



U.S. DEPARTMENT
of ENERGY

AID2E News and Updates


Wen Guan, Amit Bashyal

September 9, 2026

X f i @BrookhavenLab

AID2E Introduction

- AID(2)E = AI Assisted Detector Design at the EIC
- Scalable and Distributed Framework to design the EIC detectors
- Instead of designing and optimizing each detector subcomponent separately before integration, holistic optimization of the entire detector system within global constraints
- Multi Objective optimization
- Objectives
 - Detector efficiency, resolution, costs etc



Phase I

Phase I Goals

- Framework that supports optimization of at least 2 ePIC detector subsystems
- Support ax and pymoo optimization algorithm
- Optimization problem configurable as DAG based workflow that can be executed in local, remote resources

What has been done so far

Currently happening

Closure Tests (~2024)

Tests done with optimization libraries ([ax-platform](#) botorch) on simple multi objective problems using Multi Objective Bayesian Optimization (MOBO)

[Closure Test](#): Use PanDA/iDDS to perform those tests using distributed computing resources

Detector Optimization Software (2024 - 2025)

Frameworks for various EIC sub detectors were being developed by different groups.

Development of [job scheduler](#) that supported running optimization job locally and remotely via PanDA.

Integration of [PanDA job runner for dRICH detector](#)

Integration of [PanDA job runner for BIC detector](#)

[AID2E Framework](#)

Python based optimizer interface to integrate various detector optimization software

Refactoring of detector simulation software to interface with the framework

Refactoring of job schedulers

Exploration of other optimization algorithms like [pymoo](#) within the framework

AID2E Framework

- A unified python framework to run detector-design optimization workflows for EIC/ePIC.
- Provides resources to:
 - Describe a detector design problem (optimization problem, detector parameters and constraints))
 - Interface with optimization software like ax, pymoo
 - Interface with Scheduler to schedule optimization job (PanDA, Slurm, Local)
- Framework does:
 - Generate candidate design points
 - Execute evaluation based on user specified optimization
 - Collect objective metrics
 - Feed results back into the optimizer for next iteration

AID2E Framework workflow

```
(.venv) (base) amitbashyal@lpo-178574:~/Documents/BNL-AID2E/ncp-server/AID2E-framework$ aid2e
Usage: aid2e [OPTIONS] COMMAND [ARGS]...

AID2E - AI assisted Detector Design for EIC.

A framework for optimization of detector designs and other complex systems.

Commands are organized into categories:

Configuration Inspection:
  describe - Quick summary of config files (auto-detects type)
  inspect  - Detailed configuration view with section filtering
  validate - Validate configuration syntax and structure


Workflow Execution:
  optimize - Run optimization from configuration

Utilities:
  list     - Show available optimizers/templates/problems
  version  - Display version information

Options:
  --version Show the version and exit.
  --help    Show this message and exit.

Commands:
```

- CLI based user interface
- Layout the optimization strategy in a yaml file
- ***aid2e optimize yamlfile.yaml*** → Entry point



An Optimization in AID2E Framework

aid2e optimize
yamfile.yaml is the entry
point

Problem is the
optimization problem
with design
parameters and
objectives

```
AID2E-framework > ! input-panda-multistage.yaml
1 # DTLZ2 + Ax + PanDA/iDDS stage-native dataset-dependency template.
2 #
3 # This is the issue-61 target shape: sim/reco, analysis, and final objective
4 # calculation are separate AID2E stages. Stage dependencies are declared in job
5 # payload metadata so the PanDA/iDDS scheduler can translate them into
6 # parent_internal_id, output dataset, and input dataset contracts. The dataset
7 # dependency shape follows scheduler_epic's tested pattern: sim/reco and ana use
8 # matching one2one child keys, and the local final stage performs all2one result
9 # fan-in.
10 #
11 # Intended command:
12 #
13 # aid2e optimize input-panda-multistage.yaml
14 #
15 # Note: PanDAiDDSScheduler now detects optimizer trial wrapper jobs for this
16 # shape and packages these stage records into a scheduler-local
17 # panda_multistage_graph payload. DAGExecutor remains scheduler-agnostic.
18
19 problem:
20   name: "DTLZ2 CLI PanDA Multi-Stage Dataset Smoke Test"
21   problem_type: "DTLZ2"
22   output_location: "/tmp"
23   work_location: "/tmp"
24   design_parameters_file: "examples/optimizers/dtlz2_design.params"
25   objectives: --
26
27 optimizer:
28   name: "ax"
29   type: "bayesian"
30   parameters: --
31
32 scheduler:
33   runner_type: "PanDAiDDSRRunner"
34   parameters: --
35   max_retries: 1
36   output_location: "./aid2e_panda_scheduler_output"
37   monitor_interval: 5
38
39 workflows:
40   workflows:
41     - name: "dtlz2_panda_multistage_eval"
42       description: "Stage-native PanDA/iDDS template: simreco datasets feed analysis through PanDA-managed dependencies."
43       branches:
44         - name: "main"
45           stages:
46             - name: "simreco"
47               parallelism:
48                 max_concurrent: 4
49                 retry_max: 0
50                 timeout_sec: 600
51             jobs:
52               - name: "simreco_eta_0p1_pi" --
53               - name: "simreco_eta_0p1_kaon" --
54               - name: "simreco_eta_0p2_pi" --
55               - name: "simreco_eta_0p2_kaon" --
56             - name: "ana" --
57             - name: "final" --
```

An Optimization in AID2E Framework

aid2e optimize
yamfile.yaml is the entry
point

Optimizer layer to
setup optimization
algorithm.

```
AID2E-framework > ! input-panda-multistage.yml
1 # DTLZ2 + Ax + PanDA/iDDS stage-native dataset-dependency template.
2 #
3 # This is the issue-61 target shape: sim/reco, analysis, and final objective
4 # calculation are separate AID2E stages. Stage dependencies are declared in job
5 # payload metadata so the PanDA/iDDS scheduler can translate them into
6 # parent_internal_id, output dataset, and input dataset contracts. The dataset
7 # dependency shape follows scheduler_epic's tested pattern: sim/reco and ana use
8 # matching one2one child keys, and the local final stage performs all2one result
9 # fan-in.
10 #
11 # Intended command:
12 #
13 # aid2e optimize input-panda-multistage.yml
14 #
15 # Note: PanDAiDDSScheduler now detects optimizer trial wrapper jobs for this
16 # shape and packages these stage records into a scheduler-local
17 # panda_multistage_graph payload. DAGExecutor remains scheduler-agnostic.
18
19 problem:
20   name: "DTLZ2 CLI PanDA Multi-Stage Dataset Smoke Test"
21   problem_type: "DTLZ2"
22   output_location: "/tmp"
23   work_location: "/tmp"
24   design_parameters_file: "examples/optimizers/dtlz2_design.params"
25 > objectives:--
30
31 optimizer:
32   name: "ax"
33   type: "bayesian"
34 > parameters:--
41
42 scheduler:
43   runner_type: "PanDAiDDSRRunner"
44 > parameters:--
60   max_retries: 1
61   output_location: "./aid2e_panda_scheduler_output"
62   monitor_interval: 5
63
64 workflows:
65   workflows:
66     - name: "dtlz2_panda_multistage_eval"
67       description: "Stage-native PanDA/iDDS template: simreco datasets feed analysis through PanDA-managed dependencies."
68       branches:
69         - name: "main"
70           stages:
71             - name: "simreco"
72               parallelism:
73                 max_concurrent: 4
74                 retry_max: 0
75                 timeout_sec: 600
76             jobs:
77 >               - name: "simreco_eta_0p1_pi"--
92
93 >               - name: "simreco_eta_0p1_kaon"--
108
109 >               - name: "simreco_eta_0p2_pi"--
124
125 >               - name: "simreco_eta_0p2_kaon"--
140
141 >             - name: "ana"--
214
215 >             - name: "final"--
237
```

An Optimization in AID2E Framework

aid2e optimize
yamlfile.yaml is the entry
point

User defines the
workflow layer.

```
AID2E-framework > ! input-panda-multistage.yaml
1 # DTLZ2 + Ax + PanDA/iDDS stage-native dataset-dependency template.
2 #
3 # This is the issue-61 target shape: sim/reco, analysis, and final objective
4 # calculation are separate AID2E stages. Stage dependencies are declared in job
5 # payload metadata so the PanDA/iDDS scheduler can translate them into
6 # parent_internal_id, output dataset, and input dataset contracts. The dataset
7 # dependency shape follows scheduler_epic's tested pattern: sim/reco and ana use
8 # matching one2one child keys, and the local final stage performs all2one result
9 # fan-in.
10 #
11 # Intended command:
12 #
13 # aid2e optimize input-panda-multistage.yaml
14 #
15 # Note: PanDAiDDSScheduler now detects optimizer trial wrapper jobs for this
16 # shape and packages these stage records into a scheduler-local
17 # panda_multistage_graph payload. DAGExecutor remains scheduler-agnostic.
18 #
19 problem:
20   name: "DTLZ2 CLI PanDA Multi-Stage Dataset Smoke Test"
21   problem_type: "DTLZ2"
22   output_location: "/tmp"
23   work_location: "/tmp"
24   design_parameters_file: "examples/optimizers/dtlz2_design.params"
25   objectives: --
30
31 optimizer:
32   name: "ax"
33   type: "bayesian"
34   parameters: --
41
42 scheduler:
43   runner_type: "PanDAiDDSRunner"
44   parameters: --
60   max_retries: 1
61   output_location: "./aid2e_panda_scheduler_output"
62   monitor_interval: 5
63
64 workflows:
65   workflows:
66     - name: "dtlz2_panda_multistage_eval"
67       description: "Stage-native PanDA/iDDS template: simreco datasets feed analysis through PanDA-managed dependencies."
68       branches:
69         - name: "main"
70           stages:
71             - name: "simreco"
72               parallelism:
73                 max_concurrent: 4
74                 retry_max: 0
75                 timeout_sec: 600
76             jobs:
77               - name: "simreco_eta_0p1_pi" --
92
93               - name: "simreco_eta_0p1_kaon" --
108
109               - name: "simreco_eta_0p2_pi" --
124
125               - name: "simreco_eta_0p2_kaon" --
146
147               - name: "ana" --
214
215               - name: "final" --
234
```

An Optimization in AID2E Framework

aid2e optimize
yamfile.yaml is the entry
point



Execution of the
workflow via
Scheduler Interface

```
AID2E-framework > ! input-panda-multistage.yml
1 # DTLZ2 + Ax + PanDA/iDDS stage-native dataset-dependency template.
2 #
3 # This is the issue-61 target shape: sim/reco, analysis, and final objective
4 # calculation are separate AID2E stages. Stage dependencies are declared in job
5 # payload metadata so the PanDA/iDDS scheduler can translate them into
6 # parent_internal_id, output dataset, and input dataset contracts. The dataset
7 # dependency shape follows scheduler_epic's tested pattern: sim/reco and ana use
8 # matching one2one child keys, and the local final stage performs all2one result
9 # fan-in.
10 #
11 # Intended command:
12 #
13 # aid2e optimize input-panda-multistage.yml
14 #
15 # Note: PanDAiDDSScheduler now detects optimizer trial wrapper jobs for this
16 # shape and packages these stage records into a scheduler-local
17 # panda_multistage_graph payload. DAGExecutor remains scheduler-agnostic.
18
19 problem:
20   name: "DTLZ2 CLI PanDA Multi-Stage Dataset Smoke Test"
21   problem_type: "DTLZ2"
22   output_location: "/tmp"
23   work_location: "/tmp"
24   design_parameters_file: "examples/optimizers/dtlz2_design.params"
25   objectives:--
26
27
28
29
30
31 optimizer:
32   name: "ax"
33   type: "bayesian"
34   parameters:--
35
36
37
38
39
40 scheduler:
41   runner_type: "PanDAiDDSRunner"
42   parameters:--
43   max_retries: 1
44   output_location: "./aid2e_panda_scheduler_output"
45   monitor_interval: 5
46
47
48 workflows:
49   workflows:
50     - name: "dtlz2_panda_multistage_eval"
51       description: "Stage-native PanDA/iDDS template: simreco datasets feed analysis through PanDA-managed dependencies."
52       branches:
53         - name: "main"
54           stages:
55             - name: "simreco"
56               parallelism:
57                 max_concurrent: 4
58                 retry_max: 0
59                 timeout_sec: 600
60             jobs:
61               - name: "simreco_eta_0p1_pi"--
62               - name: "simreco_eta_0p1_kaon"--
63               - name: "simreco_eta_0p2_pi"--
64               - name: "simreco_eta_0p2_kaon"--
65             - name: "ana"--
66             - name: "final"--
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
```

Job Scheduler for AID2E optimization jobs

- Interface to manage AID2E optimization trial jobs
 - Based on previous work on [scheduler_epic](#)
 - Core AID2E-framework by default provides support for the local runner
 - Support to PanDA and slurm runners are abstracted as plugins to the scheduler interface
- PanDA Runner that was developed and tested with dRICH-MOBO and BIC-MOBO optimization is refactored to provide AID2E Framework support

PanDA job Scheduling in AID2E Framework



- AID2E provides support for the PanDA/iDDS job scheduler to use the distributed compute resources
- AID2E can configure PanDA/iDDS to submit workflow as **one2one** or **many2one** through a PanDA specific multi graph interface. See implementation details [here](#)

Scheduling AID2E jobs via PanDA

39180 39175	3579	user.abashyal2.a2e.376376ecd784.000005.0a77ca93_3579_32311 iDDS iDDS Errors None	done 2	2 100%	
39179 39176	3579	user.abashyal2.a2e.376376ecd784.000006.41f2710b_3579_32312 iDDS iDDS Errors None	done 2	2 100%	
39176	3579	user.abashyal2.a2e.376376ecd784.000002.a921f9b1_3579_32308 iDDS iDDS Errors None	done 2	2 100%	200 / 0
39175	3579	user.abashyal2.a2e.376376ecd784.000001.12a2bfb6_3579_32307 iDDS iDDS Errors None	done 2	2 100%	200 / 0
		user.abashyal2.a2e.0203c9c142e3.000004.8f3eacaf_3578_32305			



Screenshot of AID2E completed jobs from pandamon

Ana jobs (top) are the child of sim/reco jobs (bottom).

Running dRICH Optimization with aid2e framework



Scheduler: workflow job distribution

In dRICH example:

4 regions, 2 particles

```
particles: ["pi+", "kaon+"]
analysis_executable: "script/dRICHAna_bootstrap"
output_name: "drich_optimization_toy"
scan_points:
- {p: 15, eta_min: 1.5, eta_max: 2.0, radiator: 0}
- {p: 15, eta_min: 2.0, eta_max: 2.5, radiator: 0}
- {p: 45, eta_min: 2.5, eta_max: 3.0, radiator: 1}
- {p: 45, eta_min: 3.0, eta_max: 3.5, radiator: 1}
```

Different scheduler can be configured in each workflow stage



Taken from Cynthia Nunez's update

C. Nunez showed that AID2E framework now supports dRICH optimization using joblib and slurm runner.

Ongoing tests with PanDA runner

AID2E Version 0.0 Goal

- Support the optimization of **BIC** and **dRICH detectors** of the ePIC using AID2E framework
- Derek from JLab is finalizing the BIC example based on C. Nune's work
- Paper on the AID2E design and performance ongoing

Road to Version 0.0

The screenshot displays the 'Version 0.0.0 Tracker' application interface. At the top, there's a navigation bar with tabs for 'Tracker', 'Priority board', 'Team items', 'Roadmap', and 'My items'. A search bar is located below the navigation bar. The main area is a Kanban board with five columns: 'To-Do' (1/5 items), 'In progress' (1 item), 'In review' (3 items), and 'Done' (11 items). Each column has a description of its state and a list of tasks. The 'In review' column is highlighted, showing tasks like 'Add BIC example' and 'Add CLI PanDA optimization support'. The 'Done' column shows completed tasks like 'CLI Optimize Needs to b' and 'Refactor the job hierarch'.

Version 0.0.0 Tracker

Navigation: Tracker | Priority board | Team items | Roadmap | My items | + New view

Filter by keyword or by field

To-Do 1 / 5 Estimate: 0
This item hasn't been started

- AID2E-framework #66
Add Option to Load Previous Ax Experiment and Save Output

In progress 1 Estimate: 0
This is actively being worked on

- AID2E-framework #62
Review doc and examples

In review 3 Estimate: 0
This item is in review

- AID2E-framework #27
Add BIC example
#28
1 / 2 50%
- AID2E-framework #52
epic_imaging_optimization.yml should run out-of-the-box
#53
- AID2E-framework #65
Add CLI PanDA optimization support

Done 11 Estimate: 0
This has been completed

- AID2E-framework #56
CLI Optimize Needs to b
0 / 1
- AID2E-framework #61
Refactor the job hierarch such that steps are repla
#65
- AID2E-framework #8
Implement epic_problen
#57
- AID2E-framework #36
Ergonomic Wiring Betwe
Layers
#55
- AID2E-framework #1
Implement epic workflo
#53
- AID2E-framework #44
Slurm Scheduler Scope
#57
- AID2E-framework #6
Integrate scheduler (Gitl
#22
2 / 2

Finishing up the review of BIC example and PanDA integration → Version 0.0 release.

Phase II

Project II for AID2E

- Holistic Optimization
 - Optimization of the whole epic detector as a single design point
- AID2E Framework → AID2E Agentic Framework
 - [Agent for Code Assistance](#) (Williams & Mary)
 - [Agent for Job Control/Monitoring](#) (BNL)

Agent for Code Assistance

- MCP tools that will help the user in planning, understanding, developing and using the AID2E code
- Can be Integrated within the VS Code
 - Copilot uses the AID2E Code Assistance MCP tools
- MCP tools are configured to explore the framework in structured way on demand
- Uses [Jinja](#) to leverage open models via vLLM through VS-Code
 1. VS Code sends Open AI compatible HTTP request
 2. vLLM parses the message
 3. Jinja renders messages based on templates (reduces token usage by constraining to AID2E specific context only)
 4. Model specific prompt is produced
 5. LLM reads the Jinja rendered prompt


Agent to Steer/Monitor Jobs

- Ongoing work
- Agent is currently in the Exploratory and Design Phase
- Fleshed out core requirements of the Agent
- Putting together the tools and technologies required to perform tasks


AID2E Job Steering/Monitoring Agent

- Agent to monitor and steer AID2E optimization jobs
- Currently being developed with PanDA job scheduler
 - Eventually expand to all supported job schedulers
- Humans provide configuration template
- Agent does:
 - Configuration yaml file sanity check based on past error, pre-defined criterions
 - Launch AID2E jobs using PanDA
 - Monitor PanDA jobs
 - Store completed PanDA job status in the memory for future reference (memory)
 - Failure/success
 - (Possible)Error identification in codebase
 - Job summary report
 - Optimization result
 - Summary of optimization (graphical, reports)
 - By retrieving past and present information from the memory

AID2E Job Steering/Monitoring Agent




AID2E Job Steering/Monitoring Agent



- MCP tools/server to access job information via PanDA server
- MCP server to interact with AID2E software stack
- Any other tools needed in future


AID2E Job Steering/Monitoring Agent

- Agent writes and reads information related to the optimization job (logs, reports, summary) to build the context
- OpenViking server deployed at aipanda@cern machine




AID2E Job Steering/Monitoring Agent

- Readonly database based on AID2E codebase
- iDDS codebase
- Leverage W&M work




AID2E Job Steering/Monitoring Agent



- Frontier models (Claude, OpenAI models)
- Open Models (Qwen, Llama etc) that run on remotely running vllm engine
- Tested with vllm engines running as Perlmutter Slurm Job via Spin Server
 - See more details in Tianle's update ([here](#))
- Planning to test with models provided by SDCC (setup by Meifen, Atif, Tianle et al)

Framework configured to work with local models



- An agent orchestration layer that decides and delegates MCP server/tool usage.
- Possible ways to optimize and bring down LLM token usage if necessary.

```

You: Give me the summary of the job 958076. And also give me the token usage to process and respond to this question.
Agent: Using tool 'panda_study_job'...
Agent: pandaId: 958076
monitor_url: https://pandamon01.sdcc.bnl.gov/job?pandaId=958076

files: 3
| lfn | type | guid | scope | fsize | status | dataset | destinationdblock |
| --- | --- | --- | --- | --- | --- | --- | --- |
| group.EIC.26.06.0.epic_craterlake.SINGLE.pi0.100MeV.3to50deg.37258.001603.log.tgz | log | 1a0a1aac-592b-432e-b762-5cd03e21bce8 | group.EIC | 446398 | ready | group.EIC.26.06.0.epic_craterlake.SINGLE.pi0.100MeV.3to50deg_log.108585.108585 | group.EIC.26.06.0.epic_craterlake.SINGLE.pi0.100MeV.3to50deg_log.108585.108585 |
| 1603 | pseudo_input | a352346b-a6bc-45e0-bd92-2c4e39dd4f18 | | 0 | unknown | seq_number | |
| pseudo_lfn | pseudo_input | 9100a04d-b15d-4fdc-9934-7f99b4805aa5 | | 0 | unknown | pseudo_dataset | |

Token usage:
- Prompt tokens: 0
- Response tokens: 218
- Total tokens: 218
- Source: estimated
- Note: no LLM prompt was sent; the answer was produced by deterministic tool routing/formatting.

```

```

You: Give me the status of the job 958076. And also give me the token usage to process and respond to this question.
Agent: Using tool 'panda_study_job'...
Agent: pandaId: 958076
monitor_url: https://pandamon01.sdcc.bnl.gov/job?pandaId=958076

files: 3
| lfn | type | guid | scope | fsize | status | dataset | destinationdblock |
| --- | --- | --- | --- | --- | --- | --- | --- |
| group.EIC.26.06.0.epic_craterlake.SINGLE.pi0.100MeV.3to50deg.37258.001603.log.tgz | log | 1a0a1aac-592b-432e-b762-5cd03e21bce8 | group.EIC | 446398 | ready | group.EIC.26.06.0.epic_craterlake.SINGLE.pi0.100MeV.3to50deg_log.108585.108585 | group.EIC.26.06.0.epic_craterlake.SINGLE.pi0.100MeV.3to50deg_log.108585.108585 |
| 1603 | pseudo_input | a352346b-a6bc-45e0-bd92-2c4e39dd4f18 | | 0 | unknown | seq_number | |
| pseudo_lfn | pseudo_input | 9100a04d-b15d-4fdc-9934-7f99b4805aa5 | | 0 | unknown | pseudo_dataset | |

Token usage:
- Prompt tokens: 1224
- Response tokens: 906
- Total tokens: 2130
- Source: provider

You: 

```

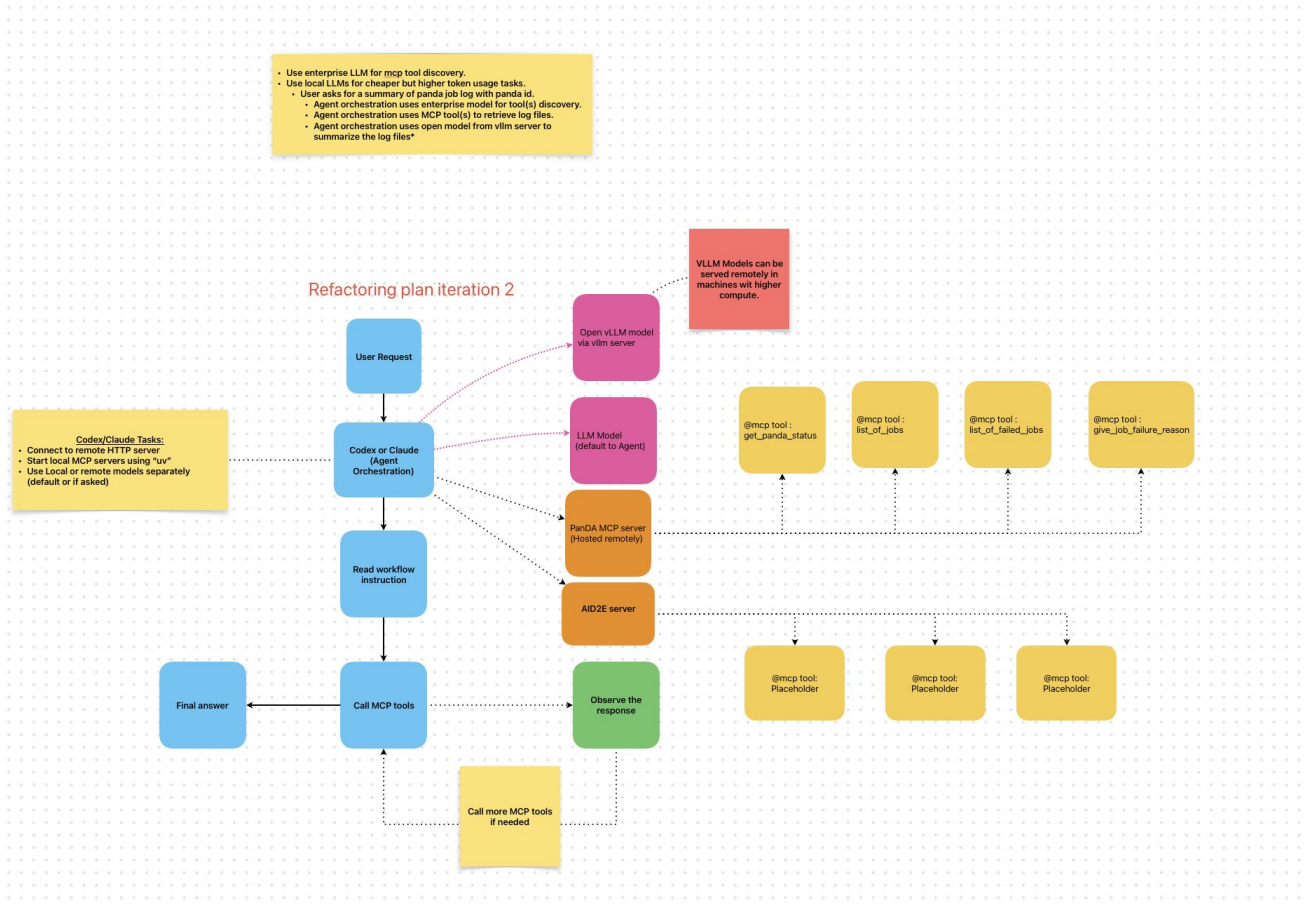
LLM model : GPT-OSS (20 B parameters) only for testing tool discovery

Token usage: Given by the model

Using strong prompting from the .md files for tool discovery

Plugging MCP servers into Codex (or Claude)

- Use Claude or Codex or other enterprise layer for orchestration
- Rely completely on native model (OpenAI models if Codex for example)
- OR Use a combination of native and open models (WIP)



› Give me the summary of the job 958076

△ Falling back from WebSockets to HTTPS transport. stream disconnected before completion: invalid peer certificate: UnknownIssuer

• I'll look up the available production/job tools and pull the record for job 958076, then summarize the useful fields rather than dumping raw output.

• Called `swf-testbed.panda_study_job({"pandaaid":958076})`

```
L {"pandaaid": 958076, "job": {"pandaaid": 958076, "jeditaskid": 37258, "reqid": 37258, "jobname":  
  "group.EIC.26.06.0.epic_craterlake.SINGLE.pi0.100MeV.3to50deg.958076", "produsername": "Sakib Rahman", "jobstatus": "finished",  
  "computingsite": "NERSC_Perlmutter_epic_test", "transformation": "https://pandaserver-doma.cern.ch/trf/user/runGen-00-00-02",  
  "processingtype": "epicproduction", "creationtime": "2026-06-24T05:58:34.724258", "starttime": "2026-06-25T15:31:34", "endtime":  
  "2026-06-25T18:05:58", "modificationtime": "2026-06-25T18:06:04.193183", "corecount": 1, "actualcorecount": 2, "nevents": 0, "maxrss":  
  2934432, "maxpss": 2551405, "maxvmem": 7019288, "maxswap": 0, "maxwalltime": 34...
```

• Job 958076 finished successfully.

Key details:

- Task: 37258
- Task name: group.EIC.26.06.0.epic_craterlake.SINGLE.pi0.100MeV.3to50deg
- Owner: Sakib Rahman
- Status: finished
- Site: NERSC_Perlmutter_epic_test
- Started: 2026-06-25 15:31:34
- Ended: 2026-06-25 18:05:58
- Runtime: about 2h 34m
- Transform exit code: 0
- Errors: none recorded across pilot, executor, DDM, dispatcher, supervisor, brokerage, or taskbuffer

What happened	Input tokens	Cached input	Output tokens	Reasoning output	Total tokens
Tool search	14,505	4,480	94	22	14,599
panda_study_job	18,095	2,432	29	0	18,124
Final answer	20,925	17,792	303	0	21,228
Full turn total	53,525	24,704	426	22	53,951

Open Viking for the Agent Context

- Wen setup OpenViking Server in aipanda machine @ CERN
- Will serve as the memory for the Agent to read and write information
- Created MCP server that provides retrieve/write tool for the OpenViking Server
- Open Viking MCP server is spawned locally at the moment

Demo : PanDA MCP Tools + SWF Agents + OpenViking

Prompt : Give me the summary of all jobs submitted by iDDS in the last 7 days

Task context

- 30 JEDI tasks, all in status done
- All ran at BNL_PanDA_1
- All used transformation `run_workflow_wrapper`
- `workinggroup` and `processingtype` are null/None for these iDDS automation jobs
- The 3 failed jobs were retried or otherwise absorbed: task-level final failure rate is 0%

By request ID

- reqid 3578: 8 jobs, 8 finished
- reqid 3579: 14 jobs, 12 finished, 2 failed
- reqid 3581: 12 jobs, 12 finished
- reqid 3582: 13 jobs, 12 finished, 1 failed
- reqid 3583: 2 jobs, 2 finished

Failures

The 3 failed jobs were in tasks 39176, 39178, and 39197, with representative PandaIDs:

- 2716851
- 2716858
- 2717847

The common reported errors were:

- DDM error 200: could not add files to DDM due attachment validation involving NULL
- Pilot error 1305: payload execution failure, beginning with WARNING: group: unknown groupid 69997

All activity in this 7-day window occurred between 2026-09-03 23:52 and 2026-09-04 19:19 by the returned PanDA timestamps.

Uses the SWF PanDA tools

Demo : PanDA MCP Tools + SWF Agents + OpenViking

Prompt : Write the summary of this report to the openviking server

```
• Called openviking.openviking_add_resource({"path": "/home/amtibashyal/Documents/BNL-AID2E/mcp-server/AID2E-Agentic-Framework/reports/idds_panda_jobs_last_7_days_2026-09-08.md", "wait": false})
  ↳ {"ok": true, "operation": "add_resource", "arguments": {"path": "/home/amtibashyal/Documents/BNL-AID2E/mcp-server/AID2E-Agentic-Framework/reports/idds_panda_jobs_last_7_days_2026-09-08.md", "to": null, "parent": null, "wait": false, "timeout": null}, "data": {"status": "success", "task_id": "b468718b-4b02-4a64-9242-42ae563afa65", "root_uri": "viking://resources/idds_panda_jobs_last_7_days_2026-09-08_2"}}

• Wrote the summary to OpenViking.

Resource URI: viking://resources/idds_panda_jobs_last_7_days_2026-09-08_2
Ingestion task ID: b468718b-4b02-4a64-9242-42ae563afa65

I also saved the source report locally at reports/idds_panda_jobs_last_7_days_2026-09-08.md.
```

Uses the OpenViking MCP server tools to write the summary.

Follows pre-written instruction to make a copy on local machine.

Demo : PanDA MCP Tools + SWF Agents + OpenViking

Prompt : Use today's and past report to give me error summary

• Using the latest 49-job report as the authoritative total, and the past reports to confirm the earlier 3578 and 3579 failure details:

Category	Jobs affected	% of total submitted jobs
No failure / finished successfully	46	93.88%
Any failed job	3	6.12%
DDM attachment validation failure, <code>ddmerrorcode=200</code>	3	6.12%
Pilot payload execution failure, <code>piloterrorcode=1305</code>	3	6.12%

Important detail: the DDM and Pilot categories are **not independent job sets**. The same 3 failed jobs carried both error types:

- 2716851, task 39176
- 2716858, task 39178
- 2717847, task 39197

So the mutually exclusive operational categorization is:

Failure signature	Jobs	% of total submitted jobs
<code>ddmerrorcode=200 + piloterrorcode=1305</code>	3	6.12%
No failure	46	93.88%

Past report alignment:

- reqid 3578: 8 jobs, 0 failures
- reqid 3579: 2 failed jobs with the same DDM + Pilot pattern
- Latest report adds 1 more failed job under reqid 3582, also same pattern.

Uses the OpenViking MCP server tools to read the summaries from the past.

Follows pre-written instruction to present reports as table.

Roadmap

- Finalize the Open Model Agent Orchestration layers
 - Stable Perlmutter Spin Server hosted vllm gateways
 - Models provided by SDCC
- Organization of the Agent Memory in the OpenViking server
 - Rules to categorize and organize written reports
 - Campaign dates, AID2E optimization problem, optimizer, trial summary, PanDA job summary (status, compute time, memory required, resource used, time spent in idle etc)
 - Rules on persistence, volatile records of optimization campaign
- MCP tool for AID2E framework (based on W&M work)
- Deploy currently locally hosted MCP server(s) to the NPPS machine(s)
- Comparison of performance, token usage with available models
- Concretize list of tasks/skills/tools for the AID2E Agent
- Target : First stable release of AID2E Job Steering/Monitoring Agent for production style optimization campaign