



Identification of Network Data Transfer **Bottlenecks in HPC Systems**



Karen Tu¹, Alex Sim (Advisor)², John Wu (Advisor)² ¹University of California Berkeley, ²Lawrence Berkeley National Laboratory



- High variation in throughput values suggest there are other parameters affecting throughput not considered in this study
- A higher number of OSTs is a bottleneck in network data transfer rates
- While file system stripe size and data transfer buffer size and parallelism were not bottlenecks, it is odd that there was no way to tune them to increase transfer performance

FUTURE WORK-

- Work with file system team to run controlled case studies
- Control for file system activity
- Collect data using different file sizes/structures
- Acquire more information about DTNs

ACKNOWLEDGMENTS —

I would like to thank my mentors Alex Sim and John Wu for their guidance and feedback. Specials thanks goes to Prof. Jinoh Kim and Mr. Jinhwan Kim at Texas A&M, Commerce for their guidance on methods. This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship (SULI) program. This work was supported by the U.S. Department of Energy, under Contract No. DE-AC02-05CH11231 and used resources of the National Energy Research Scientific Computing Center.

DISCUSSION

- If a file system is busy with reading/writing, throughput of data transfers to that system should be lowered – this is not the case
- Throughput is related to the number of OSTs a file is striped across
- Default of 1 OST yields best performance, 248 OSTs yields worst performance
- Bottlenecks for network data transfers occur in different places depending on number of OSTs
 - 248 OSTs: writing
 - 1, 4, or 16 OSTs: reading
- Tuning various file system and data transfer parameters had no effect on throughput rates
- Combinations of those parameters and holding certain values constant yielded similar looking scatterplots; no patterns

