

NUG Webinar

February 23, 2017



Allocations for Cori KNL

Queue Wait Times


Open Q&A

HDF5 and H5py mini-tutorial, Jialin Liu, NERSC Data and Analytics Group

We're using a ZOOM product that only allows broadcast. Please use Q&A feature in ZOOM to type questions.

Cori

9,300 Intel Xeon Phi “KNL” manycore nodes
2,000 Intel Xeon “Haswell” nodes
700,000 processor cores, 1.2 PB memory
Cray XC40 / Aries Dragonfly interconnect
30 PB Lustre Cray Sonexion scratch FS
1.5 PB Burst Buffer



#5 on list of Top 500 supercomputers in the world



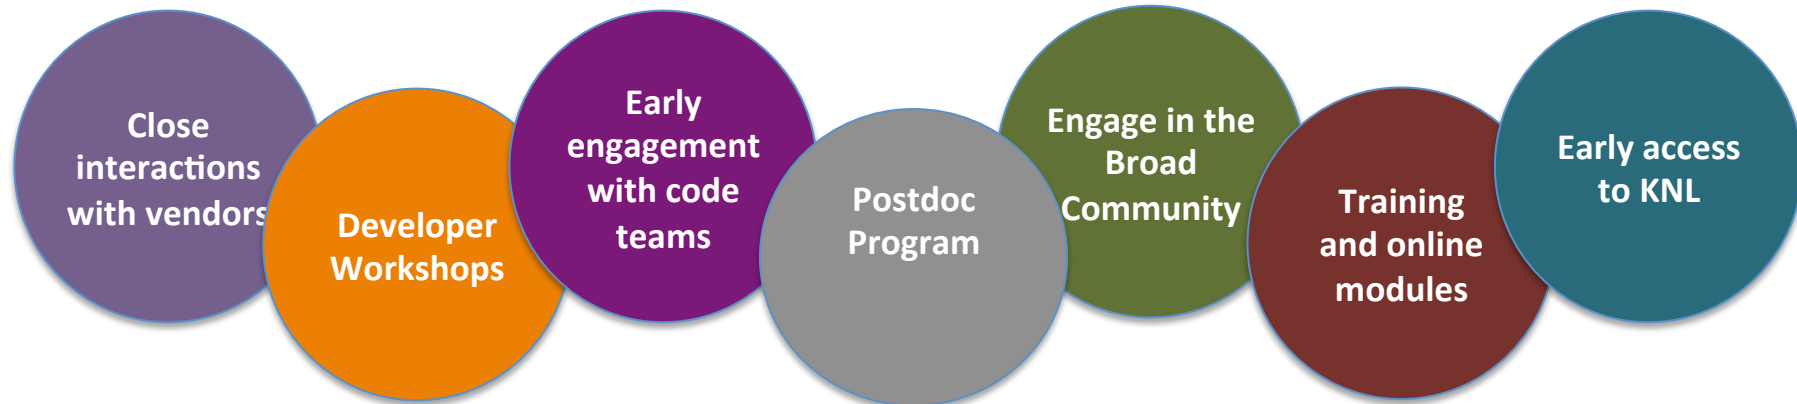
Edison

5,560 Ivy Bridge Nodes / 24 cores/node
133 K cores, 64 GB memory/node
Cray XC30 / Aries Dragonfly interconnect
6 PB Lustre Cray Sonexion scratch FS


Codes need to be prepared to run on Cori Intel Xeon Phi (“KNL”) nodes

NESAP Goal: Prepare DOE Office of Science users for Cori’s manycore CPUs

Partner closely with ~20 application teams and apply lessons learned to broad NERSC user community




Application Readiness: NESAP



ASCR

Almgren (LBNL)
Trebotich (LBNL)


BoxLib
Chombo-crunch



HEP

Vay (LBNL)
Toussaint(Arizona)
Habib (ANL)




WARP & IMPACT
MILC
HACC



NP

Maris (Iowa St.)
Joo (JLAB)
Christ (Columbia/BNL)


MFDn
Chroma
DWF/HISQ



BES

Kent (ORNL)
Deslippe (LBNL)
Chelikowsky (UT)
Bylaska (PNNL)
Newman (LBNL)



Quantum Espresso
BerkeleyGW
PARSEC
NWChem
EMGeo



BER

Smith (ORNL)
Yelick (LBNL)
Ringler (LANL)
Johansen (LBNL)
Dennis (NCAR)


Gromacs
Mercurulous
MPAS-O
ACME
CESM



FES

Jardin (PPPL)
Chang (PPPL)

M3D
XGC1



- NESAP teams are testing code and running science problems up to full scale on Cori KNL nodes with 24 hour time limit.
- All NERSC users are enabled to run on up to 512 nodes for two hours at a time.
- Users can gain access to the full system upon request and demonstration of code readiness for KNL beginning March 1 through the KNL Early Access Program (KNLEAP)
 - Go to <https://my.nersc.gov/knleap.php>
- Usage is not being charged through June 30, 2017.
- System still has a few (software) bugs, which users are helping us discover and solve.

Year	DOE Production (M NERSC Hrs)	ALCC	Directors Reserve
2014	2,400	300	300
2015	2,400	300	300
2016	2,400	300	300
2017	4,800	600	600
	2,400 (January)	150 (January)	300 (January)
	2,027 (July for KNL)	450 (July)	300 (July)
2018	~6,800	~850	~850

- Additional hours given to program (DOE allocation) managers on June 1, 2017
- 2.027 billion additional hours for use on Cori KNL.
- No new ERCAP request needed. Allocation managers will distribute based on existing 2017 requests. (You can submit a new request at any time.)
- At this time, we are planning to have NERSC hours fungible among Cori KNL, Cori Haswell, Edison. Additional allocation must go only to projects that can run on KNL.
- NERSC will send pre-production KNL usage report by repo to DOE on June 1.
- Charging on KNL starts July 1.
- NERSC will monitor and take action if hours are being used on Edison or Haswell.

The currency for charging at NERSC is the the “NERSC Hour”.
1 NERSC Hour ~ 1 core hour on Hopper (retired 2015)

Node Type	Cores per Node	Charge per Node per Hour
Cori KNL	68	96*
Cori Haswell	32	80
Edison	24	48
Hopper (retired)	24	24

We also have the ability to incentivize (or not) usage with queue charge factors.

* - Production value TBD

Science Area	KNL NERSC Hours (M)
Lattice QCD	200
Materials Science	87
Application Readiness	58
Climate Research	46
Astrophysics	23
High Energy Physics	20
Biosciences	15
Geosciences	12
Computer Science	11

- Large job discount created extra allocation: ~500 M hours equivalent
- Cori downtime in September for KNL node integration
- Placed huge pressure on Edison, resulting in long wait times
- Payoff will be ~2X hours in 2017 and ~3X in 2018.

Edison Queue Backlog

Queue Backlog Over Time



Queue Wait Relief

The Problem


40% discount → 500 M Hours extra in 2016
 500 M Hours extra → Long queues
 Big jobs in 2017 → Cori KNL

Addressing the Problem

Edison big job discount:
 40% → 20% on March 1
 20% → 0% on July 1

Propose eliminating “low”
 Projects out of time run free in “scavenger”

Edison Big Job (>16K cores) Usage



- Cori KNL nodes are free through July 1, 2017
- Charging on Cori KNL nodes starts July 1, 2017
- Additional allocation distributed to PMs on June 1
- Program managers will distribute additional time
- Codes need to be ready to use the Xeon Phi and program managers need to consider readiness in allocation decisions
- NERSC will send KNL usage report and other advice to DOE
- Participate in the KNLEAP program
 - <https://my.nersc.gov/knleap.php>