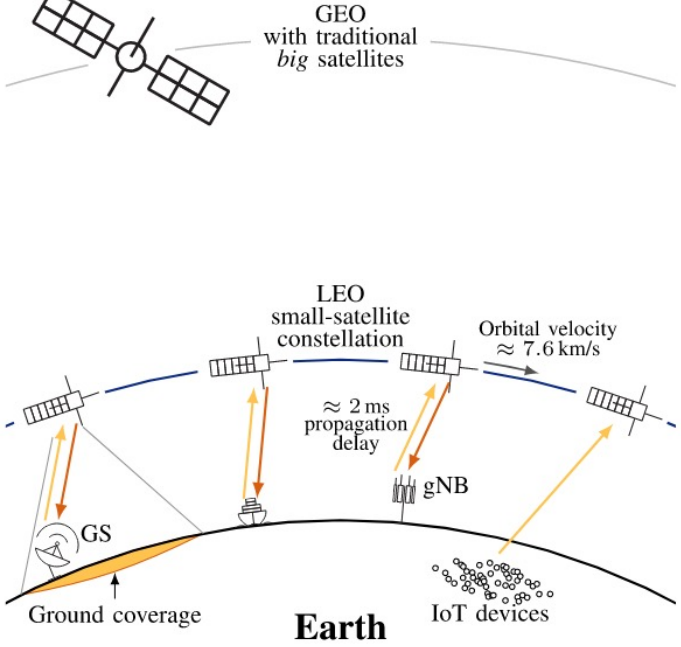


**FIGURE 1.** Overview of the unique characteristics of LEO small-satellite constellations with respect to traditional GEO satellites.

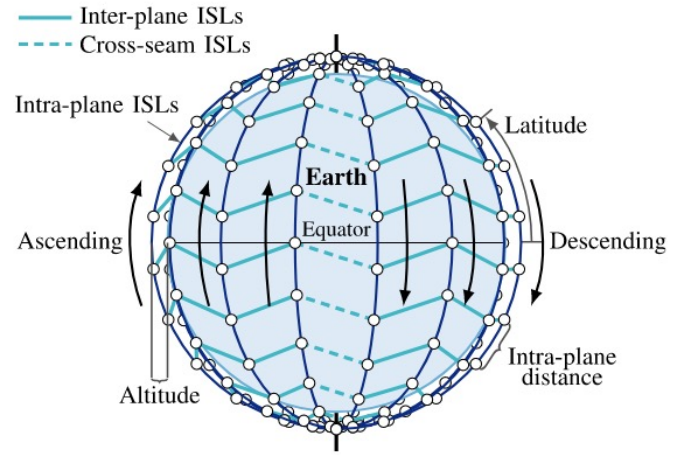


**FIGURE 2.** Diagram of a Walker star LEO constellation with the established intra- and inter-plane (including cross-seam) ISLs.

I. Leyva-Mayorga et al.: LEO Small-Satellite Constellations for 5G and Beyond-5G Communications

# Thank you!

# Questions?