



# **EAF status and updates**

Maria P. Acosta (FNAL)


SCD projects meeting

Sept 8<sup>th</sup> 2022

# Elastic Analysis Facility Team

- Burt Holzman – Project lead
- Maria Acosta – Technical lead for applications
- Chris Bonnaud – Technical lead for infrastructure
- Joe Boyd
- Glenn Cooper
- Lindsey Gray
- Farrukh Khan
- Ed Simmonds
- Nick Smith
- Elise Chavez

# We're live!! (on Beta) <https://analytics-hub.fnal.gov>



Welcome to JupyterHub @ the Fermilab Elastic Analysis Facility

Use your Fermi SERVICES domain credentials to log in

If you have an existing environment and want to run it as a notebook, go to [EAF BinderHub](#)

EAF is in beta testing phase. This is the point where we need your help:

- Please note that GPU availability is on a first come, first serve basis. If you request a notebook with a GPU and it times out, please try again later.
- Inactive/idle notebooks will be automatically stopped after 8 hours
- To report your feedback please visit the following [GitHub issue](#), open as a safe feedback space.
- If you uncover a security issue, please report it privately by emailing [eaf-admins@fnal.gov](mailto:eaf-admins@fnal.gov)
- If you find any other regressions, please open an issue in the [EAF GitHub repository](#)
- If you don't find any issues, we also appreciate positive input. Make sure to add the successful update on the [feedback space](#).

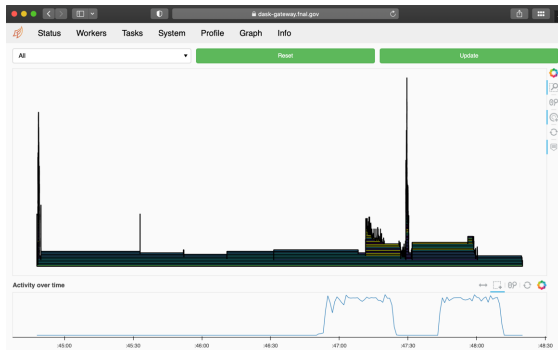
Sign in

Username:

Password:

Sign in

- 39 Beta users (thank you!)
- 26 Notebook flavors
- 2.6 Tb Ceph persistent storage allocated (of 45TB)



# September News

- Completed Helm chart and JupyterHub application updates to latest releases. New capabilities include Token-based access control, upgraded SQLite DB schemas and latest installation of [Kubespawner](#).
- Thanks to infrastructure team, A100 GPU installation into the cluster was completed.
- Iterated over several MIG partition configurations, 20Gb vGPU mem seems to be the ideal setup for current use cases.
- Burt facilitated agreement between the Lab and Anaconda to restore previously blocked URLs due to licensing issues, this enabled the BinderHub service to be operational again.
- CVMFSexec is working as expected with individual user caches, R&D is ongoing for alternative Cache setups.



## Fundamental principles:

- Create a user-oriented analysis facility based on our own experiences supporting scientists on traditional grid technologies.
- Explore, deploy and collaborate on industry-level technologies and strategies for optimizing data analysis in preparation for HL-LHC and upcoming experiments with large data demands such as DUNE.
- Support science groups across the lab on advancing data analysis techniques and expand the toolset and hardware available for Fermilab experiments.
- Foster collaboration with experiments and research groups within the Lab in order to better understand science analysis needs and provide computing solutions accordingly.

Secure

Integrated & functional

Multi-VO

DevOps (operational  
sustainability)

Active collaboration

# Status and upcoming (Sept 2022)

Secure

Integrated & functional

Multi-VO

DevOps (operational  
sustainability)

Active collaboration

## Status

- Dev instance running with FedID capabilities via CILogon which uses Fermilab SSO for user authentication.
- Started integrating FERRY/EAF to pull ACCEL-AI users, as first step to detach from SSI users. Pending conversation with Steve White for adding EAF certs into FERRY's auth list

## Upcoming:

- Decoupling gpgrid and LPC users from ssi-users is next.
- FedID is implemented but needs extensive testing.

# Status and upcoming (Sept 2022)

Secure

Integrated & functional

Multi-VO

DevOps (operational  
sustainability)

Active collaboration

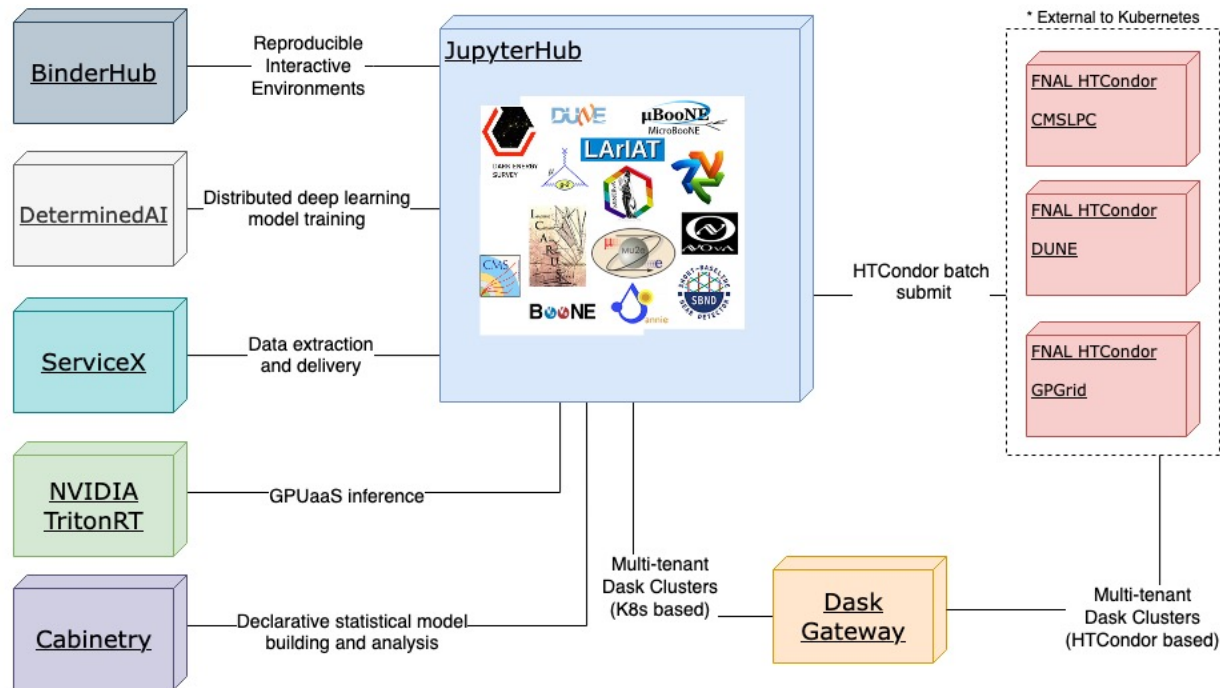
## Status

- New tech -> new challenges. With the addition of A100 GPUs, drivers supporting older models stopped working. The project has decided to drop support for the older K40s and repurposed the T4 outside of OKD.
- BinderHub is back in service, currently in alfa stage. We are ramping up after anaconda block was lifted.
- Our Dask deployment started failing after upgrade. Older libraries are not supported with the current Kubernetes version, we are working on a clean upgrade to the refactored release of Dask Gateway.

## Upcoming:

- Delays on Dask workers joining clusters because they are subject to scheduling limitations when they enter the batch system as jobs. We're working on a solution with T1/LPC admins.
- Phasing out old version of Dask Gateway and pushing external backend changes upstream.

# Current applications Ecosystem





# Status and upcoming (Sept 2022)

Secure

Integrated & functional

Multi-VO

DevOps (operational sustainability)

Active collaboration

## Status

- NEW documentation site:  
<https://eafjupyter.readthedocs.io/en/latest/>
- Working heavily with current beta users to bring several current use cases to production-grade quality.
- Catalog changes/modernization is coming up, already testing some ideas on dev.

## Upcoming:

- Reach out to experiment Liaisons and define functional and non-functional requirements for specific experiments/science groups
- Provide reproducible curated environments for users and allow them to develop and deploy their own notebook environments.

# Status and upcoming (Sept 2022)

Secure

Integrated & functional

Multi-VO

DevOps (operational sustainability)

Active collaboration

Current Catalog

Prototype

Server Options

- LBNF/DUNE/ProtoDUNE**
  - SL7 Interactive General Purpose Notebook
  - GPU SL7 Interactive (NVIDIA Tesla T4)
  - GPU SL7 Interactive (NVIDIA Tesla K40m)
- ACORN**
  - ACORN CPU Only SL7 Interactive
  - ACORN GPU SL7 Interactive (NVIDIA Tesla T4)
  - L-CAPE CPU Only SL7 Interactive
  - L-CAPE GPU SL7 Interactive (NVIDIA Tesla K40)
  - L-CAPE GPU SL7 Interactive (NVIDIA Tesla T4)
  - READS CPU Only SL7 Interactive
  - READS GPU SL7 Interactive (NVIDIA Tesla K40)
  - READS GPU SL7 Interactive (NVIDIA Tesla T4)
- FIFE/Neutrinos**
  - SL7 Interactive General Purpose Notebook
  - GPU SL7 Interactive (NVIDIA Tesla K40m)
  - GPU SL7 Interactive (NVIDIA Tesla T4)
- Cosmic Frontier**
  - SL7 Interactive General Purpose Notebook
  - GPU SL7 Interactive (NVIDIA Tesla K40m)
  - GPU SL7 Interactive (NVIDIA Tesla T4)
- Fermi generic SL7/CC8**
  - Basic SL7 Interactive
  - Basic OCS Interactive
- CMSLPC**
  - SL7 Interactive
  - COFFEA-DASK SL7 Interactive
  - GPU SL7 Interactive (NVIDIA Tesla K40m)
  - GPU SL7 Interactive (NVIDIA Tesla T4)
  - GPUaaS - Boosted Decision Trees SL7 Interactive (NVIDIA Tesla T4)

Start

CMS  
CMS collaboration

Environment

SL7 Vanilla

CMS LPC SL7 Interactives



CMS  
CMS collaboration

Environment

- ✓ SL7 Vanilla
- SL7 COFFEA-Dask
- SL7 Boosted Decision Trees (GPUaaS)
- SL7 with NVIDIA T4 GPU

# Status and upcoming (Sept 2022)

Secure

Integrated & functional

Multi-VO

DevOps (operational  
sustainability)

Active collaboration

## Status

- Monitoring was ported properly and redesigned for better insight on Jupyter metrics.
- (From July slides) OKD cluster upgrade is happening next week (Mon-Tues) to latest OKD (4.10). Will solve a performance problem of starting notebooks when users have many files ~1M. This upgrade was successful, we were able to restart application operations after a few minutes of downtime.
- Operations learning: Checklists, git repo documentation, pre-upgrade spreadsheets, checks and teamwork helped ensure a clean install.

## Upcoming:

- Proactive monitoring alarms via CheckMK
- Re-evaluating the purpose and state of our public GitHub repository - currently operating from private GitLab

# Status and upcoming (Sept 2022)

Secure

Integrated & functional

Multi-VO

DevOps (operational sustainability)

Active collaboration

## Monitoring for JupyterHub



- NEW external proxy monitoring



- Deeper insight on spawning process duration and outcomes for each step: poll, spawn, stop.

# Status and upcoming (Sept 2022)

Secure

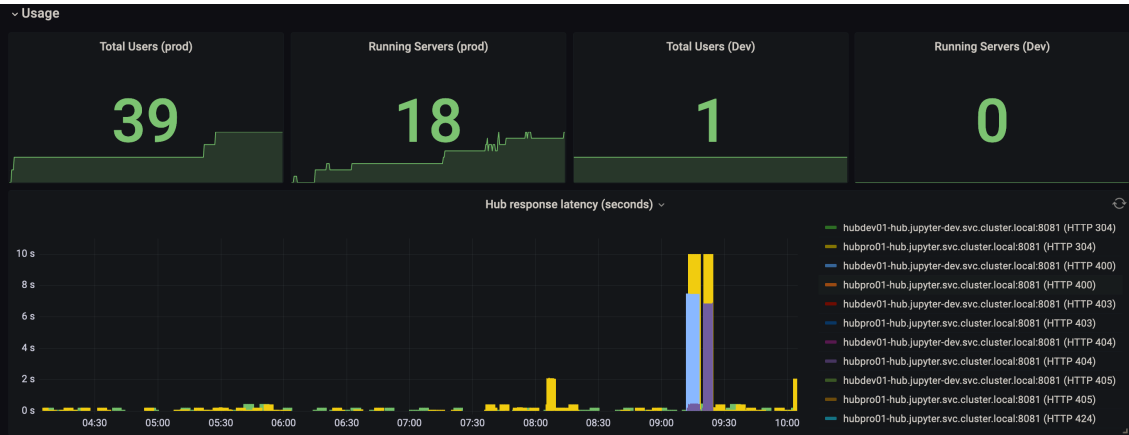
Integrated & functional

Multi-VO

DevOps (operational sustainability)

Active collaboration

## Monitoring for JupyterHub



- Spawning process duration and Hub (application) startup time



- Dev and Prod instances scraped for metrics every 30s
- Response latency and HTTP response codes

# Status and upcoming (Sept 2022)

Secure

Integrated & functional

Multi-VO

DevOps (operational  
sustainability)

Active collaboration

## Status


- Getting data in and out of the facility is a big priority and a general question from all our users. We are working on a case-to-case basis to accommodate current users' needs. Starting to communicate with storage R&D experts.
- Tutorial-style guides and example jupyter notebooks will be included as part of our documentation effort.

## Upcoming:

- Project is currently outlining plans for user communication and support channels as well as feedback spaces

# We need your input

- Have a project that could benefit from JupyterHub?
- Is there a computing need or requirement that fits the AF model?



## Contact us!

- Email me ([macosta@fnal.gov](mailto:macosta@fnal.gov)) and Burt Holzman ([burt@fnal.gov](mailto:burt@fnal.gov)) with your thoughts!
- Visit the **NEW** Documentation site: <https://eafjupyter.readthedocs.io>
- If you uncover a security issue, please report it privately by emailing [eaf-admins@fnal.gov](mailto:eaf-admins@fnal.gov).
- If you find any other regressions, please open an issue in the [EAF GitHub repository](#).
- If you don't find any issues, we also appreciate positive input. Make sure to add the successful update on the [feedback space](#).

## Summary and outlook:

- **GPUs are in demand! (again)** Thanks to our users for feedback during GPU partitioning, many interesting projects going on.
- How can we effectively bring data in and out of the facility? Still evaluating possibilities but no clear path forward yet.
- The **facility is gaining traction and interest** from multiple groups and institutions, working round the clock to bring it to a production-ready state.
- **Big focus on documentation**, user channels, feedback and other activities in preparation for inviting a broader set of users to use the facility.
- Thank you to our beta users for the amazing feedback! Please keep it coming

Thanks 😊 Questions?

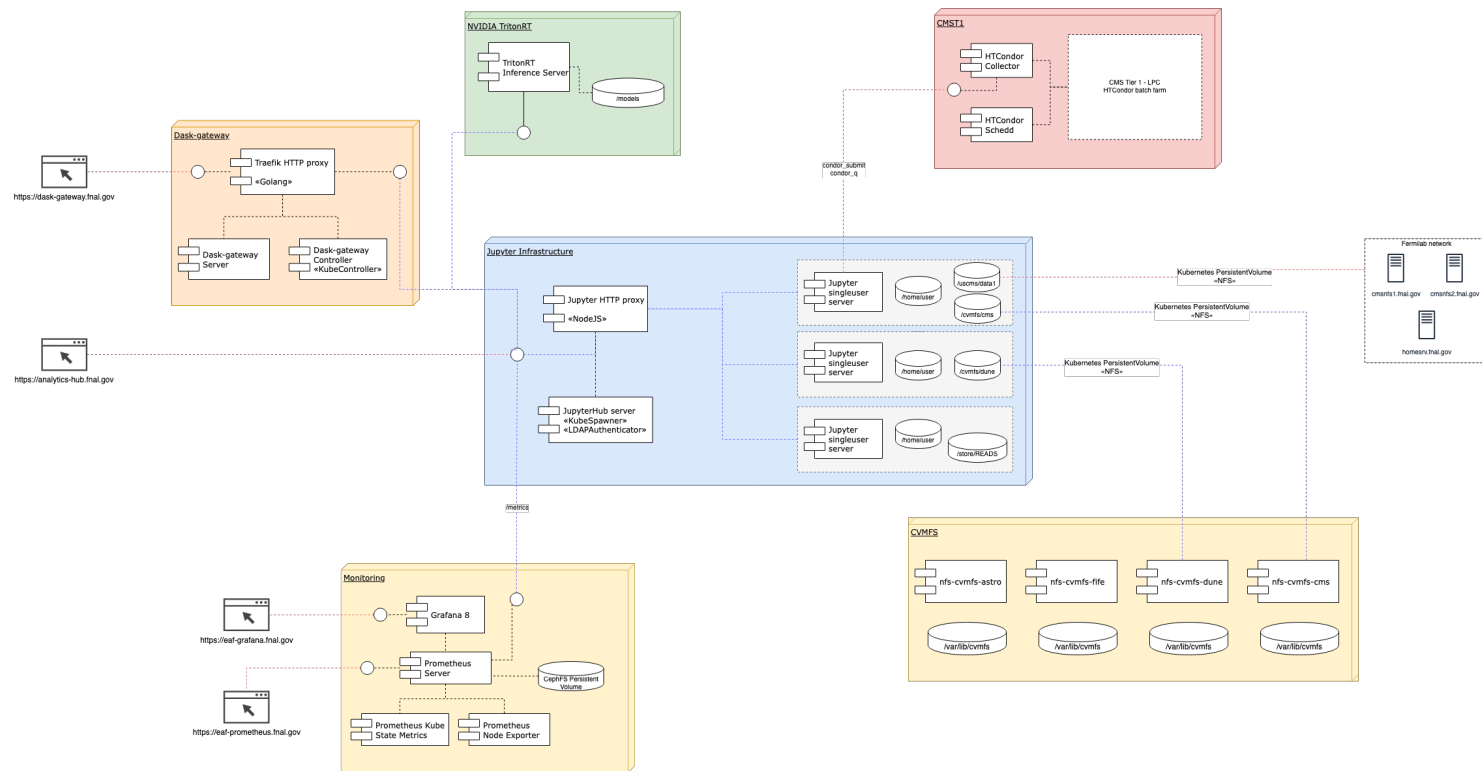
Maria Acosta - ACORN, EAF

[macosta@fnal.gov](mailto:macosta@fnal.gov)

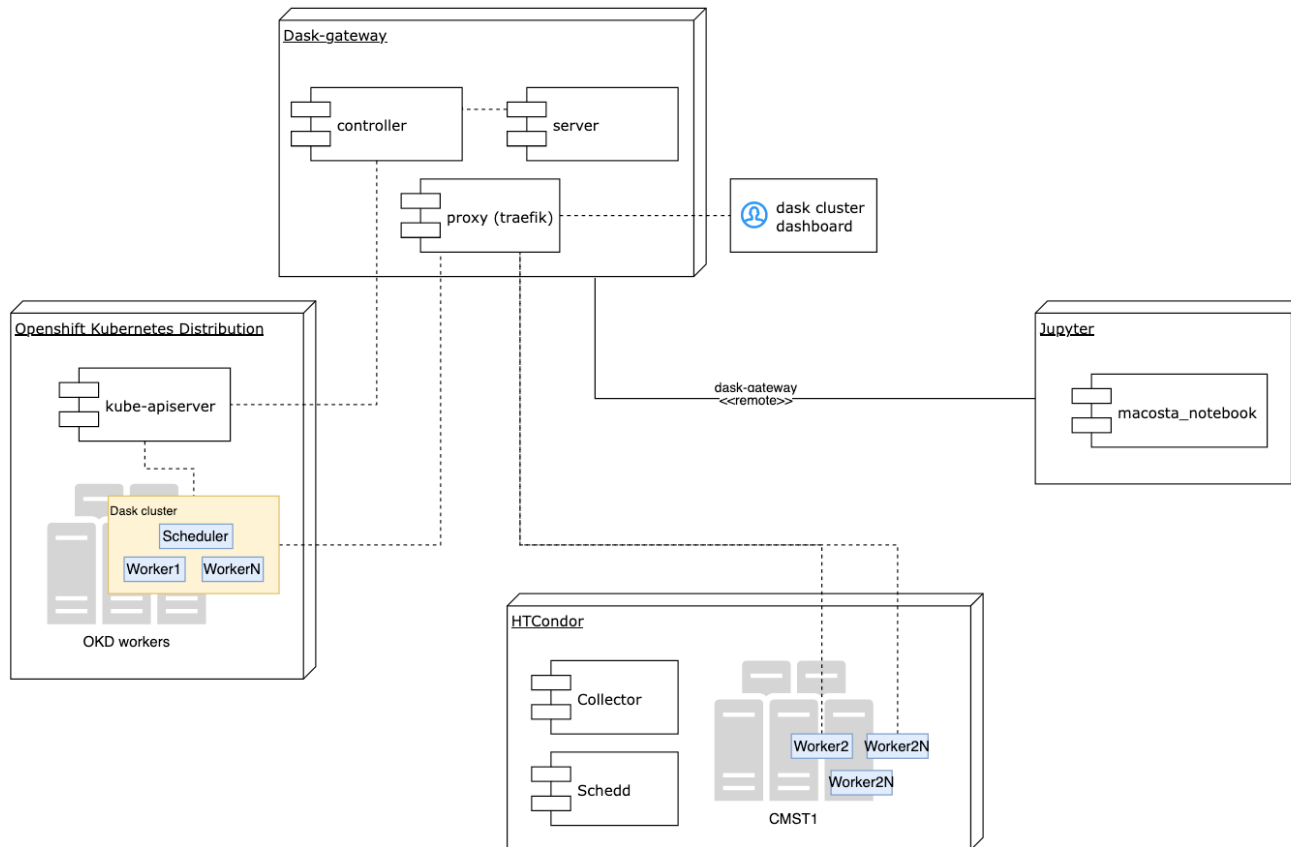
@macosta on Slack



# Backup – detailed component diagram



# Backup - Dask cluster burst – the ‘Elastic’ side of the facility



# Backup – login page (<https://analytics-hub.fnal.gov>)

The screenshot shows a web browser window with the address bar displaying `analytics-hub.fnal.gov`. The page header includes the JupyterHub logo and the Fermilab logo. The main content area contains a welcome message and a sign-in form. The sign-in form has an orange header with the text "Sign in". Below the header, there are two input fields: "Username:" with the value "macosta" and "Password:" with a masked password ".....". A "Sign in" button is located below the password field. To the left of the sign-in form, there is a message about OpenShift Kubernetes cluster upgrades and a list of beta testing instructions.

**Welcome to JupyterHub @ the Fermilab Elastic Analysis Facility**

Use your Fermi SERVICES domain credentials to log in

The OpenShift Kubernetes cluster will have upgrades installed on 2/16 from 10:00 – 11:00 am. As a result, Jupyter notebook pods may be killed and re-launched during this time.

If you have an existing environment and want to run it as a notebook, go to EAF [BinderHub](#)

EAF is in beta testing phase. This is the point where we need your help:

- Please note that GPU availability is on a first come, first serve basis. If you request a notebook with a GPU and it times out, please try again later.
- Inactive/Idle notebooks will be automatically stopped after 8 hours
- To report your feedback please visit the following [GitHub issue](#), open as a safe feedback space.
- If you uncover a security issue, please report it privately by emailing [eaf-admins@fnal.gov](mailto:eaf-admins@fnal.gov)
- If you find any other regressions, please open an issue in the [EAF GitHub repository](#)
- If you don't find any issues, we also appreciate positive input. Make sure to add the successful update on the [feedback space](#).

**Sign in**

**Username:**

**Password:**

**Sign in**

# Backup – named servers

analytics-hub.fnal.gov

jupyterhub Home Token Help/FAQ Admin Services macosta Logout

Stop My Server

My Server

## Named Servers

In addition to your default server, you may have additional 5 server(s) with names. This allows you to have more than one server running at the same time.

Server name	URL	Last activity	Actions
<input type="text" value="Name your server"/>	<a href="#">Add New Server</a>		
ad		21 days ago	<a href="#">start</a> <a href="#">delete</a>
dask	<a href="/user/macosta/dask">/user/macosta/dask</a>	a day ago	<a href="#">stop</a>
dune		a month ago	<a href="#">start</a> <a href="#">delete</a>
fife		3 months ago	<a href="#">start</a> <a href="#">delete</a>
lpc		5 days ago	<a href="#">start</a> <a href="#">delete</a>

# Backup – a CMSLPC notebook running AGC COFFEA analysis

File Edit View Run Kernel Git Diagram Tabs Settings Help

macosta@jupyter-macost x coffea.ipynb

Markdown git Python 3 (ipykernel)

Filter files by name

/ ... / analyses / cms-open-data-ttbar /

Name	Last Modified
figures	a day ago
histograms	a day ago
utils	a day ago
Y: cabinetry_conf...	a day ago
coffea_fnal_eaf...	a day ago
coffea.ipynb	seconds ago
histograms.root	2 minutes ago
ntuples.json	a day ago
workspace.json	2 minutes ago

```
[8]: all_histograms[:, "4j2b", :, "nominal"].stack("process")[::-1].plot(stack=True, histtype="fill", linewidth=...
plt.legend(frameon=False)
plt.title(">= 4 jets, >= 2 b-tags")
plt.xlabel("$m_{bj}$ [GeV]");
```

>= 4 jets, >= 2 b-tags

single\_top\_t\_chan  
single\_top\_s\_chan  
wjets  
single\_top\_tW  
t\_tbar

$m_{bj}$  [GeV]

Our top reconstruction approach ( $bjj$  system with largest  $p_T$ ) has worked!

Let's also have a look at some systematic variations:

- b-tagging, which we implemented as jet-kinematic dependent event weights,
- jet energy variations, which vary jet kinematics, resulting in acceptance effects and observable changes.

Simple 2 0 Python 3 (ipykernel) | Idle Mode: Command Ln 1, Col 1 coffea.ipynb