

What follows are some IOR tests run at Sandia on a 160-OSS/320-OST Lustre file system. This file system had just been reformatted, prior to the runs.

The following issues seem key ones:

- the single shared file is a factor 4-5 too slow, what is the overhead?
- why are reads so slow?
- why is there a significant read dropoff?
- why is two cores so much slower than single core?

Figure 1: File per process, using one core per socket follows

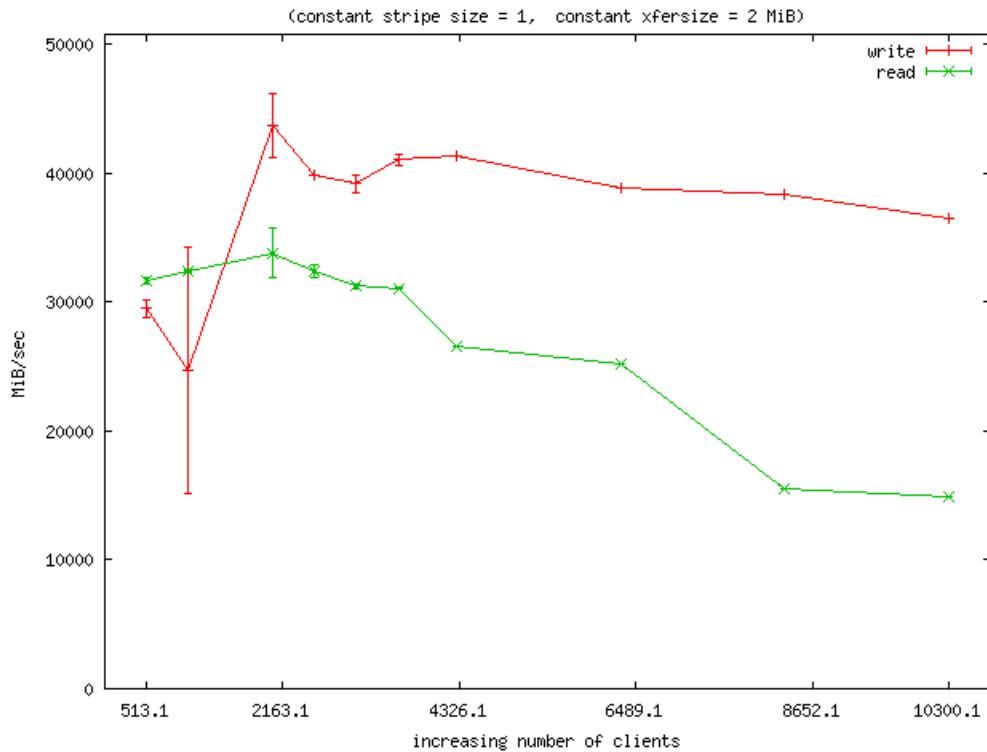


Figure 2: File per process, using two cores per socket:

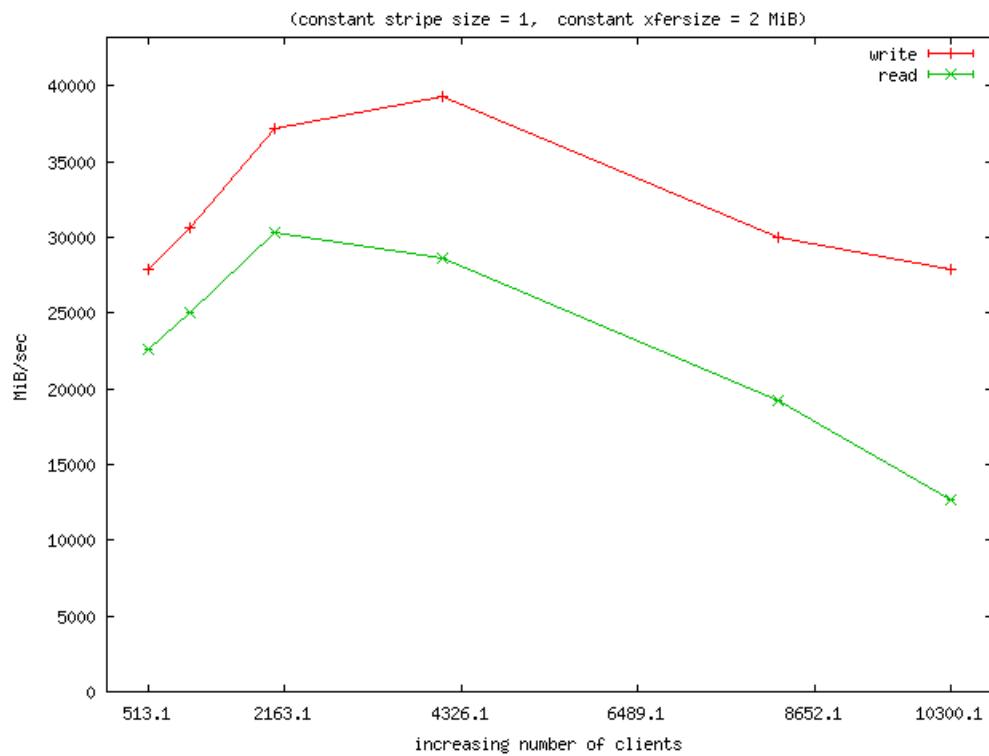


Figure 3: Single file, using one core per socket follows:

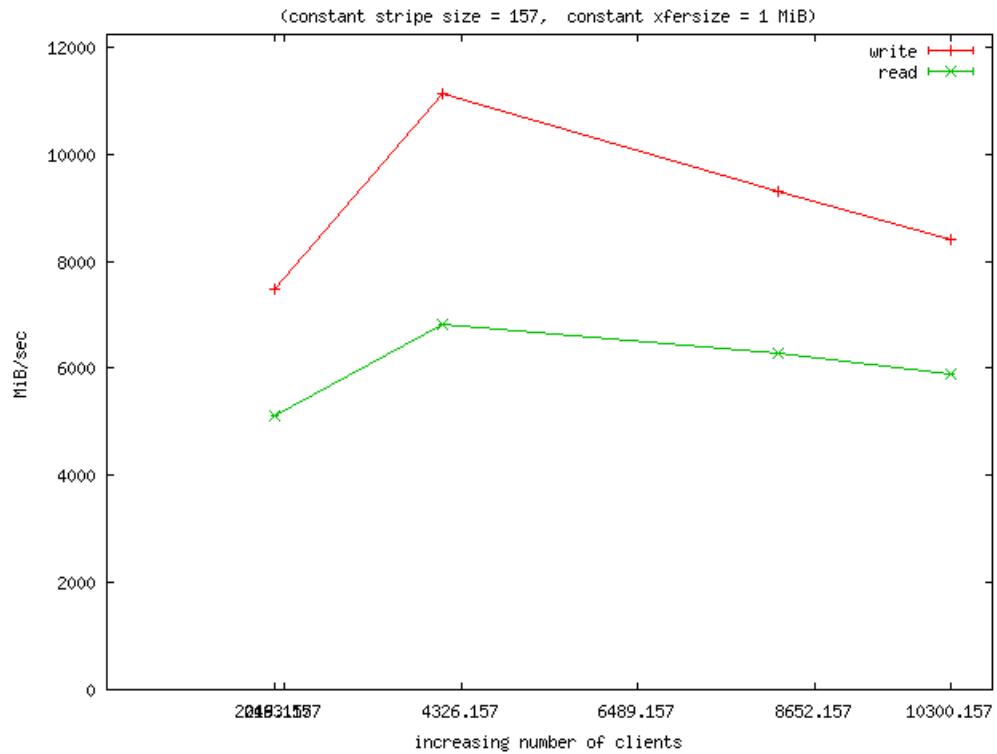


Figure 4: Single file, using two cores per socket follows:

