

HPC Performance Considerations

Storage

The following points/suggestions should be considered when running jobs concurrently:

- HPC storage has been purchased to deliver expected capacity requirements - evaluation of vendor responses is dominated by \$/TB.
- Our current (2021) general use storage platform is very much entry level capabilities when it comes to performance.
- Researchers should avoid running parallel/concurrent `tar`, `gzip`, etc. commands. Utilities such as `pbzip2` should also be avoided.
 - If your computational research is I/O intensive, ensure that it is configured to use local scratch space (`/fast/tmp`)
 - Consider other researchers and create a separate, sequential operations job for all your post-job I/O transactions.
- Compression tools don't provide much benefit when working on binary data files - it's generally better to leave binary files uncompressed

Performance versus block size (read/write)

The following output shows the performance achievable, on SATA SSDs, for reads/writes of differing block sizes.

```
Block size 4K Write: 256000000 bytes (256 MB, 244 MiB) copied, 2.97716 s, 86.0 MB/s
Block size 4K Read : 256000000 bytes (256 MB, 244 MiB) copied, 2.54667 s, 101 MB/s
Block size 64K Write: 255983616 bytes (256 MB, 244 MiB) copied, 0.287594 s, 890 MB/s
Block size 64K Read : 255983616 bytes (256 MB, 244 MiB) copied, 0.248797 s, 1.0 GB/s
```

The story told by the brief output above is universal - small files move around at much slower rates than large files.