

PLFS: The Parallel Log-structured Filesystem

John Bent, Garth Gibson, Gary Grider, Ben McClelland, Paul Nowoczynski, James Nunez, Milo Polte, and Meghan Wingate

Parallel applications running on high performance computing clusters across thousands of processors rely on checkpointing to protect themselves from failures. The process of writing a checkpoint must be completed quickly so that applications may return to useful work. In this poster we present the Parallel Log-structured Filesystem (PLFS), a middleware layer that dramatically accelerates scientific application checkpoints. PLFS transparently decouples concurrent checkpoints into a filesystem-efficient access pattern of independent writes to individual log-files. By decoupling writes in this manner, PLFS removes points of contention, utilizes more filesystem resources, and dramatically decreases the time required to perform the checkpoint. PLFS is implemented as a FUSE-based filesystem requiring no changes to either application code or the underlying parallel filesystem. Our evaluation demonstrates that PLFS provides 2x-150x speedups for application checkpointing, with greater greater benefits at larger scale.