HPCs in Europe

Andrej Filipcic

Existing EU HPCs

- All that are tagged HPC in dashboards (some US, RU)
- Usage since last October
  - 10k cores on average with peaks up to 30k cores
- Many HPCs participated but none stands out with large resource allocation
- Missing IT4I with 2k cores on average

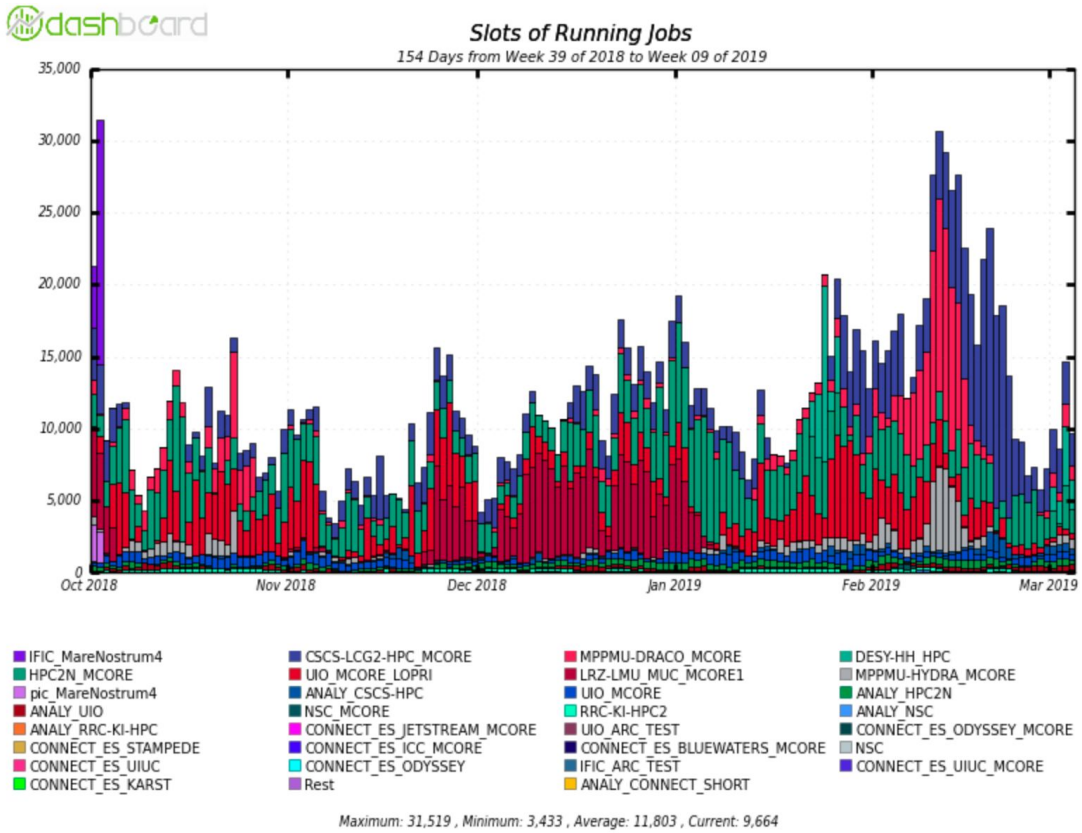

CPU consumption

- In total, ~30 Mhours in 5 months
  - But >50% is on pledged HPCs
- About 15% wallclock in failed jobs

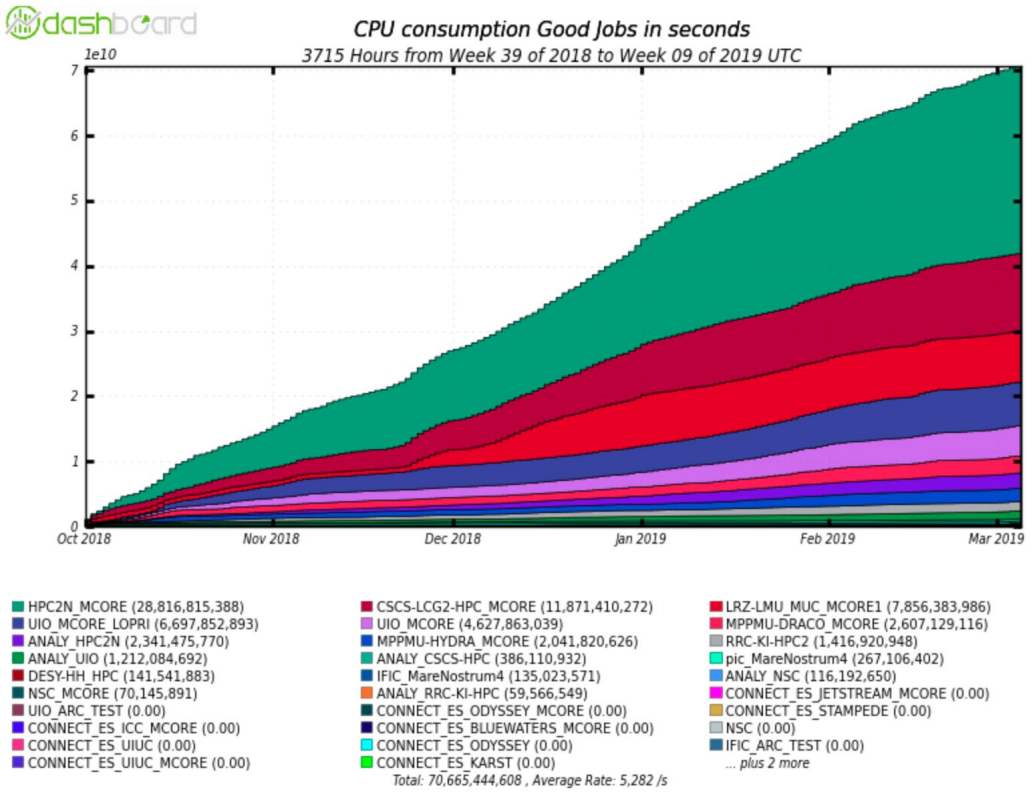

DE

- SuperMuc @ Munich
  - Phase-1 gone
  - Phase-2 no significant allocation yet, single node testing for now - expected to use resources again since April
  - Charliecloud for containers
    - https://hpc.github.io/charliecloud/
- Draco @ Munich
  - Full simulation on up to 50 nodes
- Maxwell @ DESY:
  - Mostly used by XFEL
  - Running ES in backfill mode, preemptive
  - Likely to use more resource when not in datataking mode

ES

- MareNostrum4:
  - Successfully running in 2018, offline now
  - New allocation of 700k Mhours to continue this year
  - Under discussion to use it continuously without allocation
- IFIC Lusitania
  - Successful in 2018
  - Offline now
  - 2M hours allocation

CZ

- IT4I Salomon and Anselm (to be extended)
  - Running steadily up to 2k cores
  - Access to cvms and local squid

CH

- PizDaint
  - a week ago, the Phoenix cluster was completely decommissioned
  - All the CPU resources for ATLAS and other VOs are now on Cray
  - No opportunistic mode
  - 40% of 10k cores pledged to ATLAS, soon to be extended

IT

- Marconi @ Cineca
  - Waiting for approval of the ATLAS application (20Mhours)
  - No resources yet

FR

- IDRIS
  - successful tests in 2018
  - But now offline

EU/PRACE HPC allocations and usage

- Nordugrid: Abisko, Abel, NSC pledge ~5k cores
  - Many more used on Abisko (opportunistic) and Abel (preemptive mode)
  - But Abel is end of life this year, resources will move to cloud
- SuperMuc, Hydra
  - opportunistic/preemptive usage on ~6000 cores, or ~50M hours/year
- MareNostrum4:
  - 700k hours allocated
- IT4I:
  - Backfill usage of 2k cores
- Cineca: 20M hours planned per year

Integration in ADC

- SW distribution:
  - Cvmfs normal mode on pledged, with custom squid on IT4I
  - Cvmfs with parrot (DE)
- All have ARC-CE
  - CE moves data on most
  - Some can run truepilot or pull mode (Maxwell)
- aCT, migrating to Harvester+aCT
  - All tested with harvester
  - Some connectivity issues to CEs with firewall opened to specific aCT machines
- Pledged HPCs run all jobs
  - Analysis, all production types

Prospects with existing HPC centers

- Many were successfully tested in the past, but were not put into steady production
  - Some are small
  - Many cannot provide significant resources without approved allocation
- Allocations in PRACE are limited in time and scope - mostly tuned for single scientists or small groups
  - PRACE allocated 16B hours in total in 10 years, ATLAS uses 3B/year
  - We cannot expect large allocations in near future

EuroHPC Joint Undertaking

- Goals till end of 2020:
  - Build two pre-exascale machines with ~200PFlops each, ~500 Meuro per machine
  - Build 2-4 petascale machines (10-20PFlops), EU budget 30M euro for all
- 25 EU countries are members of the Governing Board
- Exascale: starting 2021
  - Total budget ~2.7B euro for 6 years

Calls for HPCs

- Both pre-exascale and peta-scale calls are out
  - Deadline 4.4 (pre-exa), 14.4 (peta)
  - Selection Q2 2019
- Likely applicants for pre-exa:
  - Finnish consortium: all Scandinavian states + some continent states
  - Italian consortium
  - Spanish consortium
- Applicants for peta-scale
  - ~7 countries will likely apply
  - Some extending already planned centers (LX, SI)

Horizon2020 Research & Innovation calls

- Total budget ~400M euro, 50% contribution from member states, some coming out Q1-Q3 2019
- Extreme Scale Technologies and Applications (90M)
  - Extreme scale computing technologies (hardware, software, methods and algorithms for key applications)
  - HPC applications to ensure European leadership
- Widening HPC skills and use (30M)
  - increase the knowledge and human capital and upraise HPC capabilities, including through the creation of national HPC Competence Centres and their networking and coordination across the Union
  - Federating European supercomputing resources
  - Support to SMEs
- European Processor Initiative (40M)
- Competence centers (2M/country)

Architecture

- Significant funding from most members states
  - Some in procurement
  - Some in-kind, manpower
  - WLCG funding will be affected
- Not clear yet what architecture, possibly a mixture of
  - CPU partitions
  - GPU + FPGA partitions
  - To address intense CPU compute, AI and BigData processing
- The HPC machines need to be on top of TOP500 list
  - Uniform CPU+GPU machines are not at all excluded (Summit), although the CPUs will likely be x86_64

Opportunities

- HPCs will be connected with each other:
  - Big machines will be built by consortia of many countries
  - Opportunity for ATLAS distributed computing institutions to participate in R&I calls, especially data management
- 50% of resources will be managed by JU, 50% by participating countries
  - Participating countries could grant a significant allocation to ATLAS

Conclusions

- HPC resources in EU have steady contribution, but not much growth
- With current resources, the contribution is not expected to grow significantly
- EuroHPC pre-exascale and peta-scale computers will be available end of 2020
  - Some resources even in 2019
- EuroHPC machines might bring
  - More resources for ATLAS
  - Projects for ATLAS software & computing development
  - Some sites might provide the pledge through HPCs