



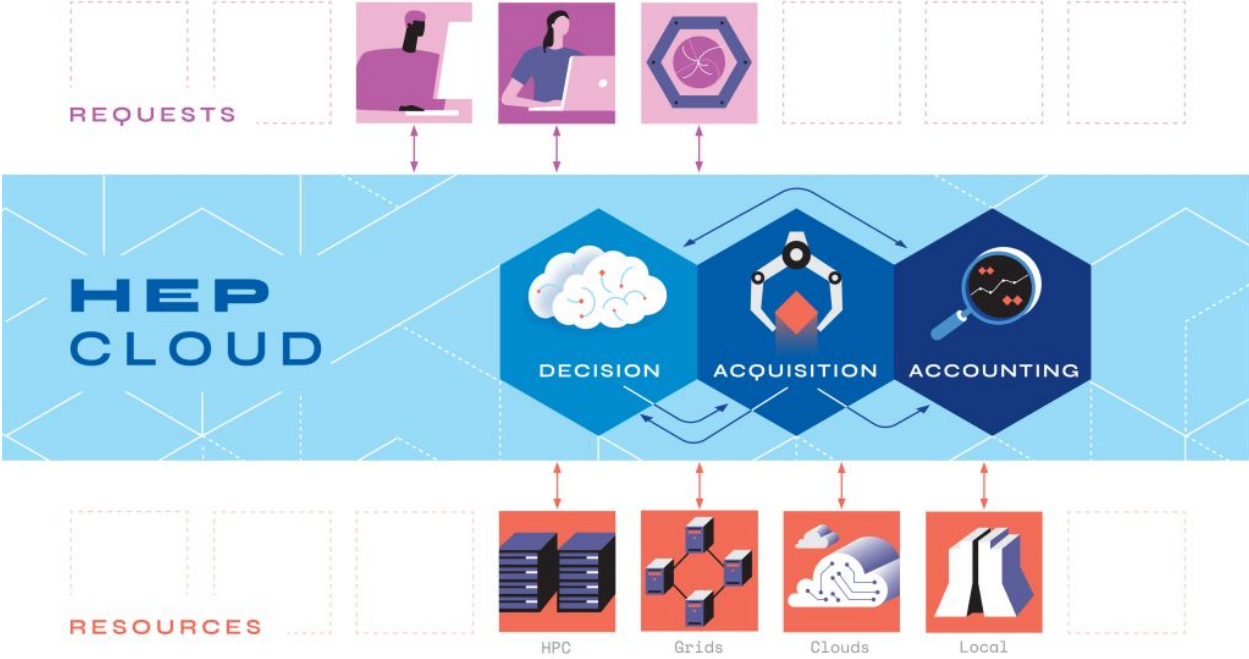
HEPCloud Operations

Vito Di Benedetto

HEPCloud Stakeholders Meeting

September 25th, 2024

HEPCloud <https://hepcloud.fnal.gov/>



HEPCloud Operations review 2023-2024

- HEPCloud infrastructure upgrade:
 - upgrade to DecisionEngine 2.0
 - OS migration to EL9
 - almost complete (thanks Steve Timm!)
- USCMS allocations:
 - NERSC facility
 - ACCESS (XSEDE) allocation program
 - TACC Frontera allocation program
- FIFE (DUNE, NOvA, gm2, mu2e, ...) allocation:
 - NERSC facility
- DUNE (starting on 2024) allocation:
 - NERSC facility


HEPCloud Operations review 2023-2024

- NERSC
 - CORI: 1,600 CPU nodes (Haswell)
9,300 CPU nodes (KNL)
 - Perlmutter: 3,072 CPU nodes (2xAMD EPYC 7763)
1,792 GPU nodes (4xNVIDIA A100 - AMD EPYC 7763)
- ACCESS (XSEDE):
 - TACC Stampade2: 4,200 CPU nodes (KNL)
1,736 CPU nodes (Intel Xeon Skylake)
 - SDSC Expanse: 732 CPU nodes (AMD Rome)
52 GPU nodes (NVIDIA A100 - Intel Xeon 6248 “Cascade Lake”)
 - PSC Bridges2: 504 CPU nodes (2xAMD EPYC 7742)
24 GPU nodes (8xNVIDIA Tesla V100 - 2xIntel Xeon 6248 “Cascade Lake”)
 - Purdue Anvil: 1,048 CPU nodes (2xAMD Milan)
- TACC Frontera: 8,368 CPU nodes (Intel 8280 “Cascade Lake”)
360 GPU nodes (4xNVIDIA Quadro RTX 5000 - Intel Xeon CPU E5-2620 v4)

USCMS HEPCloud Jobs 2023-2024

Claimed Cores by Site.Entry

timeshift -5m



- NERSC-Perlmutter-CPU.CMSHTPC_T3_US_NERSC_Perlmutter_CPU_RHEL8
- NERSC-Perlmutter-CPU.CMSHTPC_T3_US_NERSC_Perlmutter_CPU_SL7
- PSC-Bridges2.CMSHTPC_T3_US_Bridges2
- Purdue-Anvil.CMSHTPC_T3_US_PURDUE-Anvil
- SDSC-ExpansE.CMSHTPC_T3_US_SDSC-ExpansE
- T3_US_NERSC.CMSHTPC_T3_US_NERSC_Cori_SL7
- TACC-Frontera.CMSHTPC_T3_US_TACC_FRONTERA

	min	max	avg	current	total
NERSC-Perlmutter-CPU.CMSHTPC_T3_US_NERSC_Perlmutter_CPU_RHEL8	0	85565	13145	0	8123361
NERSC-Perlmutter-CPU.CMSHTPC_T3_US_NERSC_Perlmutter_CPU_SL7	0	57720	5491	9920	3393334
PSC-Bridges2.CMSHTPC_T3_US_Bridges2	0	24671	2302	0	1422743
Purdue-Anvil.CMSHTPC_T3_US_PURDUE-Anvil	0	8137	2418	0	1494249
SDSC-ExpansE.CMSHTPC_T3_US_SDSC-ExpansE	0	4092	2060	0	1273253
T3_US_NERSC.CMSHTPC_T3_US_NERSC_Cori_SL7	0	65039	2932	0	1812197
TACC-Frontera.CMSHTPC_T3_US_TACC_FRONTERA	0	10961	3100	6272	1915641

USCMS Details - NERSC

-> Details for 2023

- Total allocation was 800,000 NERSC-Hours
 - 1 Perlmutter node (256 threads) = 1 NERSC Hour
 - 1 Cori-Haswell Node (64 threads) = 0.35 NERSC-Hour
- Used a mix of Cori-Haswell (until 2023-04) and Perlmutter
- Total raw node-hours: 1,271,971.28
- Allocation used: 90.6%

-> Details for 2024

- Max length jobs 24h (was 12h)
- Total allocation is 1,318,000 NERSC-Hours
 - 1 Perlmutter node (256 threads) = 1 NERSC Hour
- Total raw node-hours: 929,977.87
- Current usage ~70.6%, the balance can be used by Jan 2025

USCMS Details - ACCESS (XSEDE) and TACC Frontera

-> Details for 2023

- We use 4 ACCESS (XSEDE) resources, (July 1, 2022 - June 30, 2023);
 - TACC Stampede2: 75k Core-Hours; used 100%
 - SDSC Expanse: 7.5M Core-Hours; used 100%
 - PSC Bridges2: 11.4M Core-Hours; used 100%
 - Purdue Anvil: 26.3M Core-Hours; used 100%
- TACC Frontera (Mar 1, 2023 - May 31, 2024): 732k Core-Hours; used 100%

-> Details for 2024


- We use 3 ACCESS (XSEDE) resources, (July 1, 2023 - Sept 30, 2024):
 - SDSC Expanse: 30.3M Core-Hours; used 100%
 - PSC Bridges2: 29.3M Core-Hours; used 100%
 - Purdue Anvil: 33.4M Core-Hours; used 100%
- TACC Frontera (June 1, 2024 - May 31 2025): 640k Core-Hours; used 36%

FIFE NERSC Jobs CY 2023

Claimed Cores by Site Entry

timeshift -5m


NOvA



Claimed Cores by Site Entry

timeshift -5m


DUNE



Claimed Cores by Site Entry

timeshift -5m

mu2e



- Allocation: 77,849.68 NERSC-Hours
- used 54,065.9 (~70%)
- significant allocation fraction used through direct submission
- raw node hours usage:
 - DUNE: 28,505.10
 - NOvA: 16,303.97
 - gm2: 9,316.01
 - mu2e: 897.48

FIFE/DUNE NERSC Jobs CY 2024

FIFE:

- Allocation: 47,032.02 NERSC-Hours
- used: 28,382.6 (~60.3%) so far
- raw node hours usage: 24,695.91
- significant allocation fraction used through direct submission

DUNE:

- Allocation: 45,336.69 NERSC-Hours
- used: 2,687.88 (~6%) so far
- allocation can be used by Jan 2025

Conclusions

USCMS:

- Usage:
 - HEPCloud usage in 2024 has increased from 2023:
 - ACCESS (XSEDE) resources: a factor of ~2
 - NERSC: a factor of ~1.65 (projection)
- Allocation:
 - New allocation for ACCESS program approved, will become active on Oct 1st

FIFE/DUN E:

- Allocation stable over time

Thank you!

Questions?