Data Management

Julian Borrill

Interim Data Management L2 Scientist

CMB-S4

Range

- Receive raw data from Data Acquisition
    - The raw data become DM’s when they hit site storage
- Deliver science-quality intermediate data products to the collaboration
    - “Intermediate data products” are single-frequency maps and transient alerts
    - “Science-quality” includes all documentation and ancillary data products required to analyze the data
- Deliver science-quality intermediate and final data products, and the software used to generate them, to the community
- The Data Management construction project must:
    - Support the optimization and validation of the experiment design
    - Be ready to transition to operations at first telescope commissioning

Scope

- Data registration on receipt from DAQ
- Data movement from sites to US and between US data centers
- Archival storage of raw data and derived data products
- Production of daily single-frequency maps from all telescopes
- Identification of transients in daily maps & issuing of science alerts
- Monitoring of data quality in daily maps & issuing of operational alerts
- Characterization of the experiment (instrument + observation) from design, laboratory & field data
- Production of bulk single-frequency maps, including systematics mitigation
- Characterization of bulk single-frequency maps
- Production of mock datasets for design validation & data characterization
- Delivery of science-grade intermediate data to the collaboration
- Receipt of data quality/sufficiency feedback from the collaboration
- Delivery of science-grade intermediate and final data to the community

Work Breakdown Structure

**interim*

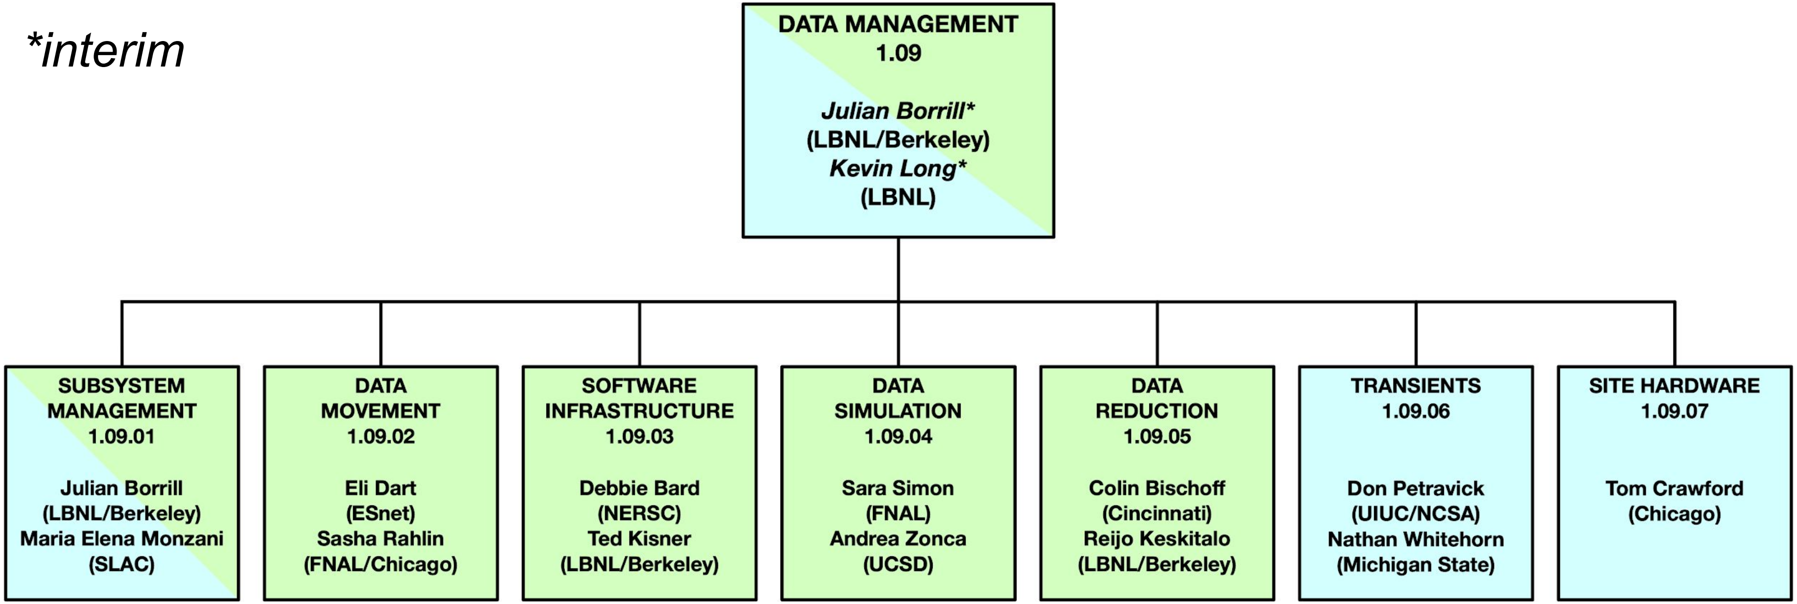

- Intentionally distributed leadership to leverage the full range of Stage 3 expertise and to interface with the collaboration as widely as possible.

Design Drivers - Data Rates

TELESCOPES	DETECTORS	SAMPLING FREQUENCY (Hz)	DATA RATE	
			RAW (Samples/Sec)	COMPRESSED (Gbps)
CHLAT	243,520	400	9.7×10^7	1.09
SPLAT	114,432	400	4.6×10^7	0.51
SAT	153,232	100	1.5×10^7	0.17

Data rates set site bandwidth and local storage requirements.

Design: sufficient network bandwidth (1.1 Gbps) from Chile; insufficient network bandwidth (0.7 Gbps) from South Pole.

Design: 1 month backup (382 TB) in Chile; 1 year backup (5.4 PB) at South Pole.

Design Drivers - Data Volumes

TELESCOPES	DAILY DATA (TB)	SPINNING DATA (PB)	ARCHIVAL DATA (PB)
CHLAT	11.8	4.3	30.1
SPLAT	5.5	2.0	14.2
SAT	1.9	0.7	4.7

Data volumes set US data center storage requirements

Design: 1 year of raw data + 7 years of science data (17 PB) spinning at each data center

Design: 7 years raw data (49 PB) archived at each data center

Design Drivers - Daily Data Processing

TELESCOPES	CYCLES (TFLOP)	PEAK MEMORY (TB)	SINGLE DAILY MAP DATA (GB)	TOTAL DAILY MAP DATA (TB)
CHLAT	0.84	11.5	54	138
SPLAT	0.40	5.7	4.5	12
SAT	0.13	1.2	0.02	0.05

Daily data processing drives the fast-access computational requirements

Design: CHLAT processing in the US; SPLAT+SAT processing at South Pole

Design Drivers - Bulk Data Processing

Simulating and reducing the entire dataset requires:

- Cycles: 21.9 EFLOP
- Peak memory: 11.7 PB
- Peak scratch: 2.1 PB

Design: bulk computational resources must be allocated at national computing centers

Design: balance having sufficient centers to accommodate down-times & support diverse approaches (high performance + high throughput) against the cost per center of maintaining/optimizing the software stack.

PBD Hardware Schematic

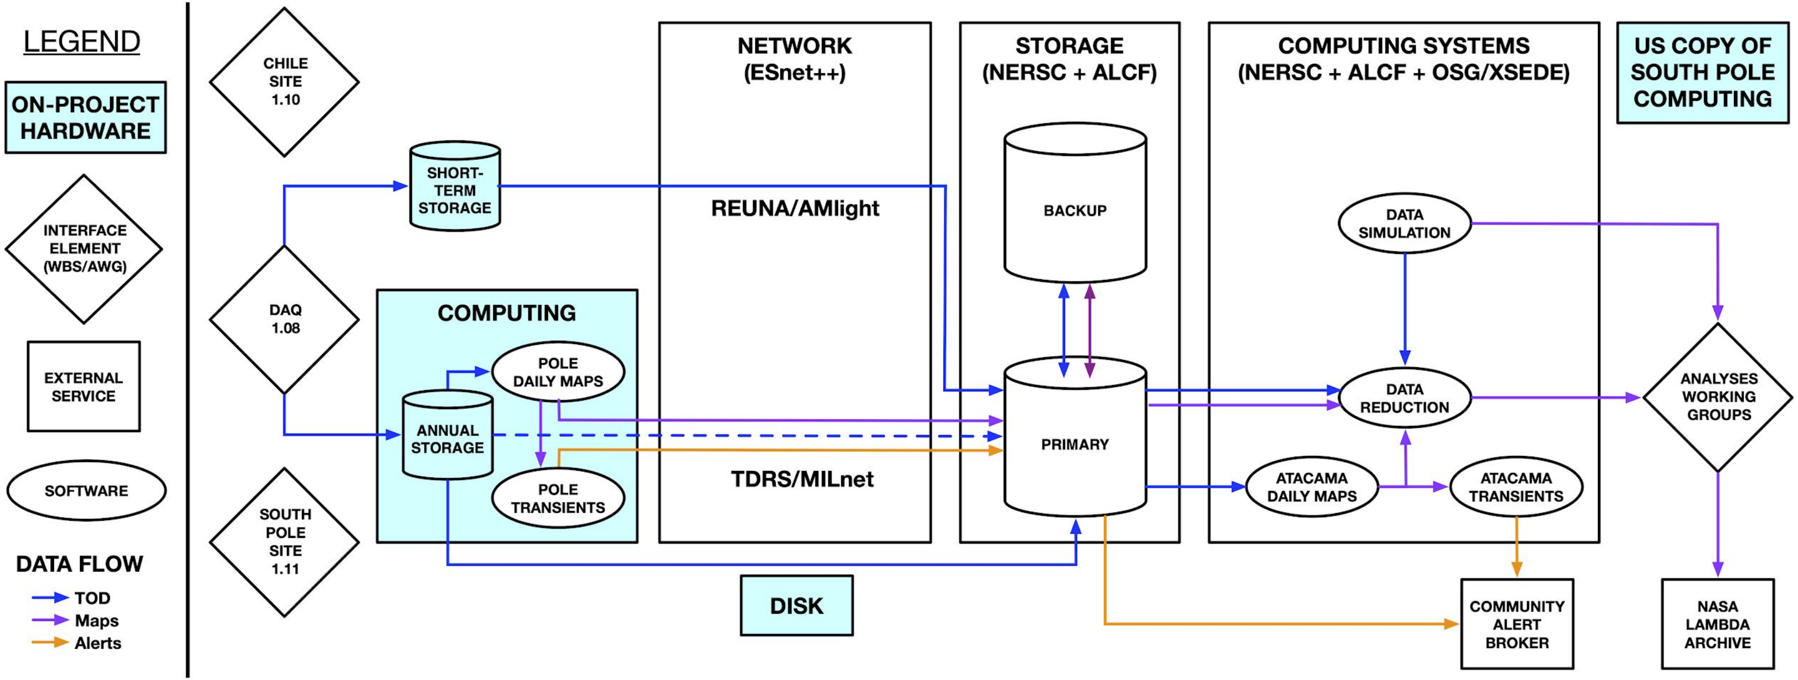

Allocated computational resources are planned, not confirmed.

PBD Software Schematic

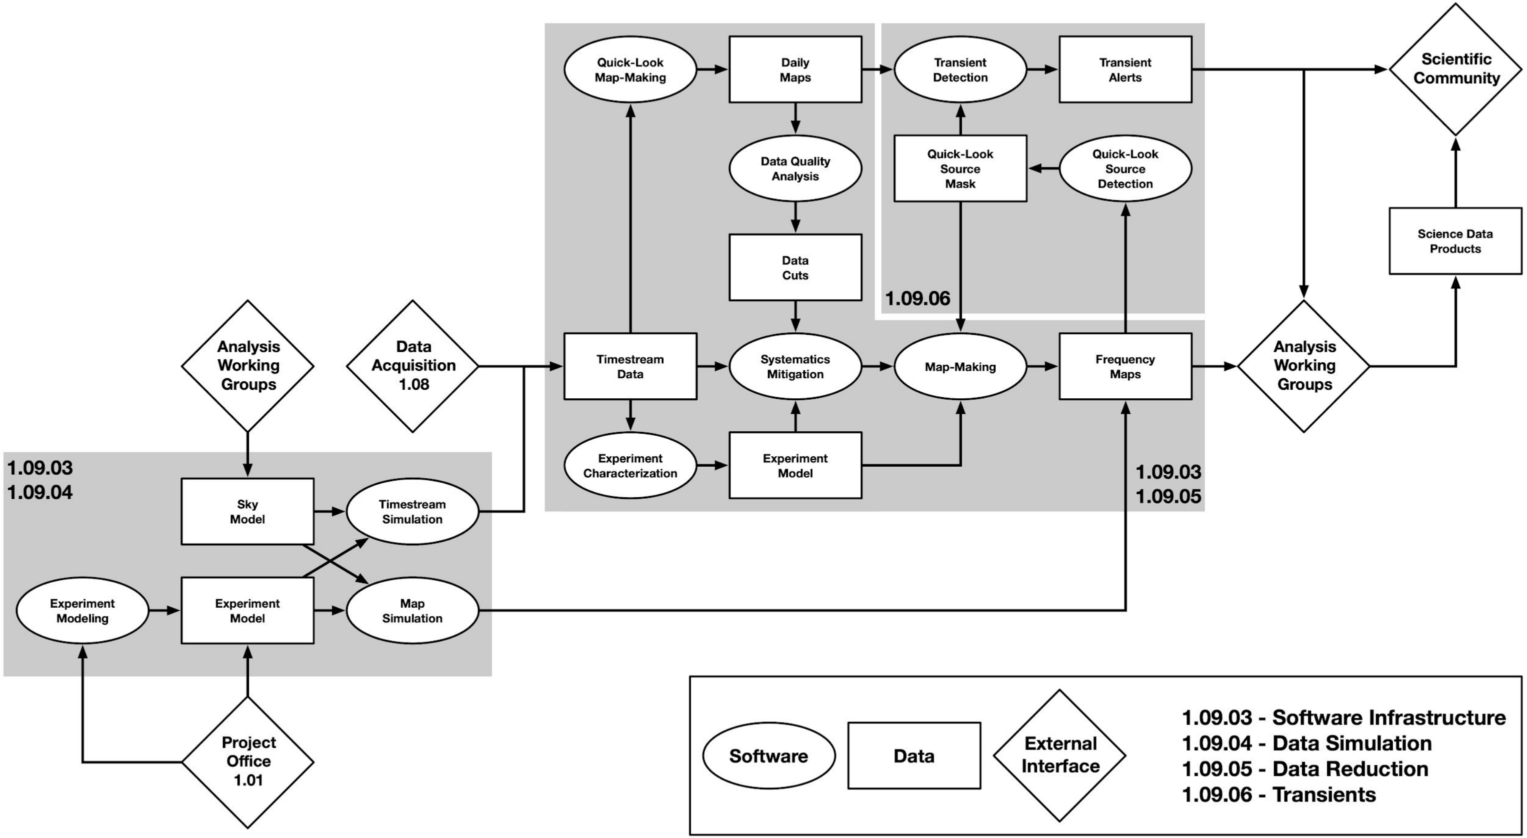

Interfaces

All interfaces must be documented:

- Interface Control Documents with other L2 Subsystems
- Memorandum of Agreement with the Collaboration

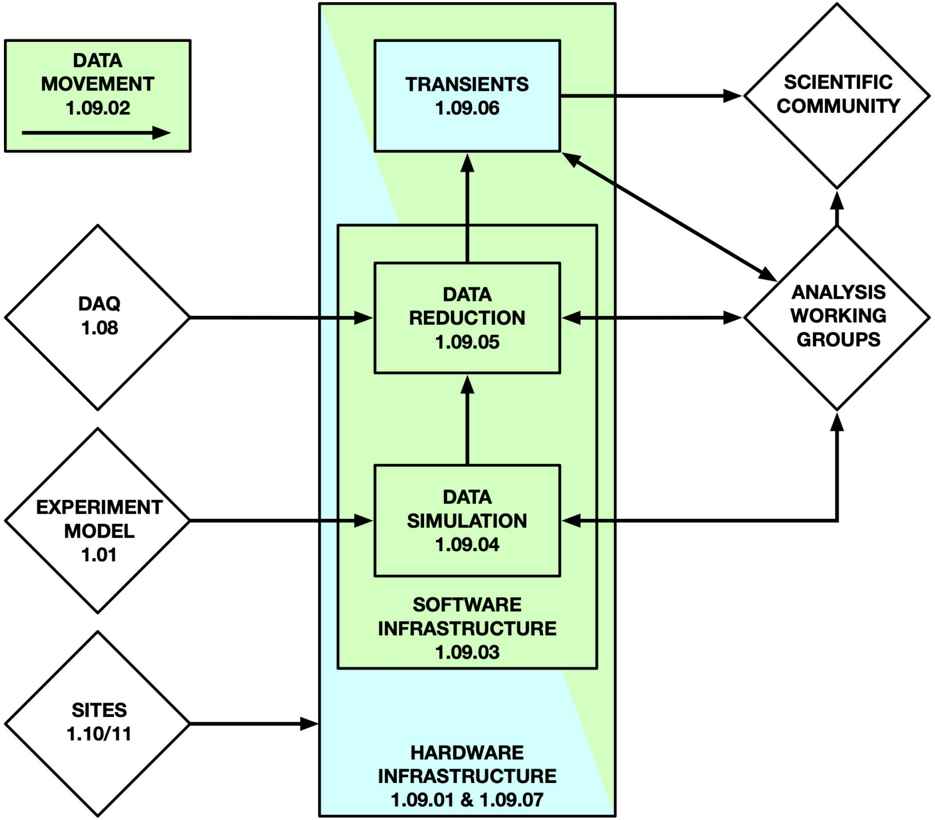

CMB-S4

Data Challenges

- Data Challenges are at the core of the DM construction project:
  - Experiment (Instrument + Observation) design validation
  - Data management subsystem validation
    - Including sufficiency of allocated computational resources
  - Analysis pipeline validation
- Each agency review features an enhanced/matured design to be validated.
- Each review is informed by a preceding data challenge.
- Each Data Challenge is a 6-month process.

Data Challenge Schedule/Process

Months Before Deadline	1.09: Data Management							AWGs	TWGs
	1.09.01 Subsystem Management	1.09.02 Data Movemement	1.09.03 Software Infrastucture	1.09.04 Data Simulation	1.09.05 Data Reduction	1.09.06 Transients	1.09.07 Site Hardware		
12		Develop and deploy overall software stack, including system-specific porting & optimization Support design optimization						Optimize experiment design	
11									
10									
9									
8									
7									
6	1. Define scope								
5				2. Freeze simulation & reduction modules					
4			3. Construct and execute pipelines						
3	6. Deliver	5. Populate data registry & archive		4. Verify outputs					
2	7. Review internally								
2		8. Distribute data			9. Execute data quality/experiment characterization analyses + transient identification			Execute science analyses & document in design report Provide feedback to DM	Execute technical analyses & document in design report Provide feedback to DM
1									
0	10. Director's Review								
-1	11. Agency Review								

Parallel Session

Presentation by each L3 team

- More detailed dive into each L3 subsystem
- Open issues and questions
- Path to CD-1/PDR

Note also the “Design Validation: Technical to Measurement” theme which will focus on Data Challenge 1 and validating the Preliminary Baseline Design.

Open Questions

- Do we have the full set of data product requirements from the collaboration?
- What is the exact boundary between project and collaboration wrt transients?
- What is the optimal map-making approach when computational cost is included? Does it vary with science case?
- Are there other computational resources we should be looking to use?
  - FABRIC in network computing
  - Joint South Pole computational infrastructure with IceCube
- How best do we deliver data to the collaboration? Can/must we deliver computational resources too?
- Do we have the right feedback loops with TWGs and AWGs?