NERSC is the mission HPC facility

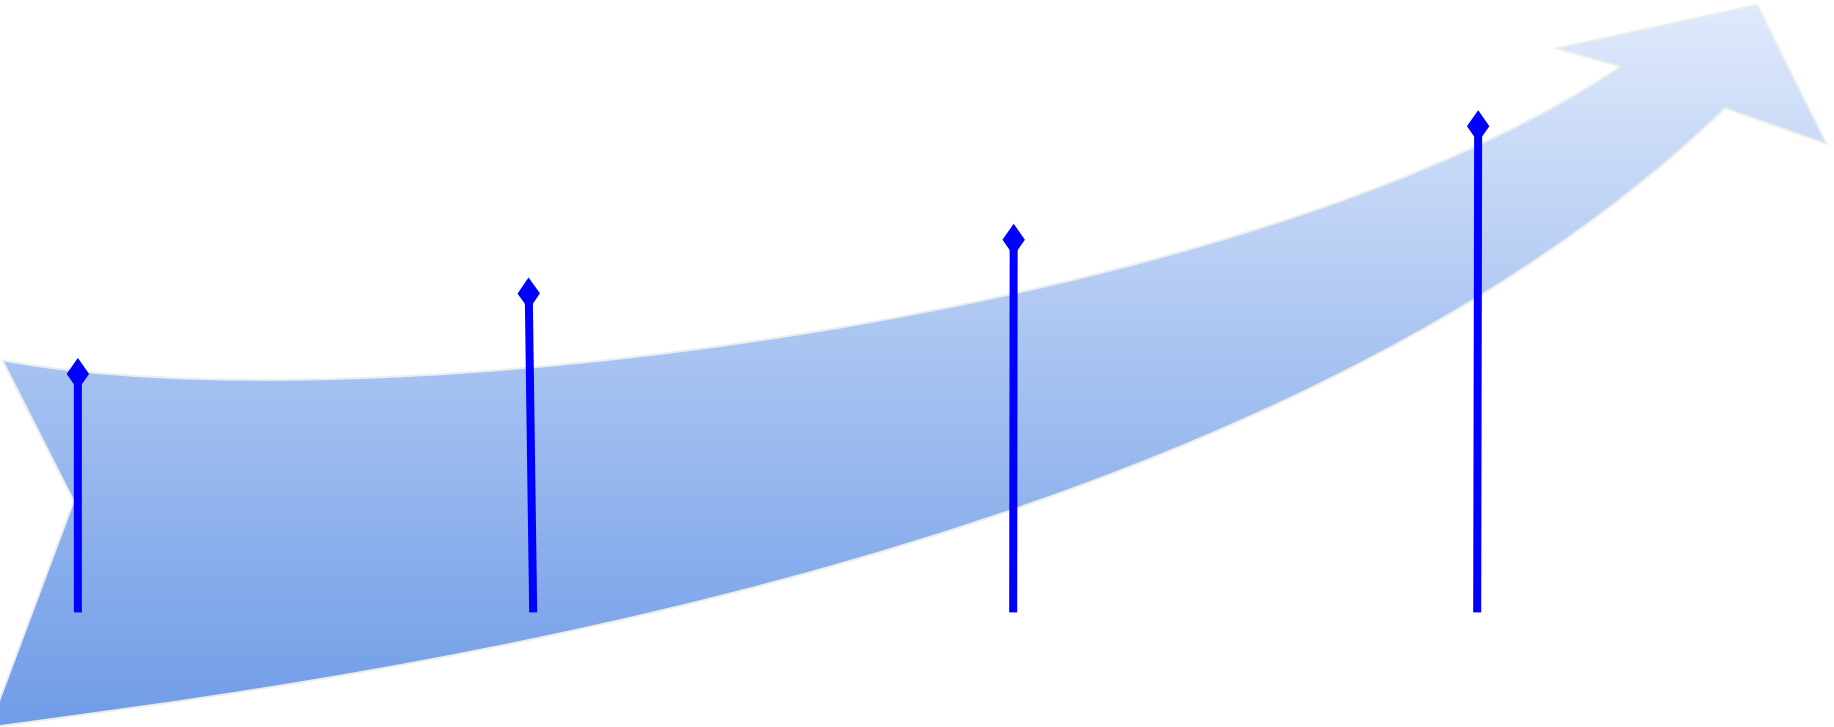

NERSC-9 Project Major Scope Items

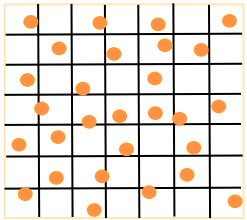

NERSC-9: A System Optimized for Science

From the start NERSC-9 had requirements of simulation and data users in mind

NERSC Exascale Scientific Application Program (NESAP)

- Prepare DOE SC users for advanced architectures like Cori
- Partner closely with ~20 application teams and apply lessons learned to broad NERSC user community.

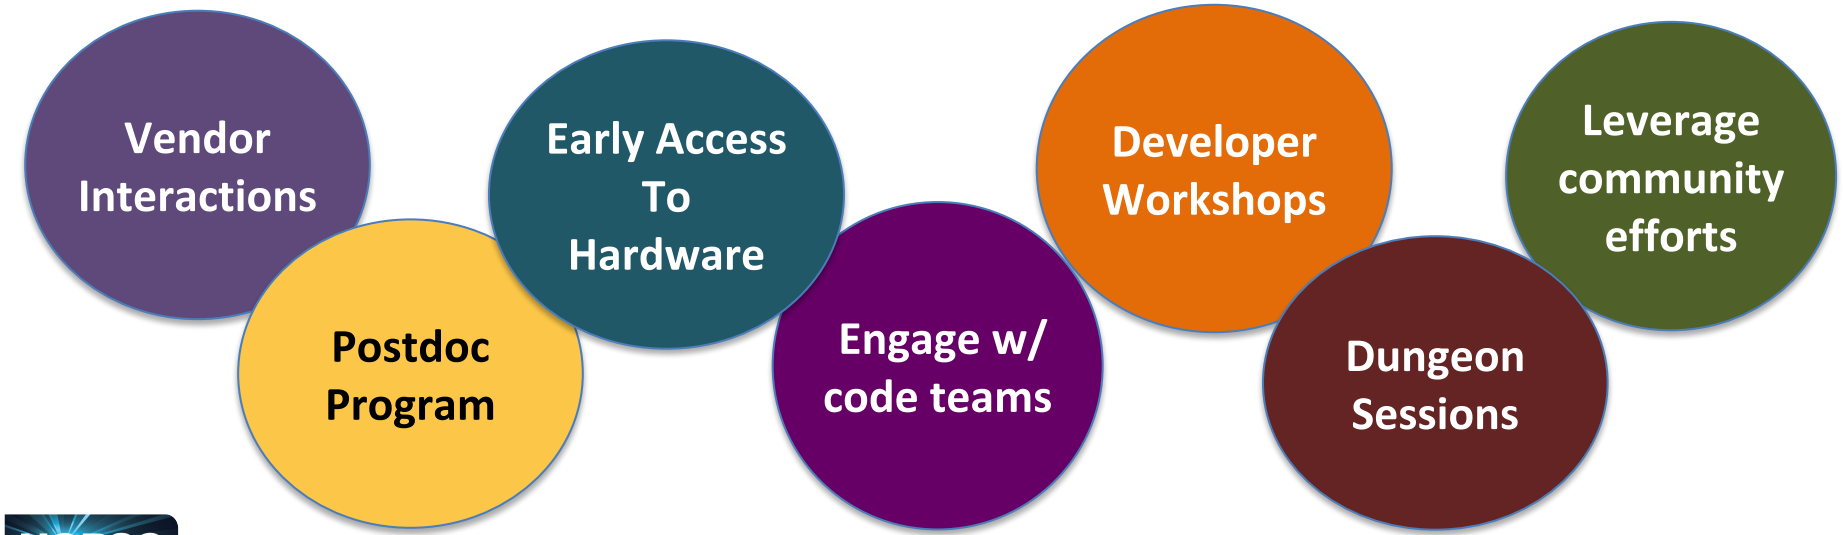

Transitioning From KNL to AMD Processors

Codes optimized on Xeon Phi (KNL) will run well on Perlmutter

Many KNL architecture features are present on Perlmutter CPUs

- Many-Core
- MPI+OpenMP Programming Model Will Continue

Easier Onramp to "Many-Core" with Perlmutter CPUs than with KNL

- More Traditional Cores
- Single Memory Technology

NERSC already supports a large number of users and projects from DOE SC's experimental and observational facilities

Palomar Transient Factory
Supernova

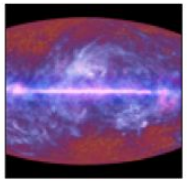

Planck Satellite
Cosmic Microwave Background
Radiation

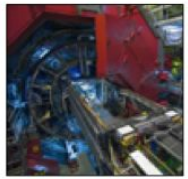

Alice
Large Hadron Collider

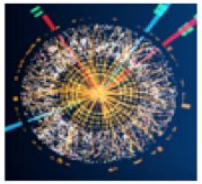

Atlas
Large Hadron Collider

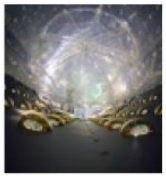

Dayabay
Neutrinos

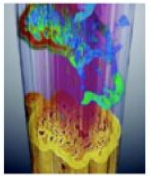

ALS
Light Source

LCLS
Light Source

Joint Genome Institute
Bioinformatics

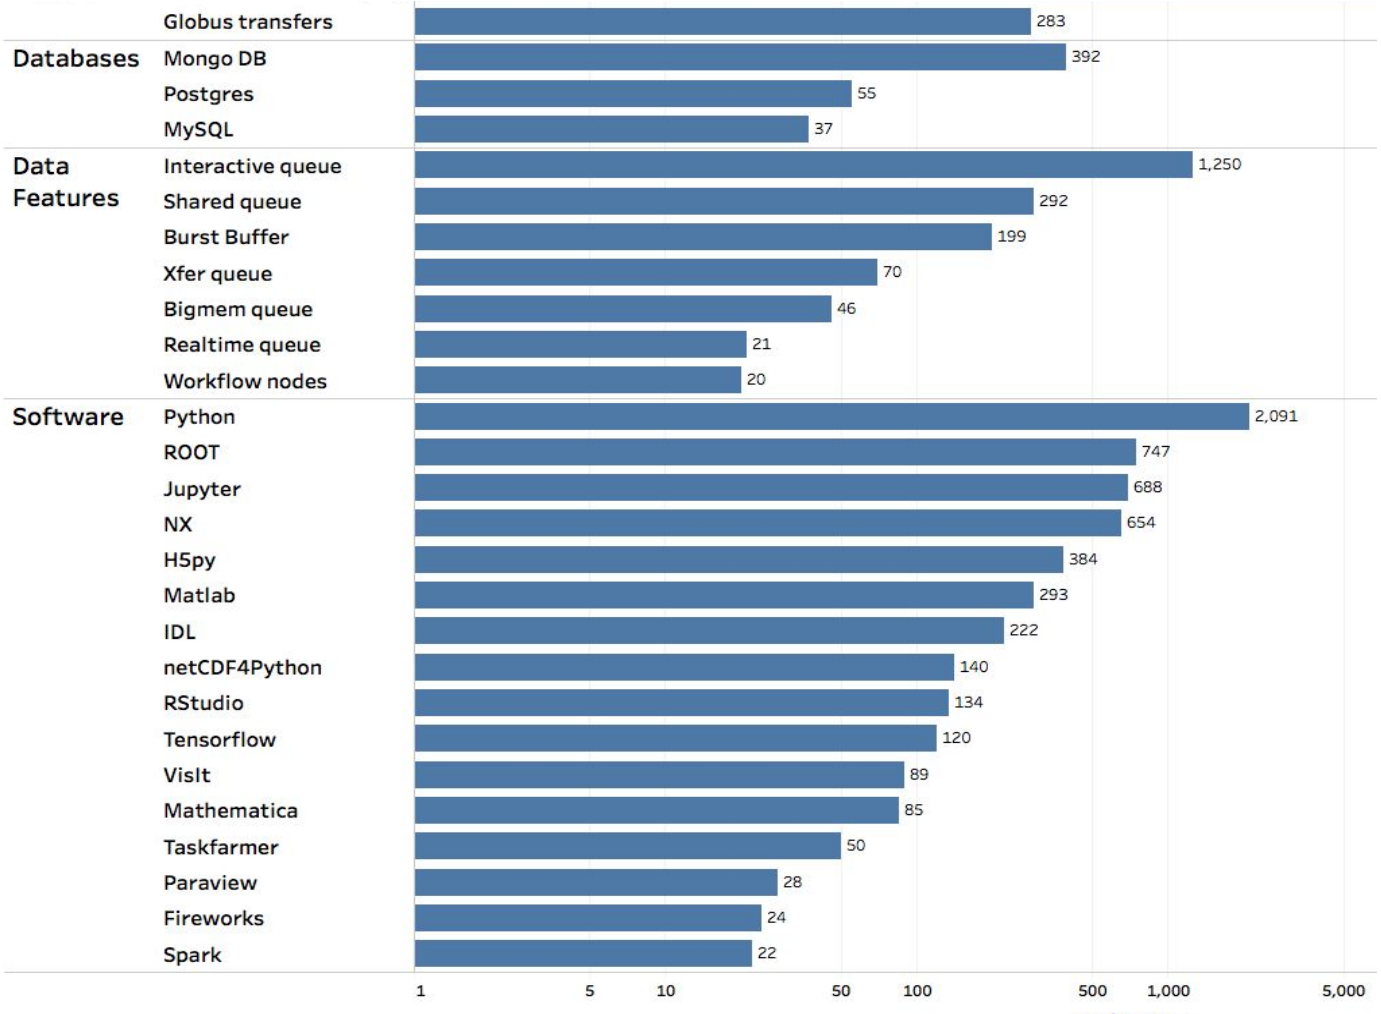
Globus transfers 283
Databases
Mongo DB 392
Postgres 55
MySQL 37
Data Features
Interactive queue 1,250
Shared queue 292
Burst Buffer 199
Xfer queue 70
Bigmem queue 46
Realtime queue 21
Workflow nodes 20
Software
Python 2,091
ROOT 747
Jupyter 688
NX 654
H5py 384
Matlab 293
IDL 222
netCDF4Python 140
RStudio 134
Tensorflow 120
VisIt 89
Mathematica 85
Taskfarmer 50
Paraview 28
Fireworks 24
Spark 22
1
5
10
50
100
500
1,000
5,000

Data Features	Cori experience	N9 enhancements
I/O and Storage	Burst Buffer 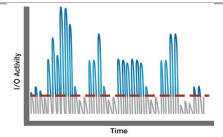	All-flash file system: performance with ease of data management 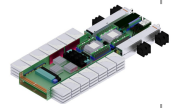
Analytics - **Production stacks** - **Analytics libraries** - **Machine learning**	 User defined images with Shifter NESAP for data New analytics and ML libraries 	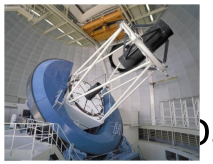 Benchmark Production Analytics workflows. Data apps in NESAP at outset Optimised analytics libraries and deep learning application benchmarks
Workflow integration	 Real-time queues	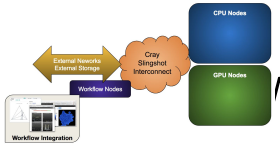 SLURM co-scheduling Workflow nodes integrated
Data transfer and streaming	SDN 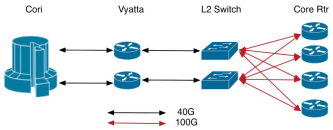	Slingshot ethernet-based converged fabric 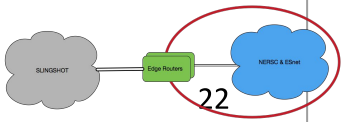

GPU Programming Models

We will support and engage our user community where their apps are:

CUDA: MILC, Chroma, HACC ...

CUDA FORTRAN: Quantum ESPRESSO, StarLord (AMREX)

OpenACC: VASP, E3SM, MPAS, GTC, XGC ...

Kokkos: LAMMPS, PELE, Chroma ...

Raja: SW4

OpenMP

Engaging around Performance Portability

NERSC is working with PGI/NVIDIA to enable OpenMP GPU acceleration

NERSC Hosted Past C++ Summit and ISO C++ meeting on HPC.

NERSC Will Pursue Membership in 2018

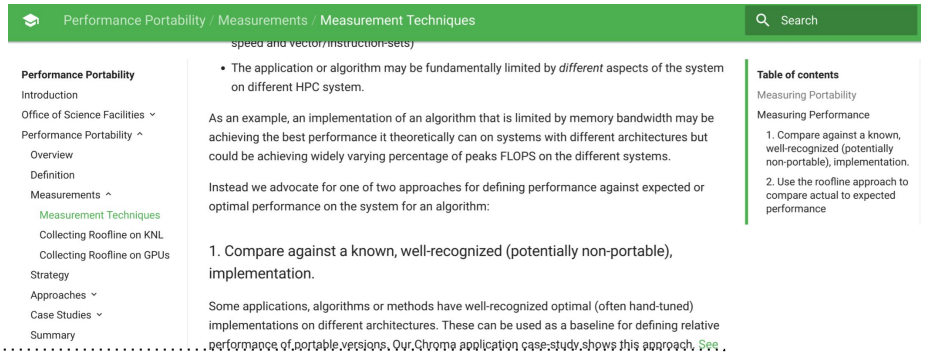

NERSC is leading development of performanceportability.org

Doug Doerfler Leading Accepted Performance Portability Workshop at SC18. and 2019 DOE COE Perf. Port. Meeting