

# **My personal experience with ePIC's CI/CD workflow**

---

Nathan Brei

Sept 19, 2025

ePIC Software Review Committee

# Agenda

1. Context
2. Compatibility requirements
3. Code quality checks
4. Benchmarking and validation
5. Recommendations

# Context



- I'm a computer scientist, not a physicist
- I've been working on the JANA2 reconstruction framework since I joined JLab in 2019
- I've been the lead developer of JANA2 ever since David became too busy
- Much of my work has been driven by requirements and desiderata from ePIC
- However, I **rarely** contribute code directly to EICrecon
- Improving ePIC's overall developer experience requires us taking a systematic approach
- That is **not** what this is
- This is merely one **very opinionated** data point on an undulating surface

# Compatibility requirements

---

## Idea

- ElCrecon's Github CI builds against a very thorough test matrix. It is roughly like so:

$$\begin{aligned} \text{test\_matrix} = & \{\text{gcc}, \text{clang}\} \times \\ & \{\text{Debug}, \text{Release}\} \times \\ & \{(\text{None}), \text{USE\_ASAN}, \text{USE\_TSAN}, \text{USE\_UBSAN}\} \times \\ & \{\text{eic\_container:nightly}\} \\ & \cup (\text{clang}, \text{Release})\{\text{eic\_container:v25.01.1-stable}, \text{eic\_container:v25.05.0-stable}\} \end{aligned}$$

- Note that each `eic_container` version includes a single version of gcc, clang, JANA2, edm4eic, etc.
- The `eic_container` environment leverages Spack for reproducible package management

# Compatibility requirements

## Problems

- 6-month backwards compatibility requirement is a formidable obstacle to introducing new framework features
- Backwards compatibility requirement pushes dependency version logic into preprocessor directives and `constexpr` if blocks
- This also makes data model changes much harder!
- I don't see the benefit to supporting older environments, apart from nightly and the latest release.

```
1  if constexpr (10000 * JVersion::major +
2                100 * JVersion::minor +
3                1 * JVersion::patch > 20403) {
4      default_plugins.emplace_back("splitting");
5  }
```

```
1  #if ((JANA_VERSION_MAJOR == 2) && (JANA_VERSION_MINOR >= 3))
2      || (JANA_VERSION_MAJOR > 2)
3  #define JANA_NEW_CALLBACK_STYLE 1
4  #else
5  #define JANA_NEW_CALLBACK_STYLE 0
6  #endif
7
8  class JEventSourcePODIO : public JEventSource {
9  public:
10     JEventSourcePODIO();
11     ~JEventSourcePODIO() override;
12     void Open() override;
13     void Close() override;
14
15     #if JANA_NEW_CALLBACK_STYLE
16         Result Emit(JEvent& event) override;
17     #else
18         void GetEvent(std::shared_ptr<JEvent>) override;
19     #endif
```

# Compatibility requirements: Effect on development workflow

- The previous examples kept the preprocessor logic in a limited area, but making a change across the entire codebase is basically impossible, e.g. removing the `app` argument from `JOmniFactoryGeneratorT`:
- The next major feature is external wiring, which by definition touches every single factory. For it to be mergeable, it effectively needs to have zero visible API changes.
- Features requested by ePIC have to be deeply tested outside of EICrecon before introduction. No co-evolution is possible.
- To compensate, I'm building a complete reconstruction chain inside the JANA examples

```
1
2  app->Add(new
3      JOmniFactoryGeneratorT<InclusiveKinematicsML_factory>(
4          "InclusiveKinematicsML",
5          {"InclusiveKinematicsElectron", "InclusiveKinematicsDA"},
6          {"InclusiveKinematicsML"});
7
8  // Literally hundreds of these app->Add() calls across EICrecon
9  // A previous version of JANA2 made the extra `app` parameter
10 // unnecessary, but we can't remove it due to backwards compat
11
```



## Code quality checks

---

# Automated code quality checks: clang-tidy, etc

## Idea

- Leverages heavy-duty tooling to enforce code style, correctness, and performance
  - Clang-tidy
  - `-Werror -Wall -Wextra -Wundef`
  - Asan, Tsan, UBSan
  - Profiling and code coverage

## Problems

- Clang-tidy is configured quite sensibly
- Previously had problems with `bugprone-easily-swappable-parameters` in one of my interfaces
- Have incorporated some of these in JANA2's CI, plan on incorporating more

```
1  HeaderFilterRegex: ''
2  Checks: '
3    bugprone-*,
4    concurrency-*,
5    cppcoreguidelines-*,
6    modernize-*,
7    portability-*,
8    readability-*,
9    -bugprone-easily-swappable-parameters,
10   -bugprone-macro-parentheses,
11   -bugprone-narrowing-conversions,
12   -bugprone-unchecked-optional-access,
13   -modernize-use-trailing-return-type,
14   -modernize-avoid-c-arrays,
15   -modernize-use-nodiscard,
16   -readability-function-cognitive-complexity,
17   -readability-identifier-length,
18   -readability-magic-numbers,
19   -readability-math-missing-parentheses,
20   -cppcoreguidelines-avoid-c-arrays,
21   -cppcoreguidelines-avoid-magic-numbers,
22   -cppcoreguidelines-narrowing-conversions,
23   -cppcoreguidelines-owning-memory
24  '
25  FormatStyle: file
```

# Automated code quality checks: clang-format

## Idea

- Ends inconsistent and misleading formatting (e.g. mixed tabs and spaces) for once and for all

## Problems

- The clang-format settings mangle the diffs
  - Lists of strings (we have a lot of these) wrap instead of having one-per-line
  - Assignments get lined up with adjacent lines
  - Can be disabled via `// clang-format off`
- The automated commits create merge conflicts
  - Diff blocks become much larger than necessary, often the entire file
  - This happens on **every** push to the PR
- The clang-format settings change frequently

```
43 // "TaggerTrackerRecHits_aligned", "TrackerEndcapRecHits_aligned",
44 // "VertexBarrelHits_aligned");
45
46 std::vector<std::vector<std::string>> m_simtrackerhit_collection_names_aligned = {{
47     "B0TrackerRecHits_TK_aligned", "BackwardMPGDEndcapRecHits_TK_aligned",
48     "DIRCBarRecHits_TK_aligned", "DRICHRecHits_TK_aligned",
49     "ForwardMPGDEndcapRecHits_TK_aligned", "ForwardOffMTrackerRecHits_TK_aligned",
50     "ForwardRomanPotRecHits_TK_aligned", "LumiSpecTrackerRecHits_TK_aligned",
51     "MPGDBarrelRecHits_TK_aligned", "OuterMPGDBarrelRecHits_TK_aligned",
52     "RICHEndcapNRecHits_TK_aligned", "SiBarrelTrackerRecHits_TK_aligned",
53     "TOFBarrelRecHits_TK_aligned", "TOFEndcapRecHits_TK_aligned",
54     "TaggerTrackerRecHits_TK_aligned", "SiEndcapTrackerRecHits_TK_aligned",
55     "SiBarrelVertexRecHits_TK_aligned"}}};
56
57 std::vector<std::vector<std::string>> m_simtrackerhit_collection_names = {
58     {"B0TrackerRecHits_TK_aligned", "BackwardMPGDEndcapRecHits_TK_aligned",
59     "DIRCBarRecHits_TK_aligned", "DRICHRecHits_TK_aligned",
60     "ForwardMPGDEndcapRecHits_TK_aligned", "ForwardOffMTrackerRecHits_TK_aligned",
61     "ForwardRomanPotRecHits_TK_aligned", "LumiSpecTrackerRecHits_TK_aligned",
62     "MPGDBarrelRecHits_TK_aligned", "OuterMPGDBarrelRecHits_TK_aligned",
63     "RICHEndcapNRecHits_TK_aligned", "SiBarrelTrackerRecHits_TK_aligned",
64     "TOFBarrelRecHits_TK_aligned", "TOFEndcapRecHits_TK_aligned",
65     "TaggerTrackerRecHits_TK_aligned", "SiEndcapTrackerRecHits_TK_aligned",
66     "SiBarrelVertexRecHits_TK_aligned"}}};
67
68 std::vector<std::vector<std::string>> m_simtrackerhit_collection_names = {
69     {"B0TrackerRecHits_TK", "BackwardMPGDEndcapRecHits_TK", "DIRCBarRecHits_TK",
70     "DRICHRecHits_TK", "ForwardMPGDEndcapRecHits_TK", "ForwardOffMTrackerRecHits_TK",
71     "ForwardRomanPotRecHits_TK", "LumiSpecTrackerRecHits_TK", "MPGDBarrelRecHits_TK",
72     "OuterMPGDBarrelRecHits_TK", "RICHEndcapNRecHits_TK", "SiBarrelTrackerRecHits_TK",
73     "TOFBarrelRecHits_TK", "TOFEndcapRecHits_TK", "TaggerTrackerRecHits_TK",
74     "SiEndcapTrackerRecHits_TK", "SiBarrelVertexRecHits_TK"}}};
75
76 InitJANAPLugin(app);
77
78 app->Add(new JOmniFactoryGeneratorT<timeAlignmentFactory>({
79     jana::components::JOmniFactoryGeneratorT<timeAlignmentFactory>::TypedWiring{
80         .tag = "timeAlignment",
81         .level = JEventLevel::Timeslice,
82         .variadic_input_names = m_simtrackerhit_collection_names,
83         .variadic_output_names = m_simtrackerhit_collection_names_aligned});
84     .tag = "timeAlignment",
85     .level = JEventLevel::Timeslice,
86     .variadic_input_names = m_simtrackerhit_collection_names,
87     .variadic_output_names = m_simtrackerhit_collection_names_aligned});
88
89 // Unfolder that takes timeframes and splits them into physics events.
90 app->Add(new TimeframeSplitter());
```

# Automated code quality checks: include-what-you-use

## Idea

- C++'s `#include` mechanism is unhygienic
- Build times were getting too slow
  - Partially due to a large “datamodel glue” header that was being (transitively) imported everywhere: [#628](#)
  - Future versions of JANA2 won't need datamodel glue at all [#446](#)
  - Current build time is ~5 mins (depending on compiler flags)

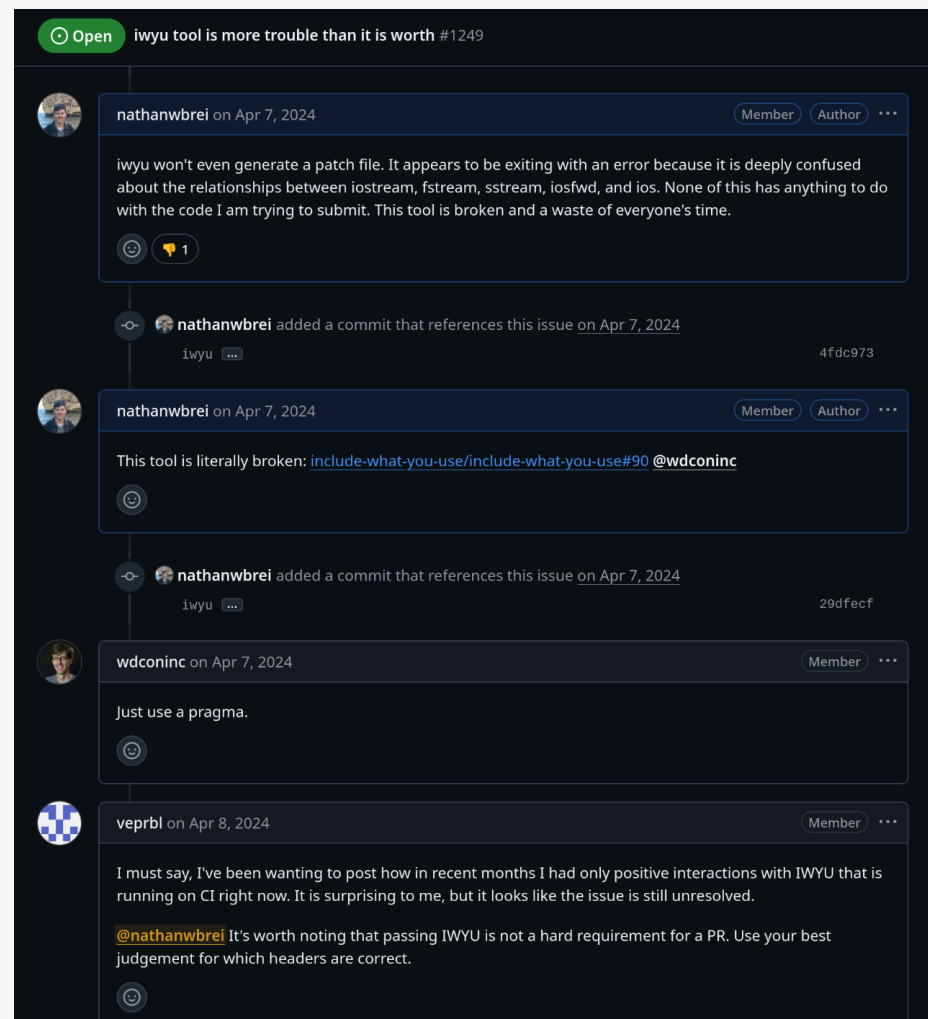
## Problems, part 1

- The `iwyu` tool produces a patch file, rather than automatically committing its changes
- The user must navigate to the “Artifacts” tab, download the patch file, apply the patch file, commit, push, and wait
- The `iwyu` tool gets caught in cycles, where it rejects its own patch and generates a new one
- It's complaints are not limited to the user's own changes
  - Upgrading a dependency can cause `iwyu` to rearrange includes throughout the entire codebase

# Automated code quality checks: include-what-you-use

## Problems, part 2

- iwyu struggles to understand forward declarations...
  - including 'iosfwd' in the standard library
  - The iwyu developers aren't concerned: [#144](#)
  - iwyu breaks every now and then for ePIC due to changing compiler/stdlib versions: [#1984](#)
- The solution is, once again, pragmas
  - There is a rich and sophisticated family of pragmas available ([link](#))
  - Sometimes all you need is `// IWYU pragma: keep`
  - Otherwise, you'll need deep knowledge of both C++ semantics and your dependencies' internals, and lots of trial and error



## Problems, part 3

- iwyu creates problems for me as a framework developer
- Template methods are an important part of my interface, but they create problems with circular includes
- To fix these, I split the problematic headers into
  - “Forward headers” containing the class and template declarations
  - “Full headers”, which import the forward header and also provide the template definitions
- iwyu insists on my users providing both!
  - Breaks the separation between my interface and internals
  - Forces me to annotate JANA2 code with iwyu pragmas

```
1  #include <DD4hep/DetElement.h>
2  #include <DD4hep/Detector.h>
3  #include <Evaluator/DD4hepUnits.h>
4  #include <JANA/JApplication.h>
5  #include <JANA/JApplicationFwd.h>
6  #include <JANA/Utils/JTypeInfo.h>
7
8  #include <functional>
9  #include <map>
10 #include <memory>
11 #include <string>
12 #include <vector>
13
14 #include "algorithms/digi/PhotoMultiplierHitDigiConfig.h"
15 #include "algorithms/pid/IrtCherenkovParticleIDConfig.h"
16 #include "algorithms/pid/MergeParticleIDConfig.h"
17 #include "algorithms/tracking/TrackPropagationConfig.h"
18 #include "extensions/jana/JOmniFactoryGeneratorT.h"
19 #include "factories/digi/PhotoMultiplierHitDigi_factory.h"
20 #include "factories/pid/IrtCherenkovParticleID_factory.h"
21 #include "factories/pid/MergeCherenkovParticleID_factory.h"
22 #include "factories/pid/MergeTrack_factory.h"
23 #include "factories/pid/RichTrack_factory.h"
24
25 #include "services/geometry/richgeo/ActsGeo.h"
26 #include "services/geometry/richgeo/RichGeo.h"
27 #include "services/geometry/richgeo/RichGeo_service.h"
```

## Automated code quality checks: Effect on workflow

- I know of no way to run the CI test suites locally automatically
- Instead, I write my code against my own unit and integration tests
- Once satisfied, I download an npsim test data file from a recent CI job, and run it against EICrecon locally, until I observe no crashes
- Then I push my code and create a pull request
- Almost every single time I push, the CI adds an additional commit changing around whitespace and includes, which clutter my history and introduce merge conflicts
- PRs take several days in the best case, and up to a year in the worst case, to be merged
- As an developer who prefers rebasing and stacking small PRs, I'm instead forced to split the automated commits, squash them into my clean commits, cherry pick my clean commits on top of main, and then force-push my branch.
- Developers less comfortable with git will repeatedly merge main back into their branch and accumulate a long messy history

# Benchmarking and validation

---



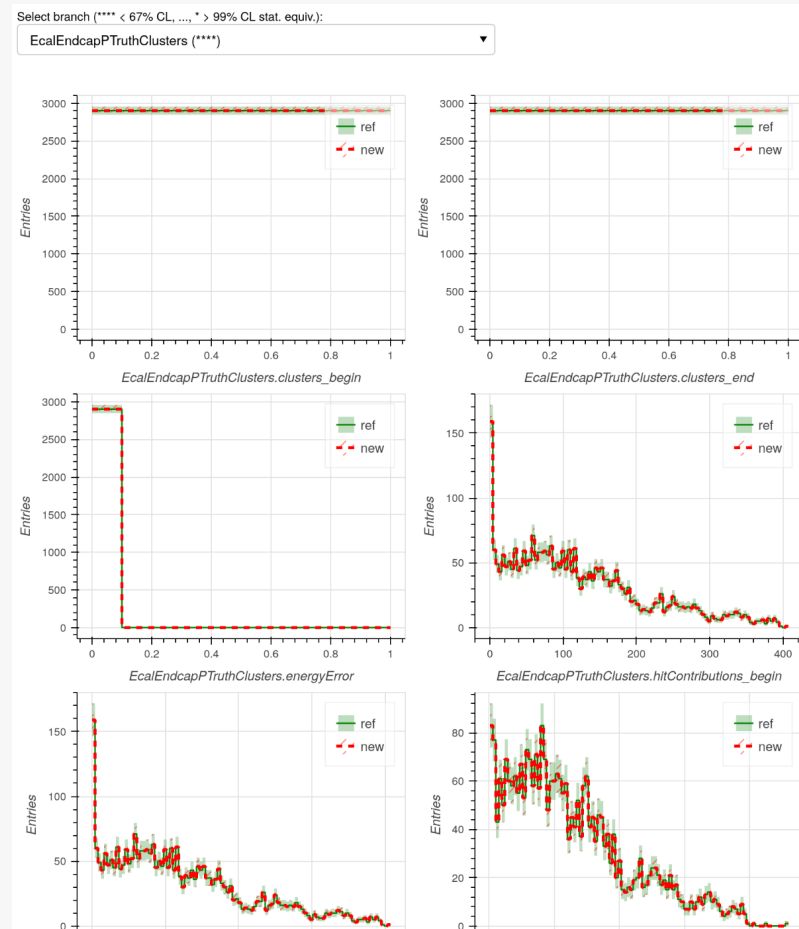
# Benchmarking and validation

## Idea

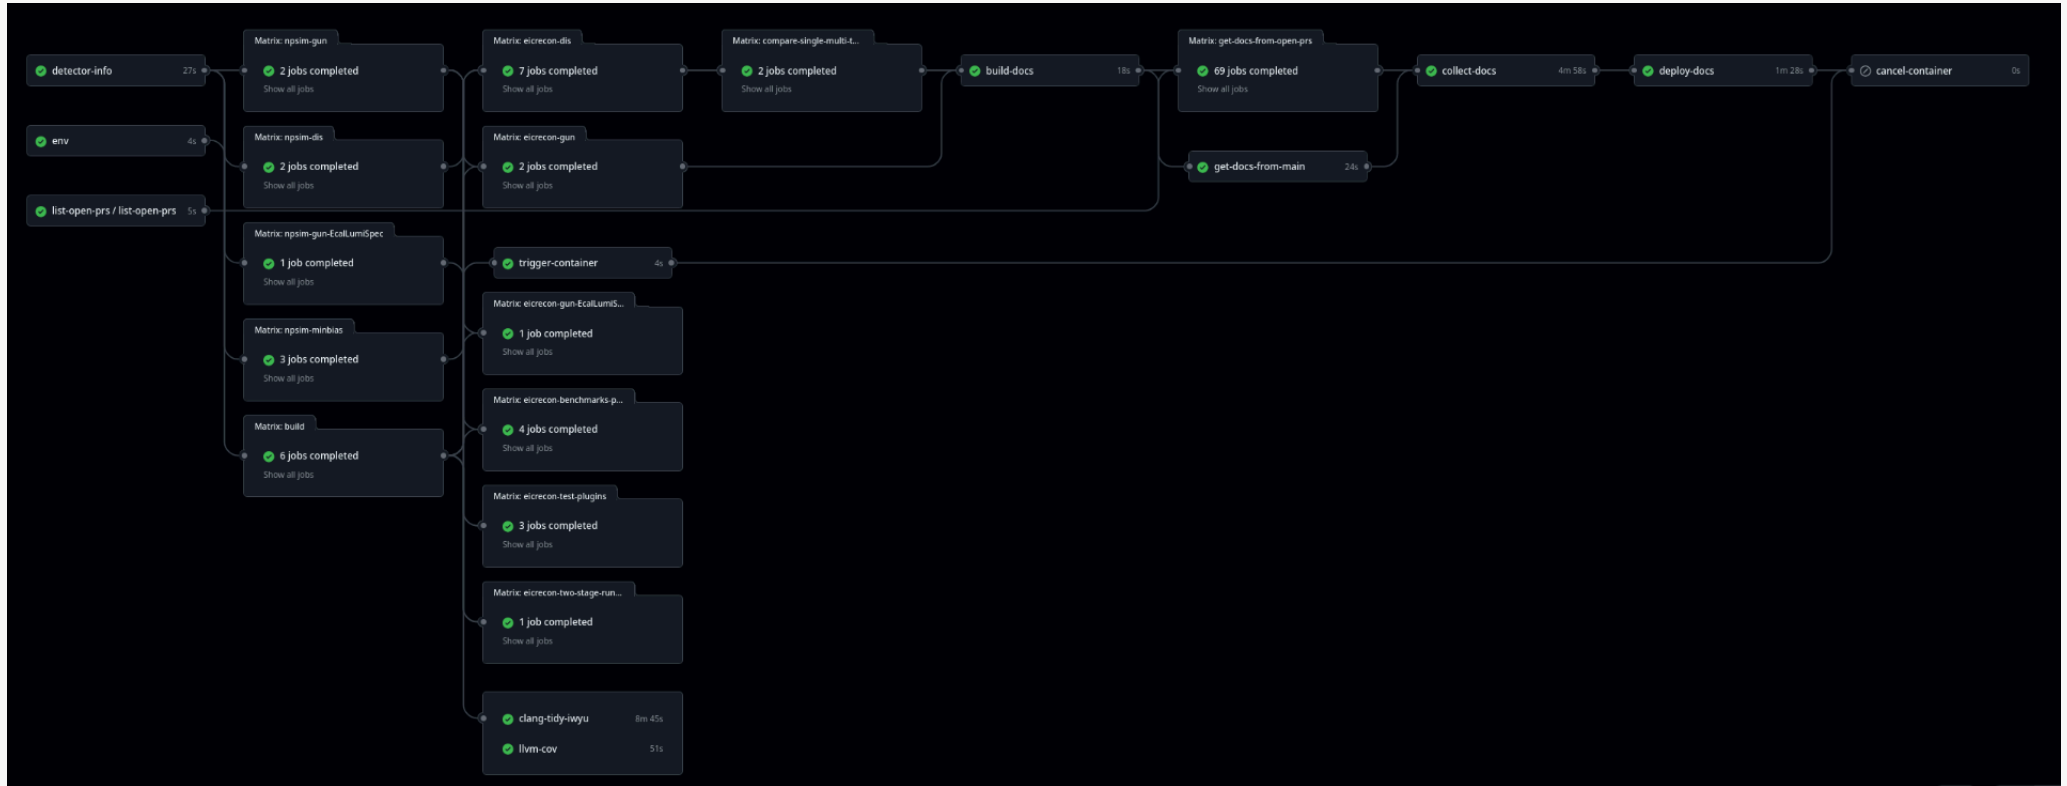
- Physics benchmarks are run for every PR, and returns a diff of each histogram. A CI job will fail if the diff is above some threshold.
- I love this. This is amazing!

## Problems

- (Most) validation does **not** happen on <https://github.com/eic/EICrecon>, but rather [https://eicweb.phy.anl.gov/containers/eic\\_container](https://eicweb.phy.anl.gov/containers/eic_container)
- Does validation still use Juggler?
- Are benchmarks deterministic? What happens when a code change which cannot change a benchmark result still changes a benchmark result?



# The public GitHub CI pipeline



# The ANL GitLab CI pipeline

The screenshot displays the GitLab CI pipeline interface for the project 'eic\_container: eic/ElCrecon: Flag clusters as EMCal or HCal'. The pipeline is titled 'Pipeline #128763 - EIC / ✕' and is currently in a 'running' state. The URL shown is 'https://eicweb.phy.anl.gov/EIC/benchmarks/detector\_benchmarks/-/pipelines/128763'. The interface includes a sidebar with navigation options like 'Project', 'Pipelines', 'Jobs', 'Pipeline schedules', 'Artifacts', 'Deploy', 'Operate', 'Monitor', and 'Analyze'. The main content area shows the pipeline stages and their corresponding jobs. The stages are: 'eic\_container', 'status-pending', 'config', 'simulate', 'calibrate', 'benchmarks', 'collect', 'deploy', and 'status-report'. The 'simulate' stage is expanded, showing a list of jobs with their status (green checkmark for success, orange exclamation mark for failure) and duration. The 'collect' stage is also expanded, showing a list of jobs with their status and duration. The 'status-report' stage shows a summary of the pipeline's overall status.

**Project:** eic\_container: eic/ElCrecon: Flag clusters as EMCal or HCal

**Pipeline:** #128763

**Jobs:** 280

**Failed Jobs:** 3

**Tests:** 0

**Stages:** eic\_container, status-pending, config, simulate, calibrate, benchmarks, collect, deploy, status-report

**Jobs in 'simulate' stage:**

- sim\_b0\_tracker
- sim\_backgrounds
- sim\_backward
- sim\_backward\_emcal
- sim\_beamline
- sim\_calo\_pid
- sim\_emcal\_gaps
- sim\_emcal\_barrel\_electrons
- sim\_emcal\_barrel\_photons
- sim\_emcal\_barrel\_pi0
- sim\_emcal\_barrel\_pion\_rejection
- sim\_emcal\_barrel\_pions
- sim\_femc\_electron
- sim\_femc\_photon
- sim\_femc\_pi0
- sim\_forward
- sim\_insert\_muon
- sim\_insert\_neutron
- sim\_insert\_tau
- sim\_ithcal
- sim\_ithcal\_acceptance
- sim\_ithcal\_basic\_distribution
- sim\_ithcal\_basic\_distribution\_full
- sim\_tracking\_performance
- sim\_tracking\_performance\_dis
- sim\_zdc
- sim\_zdc\_lambda
- sim\_zdc\_lyso
- sim\_zdc\_neutron
- sim\_zdc\_photon
- sim\_zdc\_pi0
- sim\_zdc\_sigma
- tracking\_detectors/sim\_track\_hls

**Jobs in 'collect' stage:**

- collect\_results\_backgrounds
- collect\_results\_backward\_emcal
- collect\_results\_backward\_emcal
- collect\_results\_beamline
- collect\_results\_calo\_pid
- collect\_results\_campaign
- collect\_results\_emcal\_gaps
- collect\_results\_femc\_electron
- collect\_results\_femc\_photon
- collect\_results\_femc\_pi0
- collect\_results\_insert\_muon
- collect\_results\_insert\_neutron
- collect\_results\_insert\_tau
- collect\_results\_ithcal
- collect\_results\_ithcal\_campaigns
- collect\_results\_material\_scan
- collect\_results\_ithcal\_acceptance
- collect\_results\_ithcal\_basic\_distribution
- collect\_results\_pid
- collect\_results\_rich
- collect\_results\_tracking\_performance
- collect\_results\_tracking\_performance\_campaigns
- collect\_results\_tracking\_performance\_dis
- collect\_results\_zdc
- collect\_results\_zdc\_lambda
- collect\_results\_zdc\_lyso
- collect\_results\_zdc\_neutron
- collect\_results\_zdc\_photon
- collect\_results\_zdc\_pi0
- collect\_results\_zdc\_sigma
- results\_b0\_tracker

## Problems

- Running the full CI pipeline takes a long time even if everything passes (at least an hour)
- Many of these CI jobs appear to be redundant or bizarre
  - Running sim jobs when none of the sim code has changed
  - Merging GitHub Pages documentation for all open PRs (~60)
- I've **never** had a PR pass the full pipeline without at least a day's work
  - My workflow: Remember to pull my code to pick up any automatic whitespace fixes, make my fix, push it, set a timer for 15 minutes, go work on something else, repeat
  - Each `build` job terminates its sibling `build` jobs on failure, forcing more iterations
  - Oftentimes something breaks where I have no knowledge how (or the permissions) to fix it
    - Example: Segfault in the `#pragma once` analyzer
    - Example: Missing CI tokens
    - Example: Nonsensical errors in benchmark output
  - I almost always end up asking for help from D.K. or W.D.

# Recommendations



# Recommendations

## Low-hanging fruit

- Get rid of iwyu completely
- Use sensible defaults with clang-format and don't change them
- Get rid of the overengineered documentation merging
- Any CI failures that can be fixed via a simple pragma need to advertise this fact loudly in their error logs, and link to the corresponding documentation

## Medium-hanging fruit

- Create repository for sim data rather than always regenerating (and other testing data too!)
- Create a separate “release” CI pipeline which runs clang-format, etc, getting it off of the critical path for pull requests completely
- Reconsider backwards compatibility requirement
- CI `build` jobs shouldn't kill the others when one fails

## High-hanging fruit

- Add a mechanism for running the full GitHub CI pipeline locally

**Thank you for listening!**