Developing an Architecture for Global SKA Communications

Formation of the Technical Working Group on the Global Network Architecture for SKA

Richard Hughes-Jones

SKA-NREN Forum Meeting 4

17 May 2023

Agenda

- Introduction to Global SKA Communications
- Input from Recent WLGC LHCOPN-LHCONE and SIG-NGN Meetings
- Terms of Reference for the
  Technical Working Group on the Global Network Architecture for SKA
- Current Status

Global Network & Paths of Interest to SKA

- Dedicated Primary (**red lines**) & Backup links (**yellow lines**) from both telescopes
- Use of the shared academic network (**blue lines**).
- 1 PetaByte/day pushed by SDP from each Telescope to the SRCs → 100 Gigabit/s
- SRC-SRC data transfers
- Costs based on 10 to 15 year IRU per 100 Gbit circuit projected to 2025 prices

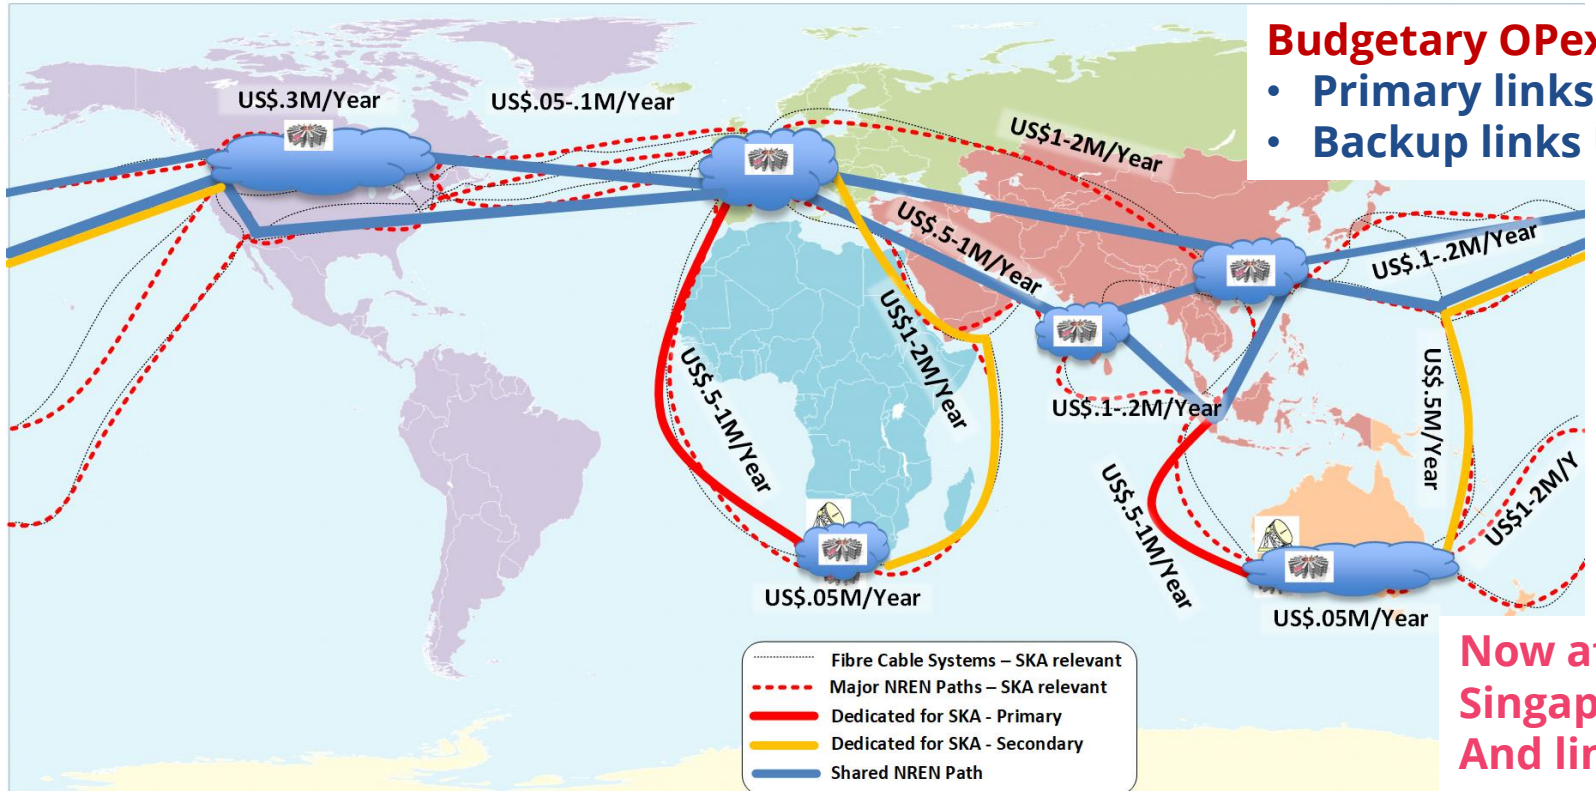

Budgetary OPex Costs (2020 estimates for 2025)

- **Primary links USD 1.7 – 2.3 M per year**
- **Backup links USD 2.3 – 3.3 M per year**

Now affordable paths Singapore – Europe direct, And links to Latin America

Global Connectivity for SKA: Telescopes to SRC and SRC to SRC

- Global VRF based overlay with peering linked over the shared academic network for bulk data transfers.

- Isolation of SKA traffic from other users
- Easier for NRENs to implement the routing, policies and monitoring
- SKA traffic can be engineered
  - Use specific paths & routes
- Layer 3 routing provides isolation
  - any network configuration issues
  - strictly limits broadcast storms
- Layer 3 will re-route traffic as long as there is an alternative network path
- Configuration actions have to be undertaken by the NREN and a Site to join the SKA VRF, which provides an extra layer of security

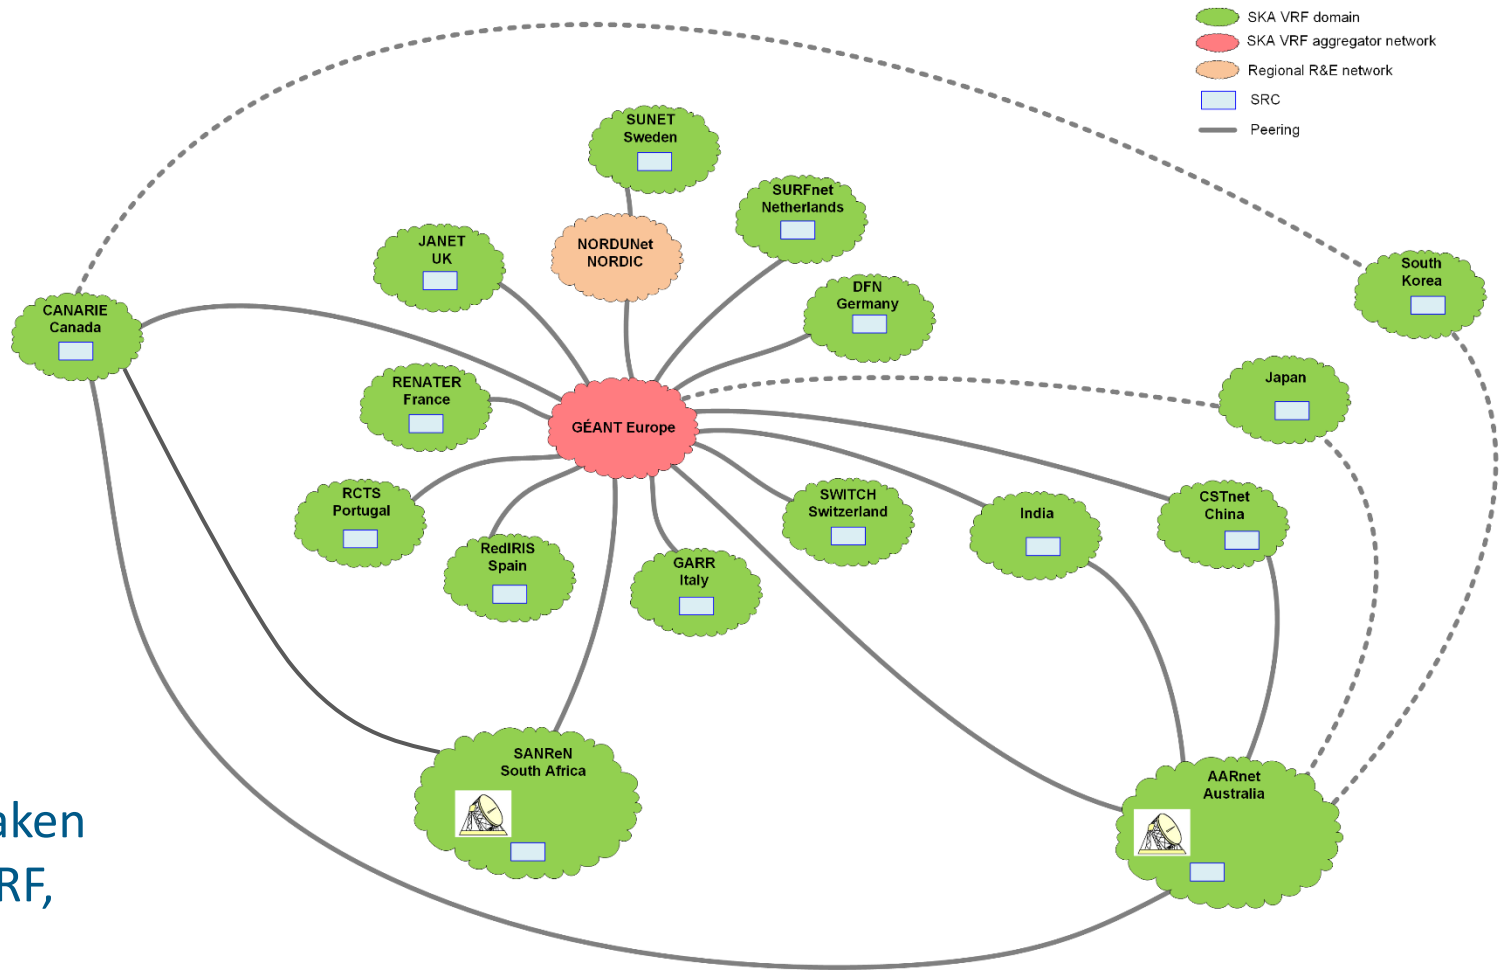

Global Paths of the Data Flows Pushed to the SRC 1 Replica

- Five flows on the submarine cable from Perth to Singapore .
- Then join the general purpose routed IP academic network.
- Single flows on the routes to Canada, China and India, Australia is local, and
  two 20 Gbit/s flows would be carried to London to reach SRCs in Europe and South Africa.

- Five flows on the submarine cable from Cape Town to London.
- Then join the general purpose routed IP academic network.
- Different submarine cables used to reach India and Australia, Europe is local, and
  two 20 Gbit/s flows cross the Atlantic to SRC in Canada and China.

SKA1-LOW Australia

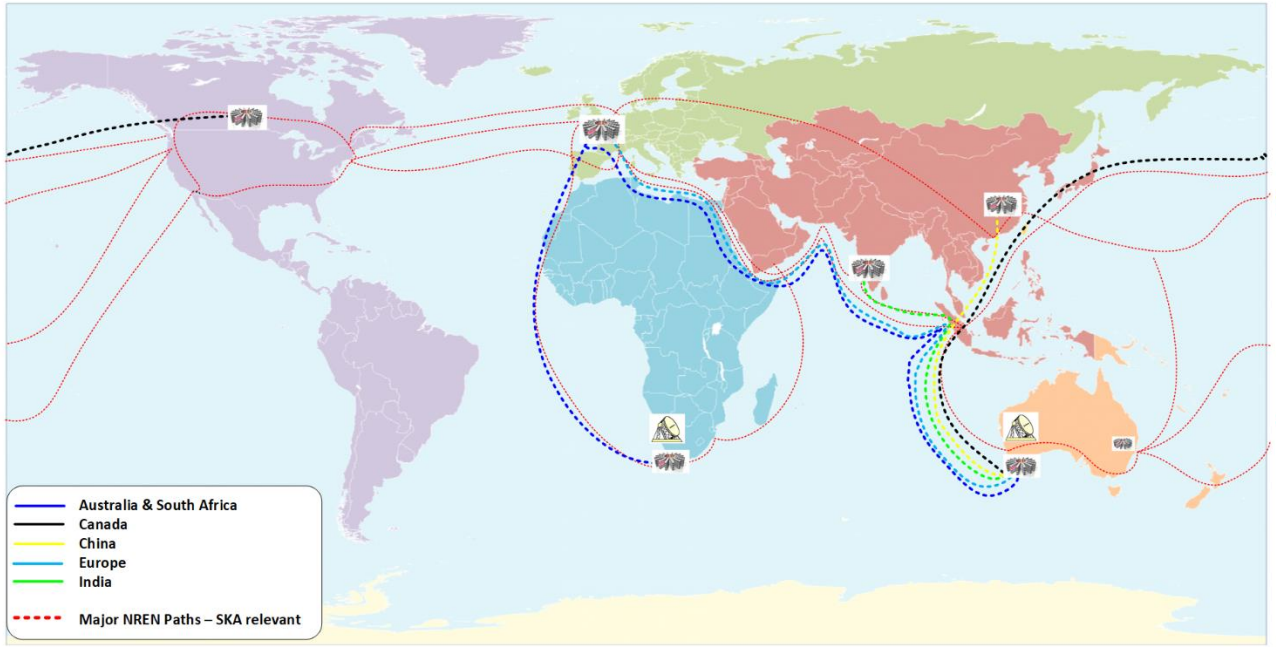

SKA1-MID South Africa

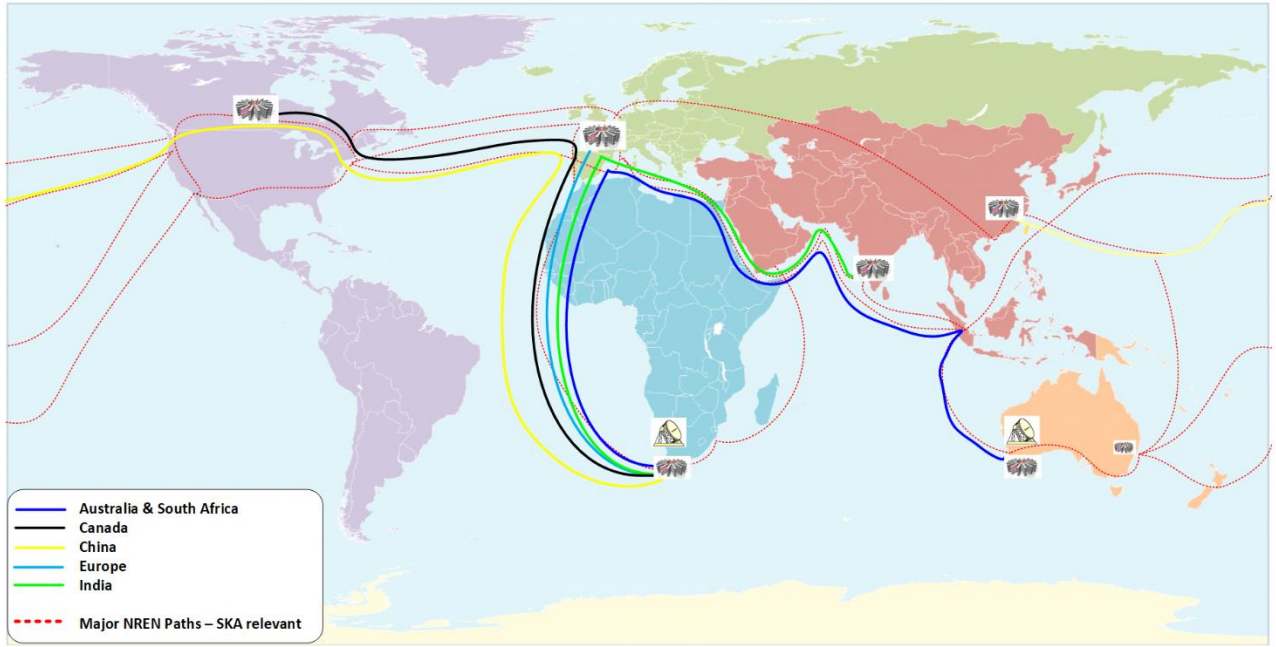

Models of Data Flows from the Telescopes to a SRC

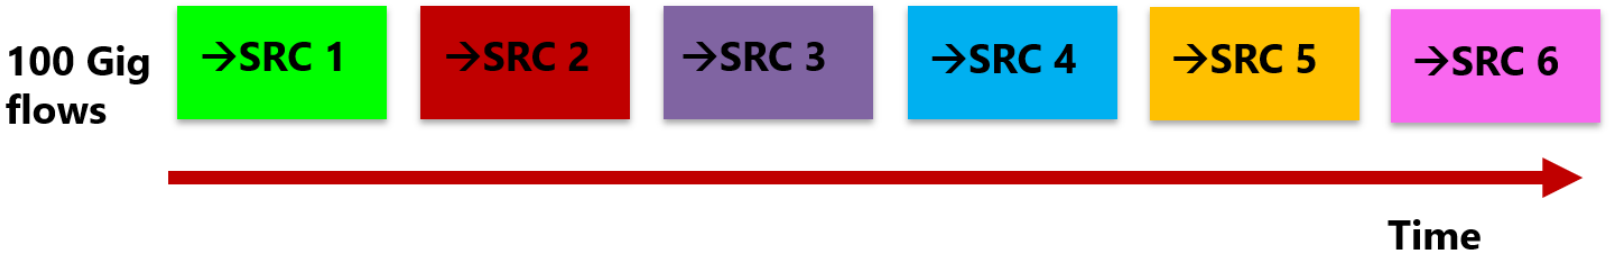

- Data sent to each SRC in turn.
- Requires clean 100Gbit/s paths to each SRC.
- Worst case the SRC would require a 200 Gbits/s link.

- Several flows to different SRC taking place in parallel.
- Operationally more realistic.
- Makes efficient use of the network to each SRC.
- WP4 demonstrated stable 28 Gbit/s flows – TCP limit.

The SDP push model gives the advantage of scheduling the use of the bandwidth on the telescope access link .

Input from Recent WLGC LHCOPN-LHCONE and SIG-NGN Meetings

IPv6 in WLCG

- HEPiX working Group started 2011 !
- Phase 1 - 2011-2016 - full analysis, investigations, testbed
- Phase 2 - 2017-2023 - deploy dual-stack storage on WLCG
- Phase 3 - 2019-onwards - plan for IPv6-only

- Sites have to configure both IPv4 and IPv6.
- Dual-stack has more security threat vectors.

- IPv6 a big help for “multiONE”
  (several LHCONE for different communities)

- ***Message to new research communities***
  - build on IPv6 from start

Good news (IPv6 on WLCG) after removing several “obstacles” during the last year

LHCOPN network (at CERN) ~95% IPv6 last 30 days

Storage - Tier-1 (100%) and Tier-2 (93%)

LHCONE network at CESNET (CZ)
- last 30 days
Ingress ~93% IPv6
Egress ~90% IPv6

Dave Kelsey CHEP2023

“Fixed” plots from CERN (>5 Mar 23)

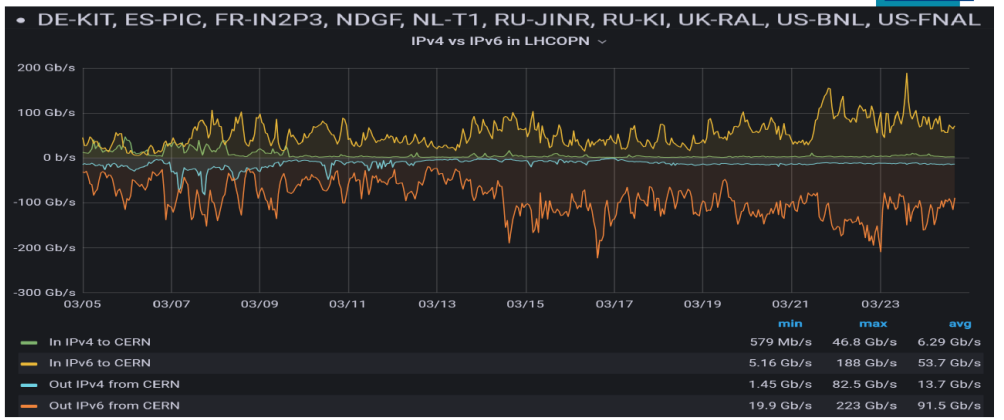

https://monit-grafana-open.cern.ch/d/cumEJJb4z/lhcopn-one-ipv6-vs-ipv4

Bruno Hoeft SIG-NGN April 2023

Use of Jumbo Frames

- Makes a big improvement to data transfers:
  - Throughput
  - Recovery time
- Concern about mixing Jumbo & 1500 Byte PCs
- Tests show 9000 to 1500 Byte MTU transfers do work for TCP.
- Care needed when configuring sites & PCs:
  - Path MTU discovery (PMTUD ICMP) needs to work
  - net.ipv4.tcp_mtu_probing=1 for IPv4

Network test data

•Iperf (Raul from Jisc)

Source	Destination	RTT	9000	1500
SURF (NL)	RNP (Brazil)	100ms	31 Gbit/s	20 Gbit/s
Jisc (London)	BNL (USA)	100ms	14 Gbit/s	6 Gbit/s
SURF (NL)	Jisc (London)	7.2 ms	23 Gbit/s	6 Gbit/s

•Tcpmon (Richard Hughes-Jones – Geant)

Source	Destination	RTT	9000	1500
London	Cambridge	3ms	37 Gbit/s	15.8 Gbit/s
London	AARnet	262ms	21 Gbit/s	3.4 Gbit/s

4

Chris Walker LHCOPN-LHCONE April 2023

Jisc

Throughput Europe to Australia RTT 262 ms

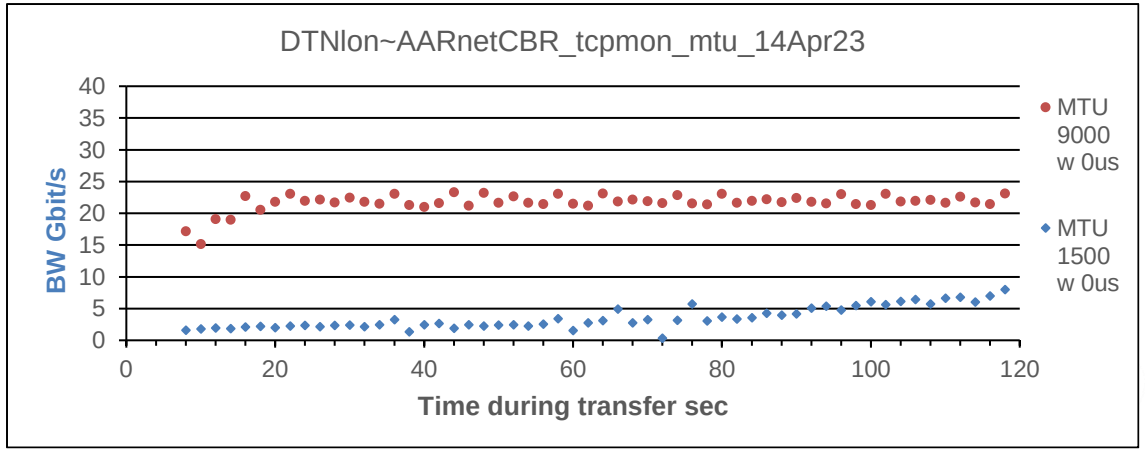

Terms of Reference for the Technical Working Group on the Global Network Architecture for SKA

Draft Terms of Reference: Goals and Connectivity

- The goals of the Technical Working Group on the Global Network Architecture for SKA are:
  - To document the agreed global connectivity required for the SKA community together with the associated operational parameters that will be used within this infrastructure.
  - To encourage early adoption.
  - Support new stakeholders as they connect.
  - Facilitate the exchange of performance evaluation.
- The network connectivity requirements include:
  - The transmission of the Observatory Data Products from the telescopes to the SKA Regional Centres.
  - The movement of data between SRCs e.g. Observatory and Advanced Data Products.
  - The provision of access to the federated distributed computing platform for the science users.

Draft Terms of Reference: Some Areas for Consideration

- Update on the available global academic connectivity.
- The use of peered VRFs overlay for the data transport.
- Proposed use of IPv6 only.
- The use of Jumbo 9000 Byte frames.
- Recommendations for DMZ and DTNs on SKA sites.
- Impact of the expected data placement and data flow models.
- Impact of the expected SRC locations.
- Consideration of SKA construction road map.
- Facilitation of performance evaluations using simulated SRC data flows.

Draft Terms of Reference: The Key is Collaboration

- The end-to-end network connectivity will involve several different domains such as:
  - The SRC compute sites
  - The NRENs
  - The Regional Networks
  - The inter-continental links

- The detailed design, implementation, and operation of each domain is the responsibility of that domain.

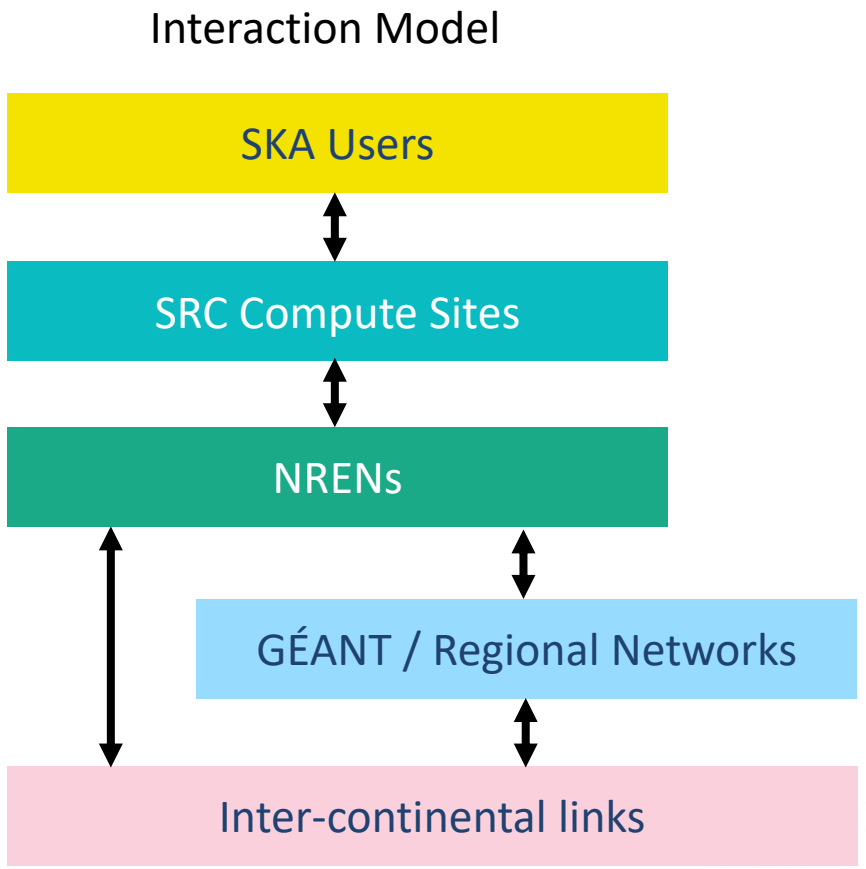

Technical Working Group: Status

- There have been positive responses to join the Network Architecture WG from:
  - AARNet
  - CANARIE
  - GÉANT
  - STFC RAL
  - TENET
- The timescale is to have first recommendations by September 2023.
- The outcomes will be presented in SKA-NREN Forum meetings.

- Kickoff meeting by end of June 2023

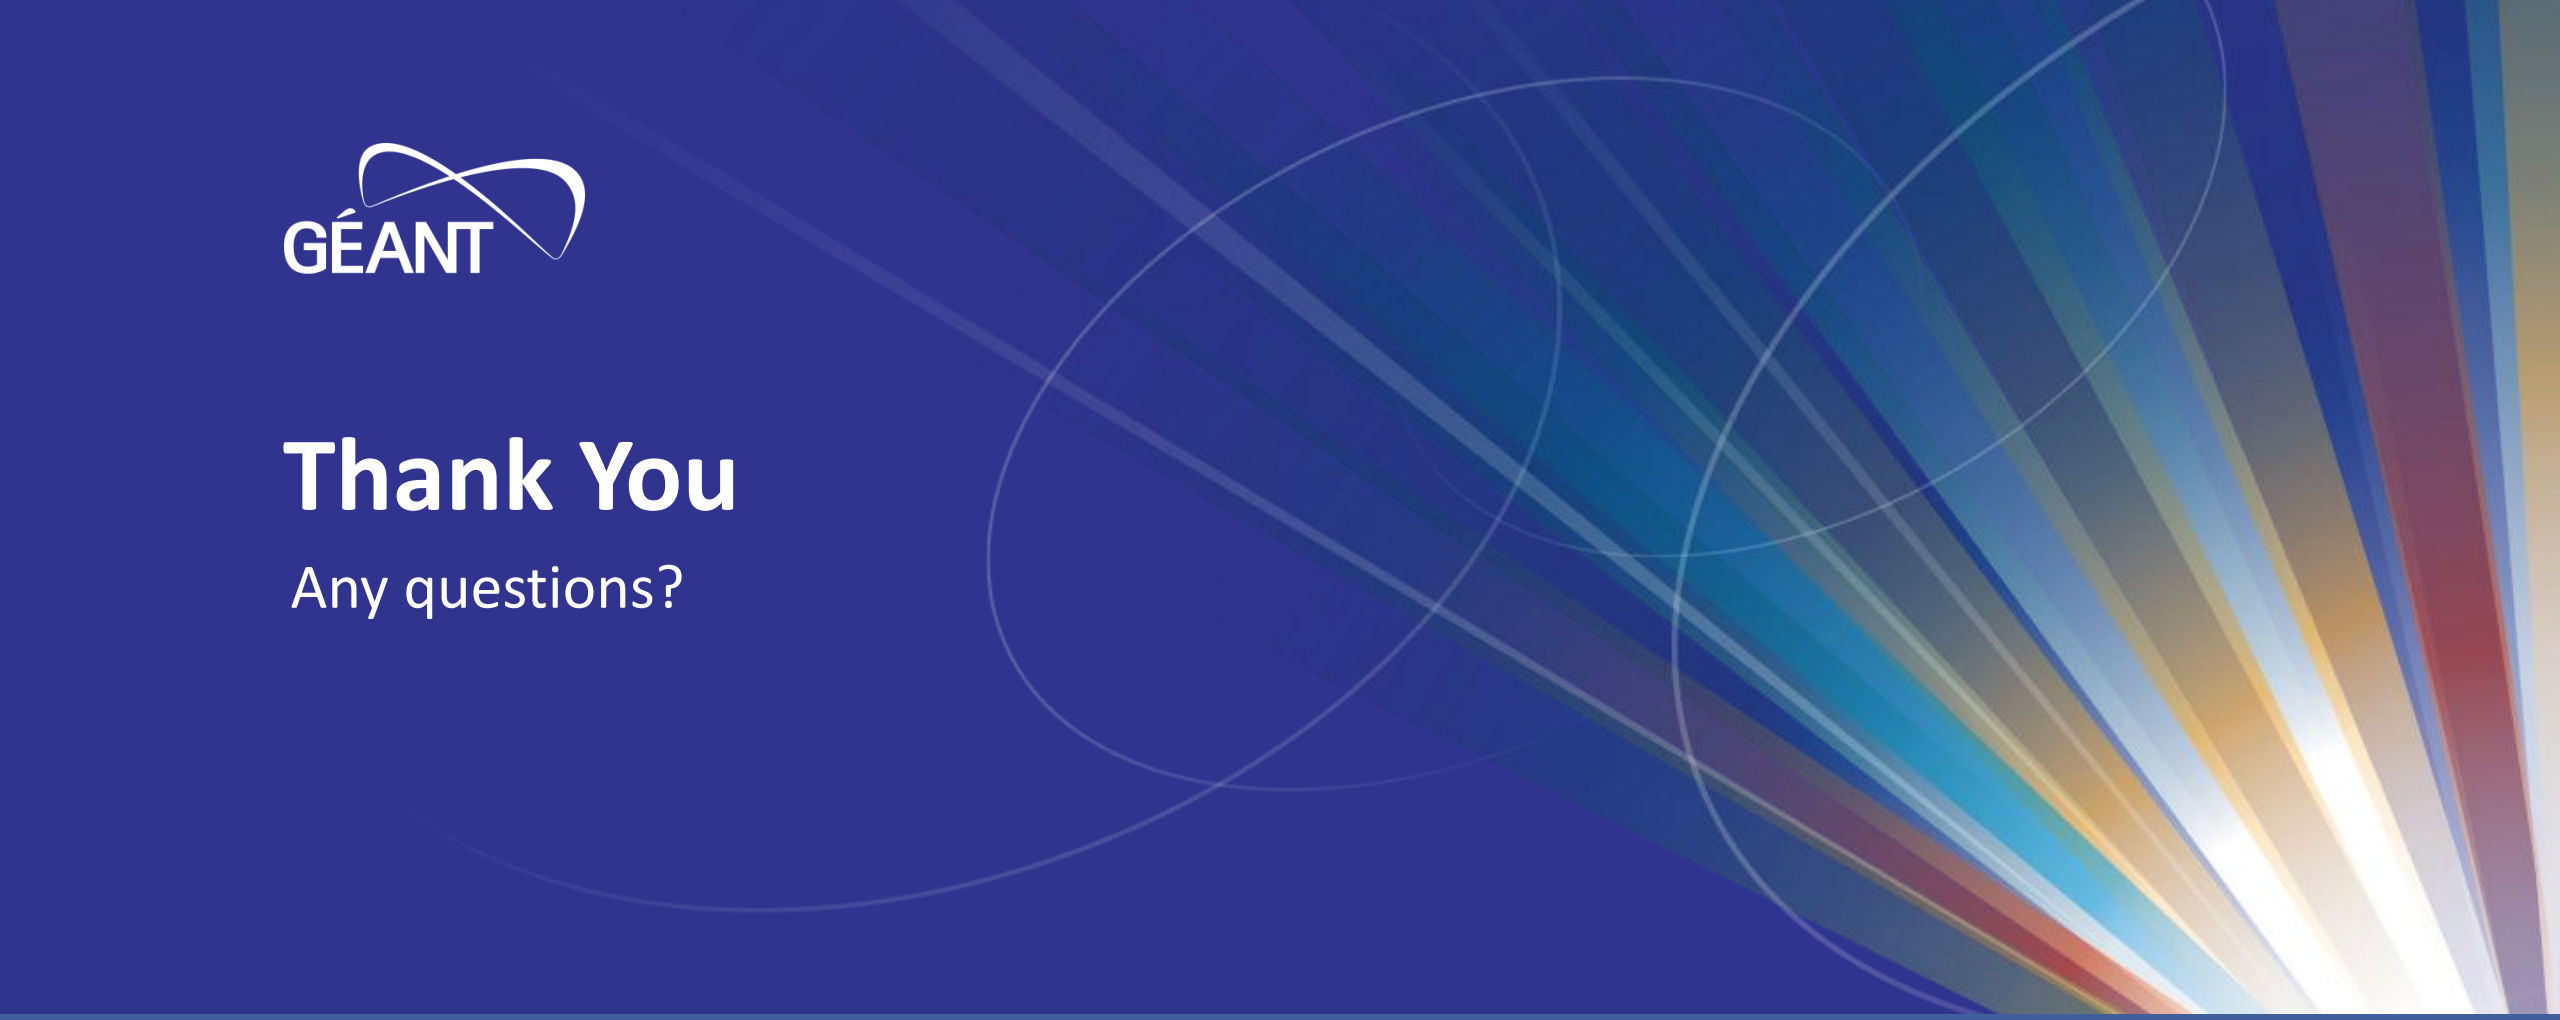

Thank You

Any questions?

www.geant.org

Co-funded by
the European Union