


MVAPICH2 at NERSC



Shahzeb Siddiqui (LBNL)
Sameer Shende (University of Oregon)
Prathmesh Sambrekar (LBNL/Arizona State University)

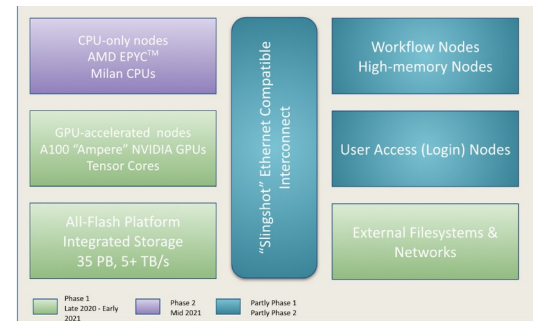
- National Energy Research Scientific Computing Center (NERSC) is scientific computing facility as part of the [Office of Science](#) in the [U.S Department of Energy](#)
- At NERSC we support up to 8000+ users that span a wide range of disciplines including climate modeling, simulation of universe, high energy physics experiments, protein simulation, and many more.
- NERSC most recent supercomputer called [Perlmutter](#) is named after astrophysicist [Saul Perlmutter](#) for his Nobel Prize in Physics for discovering the rate of expansion of the universe



Perlmutter



- Perlmutter is a HPE Cray Shasta platform composed of CPU-only and GPU accelerated nodes.
- The system will be delivered in two phases
- The system contains 1536 GPU accelerated nodes with NVIDIA A100 GPUs and 3072 CPU-only nodes powered with AMD EPYC Milan CPUs
- Perlmutter is currently ranked #7 on the [Top500 list](#) with a peak performance of 93.75 PFlop/s



7	Perlmutter - HPE Cray EX23Sn, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-10, HPE DOE/SC/LBNL/NERSC United States	761,856	70.87	93.75	2,589
---	---	---------	-------	-------	-------

NERSC

- FIGURE 23.** How to access e4s/22.02 stack built with mvapich2-gdr

MVAPICH2 Modules on Perlmutter



- We have modulefiles for mvapich2 on shared filesystem, one must add the directory to MODULEPATH since we don't expose these modules to all users

```
~/ module use /global/cfs/cdirs/m3896/shared/modulefiles
~/ module av

----- /global/cfs/cdirs/m3896/shared/modulefiles -----
mvapich2/3.0a
```

Common Errors during spack builds



- We have noticed a few spack packages fail to build with cray-mpich due to proper support with cray compilers.
- Some packages expect MPI wrappers such as mpicc but on Cray we expect to use cray wrappers **cc**, **CC**, **ftn** when invoking MPI programs

```
# configure: error: Failed to find C MPI Wrapper. see https://cdash.spack.io/buildSummary.php?buildid=104940  
# - rempi@1.1.0
```

```
#- mpifileutils@0.11.1 ~xattr # failed to install libcircle Unable to find suitable MPI Compiler. Try setting MPICC.
```

MVAPICH2 Installation on Perlmutter



- MVAPICH2 3.0a supports SLURM and Slingshot 11 network interfaces on both CPU and GPU nodes of Perlmutter
- To use MVAPICH2:
 - module use /global/cfs/cdirs/m3896/shared/modulefiles
 - module load mvapich2/3.0a
 - Use mpicc, mpicxx, mpif90 as compilers
 - mpicc foo.c -o foo
 - Use srun to launch the binary
 - srun -n 4 ./foo

Building E4S on Perlmutter



- **The Extreme-scale Scientific Software Stack (E4S) is installed on Perlmutter using the mvapich2/3.0a module.**
- **For more details regarding this deployment, please refer to E4S documentation page <https://e4s.readthedocs.io/en/latest/deployment.html#perlmutter>**

- module use /global/cfs/cdirs/m3896/shared/modulefiles
- module avail e4s
e4s/22.05/mvapich2-3.0a ...

- module load e4s/22.05/mvapich2-3.0a
- module avail

E4S 22.05 built with MVAPICH2/3.0a



sameer@perlmutter:login24:~> module avail

```
----- /global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a-slurm/spack/share/spack/lmod/cray-sles15-x86_64/mvapich2/3.0a-es35auw/Core -----
adios/1.13.1                fortilinos/2.0.0                nccmp/1.9.0.1                py-warpx/22.05-dims3
adios2/2.8.0-cuda80         globalarrays/5.8                nco/5.0.1                   py-warpx/22.05-dimsRZ      (D)
adios2/2.8.0                (D)                             hdf5/1.10.7                 omega-h/9.34.1             scr/3.0rc2
amrex/22.05                 heffte/2.2.0-cuda80            openpmd-api/0.14.4          slate/2021.05.02-cuda80-openmp
arborx/1.2                  heffte/2.2.0                   (D)                         slate/2021.05.02-openmp    (D)
axom/0.6.1-openmp          hpx/1.7.1-cuda80               parsec/3.0.2012             slepc/3.17.1-cuda80
butterflypack/2.1.1         hpx/1.7.1                     (D)                         slepc/3.17.1              (D)
cabana/0.4.0               hypre/2.24.0                  petsc/3.17.1                (D)                         strumpack/6.3.1-openmp
caliper/2.7.0-cuda80        kokkos-kernels/3.6.00-cuda80   precice/2.4.0               sundials/6.2.0
caliper/2.7.0              (D)                             pumi/2.2.7                 tasmanian/7.7-openmp
darshan-runtime/3.3.1      libquo/1.3.1                  py-cinemas/1.7.0           tau/2.31.1-cuda
datatransferkit/3.1-rc3    mercury/2.1.0                 py-libensemble/0.9.1       veloc/1.5
dyninst/12.1.0-openmp      metall/0.20                   py-petsc4py/3.17.1
faodel/1.2108.1            mfem/4.4.0                    py-warpx/22.05-dims2
```

```
----- /global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a-slurm/spack/share/spack/lmod/cray-sles15-x86_64/openmpi/4.1.3-gw3a4bv/Core -----
gptune/3.0.0
```

```
----- /global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a-slurm/spack/share/spack/lmod/cray-sles15-x86_64/Core -----
aml/0.1.0                  flux-core/0.38.0-cuda          legion/21.03.0-cuda80-cuda   plasma/21.8.29
archer/2.0.0               flux-core/0.38.0              (D)                         legion/21.03.0             (D)
argobots/1.1              gasnet/2022.3.0               magma/2.6.2-cuda80          py-jupyterhub/1.4.1
bolt/2.0                   ginkgo/1.4.0-cuda80-openmp     mpark-variant/1.4.0        qthreads/1.16
chai/2.4.0                 ginkgo/1.4.0-openmp           (D)                         mvapich2/3.0a              (D)
charliecloud/0.26         gmp/6.2.1                     nrm/0.1.0                  swig/4.0.2-fortran
cmake/3.23.1              (D)                             gotcha/1.0.3               nvhpc/22.3
darshan-util/3.3.1        kokkos-kernels/3.6.00-openmp  (D)                         papi/6.0.0.1-cuda
flit/2.1.0                kokkos/3.6.00-openmp          pdt/3.25.1
                           (D)                         raja/0.14.0-cuda80-openmp  superlu/5.3.0
                           (D)                         nvhpc/22.3                swig/4.0.2-fortran
                           (D)                         umap/2.1.0
                           (D)                         zfp/0.5.5-cuda80
```

145 products built using MVAPICH2

NERSC

1: adios2	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/adios2-2.8.0-kif4mquytlh4cybngfiuux4fn3ivrt
2: aml	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/aml-0.1.0-5gdjl3ickdtb27qq62kopfpdpxmxi3t
3: amrex	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/amrex-22.05-ft3rksivlgipquadvsh7gxuis4izq6m
4: arborx	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/arborx-1.2-tc416d3k6pipywty6shbq2wba44tp
5: archer	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/archer-2.0.0-pa33i3csxzhls5t5nfcw64taxisr7wt
6: argobots	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/argobots-1.1-nsmph2wbkkrus4zx4jemezbabx7qgr7y
7: axom	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/axom-0.6.1-gdjebwss54lr2capdigrhh5zt6xd4u3
8: bolt	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/bolt-2.0-7hf4fziabmplaj7hryhps5cozunhyv
9: butterflypack	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/butterflypack-2.1.1-ooghq5wo5xszwmojow647rscnzwmsv
10: cabana	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/cabana-0.4.0-l2bj6kvhvxud3xabnr3p37270q2d6b5v
11: caliper	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/caliper-2.7.0-svgn26f3urqah5626zeftyf2mbrkkipis
12: chai	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/chai-2.4.0-w6v54m7r4r2hxaawg63yrfj55aikkxqx
13: charliecloud	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/charliecloud-0.26-czgnnskv4k7jrfjcy3ndy62lwk44skb
14: conduit	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/conduit-0.8.3-wkjkgns0x3zjt63gg3heby3kp72btef
15: darshan-runtime	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/darshan-runtime-3.3.1-44lbnty4ftuccbf7gowl262ztdvkoxvxt
16: datatransferkit	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/datatransferkit-3.1-rc3-livrocqm2xowtjvtvyqoj5a5r6k06tuec7
17: dyninst	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/dyninst-12.1.0-wickbpnislxl6iztkpnnigte4jtiiych
18: faodel	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/faodel-1.2108.1-isf6lqapcman3r73lk64d5jaacngw3j
19: flit	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/flit-2.1.0-bhhsfeffdpfuewy53ztrrqr53vsn12
20: forttilinos	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/forttilinos-2.0.0-zlzmeycesulqwidlgwm2juplygs4kg
21: gasnet	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/gasnet-2022.3.0-bwmz3brqu4oqctgs61bkeztmxyqtxjm
22: ginkgo	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/ginkgo-1.4.0-qvks5kj03gpqzslompq27ddbjd31rjtd
23: globalarrays	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/globalarrays-5.8-rhohdqk3vg3icaoxnt6q46wia3z13lbwp
24: gotcha	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/gotcha-1.0.3-s2xkdnmdusbrxwuhssvqgrcsh7rnhucn
25: gptune	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/gptune-3.0.0-k4q6zfkhtcseol46lvbhek54ulwgrw6
26: hdf5	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/hdf5-1.12.2-bubavn15ec5lmfxc3d3c3g4jtjtm7fsk
27: heffte	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/heffte-2.2.0-wgagrczlssnki6aqtcmqpsr2jtyktjl
28: hpx	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/hpx-1.7.1-xuw726mdztobnu7ap16jowdhgrc6vgz4
29: hypre	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/hypre-2.24.0-2okdvdy7gxaanlhtwnvum2uv03ntjmhw
30: kokkos	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/kokkos-3.6.00-7eud56ps23wkwkoi27272f7q645v1
31: kokkos-kernels	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/kokkos-kernels-3.6.00-lcqcx55m2chvnyizpfboscsivhialepe
32: lammps	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/lammps-20220107-uk74xi2eio4yvgwk5e4rxws5coervpz
33: legion	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/legion-21.03.0-53f3b6ytkfw15fnux5yz3uka3b15cshr
34: libquo	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/libquo-1.3.1-5ju66szpfhrwpkdb6o2lwhhkitrnwy5v
35: magma	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/magma-2.6.2-zwuh6mm62fc50e6ypvkmgajbizz4hkzq
36: mercury	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/mercury-2.1.0-1eykx65oexkjk6xh4nf5mjgh40q7tcho
37: metall	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/metall-0.20-ngcmficzomrownqbogxvinm4tz5exdfas
38: mfem	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/mfem-4.4.0-iyvynznmxdoxacwoynhfbz5ru2clxwve



U.S. DEPARTMENT OF
ENERGY

Office of
Science



E4S products built using E4S (contd.)



39: mpark-variant	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/mpark-variant-1.4.0-uvr4bft4rni7i254ctevurkikuv36z6x
40: nccmp	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/nccmp-1.9.0.1-pwjdebjxx7bk4glw573b4lckjfybuk5g
41: nco	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/nco-5.0.1-dtzaouk2bid2kqiyjzqodu773x75sqvj
42: netlib-scalapack	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/netlib-scalapack-2.2.0-dhiu3z27cn5fq13djmaya2q56sudyuyve
43: nrm	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/nrm-0.1.0-7dphpnstxxwv7tqfjqtneym2v5syq50k
44: omega-h	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/omega-h-9.34.1-mpoe4z1gqnls7hpavj6quk62luzom5lw
45: openmpi	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/openmpi-4.1.3-gw3a4bvon2tyeg55baai7pvedk45c42
46: openpmd-api	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/openpmd-api-0.14.4-dx3hlgzye2mdfuvlfovujwggkhkdzodgn
47: papi	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/papi-6.0.0.1-o5pse7e7jjala2y5bwrreiv6uto1dk4mz
48: papyrus	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/papyrus-1.0.2-twr7purknxy6ympux6xdqnfmlau5l5rq
49: parsec	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/parsec-3.0.2012-vuxyiv73cbnbg3vicfodyu4j5s53he6n
50: pdt	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/pdt-3.25.1-4wzbhsadvu6upir6eowsqby6mdf2op3a5jy
51: petsc	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/petsc-3.17.1-v5gxr3ud1padu6b24sy5mgyf73xfmynwb
52: plasma	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/plasma-21.8.29-6wzczv2jqqn23s7gc2ujxfq3yvzy4ckb
53: precice	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/precice-2.4.0-v7wscgidow3wlmn4chqgxxow4kgjdhrfn
54: pumi	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/pumi-2.2.7-f7gzmwq5b4l1uvtfxccqchqdbfbf2edf
55: py-cinemas	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/py-cinemas-1.7.0-j75bb2inxrx5piawm2wy4uhb4jfrbesf
56: py-jupyterhub	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/py-jupyterhub-1.4.1-jysyxh242wzampidb5r355lx7bd4tixwb
57: py-libensemble	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/py-libensemble-0.9.1-ppigxh7jxampyupqilgt2vjed4m3x3zyn
58: qthreads	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/qthreads-1.16-s32g45lb506207fmlkcbby4frrzvp6pug
59: raja	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/raja-0.14.0-ztcuu26vl3seibwdi6pfwx22qryfluki
60: scr	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/scr-3.0rc2-koc07bok3bqz445j524qwikzh7v
61: slate	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/slate-2021.05.02-wkdngoku4cicegrntu4pu4p5j5u7o3b5
62: slepc	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/slepc-3.17.1-t5dplaiout4kzutufoegqpnka5voxw2n
63: strumpack	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/strumpack-6.3.1-vdjhqejekftpmxz3kr2ndp7anc32om6
64: sundials	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/sundials-6.2.0-dn7f771anhsyvu7g2ybegqfbeatpige
65: superlu-dist	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/superlu-dist-7.2.0-h2asqgfihih3p3kws2clrf3z3yzkqnv5p
66: swig	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/swig-4.0.2-fortran-orersnc6ghyazliftwa6qvk2lhttpcxe
67: sz	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/sz-2.1.12-dkeddzkzqogrlazcwrs0a67icas3nc
68: tasmanian	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/tasmanian-7.7-55rvccldbgb53cwkpry72bexdqkvxvf
69: tau	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/tau-2.31.1-x75mktmygps6ew5mps7em4mvaauivgq3
70: trillinos	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/trillinos-13.0.1-yft4mfzoyqauhl3kmbqm2cmj3srwpw
71: umap	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/umap-2.1.0-5ryuhm3tfsqtf3fvd3aydegjexob2d
72: umpire	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/umpire-6.0.0-4wkon17m7w53j554ak7juxwz26fhrv
73: veloc	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/veloc-1.5-3zk6uqg36fu7ctmqnurnroof2bnykz3m
74: zfp	/global/cfs/cdirs/m3896/shared/ParaTools/E4S/22.05/mvapich2-3.0a/spack/opt/spack/cray-sles15-zen3/gcc-11.2.0/zfp-0.5.5-jha3y2d7gajabhm2ut6qfwlq3fwtieir

Running OSU Benchmark with MVAPICH2



- We have developed a test to run the OSU benchmark for the mvapich2 installation.
- Test are run via [buildtest](#), an HPC testing framework that will automate build and execution of test
- Test results are pushed to CDASH via buildtest

```
buildspecs:
  mvapich2_mpi_omb_tests:
    type: script
    executor: perlmutter.slurm.debug
    description: Run OSU microbenchmarks
    tags: [gpu, mpi]
    sbatch: ["-N 2", "-t 5", "-G 4", "-C gpu", "-A m3503.g"]
    run: |
      module use /global/cfs/cdirs/m3896/shared/modulefiles
      module load mvapich2/3.0a
      srun -n 2 osu_latency
      srun -n 2 osu_bw
      srun -n 2 osu_bibw
```

Test: mvapich2_mpi_omb_tests (Passed)
Build: gpu (perlmutter) on 2022-08-18 05:09:31
Labels: mpi, gpu

compiler	
description	Run OSU microbenchmarks
endtime	2022/08/17 22:10:32
hostname	login27
starttime	2022/08/17 22:09:31
user	e4s
Return Code	0
Test ID	4dc3d3df-1101-4943-b370-2d8c8b64ce8
build_script	/global/cfs/cdirs/m3503/buildtest/runs/perlmutter_check/2022-08-17/perlmutter.slurm.debug/mvapich2-omb/mvapich2_mpi_omb_tests/4dc3d3df/mvapich2_mpi_omb_tests_build.sh
buildenv	/global/cfs/cdirs/m3503/buildtest/runs/perlmutter_check/2022-08-17/perlmutter.slurm.debug/mvapich2-omb/mvapich2_mpi_omb_tests/4dc3d3df/build-env.txt
command	bash --noprofile --nohistory mvapich2_mpi_omb_tests_build.sh
errfile	/global/cfs/cdirs/m3503/buildtest/runs/perlmutter_check/2022-08-17/perlmutter.slurm.debug/mvapich2-omb/mvapich2_mpi_omb_tests/4dc3d3df/mvapich2_mpi_omb_tests_err
executor	perlmutter.slurm.debug
id	4dc3d3df
logpath	/global/cfs/cdirs/m3503/ci-data-dir/iddig0/Static/00155091/NERSC/buildtest-nersc/buildtest-varlog/buildtest_inmqmq.log
outfile	/global/cfs/cdirs/m3503/buildtest/runs/perlmutter_check/2022-08-17/perlmutter.slurm.debug/mvapich2-omb/mvapich2_mpi_omb_tests/4dc3d3df/mvapich2_mpi_omb_tests.out
schemafile	script.schema.json
stagedir	/global/cfs/cdirs/m3503/buildtest/runs/perlmutter_check/2022-08-17/perlmutter.slurm.debug/mvapich2-omb/mvapich2_mpi_omb_tests/4dc3d3df/stage
tags	gpu mpi
testpath	/global/cfs/cdirs/m3503/buildtest/runs/perlmutter_check/2022-08-17/perlmutter.slurm.debug/mvapich2-omb/mvapich2_mpi_omb_tests/4dc3d3df
testroot	/global/cfs/cdirs/m3503/buildtest/runs/perlmutter_check/2022-08-17/perlmutter.slurm.debug/mvapich2-omb/mvapich2_mpi_omb_tests/4dc3d3df

View GitLab CI results

#	OSU MPI Bi-Directional Bandwidth Test v5.9	
#	Size	Bandwidth (MB/s)
1		2.37
2		4.80
4		9.65
8		19.13
16		38.54
32		76.35
64		154.04
128		306.16
256		556.54
512		1092.87
1024		2189.58
2048		4386.47
4096		8648.64
8192		15222.25
16384		23030.97
32768		28904.61
65536		34304.40
131072		37953.22
262144		38336.29
524288		40981.84
1048576		41960.33
2097152		42035.10
4194304		41196.03



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Closing Remarks



- **NERSC will provide support for the vendor MPI (cray-mpich), on the contrary NERSC users will experiment with any MPI provider they have access to including OpenMPI, MVAPICH2, MPICH which can be troubling since we don't have proper support for any other MPI**
- **MVAPICH2-GDR is promising, considering it provides support with Slurm and Slingshot 11**
- **With sufficient testing, we can provide user with documentation on how to use MVAPICH2-GDR on Perlmutter**
 - Performance comparison between cray-mpich and MVAPICH2-GDR
- **MVAPICH2-GDR can provide user to experiment with spack build instead of being locked into cray-mpich, we have seen several spack packages fail to recognize MPI wrappers.**