International Centre for Radio Astronomy Research

SKA Data Flow

An observing project oriented approach to the SKA data flow design

Andreas Wicenec

2010-10-26

Random thoughts

- Data flow does not stop after the correlator
- Science with ASKAP, Meerkat, MWA, SKA is impossible without HPC
- Observing with ASKAP, Meerkat, MWA, SKA is impossible without HPC

➢ New era for both HPC and observational Astronomy

➢ SKA will need the fastest supercomputer and the fastest network on the planet just to work

➢ Level of integration of computing and observatory enters a new dimension

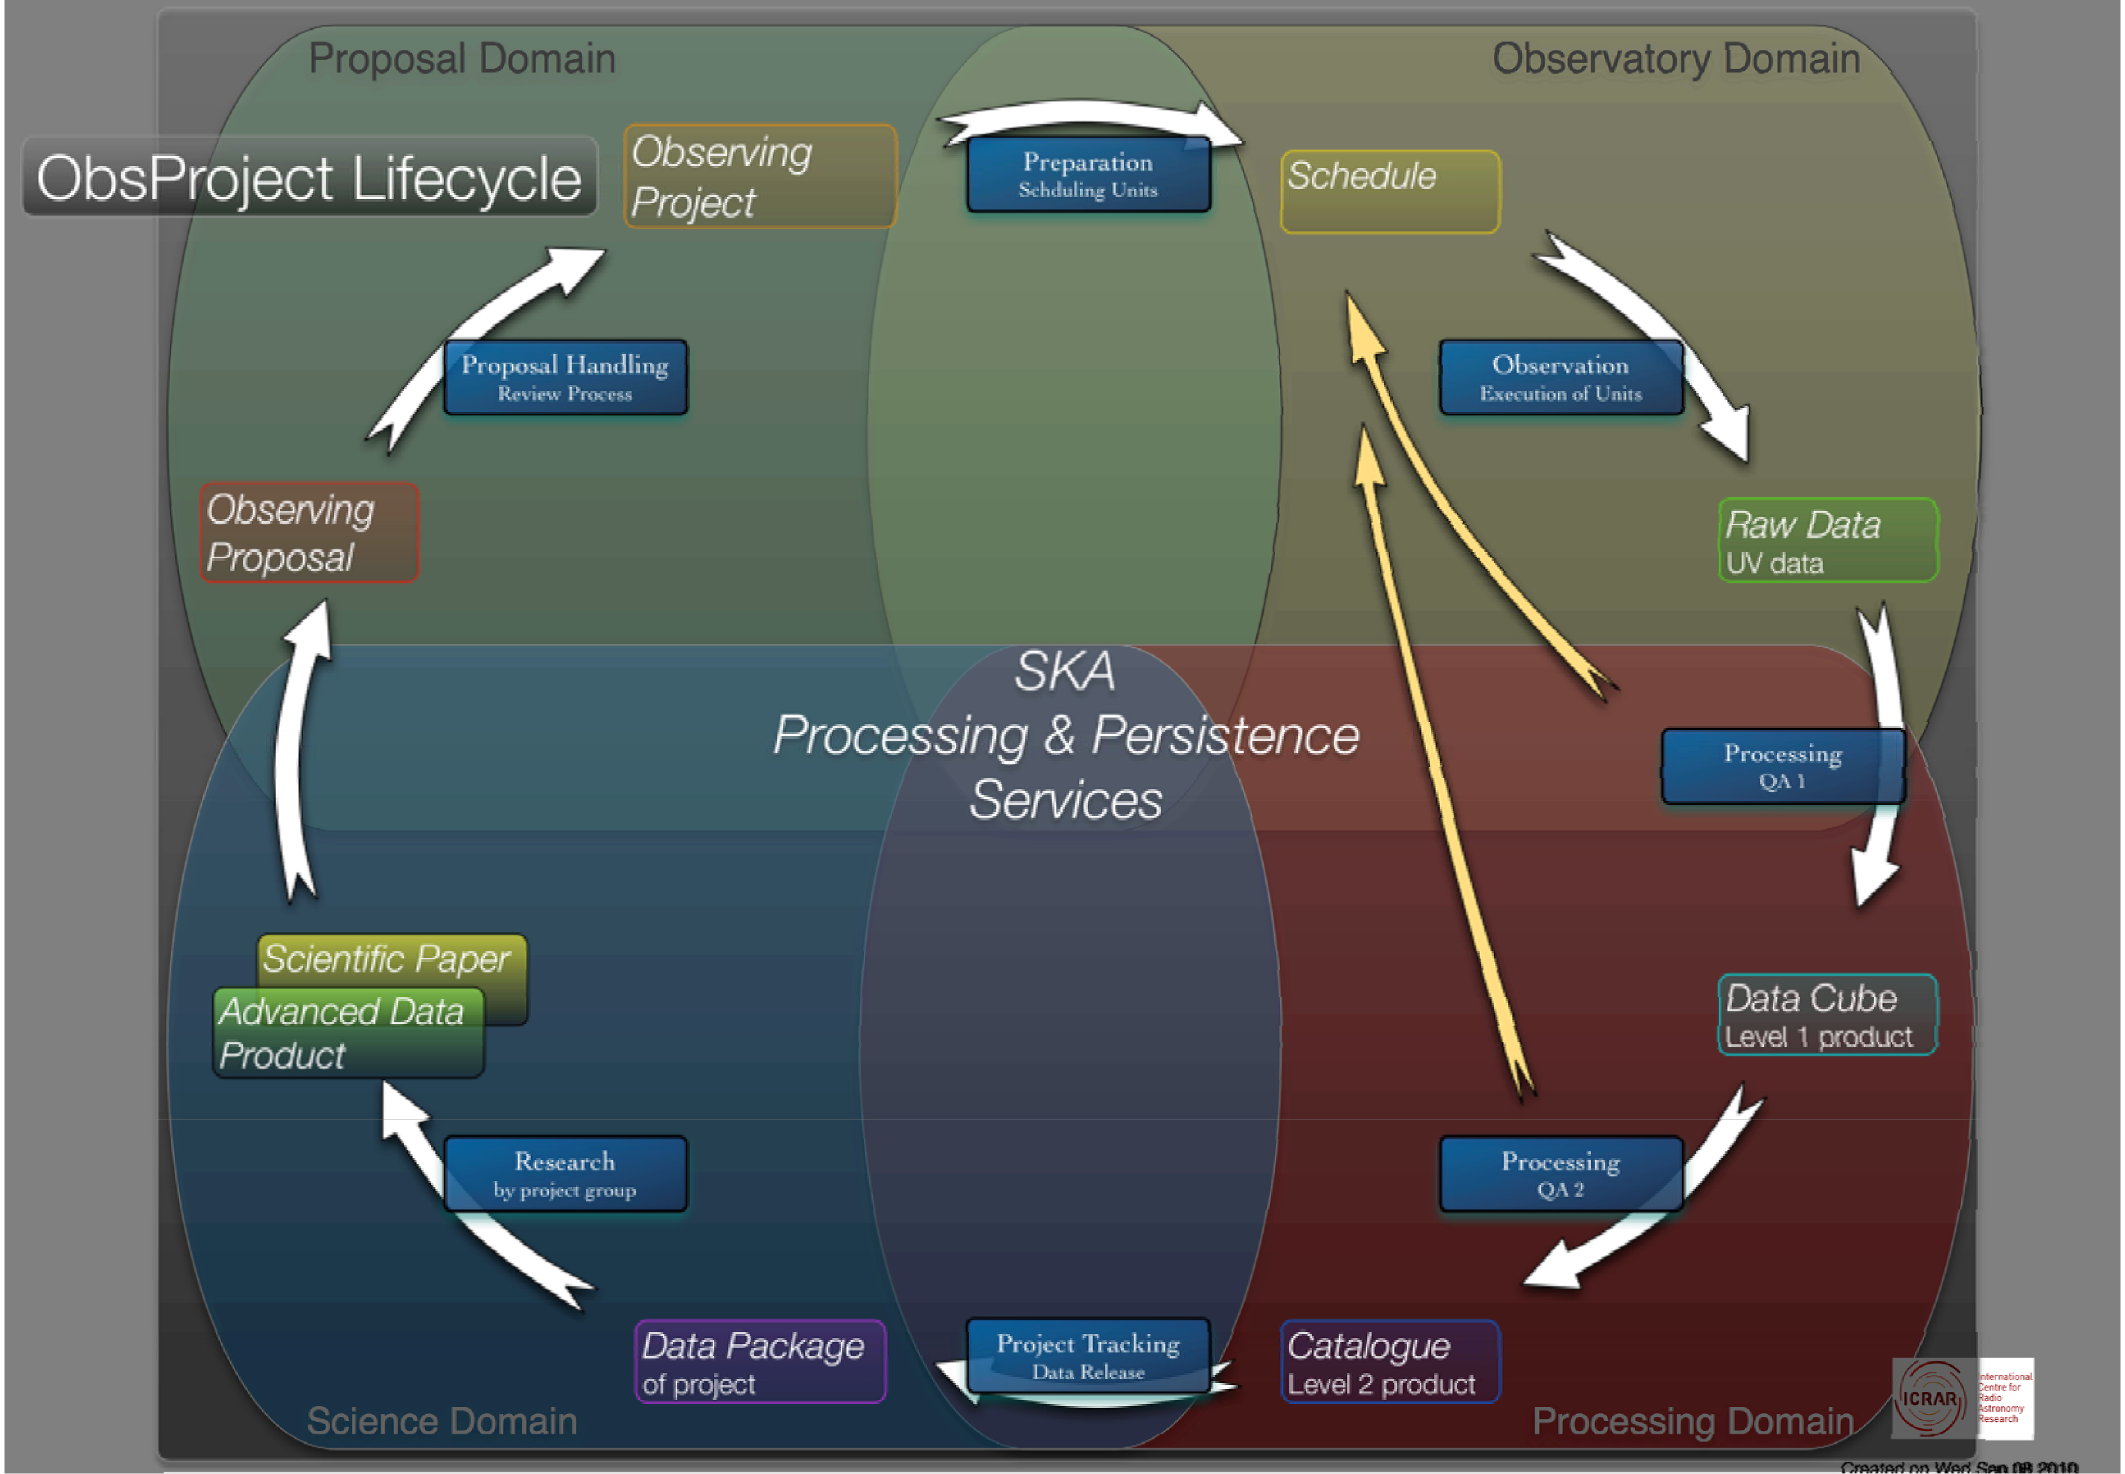
ObsProject Lifecycle
Proposal Domain
Observatory Domain
Science Domain
Processing Domain
Observing Project
Preparation
Scheduling Units
Schedule
Proposal Handling
Review Process
Observation
Execution of Units
Observing Proposal
Raw Data
UV data
SKA
Processing & Persistence
Services
Processing
QA 1
Scientific Paper
Advanced Data Product
Data Cube
Level 1 product
Research
by project group
Processing
QA 2
Data Package
of project
Project Tracking
Data Release
Catalogue
Level 2 product
ICRAR
International Centre for Radio Astronomy Research
Created on Wed Sep 08 2010

Data Flow Design

- Data flow needs to reflect operational requirements, such as:
  - Observatory operations plan including data policies.
  - Location, size and type of data sources and sinks.
  - Usage patterns of data.
  - User profiles and access patterns.
  - Subsystem requirements for temporary and persistent data storage and access.
  - Auxiliary data requirements.

Examples

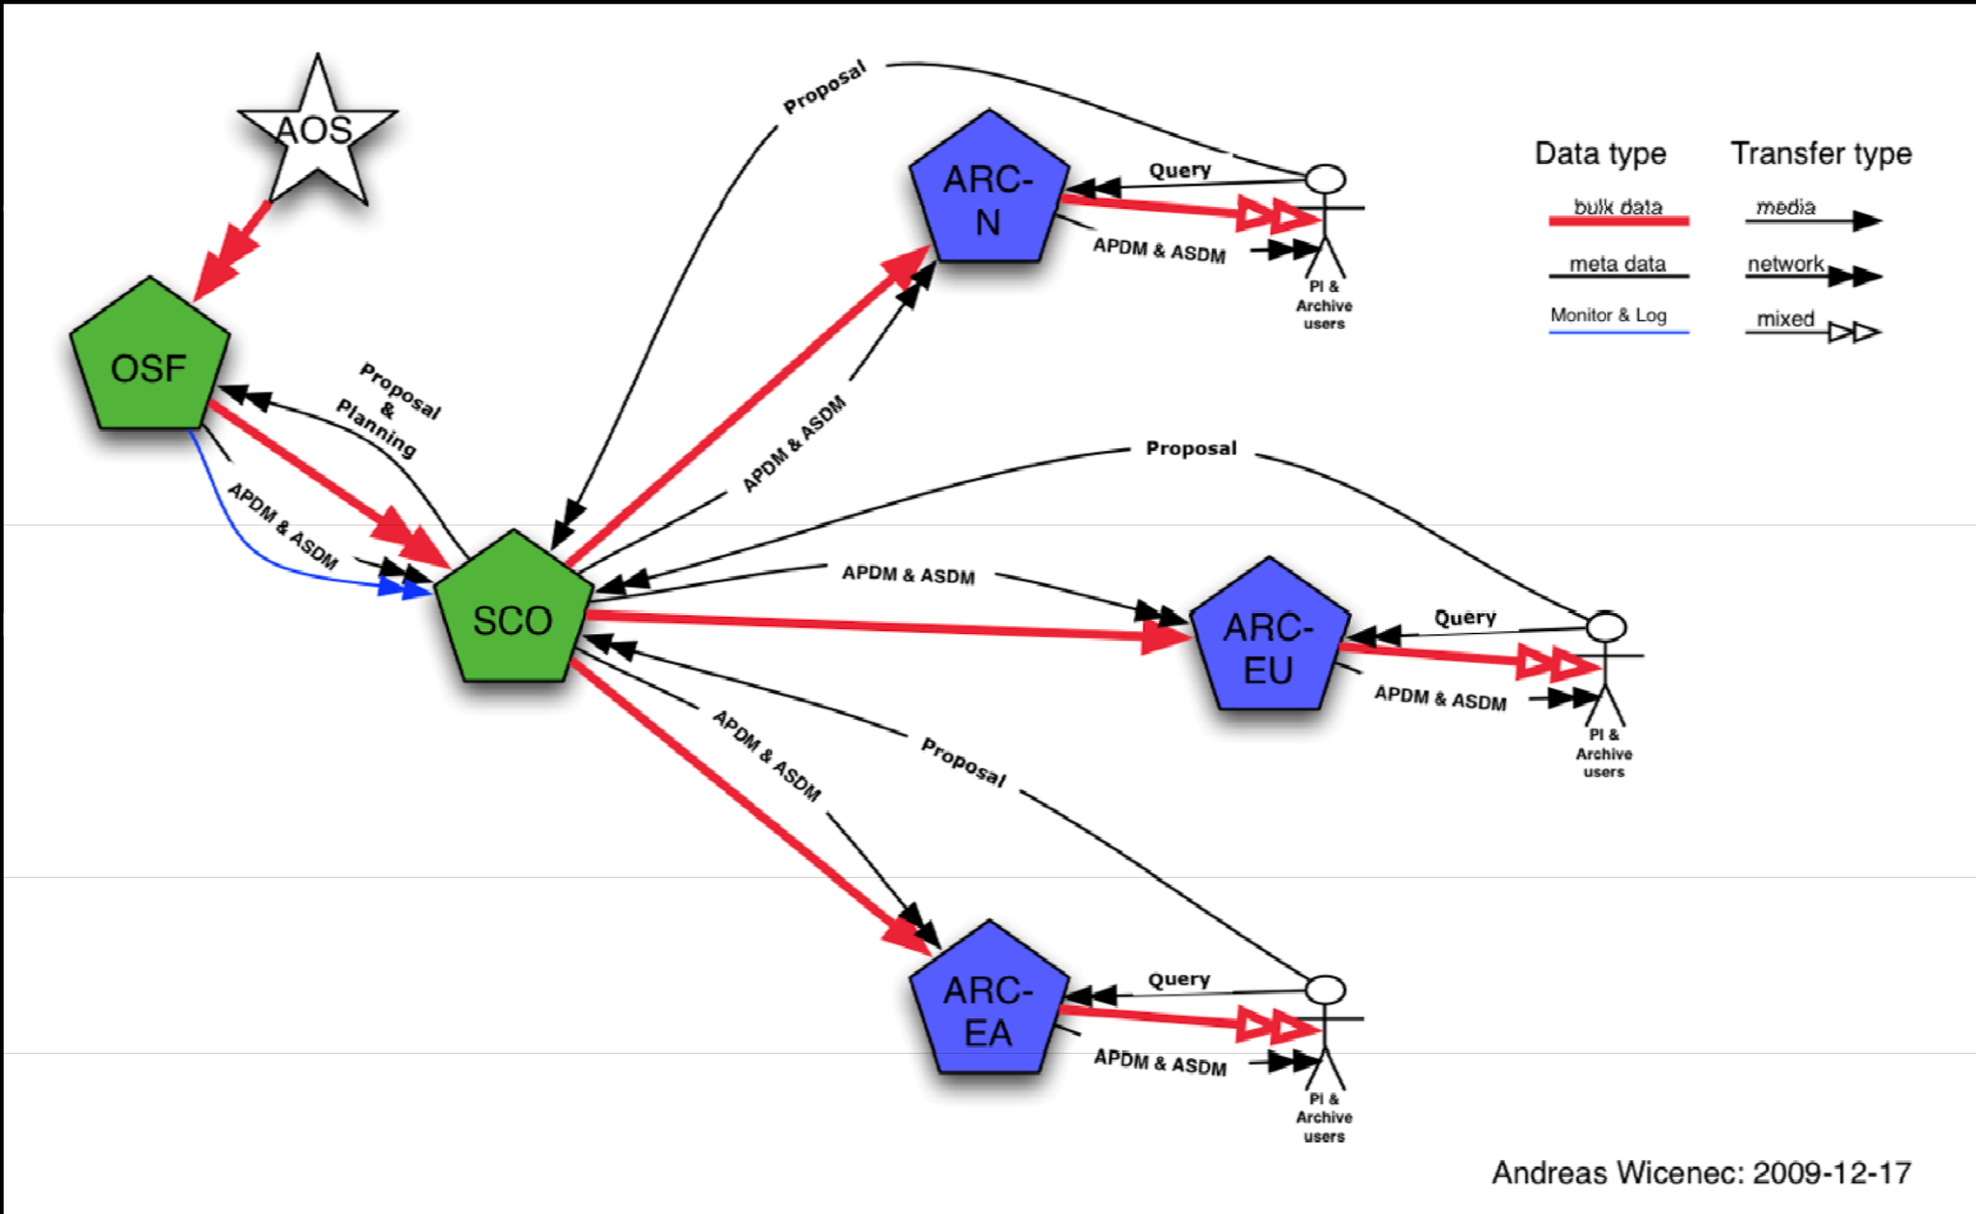

AOS
OSF
SCO
ARC-N
ARC-EU
ARC-EA
Proposal
Proposal & Planning
APDM & ASDM
Query
PI & Archive users
Data type
bulk data
meta data
Monitor & Log
Transfer type
media
network
mixed
Andreas Wicenec: 2009-12-17

ESO Data Flow Diagram

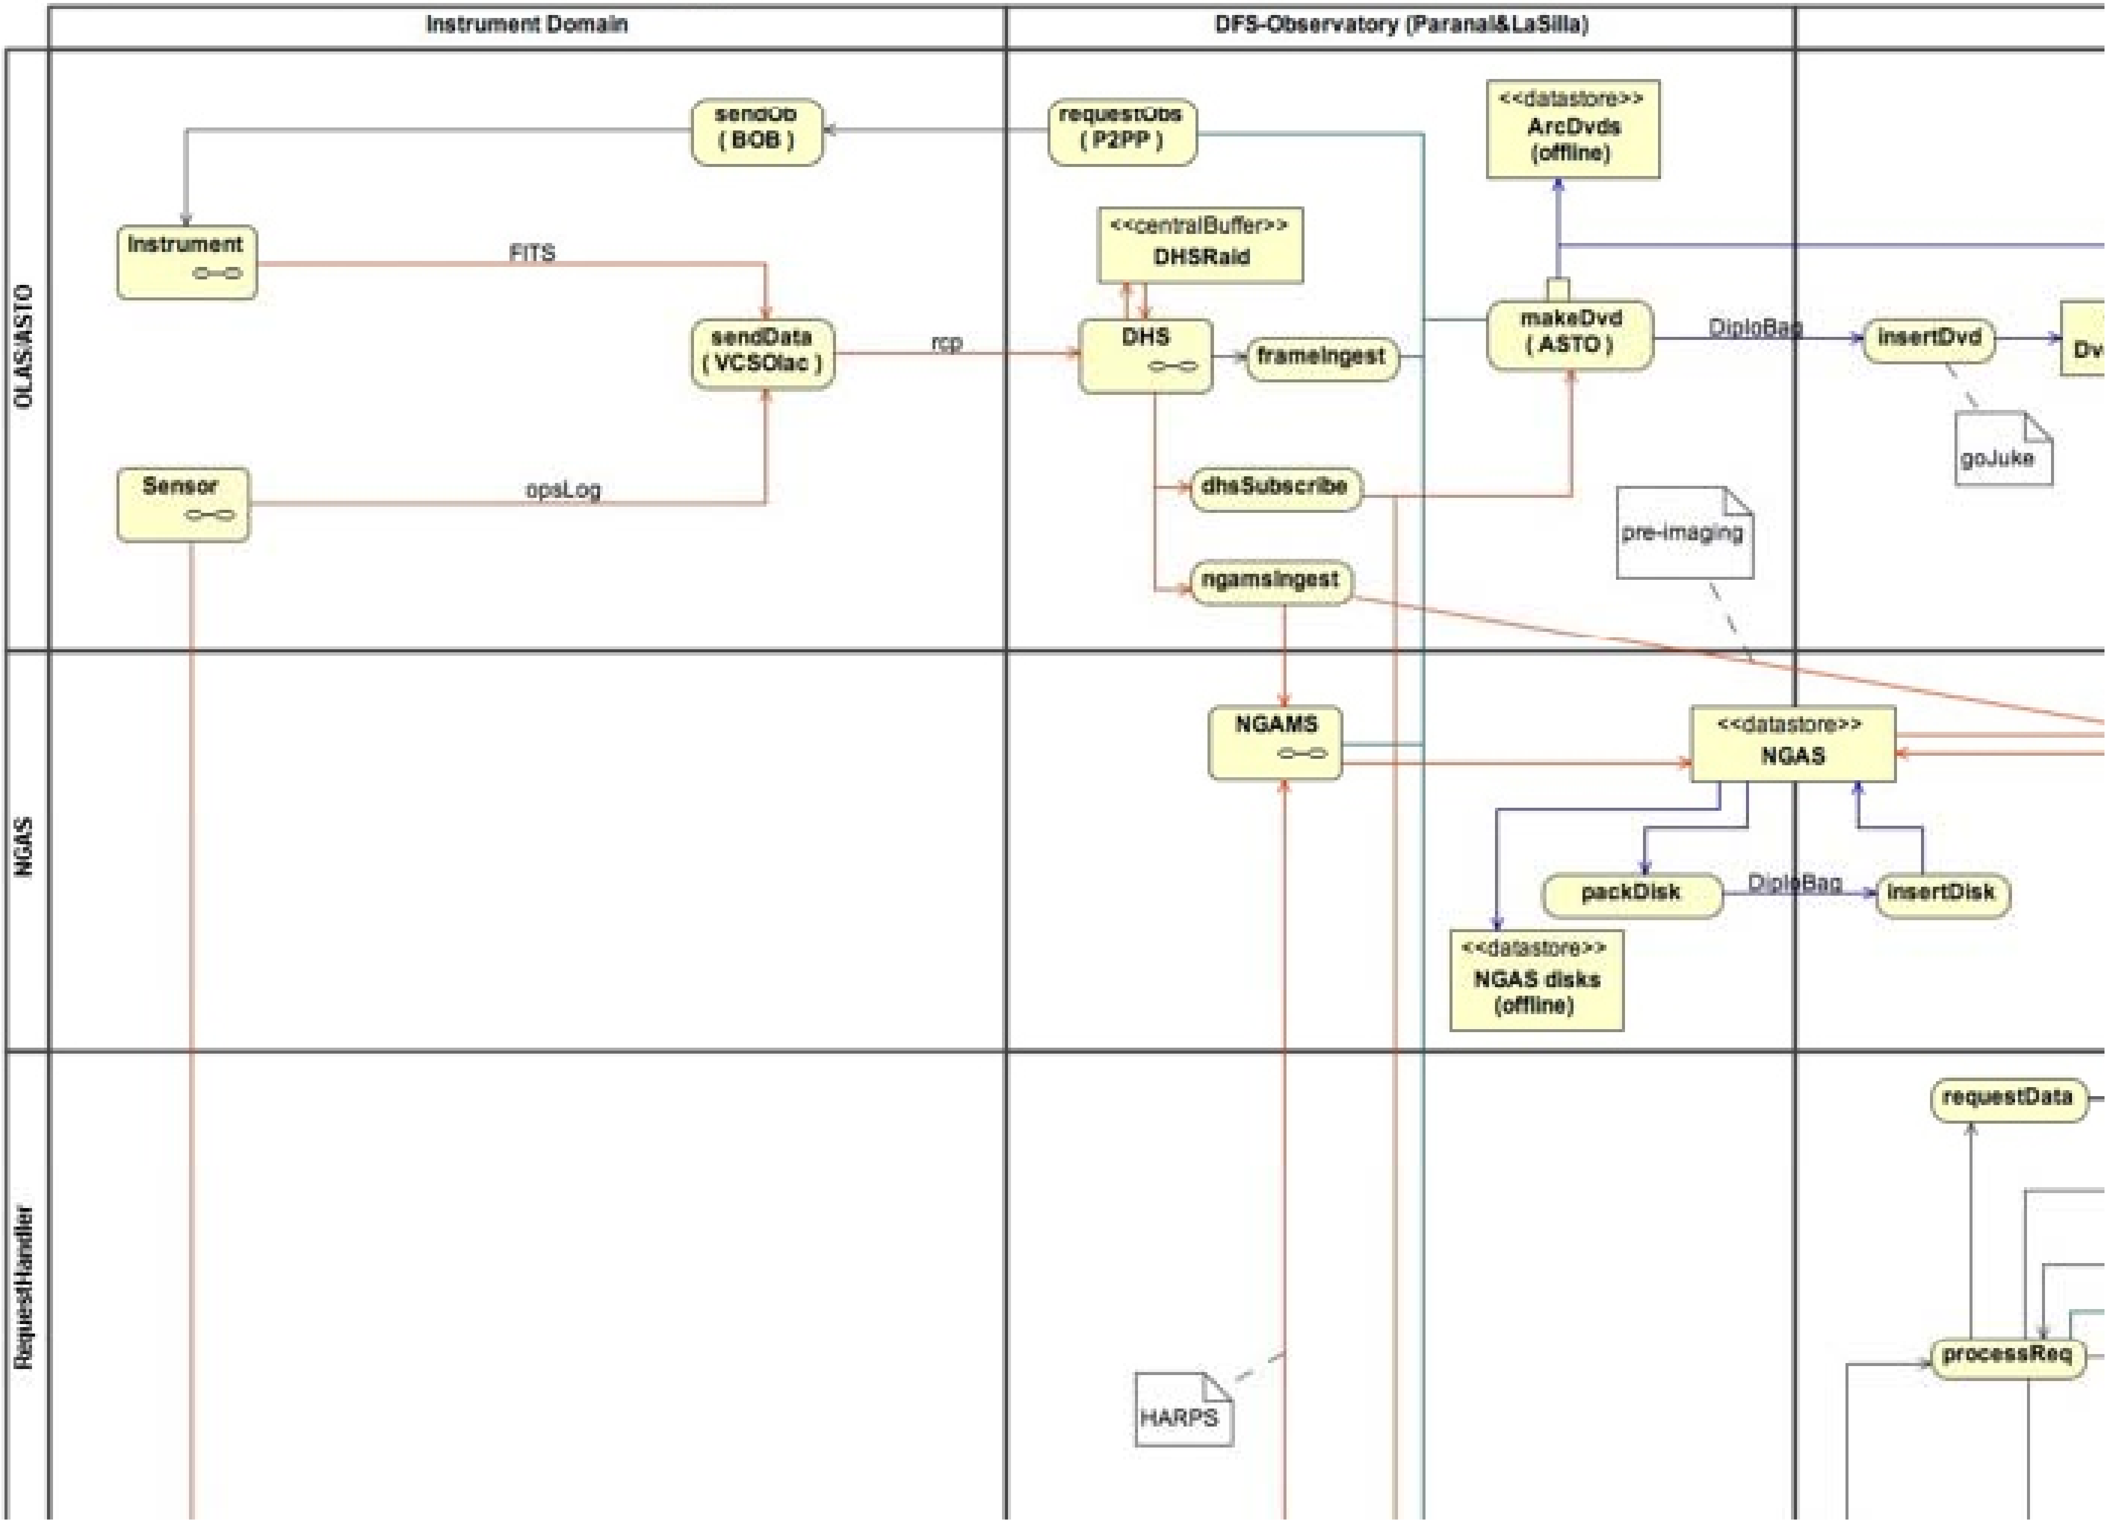

Does SKA need a data grid?

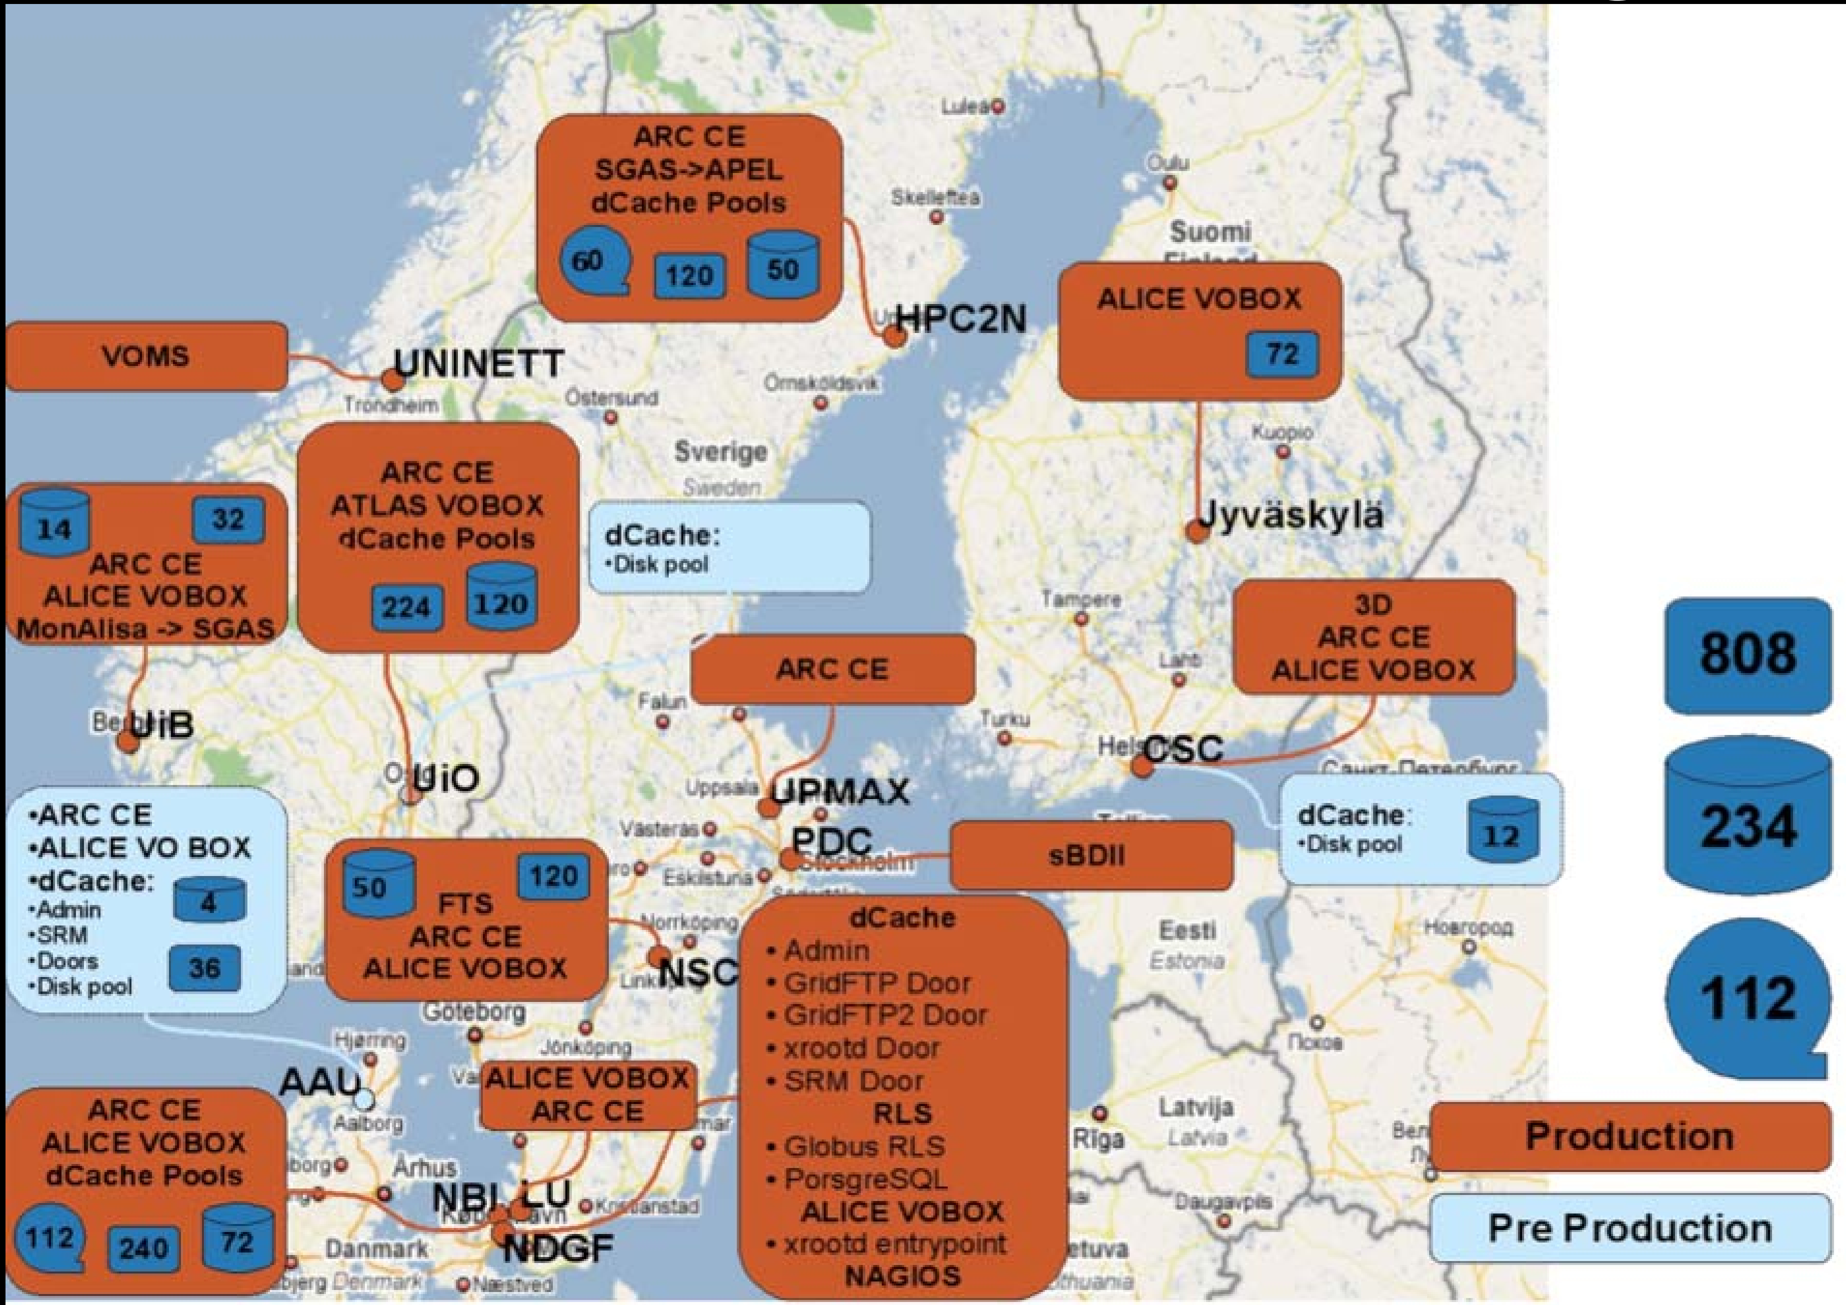

NDGF/Nordunet 2008

Science Operations Plan

- Operations plan and high level analysis are key to the design process, but are notoriously late, incomplete or lacking general agreement.
- Operations plan also main connection point between scientific intend and technical implementation.
- Operations plan has to be augmented, detailed and explained by domain specialists, aka subsystem scientists in an active process.

Operations Plan Questions

- How will the SKA announce, select and carry out observations?
  - PI driven proposals?
  - Peer-review process? Continuous or in terms?
  - Technical evaluation?
  - Short observing projects **and** long survey type projects?
  - Automatic scheduling, heuristics, long-mid-short term?
  - Calibration plan, technical time?

More Questions

- How will the SKA observatory work?
  - Technical and scientific staffing 24/7?
  - Service and/or visitor mode?
  - QA process?
  - Observing project deliverables: What does the observatory guarantee?
  - Observing modes?
  - Technical details constraining observations?
  - Policies: Data access/ownership/proprietary period?

Operations Plan vs. Observing Project

- Operations Plan needs to be analysed and turned into an high level implementation plan.
- Observing Project requires detailed definition of observatory characteristics and capabilities: If not known well enough a project could turn out not be feasible **after** acceptance.
- Change of view: Design of Observing Project could drive the detailed specifications of the operations plan. Means: Write operations plan according to science projects, not the other way around!

How

- The observing project collects all science and technical requirements and all observational and auxiliary data and scientific results:
  - Construct a representative set of SKA observing projects including calibration and technical projects.
  - The set should include projects suited for execution with a small subset of the SKA, SKA1 and SKA2.
  - Design and build the system against this set.
  - Document the translation process from science requirements into system requirements. Example: Spatial resolution of x mas translates into baseline of y km.

I/O Challenge

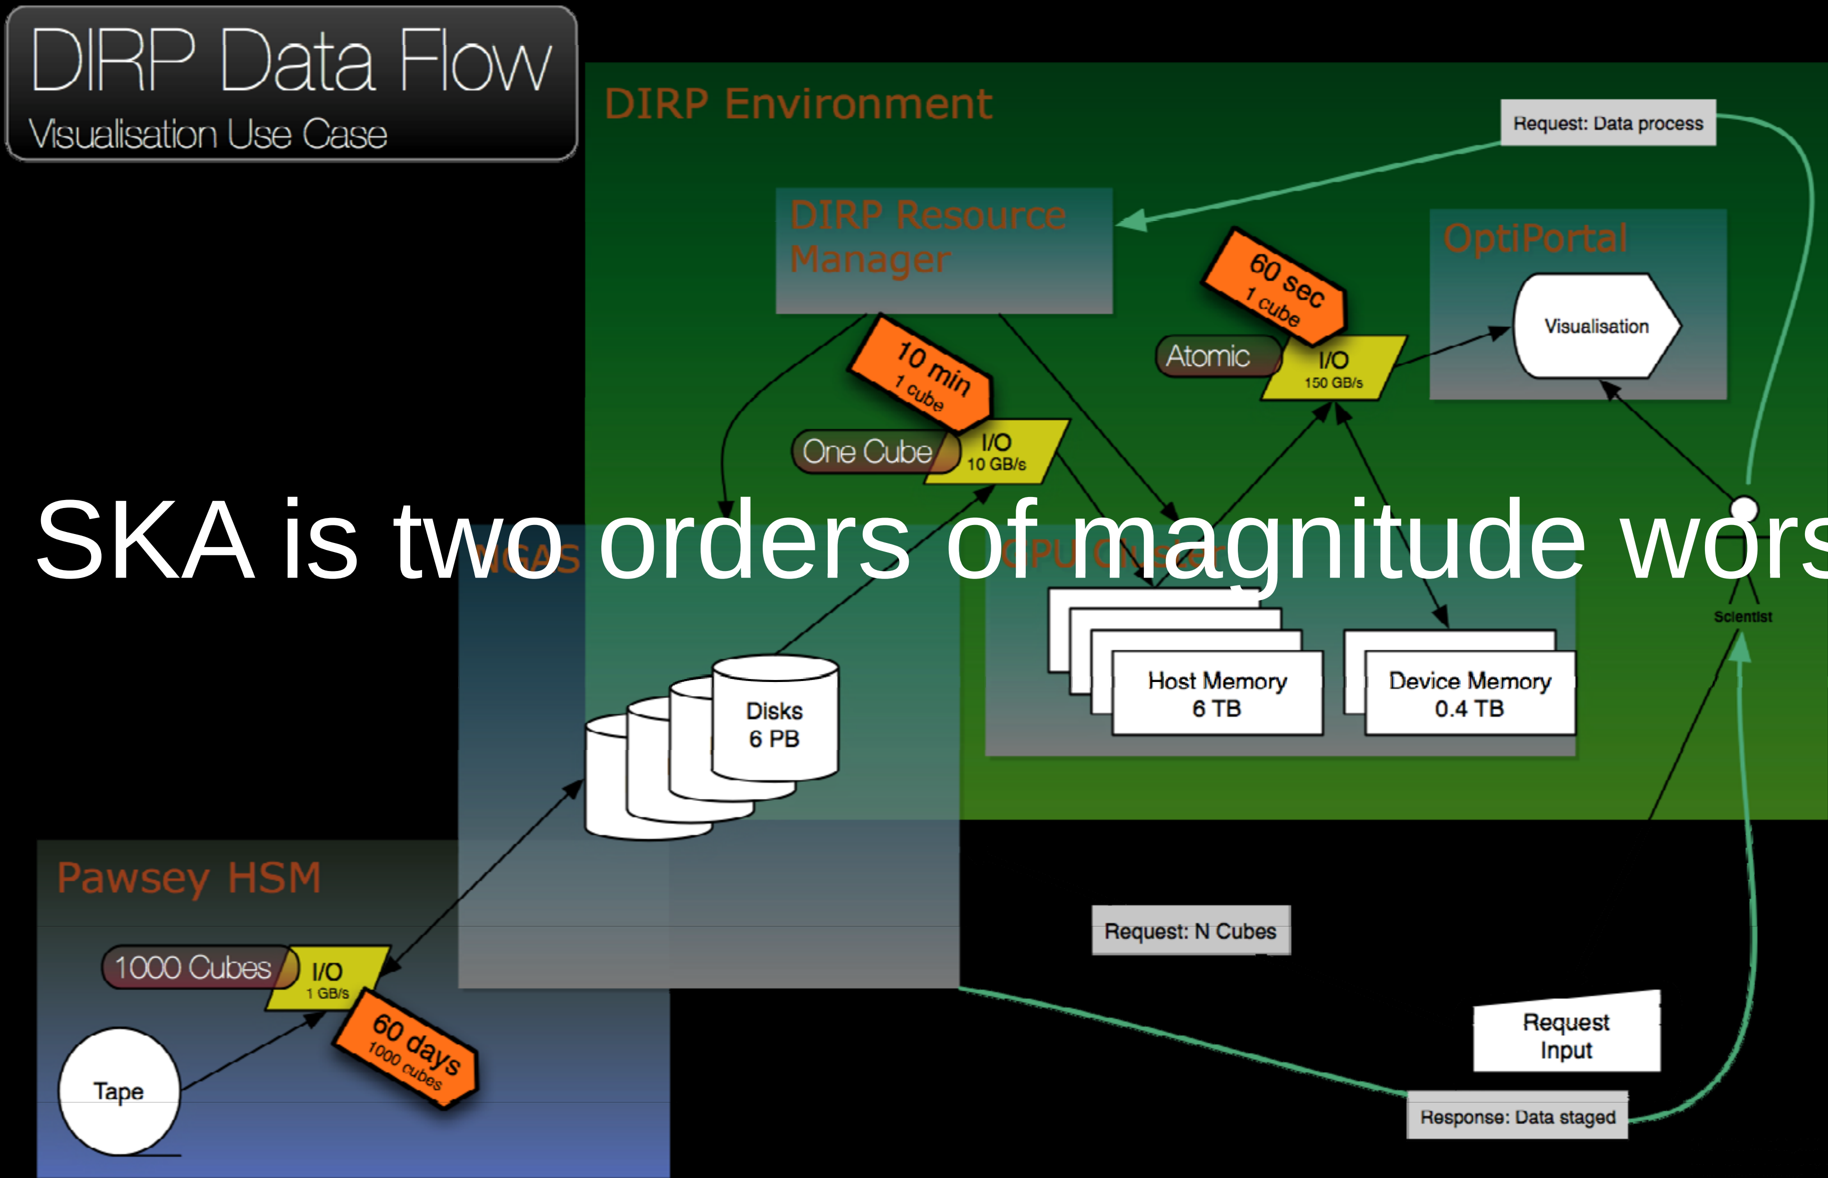

SKA is two orders of magnitude worse!

Algorithmic Challenge

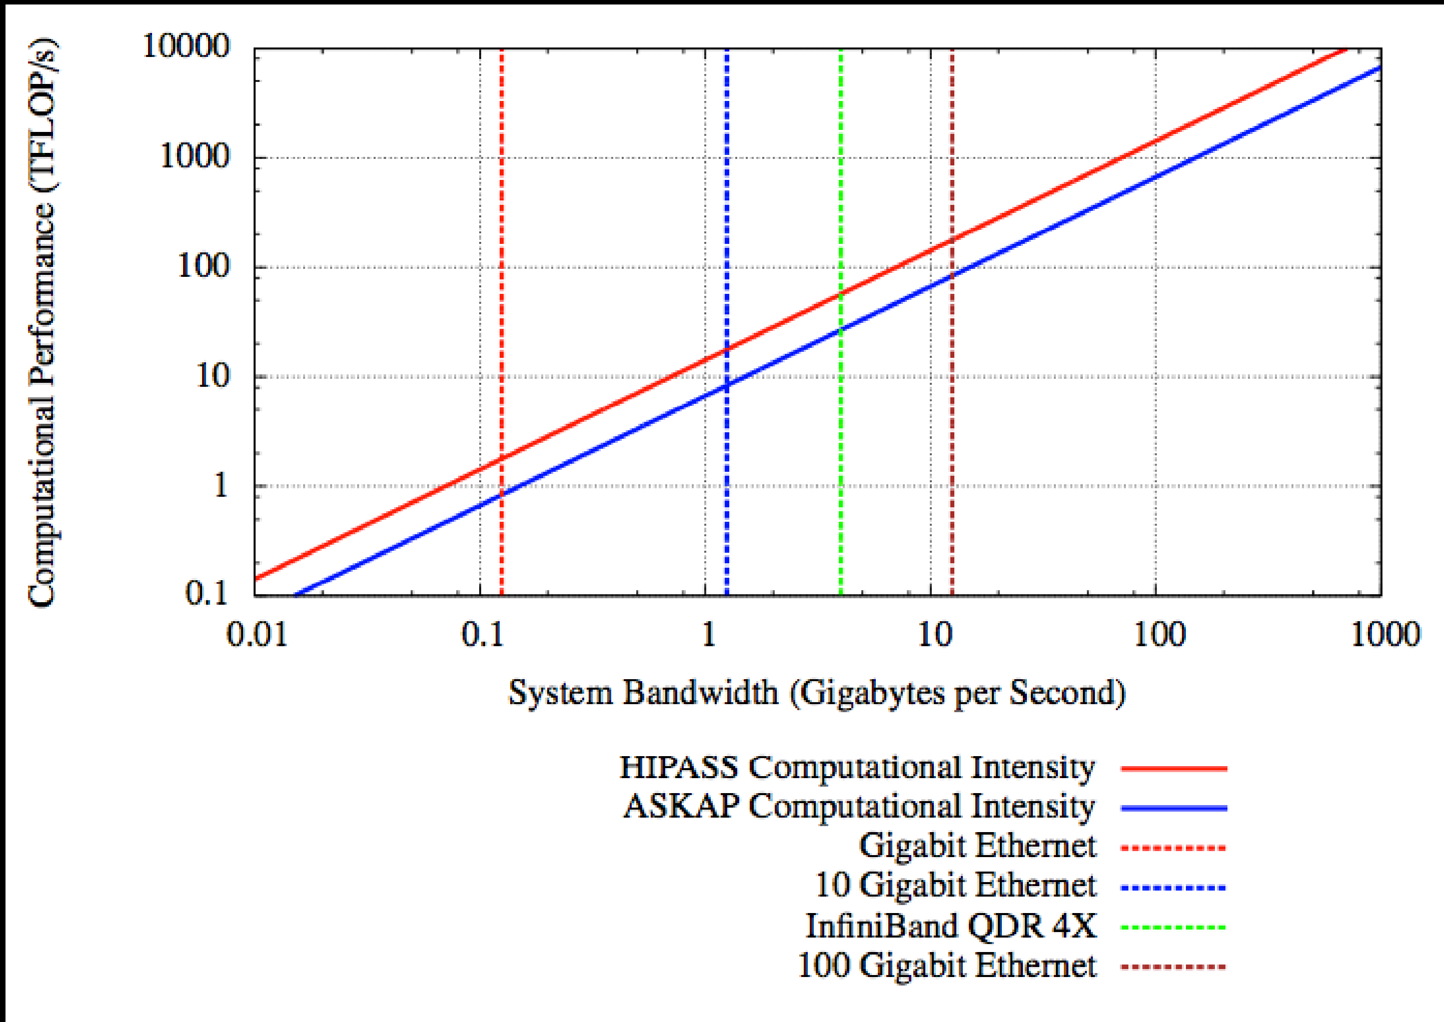

Algorithmic Challenge

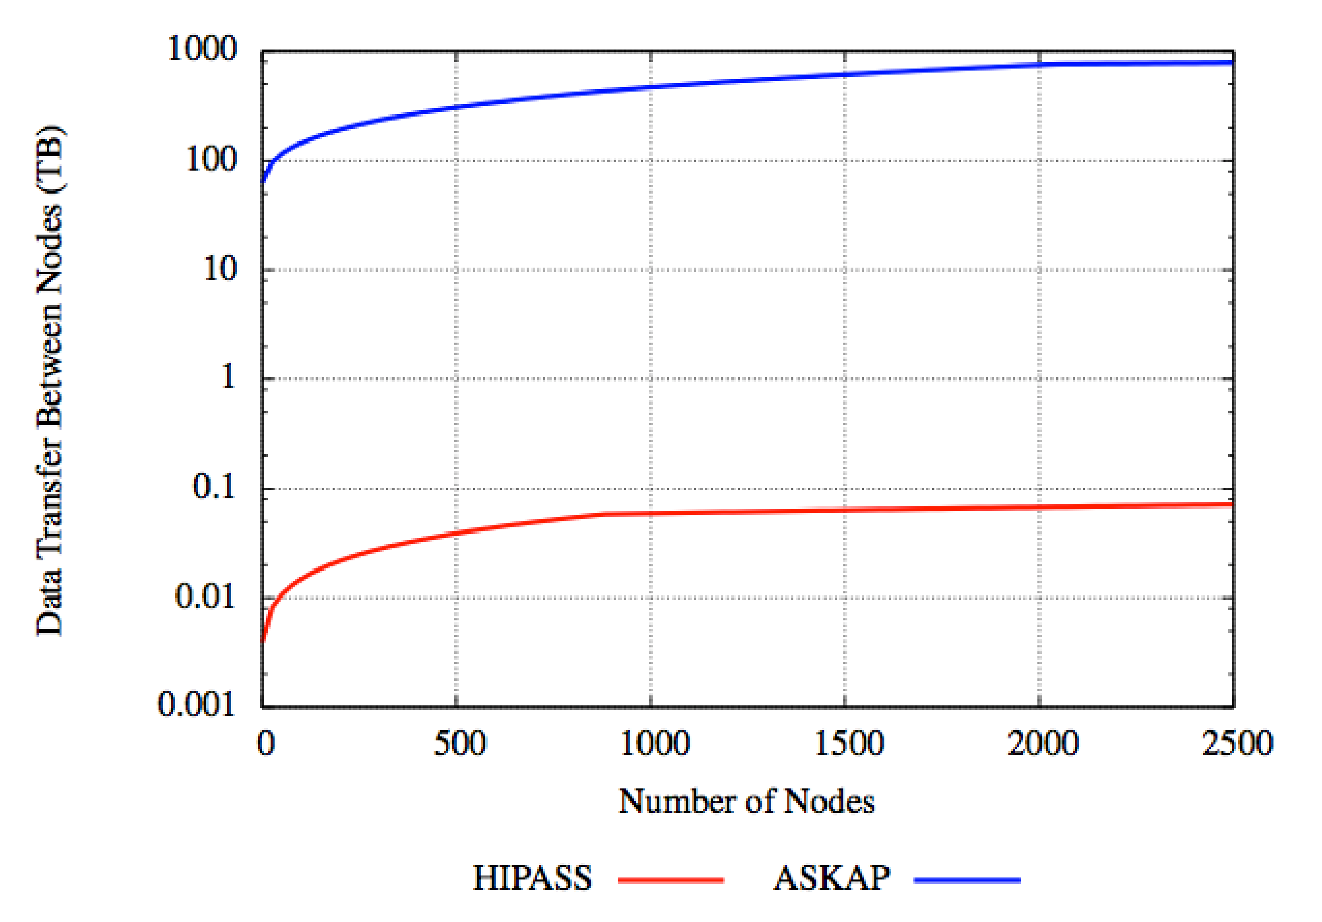

Algorithmic Challenge

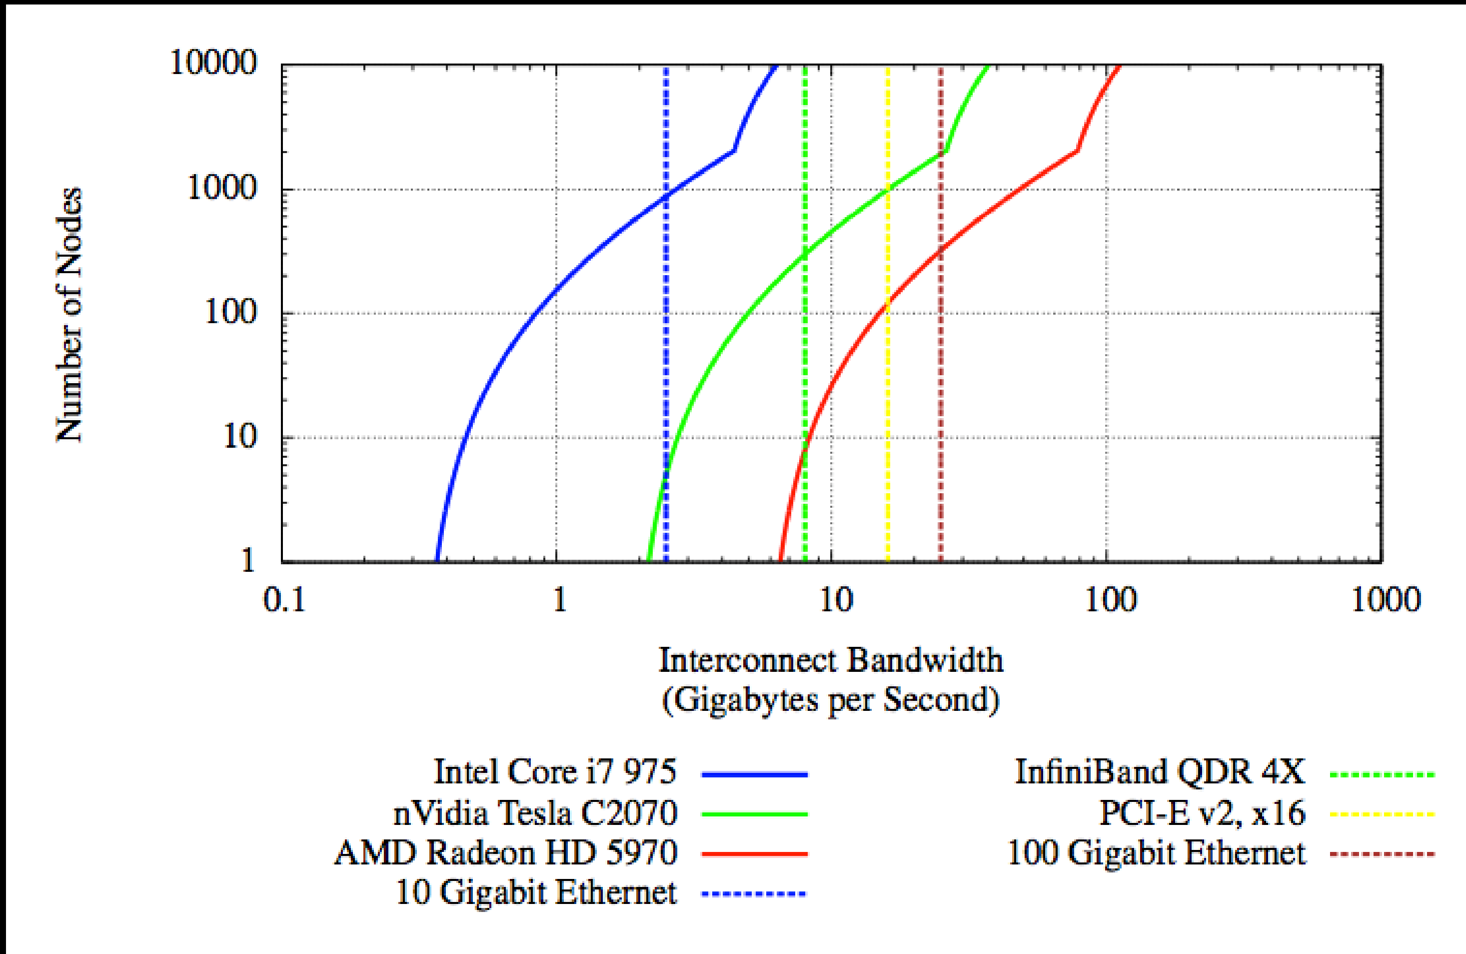

Conclusion

- Observing Project approach proven concept, but needs some adjustments for SKA to be able to cope with data flow and SKA usage pattern.
- SKA will most probably require the fastest HPC centre on the planet, at a fraction of cost and power consumption of the current top computer!
- HPC centre needs to be fully dedicated and used in streaming mode!!
- Scientific exploitation of survey data will push astronomers to a whole new world.

SKA == 40 x this

Tihanhe-1A
7168 GPUs
14336 Cores
> 4 MW

I/O Challenge

- 1 cube == 6TB
- 1 cube at 100 Gbit/s == 1 cube at 12 GByte/s ==> 8 minutes
- PCIe x 16 = 16 x 8 Gbit/s == 16 GByte/s
- DDR3 12800 memory bandwidth == 12.8 GByte/s ==> 1 cube transfer into DDR3 12800 memory takes also about 8 minutes!!
- NVidia Tesla Memory Bandwidth = 170 Gytes/s (6 GB dedicated memory) ==> about 35 seconds to read 1 cube
- NVidia Tesla 6 GB ==> 1000 Tesla cards to keep one cube in memory ==> intelligent memory management between host and GPU memory.

I/O Challenge

- NVidia Tesla 2070 requires about 0.25kW, 1 TFlops ==> 200 TFlops means 200 cards, these can keep 1.2 TB of dataSustained transfer rate
- SATA disks = 100 MByte/s, ==> requires 240 disks full speed to get to 12 GByte/s (2 x DDN 9900, Panasas)
- iODriveDuo transfer rate 1.5 GByte/s (read), 1GByte/s (write) ==> requires 8 ioDriveDuo to get to 12 GByte/s read, 12 cards for write1 ioDriveDuo == 640 GB ==> 10 cards to hold 1 cube
- BlueGen I/O nodes limited to 10 Gbit/s ==> 10 I/O nodes for 100 Gbit/s rest of the system can be scaled accordingly.

- Build simulators for everything which does not exist (yet) to have a complete life-cycle in place very early in the project.