



Operational Intelligence in the ATLAS Continuous Integration System

Alexander Undrus (BNL) on behalf of the ATLAS Collaboration

26th International Conference on Computing in High Energy & Nuclear Physics (CHEP 2023)


May 2023



The ATLAS detector is the largest volume detector at the Large Hadron Collider (LHC) at CERN (the European Organization for Nuclear Research) in Switzerland.



The ATLAS Collaboration consists of over 5500 members from more than 180 institutions in 42 countries. ATLAS takes experimental physics into unexplored territories – searching for new processes and particles that could change our understanding of energy and matter.



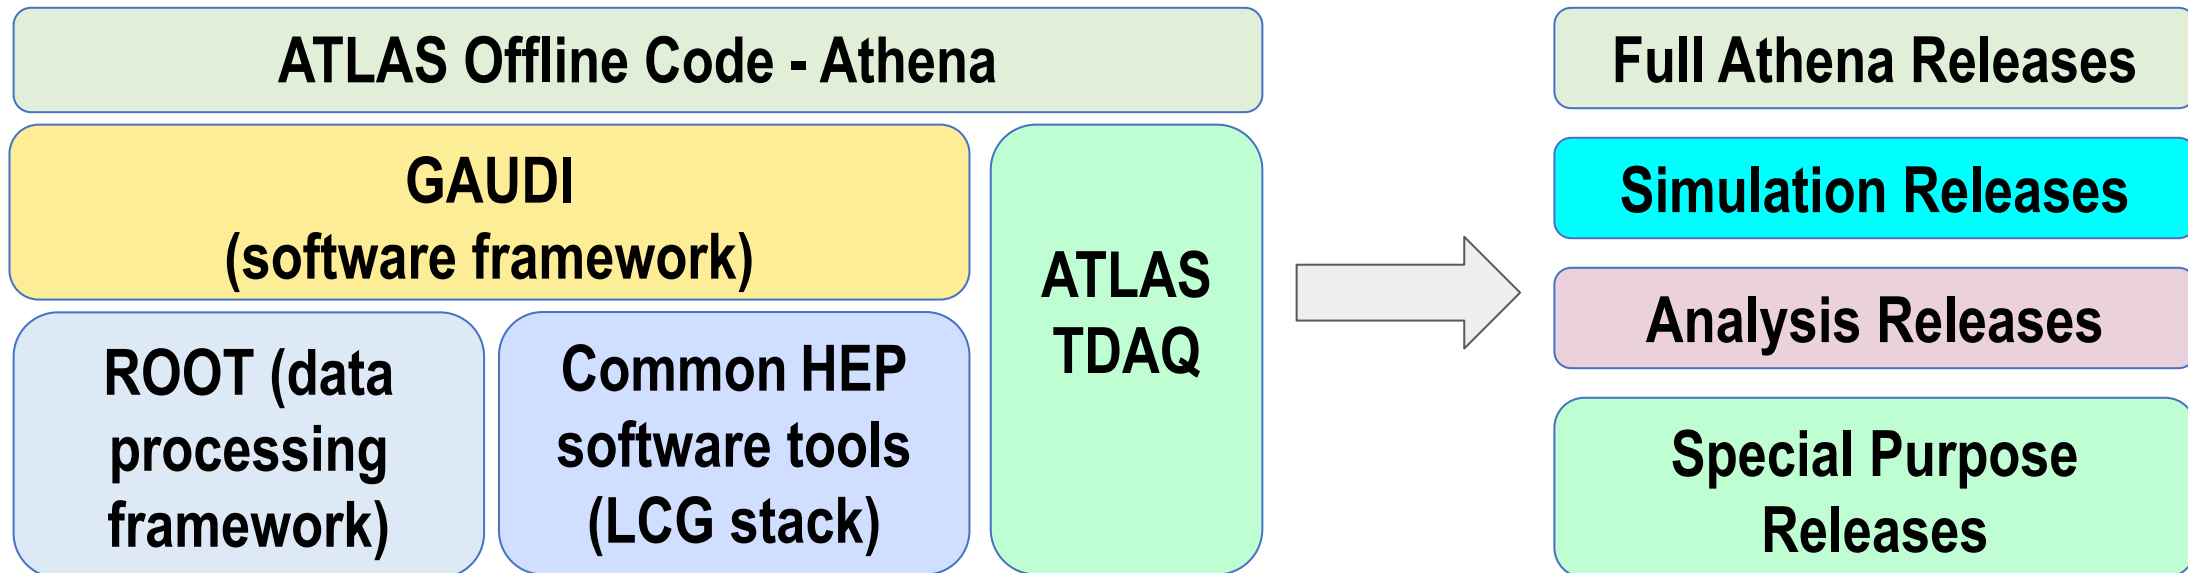
ATLAS records over 10,000 TB of new raw data per year, distributed over 130 computing centers worldwide. The ATLAS custom software totals more than 5 million lines of code. It converts detector signals into information that physicists can study.

This talk

- Describes the role of the ATLAS Continuous Integration (CI) System in the ATLAS offline software development infrastructure
- Outlines the CI system components and processes
- Details Operational Intelligence techniques to accelerate CI jobs and lower operating costs
- Explains the Directed Acyclic Graph (DAG) approach in CI pipelines
- Reports achieved improvements

ATLAS Offline Code Base


- Public under an APACHE 2.0 license
- Simulation, Reconstruction, Calibration, Analysis software
- Over 5 million lines of C++ and Python code
- Under constant development to get the latest technologies, analysis techniques
- Require over 200 external packages (mostly supplied by CERN SFT team, ATLAS TDAQ releases, GAUDI architecture framework, generators)
- Full and project-specific (e.g. Analysis) releases are available



ATLAS Developers Community

509 ATLAS Athena Developers in 2022

- Welcoming environment
- Inclusive pathways for new developers with wide range of backgrounds and experiences from all over the world
- Mentoring programs, seamless communication tools, detailed documentation:



Introductory tutorials at ATLAS Induction Days

Tutorials for subsystem's software (e.g. Analysis Tools, Trigger)

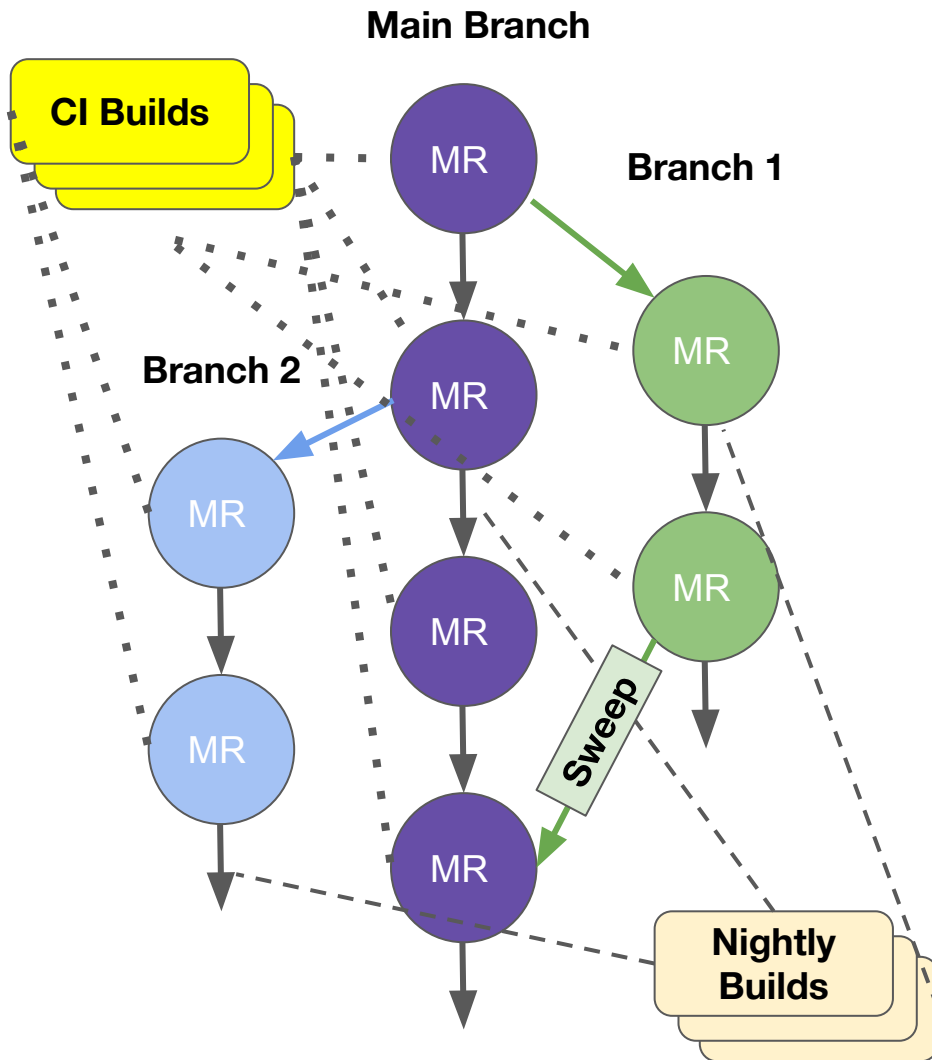
GitLab - web based code manager facilitating collaborative development

Software tutorial weeks (~ 3 times a year)

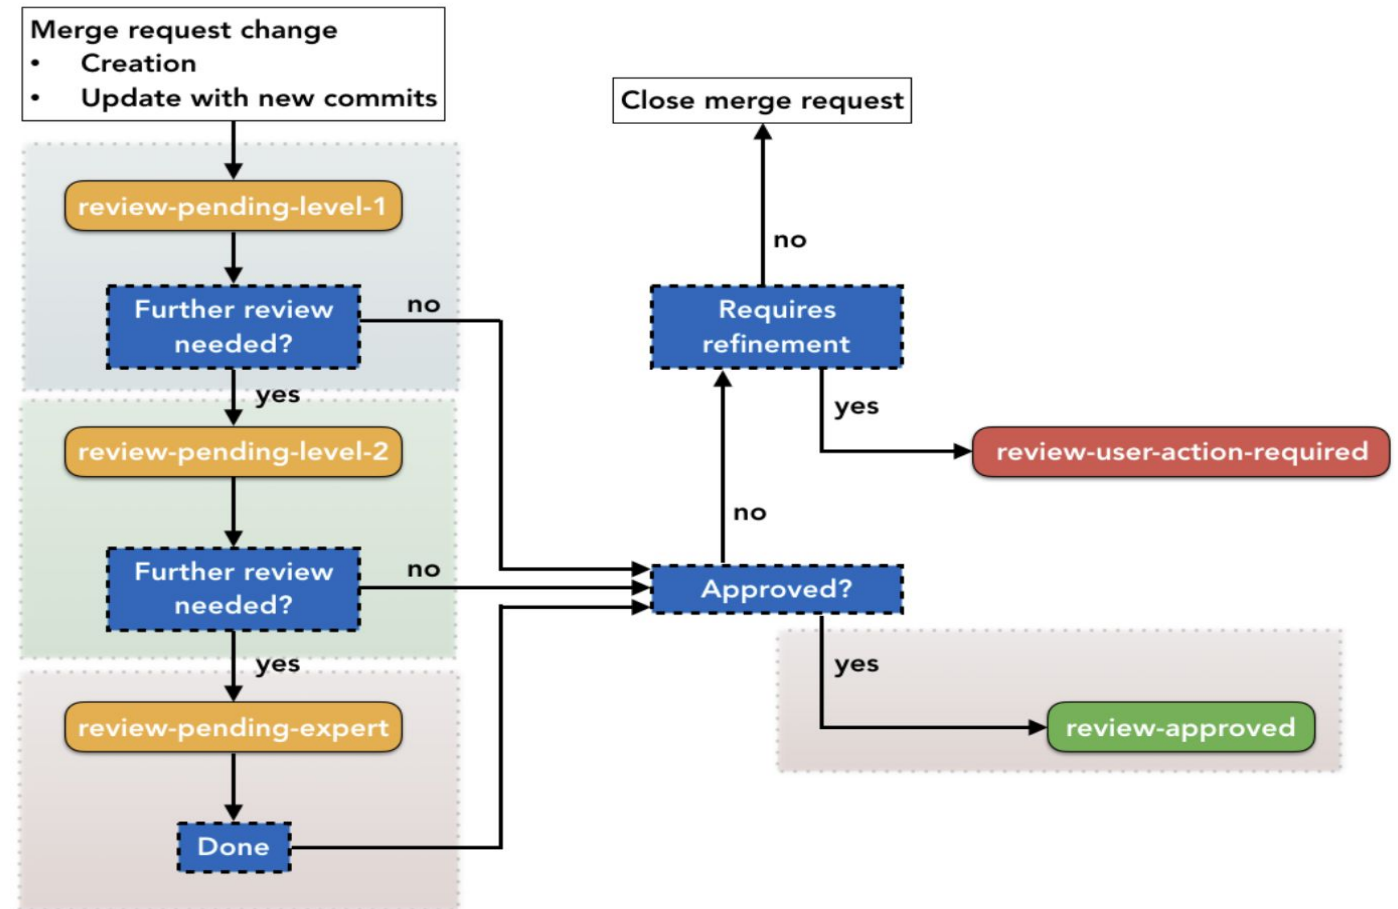
Documentation server powered by Jekyll and Bootstrap

Powerful ATLAS CI system for automatic testing and deployment of code

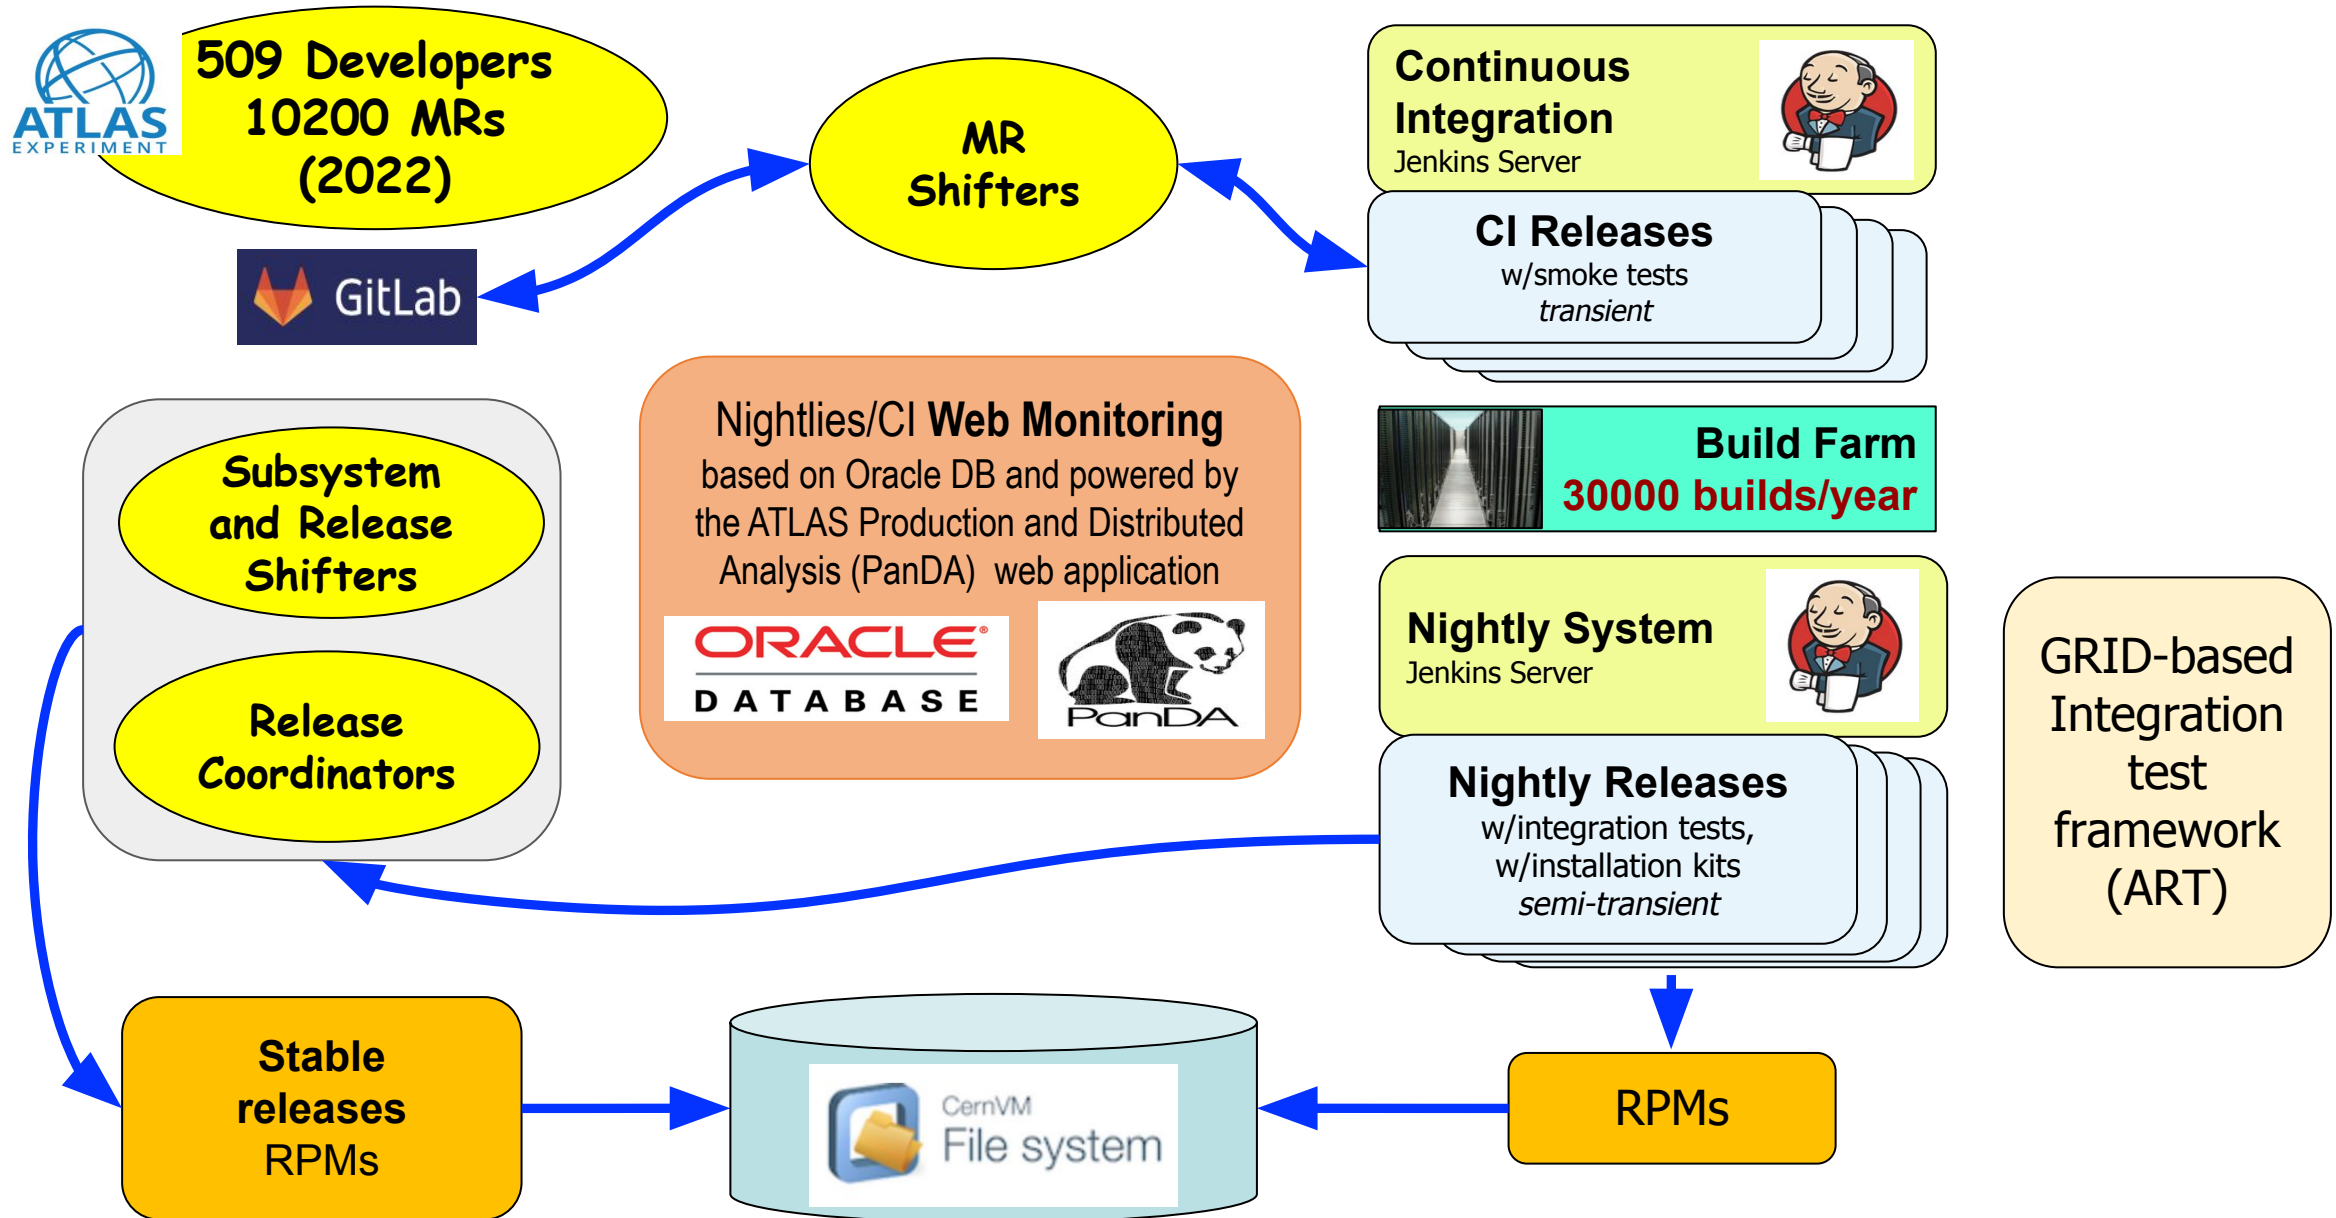
ATLAS Software Development Workflow



- CI build for each GitLab Merge Request (MR) creation/update
- Every working day, there are two level-1 and two level-2 shifters reviewing each MR targeting the main and production branches
- Every night, approved and merged MRs are validated in nightly builds



ATLAS CI and Nightly Systems at a Glance



ATLAS CI System Summary

- **Key component of ATLAS software workflow – Jenkins based build and testing systems interconnected with GitLab**
 - **Fast turnaround is crucial** for reaching development milestones on time
- **Big scale and complexity: 17600 CI jobs completed in 2022**
 - Each CI job includes compilation of up to 5 million lines of C++ code and comprehensive testing
- **Large-scale build farm of 1400 fast cores (shared with the Nightly System)**
 - ~ 7% smaller than it was 5 years ago
 - CI System consumes about 50% of farm resources
- **Multiple branches, platforms**
- **Full and project-specific (e.g. Analysis) software builds**
- **Fast and efficient pipelines with directed acyclic graph (DAG) features**
- **Rapid and comprehensive feedback to developers**
 - Dynamic monitoring is based on the Oracle DB technology and integrated with the ATLAS PanDA web service
 - Job results are posted directly to GitLab MR views

ATLAS CI Operational Intelligence Goals

- Provide faster software development workflow
- Improve decision-making by providing comprehensive feedback for proposed software updates and patches
- Promote welcoming and inclusive learning environment
- Use computing resources in environmentally responsible way by enhancing performance and efficiency

ATLAS CI Operational Intelligence Techniques

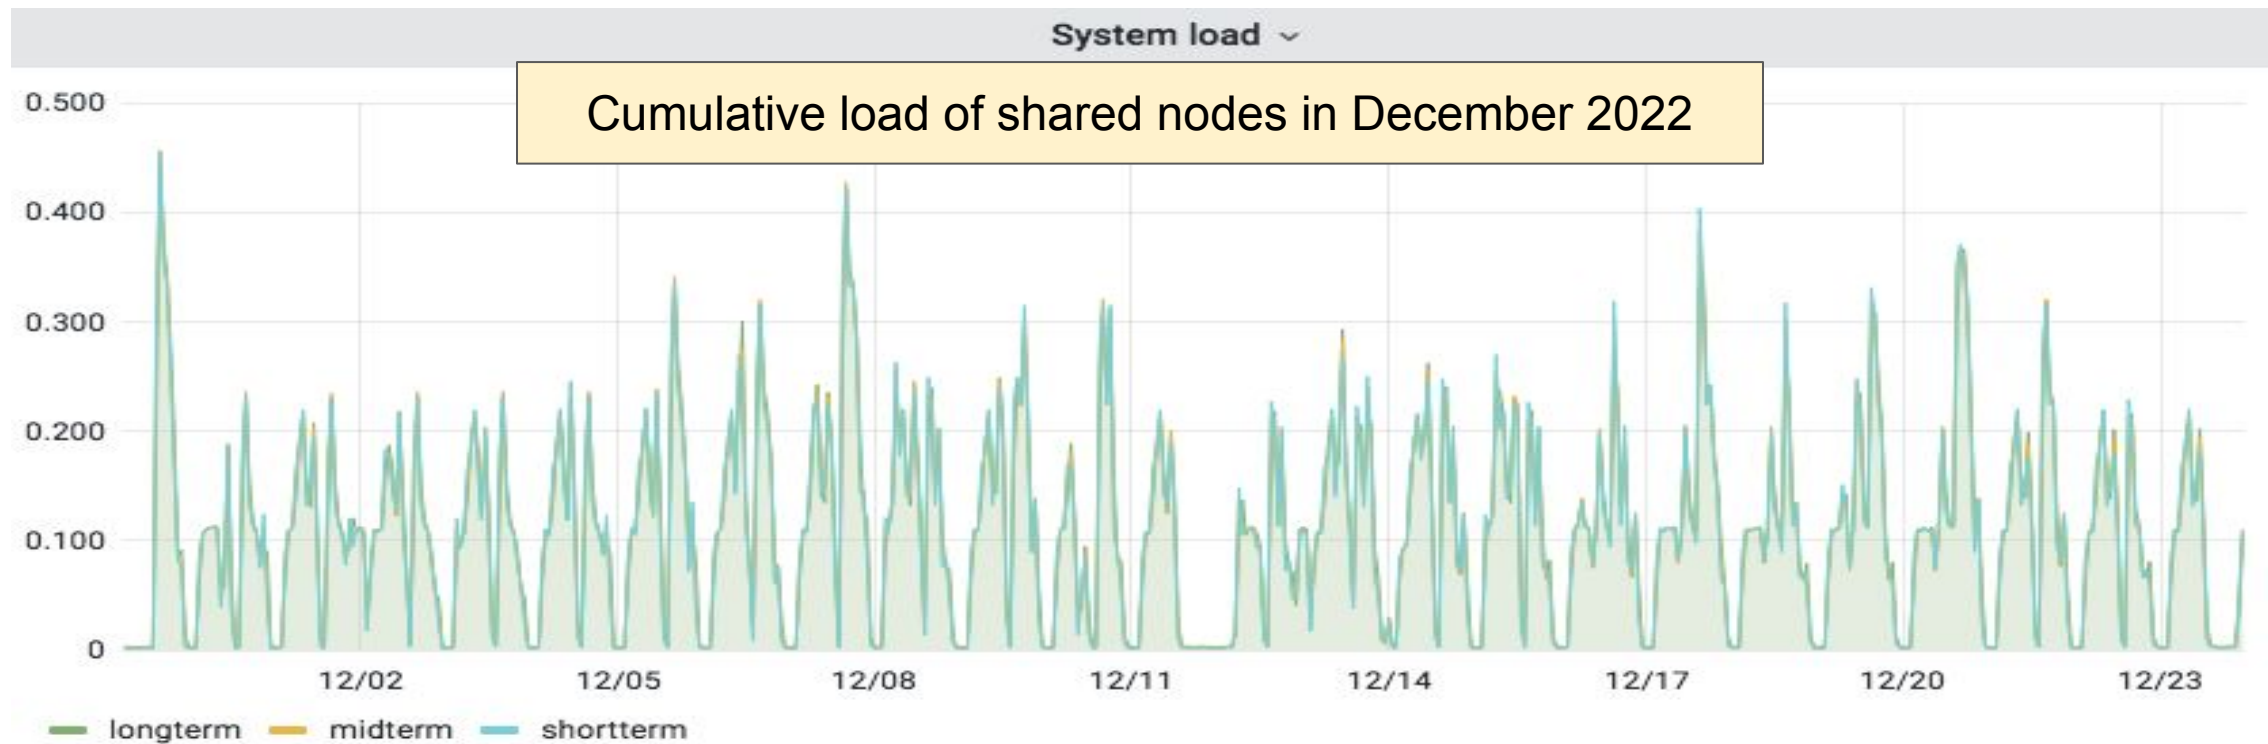
- Parallelization of operations (compilations, testing)
- Timely hardware upgrades, maintaining efficient resource utilization
- Intelligent CI pipelines
 - Directed acyclic graph (DAG) for building optimum job scenarios to achieve the quickest execution and less energy consumption
- Comprehensive collection and analysis of job results
 - Learning from accumulated data for generation of DAG scenarios
- Real-time visibility of results with dynamic monitoring and posting to GitLab

Parallelization of operations

- CMake build system with incremental and parallelization functionalities enabled
 - Rebuilding only the parts of the code affected the MR changeset -> speed up by the factor 3 - 10
 - Builds run on multi-core systems with maximum compiler and linker parallelization
- Builds are tested with the CTest tool with optimized parallelization parameters
 - Tests configurations specify number of required process slots
- In 2022 the CI farm hardware was upgraded from 16-core to 64-core nodes
 - Accelerated jobs by the factor of 3 due to enhanced parallelization

Efficient resource utilization

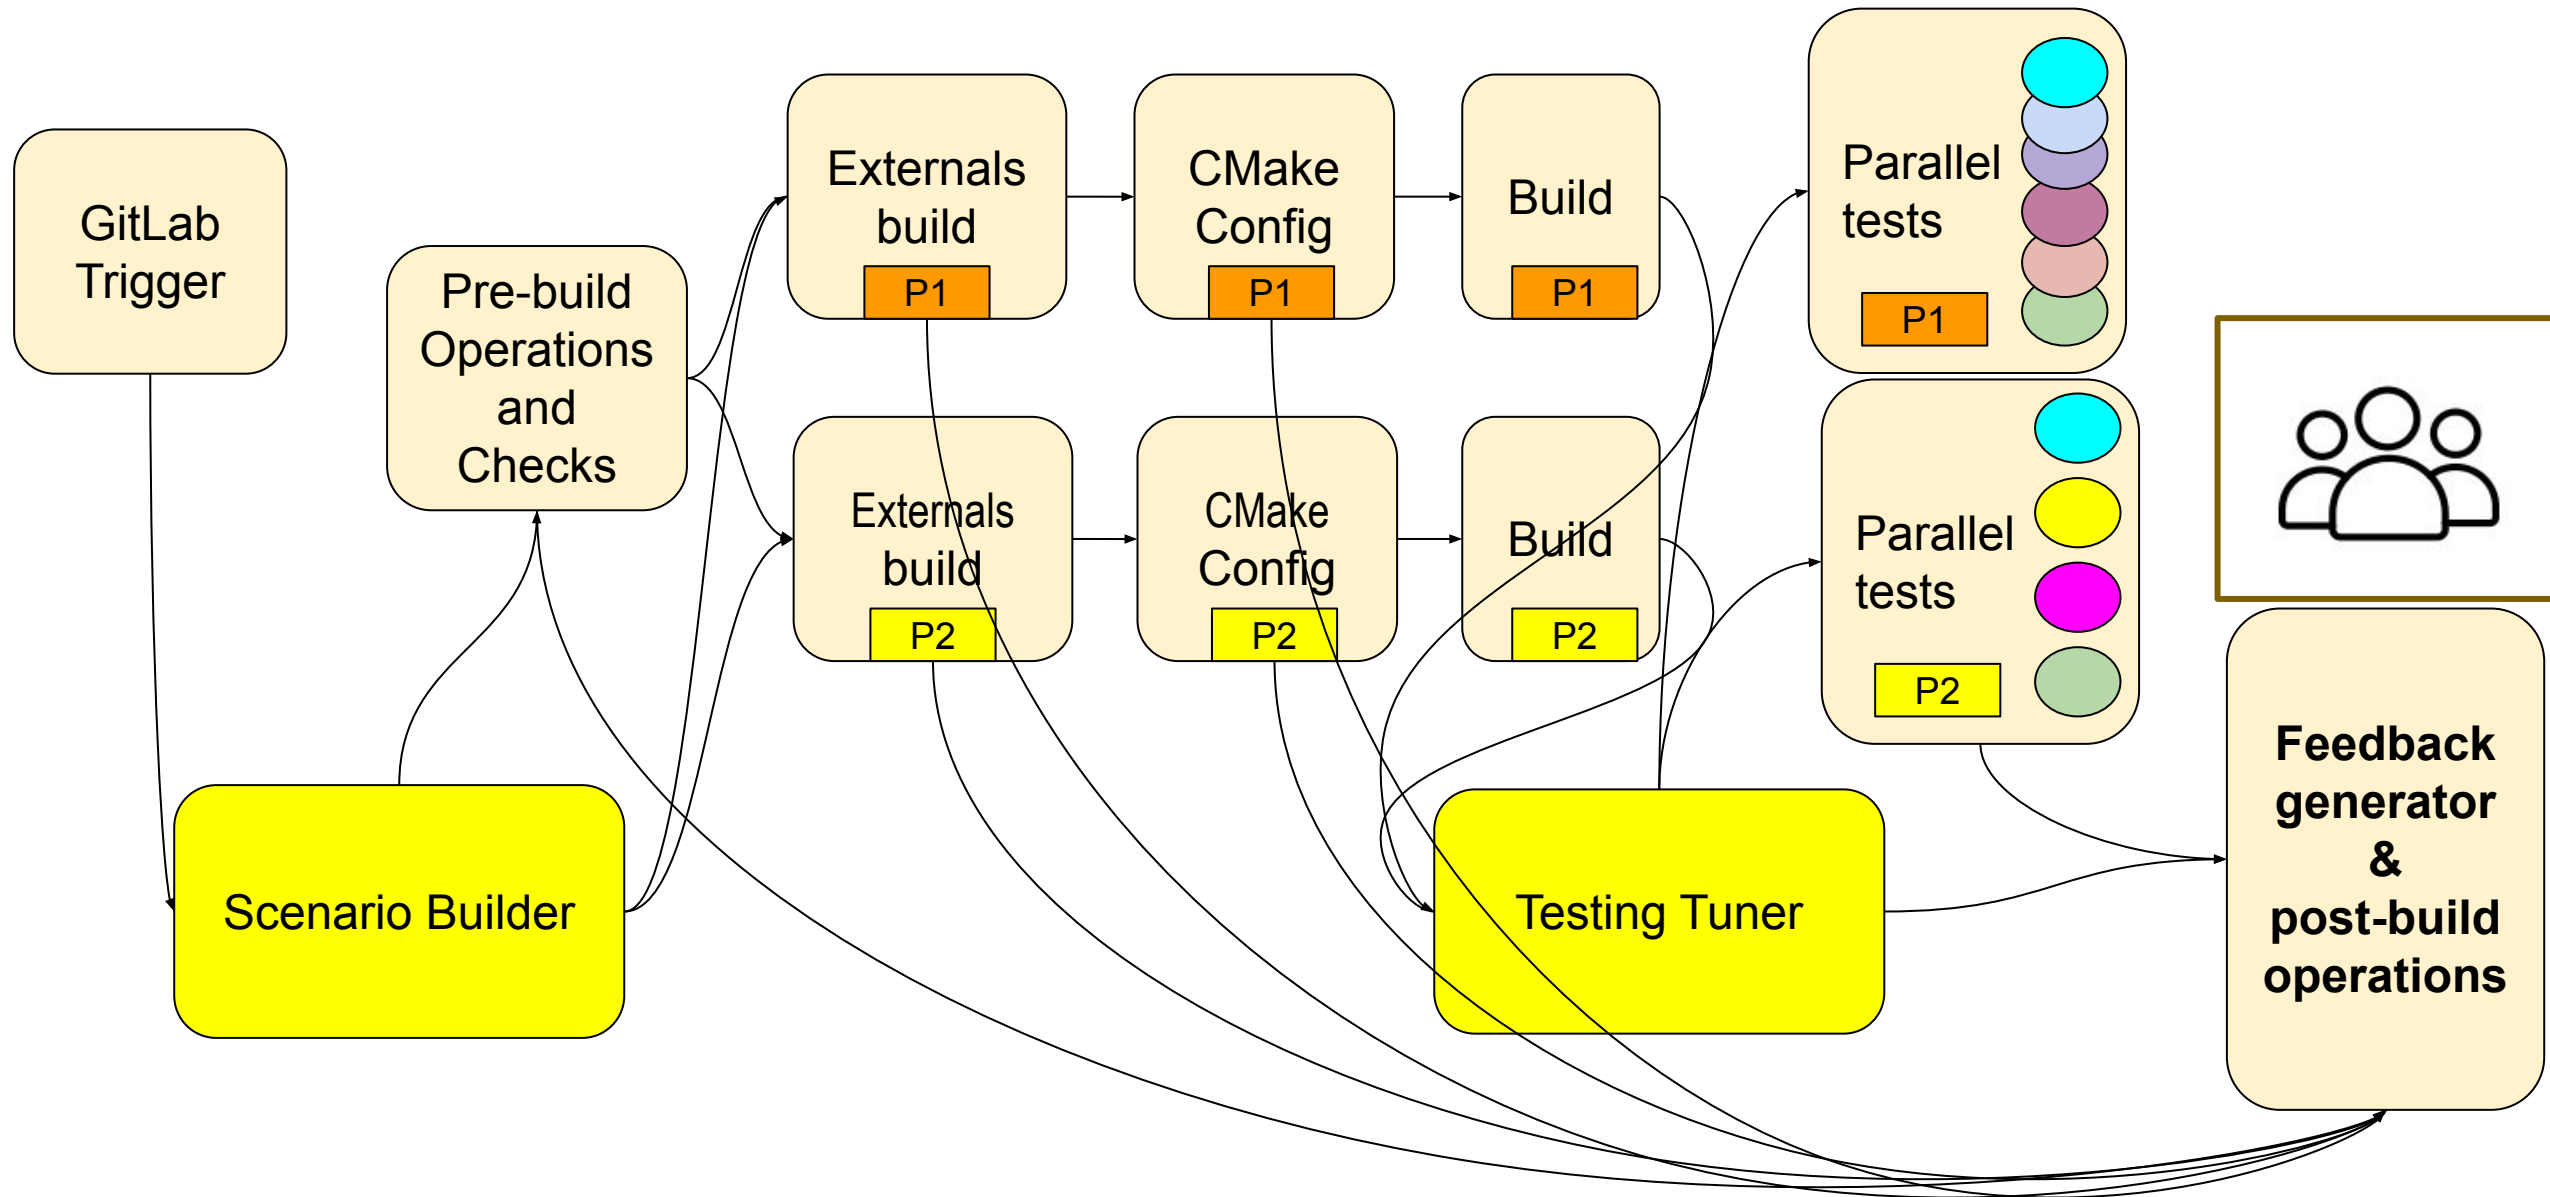
- Maximum system loads with compilations and testing parallelization
- Build nodes are shared between the CI and Nightly system
 - Most CI jobs run during day time
 - Nightly builds are performed at night time when development activity is low



DAG Concept for CI Pipelines

- Directed Acyclic Graph (DAG) is used to represent the dependencies between different stages of pipelines
 - Directed: each phase has at least one upstream or downstream task
 - Acyclic: no cycles (jobs do not retry themselves)
- DAG scenario is generated at the beginning of a CI job
 - Set of tasks to be executed in a specific order
 - Goal: select the most economical way to provide sufficient information to make a MR decision
 - Job dimensionality reduction: only build software projects affected by the changeset (from 1 to 6 projects)
 - Bypass unnecessary stages (e.g. cancellation of testing after unsuccessful compilation)
 - Tailored testing focused on the functionalities affected by the changeset
 - Testing can be canceled altogether (e.g., for documentation updates)

Example of the pipeline with 2 software projects built



Machine Learning Techniques for Test Tailoring

- Parameters and results of CI jobs are stored in the Oracle database
- Generation of a testing scenario includes analysis of historical data on test results and code changes to determine which tests are sensitive to potential code defects in the software domain(s) of the changeset
- System training jobs are run in between CI jobs (mostly on weekends when the system is not busy)
 - Sensitivity of tests to artificial artifact removals in specific software domains is measured, stored in the DB, and later used during scenario generation
- Full set of CI tests probes all code domains, includes event simulation, reconstruction and analysis and can take more than one hour
 - Test tailoring can reduce the runtime of 1/2 of ATLAS CI jobs by 30 to 50%
- Currently under verification, gradually put in production

Feedback to Developers

- Scenario decisions and job results are posted on GitLab MR pages
- Developers can override job optimization decisions by adding labels to GitLab MR
 - E.g., request the full testing or complete (non-incremental) compilations
- Database-backed monitoring displays job compilation and testing results, delivers performance plots, provides log files with error and warning messages highlighted

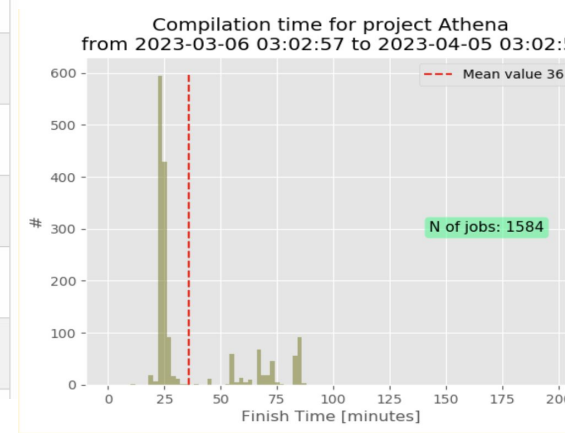
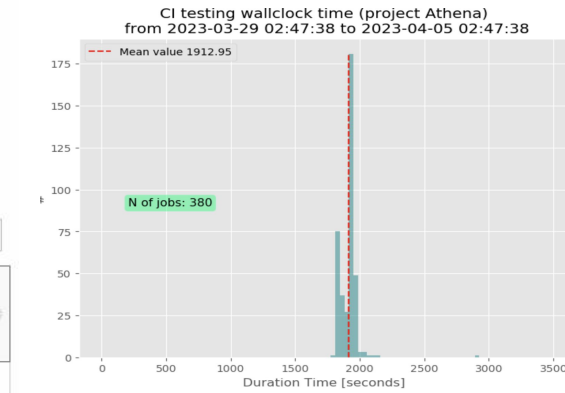
ATLAS CI Builds Summary

- [Link to global system page](#)
- [Nightly and CI systems features and policies](#)
- [Merge Request Problems Display \(from C. Suster\)](#)
- [CI analytics \(build and test time distributions\)](#)

Show 100 entries

Search:

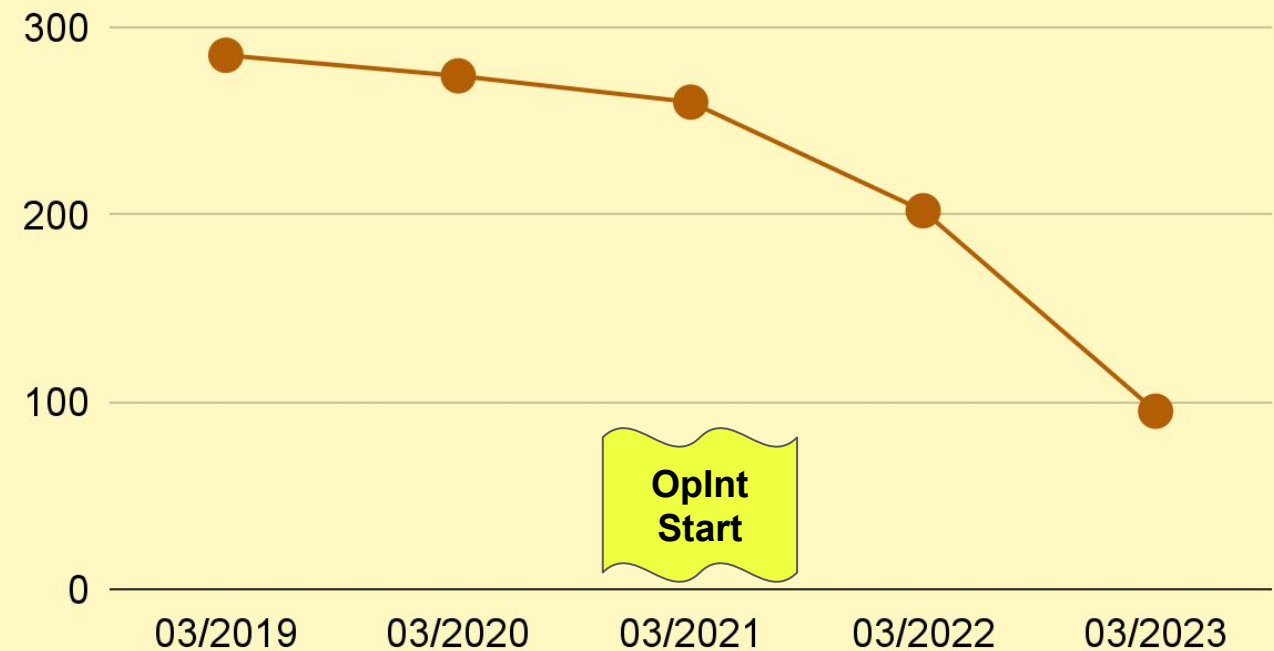
| Release | Platform | Project | git branch (link to MR) | Job time stamp | git clone | Externals build | CMake config | Build time | Comp. errors (w/warnings) | Test time | CI tests errors (w/warnings) | Host |
|---------------------------|---------------------------|---------------|-------------------------|------------------|-----------|-----------------|--------------|------------------|---------------------------|------------------|------------------------------|---------------|
| MR-61986-2023-04-04-13-35 | x86_64-centos7-gcc112-opt | Athena | master | 2023/04/04 13:35 | ✓ | ✓ | ✓ | 2023/04/04 11:32 | 0 (0) | 2023/04/04 12:28 | 1 (1) | aibuild64-015 |
| MR-62097-2023-04-04-13-25 | x86_64-centos7-gcc112-opt | Athena | 23.0 | 2023/04/04 13:25 | ✓ | ✓ | ✓ | 2023/04/04 12:44 | 0 (0) | 2023/04/04 12:50 | 0 (0) | aibuild64-011 |
| MR-62091-2023-04-04-13-22 | x86_64-centos7-gcc112-opt | AthSimulation | 23.0 | 2023/04/04 13:22 | ✓ | ✓ | ✓ | 2023/04/04 12:25 | 0 (0) | 2023/04/04 12:26 | 0 (0) | aibuild64-002 |
| MR-62091-2023-04-04-13-22 | x86_64-centos7-gcc112-opt | Athena | 23.0 | 2023/04/04 13:22 | ✓ | ✓ | ✓ | 2023/04/04 12:13 | 0 (0) | 2023/04/04 12:26 | 0 (0) | aibuild64-002 |
| MR-62094-2023-04-04-13-10 | x86_64-centos7-gcc112-opt | Athena | master | 2023/04/04 13:10 | ✓ | ✓ | ✓ | 2023/04/04 12:35 | 0 (0) | 2023/04/04 12:37 | 0 (0) | aibuild64-013 |
| MR-62058-2023-04-04-13-02 | x86_64-centos7-gcc112-opt | Athena | 23.0 | 2023/04/04 13:02 | ✓ | ✓ | ✓ | 2023/04/04 12:29 | 0 (0) | 2023/04/04 12:30 | 0 (0) | aibuild64-012 |



Summary and Outlook (1/2)

- The ATLAS Continuous Integration is a large scale building and testing system supporting development of 5 million lines codebase
- Since its start in 2021, Operational Intelligence and hardware upgrades helped to enhance the system performance and efficiently use computing resources
- Since 2019, the CI system runs ~17000 jobs annually
- The average job run time was reduced from 300 to 100 minutes
- At the same time, the core count was reduced by 7% at the shared with the nightly system build farm
- Faster feedback to developers helped reaching development milestones on schedule
- It also promotes welcoming and inclusive learning environment in the community

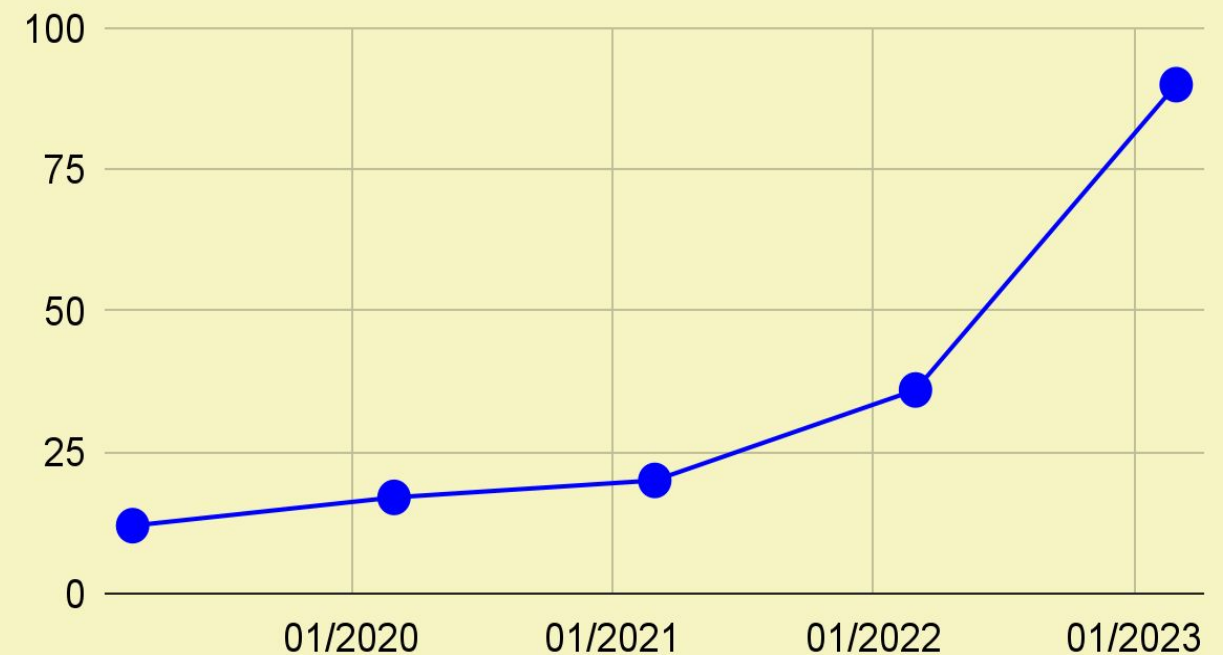
Average ATLAS CI Job Run Time, minutes



Summary and Outlook (2/2)

- Current focus is on testing optimization
- Additional testing is welcome as it helps to identify issues before they impact the production
- Intelligent test tailoring allows to accelerate jobs by up to 50%
- Care is taken to avoid skipping needed tests: test tailoring is being put in production in a piecemeal manner

Number of Integration Tests in CI Jobs



Operational Intelligence helps to accelerate ATLAS CI jobs and use computing resources in socially and environmentally responsible way. As a result, ATLAS community is enabled to deliver high-quality software products efficiently, in welcoming and comfortable environment