SKA Computing Challenges

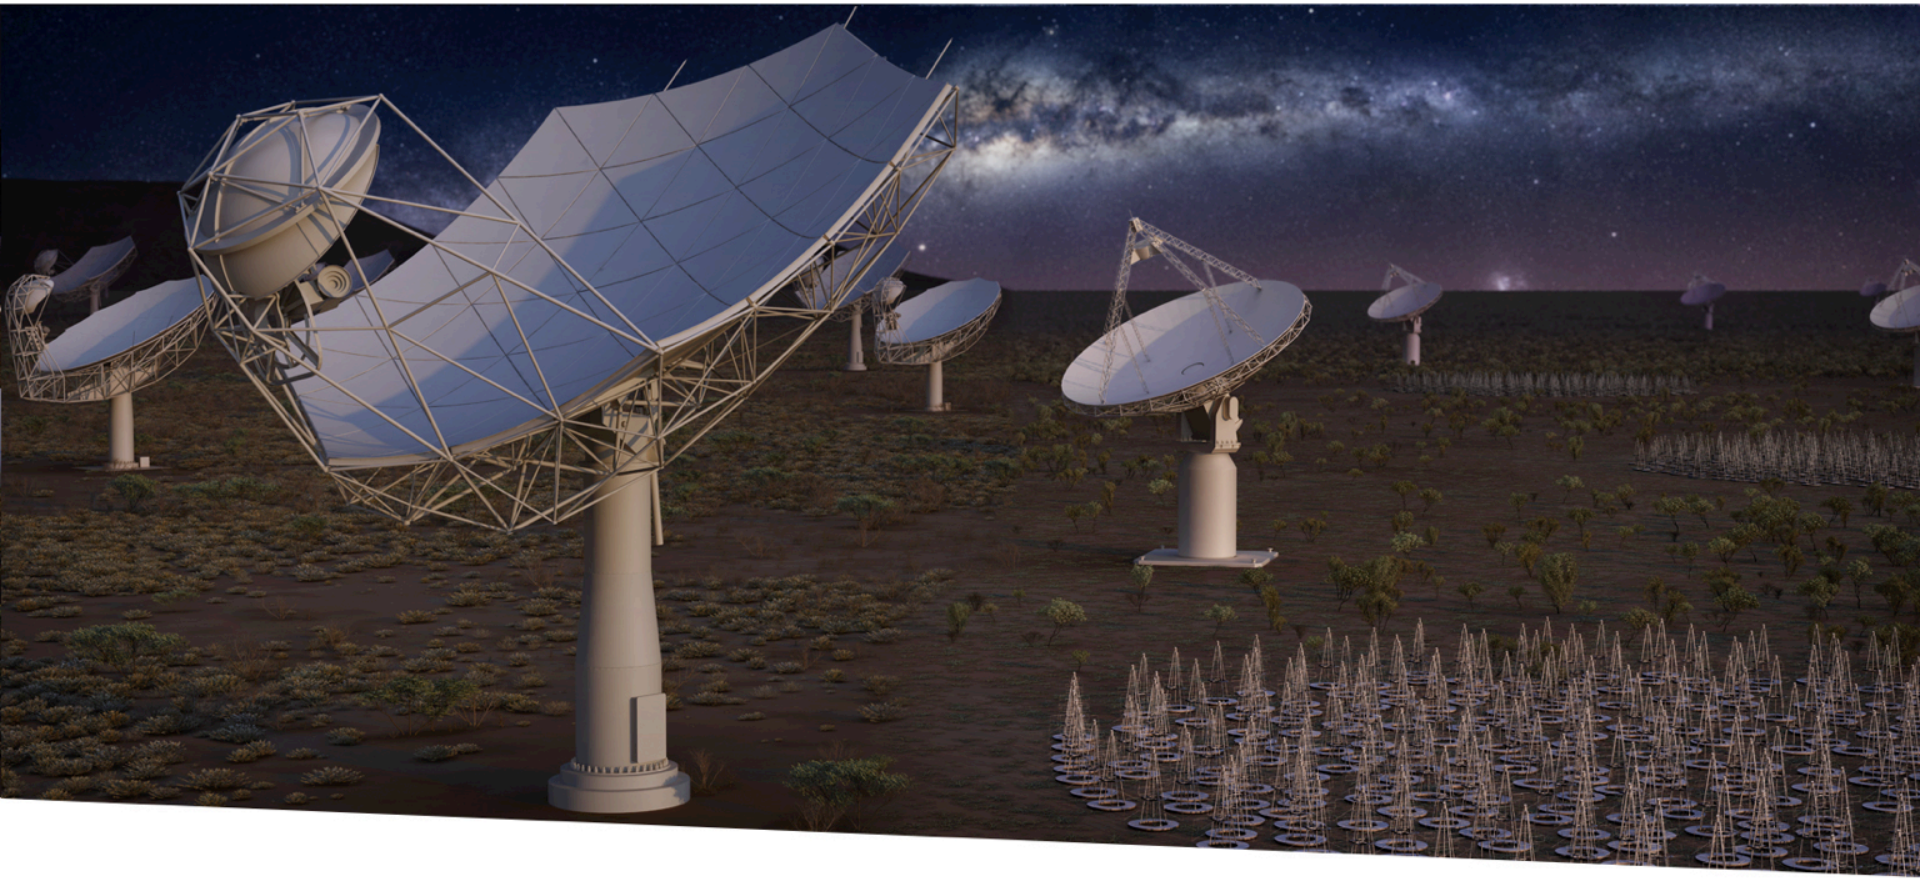

SQUARE KILOMETRE ARRAY

Exploring the Universe with the world's largest radio telescope

Miles Deegan, SKAO

Maurizio Miccolis, SKAO

SKA Status Update

5 years into pre-construction design phase

Design Consortia are going through CDRs. This is followed by system-level CDR

At the same time we are forming an IGO (like CERN, EMBL, ECMWF, …) and (hopefully) acquiring the funds to build the observatory from 2020 onwards

12 member countries to date. Spain and France are recent joiners. Hope to announce more shortly

Dataflow and processing overview

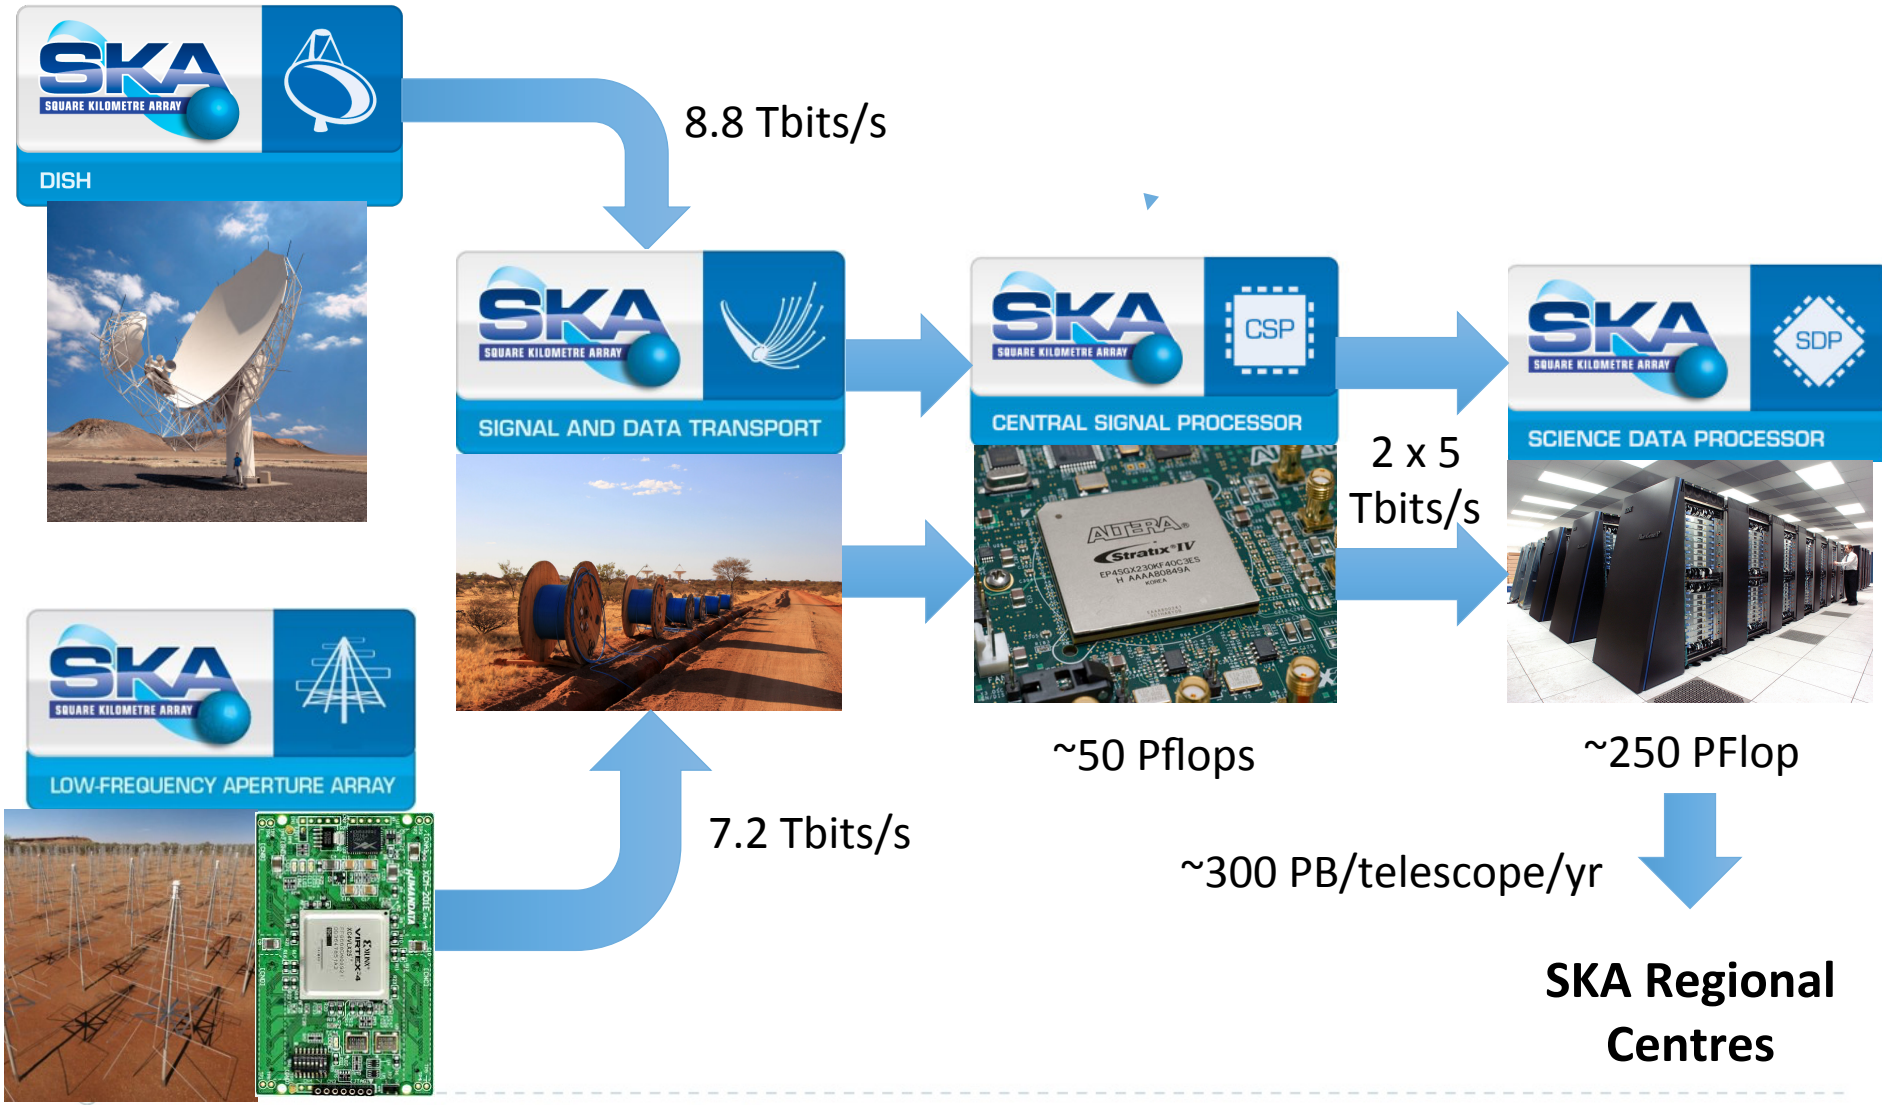

Not classical HPC…

SDP will have HPC, HPDA and cloud-like features

This is **data-intensive computing** with multiple storage tiers, complex workflows, scheduling and data management and lifecycle issues

Pipelines will likely have low operational intensity - memory bandwidth bound

A Top500/HPL oriented machine won't cut it. HPCG a more interesting benchmark: >80+% efficiency vs. <5%

Software Challenges

We will spend around 100M Euros on software

SKA will be a complex as well as expensive software project. We need to:

1) Develop a new codebase (some reuse possible)
2) Support and evolve it for 50 years (operational lifetime)
3) All with a globally distributed development and operations team

SKA is using the Scaled Agile Framework (SAFe) methodology and has produced software development standards. Trying to adopt Lean, Agile principles across the organisation. Cultural as well as technical change needed

General HPC Challenges 1: Power

- Current Exascale roadmap (US)
  - indicates ~40 MW for ExaFlop by 2022/3 (enough to power a small town)
  - Up from 20-30 MW
- Some estimates of 60 MW coming from other projects
- Technical problem: a lot of heat to get rid of
- Financial problem: SKA OpEx budget will be very stretched if vendors are not addressed

General HPC Challenges 2: Cost

- Capital investment
  - One time expectation of capital outlay of around $250M for early exascale systems
  - It could now be $500M for systems (CORAL2)
  - (and budgets include a lot of NRE)
  - Recent GPU and memory cost trends have been upward
    - The bursting of cryptocurrency bubble might bring these back down

Exascale and Co-design

Exascale Race/Technologies

IDC-Projected Exascale Dates and Suppliers

U.S.

- Sustained ES: 2023
- Peak ES: 2021
- Vendors: U.S.
- Processors: U.S.
- Initiatives: NSCI/ECP
- Cost: $300-500M per system, plus heavy R&D investments

EU

- Sustained ES: 2023-24
- Peak ES: 2021
- Vendors: U.S., Europe
- Processors: U.S., ARM
- Initiatives: PRACE, ETP4HPC
- Cost: $300-$350 per system, plus heavy R&D investments

China

- Sustained ES: 2023
- Peak ES: 2020
- Vendors: Chinese
- Processors: Chinese (plus U.S.?)
- 13th 5-Year Plan
- Cost: $350-500M per system, plus heavy R&D

Japan

- Sustained ES: 2023-24
- Peak ES: Not planned
- Vendors: Japanese
- Processors: Japanese
- Cost: $600-850M, this includes both 1 system and the R&D costs…will also do many smaller size systems

IDC-Projected Exascale Investment Levels (In Addition to System Purchases)

U.S.

- $1 to $2 billion a year in R&D (including NRE)
- Investments by both governments & vendors
- Plans are to purchase multiple exascale systems

EU

- About 5 billion euros in total
- Investments in multiple exascale and pre-exascale systems
- Investments mostly by country governments with a little from the EU

China

- Over $1billion a year in R&D
- Investments by both governments & vendors
- Plans are to purchases multiple exascale systems each year
- Already investing in 3 pre-exascale systems by 2017/18

Japan

- Planned investment of just over $1billion* (over 5 years) for both the R&D and purchase of 1 exascale system
- To be followed by a number of smaller systems ~$100M to $150M each
- Creating a new processor and a new software environment

** Note that this includes both the system and R&D*

SKA and co-design

- Project very keen to engage with exascale community
  - Efforts are underway
- In the EU there are significant investments in this area
  - EuroHPC – joint undertaking with 1.5B Euro budget
  - European Processor Initiative
  - Use ARM and RISC-V technologies
  - Look to develop not just exascale systems but technology for automated vehicles as well – BMW involvement
  - 23 partners including CINECA and U. Pisa
- SKA use cases to help drive this working with partners in member countries?

SRCs Overview

- SRCs are needed to produce advanced data products from observatory data products coming out of SDP
  - ~ 300 PB/annum/telescope (current estimate)
  - Conceptually similar to the WLCG – CERN's WW compute grid with its tiered structure and we work with CERN
- Currently not part of SKA, or funded by SKAO, but:
  - Essential to generate science
  - Coordinated with assistance from SKAO and accredited with SKAO
- Support regional astronomers with their data processing
- Act as a centre for domain expertise
- Amount of infrastructure required?
  - TBD, but expect a lot of major HPC centres to be involved globally plus use of clouds – national/regional publicly funded science clouds, plus public (private-sector) clouds.

Regional Centre Network

- Need 2x 100 Gb circuits to distribute ~300 PB/annum
- ***Guesstimate*** of Regional Centre locations
  - No decisions have been taken

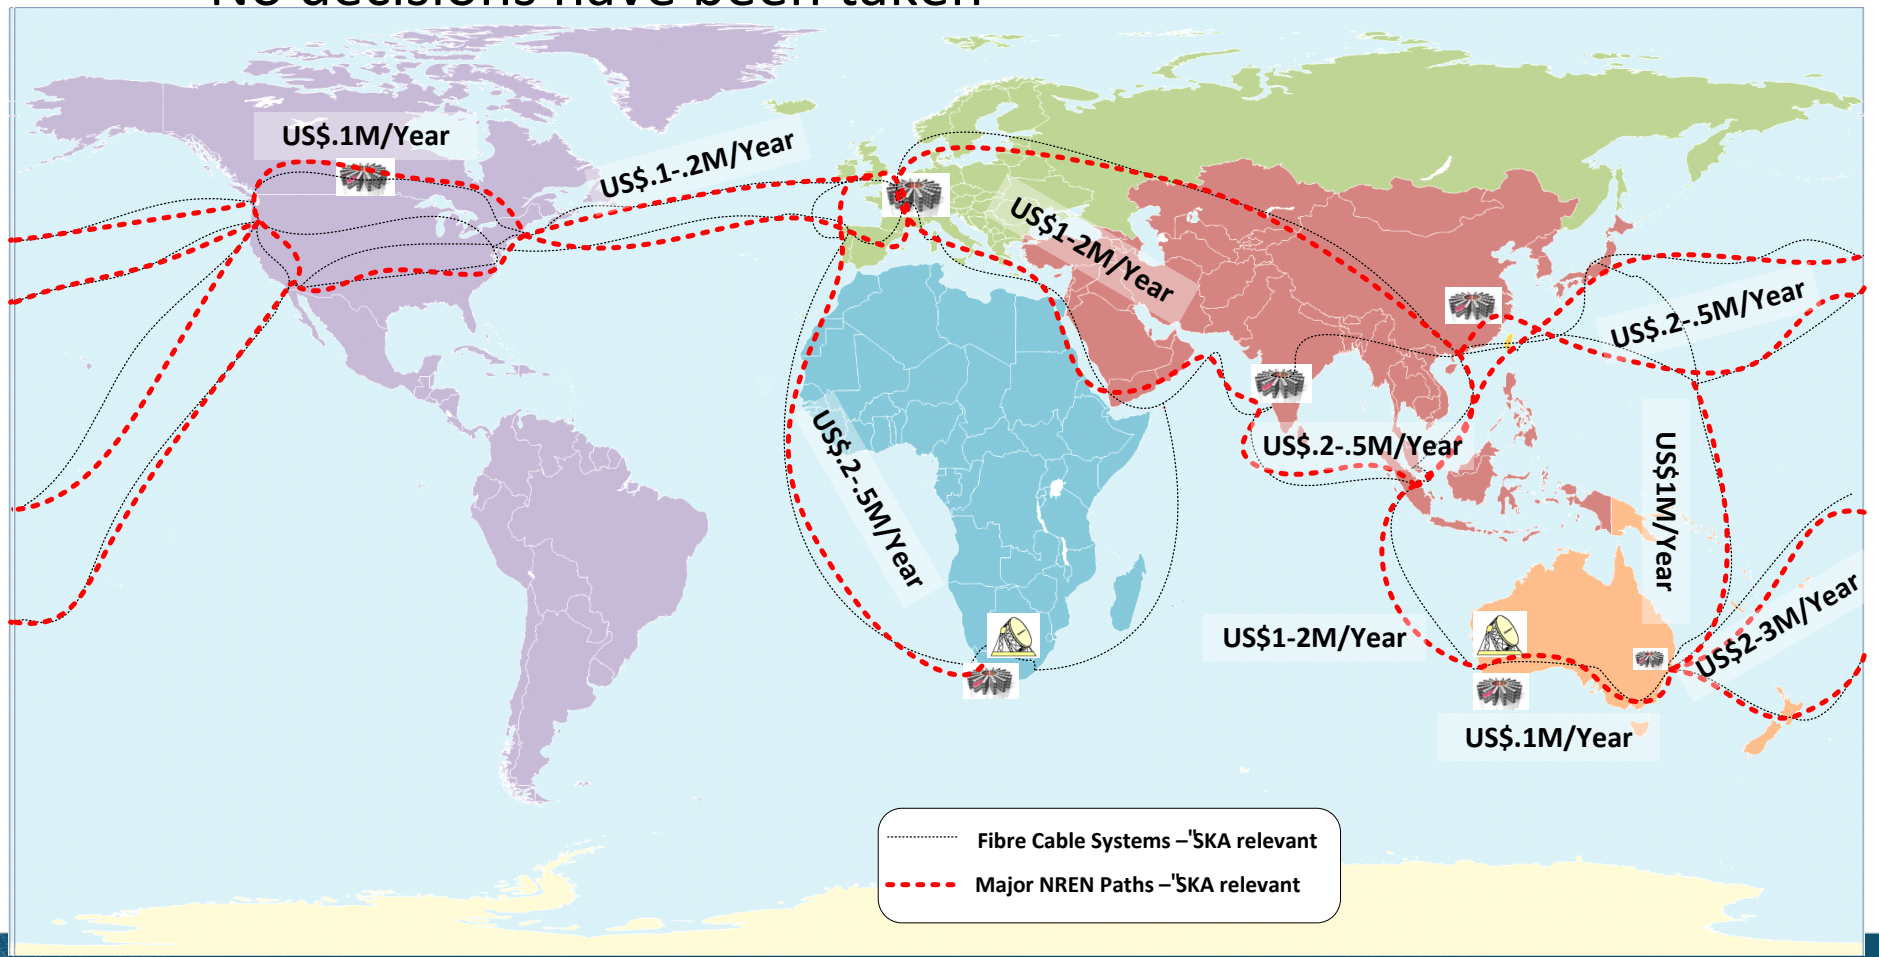

SRC Overview cont.

- Initial set of requirements has been generated
- Regions are already investing in design and prototyping efforts:
  - AENEAS H2020 project
  - ERIDANUS – Sino-Australian effort
  - Canada's NRC
- SKA is a key ESFRI in H2020 European Open Science Cloud (EOSC) plans
  - Working with ESO, CTA, CERN, + others in astronomy, astro-particle and particle physics on the ESCAPE project
  - ESCAPE will do things such as prototype data lakes for EOSC
  - Scope for involvement of Italian HPC centres

Data Challenges

- Effort based at SKA HQ
- A few themes:
  - Improving calibration and imaging algorithms
    - Robert Braun's talk mentioned a challenge that's just been launched
  - Testing SDP software as it evolves
  - Prototyping SRC workloads
  - Data management and ensuring 100s PBs can be moved efficiently across the globe
- Keen to engage the community
  - Science WGs, HPC centres, NRENs

Thank you

- Any questions?