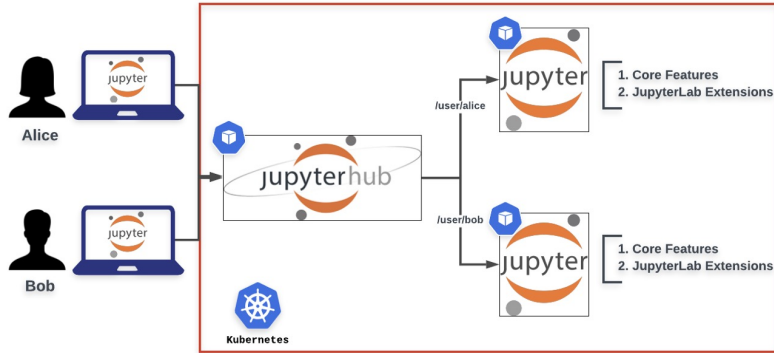



# Elastic Analysis Facility: a JupyterHub-based deployment



- Originally standalone Jupyter Notebooks
- Evolved to a self-hosted, multi-user platform for hosting multiple notebooks, kernels and **highly customizable** environments.
- Can be deployed in multiple platforms including Cloud, on prem and Kubernetes.

- ✓ Implements authentication, login pages and token-based roles
- ✓ Tracks activity and does effective resource management
- ✓ Proxying is done behind the scenes
- ✓ Pseudo-interactive (not HPC or HTPC): **launches container on single node**

# A JupyterHub-based deployment - Login and Auth



Welcome to JupyterHub @ the Fermilab Elastic Analysis Facility

Click below to login

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

Sign in with Fermilab SSO

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](#).

- Accessible from the Lab network or via VPN
- Login with SSO
- UID/GID is propagated to the notebook in order to preserve permissions

<https://analytics-hub.fnal.gov>



# Documentation: [eafjupyter.readthedocs.io](https://eafjupyter.readthedocs.io)

🏠 Fermilab Elastic Analysis Facility

latest

Search docs

**ACCOUNTS AND AUTHENTICATION**

EAF Access and Accounts

**JUPYTER - THE BASICS**

About JupyterHub

Multiple sessions

Virtual Studio Code

Help, something's broken!

**SERVER AND NOTEBOOK OPTIONS**

Getting started - Choosing a Notebook

Fermi Generic SL7/CC8

FIFE/Neutrinos

LBNF/DUNE/ProtoDUNE

Cosmic Frontier

🏠 / Welcome to the Fermilab EAF documentation!

[🔗 Edit on GitHub](#)

## Welcome to the Fermilab EAF documentation!

The **Fermilab Elastic Analysis Facility** is a kubernetes-based platform providing services for integrated data, software and computational resources to execute one or more elements of scientific analysis workflows for Fermilab experiments and scientists. These resources are centrally managed by the Scientific Computing Division at Fermilab as part of the EAF project and generally available for all Fermilab users. Check out the userdocs/index section for more information on environments, general user documentation for the analysis tools ecosystem and navigating the UI.

### 📌 Important

You can get help by:

- Filing a ticket via [ServiceNow](#)
- Joining the `#eaf-users` channel on `fnal.slack.com` (e-mail eaf-admins AT fnal.gov for an invitation if you don't already have access to the slack)
- Joining and emailing the mailing list: `eaf-users@fnal.gov`

# Support

## 📌 Important

You can get help by:

- Filing a ticket via [ServiceNow](#)
- Joining the `#eaf-users` channel on `fnal.slack.com` (e-mail eaf-admins AT fnal.gov for an invitation if you don't already have access to the slack)
- Joining and emailing the mailing list: `eaf-users@fnal.gov`

# A JupyterHub-based deployment - Current Catalog

## Server Options

GPUS (used/capacity):  
10GB (4/8), 20GB (17/20), 40GB (2/2)



**CMS**  
CVMFS, HTCCondor, COFFEA

CPU Interactives

SL7 COFFEA-Dask

NVIDIA® A100 GPU

SL7 - 10GB GPU slot



**LBNF DUNE/ProtoDUNE**  
CVMFS, LarSoft

CPU Interactives

SL7 Vanilla

NVIDIA® A100 GPU

SL7 - 10GB GPU slot



**FIFE**  
CVMFS Neutrinos/Mu2e/gm2

CPU Interactives

SL7 Vanilla

NVIDIA® A100 GPU

SL7 - 10GB GPU slot



**Astro/Cosmic Frontier**  
CVMFS, LSST kernel

CPU Interactives

SL7 LSST kernel

NVIDIA® A100 GPU

SL7 - 10GB GPU slot



**ACCEL-AI**  
Tensorflow, pyTorch

CPU Interactives

SL7 L-CAPE

NVIDIA® A100 GPU

SL7 L-CAPE - 10GB GPU slot



**ACORN**  
ACSYS python, Fortran

CPU Interactives

SL7 AIMPS (FORTRAN)

NVIDIA® A100 GPU

SL7 - 10GB GPU slot



**Fermilab**  
Generic SL7 notebooks

CPU Interactives

SL7 Vanilla

NVIDIA® A100 GPU

SL7 - 10GB GPU slot

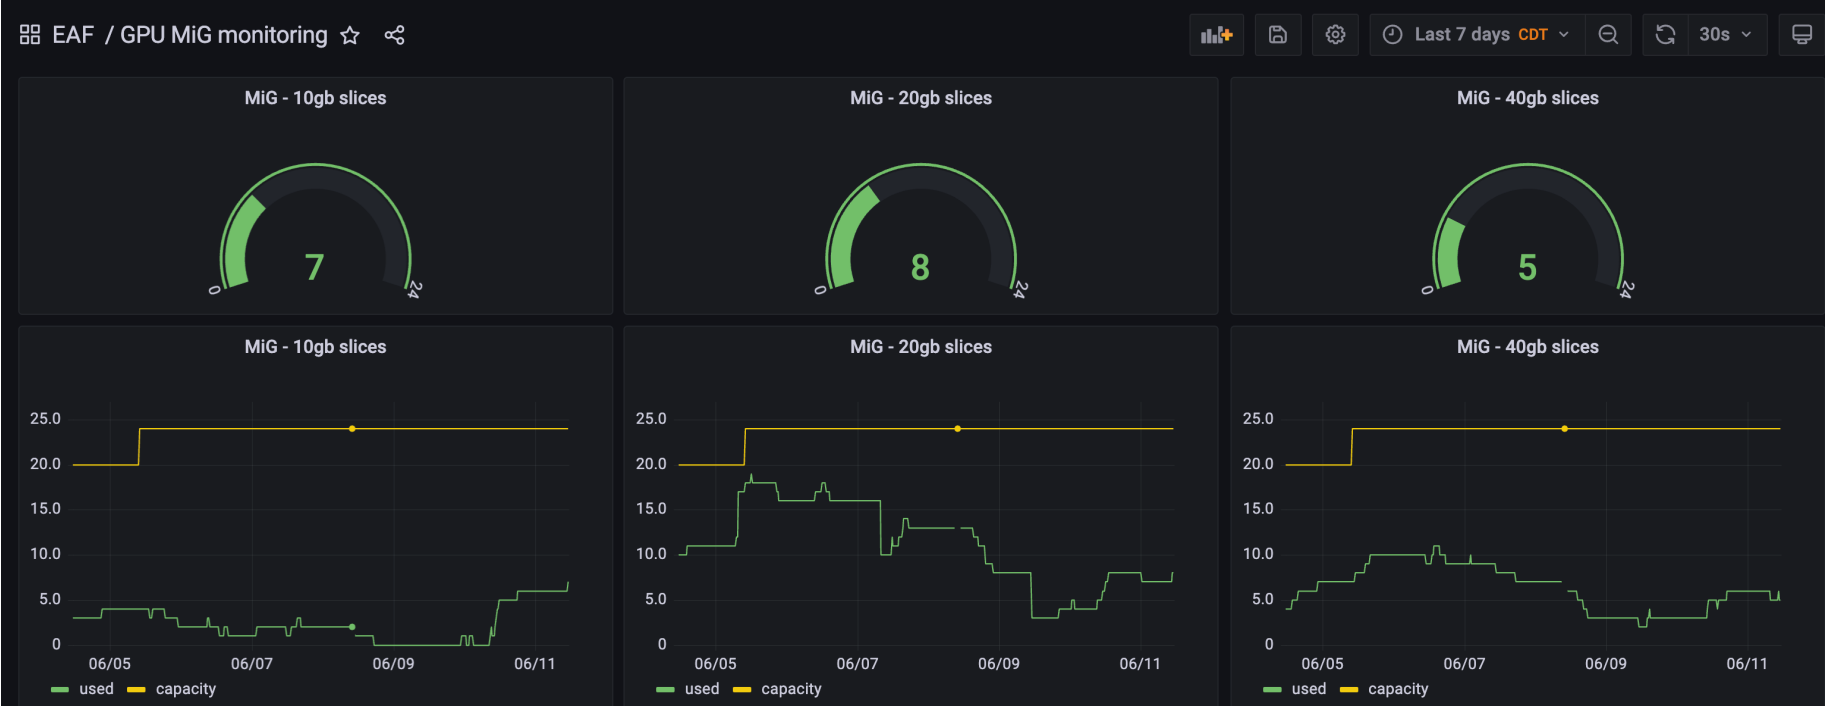
Rescue Image

Rescue Image - for use if overquota or other startup issues

# Resources

- **CPU nodes**
  - 8 cores guaranteed but full node (68-126) cores available when idle
  - 64 GB memory
- **GPU nodes**
  - A100s sliced into 10GB, 20GB, 40GB fractions using Nvidia MiG (full A100 is 80 GB)
  - Up to 90 GB memory

# GPU nodes



# Resources

- **CPU nodes (5 nodes, but can add more)**
  - 8 cores guaranteed but full node (68-126) cores available when idle
  - 64 GB memory
- **GPU nodes (11 nodes, 26 GPUs)**
  - A100s sliced into 10GB, 20GB, 40GB fractions using Nvidia MiG (full A100 is 80 GB)
  - Up to 90 GB memory
- **Storage**
  - **24 GB** EAF user disk (NVMe)
  - **10 TB** shared scratch space (Ceph)
  - **/nashome** mounted
  - **/exp (experiment areas)** mounted
  - **CVMFS** areas mounted
  - Can mount additional Ceph areas
- **Access to batch**
  - JobSub (condor\_submit if CMS)



# User Customization

- **Curated images**
- **Mamba (anaconda)** and **pip**
- **Binder** (build-your-own images) can be deployed if needed
- **VS Code** integration

# Is EAF a good fit for you as a WC user?

- Yes, if:
  - You need a single node/GPU or you need a front-end to HTPC batch systems
  - You are comfortable with notebooks or the web-based terminal
- No, if:
  - You are running multinode jobs (leveraging MPI, multiple GPUs for training, etc.)\*
  - You want to continue to use Slurm as a batch/HPC scheduler\*
  - You absolutely need to use ssh for a pure CLI experience

\* We are doing R&D on adding a Slurm GPU-focused minicluster inside EAF to cover these use cases