

Exascale challenges

- Metadata scalability
 - POSIX will not go away entirely
- Alternative storage models
 - Make better use of object stores
 - Push functionality/intelligence into storage nodes
 - Client write burst buffers
- Design for failure
 - As systems scale, no long feasible to address reliability with hardware

Paths forward

- Improve collaboration in kernel community
- Improve parallel extensions to POSIX
 - Efforts to push POSIX I/O HPC extensions (statlite, “lazy io”) into Linux have largely stalled
 - Focus on support in middleware?
- Adapt middleware to non-POSIX backends
 - e.g. NetCDF on top of an object store
- Lower barrier for testing new systems at scale