



Ask PanDA

Paul Nilsson

February 19, 2025

NPPS Weekly Meeting



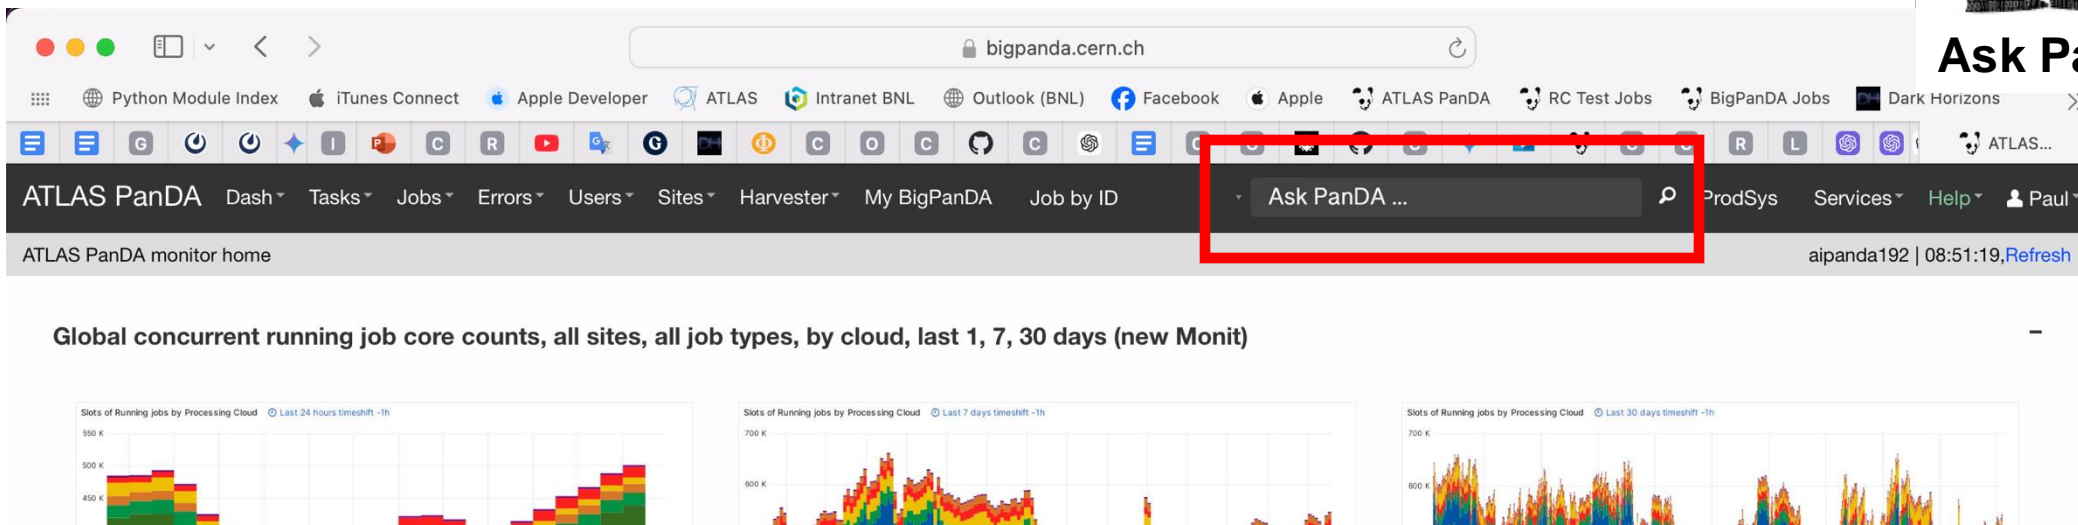
@BrookhavenLab

Introduction

- "Ask PanDA" (prel. name) is a prototype virtual assistant using OpenAI API
 - Originally intended to improve how we access project documentation but might grow into something more interesting (and useful)
 - Can answer both **static** (predefined) and **dynamic** (real-time) questions
 - "What is PanDA?" (static)
 - "How does Harvester work?" (static)
 - "Tell me about brokering in PanDA" (static)
 - "How many pilots ran yesterday at BNL?" (dynamic)
 - "Why did job XYZ fail?" (dynamic)
- Why is that interesting?
 - Reduces time spent searching for answers
 - For both users as well as the people that have to spend time looking for answers and respond
 - Operators monitoring the system trying to understand it (esp. when things go wrong)

Initial Ideas – Static Queries

- Use AI to answer questions about the PanDA documentation
 - No AI is better than the information it is trained on
 - Trick is to use well written, relevant and up-to-date documentation along with a proper prompt template (only answer relevant questions)
- Provide access via standard entry point, i.e. the PanDA monitor
 - Re-use/extend search field or add new “Ask PanDA” button somewhere



The screenshot shows the ATLAS PanDA monitor interface. The browser address bar displays `bigpanda.cern.ch`. The navigation bar includes links for `ATLAS PanDA`, `Dash`, `Tasks`, `Jobs`, `Errors`, `Users`, `Sites`, `Harvester`, `My BigPanDA`, `Job by ID`, `Ask PanDA ...` (highlighted with a red box), `ProdSys`, `Services`, `Help`, and `Paul`. Below the navigation bar, the page title is `ATLAS PanDA monitor home`. The main content area displays `Global concurrent running job core counts, all sites, all job types, by cloud, last 1, 7, 30 days (new Monit)`. Three charts are visible, each showing `Slots of Running jobs by Processing Cloud` for different time periods: `Last 24 hours timestep -1h`, `Last 7 days timestep -1h`, and `Last 30 days timestep -1h`. The charts use a color-coded bar chart format to represent job core counts across different processing clouds.

RAG Approach for Static Queries

- **Retrieval-Augmented Generation (RAG)** is a better approach for our purposes than training a new model
 - Since PanDA documentation is essentially the only source of knowledge needed for our chatbot (and thus has too little information to train an LLM from scratch)
- RAG stores the input document in an **embedding-based database** (Langchain FAISS)
 - When a user asks a question, retrieve the most relevant sections and pass them to an existing LLM (GPT)
 - RAG does not extend the LLM model but augments the answers
 - Works well for long documents since it avoids input size limitations
- Currently using GPT-4o-mini model (\$0.15 per 1M tokens input, \$0.60 per 1M output)

Footnote: Would be interesting to try GPT-4o

- Very fast and has a 128k context window which allows to process large portions of the document in one go
- Alas, it's expensive (\$2.5 per 1M tokens input, \$10 per 1M output)

Current Implementation

- PanDA documentation (readthedocs) was exported as pdf, then converted to plain text
 - Realized later that readthedocs does not export properly – used GitHub source files instead
- A python script using OpenAI and LangChain API modules
 - For testing purposes flask is currently used for creating simple web interface – but seems to be a bottleneck
- A “proper” prompt template is used

```
panda_expert_template = """
You are Dr PanDA, the expert who knows almost everything about the PanDA Project.
Your expertise is exclusively in providing information and advice about anything related to the PanDA Project.
This includes any general PanDA Project-related queries.
If the user asks about panda, it is understood that they mean PanDA and not the animal.
If a question is not about the PanDA project, respond with, "I can't assist you with that, sorry!", unless it is about CERN, BNL.
Context: {context}
Question: {question}
Answer:
"""
```

- PanDA docs are over 600 kB - can only send 15kB to OpenAI
 - Script therefore divides the text in chunks and once it gets user input, it tries to find the most likely chunk (using OpenAI embedding model via Langchain) before sending it off to the GPT model
- Tried different ChatGPT models
 - GPT-3.5 (fast but not so good answers, sometimes cut in the middle of a sentence – but is very cheap)
 - GPT-4 (much better, more expensive of course)
 - Switched to GPT-4o-mini this week

Static

Example 1

- How to run a hello world job with prun
 - PanDA client tool used to submit user jobs to the grid

Ask PanDA Chatbot

user: hi there

bot: Hello! How can I assist you with the PanDA Project today?

user: show me hot to run a hello world job with prun

bot: To run a "Hello World" job with prun, you can follow the example below:

```
```bash
$ prun --exec "pwd; ls; echo Hello-world > myout.txt" --outDS user.hoge.`uuidgen` --nJobs 3 --output myo
```
```

In this command:

- `--exec`` takes the execution string which is executed on remote computing resources. In this case, it's printing the working directory (``pwd``), listing the files (``ls``), and echoing "Hello-world" into a file named ``myout.txt``.
- `--outDS`` takes the basename of the output collection. Here, it's generating a unique identifier using ``uuidgen`` and prefixing it with ``user.hoge.``.
- `--nJobs 3`` specifies that 3 jobs should be run.
- `--output myo`` specifies the output file name.

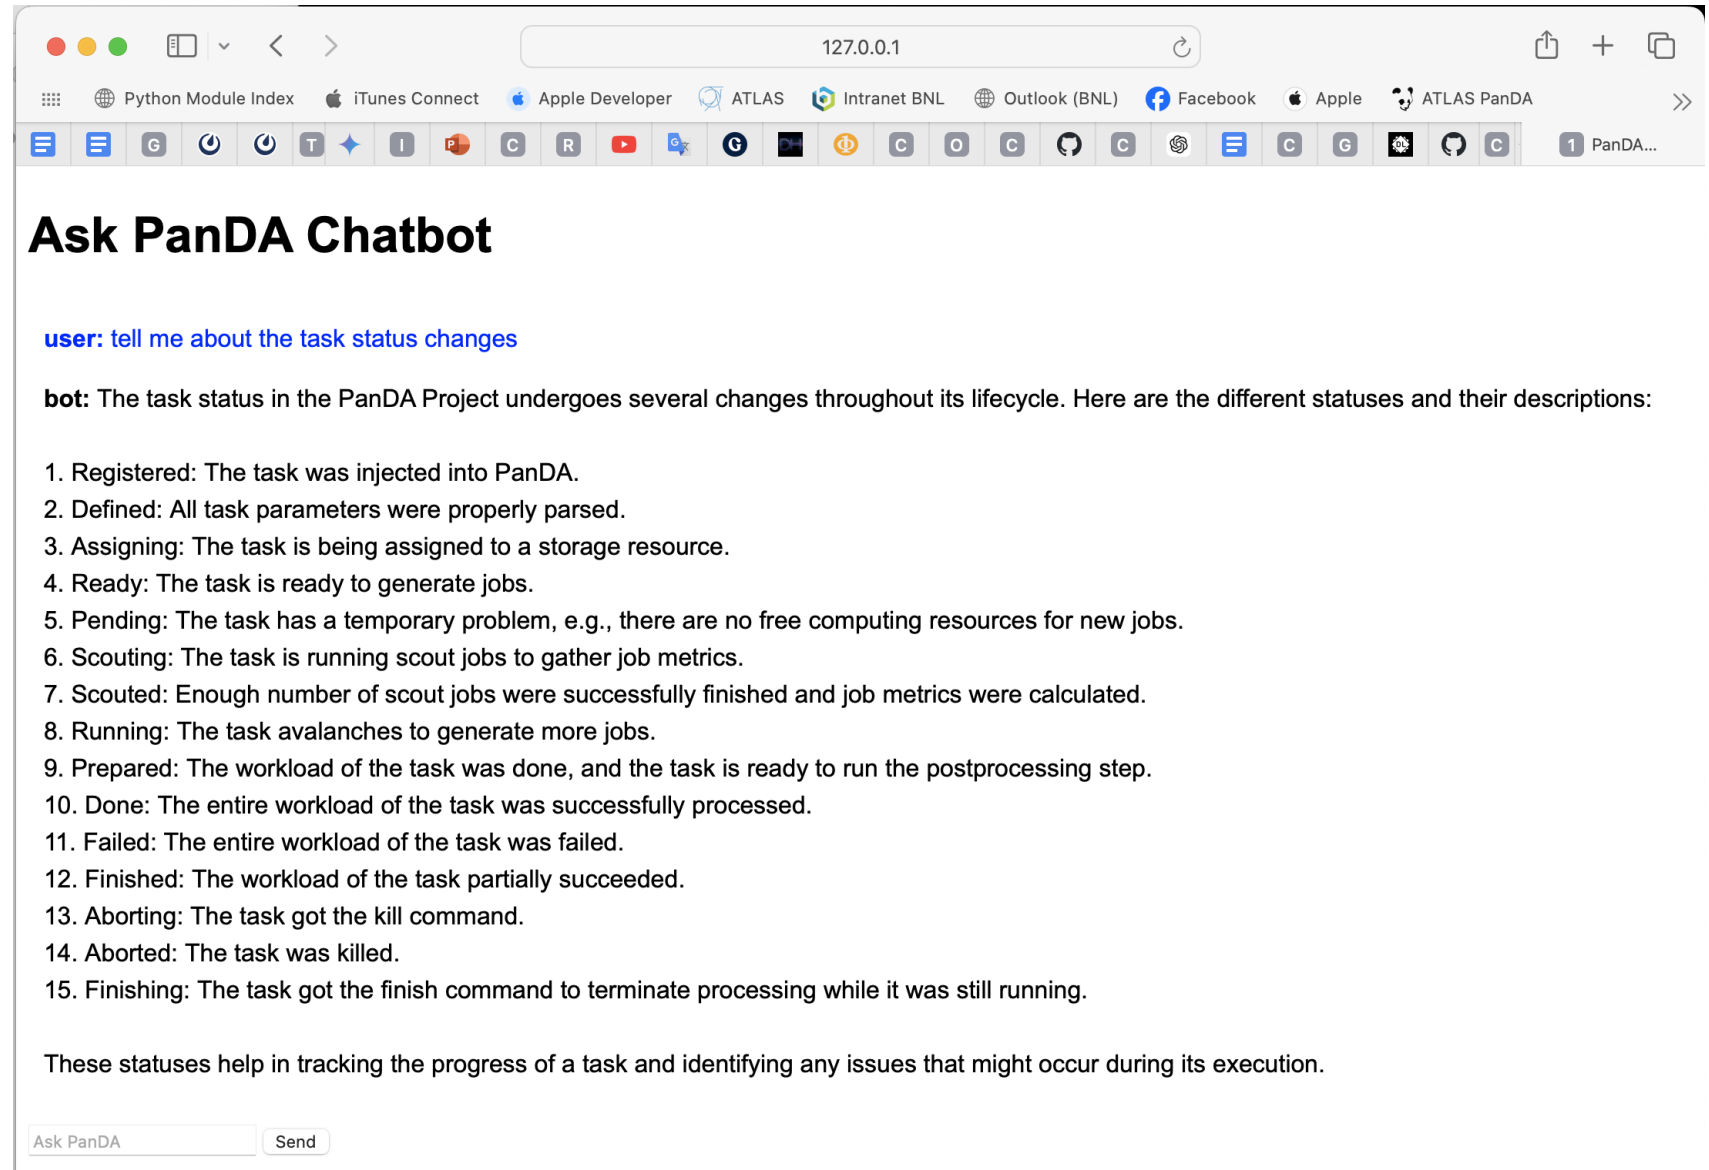
Please ensure that you have the necessary permissions and resources to run this command.

Ask PanDA

Static

Example 2

- Task status changes are explained
- Formatting and list numbering done by ChatGPT
- Note: if you ask ChatGPT (web version) the same question, you will probably get the same answer in this case
 - Because it's on the internet!
 - It might also be outdated info



Ask PanDA Chatbot

user: [tell me about the task status changes](#)

bot: The task status in the PanDA Project undergoes several changes throughout its lifecycle. Here are the different statuses and their descriptions:

1. Registered: The task was injected into PanDA.
2. Defined: All task parameters were properly parsed.
3. Assigning: The task is being assigned to a storage resource.
4. Ready: The task is ready to generate jobs.
5. Pending: The task has a temporary problem, e.g., there are no free computing resources for new jobs.
6. Scouting: The task is running scout jobs to gather job metrics.
7. Scouted: Enough number of scout jobs were successfully finished and job metrics were calculated.
8. Running: The task avalanches to generate more jobs.
9. Prepared: The workload of the task was done, and the task is ready to run the postprocessing step.
10. Done: The entire workload of the task was successfully processed.
11. Failed: The entire workload of the task was failed.
12. Finished: The workload of the task partially succeeded.
13. Aborting: The task got the kill command.
14. Aborted: The task was killed.
15. Finishing: The task got the finish command to terminate processing while it was still running.

These statuses help in tracking the progress of a task and identifying any issues that might occur during its execution.

Ask PanDA

Static

Example 3

- However, if you ask ChatGPT/Gemini/.. about a list of, say, PanDA pilot error codes it will either not know the answer or show something incoherent/scrambled and not complete
- **Using your own input document gives you full control**

Ask PanDA Chatbot

user: Please show all the pilot error codes in numeric order

bot: Here are all the pilot error codes in numeric order:

- 1008 GENERALERROR General pilot error, consult batch log
- 1098 NOLOCALSPACE Not enough local space
- 1099 STAGEINFAILED Failed to stage-in file
- 1100 REPLICANOTFOUND The rucio API function list_replicas() did not return any replicas. Check log for details.
- 1103 NOSUCHFILE No such file or directory
- 1104 USERDIRTOOLARGE User work directory too large
- 1106 STDOUTTOOBIG Payload log or stdout file too big
- 1110 SETUPFAILURE Failed during payload setup
- 1115 NFSSQLITE NFS SQLite locking problems
- 1117 QUEUEDATANOTOK Pilot found non-valid queuedata
- 1124 OUTPUTFILETOOLARGE Output file too large
- 1133 NOSTORAGE Fetching default storage failed: no activity related storage defined
- 1137 STAGEOUTFAILED Failed to stage-out file
- 1141 PUTMD5MISMATCH md5sum mismatch on output file
- 1143 CHMODTRF Failed to chmod trf
- 1144 PANDAKILL This job was killed by panda server
- 1145 GETMD5MISMATCH md5sum mismatch on input file
- 1149 TRFDOWNLOADFAILURE Transform could not be downloaded
- 1150 LOOPINGJOB Looping job killed by pilot
- 1151 STAGEINTIMEOUT File transfer timed out during stage-in
- 1152 STAGEOUTTIMEOUT File transfer timed out during stage-out
- 1163 NOPROXY Grid proxy not valid
- 1165 MISSINGOUTPUTFILE Local output file is missing
- 1168 SIZETOOLARGE Total file size too large
- 1171 GETADMISMATCH Adler32 mismatch on input file

Ask PanDA Chatbot

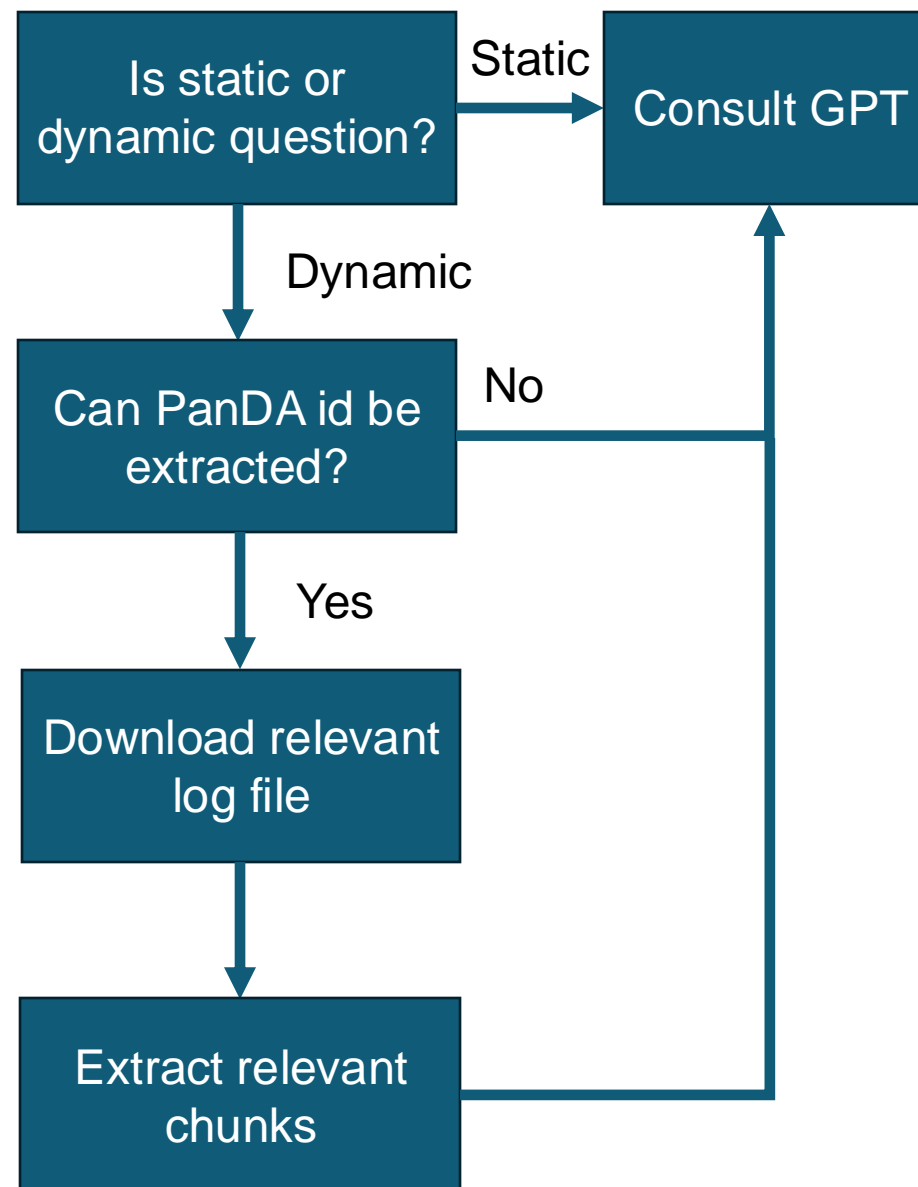
user: Please show all the pilot error codes in numeric order

bot: Here are all the pilot error codes in numeric order:

- 1008 GENERALERROR General pilot error, consult batch log
- 1098 NOLOCALSPACE Not enough local space
- 1099 STAGEINFAILED Failed to stage-in file
- 1100 REPLICANOTFOUND The rucio API function list_replicas() did not return any replicas. Check log for details.
- 1103 NOSUCHFILE No such file or directory
- 1104 USERDIRTOOLARGE User work directory too large
- 1106 STDOUTTOOBIG Payload log or stdout file too big
- 1110 SETUPFAILURE Failed during payload setup
- 1115 NFSSQLITE NFS SQLite locking problems
- 1116 QUEUEDATA Pilot could not download queuedata
- 1117 QUEUEDATANOTOK Pilot found non-valid queuedata
- 1124 OUTPUTFILETOOLARGE Output file too large
- 1133 NOSTORAGE Fetching default storage failed: no activity related storage defined
- 1137 STAGEOUTFAILED Failed to stage-out file
- 1141 PUTMD5MISMATCH md5sum mismatch on output file
- 1143 CHMODTRF Failed to chmod trf
- 1144 PANDAKILL This job was killed by panda server
- 1145 GETMD5MISMATCH md5sum mismatch on input file
- 1149 TRFDOWNLOADFAILURE Transform could not be downloaded
- 1150 LOOPINGJOB Looping job killed by pilot
- 1151 STAGEINTIMEOUT File transfer timed out during stage-in
- 1152 STAGEOUTTIMEOUT File transfer timed out during stage-out
- 1163 NOPROXY Grid proxy not valid
- 1165 MISSINGOUTPUTFILE Local output file is missing
- 1168 SIZETOOLARGE Total file size too large
- 1171 GETADMISMATCH Adler32 mismatch on output file
- 1172 PUTADMISMATCH Adler32 mismatch on output file
- 1177 NOVOMSPROXY Voms proxy not valid
- 1180 GETGLOBUSSYSERR Globus system error during stage-in
- 1181 PUTGLOBUSSYSERR Globus system error during stage-out
- 1186 NOSOFTWAREDIR Software directory does not exist
- 1187 NOPAYLOADMETADATA Payload metadata does not exist
- 1190 LFNTOOLONG LFN too long (exceeding limit of 255 characters)
- 1191 ZEROFILFSIZE File size cannot be zero
- 1196 MKDIR FAILED to create local directory
- 1200 KILLSIGNAL Job terminated by unknown kill signal
- 1201 SIGTERM Job killed by signal: SIGTERM
- 1202 SIGQUIT Job killed by signal: SIGQUIT
- 1203 SIGSEGV Job killed by signal: SIGSEGV
- 1204 SIGXCPU Job killed by signal: SIGXCPU
- 1205 USERKILL Job killed by user
- 1206 SIGBUS Job killed by signal: SIGBUS
- 1207 SIGUSR1 Job killed by signal: SIGUSR1
- 1211 MISSINGINSTALLATION Missing installation
- 1212 PAYLOADOUTOFMEMORY Payload ran out of memory
- 1213 REACHEDMAXTIME Reached batch system time limit
- 1220 UNKNOWNPAYLOADFAILURE Job failed due to unknown reason (consult log file)
- 1221 FILEEXISTS File already exists
- 1223 BADALLOC Transform failed due to bad_alloc
- 1224 ESRECOVERABLE Event service: recoverable error
- 1228 ESFATAL Event service: fatal error
- 1234 EXECUTEDCLONEJOB Clone job is already executed
- 1235 PAYLOADEXCEEDMAXMEM Payload exceeded maximum memory
- 1236 KILLEDBYSERVER Killed by server
- 1238 ESNOEVENTS Event service: no events
- 1240 MESSAGEHANDLINGFAILURE Failed to handle message from payload
- 1242 CHKSUMNOTSUP Query checksum is not supported
- 1244 NORELEASEFOUND No release candidates found
- 1246 NOUSERBARBALL User tarball could not be downloaded from PanDA server
- 1247 BADXML Badly formed XML
- 1300 NOTIMPLEMENTED The class or function is not implemented.
- 1301 UNKNOWNEXCEPTION An unknown pilot exception has occurred
- 1302 CONVERSIONFAILURE Failed to convert object data
- 1303 FILEHANDLINGFAILURE Failed during file handling
- 1305 PAYLOADEXECUTIONFAILURE Failed to execute payload
- 1306 SINGULARITYGENERALFAILURE Singularity: general failure
- 1307 SINGULARITYNOLOOPDEVICES Singularity: No more available loop devices
- 1308 SINGULARITYBINDPOINTFAILURE Singularity: Not mounting requested bind point
- 1309 SINGULARITYIMAGEMOUNTFAILURE Singularity: Failed to mount image
- 1310 PAYLOADEXECUTIONEXCEPTION Exception caught during payload execution
- 1311 NOTDEFINED Not defined
- 1312 NOTSAMELENGTH Not same length
- 1313 NOSTORAGEPROTOCOL No protocol defined for storage endpoint
- 1314 UNKNOWNCHECKSUMTYPE Unknown checksum type
- 1315 UNKNOWNTRFFAILURE Unknown TRF failure
- 1316 RUCIOSERVICEUNAVAILABLE Rucio: Service unavailable
- 1317 EXCEEDEDMAXWAITTIME Exceeded maximum waiting time
- 1318 COMMUNICATIONFAILURE Failed to communicate with server
- 1319 INTERNALPILOTPROBLEM An internal Pilot problem has occurred (consult Pilot log)
- 1320 LOGFILECREATIONFAILURE Failed during creation of log file
- 1321 RUCIOLOCATIONFAILURE Failed to get client location for Rucio
- 1322 RUCIOLISTREPLICASFAILURE Failed to get replicas for Rucio
- 1323 UNKNOWNCOPYTOOL Unknown copy tool
- 1324 SERVICEUNAVAILABLE Service not available at the moment
- 1325 SINGULARITYNOTINSTALLED Singularity: not installed
- 1326 NOREPLICAS No matching replicas were found in list_replicas() output
- 1327 UNREACHABLE NETWORK Unable to stage-in file since network is unreachable

Dynamic Queries

- For questions like “why did my job fail?”, GPT will of course not know anything
- Script therefore extracts relevant info from query and gets info from PanDA before asking GPT what went wrong
 - If the job ID can be extracted, e.g. the pilot log (only) will be downloaded from PanDA
 - Relevant chunks (error messages, warnings and such) is then sent to GPT
- Response is returned from GPT
- Currently only the pilot log is investigated but same methodology can be used for any log file
 - User could request in the query
 - Why did the payload in job XYZ fail?
→ payload.stdout/stderr gets downloaded
- Script can also download and analyze info from standard job page (JSON version)
 - Trivial - GPT is not involved in this case



Dynamic

Example 1

- NB currently not much more info is revealed than the monitor page, but if error codes can be extracted, more explanations can be provided (script has list of pilot error codes included)

ATLAS PanDA Dash Tasks Jobs Errors Users Sites Harvester My BigPanDA Job by ID Enter... ProdSys Services Help Paul

Job details for PanDA job 6518760159 aipanda194 | 10:17:53, Refresh

Job name: e07de1e5-9998-4c99-9edb-4441248e4c5b_97814

| PanDA ID | Owner | Task ID | Status | Transformation | Created Last modified | Time to start Duration [d:h:m:s] | Site | Harvester instance Worker ID | Cores | Priority | Attempt |
|------------|----------|---------|--------|-----------------|--|----------------------------------|-------|------------------------------|-------|----------|---------|
| 6518760159 | gangarbt | None | failed | runGen-00-00-02 | 2025-02-17 14:27:57
2025-02-18 06:03:46 | 0:15:34:30
0:0:01:14 | MPPMU | CERN_central_B
631308978 | 8 | 1000 | 1 |

Datasets: In: mc15_13TeV:mc15_13TeV.423202.Pythia8B_A14_CTEQ6L1_Jpsie3e13.merge.AOD.e3869_s2608_s2183_r6630_r6264_tid05536542_00
Out: hc_test.gangarbt.hc20310361.tid1079.MPPMU.62

Files summary: input: 1, size: 354.11 (MB); log: 1; output: 1

Job error summary:

Error description: **exe, 2002:** payload execution failed with 223 **pilot, 1305:** payload execution failed with 223 **trans, 223:** Proot: Input files cannot be opened

| Component | Code | Message |
|-----------|------|-----------------------------------|
| Execution | 2002 | payload execution failed with 223 |
| Pilot | 1305 | payload execution failed with 223 |

user: why did job 6518760159 fail?

bot: The job 6518760159 failed due to a non-zero exit code returned during the main payload execution, specifically exit code 223. The payload error report indicates that the execution failed with error code 2002, and the pilot discovered an additional error code 1305, which signifies a failure to execute the payload. This sequence of errors led to the job being marked as 'failed' and subsequently aborted.

Example 2

- Stage-out failure with a little bit more explanations

ATLAS PanDA Dash Tasks Jobs Errors Users Sites Harvester My BigPanDA Job by ID Enter... ProdSys Services Help Paul

Job details for PanDA job 6519134287 aipanda196 | 10:25:29, Refresh

Job name: d735573f-ccf4-4243-a431-f6fc5a60974d_11520

| PanDA ID | Owner | Task ID | Status | Transformation | Created
Last modified | Time to start
Duration
[d:h:m:s] | Site | Harvester instance
Worker ID | Cores | Priority | Attempt |
|------------|----------|---------|--------|----------------|--|--|------------|---------------------------------|-------|----------|---------|
| 6519134287 | gangarbt | None | failed | Sim_tf.py | 2025-02-17 23:25:21
2025-02-18 00:42:05 | 0:1:14:30
0:0:02:05 | RO-16-UAIC | CERN_central_A
681760491 | 8 | 1000 | 1 |

Datasets:
In: hc_test:mc15_13TeV.361106.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Zee.evgen.EVNT.e3601_tid04972714_00
Out: hc_test.gangarbt.hc20310359.tid1080.RO-16-UAIC.68

Files summary: input: 1, size: 35.91 (MB); log: 1; output: 1

Job error summary:

Error description: **pilot, 1324:** Service not available at the moment: mc15_13TeV:EVNT.04972714._000029.pool.root.1 from RO-07-NIPNE_DATADISK, Error on XrdCl:CopyProcess:Run(): [ERROR] Invalid redirect URL: (source));failed to transfer files using copytools=['rucio']

| Component | Code | Message |
|-----------|------|---|
| Pilot | 1324 | Service not available at the moment: mc15_13TeV:EVNT.04972714._000029.pool.root.1 from RO-07-NIPNE_DATADISK, Error on XrdCl:CopyProcess:Run(): [ERROR] Invalid redirect URL: (source));failed to transfer files using copytools=['rucio'] |

user: why was job 6519134287 unsuccessful?

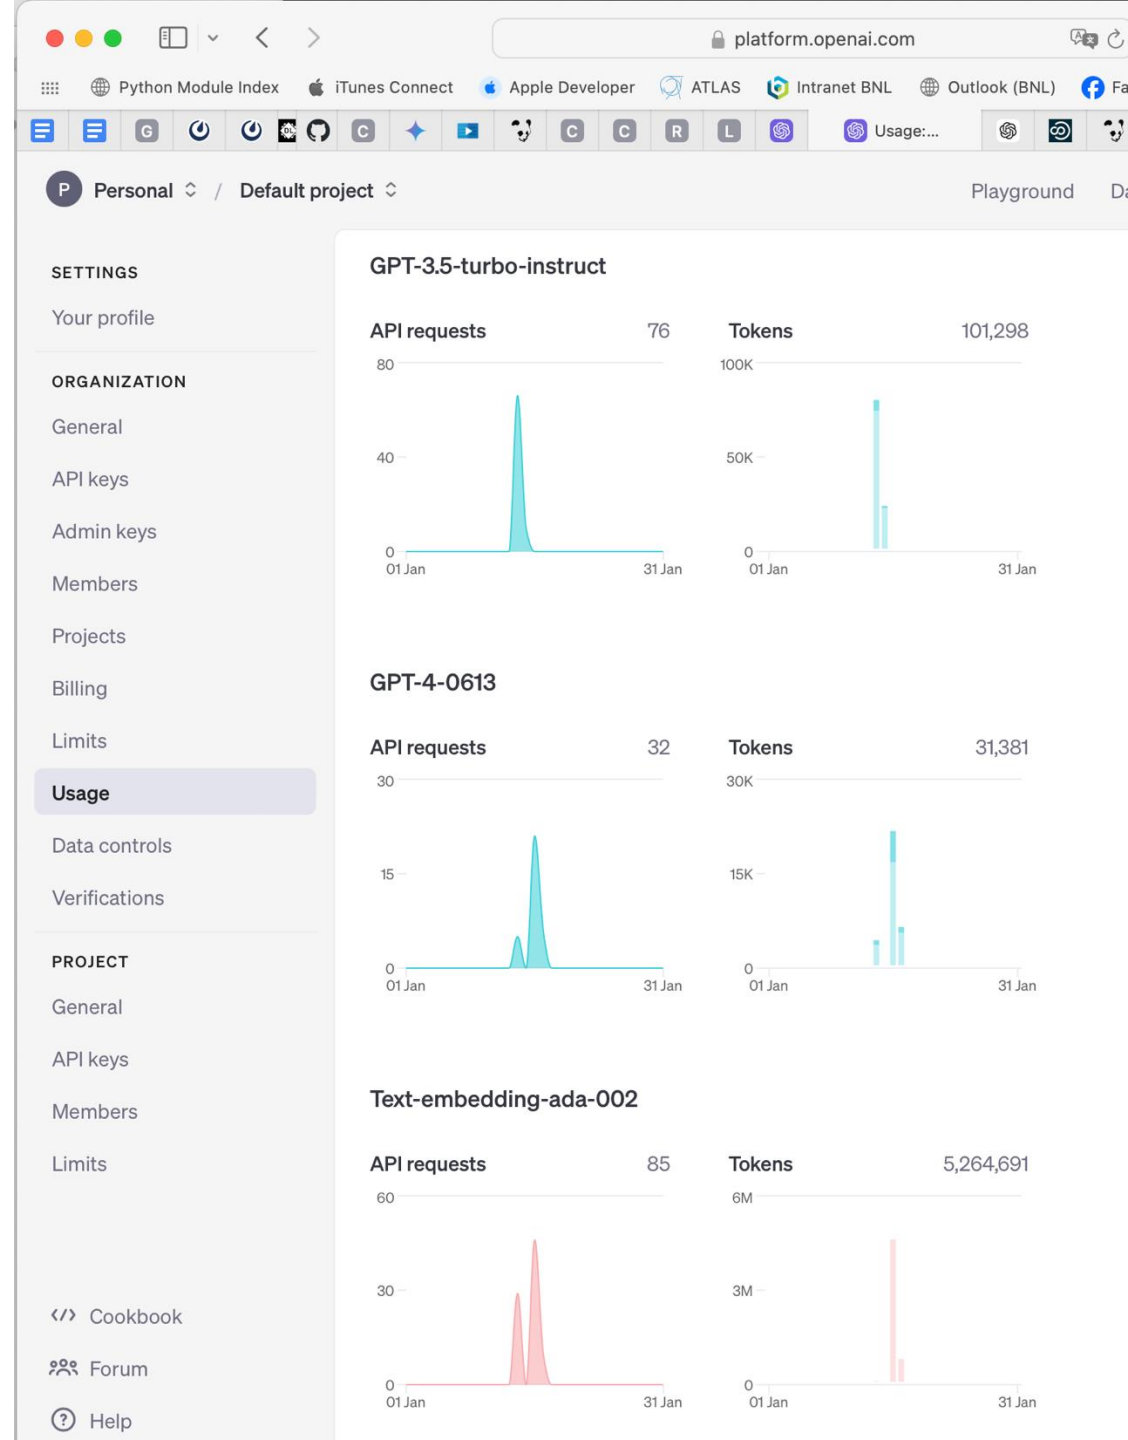
bot: The job 6519134287 was unsuccessful due to multiple issues related to file transfer and service availability. The logs indicate that there were several attempts to download files using the Rucio copytool, but these attempts failed due to an "unknown error" and specifically an "Error on Xrd."

Key points from the logs include:

- The destination file could not be deleted due to an unknown error (error code 3011).
- There were multiple download attempts that failed, with the message indicating that the requested service was not available at the moment.
- The transfer process encountered an error code 1324, which also pointed to the service being unavailable.

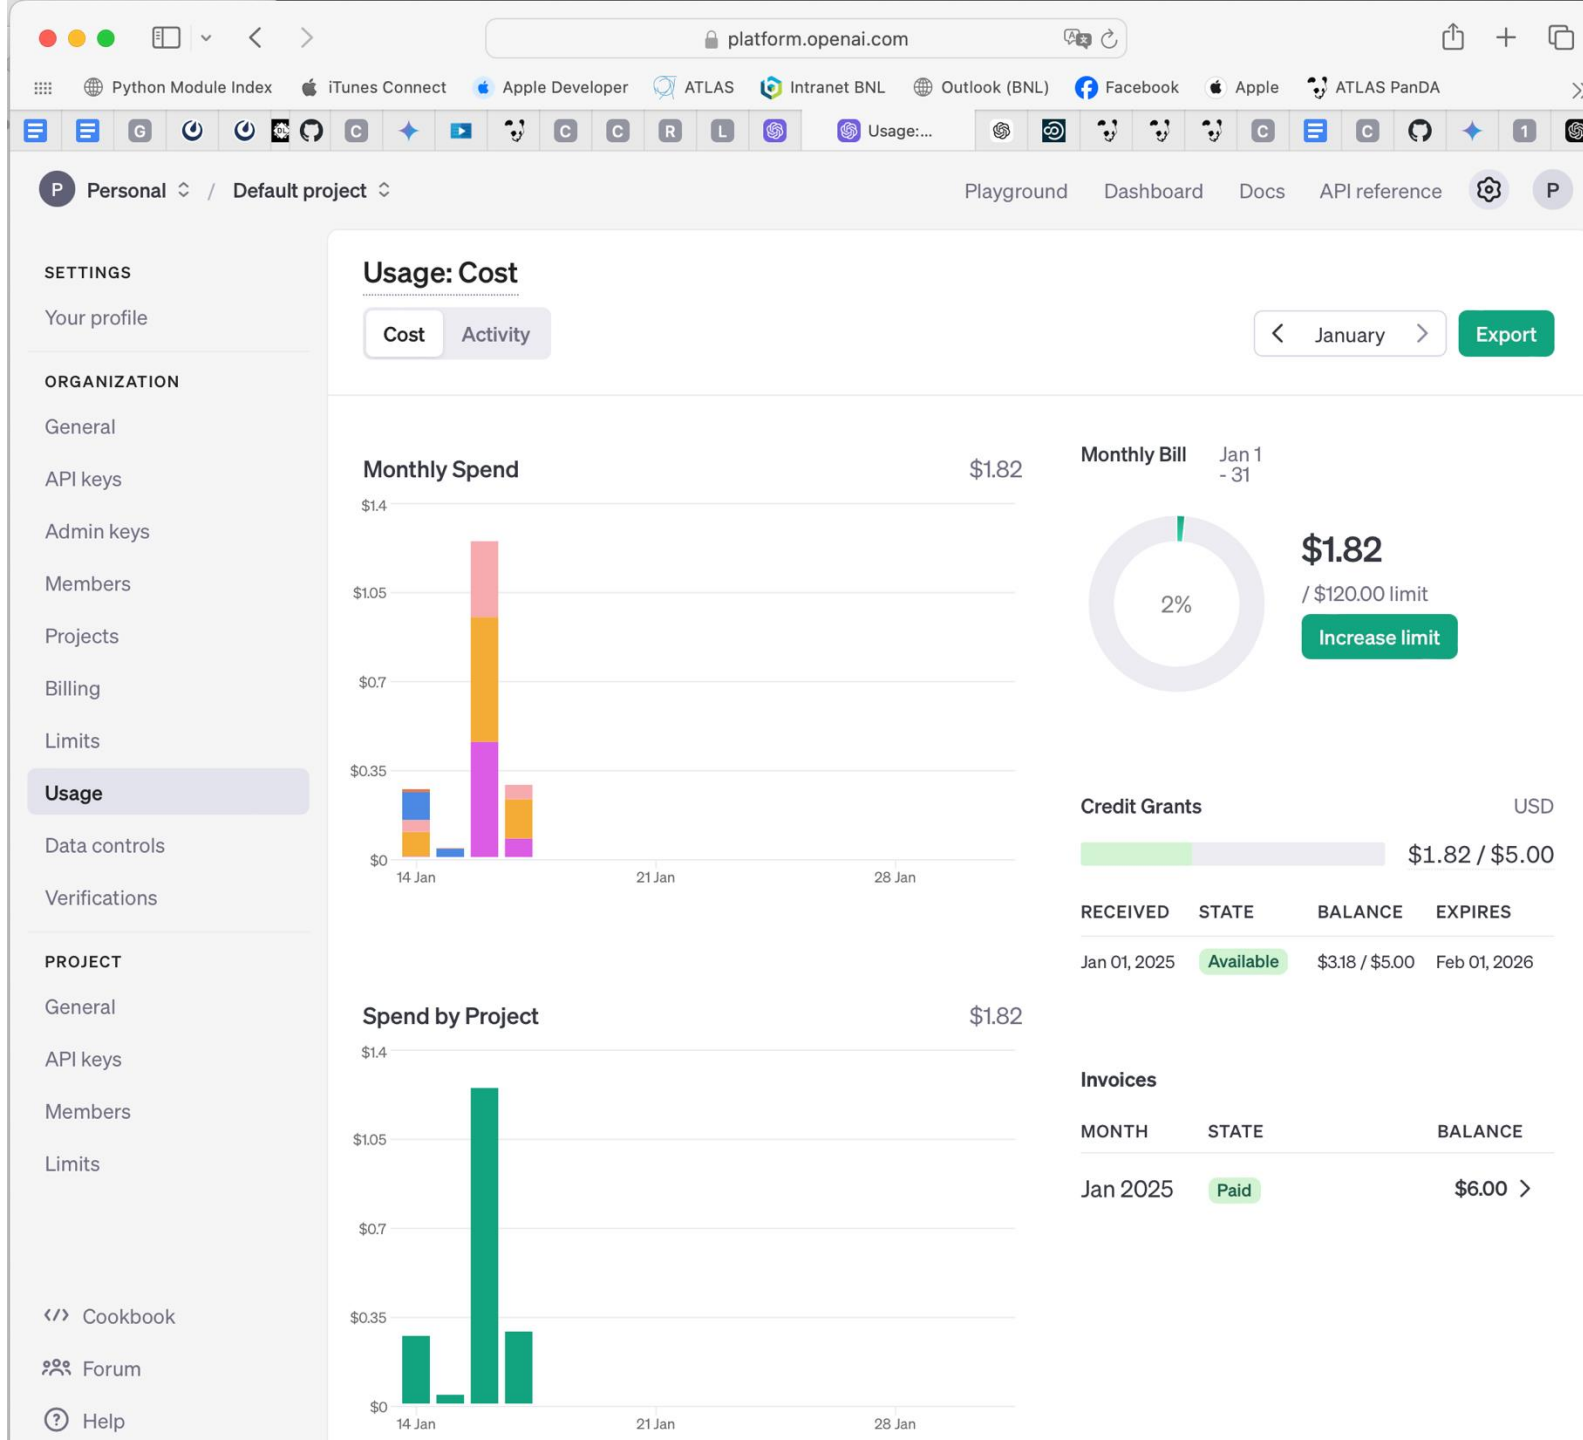
Usage statistics

- Models are priced individually
- 76 API requests with 101k tokens sent to gpt-3.5
 - User questions
- 32 API requests with 31k tokens sent to gpt-4
 - User questions
- 85 API requests with 5.2M tokens sent to embedding model
 - To find the most relevant chunk for a given question



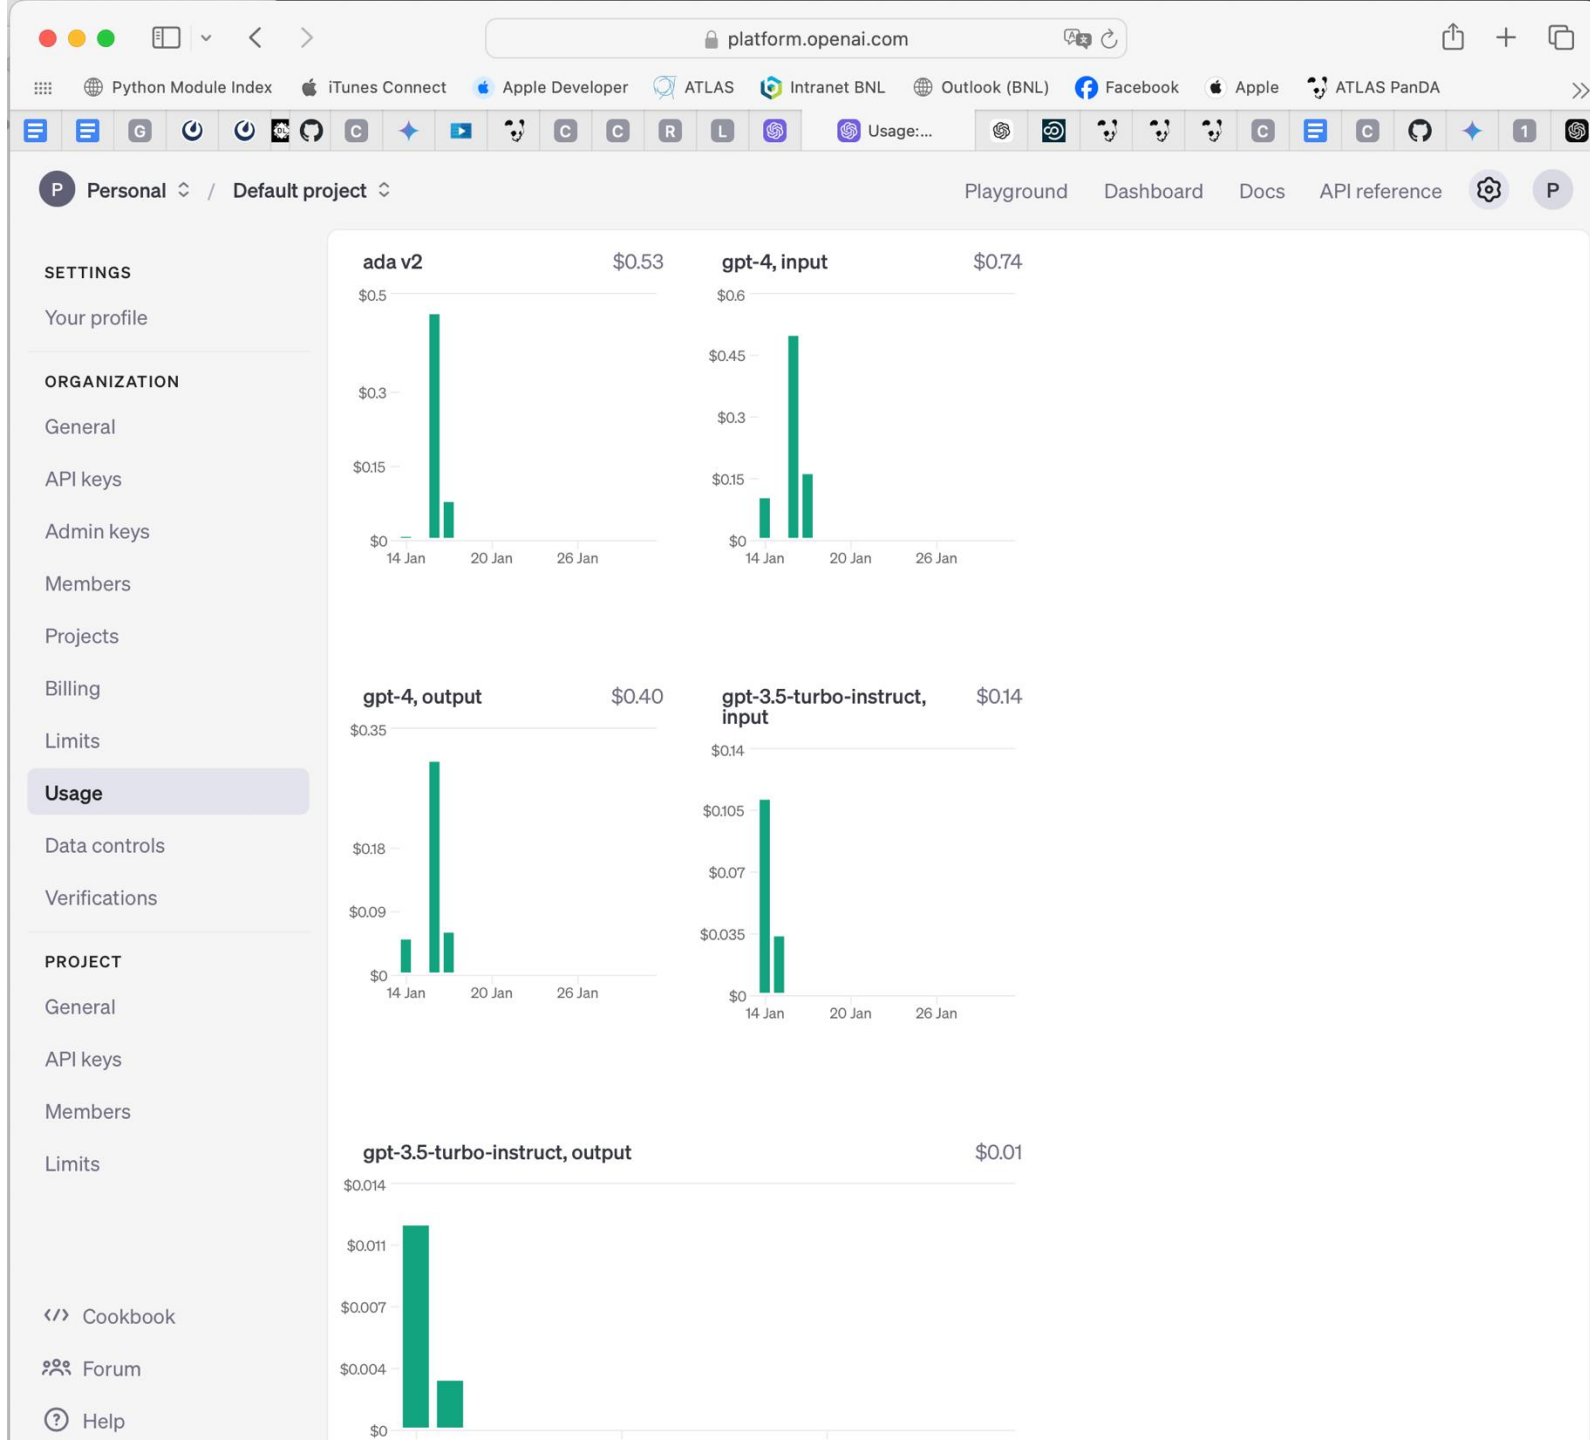
Usage Costs

- Total cost after around 100 questions: \$1.82
 - But mostly using GPT-3.5 – which is better to use if the quality of the answers is not always important
- When I switched to the better GPT-4 model, I quickly ran out of money



Breakdown

- Cost for embedding model: \$0.53
- Cost for questions + responses: \$1.29
- GPT-3.5 model really cheap, but also not so good..



Outlook

- Improve input document wherever needed, collect feedback
 - Pilot documentation not part of PanDA readthedocs (only linked), so add that and possibly other bits and pieces (in progress)
- Implement handling of more advanced queries
 - Capability of selecting different log files
 - (I'm mostly interested in error analysis – which also ties into REDWOOD project)
- Move project to Django to align with new BigPanDAMon project
- Unfortunately not for free and API not currently part of the BNL OpenAI ChatGPT license
 - Need to discuss funding at some point

