


# Producing Madgraph5\_aMC@NLO gridpacks and using TensorFlow GPU resources in the CMS HTCondor Global Pool



Brian Bockelman<sup>6</sup>, Edgar Fajardo Hernandez<sup>4</sup>, Diego Davila Foyo<sup>1</sup>, **Kenji Hurtado Anampa<sup>7</sup>**, Farrukh Aftab Khan<sup>3</sup>, Krista Larson<sup>3</sup>, James Letts<sup>4</sup>, Marco Mascheroni<sup>3</sup>, David Mason<sup>3</sup>, Antonio Perez-Calero Yzquierdo<sup>2</sup>, Todor Trendafilov Ivanov<sup>8</sup>

<sup>1</sup> Autonomous University of Puebla (MX)  
<sup>2</sup> CIEMAT & PIC (ES)  
<sup>3</sup> Fermilab (US)

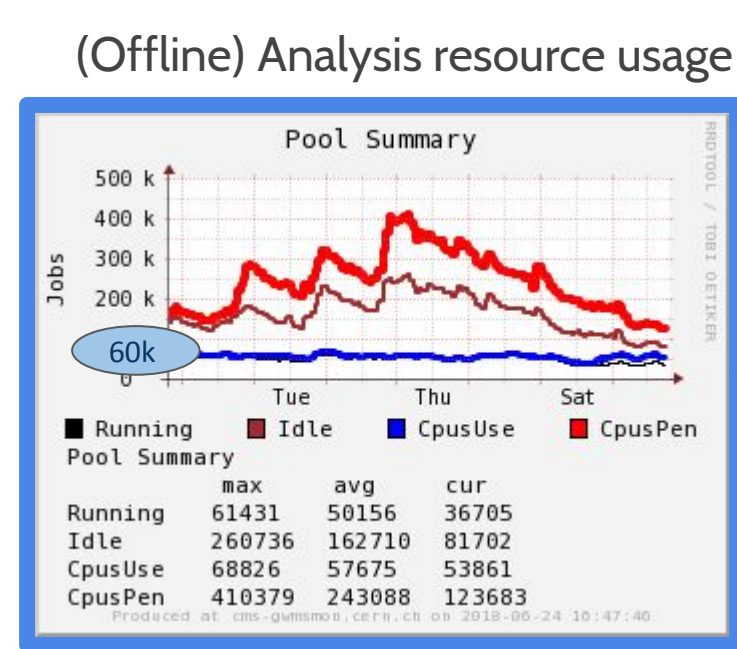
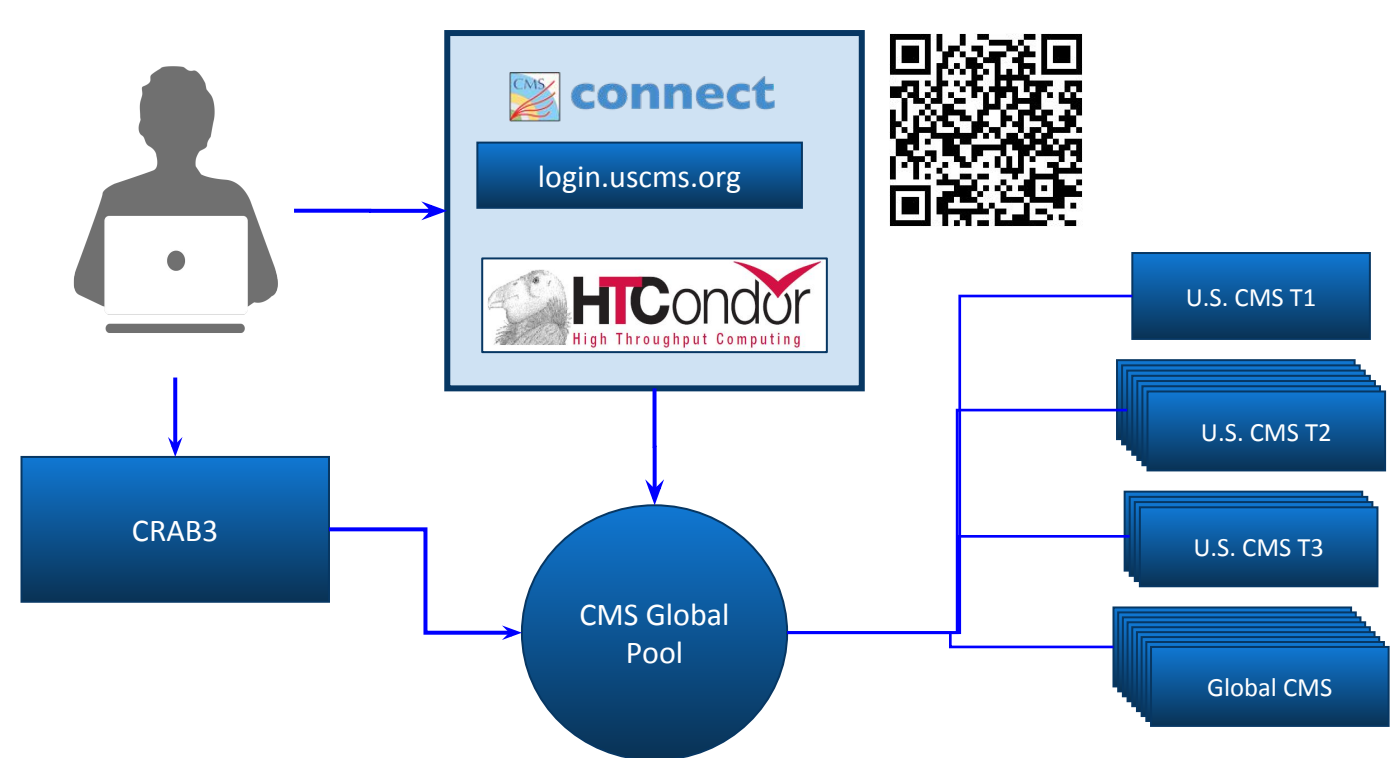
<sup>4</sup> University of California San Diego (US)  
<sup>5</sup> University of Malaya (MY)  
<sup>6</sup> University of Nebraska - Lincoln (US)

<sup>7</sup> University of Notre Dame (US)  
<sup>8</sup> University of Sofia (BG)

While submission of CMS user jobs to the Global Pool is mostly managed by CRAB3, the standard analysis workflow management tool, the generation of matrix elements for high energy physics processes via Madgraph5\_aMC@NLO and the usage of machine learning tools with GPU resources are independent use-cases that require special adaptation in order to take advantage of the Global Pool resources. This work describes the challenges and efforts performed towards adapting such workflows for it.

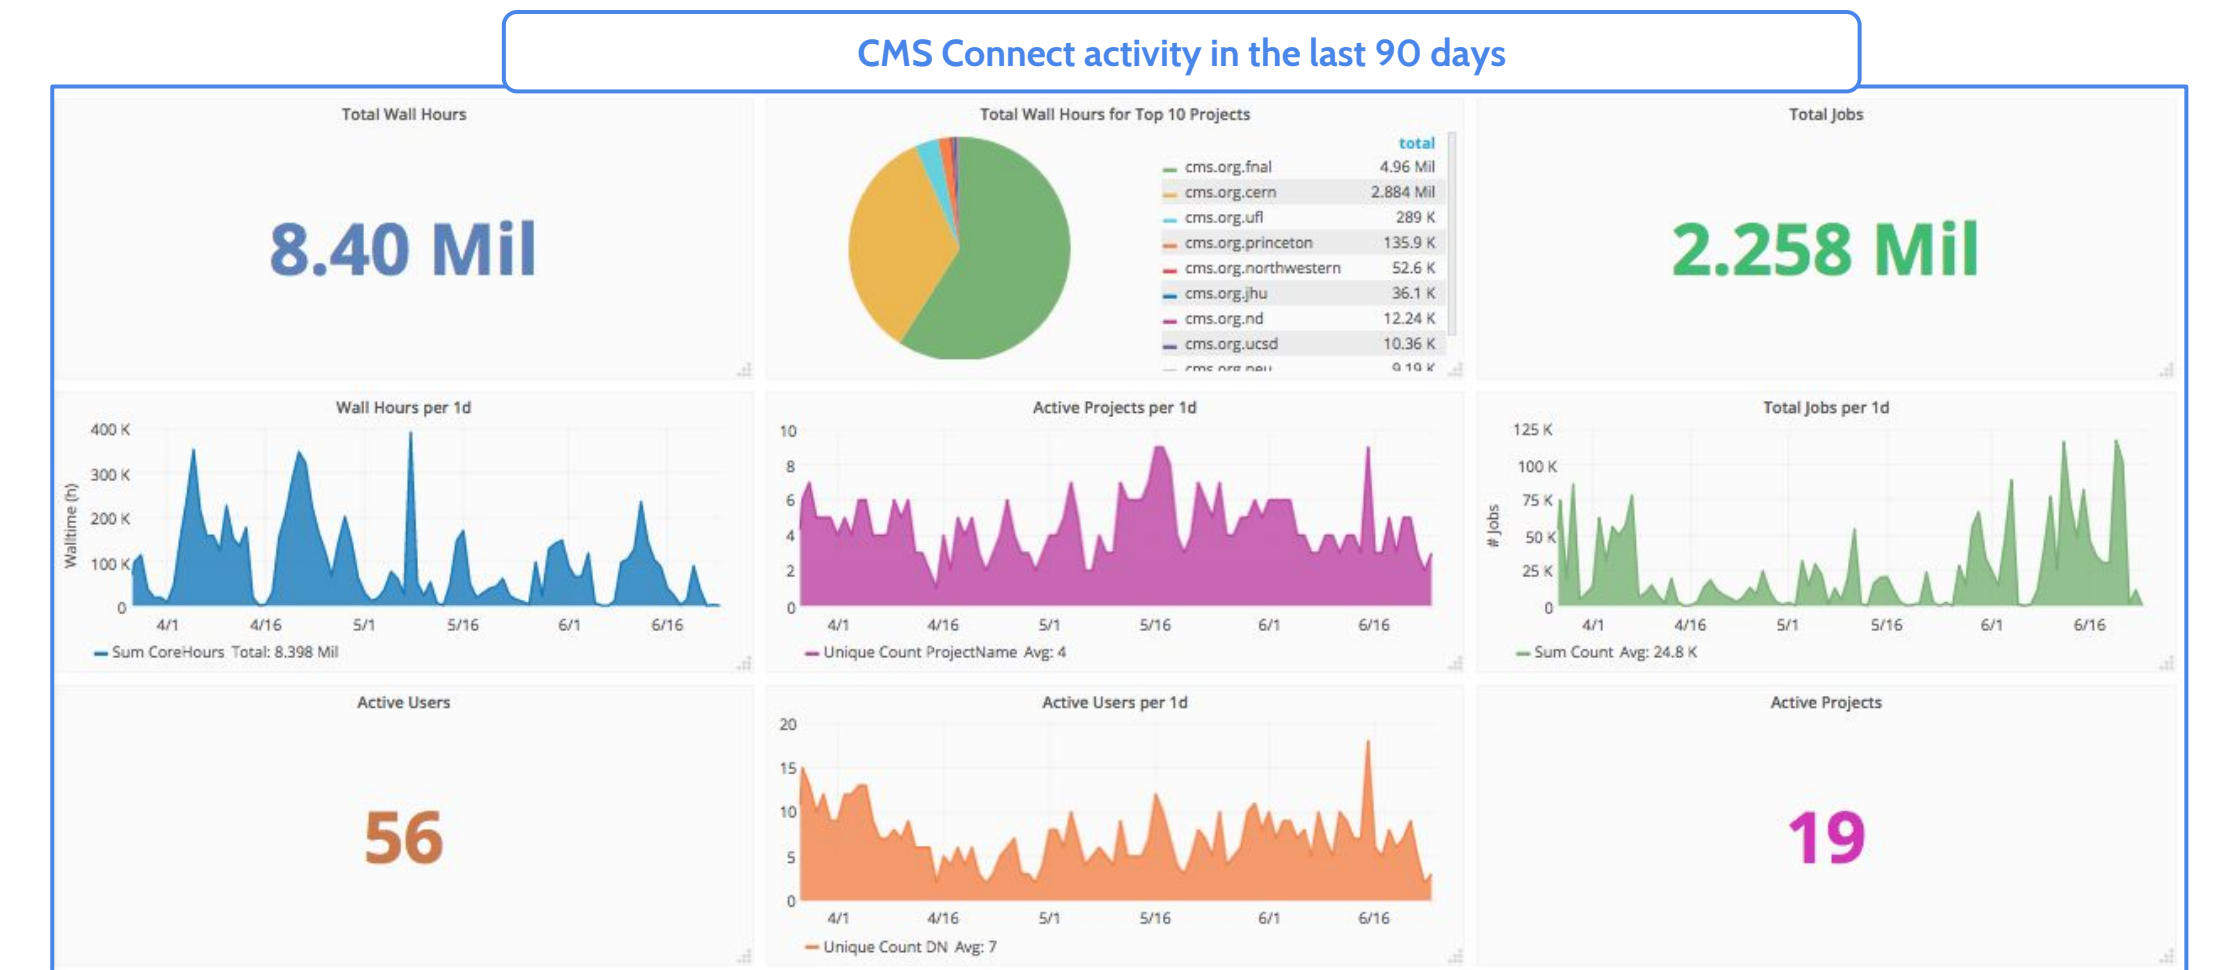
## The submission point - CMS Connect

CMS Connect provides a service with a Tier 3-Like interface where users can submit condor jobs to the CMS Global Pool, a global HTCondor pool provisioned by GlideinWMS. It complements CRAB3, dealing with a different set of analysis workflows, such as Madgraph gridpacks and the use of GPU resources with TensorFlow jobs.

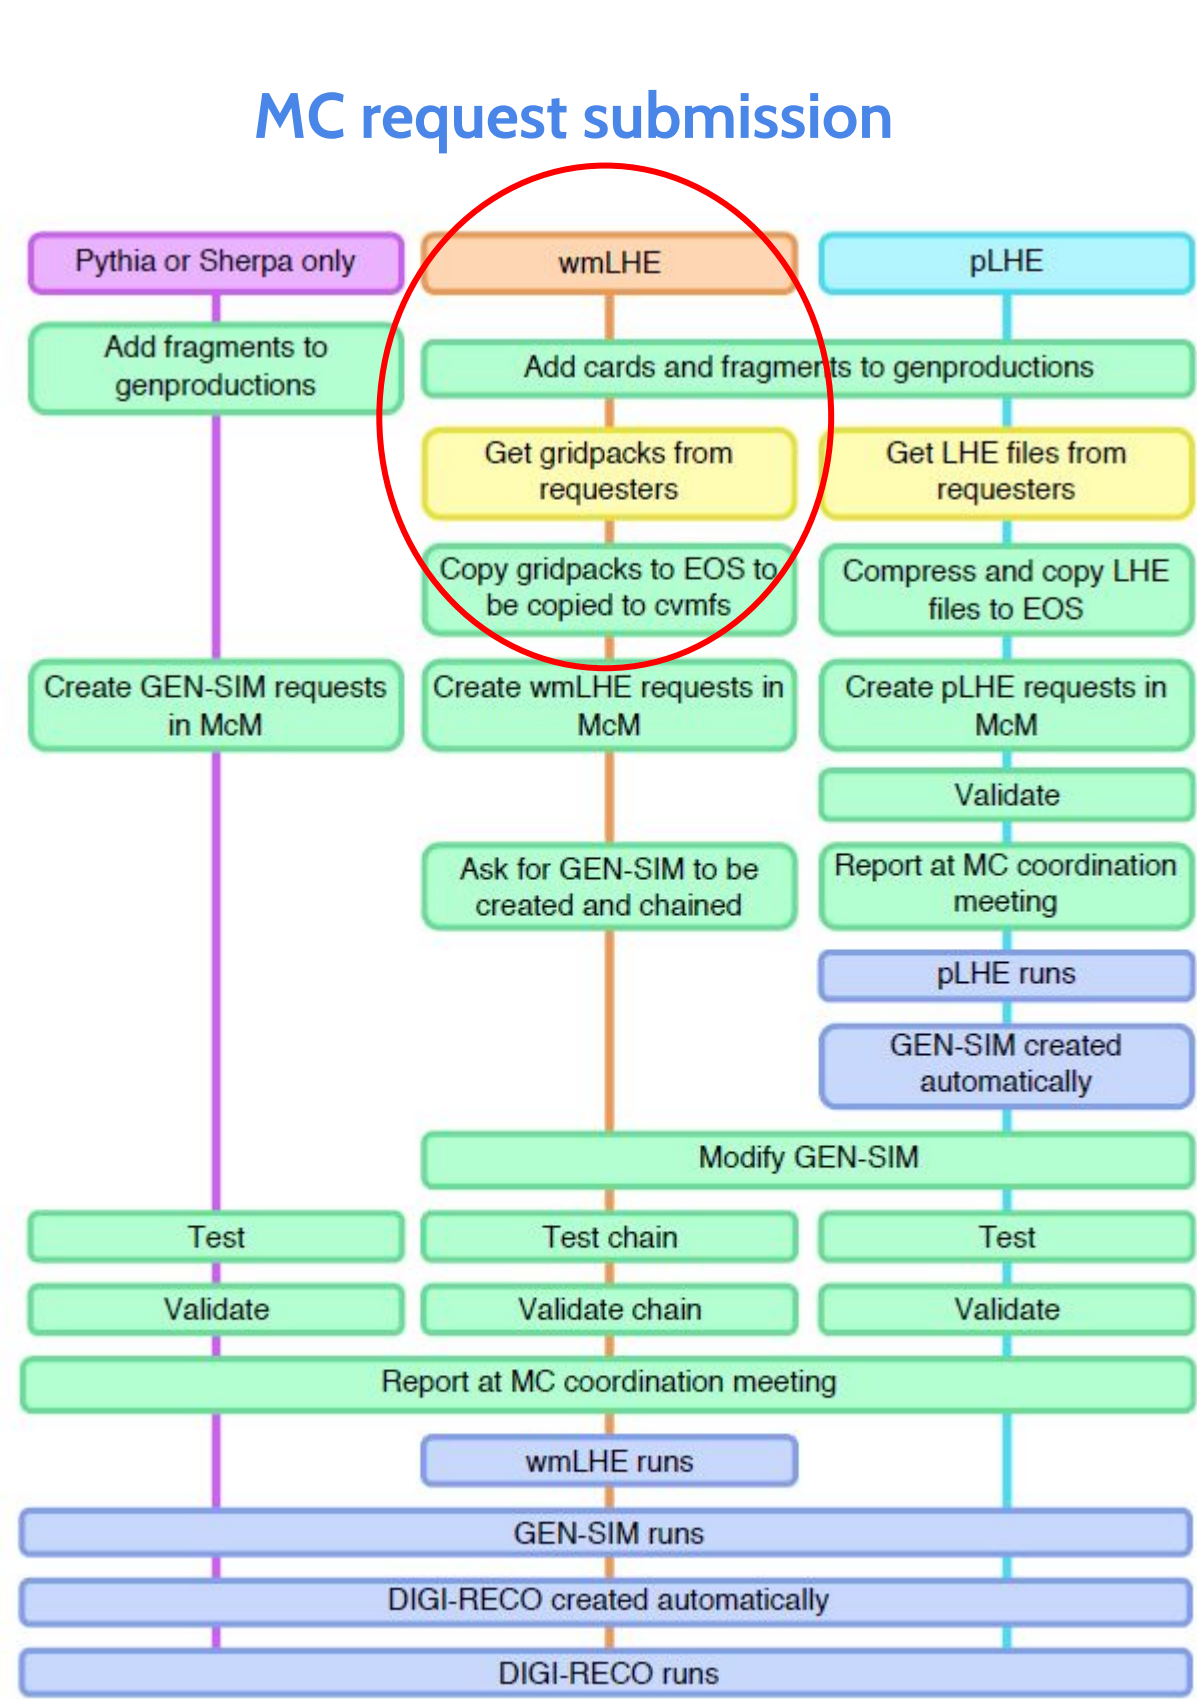


- Better suited for:
  - cmsRun jobs
    - These include a large range of analysis workflows within the CMSSW framework for simulation and data analysis.
    - Random scripts using distributed data sets, which CRAB can handle easily.

- Good for more general purpose workflows. For example:
  - Analysis scripts that users might prefer to handle directly with condor/condor\_dagman.
  - Workflows that already have a batch submission manager (e.g., Madgraph5).
  - Machine learning tools + GPU resources (TensorFlow)



## Generating Madgraph5\_amC@NLO gridpacks



Producing "gridpacks" is one of the very first steps in the Monte Carlo request submission chain.

This kind of workflow however, is not integrated with the CMS standard executable (cmsRun) or CRAB3. A generator's package is used instead to produce these gridpacks in a standalone way, using its own submission managers.


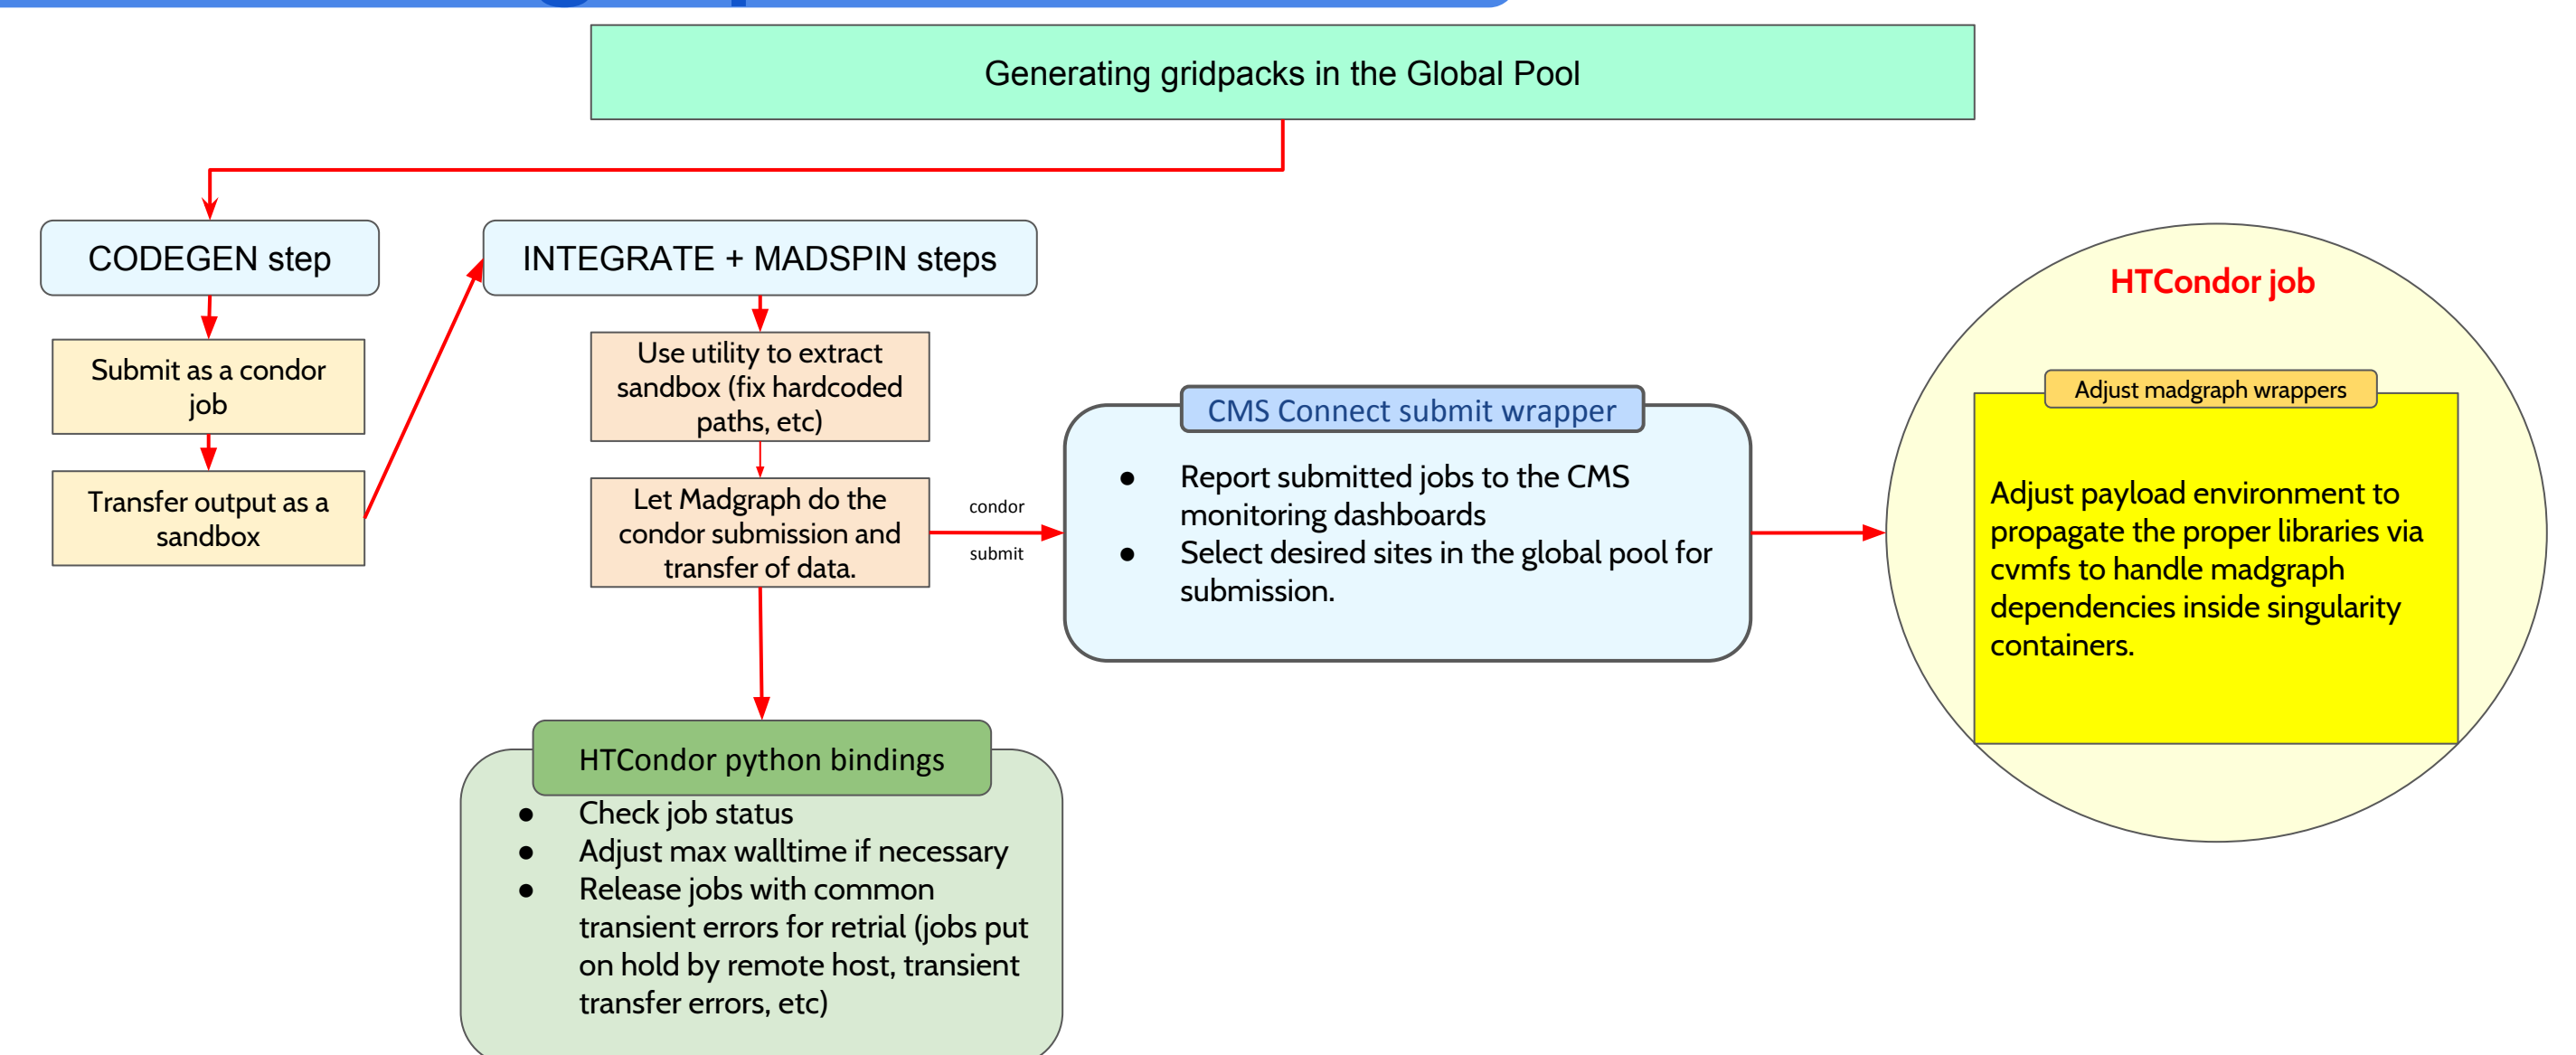
- The Monte Carlo contact persons using different local resources for producing these gridpacks didn't record the usage in CMS dashboards and had to take into account some submission differences for each resource. For example:
- CERN CAF (resources shared between all CERN-based experiments: ATLAS, CMS, ALICE, LHCb, etc):
    - Using the LSF batch system, requires working with different queue types (Ind, 2nw, etc) and submitting the whole master process as a job.
    - Using HTCondor needs to deal with AFS tokens and k5reauth.
  - FNAL LPC CAF uses HTCondor, used by US-based Institutions.

Submitting these gridpacks to the Global Pool offers:

- A uniform layer through HTCondor for these workflows, accessing all resources CMS has access to, like any other workflows handled e.g., via CRAB3.
- The activity is monitored and recorded in CMS monitoring dashboards. (Better accounting).

Some adaptations were made to make madgraph compatible with the way the Global Pool works, including:

- Dynamic adjustment of requested walltime per job when needed.
- Specifying remote CMS Sites for submission.
- Handling typical transient errors translating into jobs getting held.
- Setting specific HTCondor classads for dashboard reporting.

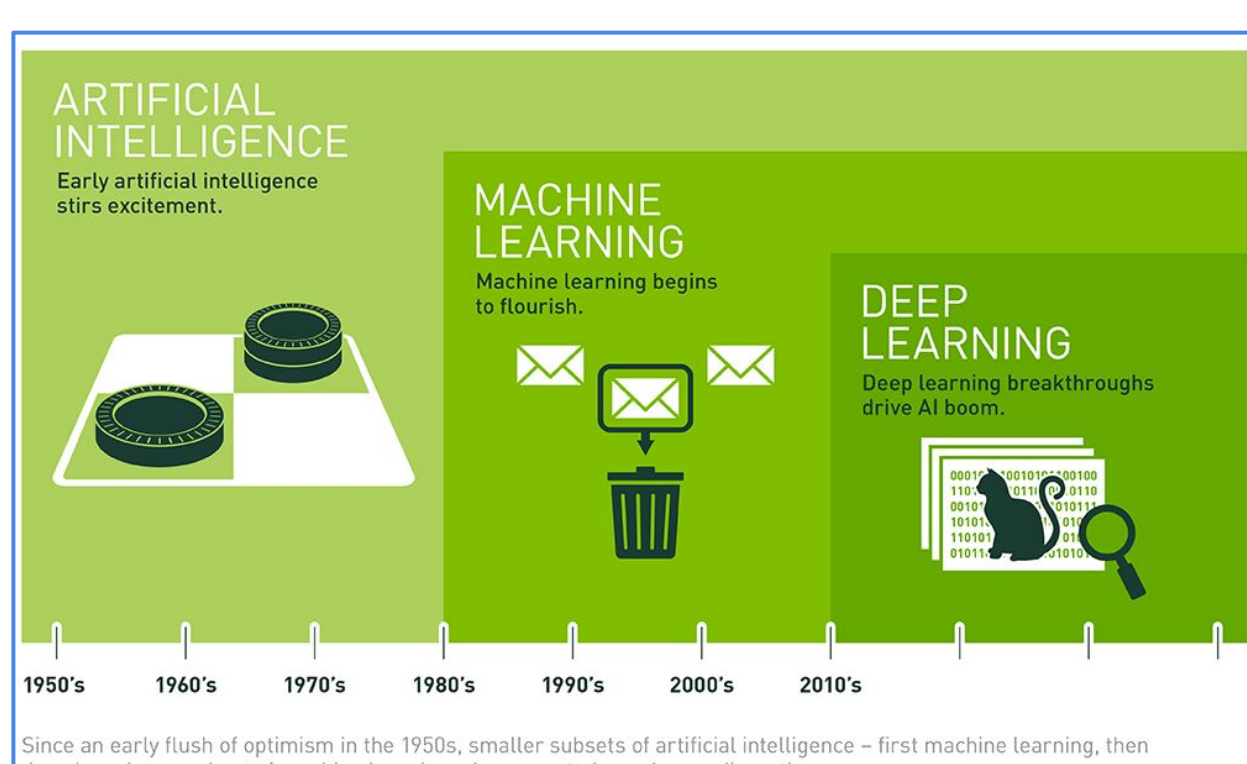


## Deep learning and GPU resources

### Deep learning

Even though Machine Learning (ML) has been a topic for decades and different ML algorithms have been used in high energy physics analysis since the nineties (e.g., Boosted decision trees, random forest, artificial neural network algorithms, etc), the boom in terms of GPU resources demand started with the training of deep neural networks (a subset of ML inspired in artificial neural networks) just few years ago.

Deep learning algorithms involve a fair amount of matrix multiplications and other operations that can be massively parallelized and thus sped up on GPUs, due to the fact that GPUs can have thousands of cores and faster bandwidth to memory. The usage of deep learning algorithms in industry has led to the development of powerful machine learning frameworks, such as TensorFlow<sup>1</sup> (developed by Google), providing APIs for programming languages such as Python and C++, two popular languages in the HEP community. As a result, more progress than even has been made driving machine learning forward.



<sup>1</sup> <https://www.tensorflow.org/>

### Using TensorFlow and GPU resources in the Global Pool

In order to use machine learning tools with Global Pool resources, the framework dependencies (i.e: TensorFlow, Keras, etc) need to be resolved first. CMS works mostly with Red Hat based Operating Systems (6 and 7), but TensorFlow officially supports Ubuntu only, so installing it by hand is not necessarily an easy task for a user.

To help with this, the CMS framework provides such dependencies via CVMFS, but its support is at the CPU-level only, since the integration with GPU resources can get tricky due to potential conflicts with GPU library dependencies. For instance, different TensorFlow versions can require specific versions of cuDNN (the Nvidia Deep Learning SDK) or the CUDA toolkit to work.

To overcome this issue on a wider scale, the OSG builds and maintains Singularity<sup>2</sup> containers based on Ubuntu for TensorFlow with GPU support. But new dependency trees arise: CUDA, CuDNN, GCC, etc. (See OSG oral presentation #49 at this conference for more details<sup>3</sup>).

The CMS Global Pool powered by HTCondor and GlideinWMS has full support for Singularity, so OSG images can be used with the infrastructure through HTCondor in a transparent way.

Condor submit example for requesting GPU resources with OSG TensorFlow images

```
Universe = vanilla

#... (request resources, define logfiles, etc)
Executable = tf_matmul_wrapper_cvmfs.sh
transfer_input_files = tf_matmul.py
request_gpus = 1
Requirements = HAS_SINGULARITY == True
+SingularityImage =
"/cvmfs/singularity.opensciencegrid.org/opensciencegrid/tensorflow-gpu:latest"
Queue 1
```

<sup>2</sup> <http://singularity.lbl.gov/>

<sup>3</sup> E. Fajardo Hernandez et. al., "OSG and GPUs: A tale of two use cases", CHEP 18 Oral Presentation on Track 3, Contribution ID #49

### Using TensorFlow and GPU resources in the Global Pool

