

# Maximizing the Performance of Scientific Data Transfer by Optimizing the Interface Between Parallel File Systems and Advanced Research Networks

Nicholas L. Mills<sup>1</sup>, F. Alex Feltus<sup>2</sup>, and Walter B. Ligon III<sup>1</sup>

<sup>1</sup>Holcombe Department of Electrical and Computer Engineering

<sup>2</sup>Genetics and Biochemistry Department

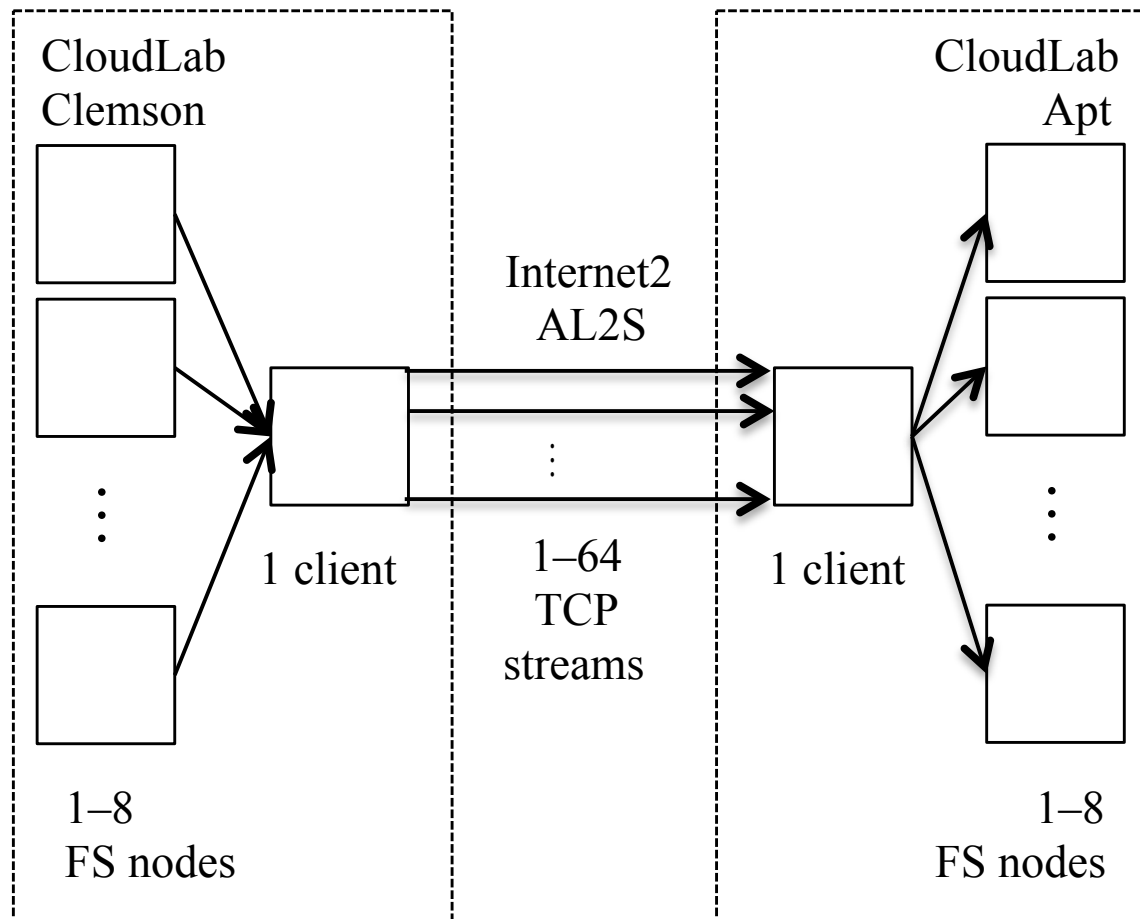
Clemson University




# Motivating questions

- What parameters can be tuned to affect the performance of parallel data transfers?
- What are the best parameters for disk-to-disk transfers of real scientific data?
- Is InfiniBand worth the extra cost/complexity?


# Dataset transfers in CloudLab





# File system benchmarks




# BeeGFS vs. OrangeFS transfers




# InfiniBand vs. TCP transfers



# 4 MiB vs. 16 MiB socket buffer



# Varying TCP streams and nodes





# Conclusion

- **Summary**

- The best throughput was found to occur with GridFTP using at least 5 parallel TCP streams with a 16 MiB TCP socket buffer size to transfer to/from 4–8 BeeGFS parallel file system nodes connected by InfiniBand.

- **Future work**

- Further tuning of kernel TCP parameters
- Alternative WAN protocols
- Effects of pipelining transfer and computation

# Acknowledgements

- **National Science Foundation**
- **CloudLab**
- **Anonymous reviewers**