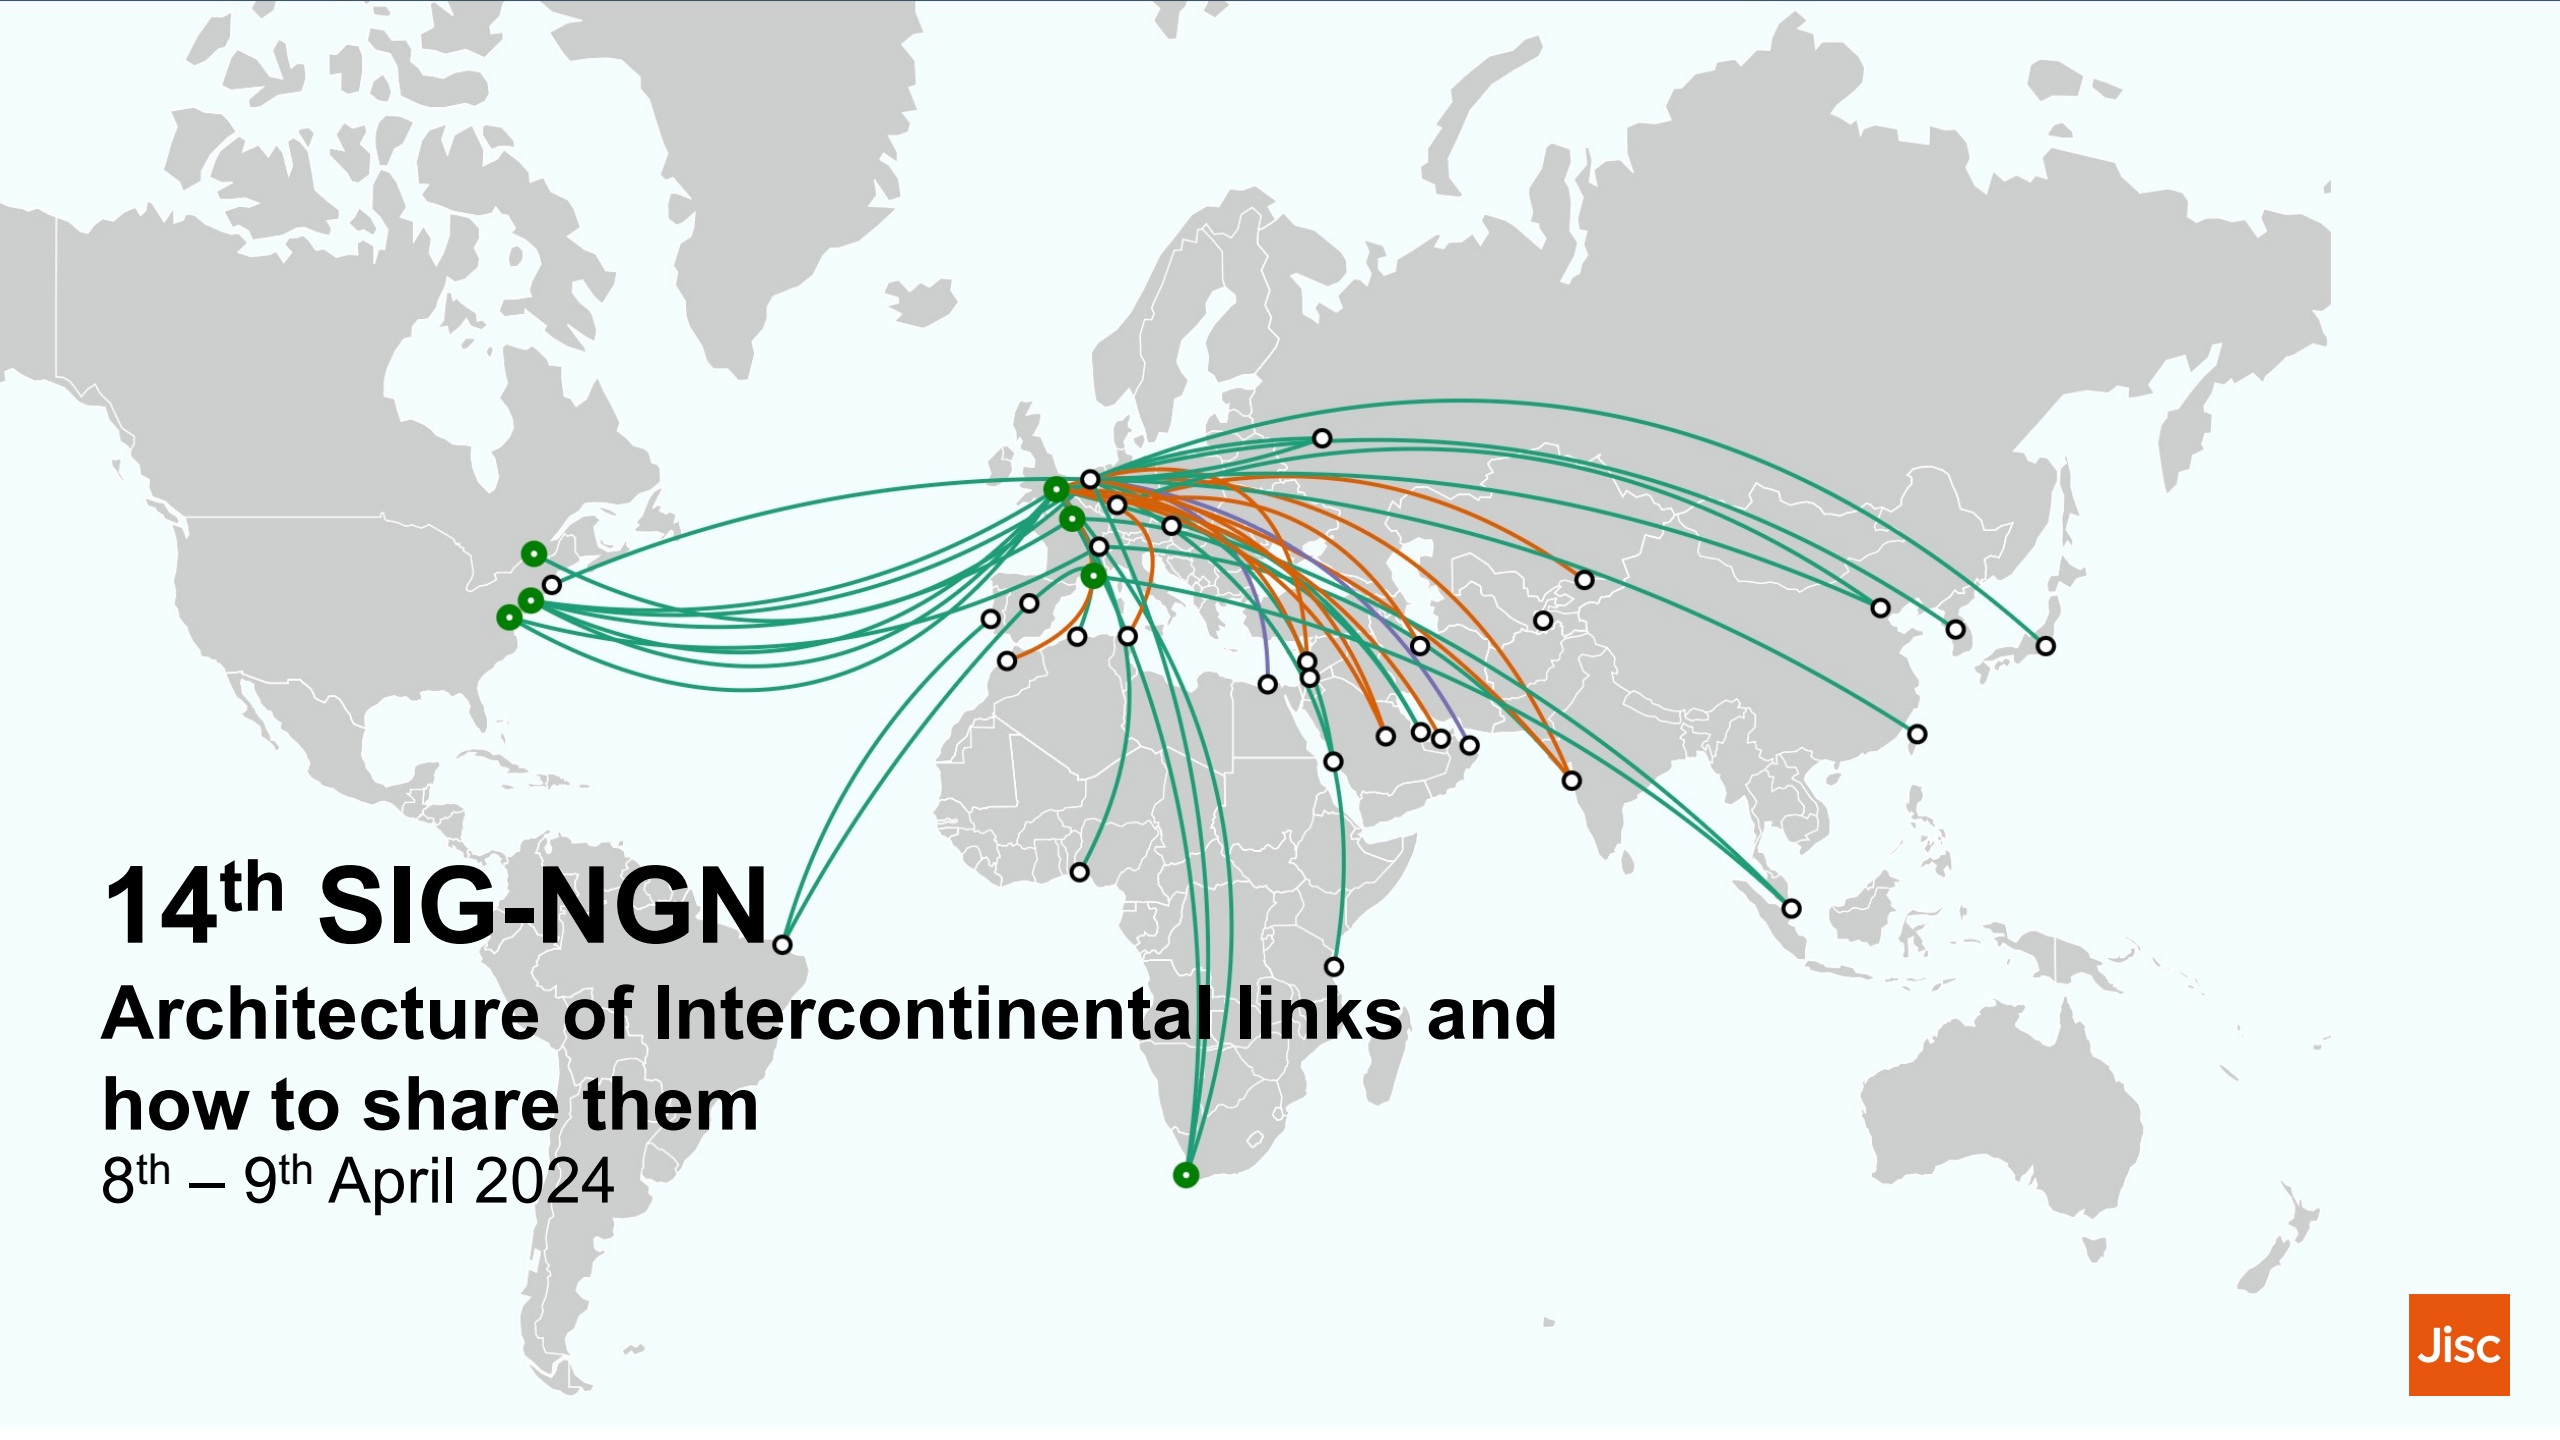


14th SIG-NGN

Architecture of Intercontinental links and how to share them

8th – 9th April 2024



Welcome to 14th SIG-NGN

- SIG-NGN Steering Committee
 - Rob Evans
 - Lars Fischer
 - Edoardo Martelli
 - Carsten Pettersson
 - Yatish Kumar
 - Mian Usman
 - Rudolf Vohnout

Cover image is from map.geant.org, there is also globalresearchmap.org and the efforts of the GNA-G mapping WG (and probably others).

Agenda

Day 1: Architecture of Intercontinental Links

1. Introduction and scene setting. Rob Evans, Jisc (20 mins)
2. How to use policy and technology to share with the community. Edward Moynihan, Indiana University (20 mins)
3. Future evolution of network and applications in Asia region. Buseung Cho, KISTI (20 mins)
4. Submarine cable technology and trends. Roeland Nuijts, Ciena (20 mins)
5. Coffee (30 mins)
6. Connecting Australia to the world. David Wilde, AARnet (20 mins)
7. MEDUSA cable. Karel van Klink, GÉANT (20 mins)
8. Planned links across the Atlantic and other intercontinental links (Bram Peeters, GÉANT)
9. Use of Spectrum sharing on $\frac{1}{4}$ fibre on intercontinental (subsea) links. Alexander van den Hil, SURF. (20 mins)
10. Open discussion (20 mins)

Agenda

Day 2: Sharing the links already built

1. Welcome. Yatish Kumar, ESnet (5 minutes)
2. Layer 2 and layer 3 overlays. Lions and tigers and bears, oh my! Rob Evans, Jisc (10 mins)
3. Evolution of LHC networking, future perspectives. Edoardo Martelli, CERN. (15 mins)
4. Spectrum sharing in the nordic NRENs. Rasmus Lund, NORDUnet. (15 mins)
5. Time & Frequency in the Netherlands. Sander Klemann, SURF. (10 mins)
6. CERN-CNAF DCI. Paolo Bolletta, GARR. (20 mins)
7. Coffee. (20 mins)
8. Infrastructure Sharing. Christian Todorov, Internet2 (20 minutes)
9. Support multiple “private” networks over WAN. Eli Dart, ESnet (30 mins).
10. Panel Discussion chaired by Lars Fischer, NORDUnet (20 mins).

Need to finish promptly to enable lunch before LHC meeting starts.

Intercontinental links, historical perspective

- Note: These days my NREN barely has any cross-border connectivity, never mind intercontinental.
- Janet had 1 x 155Mbps links to New York, upgraded to 2 x 155Mbps just after I joined, and eventually went to 6 x 155Mbps, plus a 2Mbps (!) link to China.
- TEN-155
 - DANTE World Service for commercial traffic
 - Peering with North American R&E networks
 - Dedicated links for some larger NRENs.
 - ATM PVCs for other NRENs.
- *High management overhead.*
 - Changing committed bandwidth of PVCs, moving them between links, dealing with failures.

Intercontinental links, historical perspective

- Links to the US were essential at this time to provide good-quality Internet transit.
- Not the case any longer is most of the world, yet even without that requirement the scaling of intercontinental bandwidth is increasing rapidly.
- Mbps -> Gbps -> Tbps (literally)
 - How do we provision?
 - How do we scale?
 - How do we share?

Terrestrial v Submarine

- Submarine links are expensive and take a long time to plan, build and commission.
 - Ellalink at least €50M?
 - Unlike terrestrial systems, no benefit until the entire cable is built!
- Previously the domain of large telcos and Tier 1s who would only sell circuits.
- Still expensive and take a long time to build and commission, but now consortium-led, often with one of the big content providers (Google, Facebook), and in many relevant cases for us, the EU.
- Consortia more open to NREN involvement (as long as we pay our way).
- Upgrading submarine systems a challenge.

Load-sharing

- Submarine capacity is rarely being bought between the same pairs of cities
- Link aggregation not suitable, so how is load-sharing done?
- Some knowledge of topology and BGP communities?
- Messy, manual, uneven load-sharing.
- Can we do better (do we do better?)

Resilience

- Submarine breaks can take weeks or months to fix, and are weather dependent.
- Need to expect the links to fail, because they will, and at the most inopportune moment.
 - E.g. UK Tier 1's LHCOPN links during this year's data challenge.
- Does that mean capacity has to sit idle for 95% of the time?
- Resilience within one management domain is relatively easy.
- Resilience/backup between management domains?
 - Often manual ("Help! Can we use some of your capacity?")
- Overlay networks often manually stitched between domains.

Provisioning of overlay networks

- Most popular tool still seems to be email.
- Lengthy email trails, frequent miscommunication, delays when the right person isn't available
- We can and do try to do better, but we're not there yet
- Been watching presentations with a 'ping' in one window and a few button clicks on a browser for the ping to start working for longer than I care to remember...

Global Network Architecture

- Some time back (2017) GNA-G produced a set of documents:
 - <https://www.gna-g.net/resources/general-resources/>
 - Global Exchange Points
 - Multi-Layer Services
 - Operations
- Are those documents still used/referenced?
- Has the state-of-the-art moved on, and is that documented?

On with the presentations and discussion...