International Centre for Radio Astronomy Research

The SKA Data Flow Systems Design

Hatching out of the Egg

Andreas Wicenec

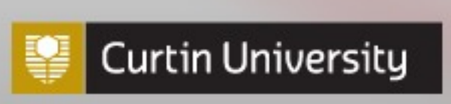

THE UNIVERSITY OF WESTERN AUSTRALIA

Achieve International Excellence

Logarithmic Growth in Flops

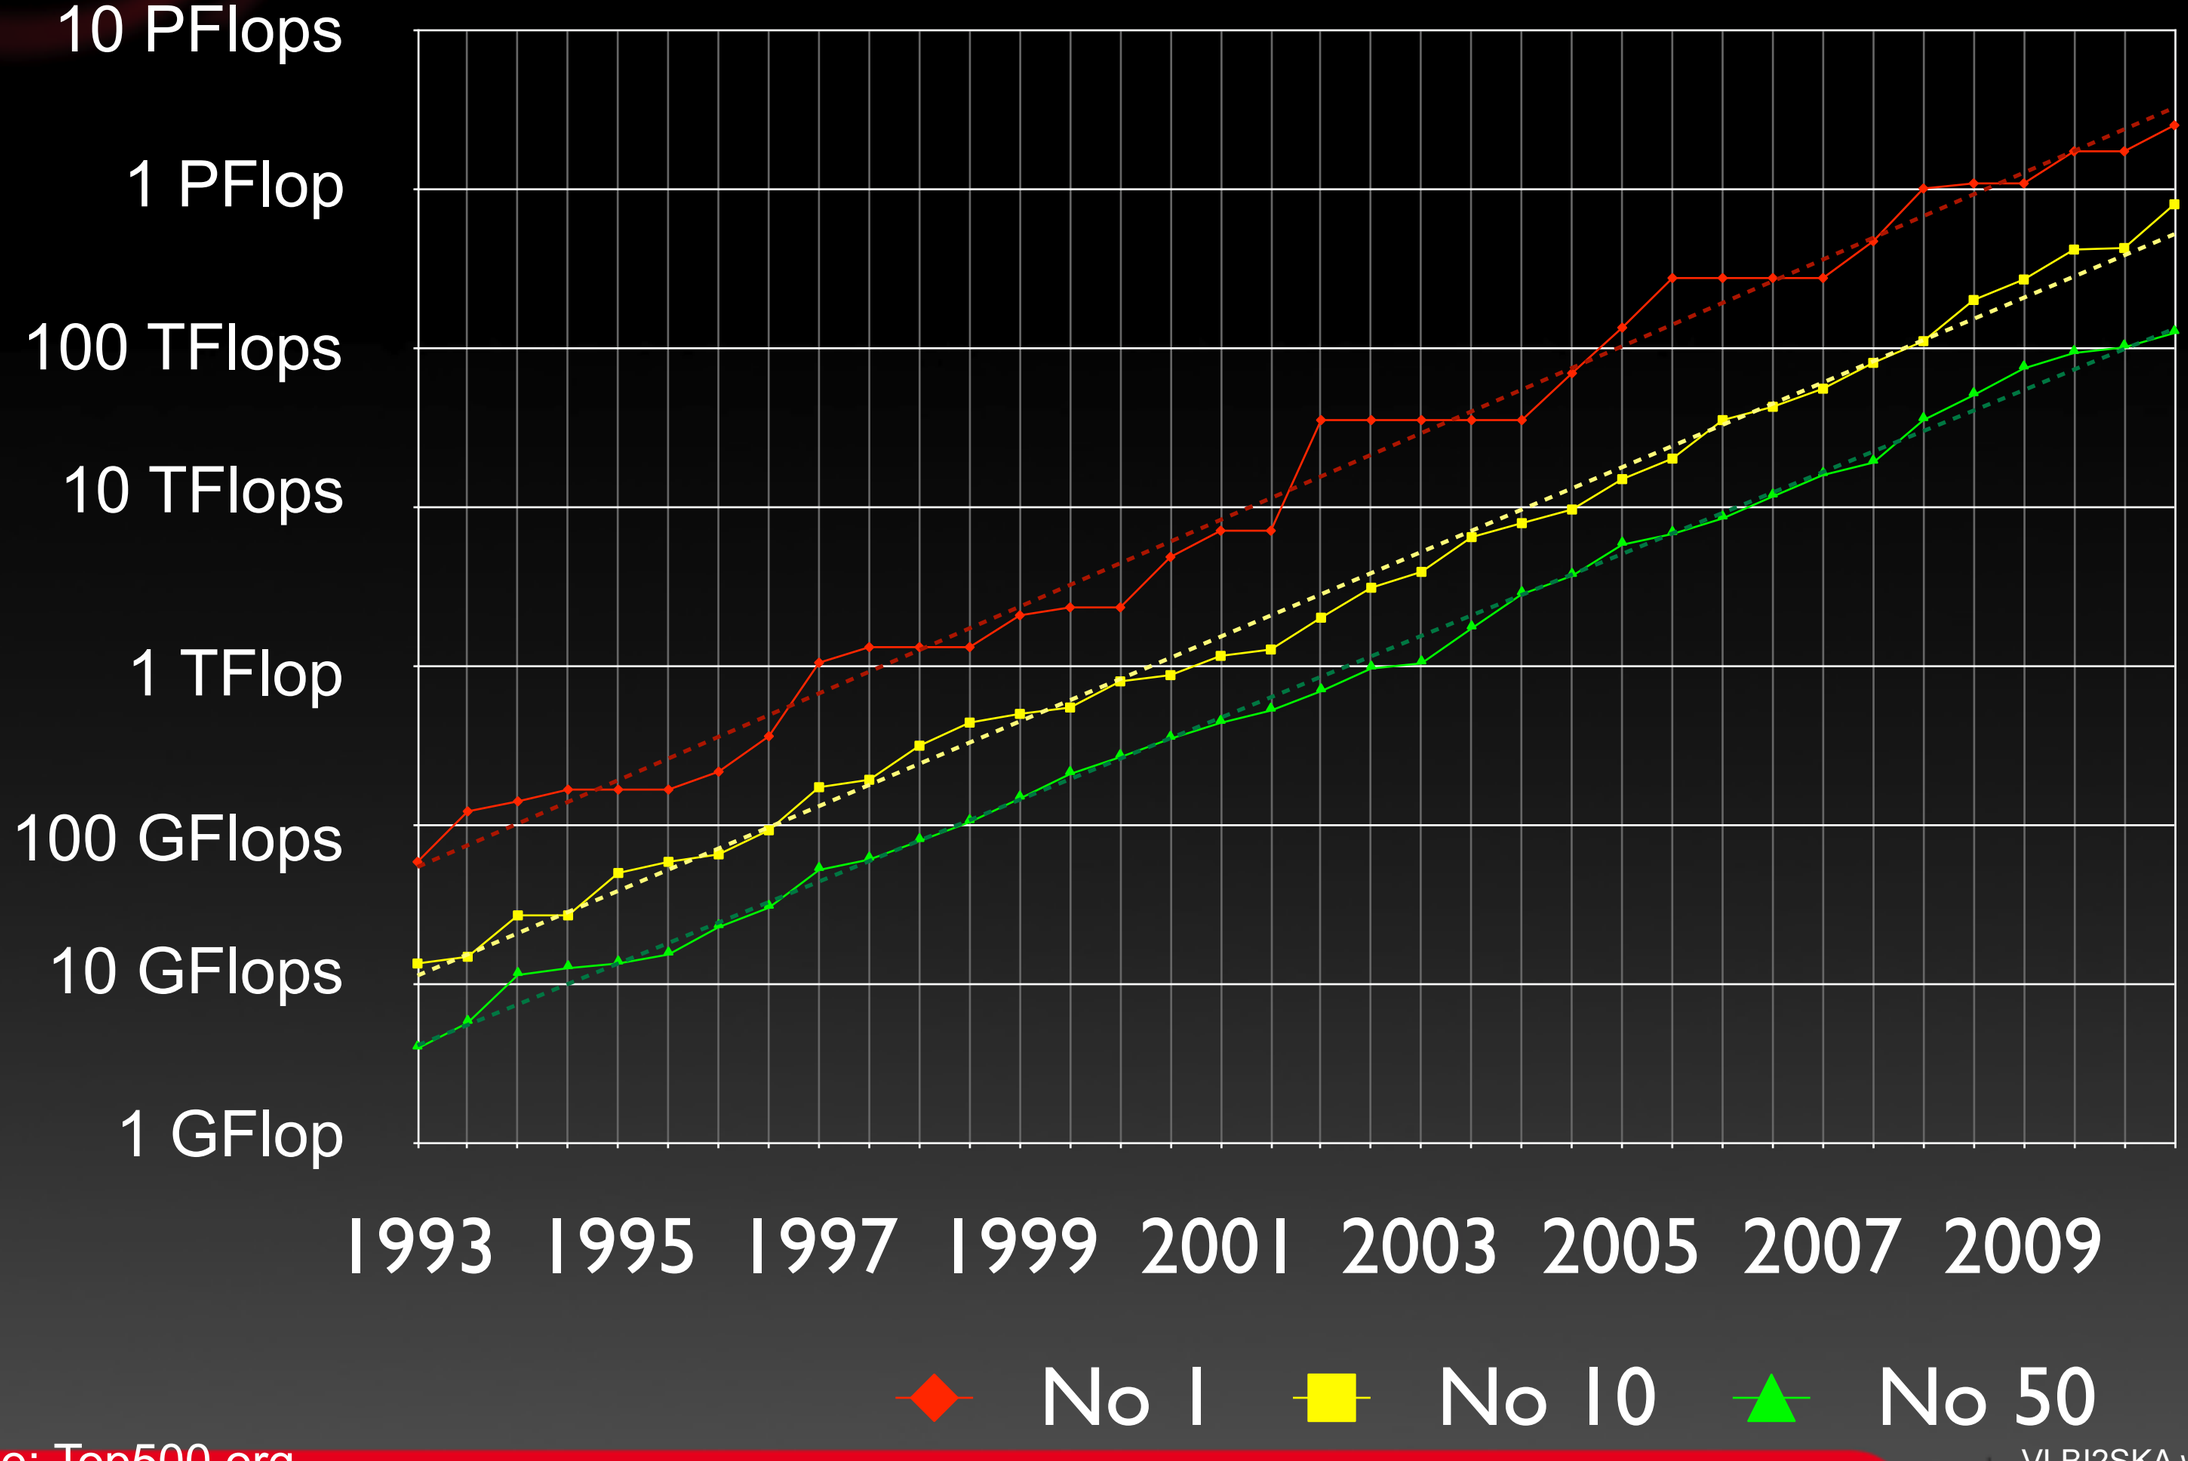

Source: Top500.org

Projected Performance

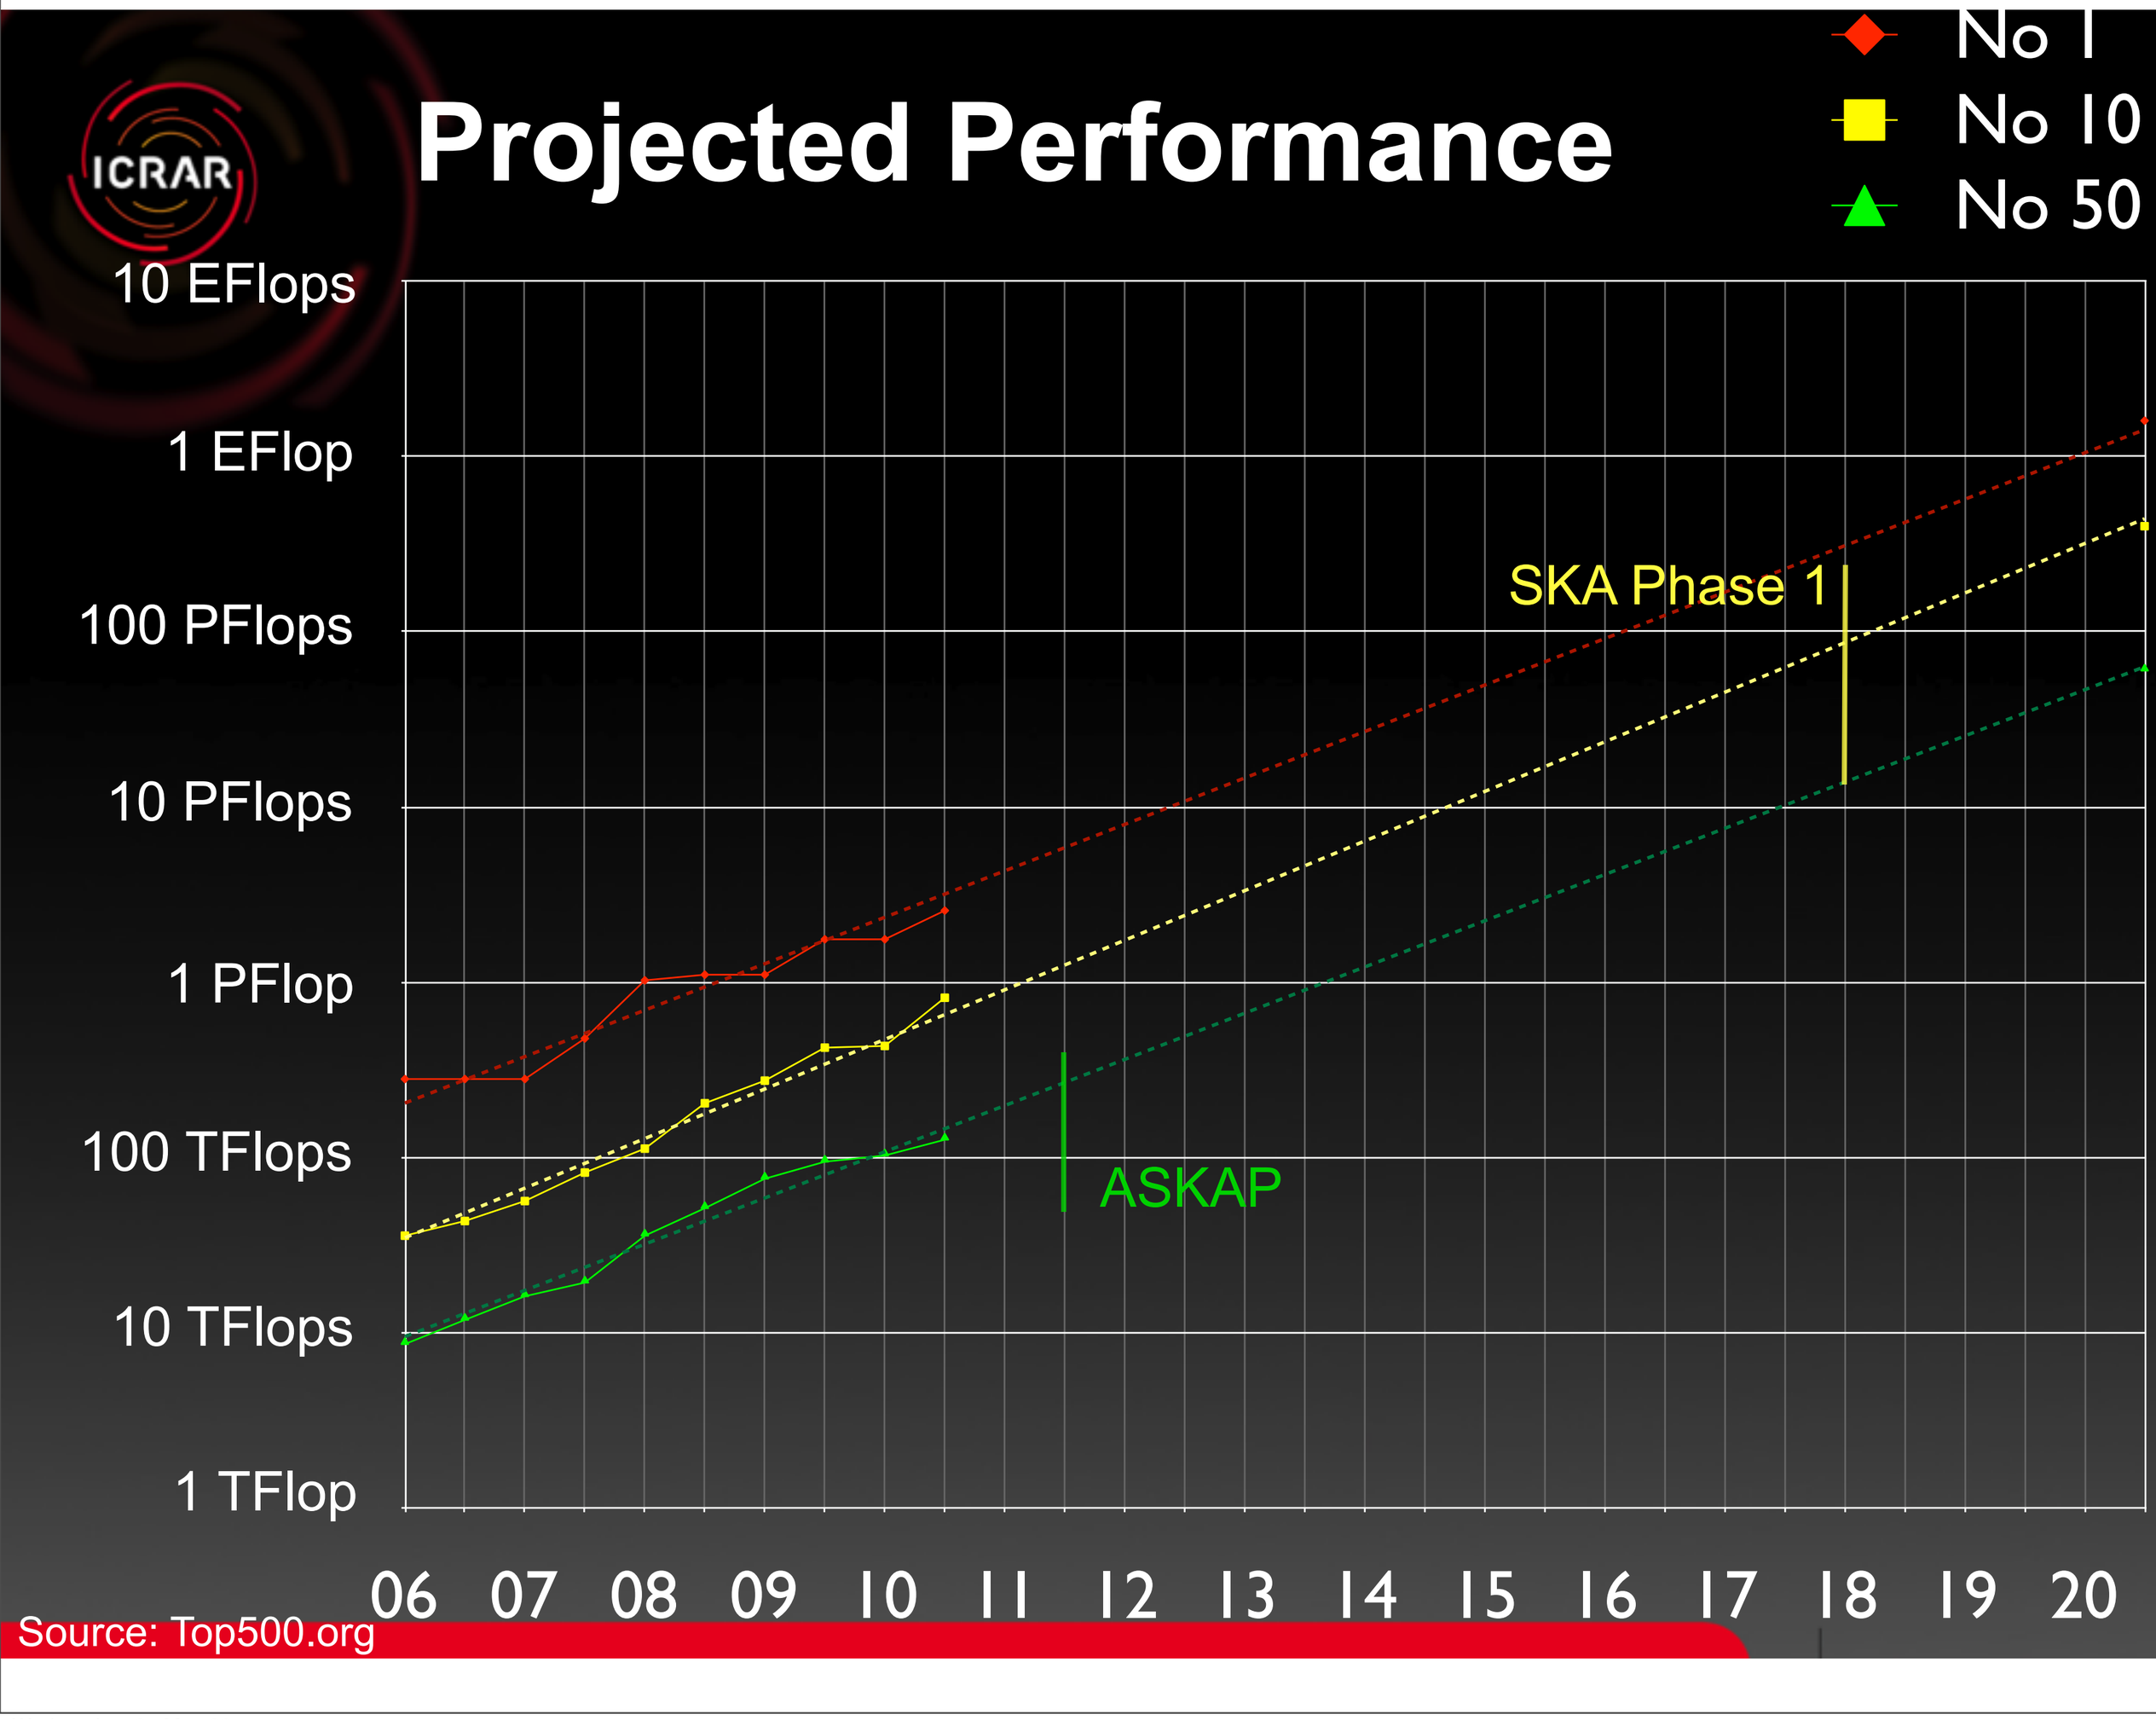

Source: Top500.org

SKA
Processing & Persistence
Services

ObsProject Lifecycle

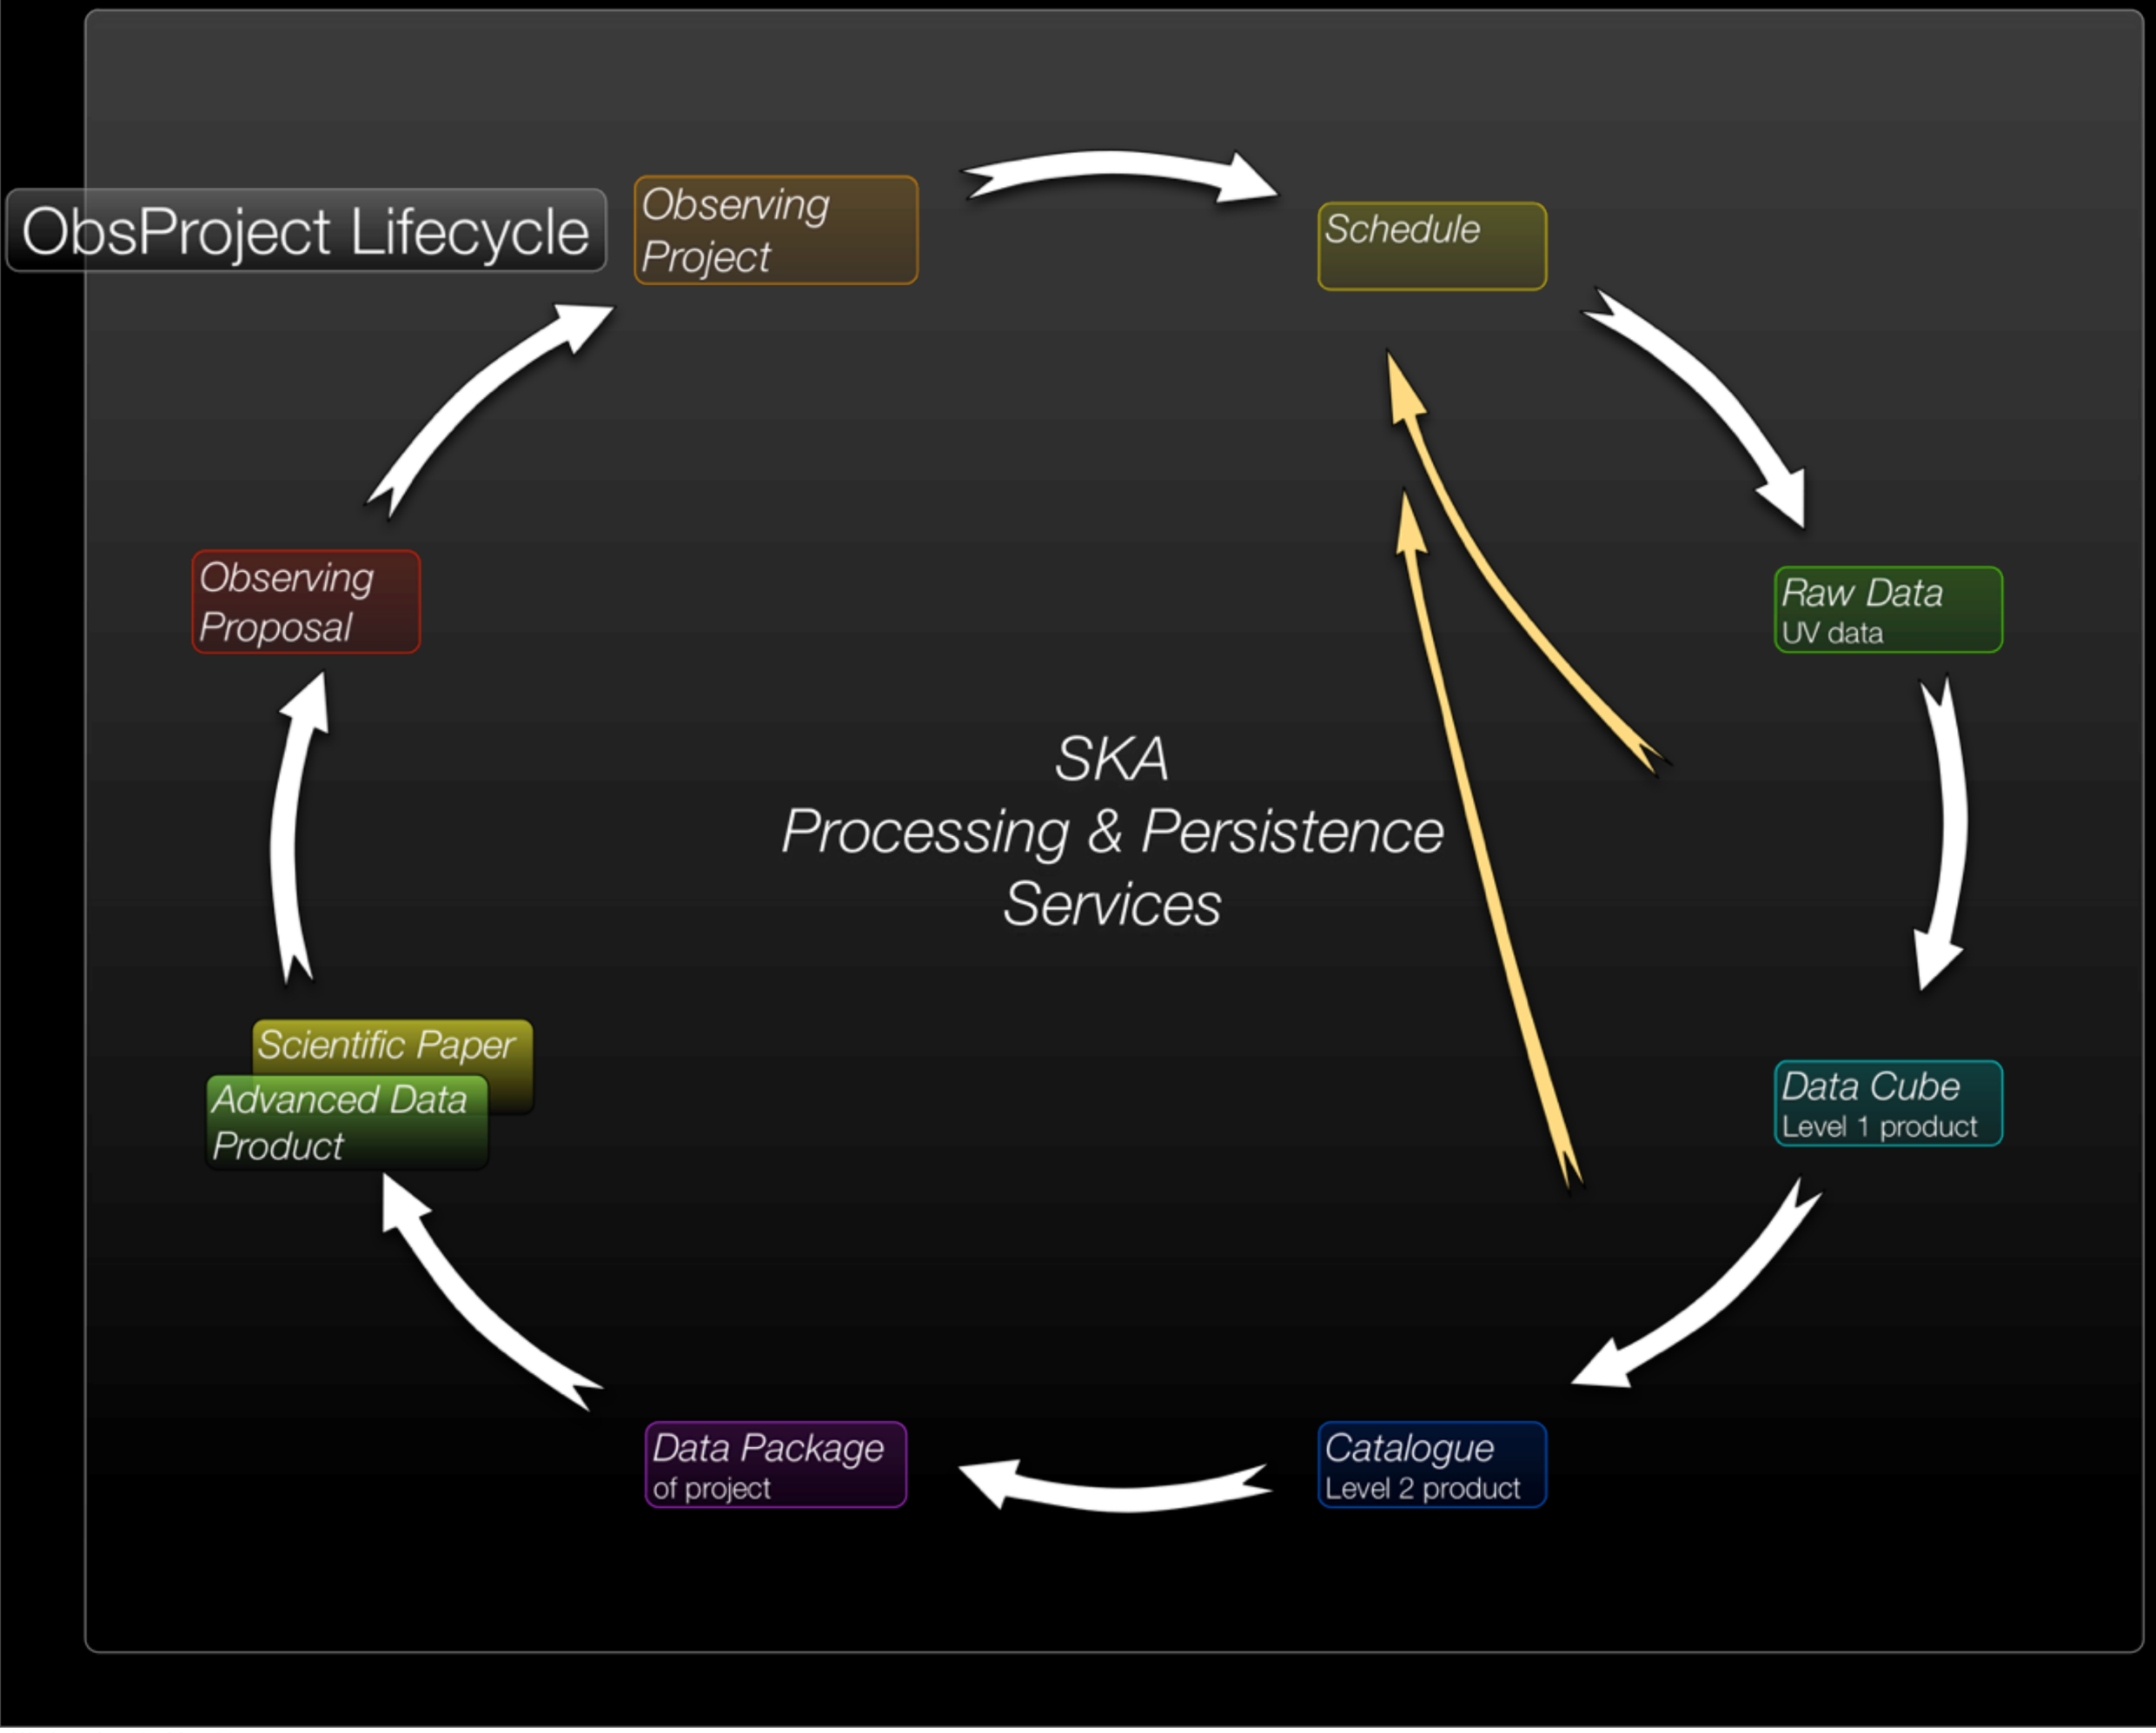

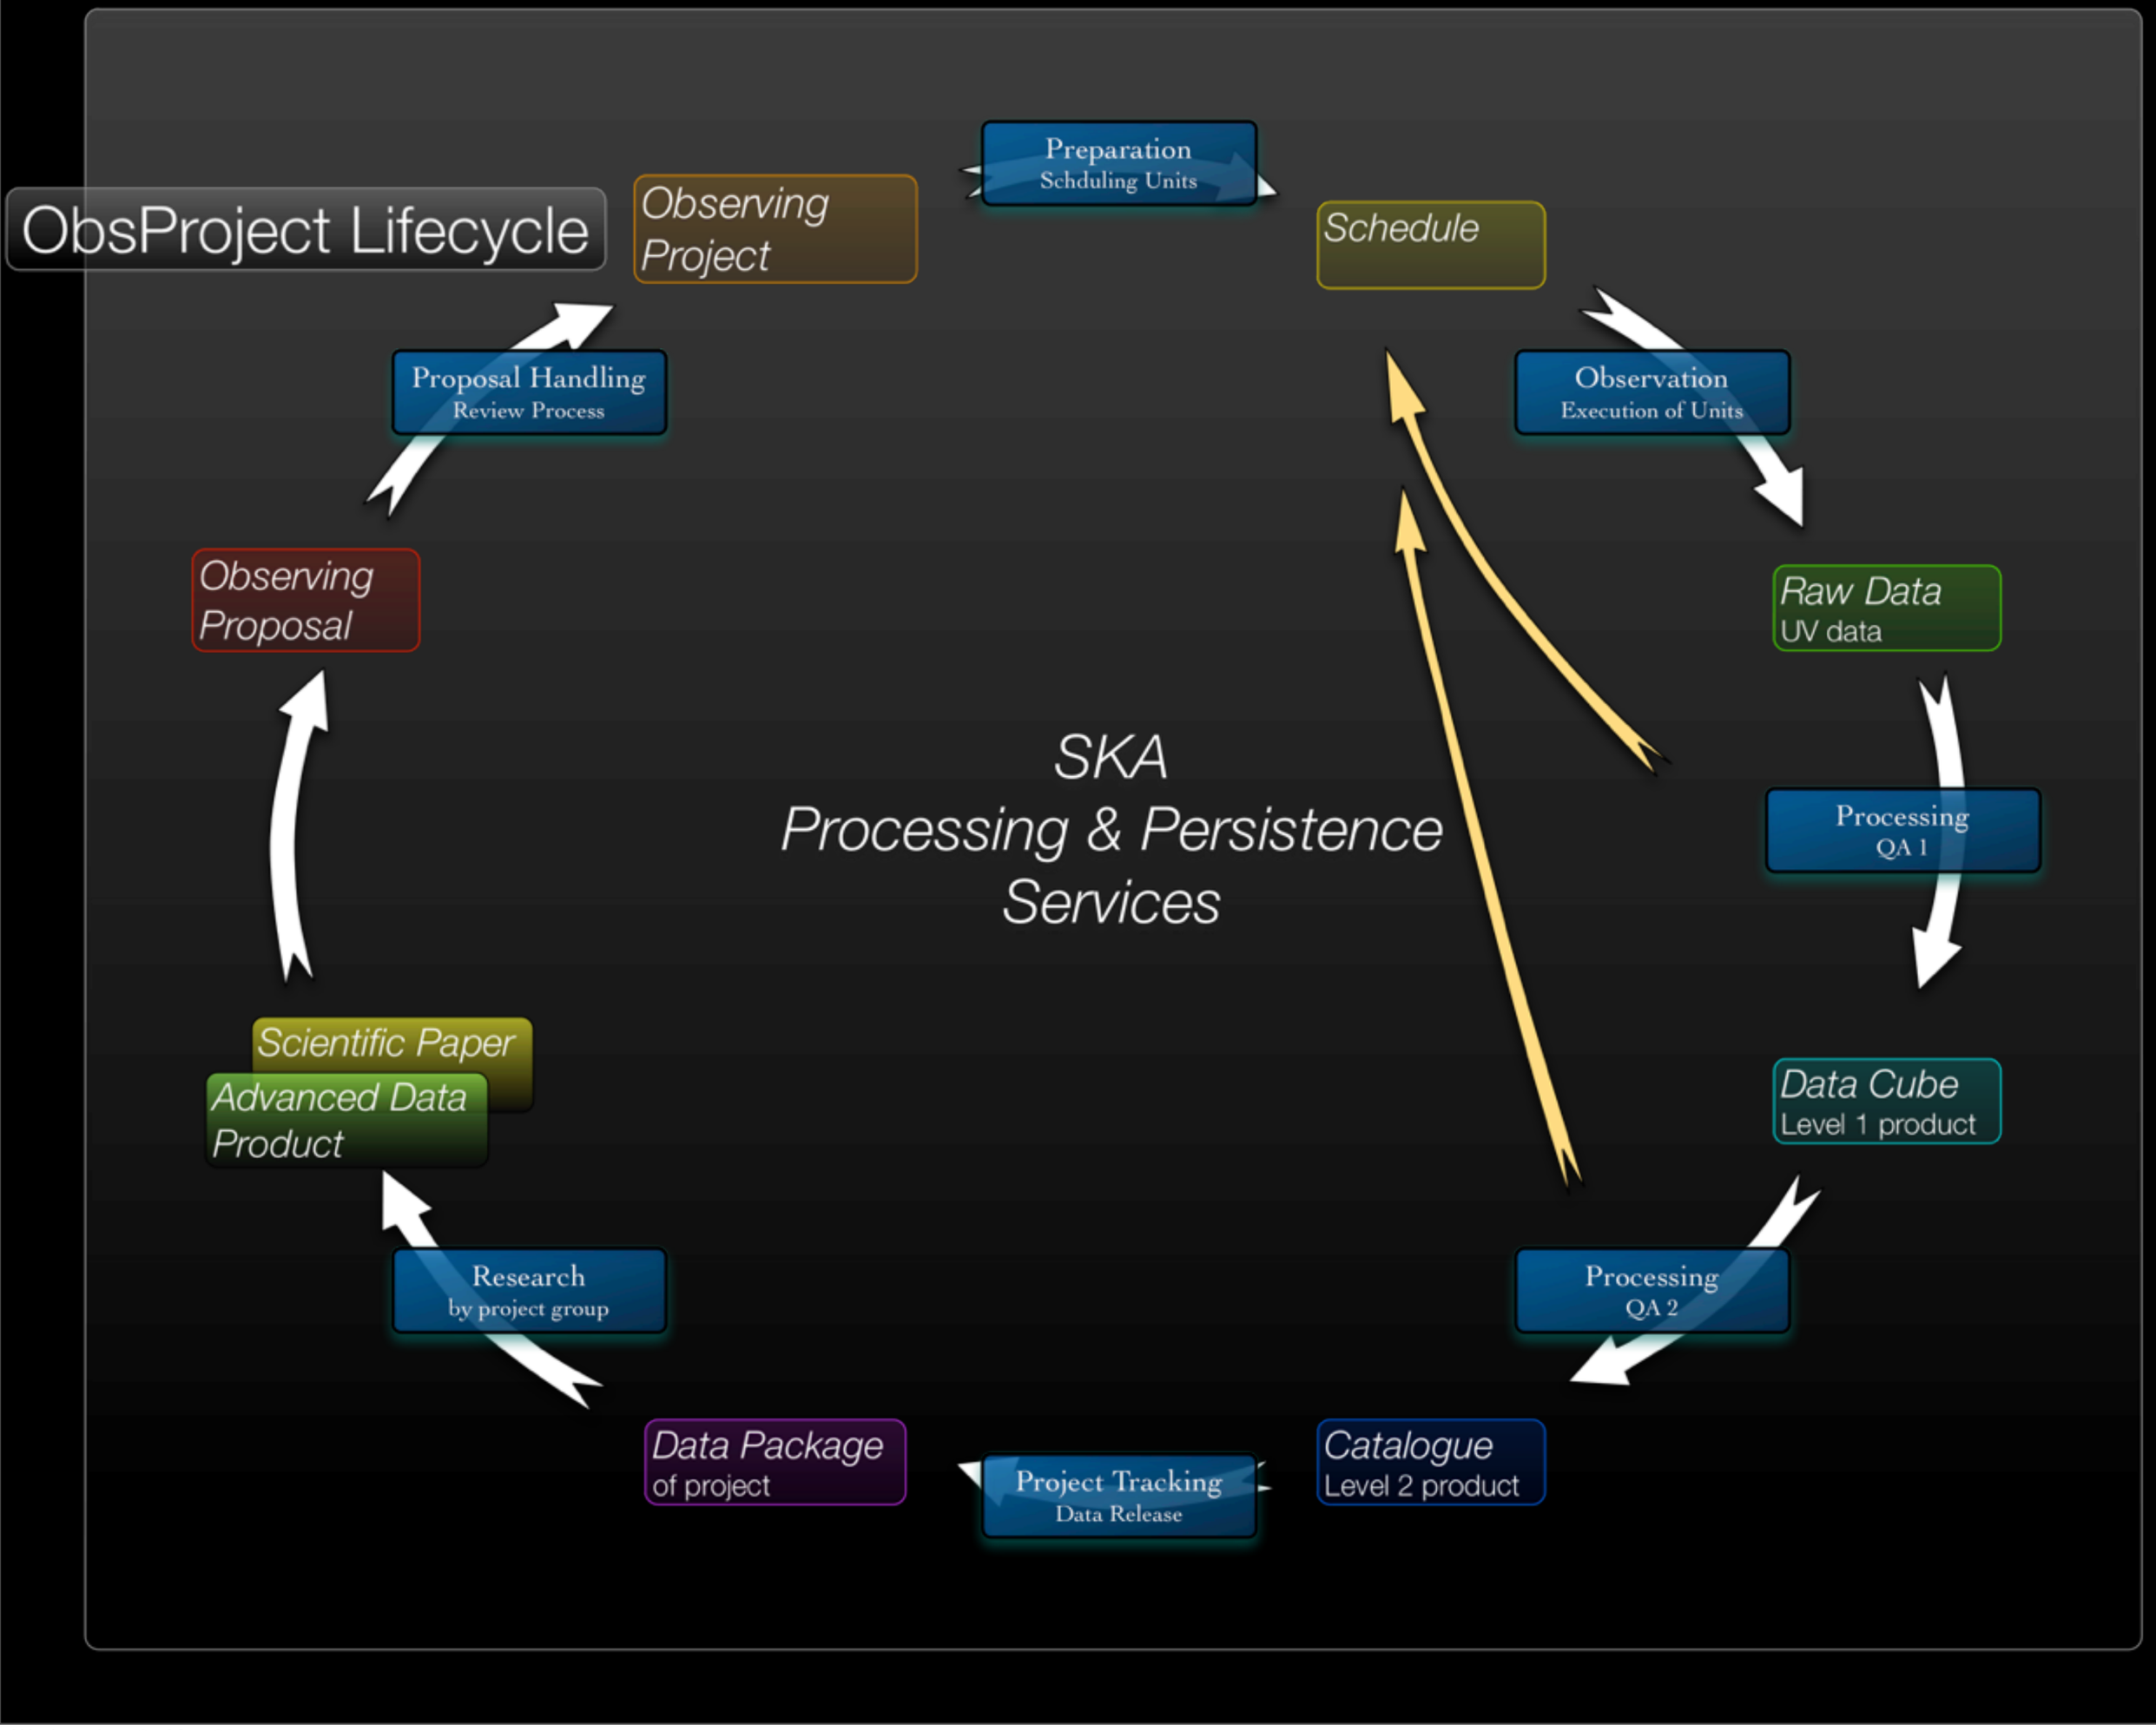
ObsProject Lifecycle
Observing Project
Preparation
Schduling Units
Schedule
Observation
Execution of Units
Raw Data
UV data
Processing
QA 1
Data Cube
Level 1 product
Processing
QA 2
Catalogue
Level 2 product
Project Tracking
Data Release
Data Package
of project
Research
by project group
Advanced Data Product
Scientific Paper
Observing Proposal
Proposal Handling
Review Process
SKA
Processing & Persistence
Services

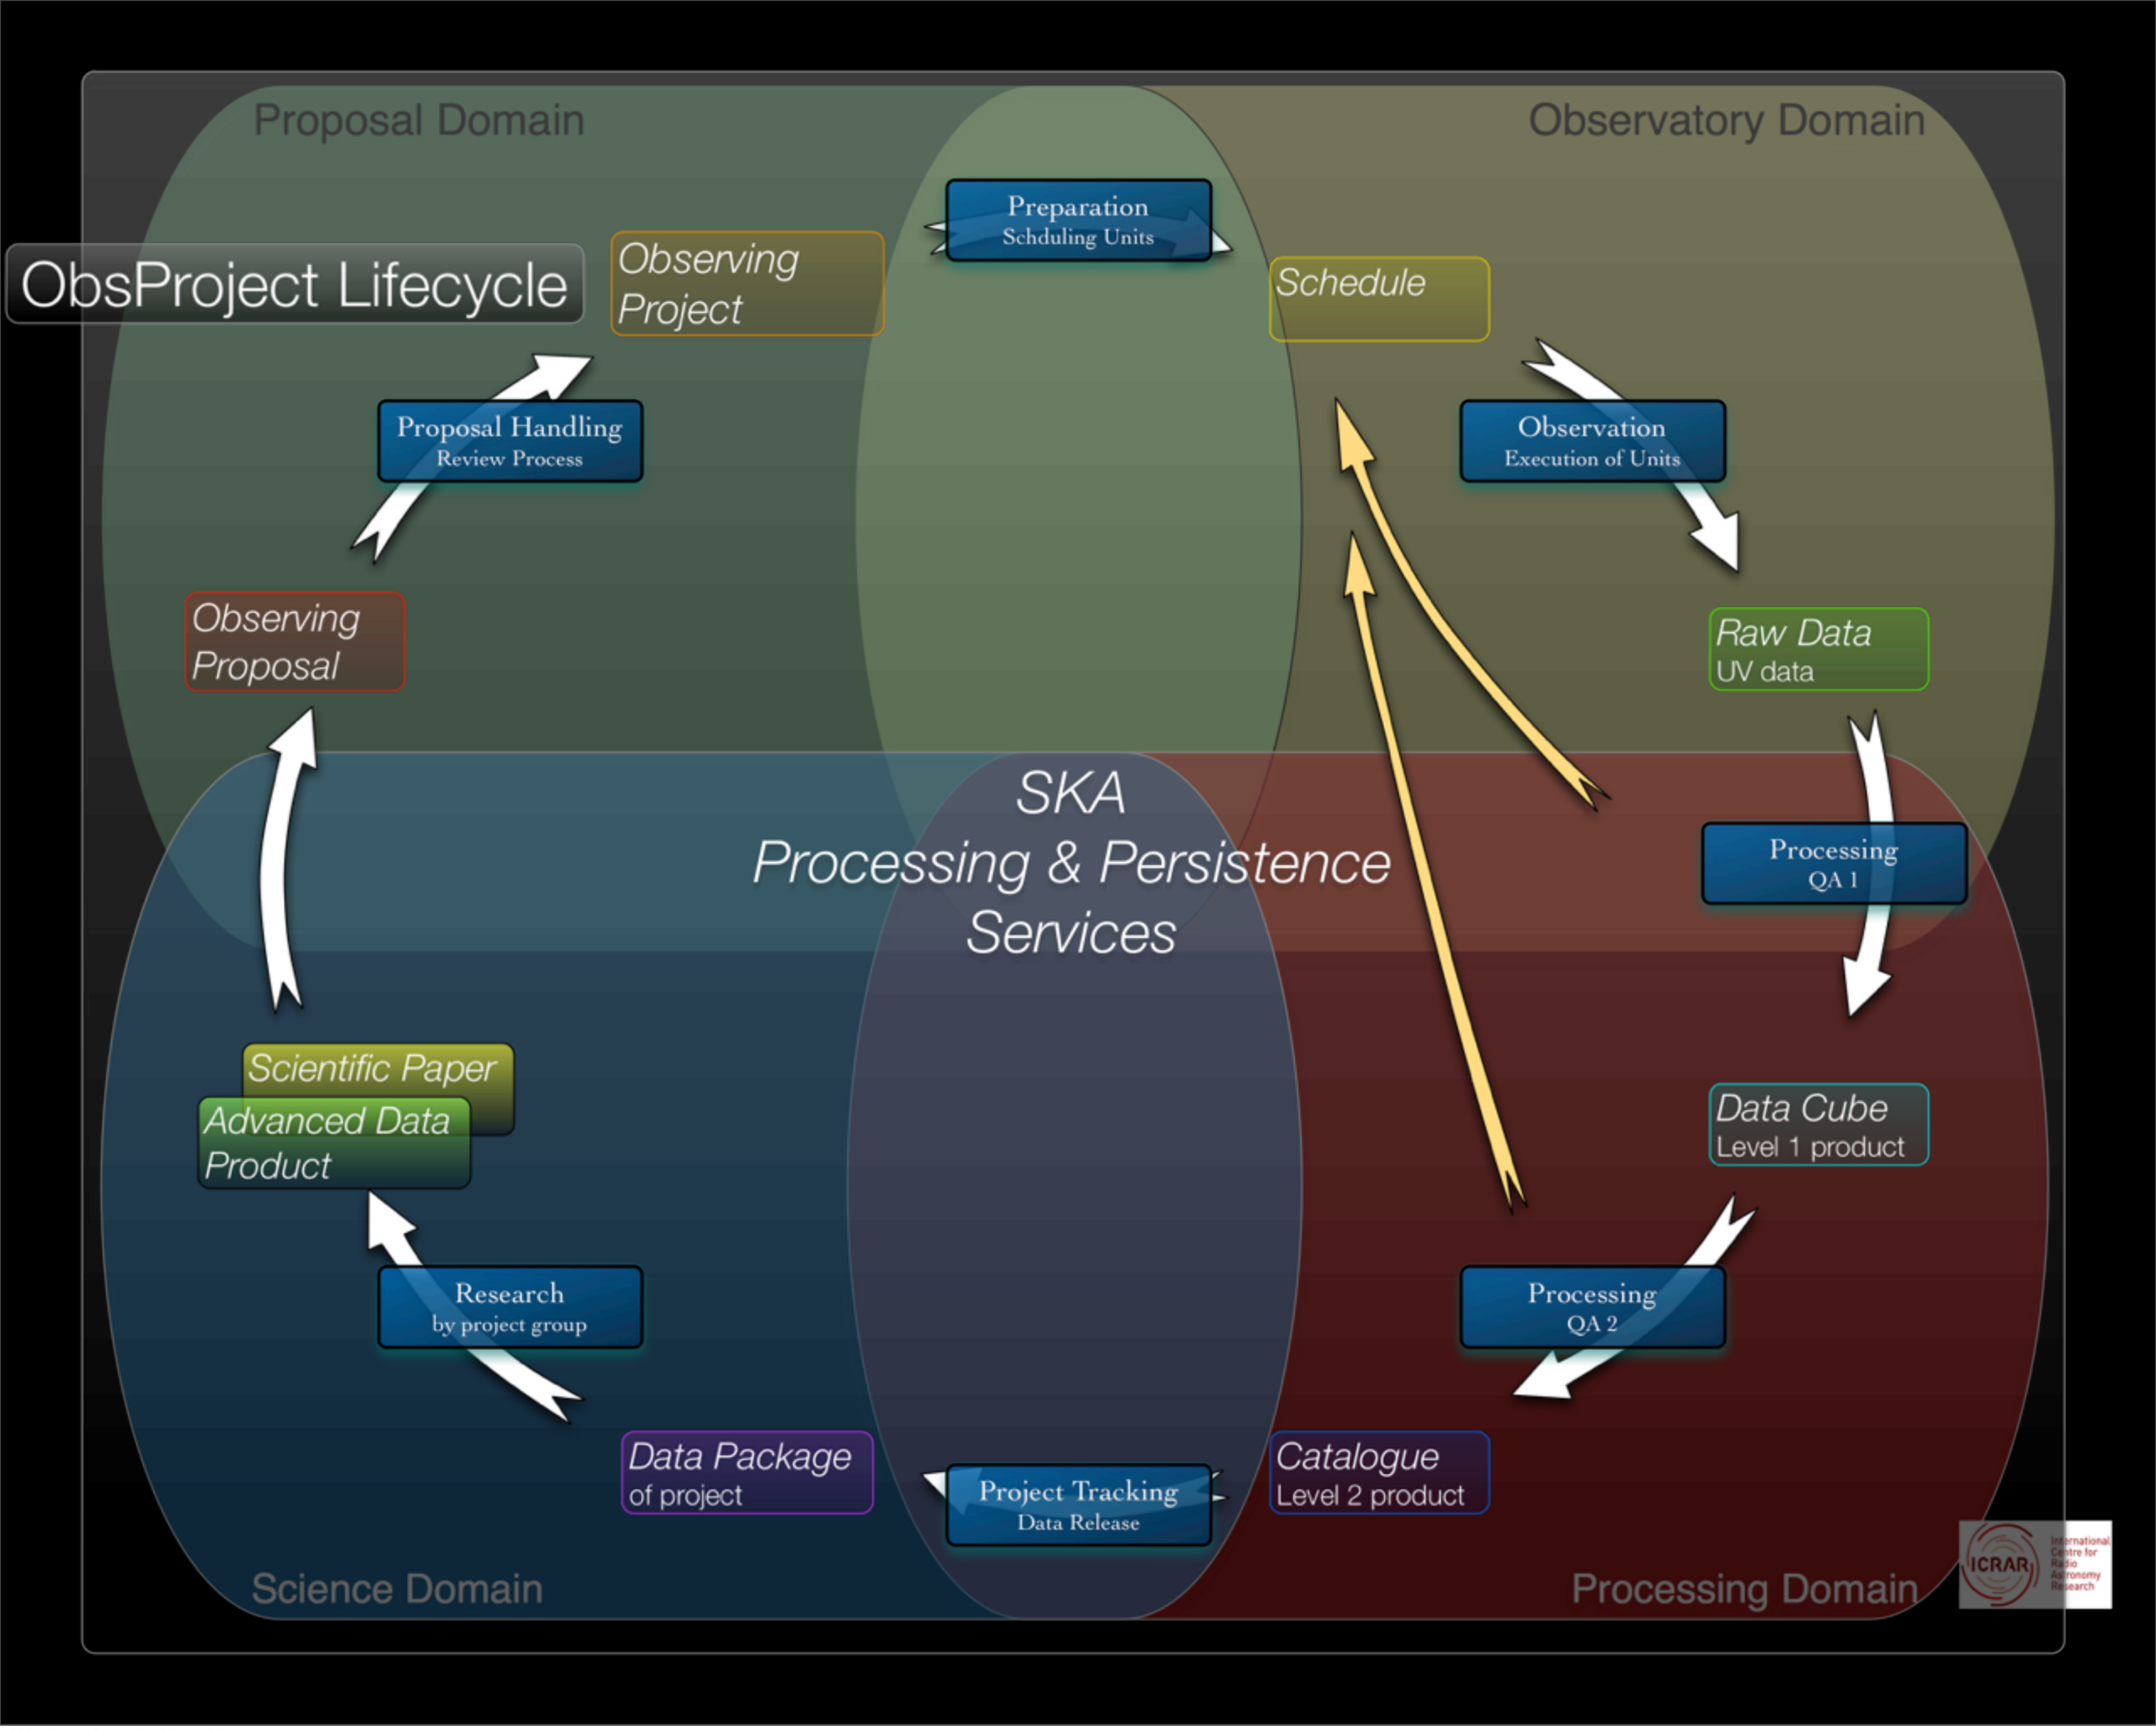
Proposal Domain
Observatory Domain
ObsProject Lifecycle
Observing Project
Preparation
Schduling Units
Schedule
Proposal Handling
Review Process
Observation
Execution of Units
Observing Proposal
Raw Data
UV data
SKA
Processing & Persistence
Services
Processing
QA 1
Scientific Paper
Advanced Data Product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

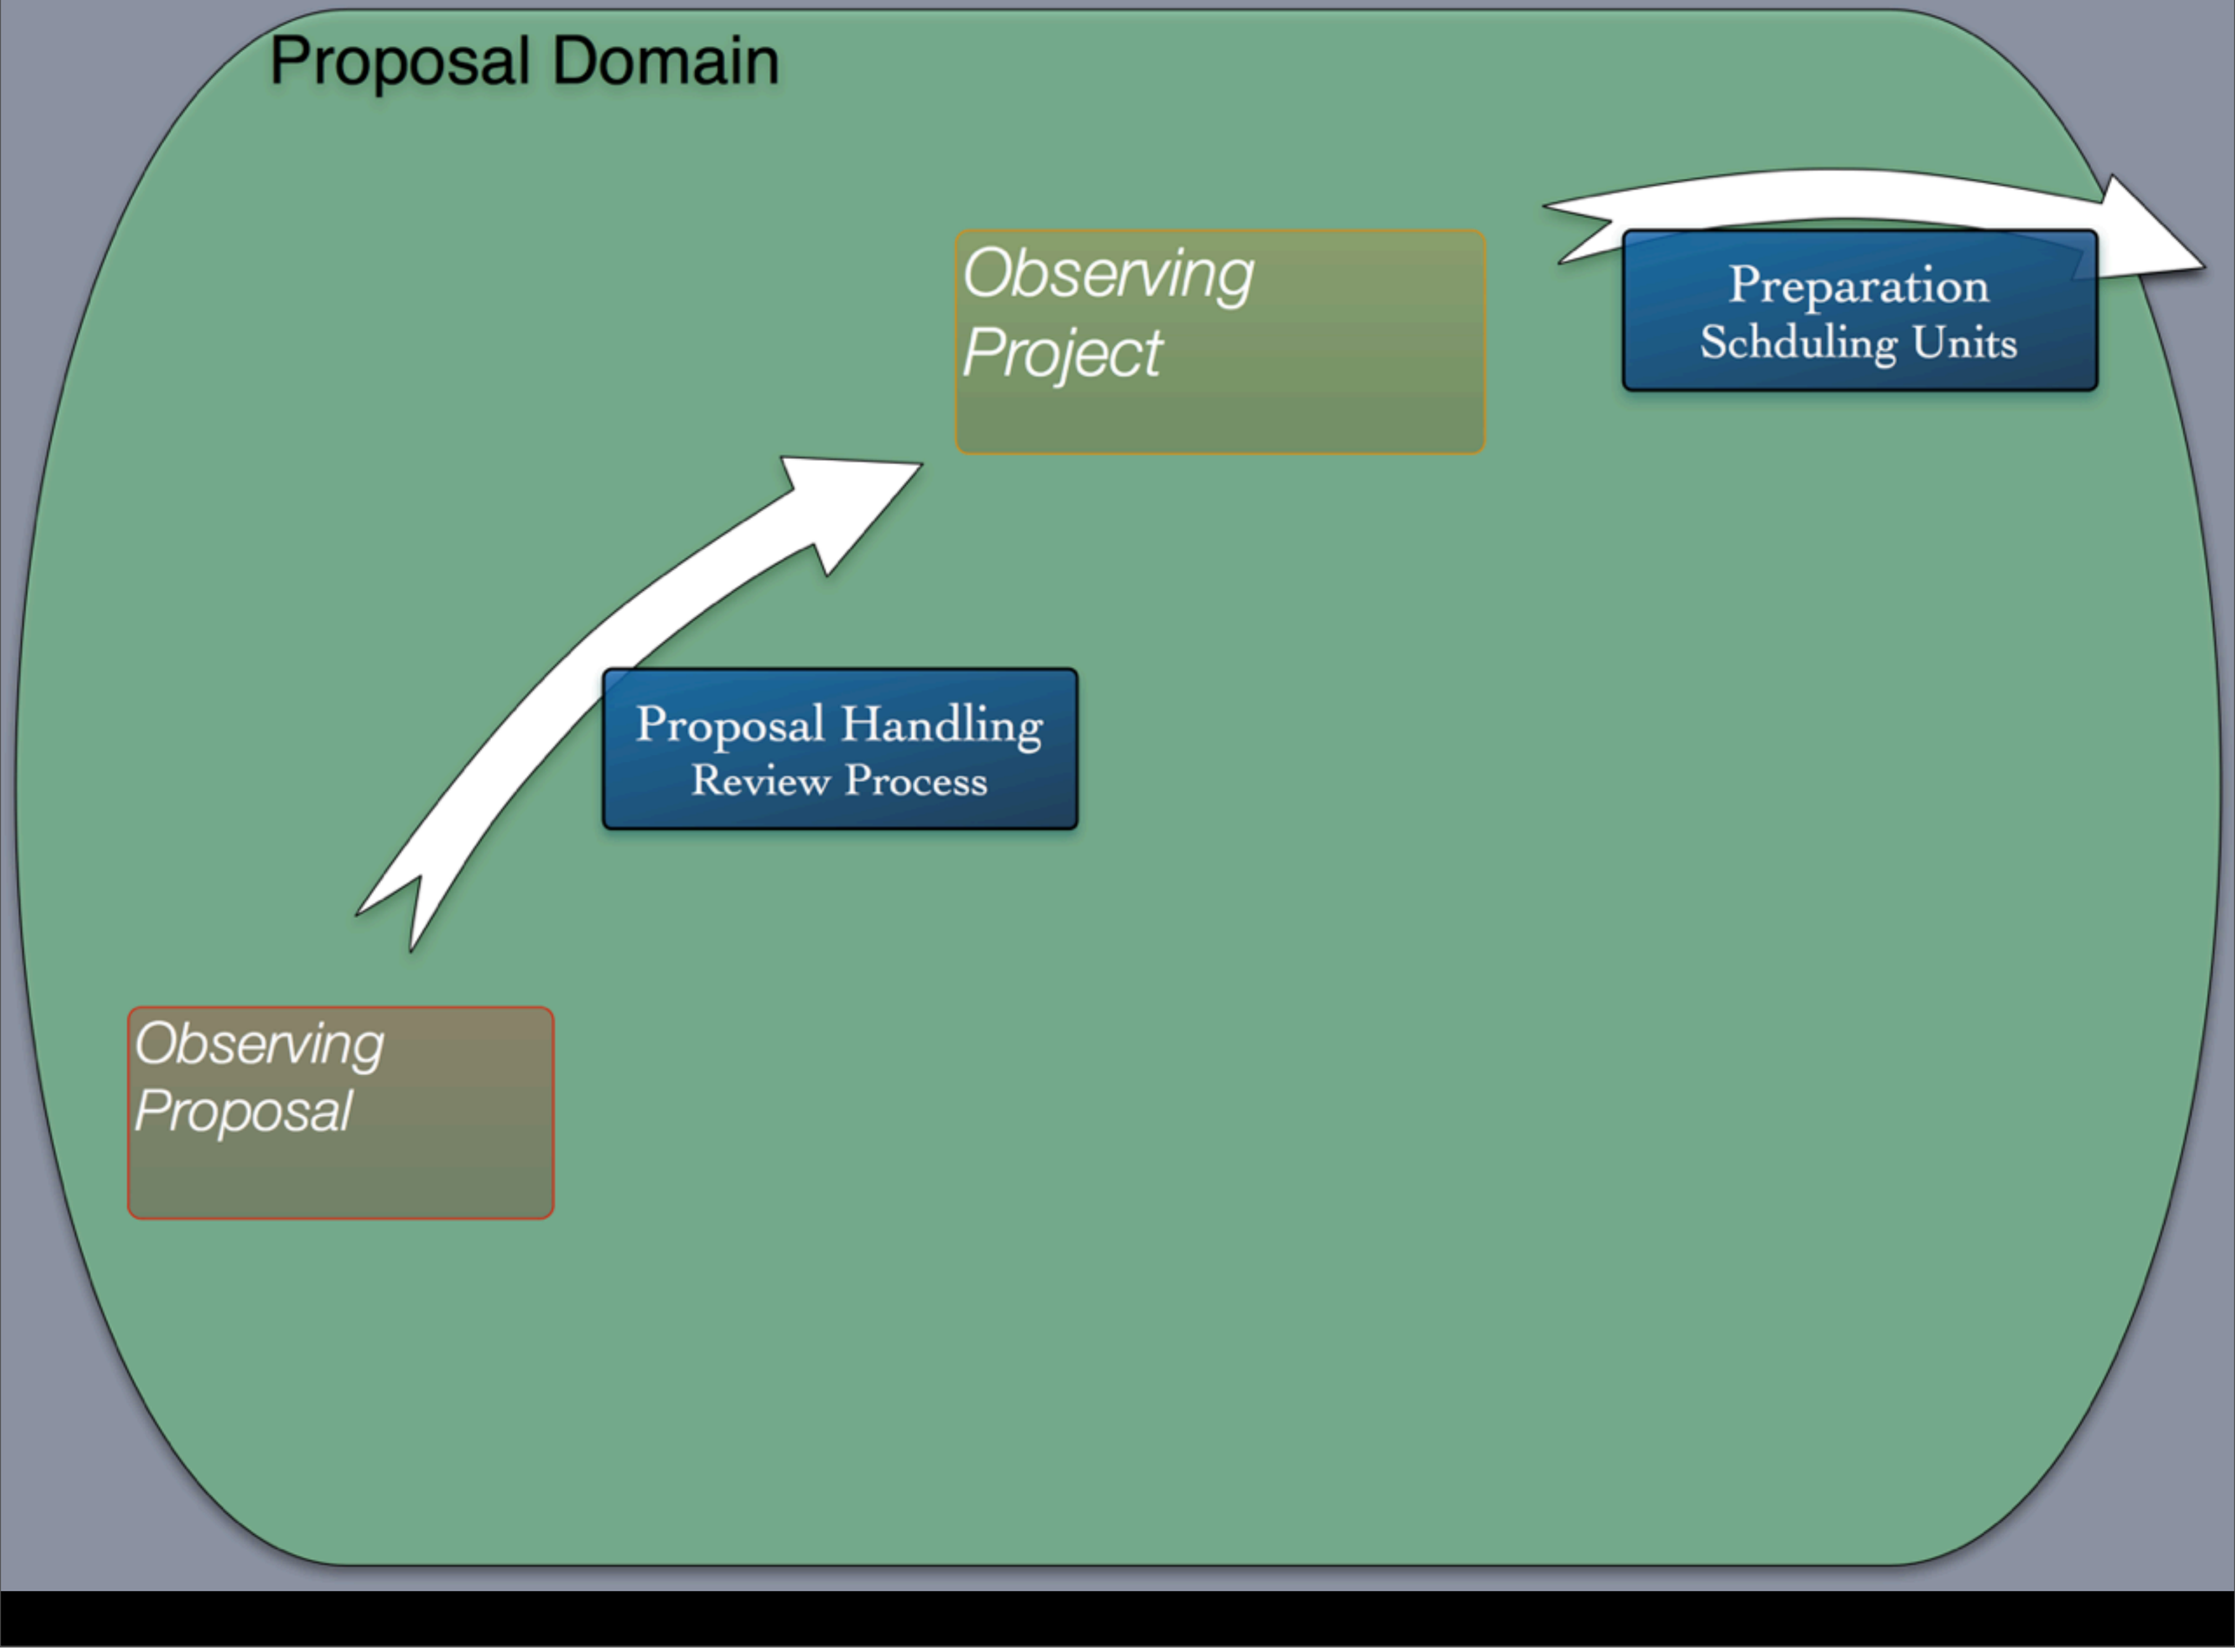
Proposal Domain
Observing Project
Preparation Schduling Units
Proposal Handling Review Process
Observing Proposal

Project State Model

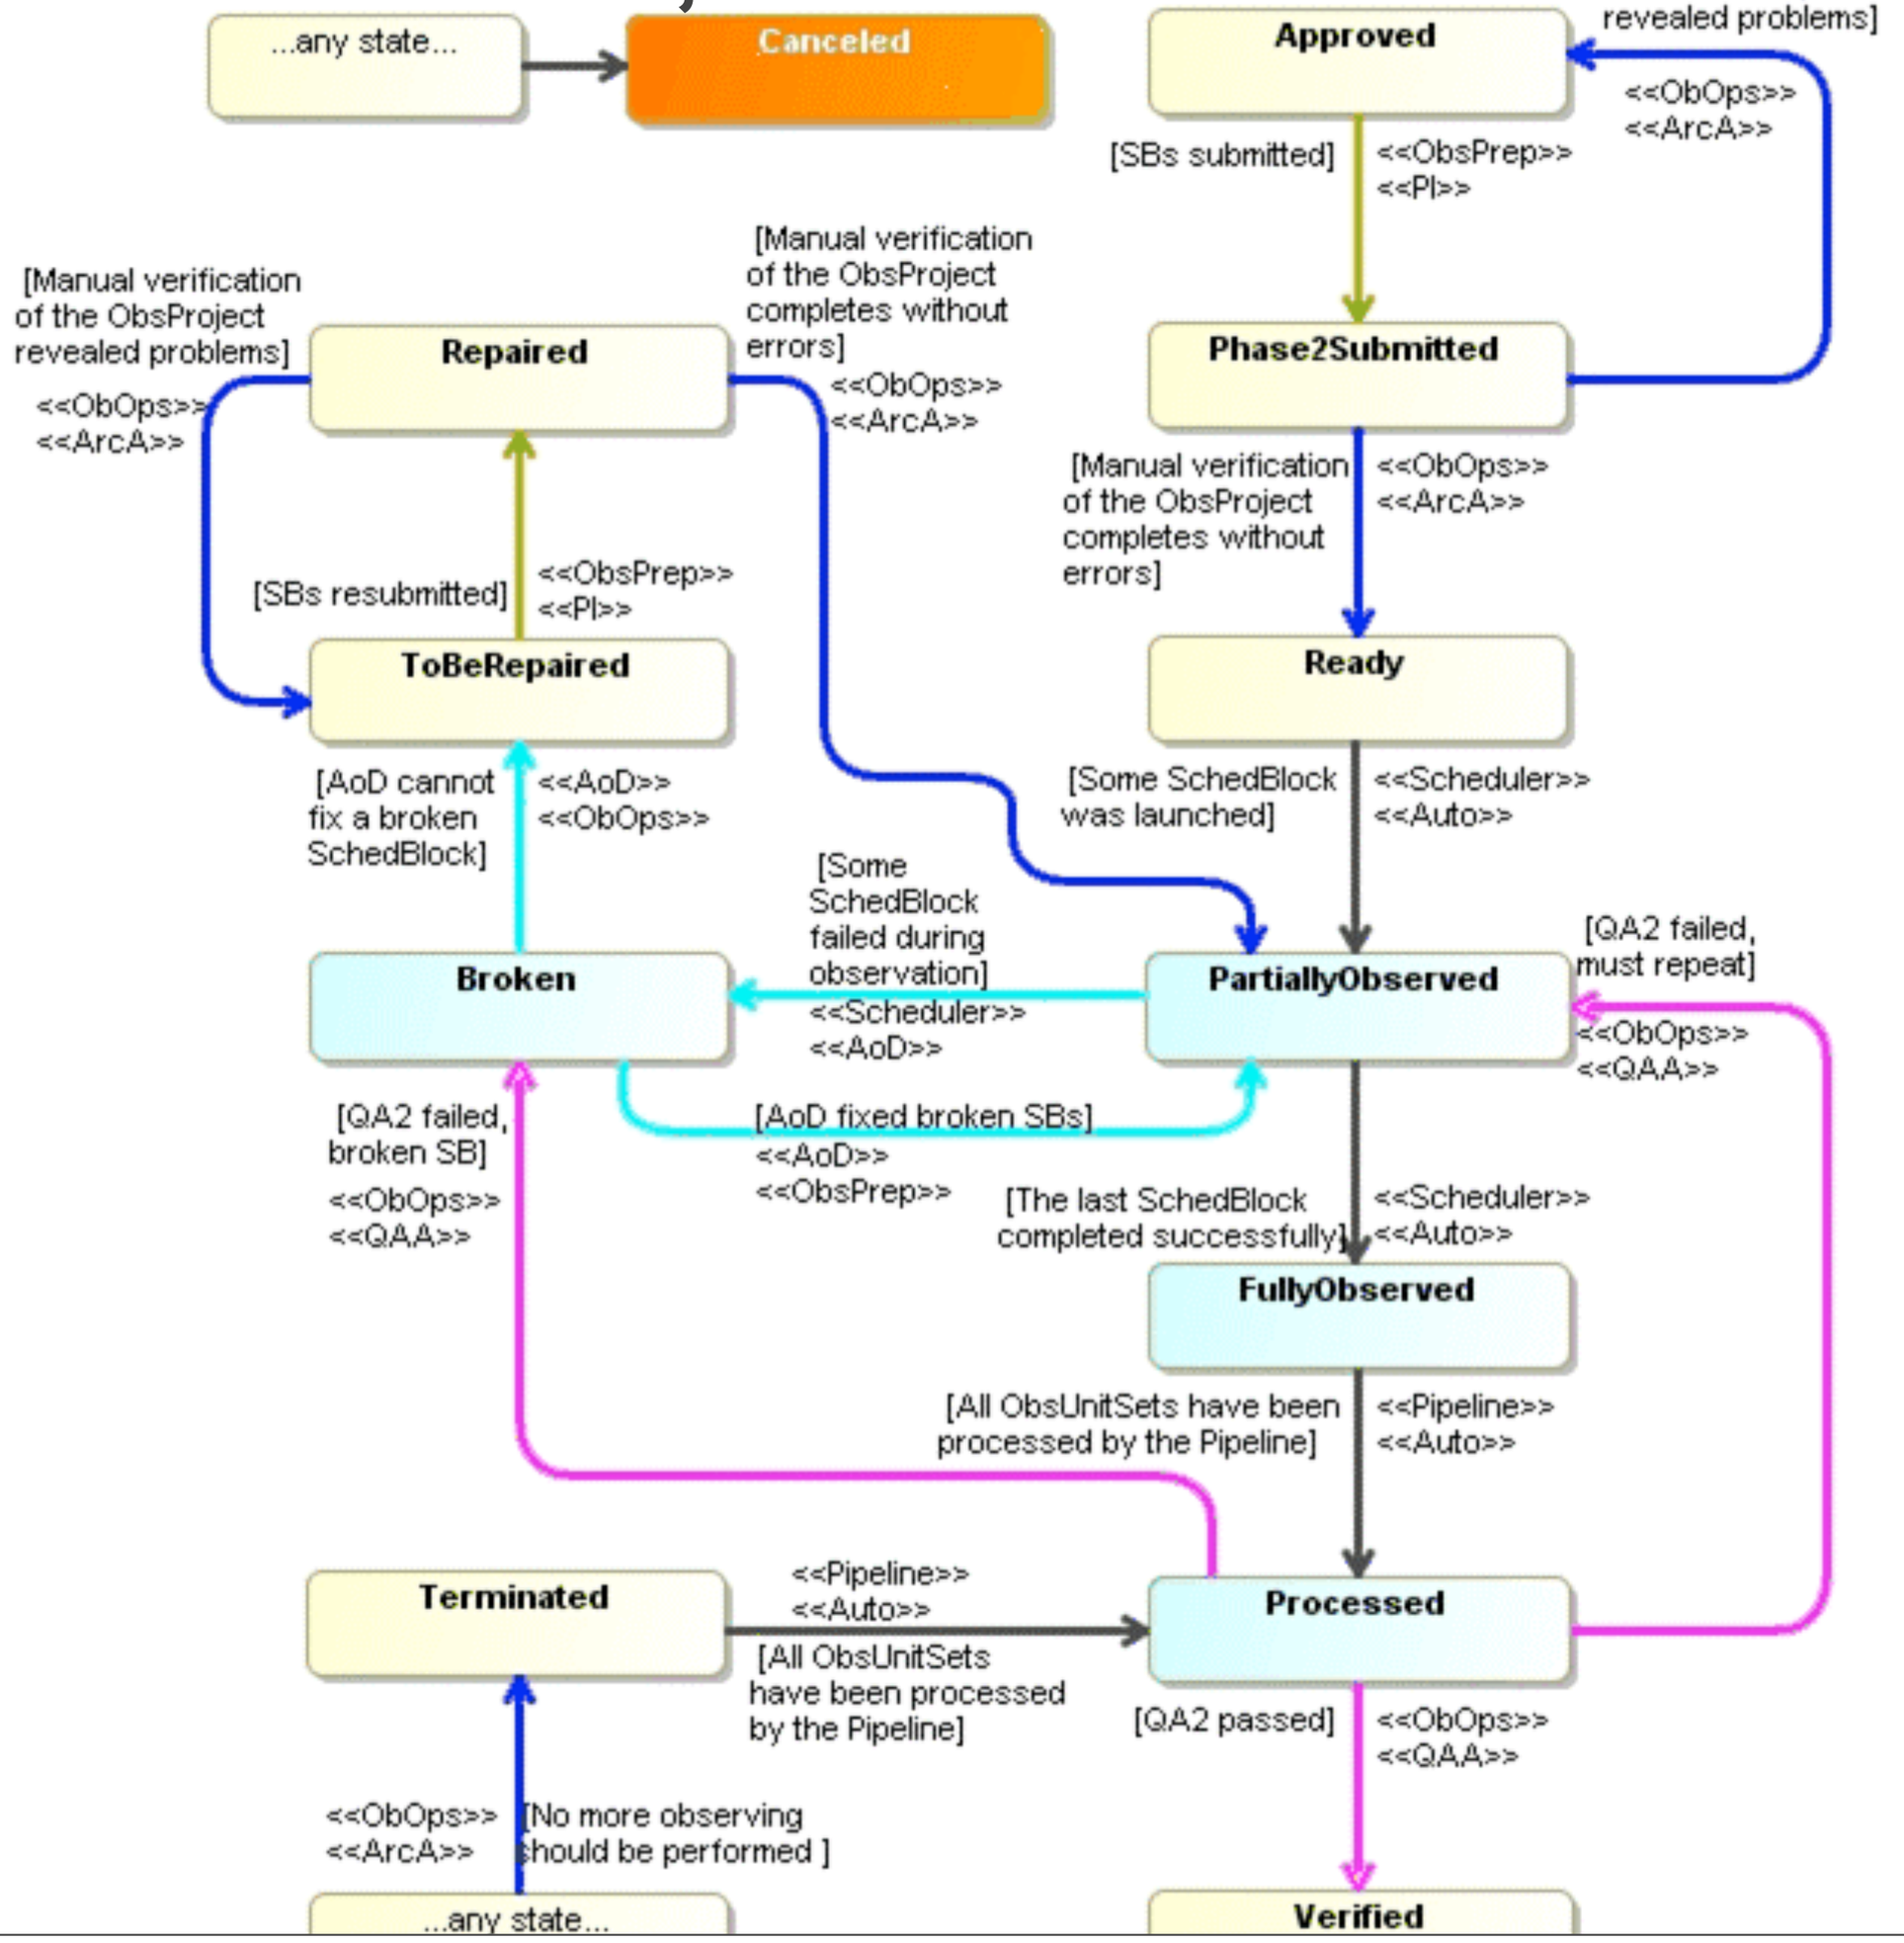

Observatory Domain

Preparation
Schduling Units

Schedule

Observation
Execution of Units

Raw Data
UV data

Processing
QA 1

SchedBlock State Model

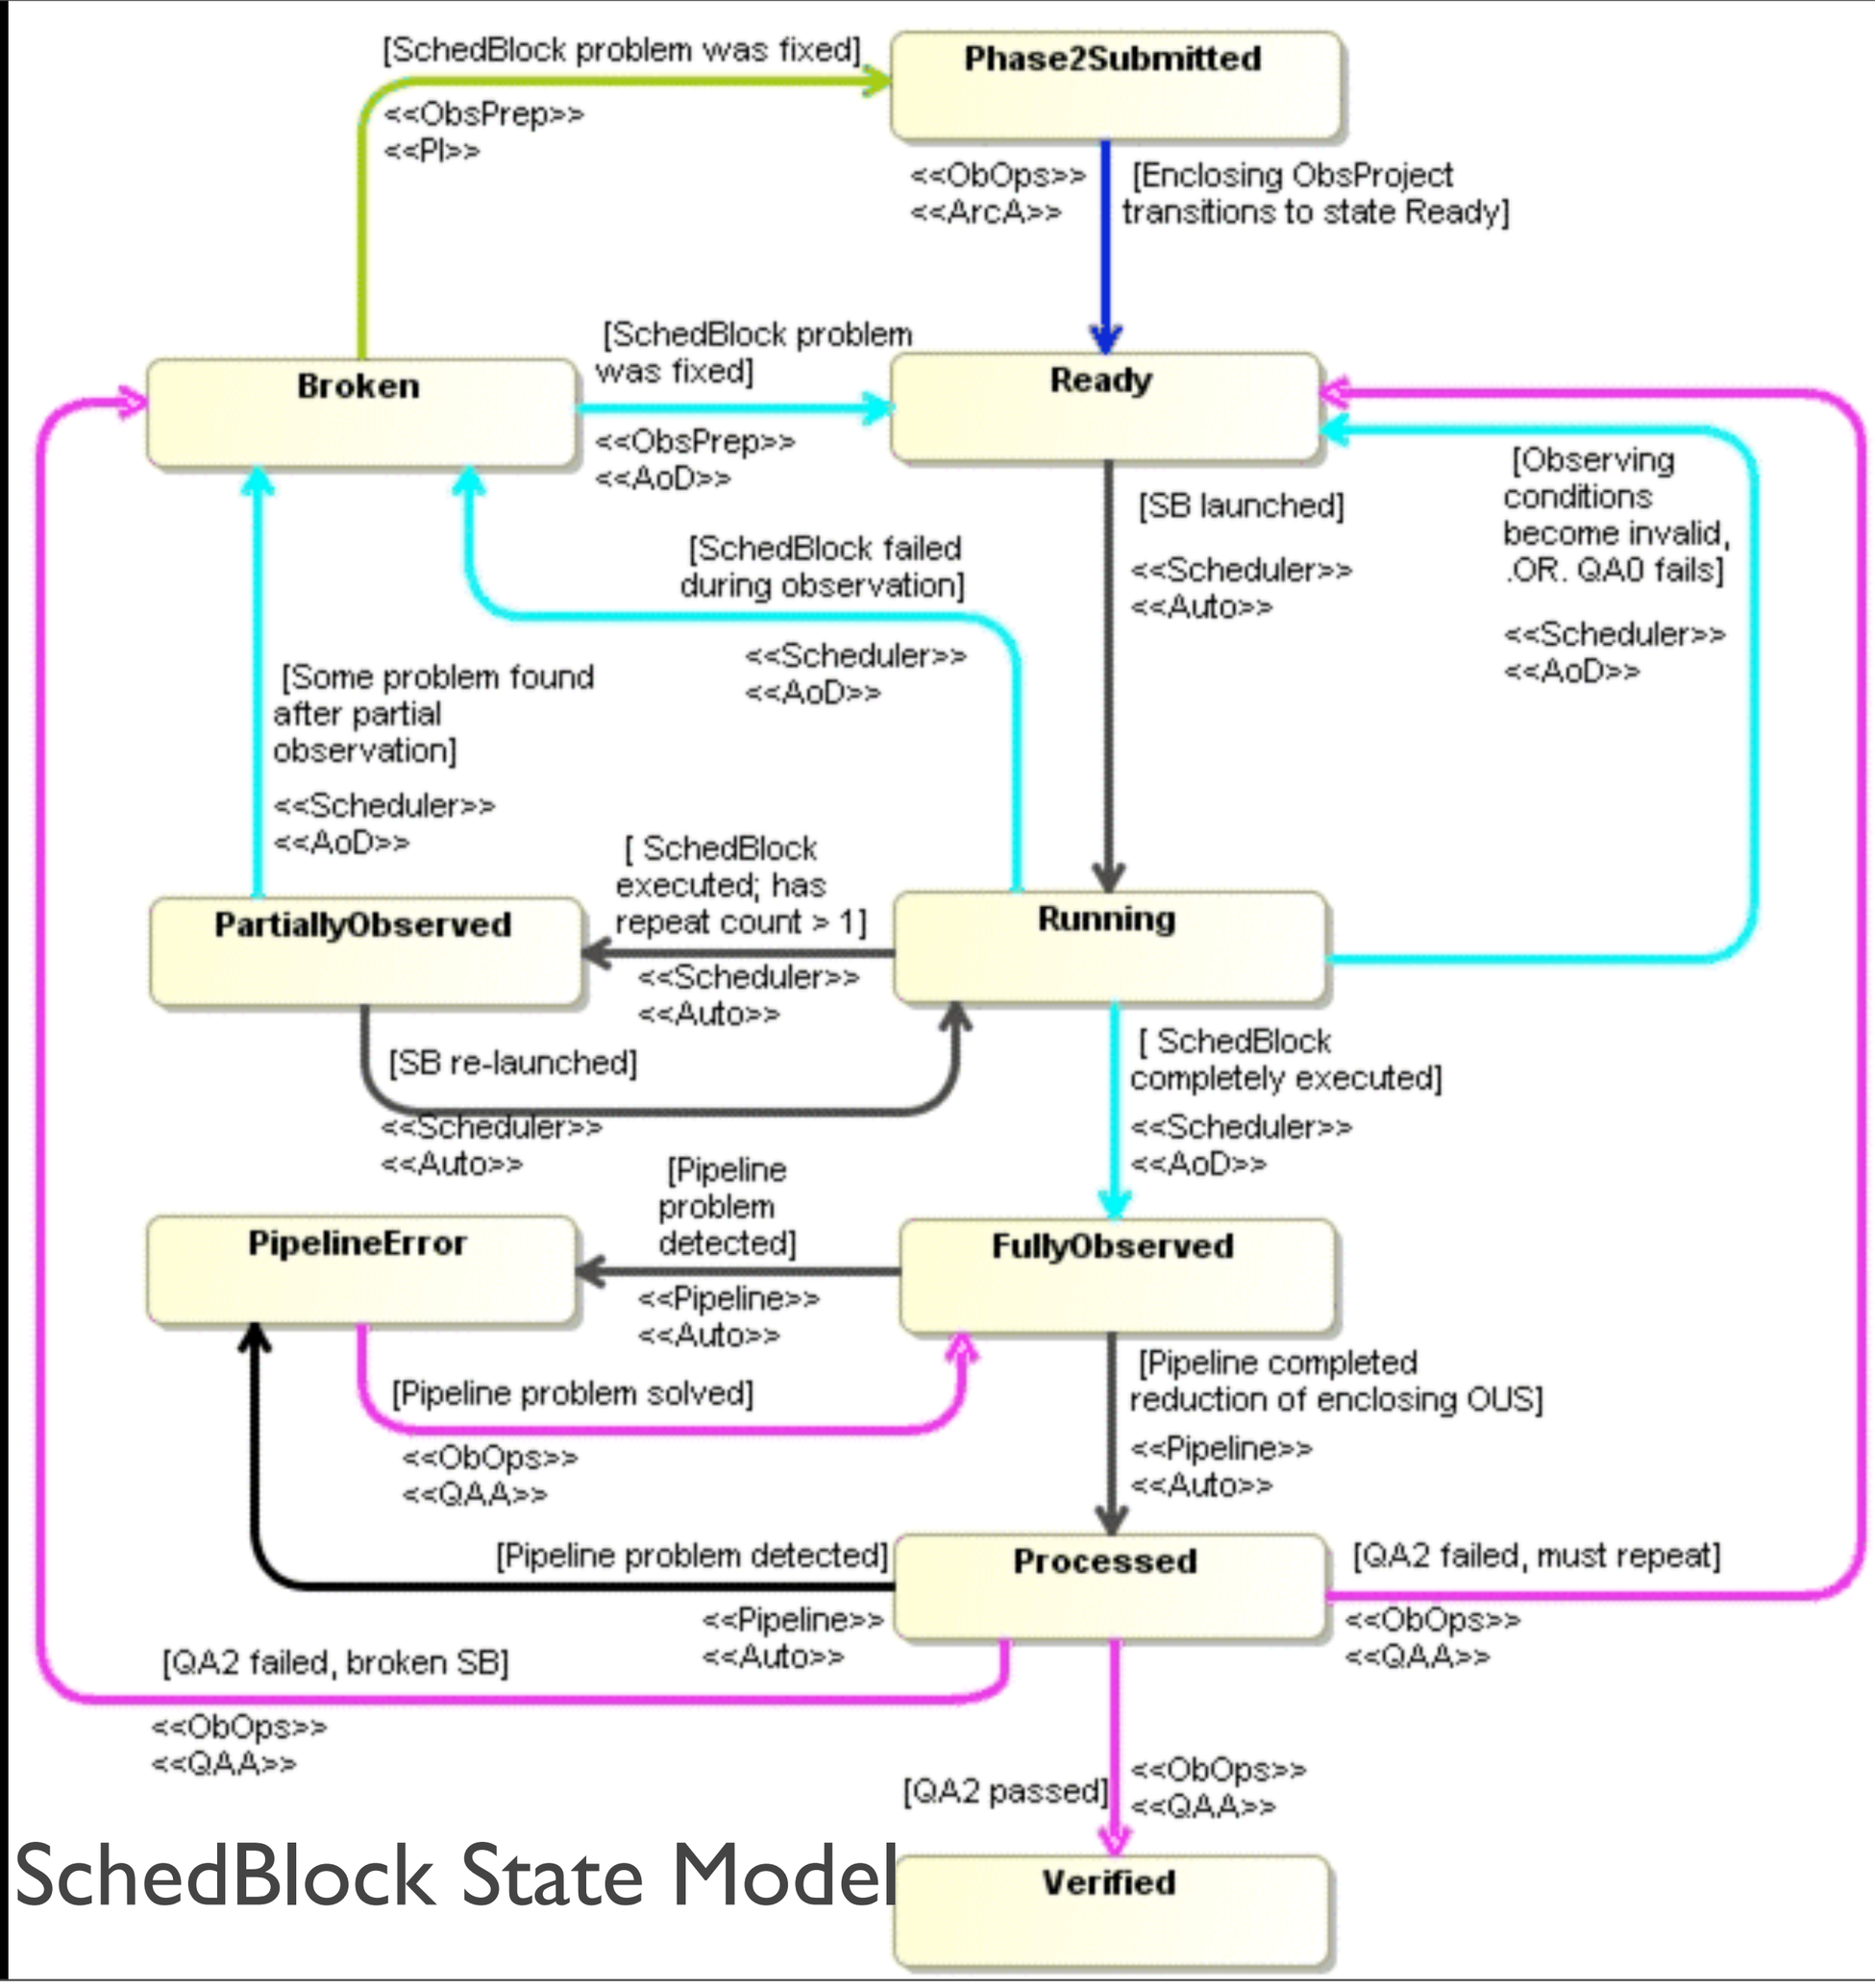

Processing
QA 1

Data Cube
Level 1 product

Processing
QA 2

Project Tracking
Data Release

Catalogue
Level 2 product

Processing Domain

DIRP Data Flow

Visualisation Use Case

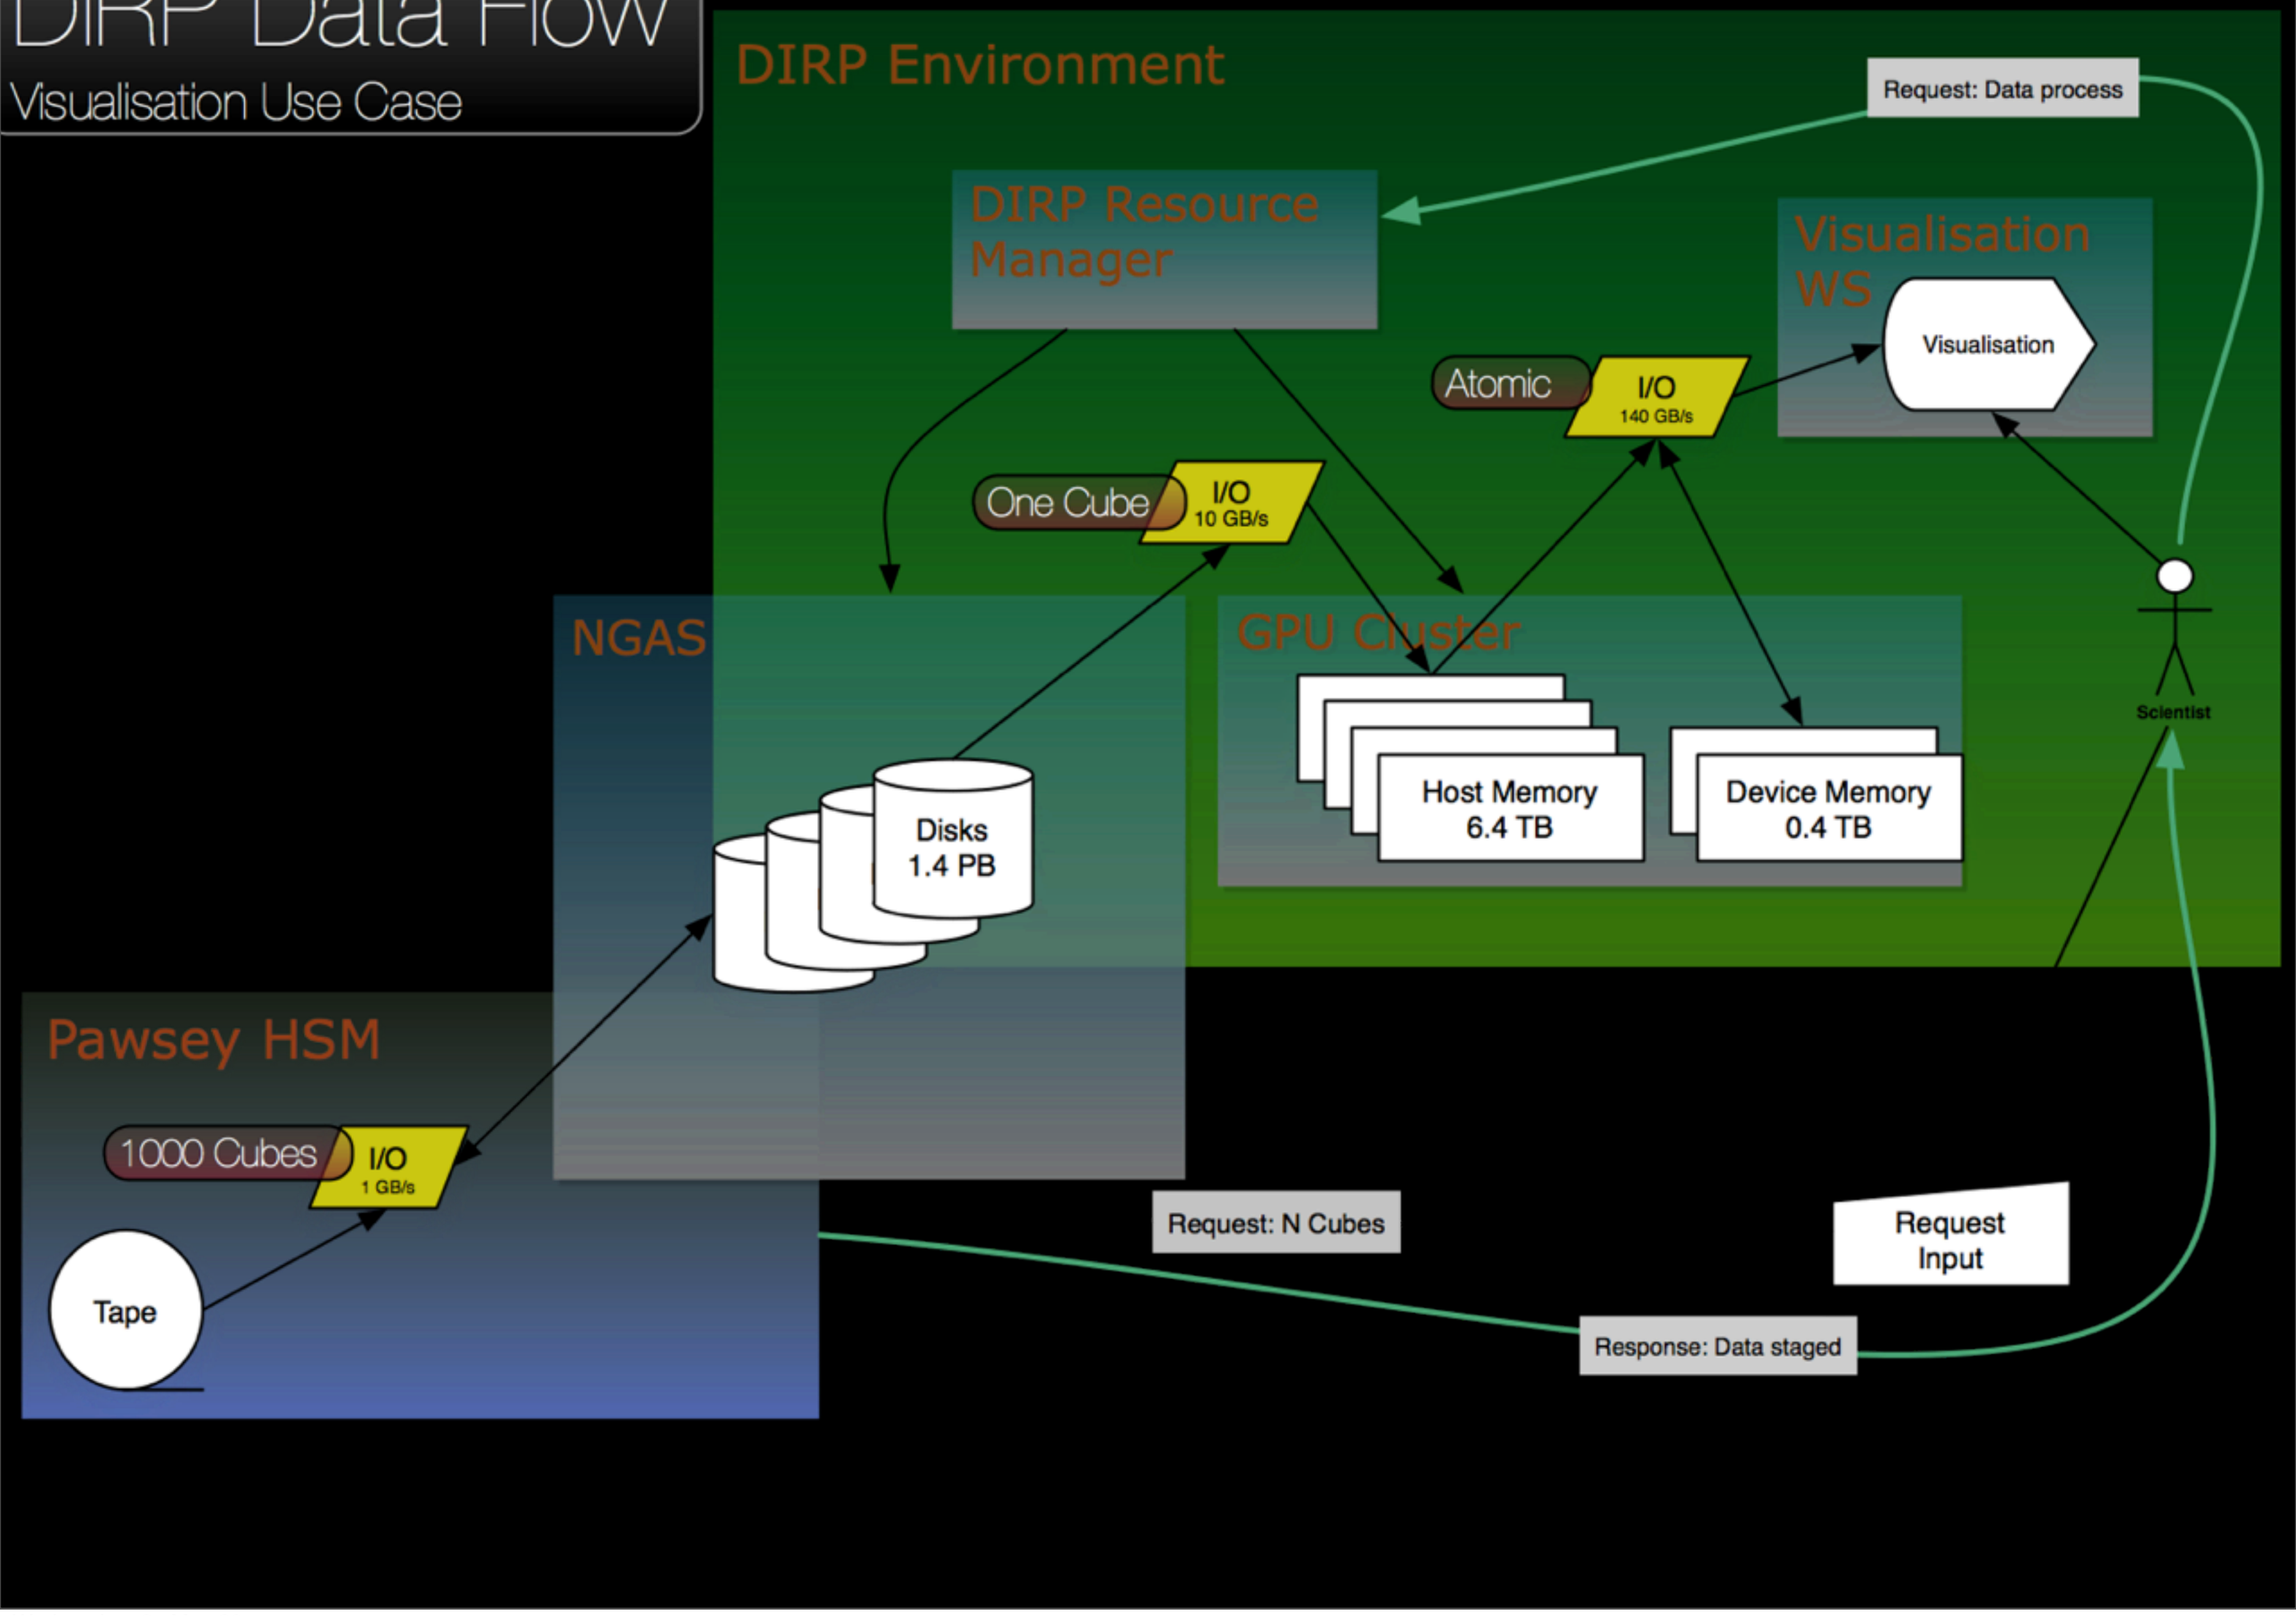

DIRP Data Flow

Visualisation Use Case

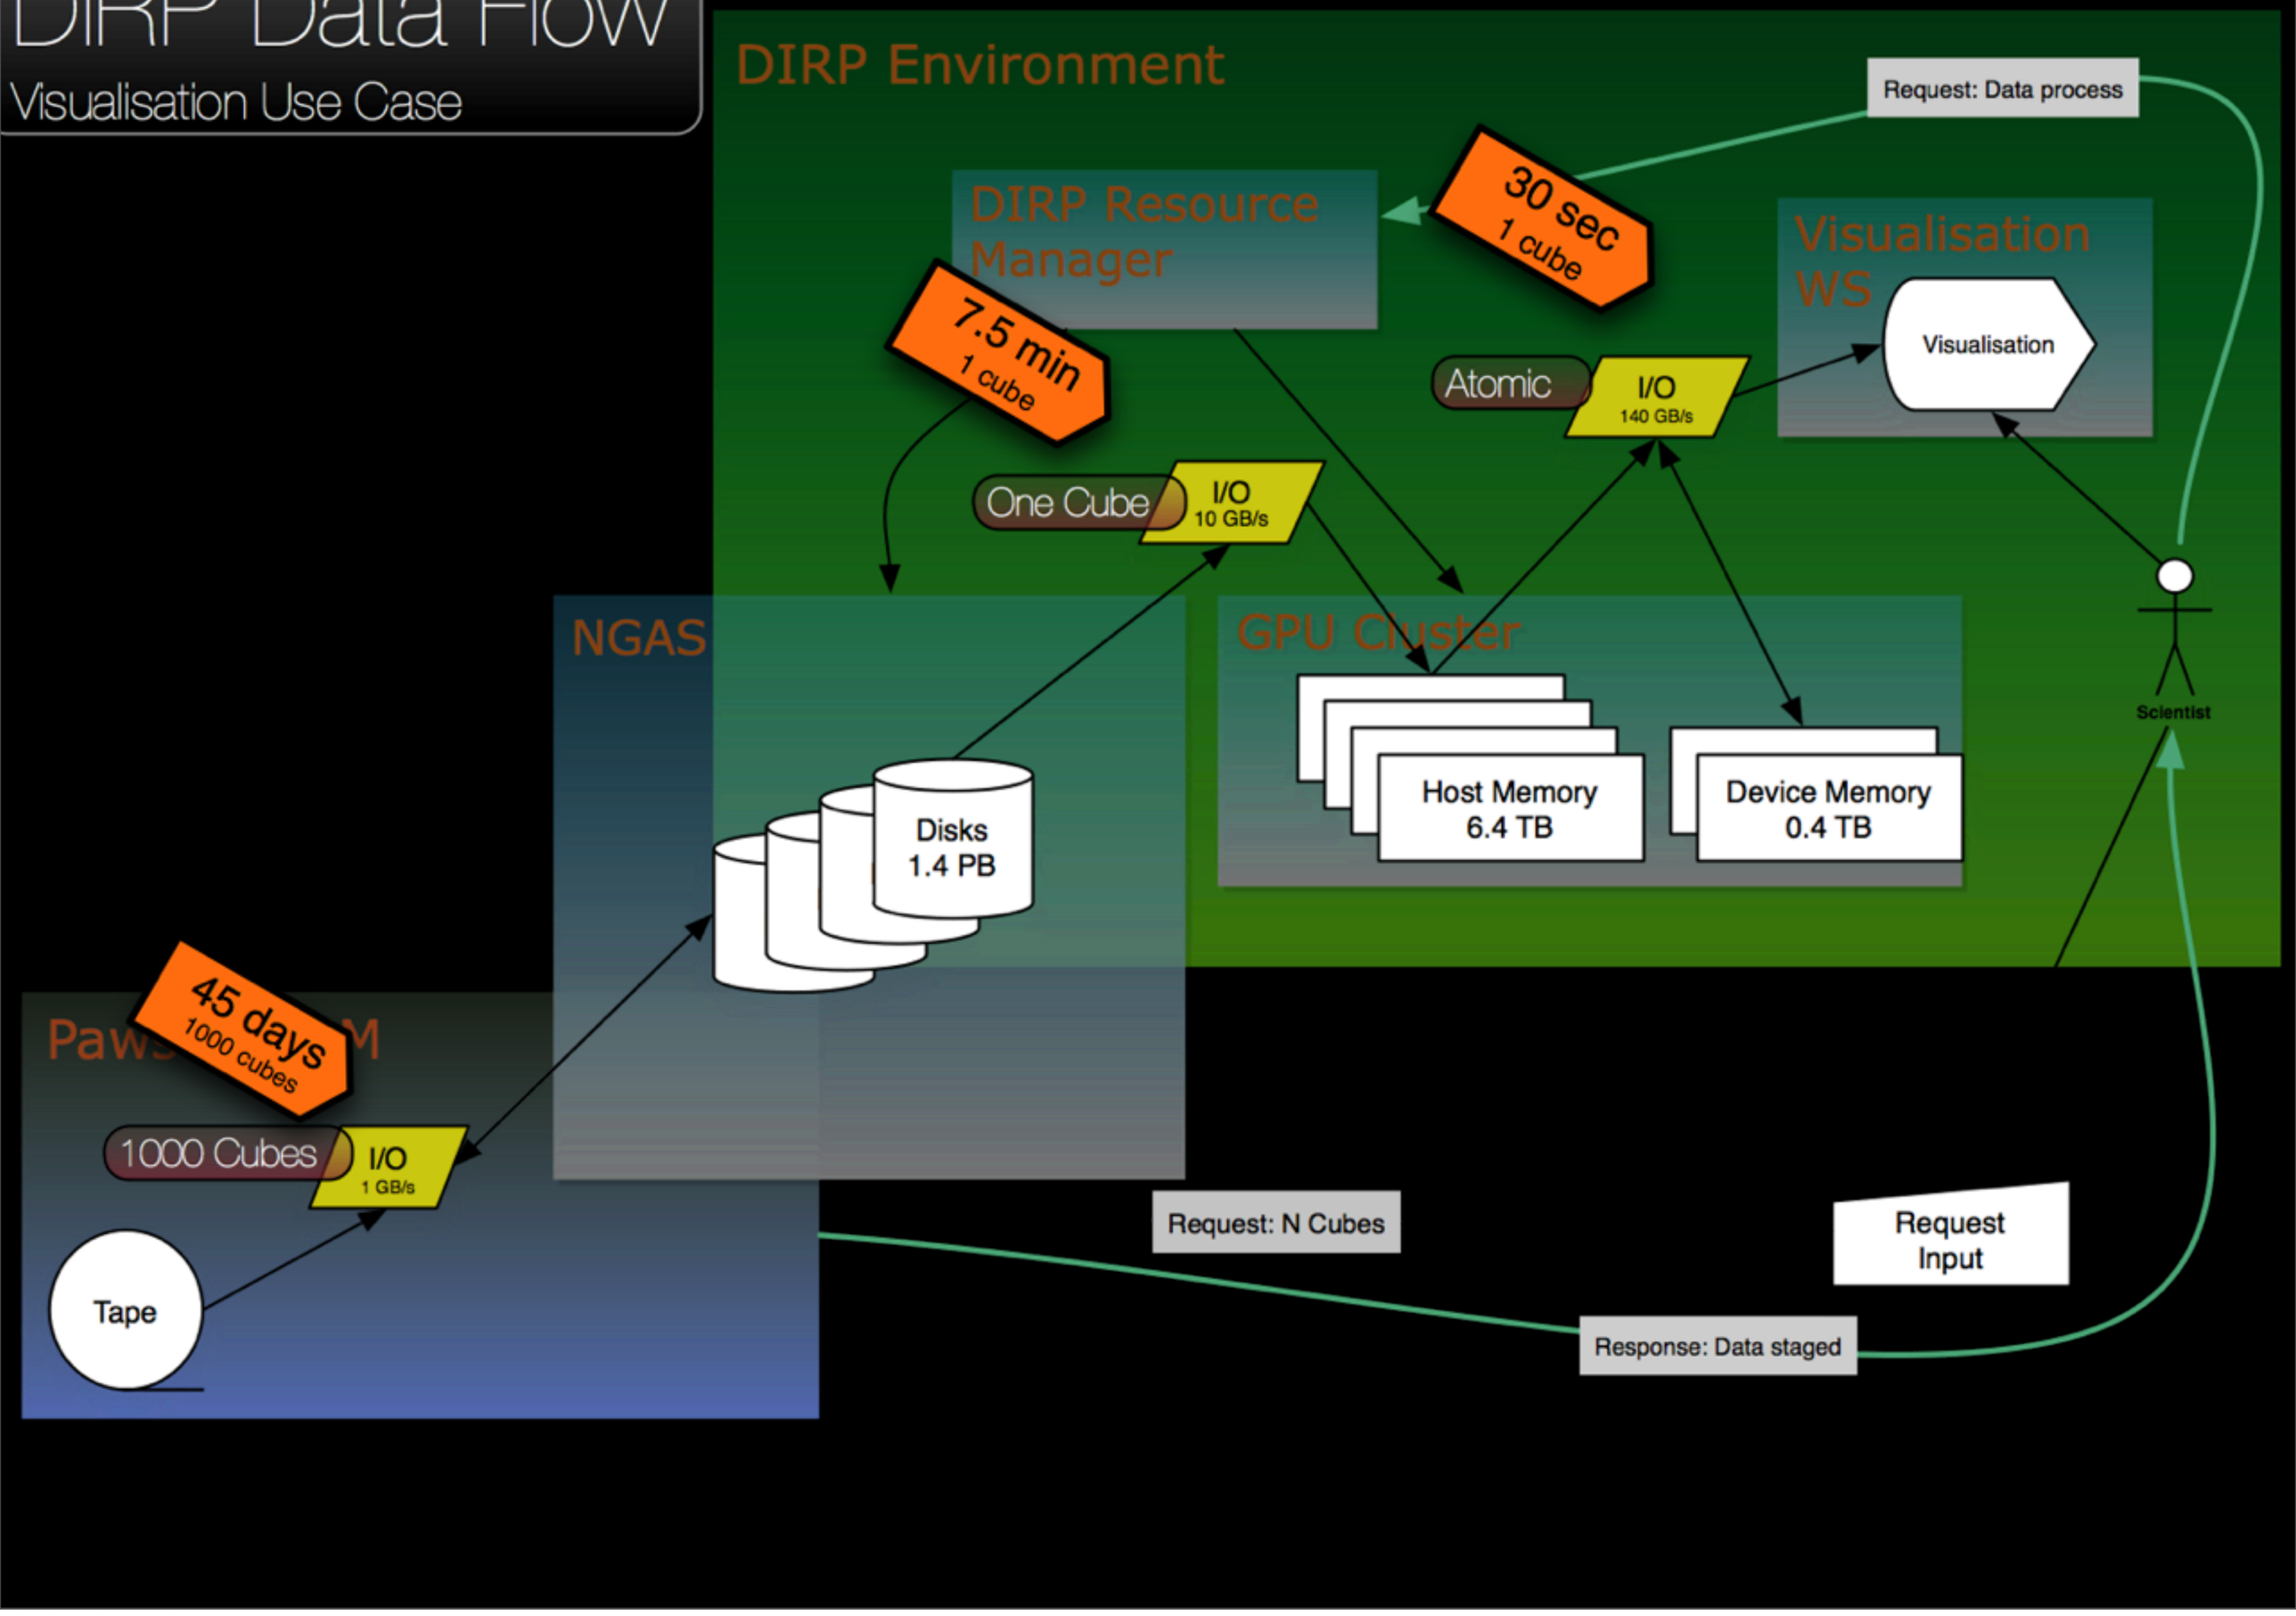

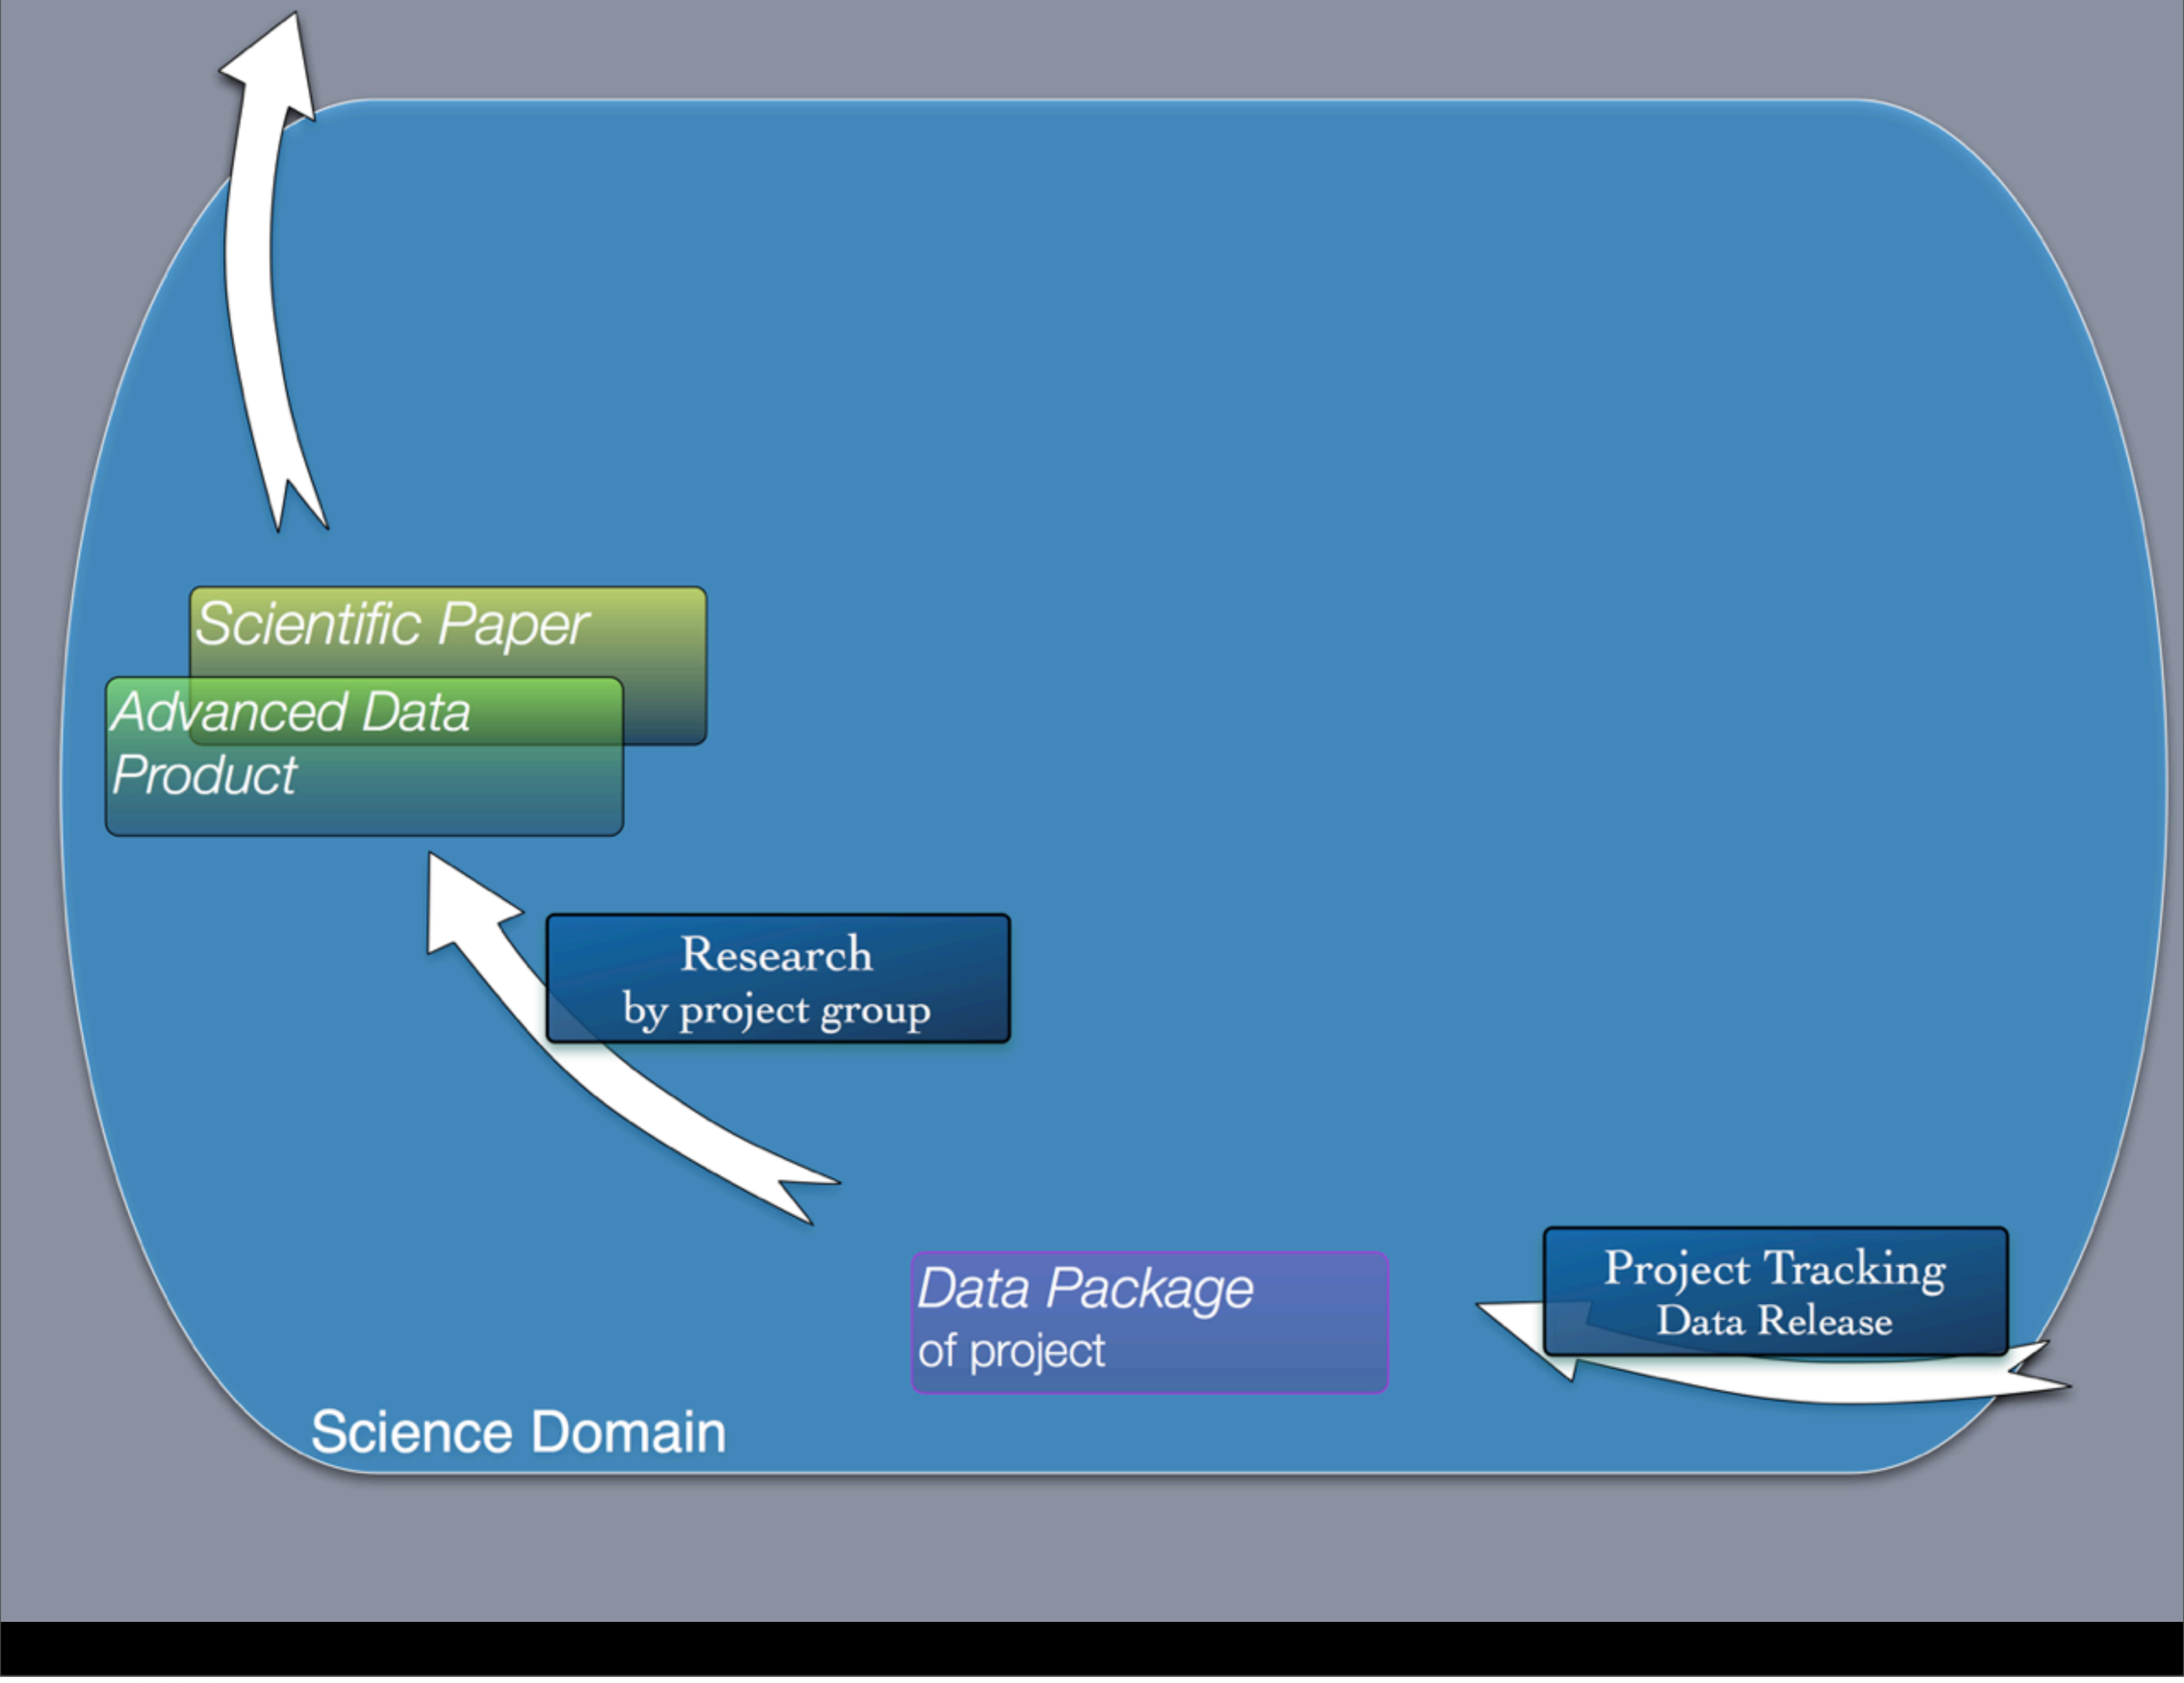
Scientific Paper
Advanced Data Product
Research
by project group
Data Package
of project
Project Tracking
Data Release
Science Domain

Science Projects

- Typical small science projects with SKA pre-cursers are many TB. Surveys are many PB.
- SKA projects likely to be a few orders of magnitude more.
- Large scale data reduction and scientific analysis only possible with dedicated vast computing resources.
- No data delivery of complete projects to individuals?
- Where are the data products of the projects stored?

Connectivity of Data Centers

- Scientific exploitation of first level data products requires many FLOPs somewhere accessible to the science teams.
- In general this is not a one-shot process, but requires multiple runs to optimize the parameters used.
- Quite different usage of HPC centers, need of research in HPC scheduling
- HPC inter-connectivity and connection between HSM (tape??) storage and processing has to be balanced.
- Much higher Amdahl number required.

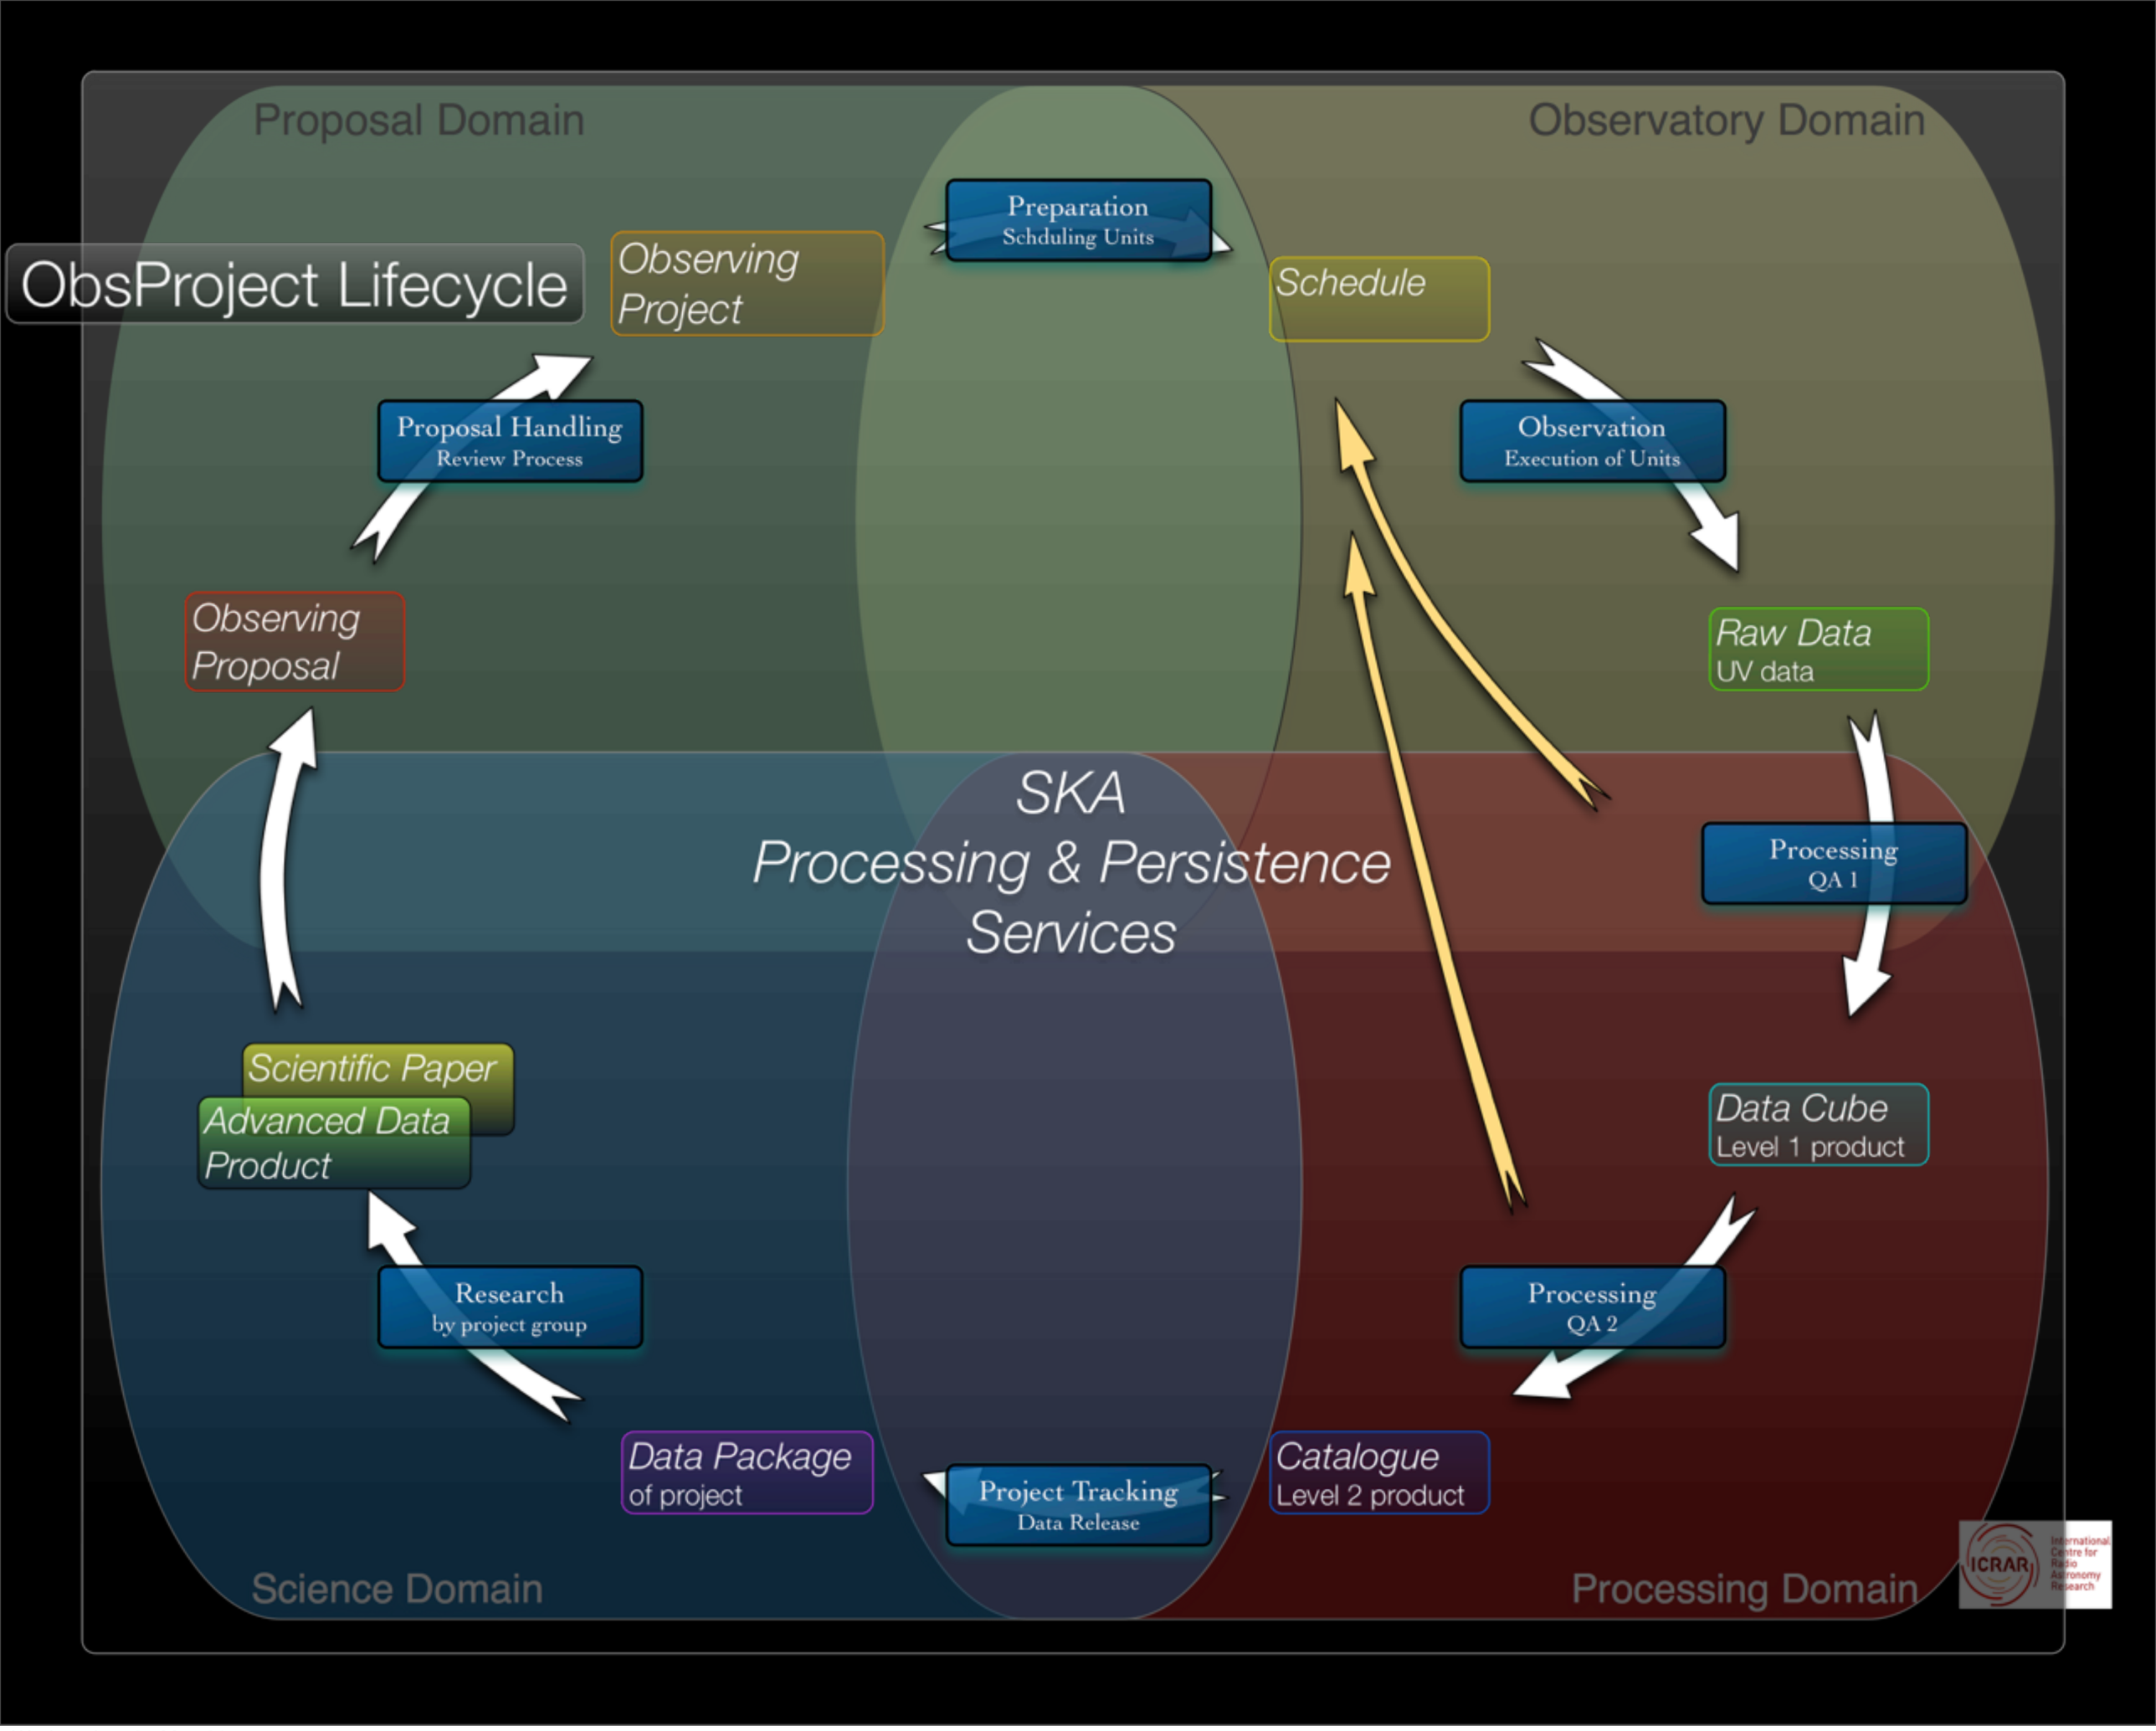
Proposal Domain
Observatory Domain
ObsProject Lifecycle
Observing Project
Preparation
Schduling Units
Schedule
Proposal Handling
Review Process
Observation
Execution of Units
Observing Proposal
Raw Data
UV data
SKA
Processing & Persistence
Services
Processing
QA 1
Scientific Paper
Advanced Data Product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

SKA
Processing & Persistence
Services

Powering the beast

- Storage power consumption 2-5 Megawatt/Exabyte

- Computing power consumption 350 MW/Exaflop with current technology

- Need GREEN solutions and better power efficiency

Persistence Layer

- Costs of persisting the operationally critical data could easily dominate the total operational costs: Need to define and implement detailed life-cycle for all data categories.
- Rigorous control of monitoring and logging activities to avoid system overload. Separate monitoring and logging from science data persistence.
- Costing, priority and operational model for long-term archiving of science data required.

Common Software

Common Software

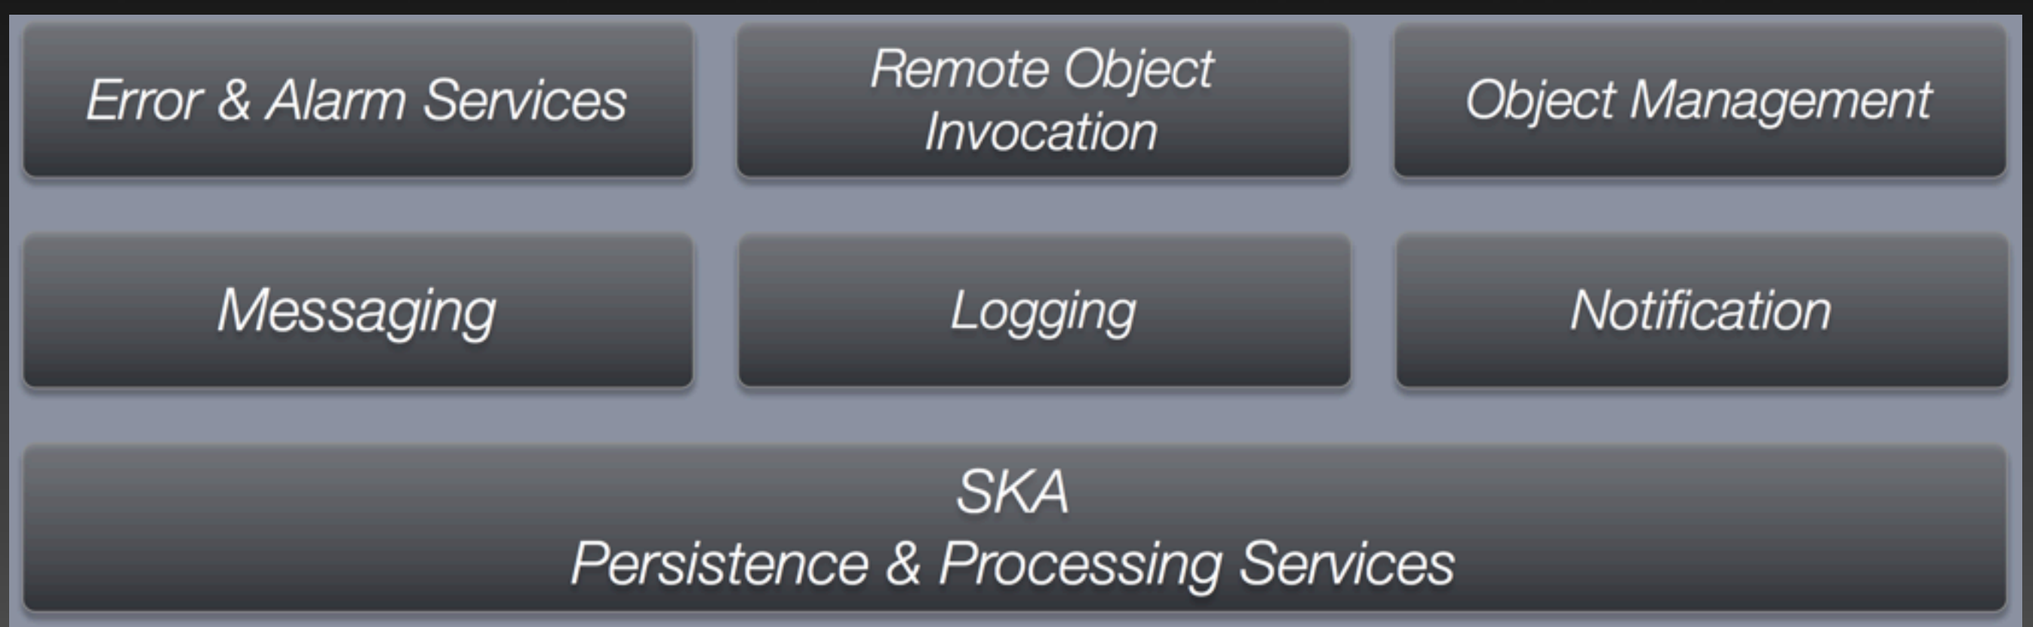

Common Software

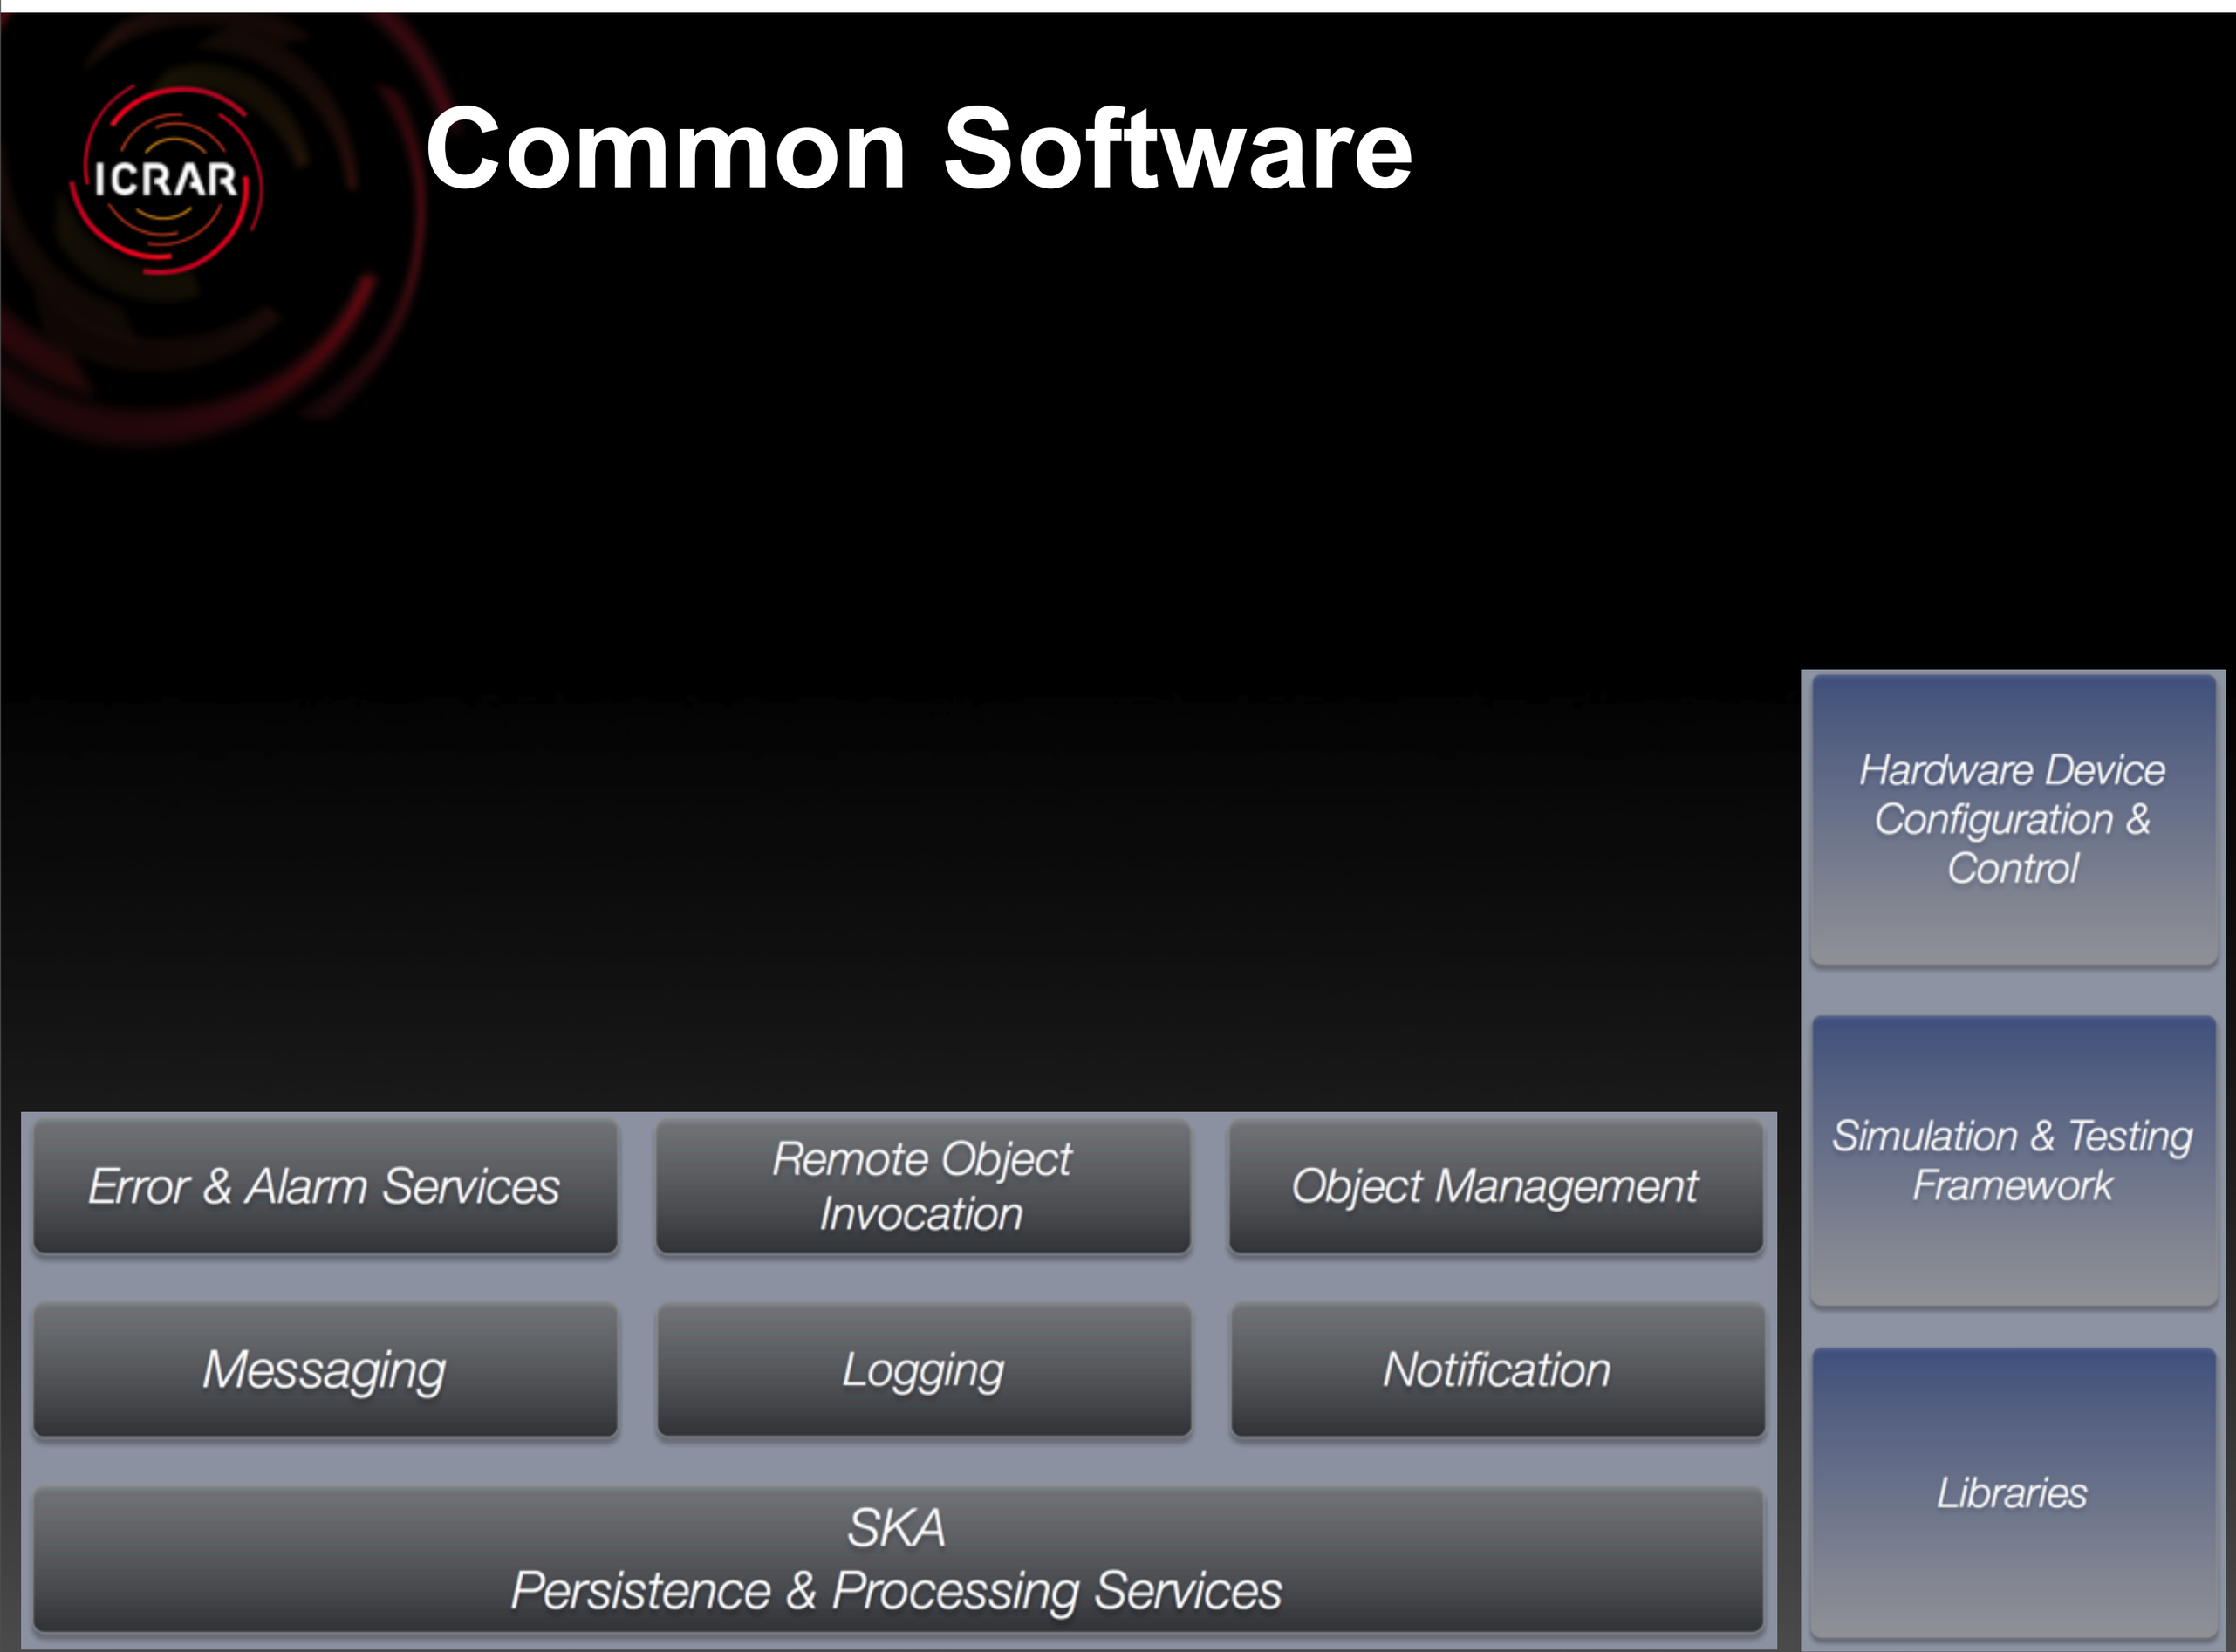

Common Software

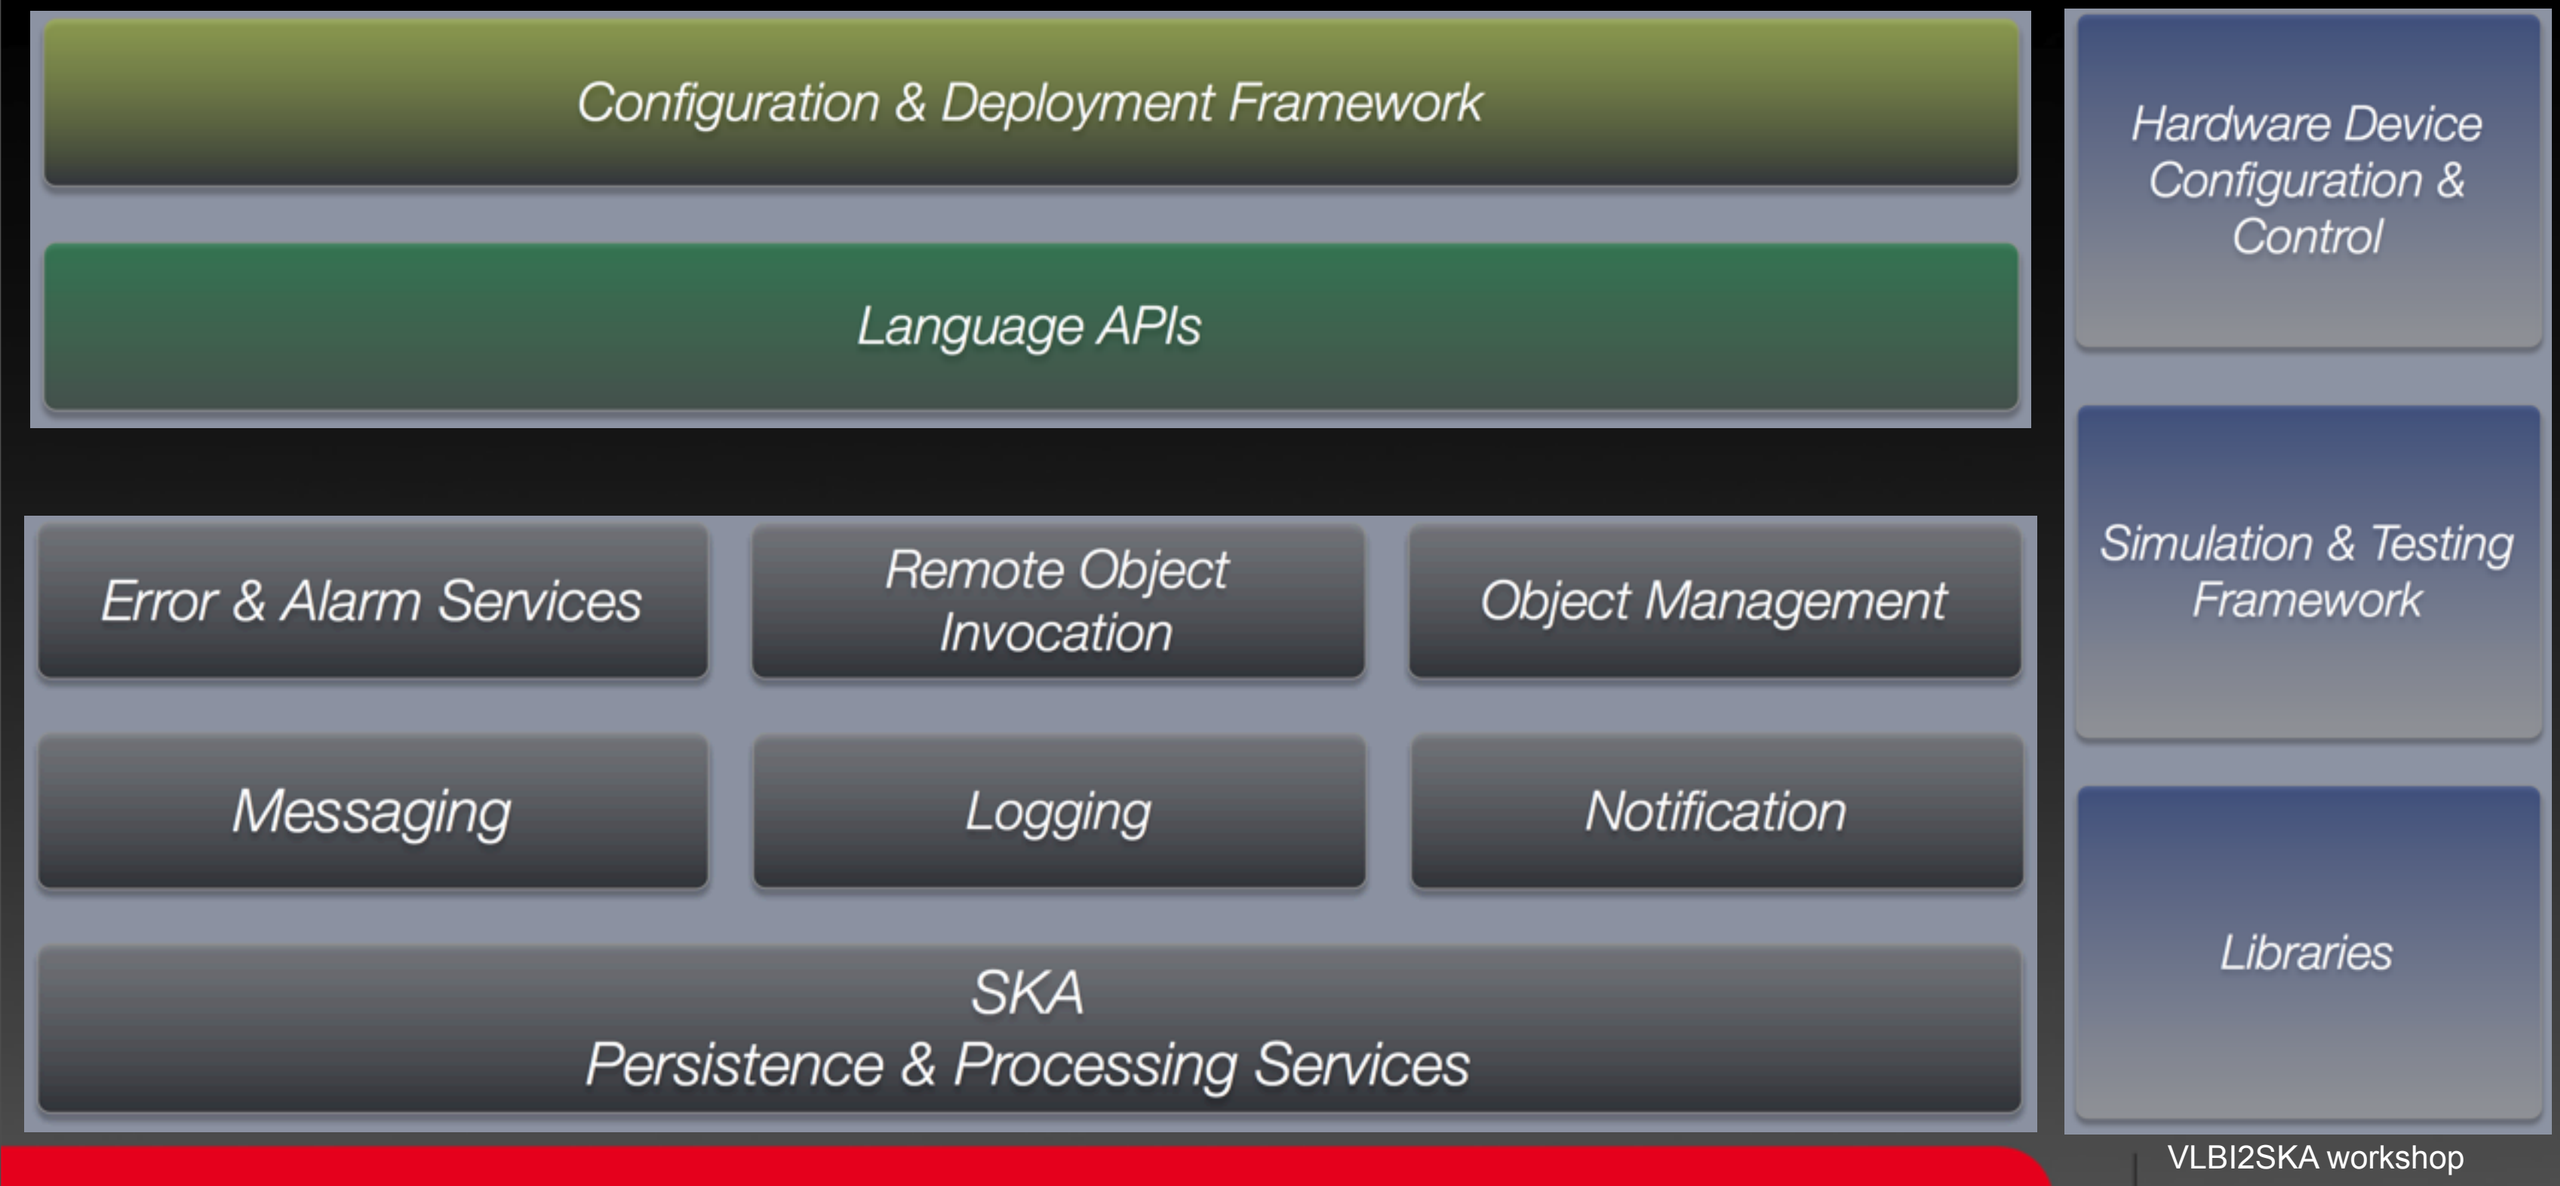

Common Software

Subsystem 1 | *Subsystem 2* | *Subsystem 3* | *Subsystem 4*

Configuration & Deployment Framework

Language APIs

Error & Alarm Services | *Remote Object Invocation* | *Object Management*

Messaging | *Logging* | *Notification*

SKA Persistence & Processing Services

Hardware Device Configuration & Control

Simulation & Testing Framework

Libraries

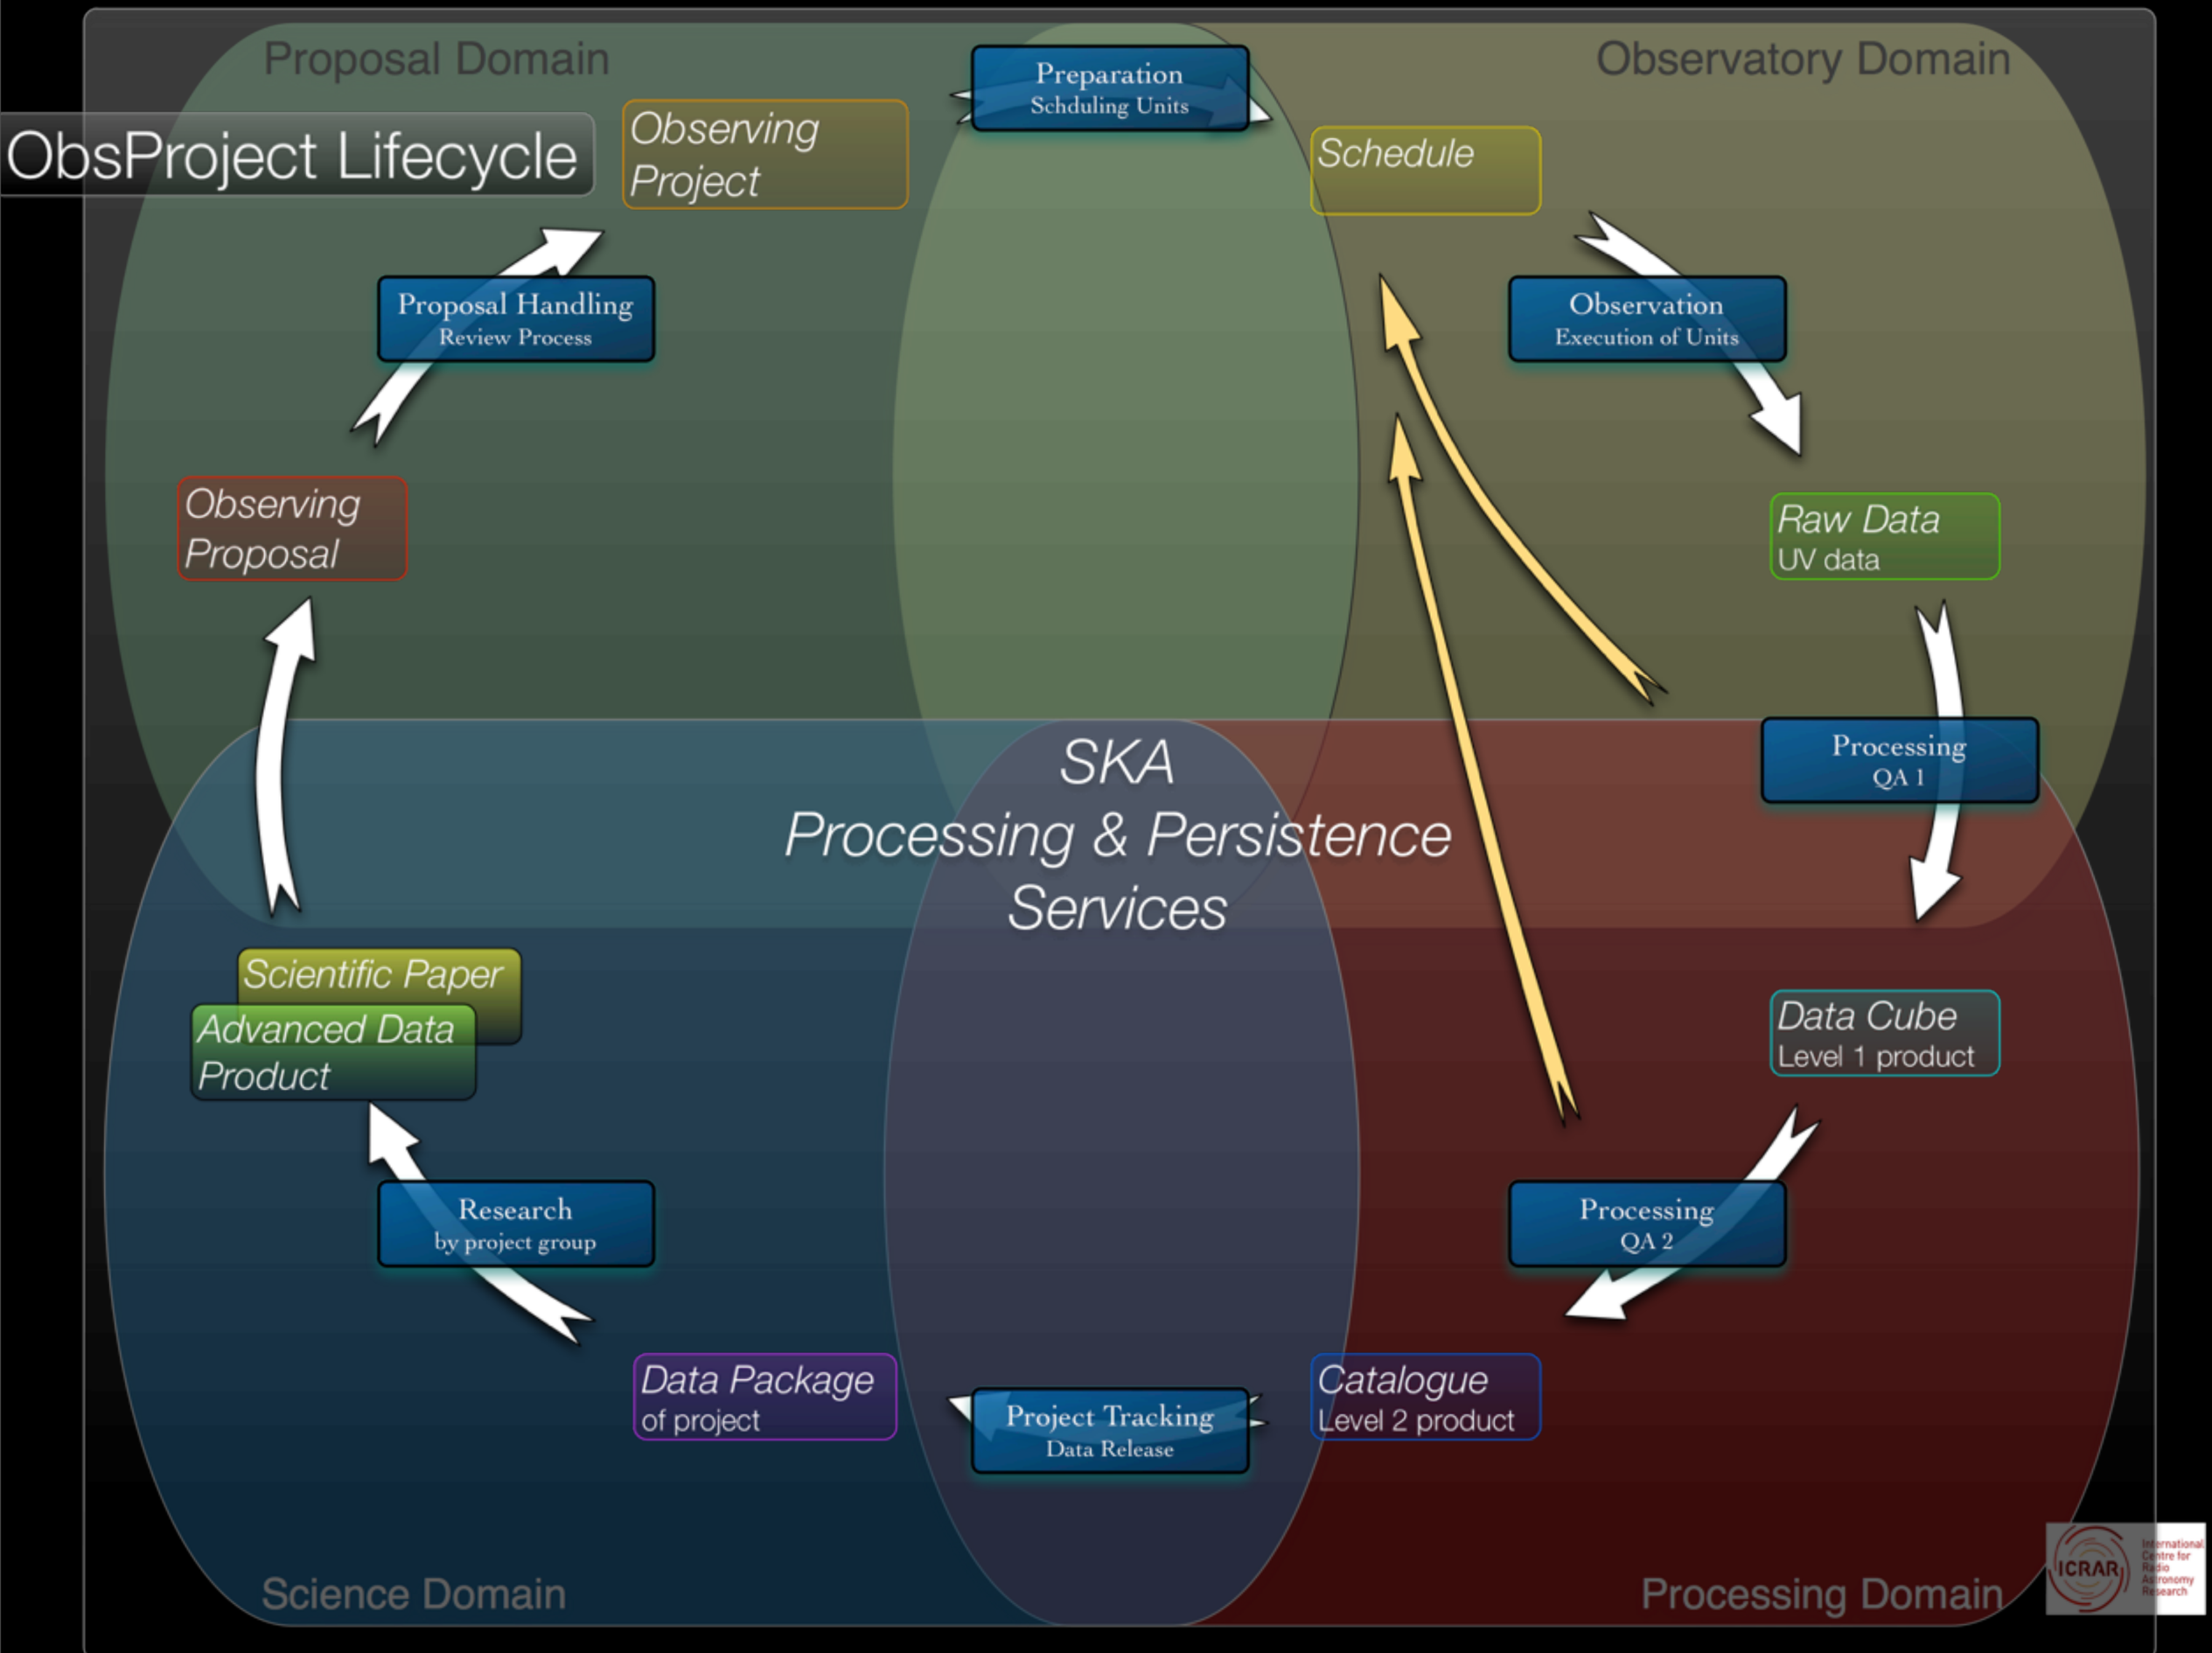
ObsProject Lifecycle
Proposal Domain
Observatory Domain
Observing Project
Preparation
Schduling Units
Schedule
Proposal Handling
Review Process
Observation
Execution of Units
Observing Proposal
Raw Data
UV data
SKA
Processing & Persistence
Services
Processing
QA 1
Scientific Paper
Advanced Data Product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

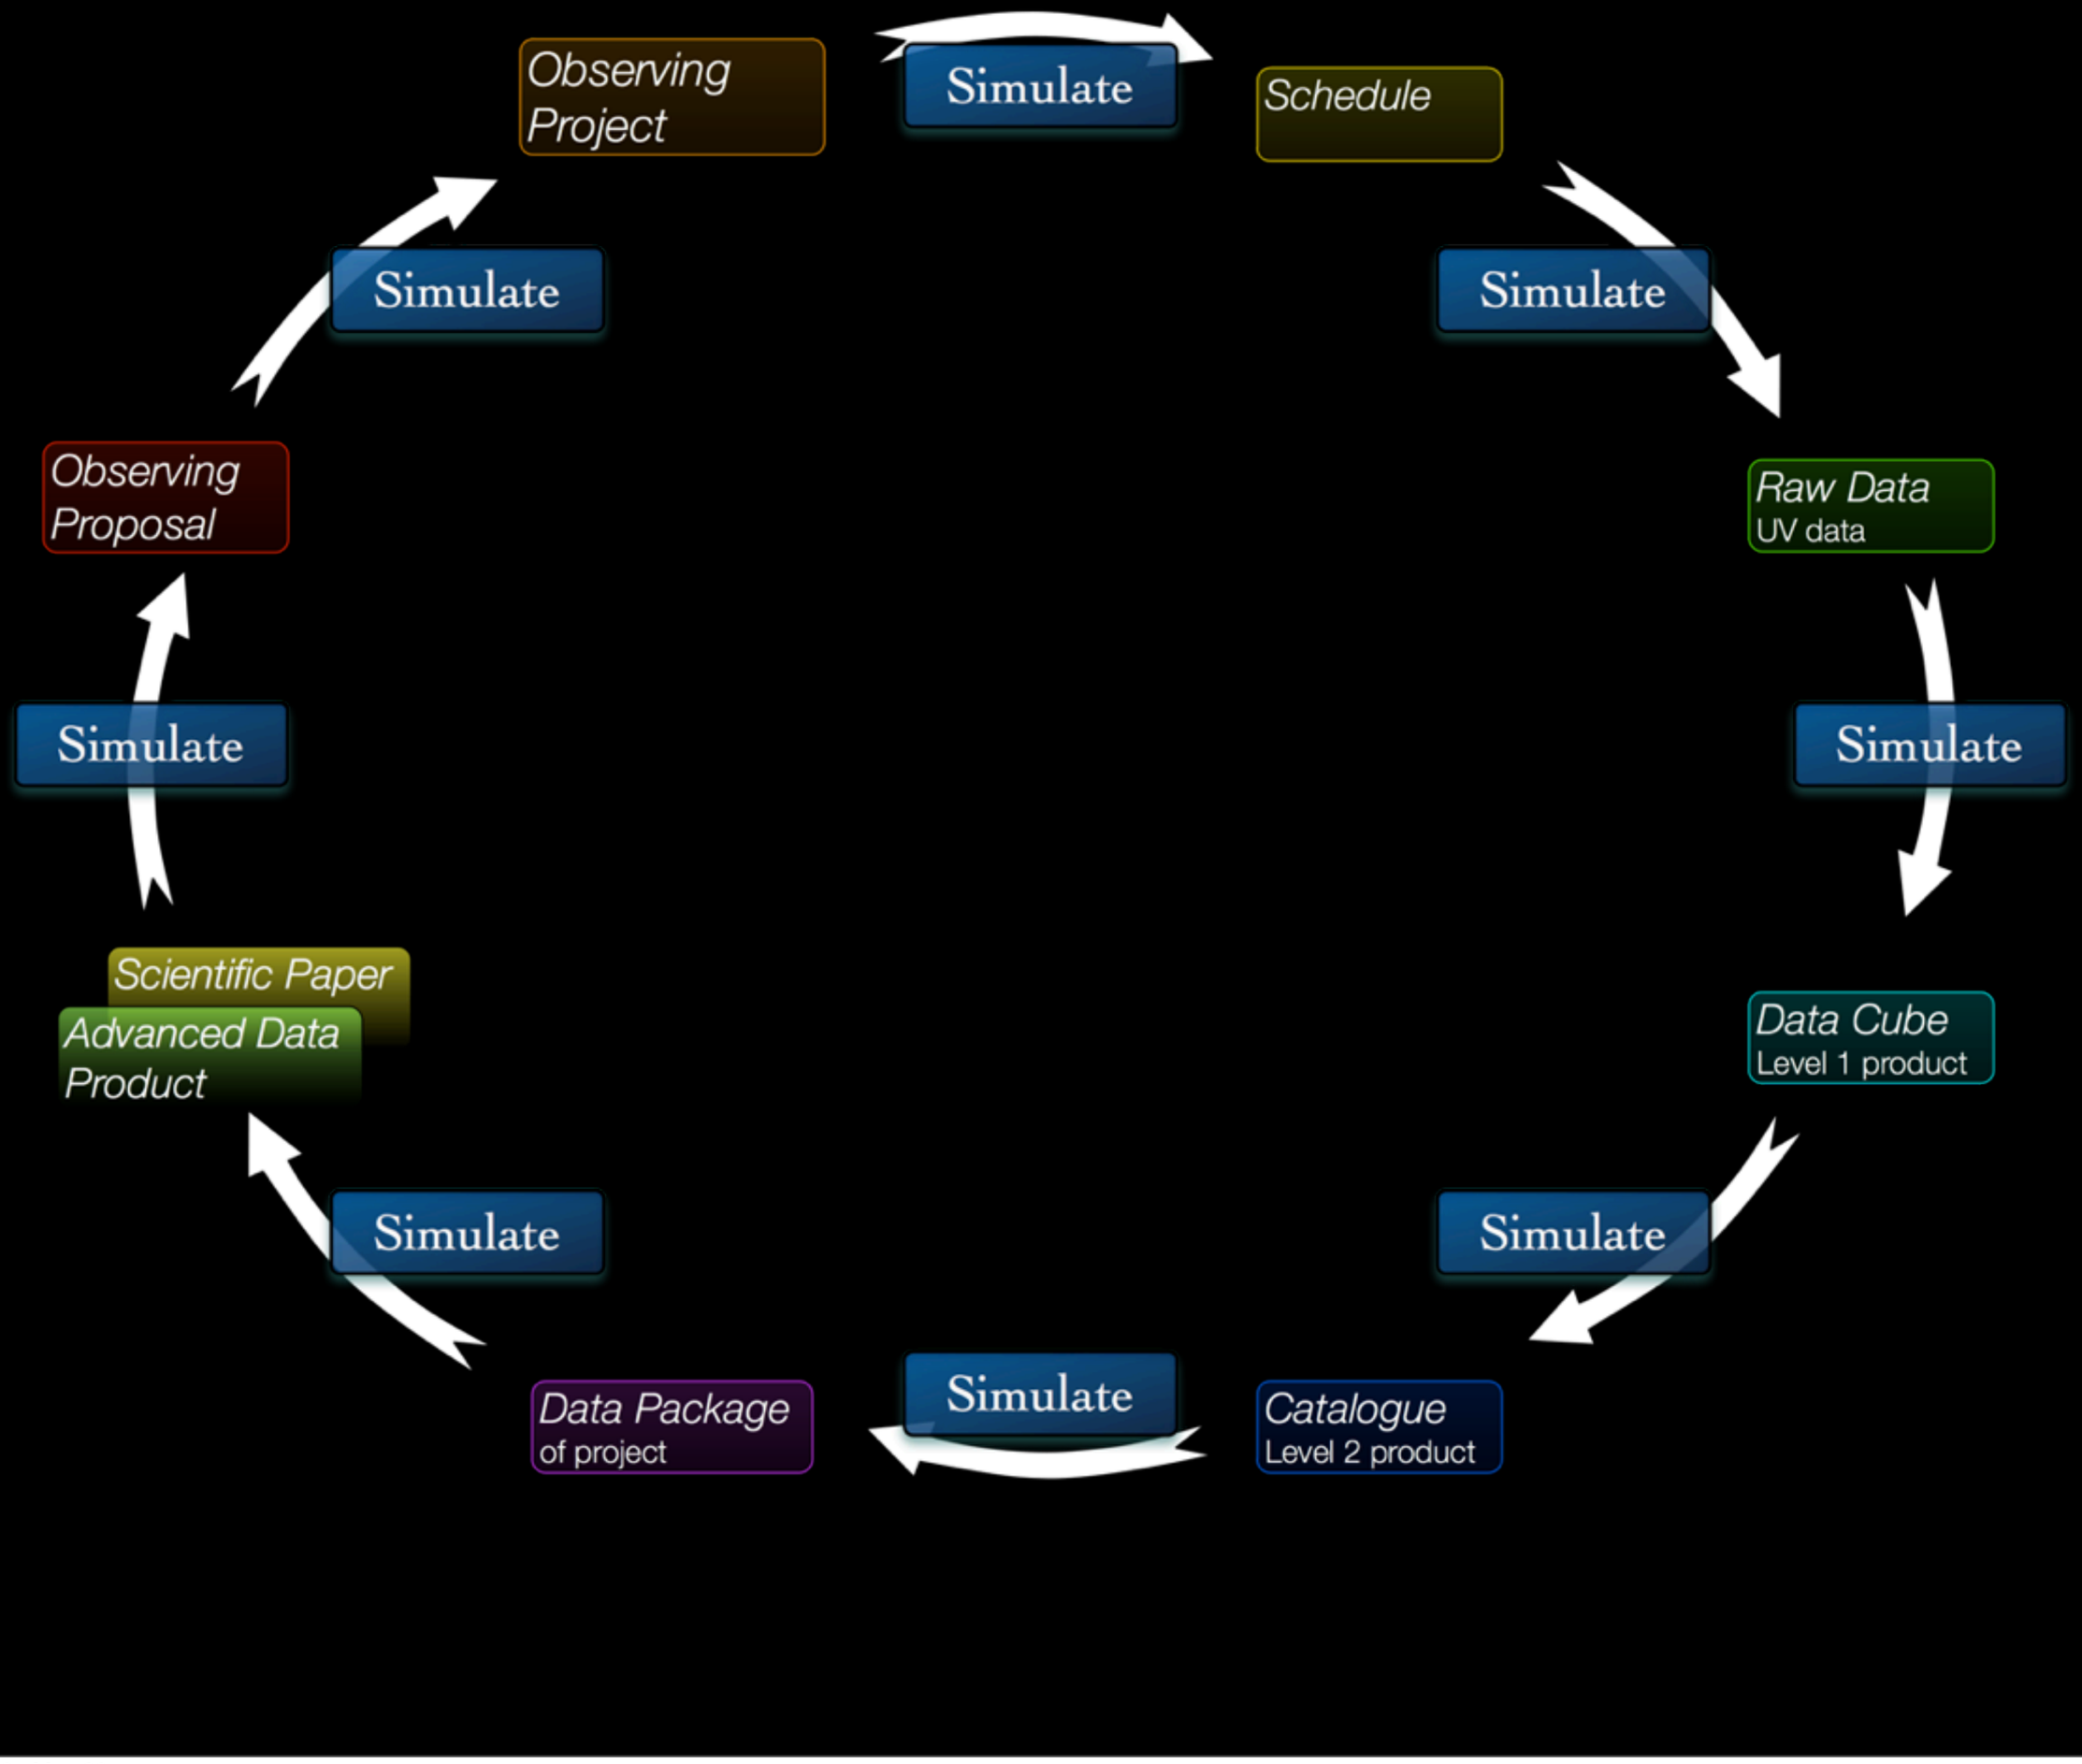
Observing Project
Simulate
Schedule
Simulate
Simulate
Raw Data
UV data
Simulate
Observing Proposal
Simulate
Scientific Paper
Advanced Data Product
Data Cube
Level 1 product
Simulate
Simulate
Data Package
of project
Simulate
Catalogue
Level 2 product

Data Flow Simulation

- Tests first approach calls for complete data flow simulation.
- Starts with simulating proposals
- Simulate every step along the way.
- Every subsystem simulates its output
- Next subsystem consumes output.
- Once simulators and interfaces are in place the whole system will work already, without any functionality...
- Simulation in some cases is a BIG effort!

Thank You!

Thanks for material: Kevin Vinsen, Tim Cornwall, Peter Quinn, ALMA Project