An analysis of the effects of the Spectre and Meltdown patches on the Lustre parallel file system

Trevor Bautista
Arizona State University
tbautis1@asu.edu

Adam Good
Dakota State University
adam.good@trojans.dsu.edu

Amaris Velez
Polytechnic University of Puerto Rico
velez_90013@students.pupr.edu

Mentors: Garrett Ransom, Christian Storer, and Scott White

Overview

In a High-Performance Computing (HPC) environment it is imperative to ensure that systems perform securely while maintaining the best performance possible. To meet these performance requirements, it is common to use a parallel file system such as Lustre. However, these performance requirements also conflict with the Spectre and Meltdown patches that are notorious for reducing performance. This research focuses on evaluating the impact these patches have on the Lustre parallel file system.

Background

Spectre and Meltdown
Spectre and Meltdown allow attackers to access data that processors temporarily make available outside the chip.
- **Spectre**: Allows the attackers to trick processors into starting speculative execution process
- **Meltdown**: Allows attackers to access the sensitive information through a computer's operating system

Lustre
Lustre is a parallel file system that separates metadata from block I/O.
- Supports hundreds of PBs of data storage and hundreds of GB/s in throughput
- Composed of 3 servers: MGS, MDS, and OSS
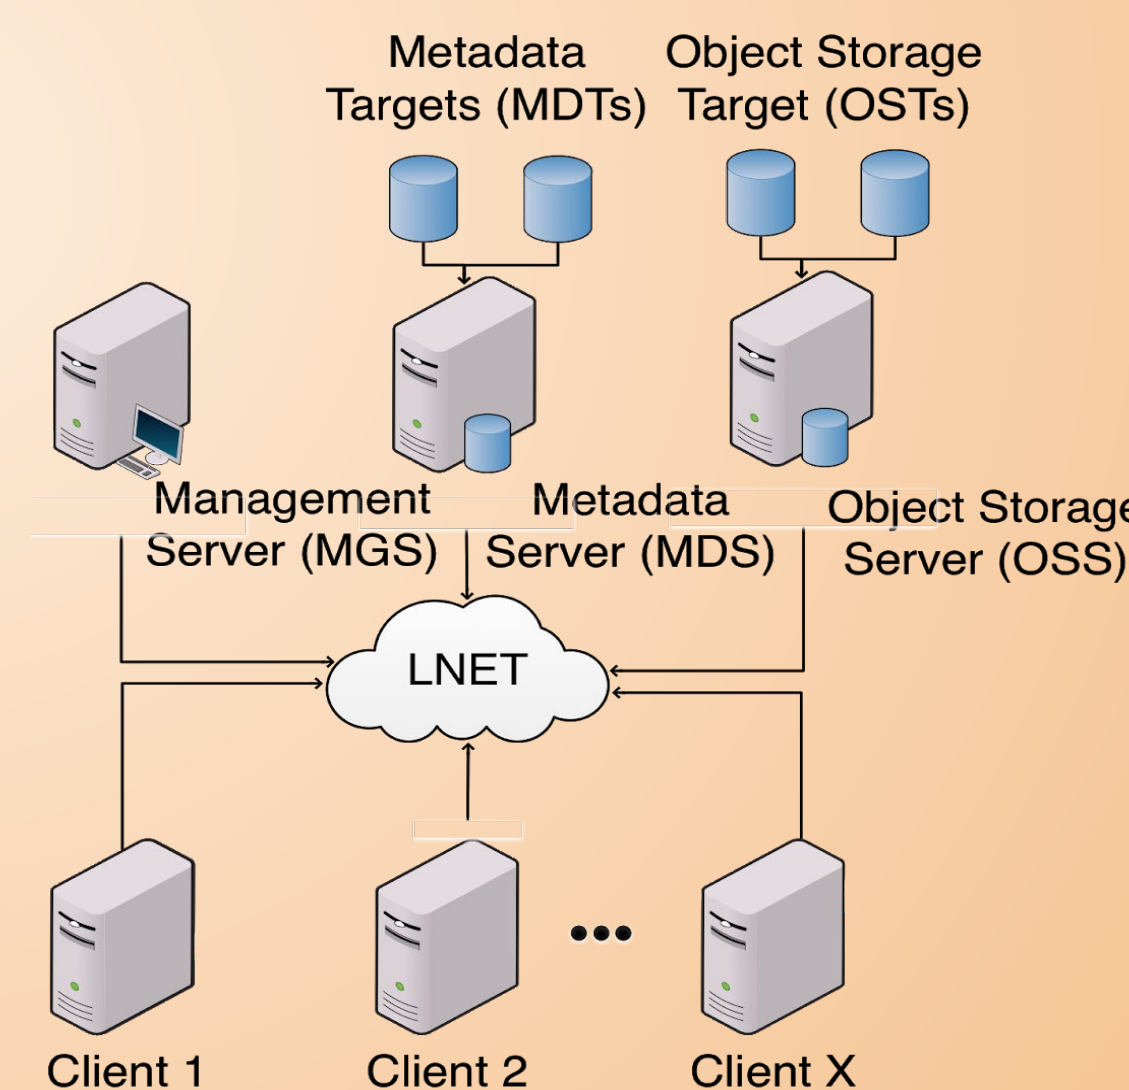
Methodology

Test cluster components: CentOS, ZFS, Lustre, Mellanox Infiniband, 1-8 clients (32 cores per client), and 3 servers (1 MGS, 1 MDS w/1 MDT, and 1 OSS w/1 OST)
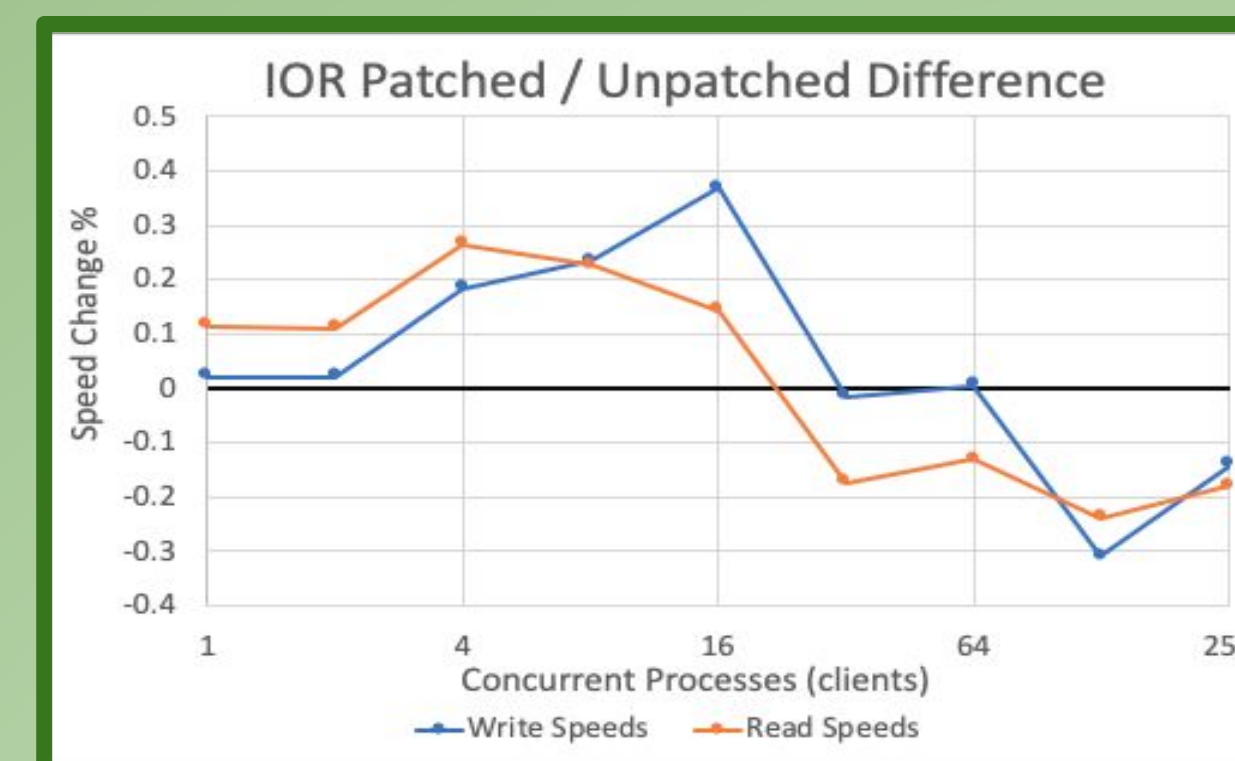
- Implemented tools: FIO, LNET-Selftest, MDtest, OBDFilter-Survey, and IOR
- Get baseline for file system performance
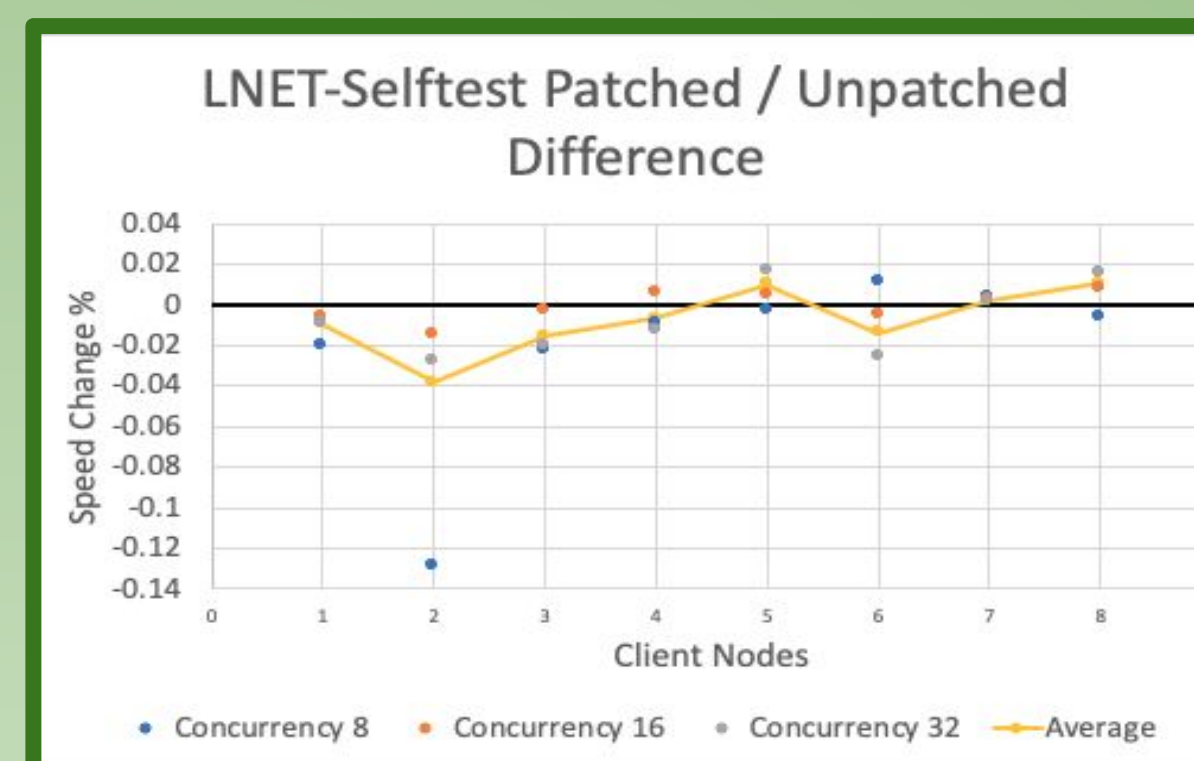- Patch file system servers and evaluate file system performance
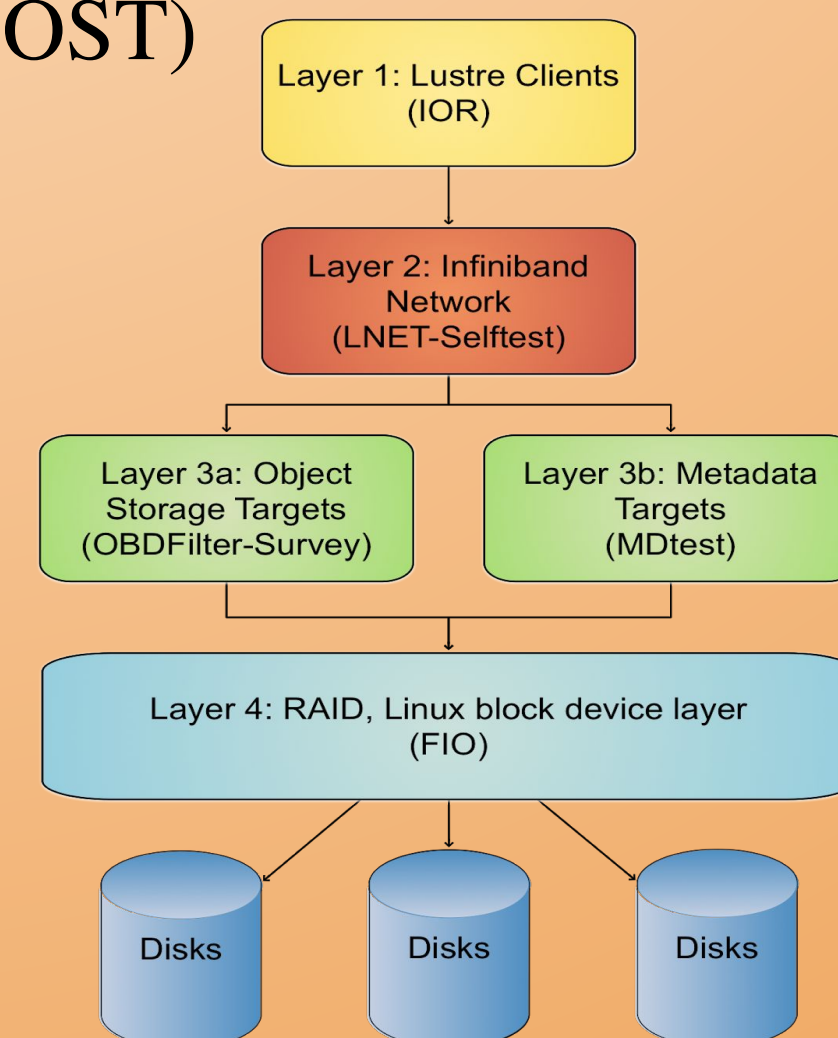
Benchmarking Results

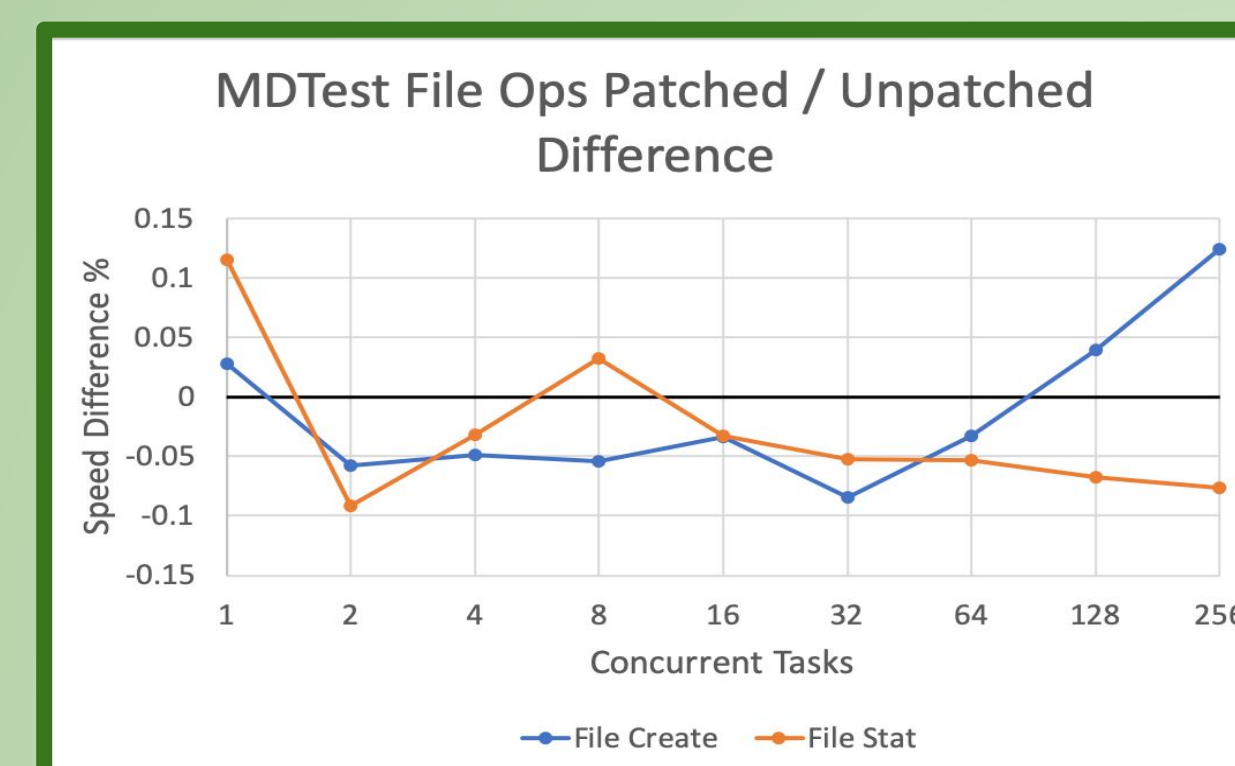The following graphs were obtained from the benchmarking procedure:
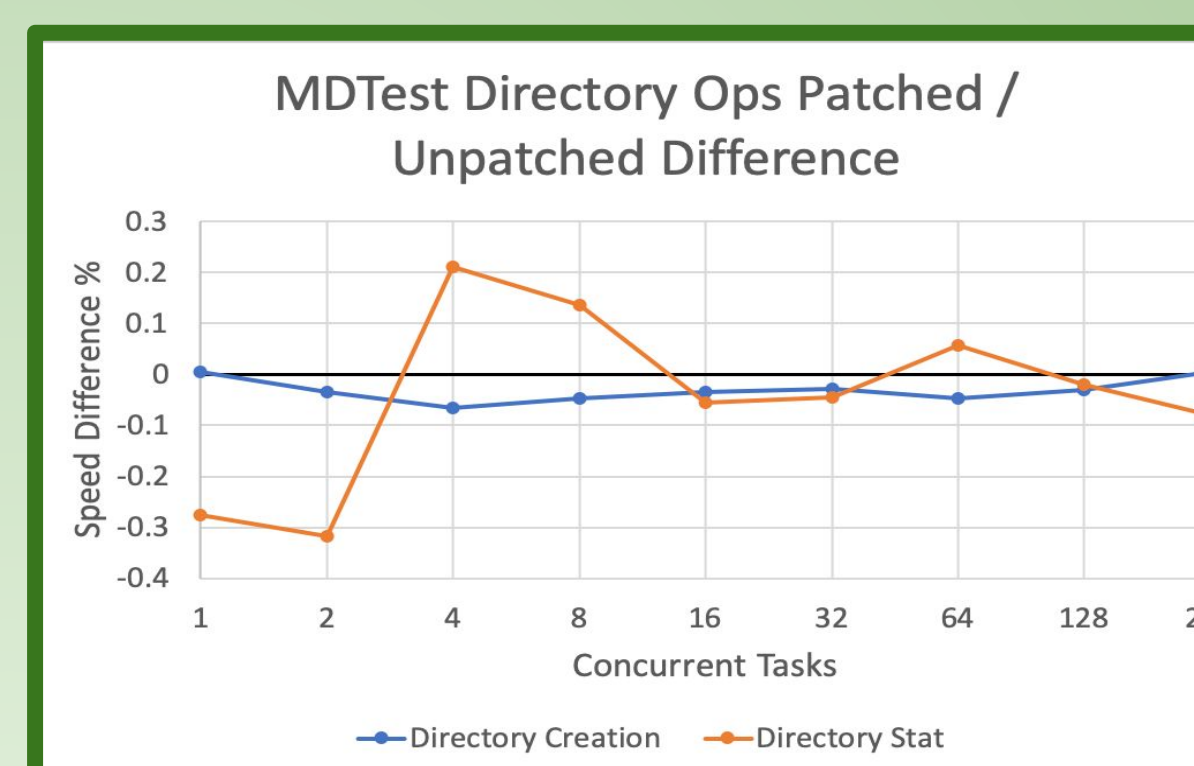
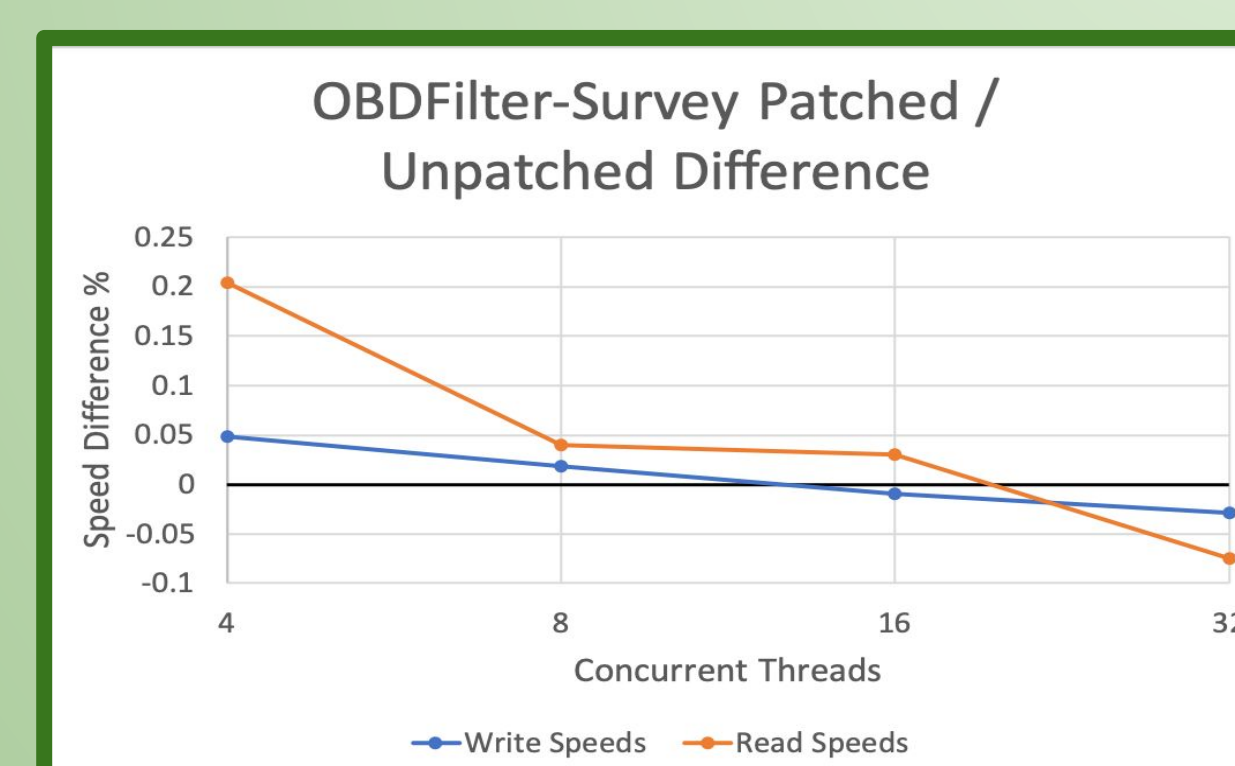### Lustre system speeds

Network speeds
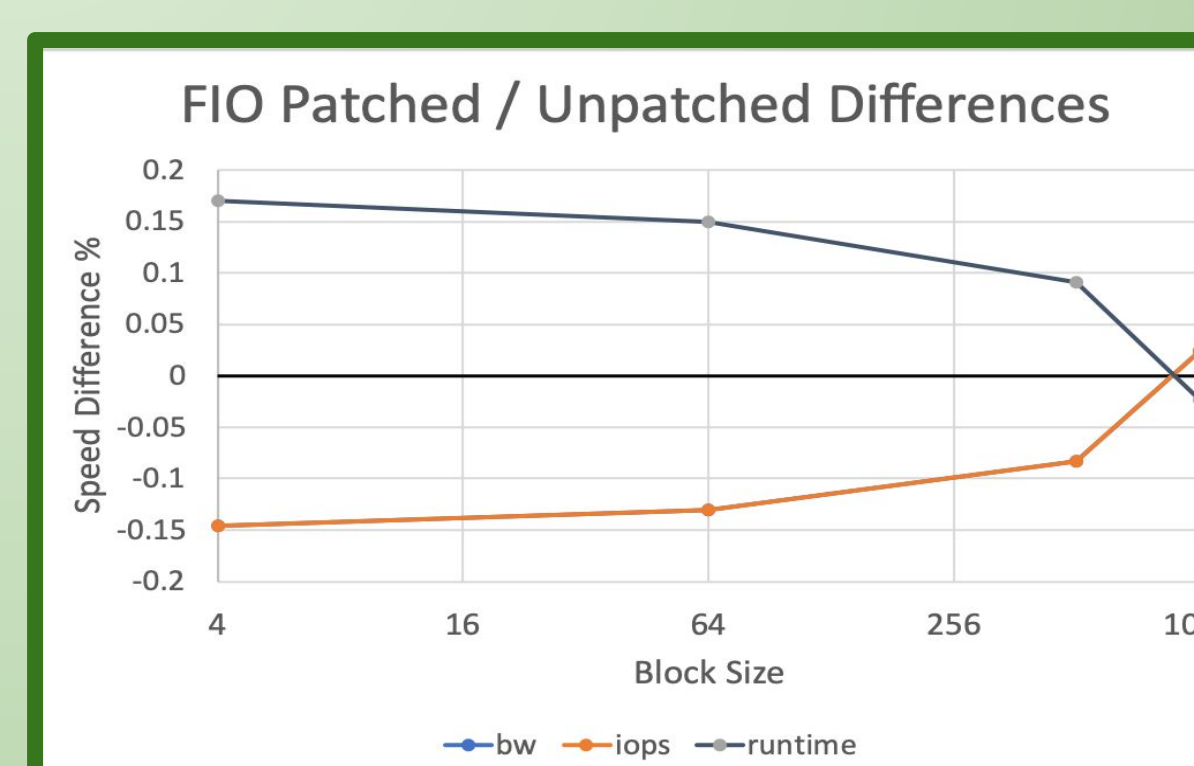

Metadata server speeds

Metadata server speeds

Object storage server speeds

Storage disks speeds

Discussion of results

The different stacks of the file system were evaluated and the following observations were made after patching:
- **IOR**: ~15% decrease in write/read performance
- **LNET-Selftest**: No significant decrease in performance
- **MDtest**: ~10% decrease in file/directory stat performance
- **OBDFilter-Survey**: ~7% decrease in read performance
- **FIO**: No significant decrease in performance

LA-UR-19-27101

Conclusions

The overall Lustre file system had a 15% decrease in performance due to the patches, one of the most significant effects observed. Among Lustre servers, the MDS suffered the greatest decrease in performance. From another point of view, the patches were found to not have a significant effect on ZFS's performance. Additionally, the effect of the patches on Lustre were more clearly observed as the number of nodes was increased.

Future Work

Due to the observation that the effect of the patches is amplified by the number of nodes, future research could focus on executing the same benchmarking processes with a greater number of nodes. An amount similar to that found on production clusters would be an appropriate choice. Additionally, the effect of the patches could be observed individually given the fact that there are many variants to Spectre and Meltdown and each patch for each variant could have a different effect on Lustre's performance. Finally, it may be beneficial to look into the cause of the increase in performance for small-scale patched systems.

Acknowledgements

Work Cited
[1] Kocher, P., Horn, Y., Fogh, A., Genkin, D., Gruss, D., Haas, W., Hamburg, M., Lipp, M., Mangard, S., Prescher, T., Schwarz, M., Yarom, Y. (2018). *Spectre Attacks: Exploiting Speculative Execution.*
[2] Simakov, N., Innus, M., Jones, M., Katz, O., White, J., Rathsam, R., Gallo, S., DeLeon, R. and Furlani, T. (2018). *Studying Effects of Meltdown and Spectre Patches on the Performance of HPC Applications Using Application Kernel Module of XDMoD*