

Red Hat container technologies for HPC

QR CODE



LINK

<https://red.ht/2HRQ2A>

p



PRODUCTS ▾

LEARN ▾

COMMUNITY ▾

SUPPORT ▾

FREE TRIAL

LOG IN ▸

Search

Search

[◀ Back to blog](#)

BLOG

Running HPC workloads with Red Hat OpenShift Using MPI and Lustre Filesystem

October 29, 2020 | by David Gray

The requirements associated with data science and AI/ML applications have pushed organizations toward using highly parallel and scalable hardware that often resemble [high performance computing \(HPC\) infrastructure](#). HPC has been around for a while and has evolved to include ultra large supercomputers that run massively parallel tasks and operate at exascale (able to perform a billion billion operations per second). Nowadays these [large systems need to run machine learning \(ML\) and other similar jobs](#) that often require use of multiple containerized applications, effectively blurring the line between traditional enterprise technologies and supercomputing applications.

Additionally, the sprawl of data has affected both scientific and commercial applications, necessitating the use of a robust distributed file system for data storage, a technology that became a standard part of operating environments at many HPC sites and enterprises.