

HPC resource integration into CMS

Computing via HEPCloud

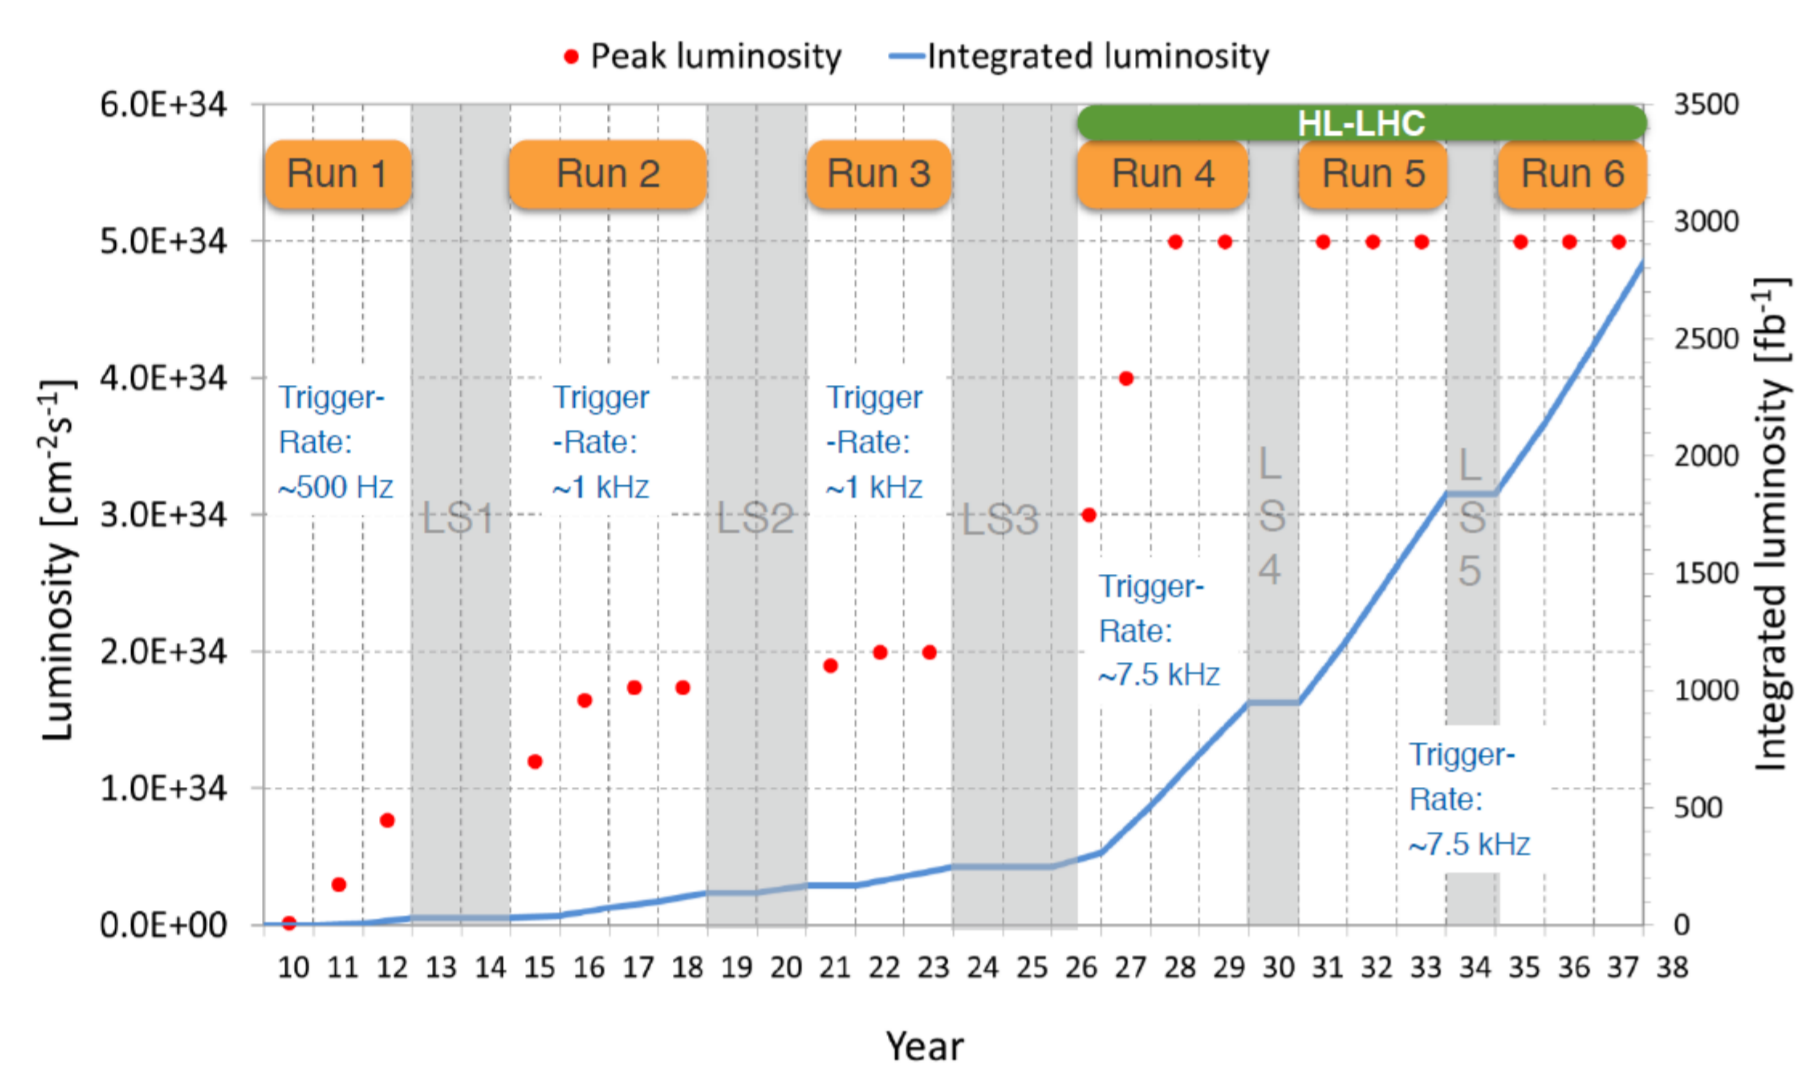
Dirk Hufnagel (FNAL)

for CMS Offline&Computing and FNAL HEPCloud Team




Why HPC ? Why now ?

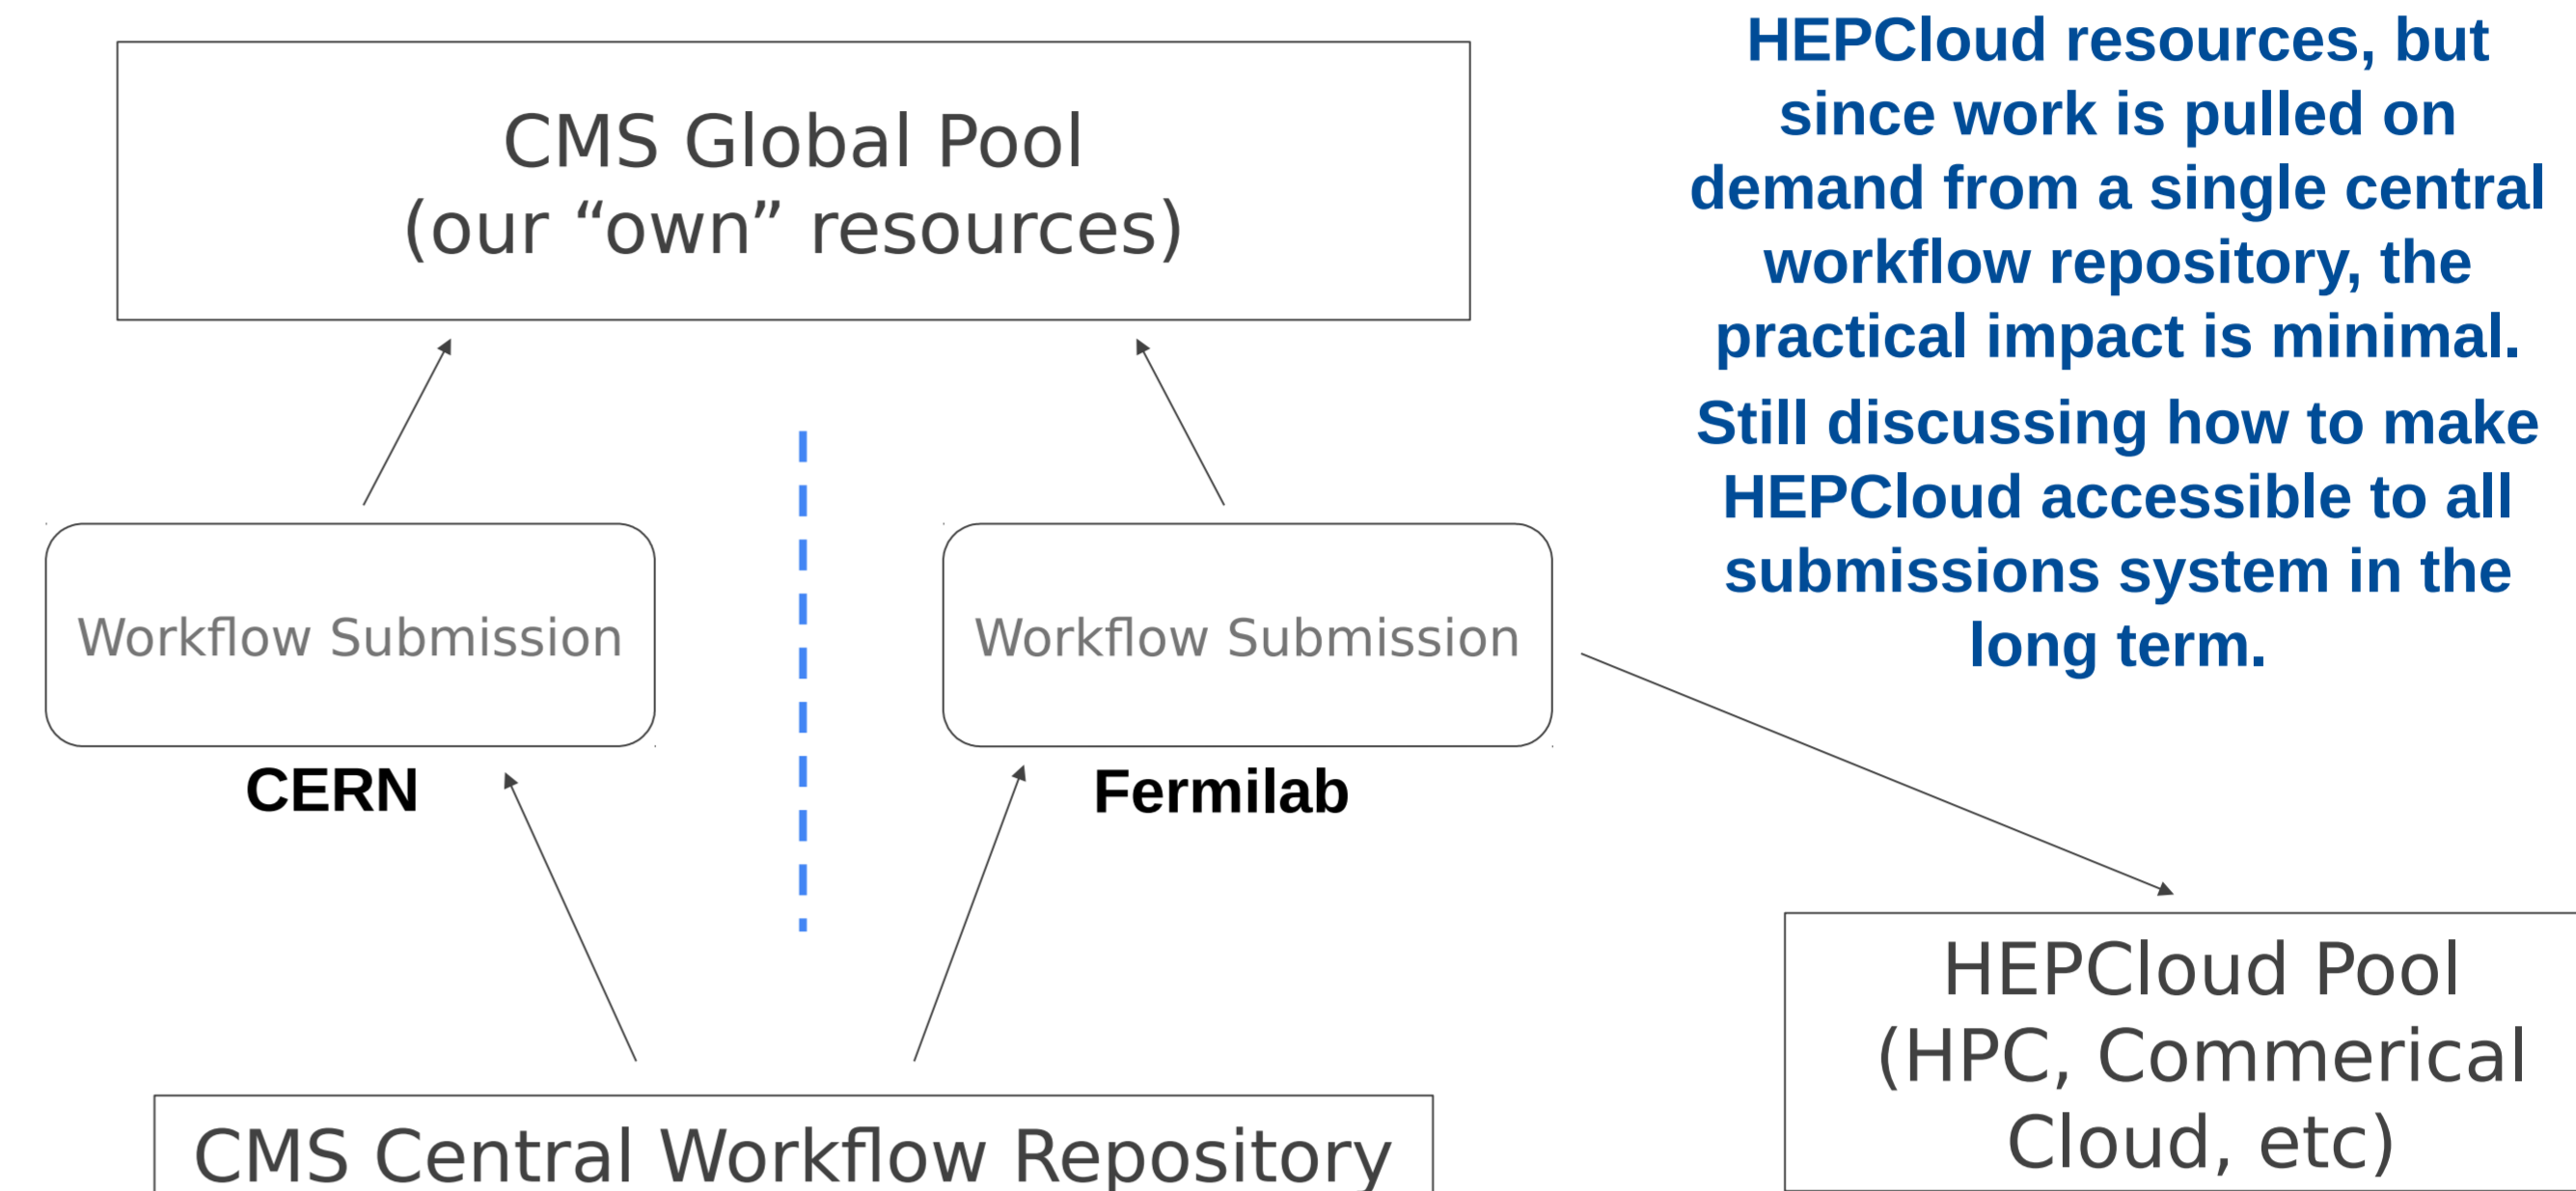
The LHC experiments have their own computing infrastructures which have been successfully used during Run1 and Run2 of LHC. Why are we looking at using HPC sites to complement our own resources now ? The reason is resource extrapolations for the planned LHC upgrade to the HL-LHC.



Trigger rates and past and expected future instantaneous and integrated luminosity of the LHC and HL-LHC.




Extrapolating from the projected luminosities to needed cpu resources for HL-LHC shows dramatic increases which we are unlikely to be able to handle with our own resources. HPC are one possible way to help fill the gap, assuming we can find a way to use them (which given their often non-standard architectures and system design is challenging).



Not all CMS submission systems can access HEPCloud resources, but since work is pulled on demand from a single central workflow repository, the practical impact is minimal. Still discussing how to make HEPCloud accessible to all submissions system in the long term.

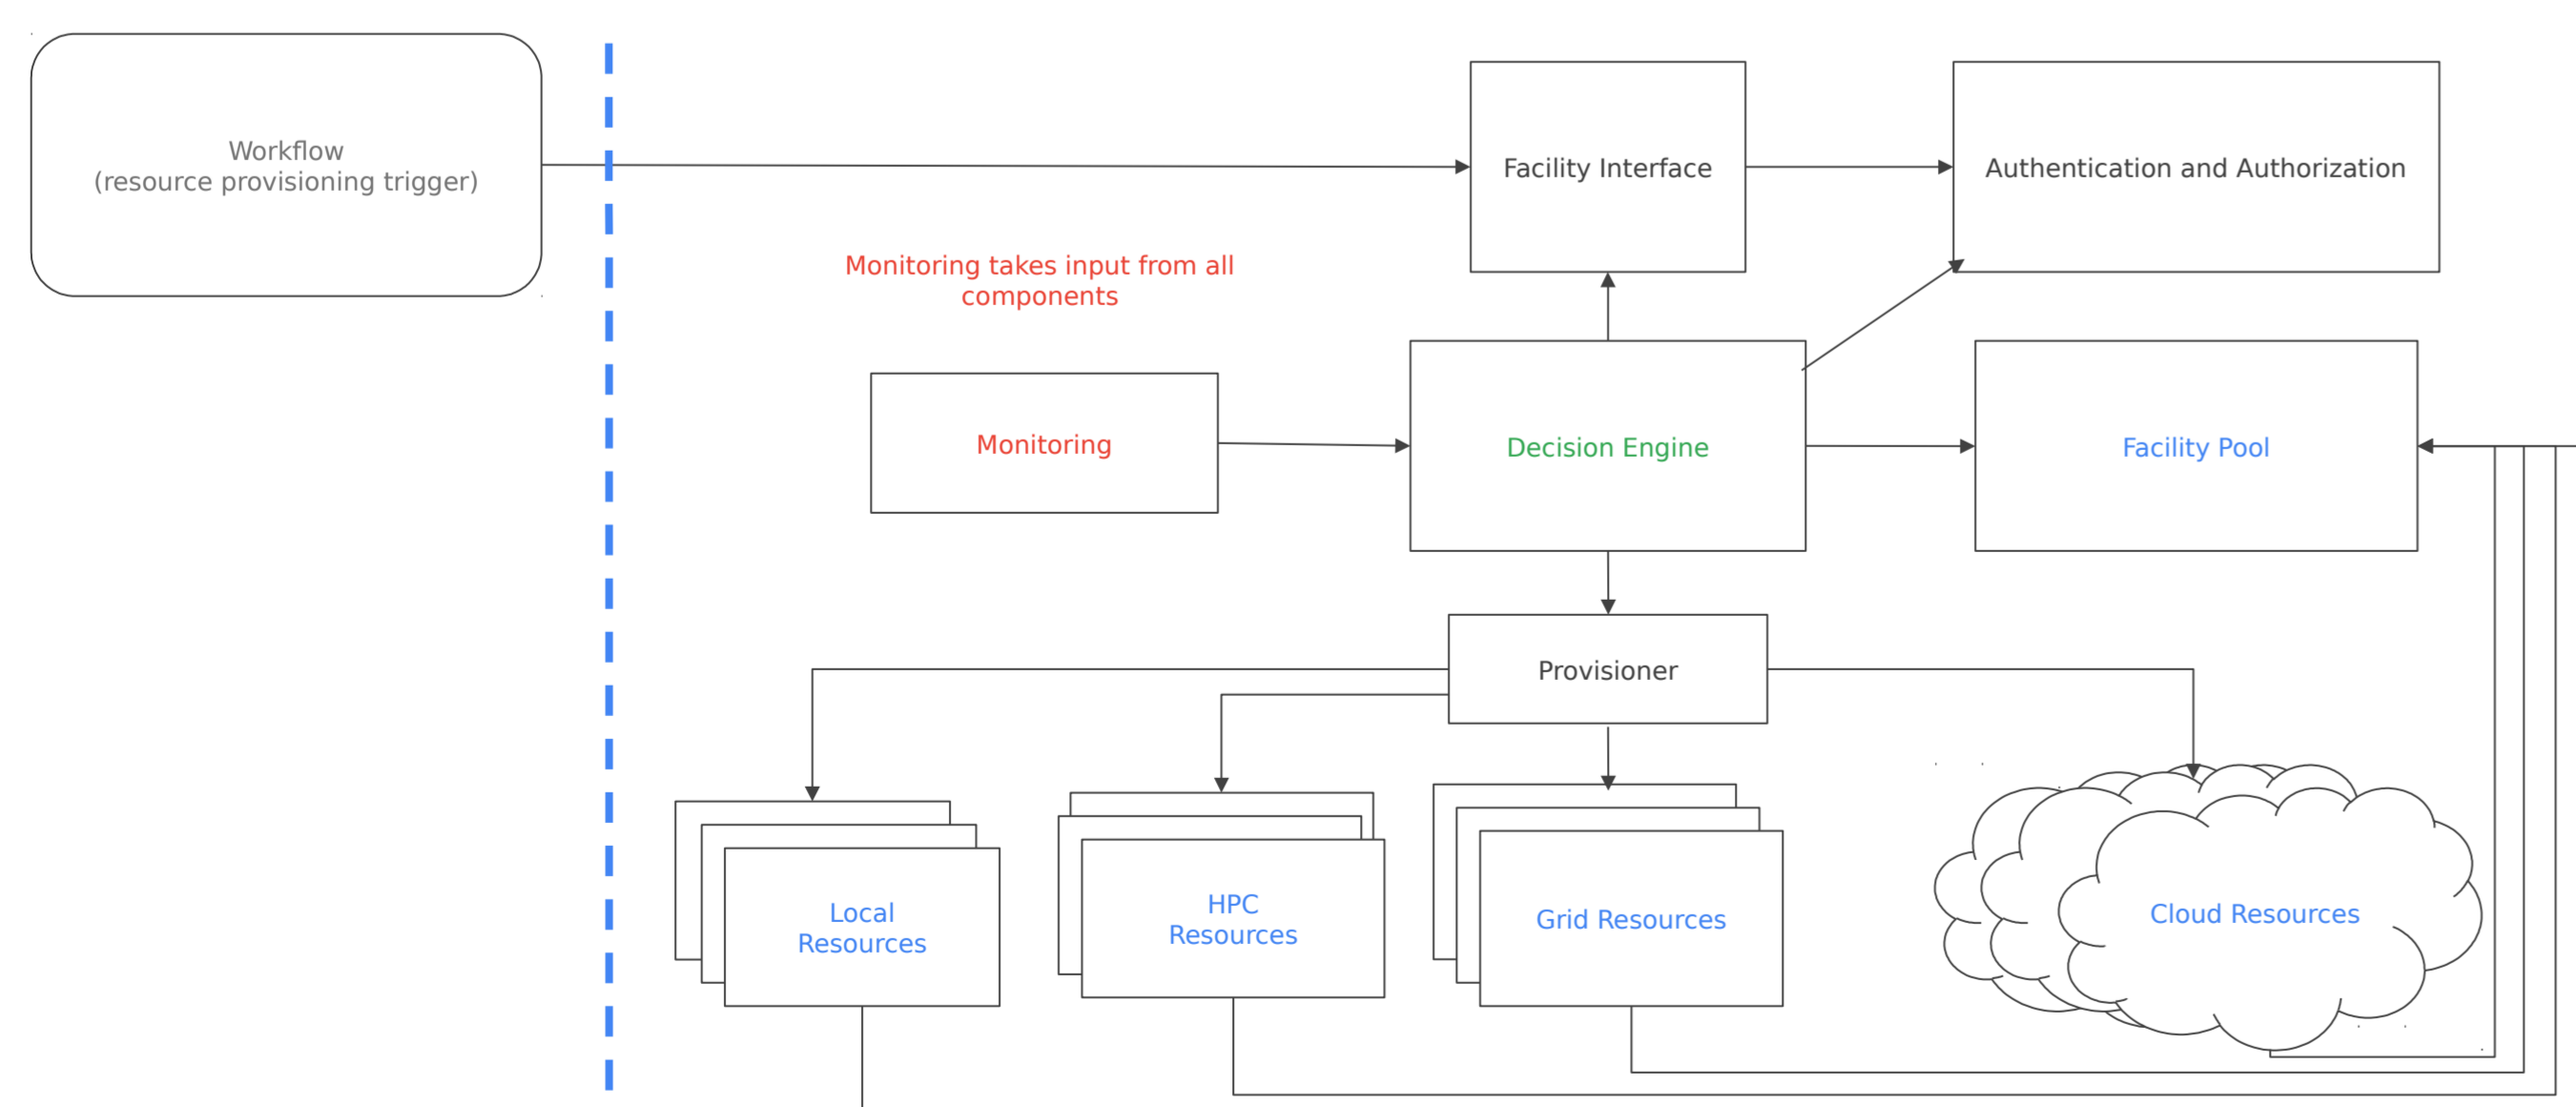
HEPCloud is planned to go into operations later in 2018. CMS has been using the existing HEPCloud prototype to send production jobs to NERSC since May 2018.

Slots (cores) provisioned at NERSC




What is HEPCloud ?

HEPCloud is envisioned as a portal to an ecosystem of diverse computing resources, commercial or academic. It will provide "complete solutions" to users, with agreed-upon levels of service. It will route to local or remote resources based on workflow requirements, cost, and efficiency of accessing various resources. USCMS is planning to use HEPCloud to provide access for CMS to US HPC.



HEPCloud Architecture

CMS has full flexibility in what jobs it can run at NERSC. Due to the tight integration into our production system, (almost) any workflow can end up running jobs at NERSC. The plot to the left shows some of the campaigns and workflows that had jobs at NERSC in the last 2 weeks (not all shown here, but almost 50 workflows in total).



Summary and Outlook

NERSC is currently usable for (almost) all CMS workflows. This means CMS can run it's highest priority workflows at NERSC as needed. If another workflow with even higher priority needs to be run, switching to it is automatic within a short time window (like at any other CMS resource).

Next steps are to increase scale of operations at NERSC and to add more HPC sites to also run CMS production jobs (XSEDE Bridges and Stampede2 are very close).

ACKNOWLEDGEMENT: This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.