



ATLAS software and computing road-map

James Catmore, UiO

NorCC Strategy kick-off meeting, 11th March 2022

HL-LHC and ATLAS computing




ATLAS road-map for HL-LHC computing

- Prepared as input to the LHCC review of computing for HL-LHC
 - Lays out a concrete plan to deliver the vision set out in the Conceptual Design Report prepared in 2020
 - Consists of a set of milestones to be met in the coming years across all areas of software and computing in ATLAS
 - Milestones are categorised as being as
 - “**Maintenance and operation**”: needed just to stand still
 - “**Conservative R&D**”: development work that can be carried out with existing person power within that domain
 - “**Aggressive R&D**”: development work that will require new person-power or existing personnel committing to new activities, not assumed to come from the relevant domain
- More aggressive milestones → potentially greater impact on resources
- Progress will be followed up internally every six months



ATLAS Software and Computing HL-LHC Roadmap



Reference:

Created: 1 October 2021

Last Modified: 22 February 2022

Prepared by: The ATLAS Collaboration

© 2022 CERN for the benefit of the ATLAS Collaboration.

Reproduction of this article or parts of it is allowed as specified in the CC-BY-4.0 license.

Roadmap topics

Project management

Reconstruction

Databases, metadata, conditions

Analysis

Core software & heterogeneous computing


Distributed computing

Detector description, simulation, digitisation

HPC and analysis facilities

Collaboration with common projects

Current or planned Norwegian involvement



Milestones

Project management

Project Organization			
MID	DID	Description	Due
PR-1		First evaluation of effort needed to deliver on HL-LHC milestones	Q2 2022
PR-2		HL-LHC Computing TDR	Q3 2024
2.1		R&D projects targeting Run 4 ("Run 4 projects" define scope and potential impact of their demonstrators, and a program of work with effort and risk estimates to the end of Phase 2.	Q4 2022
2.2		Define release, datasets and platforms to be used to evaluate Run 4 performance impact of demonstrators	Q1 2024
2.3		Run 4 projects release their demonstrators	Q2 2024
2.4		Run 4 projects evaluate the performance impact of their R&D demonstrators and estimate the effort needed to develop fully functional prototypes	Q2 2024
PR-3		Run 4 Release	Q2 2027
3.1		Run 4 projects release fully functional prototypes, estimate risks and effort needed to bring to production quality	Q2 2026
3.2		Run 4 developers tutorial	Q3 2026
3.3		Run 4 Feature Freeze	Q2 2027
PR-4		Ready for Run 4 Data Taking	Q2 2029
4.1		Run 4 projects demonstrate required functionality in release	Q3 2026
4.2		Run 4 release validated	Q1 2029
Maintenance & Operations Conservative R&D Aggressive R&D			

Databases, metadata, conditions

Database infrastructure, conditions and metadata (ADAM)			
MID	DID	Description	Due
DB-1		Relational database infrastructure consolidation	Q4 2022
1.1		Migration of application schemas used in data processing from ATLR to ATONR	Q3 2022
1.2		Setup of Frontier launchpad using ATONR_ADG (Active Data Guard) nodes.	Q4 2022
DB-2		Conditions database CREST development and integration	Q4 2026
2.1		Define the CREST development, timeline and deployment strategy, as well as COOL data migration strategy.	Q2 2022
2.2		Finalize the prototype for CREST deployment, partial COOL data migration and Athena jobs testing. Implement changes needed in IOVDbSvc.	Q4 2022
2.3		Time series data processing and storage evaluation for CREST usage (DCS data).	Q1 2023
2.4		Implement a directory structure to export selected conditions data from CREST into CVMS (access from sites without network connectivity).	Q4 2023
2.5		Deployment of a parallel infrastructure for accessing conditions via CREST for larger tests (validate scalability and experts needs).	Q1 2026
2.6		Final CREST validation for Run4 and COOL decommissioning.	Q4 2026
DB-3		Metadata evolution	Q1 2025
3.1		AMI Application migration to CERN Oracle DB.	Q4 2021
3.2		AMI Tag evolution and deployment.	Q4 2022
3.3		Fully functional and integrated metadata system to characterize datasets (data and MC) used for physics analysis.	Q1 2025
Maintenance & Operations Conservative R&D Aggressive R&D			

Core software & heterogeneous computing

Core Software, Heterogeneous Computing and Accelerators			
MID	DID	Description	Due
CS-1		Pileup-digitization in AthenaMT production ready	Q4 2022
1.1		Ensure reproducibility of MT production of presampled MB RDO files	Q2 2022
CS-2		Complete investigation of lossy compression techniques	Q4 2023
2.1		Lossy compression of the ID track covariance matrix in the primary AODs	Q4 2021
2.2		Lossy compression of DAOD	Q4 2021
2.3		Lossy compression of primary AODs	Q4 2023
CS-3		Implement I/O roadmap metadata recommendations	Q4 2022
3.1		Multi-threaded in-file metadata handling	Q2 2022
3.2		Redesign of the metadata handling infrastructure (better support for fine-grained workflows)	Q4 2022
CS-4		Evaluation of data formats well-suited for massively parallel I/O (HPCs)	Q1 2022
4.1		Storing intermediate EventService Simulation data in HDF5	Q1 2022
CS-5		Migration to ROOT 7	Q4 2026
5.1		StorageSvc capable of writing a subset of DAOD_PHYS(Lite) data to RNTuple	Q4 2023
5.2		Migration of Athena from Root 6 to 7	Q1 2025
5.3		ROOT and LCG release contributions, testing, feedback	Ongoing
CS-6		Re-evaluation of simulation data formats (including EVNT, HITS, RDO)	TBD
6.1		Technical review of file storage format, compression, etc	
6.2		Content review of data format	
CS-7		GPU Kernel scheduling	Q1 2023
7.1		Basic support for Kernel scheduling in athena	Q2 2021
7.2		Integration with Gaudi Scheduler	Q1 2023
CS-8		GPU management techniques and infrastructure in Athena	Q4 2021
CS-9		Develop Multi-algorithm heterogeneous applications	Q2 2023
9.1		ACTS-based multi-algorithm workflow	Q3 2022
9.2		FastCaloSim GPU merged into master	Q2 2022
9.3		Calorimeter clustering	Q4 2022
9.4		GPU-accelerated ML inference in athena	Q2 2023
CS-10		Infrastructure for processing data across multiple events on an accelerator	Q4 2023
10.1		Proof-of-concept prototype	Q4 2021
CS-11		GPU Memory management	TBD
11.1		First (Vecmem) prototype	Q1 2022
CS-12		Make ATLAS Data Model classes accelerator-friendly	Q4 2024
12.1		Prototype GPU-friendly xAOD classes	Q1 2022
12.2		Support for reduced/mixed precision in ATLAS EDM	Q4 2022
12.3		Decision on xAOD API evolution	Q4 2022
12.4		Event-batching and EDM	Q3 2023
12.5		Accelerator-friendly detector data model (geometry and calibration)	Q4 2024
12.6		Evaluate mechanism to offload and update detector description on device	Q4 2024
CS-13		Intra-node scheduling, targeting HPCs and grid	TBD
13.1		Raythena/HPX-based scheduler prototype	Q3 2022
CS-14		HL-LHC Technology decision: CUDA or one of its less-proprietary competitors	Q1 2024
14.1		Full parallelization pattern recommendation to collaboration	Q1 2024
14.2		Design patterns/tutorial on GPU migration	Q1 2024
Maintenance & Operations Conservative R&D Aggressive R&D			

Detector description, simulation, digitisation

Simulation			
MID	DID	Description	Due
SI-1		Updates required for MC+MC Overlay for Run4	Q2 2023
1.1		Enhance Overlay code to better deal with high pile-up environments	Q2 2023
1.2		Implement Overlay Algorithm for HGTD	Q2 2023
1.3		Implement Overlay Algorithm for ITk Pixels	Q2 2023
1.4		Implement Overlay Algorithm for ITk Strips	Q2 2023
SI-2		Track Overlay	Q2 2023
2.1		Data overlay with pre-reconstructed tracks	Q2 2023
SI-3		Review or Run 4 Simulation Strategy	Q2 2024
3.1		Evaluate Geant4 equivalents of ISF functionality	Q1 2023
3.2		Evaluate accuracy of EMEC simulation in FastCaloSim	Q4 2023
3.3		Prototype EMEC geometry using native Geant4-GPU constructs	Q4 2023
3.4		Prototype GPU version of LAr sensitive detector callback	Q2 2024
SI-4		Pile-up Digitization in AthenaMT	Q4 2023
4.1		Full AthenaMT compatible pile-up digitization	Q2 2023
4.2		Reduce memory usage of pileup digitization for very high mu	Q4 2023
SI-5		Make data overlay useable for p-p collision simulation	Q2 2026
5.1		Flexible alignment and geometry in simulation for data overlay	Q2 2024
5.2		Skimming of zero-bias Bytestream data and efficient grouping of events as input to p-p overlay	Q2 2025
5.3		Adapt Fast Chain workflow to data overlay	Q2 2028
SI-6		Fast simulation/FastChain development for Run4	Q2 2027
6.1		R&D in FastChain alternatives	Q2 2026
6.2		Development of substantial improvements to fast calorimeter simulation	Q2 2026
6.3		Development of substantial improvements to fast inner detector simulation	Q2 2026
6.4		Implementation of ITk in FATRAS	Q2 2026
6.5		Tuning fast simulation to data	Q2 2027
6.6		Improve Geant 4 interface for fast simulation	Q2 2026
SI-7		Geometry updates for the Run4 MC Campaign	Q4 2027
SI-8		Run4 optimisation of the performance of ATLAS full Geant4 simulation	Q4 2027
SI-9		Testing Geant4 Versions and configurations for Run4	Q4 2026
9.1		Validate Geant 4 v11	Q4 2023
SI-10		Validation of full and fast simulation for Run4	Q3 2028
SI-11		Feature complete digitization + overlay for all sub-systems including conditions access	Q2 2027
SI-12		Feature complete trigger simulation including compatibility with overlay	Q2 2027
Maintenance & Operations Conservative R&D Aggressive R&D			

Reconstruction

Reconstruction			
MID	DID	Description	Due
RE-1		Merge detector upgrade release (21.9) with master	Q1 2022
RE-2		Adapt Reconstruction to the Phase-II Detector	Q1 2024
RE-3		Migrate default CPU chain to ACTS	Q3 2025
3.01		Define and migrate to a new internal tracking EDM	Q1 2023
3.02		Define run-4 ATLAS EDM	Q1 2022
3.03		Migrate to the new run-4 ATLAS EDM	Q3 2022
3.04		Full chain prototype/demonstrator with all the components	Q3 2024
3.05		ITk reconstruction	Q4 2024
3.06		Muon tracking geometry and navigation	Q3 2024
3.07		Muon standalone reconstruction	Q1 2025
3.08		Muon combined reconstruction	Q3 2025
3.09		Calorimeter tracking geometry and navigation	Q3 2025
3.10		Calorimeter track extensions	Q3 2025
3.11		Electron and gamma reconstruction	Q3 2025
3.12		Particle flow reconstruction	Q3 2025
3.13		Tau reconstruction	Q3 2026
3.14		Flavour tagging	Q3 2026
3.15		Jets and Missing ET reconstruction	Q3 2026
RE-4		Accelerator and machine learning (R&D)	Q3 2025
4.1		Develop demonstrators for accelerators and new ML techniques	Q1 2024
4.2		Finalise and implement functional prototypes	Q3 2025
RE-5		Feature freeze	Q3 2026
5.1		Calorimeter reconstruction	Q1 2026
5.2		ITk reconstruction	Q1 2026
5.3		Muon reconstruction	Q1 2026
5.4		Electron and gamma reconstruction	Q1 2026
5.5		Particle flow reconstruction	Q2 2026
5.6		Tau reconstruction	Q2 2026
5.7		Flavour tagging	Q3 2026
5.8		Jets and Missing ET reconstruction	Q3 2026
RE-6		Performance freeze	Q1 2028
6.1		Calorimeter reconstruction	Q2 2027
6.2		ITk reconstruction	Q2 2027
6.3		Muon reconstruction	Q3 2027
6.4		Electron and gamma reconstruction	Q3 2027
6.5		Particle flow reconstruction	Q4 2027
6.6		Tau reconstruction	Q4 2027
6.7		Flavour tagging	Q1 2028
6.8		Jets and Missing ET reconstruction	Q1 2028
RE-7		Validation	Q3 2028
Maintenance & Operations Conservative R&D Aggressive R&D			

Analysis

Analysis			
MID	DID	Description	Due
AN-1		Baseline DAOD_PHYS(LITE) with run 3 s/w (event loop based) & TTree	Q3 2023
1.1		First bulk production of prototype DAOD_PHYS(LITE)	Q4 2021
1.2		Mechanism for evaluation of systematic uncertainties with PHYS(LITE)	Q2 2022
1.3		Demonstrator for full analysis on PHYS(LITE) for target analyses	Q4 2022
1.4		Recommendations for application of lossy compression	Q4 2022
1.5		Finalised list of PHYS(LITE) contents	Q2 2023
1.6		Development/roll-out of docs/training for run 3 analysis w/ PHYS(LITE)	Q3 2023
AN-2		PHYS(LITE) working with RNTuple	Q4 2023
2.1		Implementation of RNTuple and revised xAOD in PHYS(LITE)	Q4 2023
AN-3		Prototyping & review of columnar data operations for end-to-end analysis	Q2 2024
3.1		Tests of basic reading performance using TTree version of PHYS(LITE)	Q4 2022
3.2		Prototyping of tools for columnar CP operations and other systematics	Q4 2023
3.3		Adoption of ROOT7 data structures	Q1 2024
3.4		Performance and ease-of-use assessment leading to decision on adoption	Q2 2024
AN-4		Development of columnar analysis infrastructure	Q2 2026
4.1		Prototyping of framework for orchestrating columnar CP operations	Q3 2024
4.2		Development of columnar skimming/augmentation	Q1 2025
4.3		Demonstrate end-to-end analysis using columnar tools/fkw + dist. comp.	Q4 2025
4.4		Development & roll-out of documentation/training for run 4 analysis	Q2 2026
AN-5		Accommodate all analyses in the run 4 analysis model	Q2 2028
5.2		First assessment of run 4 analyses incompatible with PHYS(LITE)	Q2 2023
5.3		Calculate resources implications and feed back to physics coordination	Q3 2023
5.4		Set up DAODs for analyses which are unable to use the new model	Q2 2026
5.5		Updated DAOD_PHYS(LITE) to capture additional analyses	Q2 2028
Maintenance & Operations Conservative R&D Aggressive R&D			

Distributed computing

Distributed Computing			
MID	DID	Description	Due
DC-1		Transition to tokens	Q4 2025
1.1		Submission from Harvester to all HTCCondor CEs with tokens	Q1 2022
1.2		All users move from VOMS to IAM for X509	Q4 2022
1.3		All job submission and data transfers use tokens	Q4 2025
DC-2		Storage evolution	Q4 2025
2.1		No GridFTP transfers at any site	Q1 2022
2.2		SRM-less access to tape	Q4 2025
2.3		Recommended transition plan from DPM completed	Q4 2021
2.4		Transition plan from all DPM sites	Q4 2022
2.5		All sites moved away from DPM	Q2 2024
DC-3		Next operating system version	Q2 2024
3.1		Ability to run on "future OS" on grid sites	Q4 2022
3.2		Central services moved to "future OS"	Q2 2023
3.3		(CentOS 7/8 EOL)	Q2 2024
DC-4		Network infrastructure ready for Run 4	Q4 2027

4.1		Network challenge at 10% expected rate	Q4 2021
4.2		Network challenge at 30% expected rate	Q4 2023
4.3		Network challenge at 60% expected rate	Q4 2025
4.4		Network challenge at 100% expected rate	Q4 2027
DC-5		Integrating next generation of HPCs	Q2 2023
5.1		Integration of at least 2 EuroHPC sites	Q4 2022
5.2		Integration of next generation US HPCs for production	Q2 2023
DC-6		Exploratory R&D on GPU-based workflows for next generation HPC	Q4 2023
DC-7		HL-LHC datasets replicas and versions management	Q2 2024
7.1		Replicas and versions detailed accounting	Q4 2022
7.2		DAOD replicas reduction	Q4 2023
7.3		DAOD versions reduction	Q2 2024
DC-8		Data Carousel for storage optimization	Q4 2023
8.1		Investigate with sites the cost of Tape infrastructure and the estimated cost in case of sensible increase of read/write throughput	Q4 2022
8.2		Reduce the AOD on disk to 50% of the total AOD volume, using Data Carousel to orchestrate the stage from tape for DAOD production.	Q4 2023
DC-9		Disk management: secondary(cached) dataset	Q2 2023
9.1		Evaluate the impact on job brokering and task duration if disk space for secondary data is reduced	Q2 2023
Maintenance & Operations Conservative R&D Aggressive R&D			

Current involvement

Distributed computing

HPC and analysis facilities

Reconstruction

Analysis

Project management

- ATLAS distributed computing management, NorduGrid/ARC middleware development, ATLAS@Home, HPC integration
- Tau reconstruction
- Derivation framework, data formats, distributed analysis
- Development and follow-up of the milestones

Contributions to the milestones

(Non-exhaustive; based on current expertise and interests)

- GPU programming in core software
- Machine learning in reconstruction
- High performance analysis
- HPC activities, new workflows

NFR application
(Catmore & Gramstad)

Strong interest/expertise in Bergen
w.r.t. tau reconstruction

NFR application
(Catmore & Gramstad)
ATLAS QT (Langrekken)

NeIC application
(Cameron, Ould Saada, Read)

DC-5	Integrating next generation of HPCs	Q2 2023
5.1	Integration of at least 2 EuroHPC sites	Q4 2022
5.2	Integration of next generation US HPCs for production	Q2 2023
DC-6	Exploratory R&D on GPU-based workflows for next generation HPC	Q4 2023
DC-7	HL-LHC datasets replicas and versions management	Q2 2024
7.1	Replicas and versions detailed accounting	Q4 2022
7.2	DAOD replicas reduction	Q4 2023
7.3	DAOD versions reduction	Q2 2024
DC-8	Data Carousel for storage optimization	Q4 2023
8.1	Investigate with sites the cost of Tape infrastructure and the estimated cost in case of sensible increase of read/write throughput	Q4 2022
8.2	Reduce the AOD on disk to 50% of the total AOD volume, using Data Carousel to orchestrate the stage from tape for DAOD production.	Q4 2023
DC-9	Disk management: secondary(cached) dataset	Q2 2023
9.1	Evaluate the impact on job brokering and task duration if disk space for secondary data is reduced	Q2 2023

CS-12	Make ATLAS Data Model classes accelerator-friendly	Q4 2024
12.1	Prototype GPU-friendly xAOD classes	Q1 2022
12.2	Support for reduced/mixed precision in ATLAS EDM	Q4 2022
12.3	Decision on xAOD API evolution	Q4 2022
12.4	Event-batching and EDM	Q3 2023

RE-4	Accelerator and machine learning (R&D)	Q3 2025
4.1	Develop demonstrators for accelerators and new ML techniques	Q1 2024
4.2	Finalise and implement functional prototypes	Q3 2025

AN-3	Prototyping & review of columnar data operations for end-to-end analysis	Q2 2024
3.1	Tests of basic reading performance using TTree version of PHYSLITE	Q4 2022
3.2	Prototyping of tools for columnar CP operations and other systematics	Q4 2023
3.3	Adoption of ROOT7 data structures	Q1 2024
3.4	Performance and ease-of-use assessment leading to decision on adoption	Q2 2024
AN-4	Development of columnar analysis infrastructure	Q2 2026
4.1	Prototyping of framework for orchestrating columnar CP operations	Q3 2024
4.2	Development of columnar skimming/augmentation	Q1 2025
4.3	Demonstrate end-to-end analysis using columnar tools/fkw + dist. comp.	Q4 2025
4.4	Development & roll-out of documentation/training for run 4 analysis	Q2 2026

Summary

- **ATLAS people:** please read the document and consider whether new long-term commitments could contribute to delivering these milestones whilst fitting into the strategy of NorCC
- **Non-ATLAS people:** are there areas where we could meaningfully collaborate?